

Linux® Host Utilities 5.0

Installation and Setup Guide

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

Part number: 215-04688_A0
April 2009

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

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

Trademark information

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

Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

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

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

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	v
Chapter 1	Setup Overview	1
	Downloading and Installing the Host Utilities Software	4
	Introduction to Fibre Channel	6
	Introduction to Software iSCSI.	11
	Introduction to Hardware iSCSI	13
Chapter 2	Installing the HBAs for FC	21
	HBA Installation	22
	Removing the previous external HBA driver.	23
	Downloading and installing the external HBA driver	24
	Installing HBA management software	25
	Setting the required HBA and bundled driver parameters	26
	Setting the required HBA and external driver parameters	31
	Loading the bundled HBA driver on the host	34
	Loading the external HBA driver on the host	35
	Verifying the HBA settings.	36
	Recording the WWPNs of the HBAs.	38
	Unloading the bundled or external HBA driver	40
Chapter 3	Installing the HBAs for Hardware iSCSI	41
	HBAs Installation.	42
	Install the HBA management software	43
	Setting the required HBA and bundled driver parameters	44
	Loading the bundled HBA driver on the host	45
	Verifying the HBA settings.	46
	Recording the initiator nodename of the HBAs	47

	Adding targets to the host.	48
	Unloading the bundled HBA driver	49
Chapter 4	Configuring the software iSCSI Initiator	51
	Configuring Red Hat Enterprise Linux	52
	Configuring SUSE Linux Enterprise Server	62
Chapter 5	Configuring the Storage System	75
	Configuring the storage system with FC	76
	Configuring the storage system with hardware and software iSCSI	83
Chapter 6	Configuring FC Cabling and Switches.	85
	Switched fabric configurations	86
	Direct-connected configurations	87
Chapter 7	Configuring DM-Multipath.	89
	Verifying the required multipathing packages	90
	Editing the configuration file	91
	Starting DM-Multipath	98
	Verifying the configuration.	100
	Stopping DM-Multipath	102
Chapter 8	Accessing LUNs using FC.	103
	Accessing LUNs with DM-Multipath	104
	Viewing the list of LUNs.	105
	Creating and mounting a file system on a multipath device	107
	Accessing LUNs without multipathing software.	109
	Identifying and viewing information about LUNs	111
	Creating a file system	113
	Labeling and mounting the file system.	114

Chapter 9	Accessing LUNs using Software iSCSI.	.117
	Starting the iSCSI service on the host	.118
	Configuring the iSCSI service to start automatically.	.119
	Accessing LUNs with DM-Multipath on Red Hat Enterprise Linux.	.120
	Viewing the list of LUNs.	.121
	Creating and mounting a file system on a multipath device	.126
	Accessing LUNs with DM-Multipath on SUSE Linux Enterprise Server	.128
	Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 9 series	.129
	Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 10 series.	.132
	Creating and mounting a file system on a LUN on SUSE Linux Enterprise Server 9 series	.135
	Creating and mounting a file system on a LUN on SUSE Linux Enterprise Server 10 series	.137
	Accessing LUNs without multipathing software on Red Hat Enterprise Linux.	.139
	Creating a file system	.141
	Labeling and mounting the file system.	.142
	Viewing LUN information	.144
	Accessing LUNs without multipathing software on SUSE Linux Enterprise Server	.150
	Viewing information about LUNs	.151
	Creating a file system	.158
	Mounting the file system.	.159
Chapter 10	Accessing LUNs using Hardware iSCSI	.163
	Accessing LUNs with DM-Multipath	.164
	Viewing the list of LUNs.	.165
	Creating and mounting a file system on a multipath device	.167
Chapter 11	Configuring SAN Boot using FC	.169
	Configuring SAN boot on a Red Hat Enterprise Server	.170
	Configuring SAN boot on a SUSE Linux Enterprise Server.	.174
Chapter 12	Configuring SAN Boot using software iSCSI	.179
	Configuring SAN boot on Red Hat Enterprise Linux	.180

	Configuring SAN boot on SUSE Linux Enterprise Server	185
Chapter 13	Configuring SAN boot using Hardware iSCSI	193
	Configuring SAN boot on Red Hat Enterprise Server	194
Chapter 14	Veritas Array Support Library and Array Policy Module.	197
	Overview of the ASL and APM	198
	Installing the ASL and APM	199
	Removing the ASL and APM	201
	Understanding ASL error message types.	203
Chapter 15	Accessing LUNs using Veritas DMP	205
Chapter 16	Configuring SAN Boot using Veritas Storage Foundation	211
	Configuring SAN boot on Red Hat Enterprise Server	212
	Configuring SAN boot on SUSE Linux Enterprise Server.	215
Appendix A	Troubleshooting	219
	Troubleshooting—FC.	220
	Troubleshooting—iSCSI	224
	Troubleshooting—Veritas Storage Foundation	229
	LUNs not visible on the host	231
Appendix B	Default and Recommended Settings	235
Appendix C	Third-Party Copyrights and Licenses	237

Preface

About this guide

This guide describes how to install the Linux® Host Utilities 5.0 for Native OS and Veritas® on a Linux host and how to set up the host to access the storage object described by a logical unit number (LUN). It also explains how to configure the storage system to work with the host.

This guide describes how to use both the iSCSI and Fibre Channel protocols. It also describes using both protocols to connect to the same LUN. Starting with Linux Host Utilities 5.0, the previous iSCSI Linux Host Utilities and FC Linux Host Utilities are merged into a single product.

Information specific to Veritas Storage Foundation™ is provided as separate chapters in this guide.

For more information about specific Linux versions supported, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Audience

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

Terminology

This guide uses the following terms:

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

Command conventions

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

Formatting conventions

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

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

Keyboard conventions

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

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

Special messages

This guide contains special messages that are described as follows:

Note

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

Attention

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

Setup Overview

About this chapter

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

About Linux versions

Linux Host Utilities supports Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Oracle Enterprise Linux operating systems, and Veritas Storage Foundation with Red Hat Enterprise Linux and SUSE Linux Enterprise Server operating systems. 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>

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 SAN boot

SAN boot uses a SAN-attached disk, such as a LUN configured on a storage controller, as a boot device for a host. For information about the supported operating system version and Veritas Storage Foundation version, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

What Linux Host Utilities is

Linux Host Utilities 5.0 enables a Linux host to connect to your storage systems. Linux Host Utilities consists of the following components:

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

Host Utilities components

Linux Host Utilities 5.0 consists of the following software components:

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

About SAN documentation

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

Understanding device mapper multipath

Linux Device Mapper Multipathing (DM-Multipath) is supported on NetApp storage systems with Red Hat Enterprise Linux 5 series, Red Hat Enterprise Linux 4 series, SUSE Linux Enterprise Server 9 series, and SUSE Linux Enterprise Server 10 series. The DM-Multipath provides I/O failover and path

load sharing for multipathed block devices. DM-Multipath allows you to configure multiple I/O paths between host and storage controllers into a single device. If one path fails, DM-Multipath reroutes I/Os to the remaining paths.

The required kernel module for DM-Multipath are included in the Red Hat Enterprise Linux 5 series, Red Hat Enterprise Linux 4 series, SUSE Linux Enterprise Server 9 series, and SUSE Linux Enterprise Server 10 series binary kernel packages. You should install additional user space packages as part of the setup process.

When you have multiple paths to a LUN, Linux creates a SCSI device for each path. For example, a single LUN might appear as `/dev/sdd` and `/dev/sdf` if there are two paths to the LUN. DM-Multipath creates a single device in `/dev/mapper/` for each LUN that represents all the paths. For example, `/dev/mapper/360a9800043346852563444717a513571` is the multipath device created on top of `/dev/sdd` and `/dev/sdf`. You should create file systems and mount the LUN using the device in `/dev/mapper/`.

Understanding device persistence

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/`. You do not have to create a partition or label on the multipath device.

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

Understanding Veritas Dynamic 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.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Downloading and Installing the Host Utilities Software”](#) on page 4
- ◆ [“Introduction to Fibre Channel”](#) on page 6
- ◆ [“Introduction to Software iSCSI”](#) on page 11
- ◆ [“Introduction to Hardware iSCSI”](#) on page 13

Downloading and Installing the Host Utilities Software

Downloading the Host Utilities software

Download the Linux Host Utilities file (netapp_linux_host_utilities_5_0.tar.gz) 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, 32-bit compatibility library for libdevmapper is not available. You can download the 32-bit compatibility package for device-mapper-multipath from the Novell® Web site at <http://support.novell.com/techcenter/psdb>. You should install this package, to ensure that sanlun works on SUSE Linux Enterprise Server 9 SP3 x86_64 bit.

Verifying the correct version

Verify that you have the correct version of 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>

Installing the Host Utilities software

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

Step	Action
1	To remove any previous version of Linux Host Utilities, go to the directory where the previous version is installed (default is /opt/netapp/santools) and enter the following command: <code>./uninstall</code>
2	Go to the working directory to which you copied the Host Utilities file.
3	To extract the files, enter the following command: <code>tar -zxvf netapp_linux_host_utilities_5_0.tar.gz</code>

Step	Action
4	Change to the netapp_linux_host_utilities_5_0 directory. By default, this is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.
5	Enter the following command: <code>./install</code>

Note

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

Introduction to Fibre Channel

About host bus adapters

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

With Host Utilities 5.0, NetApp supports both bundled HBA drivers and external HBA drivers while using DM-Multipath or Veritas multipathing (VxDMP) on Linux hosts.

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

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

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.

ALUA support: Linux Host Utilities provides support for Asymmetric Logical Unit Access (ALUA). To see which versions of Data ONTAP® and Host Utilities supports ALUA, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability/>

Note

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

ALUA defines a standard set of SCSI commands for discovering and managing multiple paths to LUNs on FC and iSCSI SANs. DM-multipath works with ALUA to determine which paths are primary paths and which paths are secondary or partner paths to be used for failover only..

NetApp introduced ALUA support with Data ONTAP 7.2.4 and `single_image` `cfmode`.

Note

ALUA is not supported with Veritas Storage Foundation.

Data ONTAP requirements: For NetApp FC and iSCSI products for supported versions of Data ONTAP, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability/>

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 a 4-Gb target HBA, then you should manually change the cfmode.

For more information, see “[Understanding single_image mode](#)” on page 78.

About host and storage system communication: After you install and configure your FC HBA, the Linux host uses the FC protocol to access data on storage systems that run supported versions of Data ONTAP software.

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

On a host, a worldwide port name (WWPN) identifies each port on an HBA. The host WWPNs are used as identifiers when creating initiator group (igroup) on a storage system. An igroup 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-port FC connection from the HBA to the storage system through a single switch

A host is cabled to a single FC switch that is connected by cable to redundant FC ports on a high-availability storage system. A fabric-attached single-path host has one HBA.

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

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

- ◆ Direct-attached

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

Note

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

For more information (including diagrams) about the supported FC topologies, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

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
- ◆ FC 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>

Setup procedure

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

Step	Action
1	Install and configure HBAs and the Host Utilities software on the host, including the following: ◆ Installing the HBAs, if not already installed ◆ Removing the previous external HBA drivers ◆ Downloading and installing the supported external HBA driver ◆ Installing the Linux Host Utilities software ◆ Setting required HBA parameters ◆ Installing DM-Multipath packages to use the Native multipathing solution ◆ Recording the WWPNs of the HBAs ◆ Installing Veritas Storage Foundation, Array Support Library (ASL), and Array Policy Module (APM), if you want to use the Veritas DMP solution Note _____ For information about installing Veritas Storage Foundation, see the <i>Veritas Storage Foundation Installation Guide</i> for Linux that is shipped along with the Veritas Storage Foundation software. See “Veritas Array Support Library and Array Policy Module” on page 197. See “Installing the HBAs for FC” on page 21.
2	Configure the storage system, including the following: ◆ Licensing and starting the FC service ◆ Creating LUNs, creating an igroup that contains the host’s FC WWPn, and mapping the LUNs to that igroup ◆ Optional: Enabling ALUA, if required Note _____ ALUA is not supported with Veritas Storage Foundation. See “Configuring the Storage System” on page 75.
3	Optional: Migrate the storage system configuration to single_image cfmode. See “Migrating to single_image mode” on page 80.

Step	Action
4	Configure the FC switches and cabling, including the following: ◆ Verifying the firmware revision of the switches ◆ Cabling the storage system to the switches ◆ Cabling the HBAs to the switches or directly to the storage system ◆ Creating zones on the switches See “Configuring FC Cabling and Switches” on page 85.
5	Discover and access the LUNs from the host, including the following: ◆ Loading the HBA driver modules See “Installing the HBAs for FC” on page 21. ◆ Configuring DM-Multipath or Veritas DMP See “Configuring DM-Multipath” on page 89 or “Veritas Array Support Library and Array Policy Module” on page 197.

Introduction to Software iSCSI

Topics in this section

This section includes the following topics:

- ◆ “[Understanding how Data ONTAP software implements iSCSI](#)” on page 11
- ◆ “[Main features supported by Host Utilities](#)” on page 11
- ◆ “[Setting up the iSCSI protocol](#)” on page 12

Understanding how Data ONTAP software implements iSCSI

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. RFC 3720 defines the iSCSI protocol standard, and is published by the Internet Engineering Task Force. You can get a copy of the standard at www.ietf.org.

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as logical units. A Linux host running the iSCSI initiator software uses the iSCSI protocol to access Logical Unit Numbers (LUNs) on a storage system running Data ONTAP software. The iSCSI protocol is implemented over the host’s standard Gigabit Ethernet interfaces using a software driver.

The storage system does not require a hardware iSCSI HBA. The iSCSI protocol on the storage system is implemented over the storage system’s standard Gigabit Ethernet interfaces using a software driver that is integrated into Data ONTAP.

The connection between the initiator and target uses a standard TCP/IP network. No special network configuration is needed to support iSCSI traffic. The network can be a dedicated TCP/IP network, or it can be your usual public network. The storage system listens for iSCSI connections on IP port 3260.

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

Main features supported by Host Utilities

Open-iSCSI software initiator: SUSE Linux Enterprise Server 10 series or Red Hat Enterprise Linux 5 series come with the Open-iSCSI software initiator. See “[Where to go for more information](#)” on page 15 for links to Novell® and Red Hat documentation.

Setting up the iSCSI protocol

To set up the iSCSI protocol on a host and storage system, complete the following steps.

Step	Action
1	Install and configure the Host Utilities software on the host, including the following: ◆ Installing the required packages on the host, and recording or changing the host's iSCSI node name ◆ Configuring DM-Multipath support Note _____ If you are not using DM-Multipath support, skip this step. ◆ Setting the initiator parameters, including the IP address of the target on the storage system ◆ Optional: Configuring CHAP See “Configuring Red Hat Enterprise Linux” on page 52 and “Configuring SUSE Linux Enterprise Server” on page 62.
2	Configure the storage system, including the following: ◆ Licensing and starting the iSCSI service ◆ Optional: Configuring CHAP ◆ Creating LUNs, creating an igroup that contains the host's iSCSI node name, and mapping the LUNs to that igroup See “Configuring the Storage System” on page 75.
3	Access the LUNs from the host, including the following: ◆ Starting the iSCSI service ◆ Configuring the multipath service ◆ Creating file systems on the multipath device and mounting them See “Accessing LUNs using Software iSCSI” on page 117.

Note _____

You should alternate between setting up the host and the storage system, as shown in the “[Setting up the iSCSI protocol](#)” on page 12.

Introduction to Hardware iSCSI

About host bus adapters

Linux Host Utilities supports QLogic® brand host bus adapters (HBAs).

With Host Utilities 5.0, NetApp supports bundled HBA drivers while using DM-Multipath on Linux hosts.

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

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.

Data ONTAP requirements: For NetApp FC and iSCSI products for supported versions of Data ONTAP, see NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

If you do not have a supported version of Data ONTAP, upgrade to one of the supported versions.

About host and storage system communication: After you install and configure your iSCSI HBA, the Linux host uses the iSCSI protocol to access data on storage systems that run supported versions of Data ONTAP software.

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

On a host, initiator name and IP address identifies each port on an HBA. The host initiator names are used as identifiers when creating initiator group (igroup) on a storage system. An igroup permits the host access to specific LUNs on a storage system.

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
- ◆ Storage system and Data ONTAP version

You can find the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Setup procedure

The procedure for configuring the iSCSI protocol on a host and storage system follows the same basic sequence for all hosts.

Step	Action
1	Install and configure HBAs and the Host Utilities software on the host, including the following: ◆ Installing the HBAs, if not already installed ◆ Removing the previous external HBA drivers ◆ Downloading and installing the supported external HBA driver ◆ Installing the Linux Host Utilities software ◆ Setting required HBA parameters ◆ Installing DM-Multipath packages to use the Native multipathing solution See “Installing the HBAs for Hardware iSCSI” on page 41.
2	Configure the storage system, including the following: ◆ Licensing and starting the iSCSI service ◆ Creating LUNs, creating an igroup that contains the host’s iSCSI initiator name, and mapping the LUNs to that igroup See “Configuring the Storage System” on page 75.
3	Discover and access the LUNs from the host, including the following: ◆ Loading the HBA driver modules See “Installing the HBAs for FC” on page 21. ◆ Configuring DM-Multipath See “Configuring DM-Multipath” on page 89.

Where to go for more 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 this source...
Known issues, system requirements, and last minute updates	The <i>Release Notes</i> for Host Utilities 5.0, available at: http://now.netapp.com/NOW/knowledge/docs/san/
The latest supported configurations	◆ The NetApp Interoperability Matrix at: http://now.netapp.com/NOW/products/interoperability ◆ The <i>System Configuration Guide</i> at: https://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/
Configuring the storage system	◆ The <i>Best Practices for Reliability</i> at: http://now.netapp.com/NOW/knowledge/docs/docs.cgi ◆ The <i>Data ONTAP Software Setup Guide</i> at: http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml
Supported Fibre Channel SAN topologies	The NetApp Interoperability Matrix at: http://now.netapp.com/NOW/products/interoperability
Installing and configuring the HBA on 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.

If you need information about...	See this source...
Managing SAN on the storage system	◆ The <i>Command Reference</i> at: http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The <i>Data ONTAP Block Access Management Guide</i> at: http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The FilerView® online Help ◆ The <i>Best Practices for Reliability: Ongoing Maintenance</i> at: http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm.
Enabling ALUA	<i>Data ONTAP Block Access Management Guide</i> for FC at: http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. In particular, see the section “Enabling ALUA for a Fibre Channel igroup”.
Changes in this release of Host Utilities, a list of known problems, and limitations of Host Utilities	The <i>Release Notes</i> for Host Utilities 5.0.
The most current system requirements for the Linux host	The <i>NetApp Interoperability Matrix</i> at the following site: ◆ http://now.netapp.com/NOW/products/interoperability
Best practices or configuration issues	NetApp Knowledge Base at the following sites: ◆ https://now.netapp.com/eservice/kbbrowse Using partitions on Linux with NetApp LUNs may require alignment for best performance at: http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190

If you need information about...	See this source...
The supported storage system models for Data ONTAP licensed with iSCSI	◆ The <i>NetApp Interoperability Matrix</i> at the following site: ❖ http://now.netapp.com/NOW/products/interoperability ◆ <i>System Configuration Guide for Data ONTAP</i> online version at: http://now.netapp.com/NOW/knowledge/docs
Information about how to configure and manage the iSCSI service on the storage system	The following documentation for your release of Data ONTAP: ◆ <i>Data ONTAP Block Access Management Guide</i> ◆ <i>Data ONTAP Release Notes</i>
Information about Oracle Cluster File System	<i>Oracle Cluster File System (OCFS2) User's Guide</i> and <i>OCFS2 - Frequently Asked Questions</i> at: http://oss.oracle.com
Information about Oracle Real Application Cluster (RAC)	<i>Installing Oracle RAC</i> at http://www.oracle.com
NetApp Knowledge Base for Oracle RAC	The knowledge base for Oracle RAC at the following sites: ◆ http://www.netapp.com/library/tr/3423.pdf ◆ http://www.netapp.com/library/tr/3369.pdf
Information about managing the iSCSI initiator	◆ For Red Hat Enterprise Linux 5 series: ❖ The <code>iscsid.conf</code>, <code>iscsiadm</code>, and <code>iscsid</code> man pages on the host. ◆ For Red Hat Enterprise Linux 4 series Update 3 and later: ❖ The <code>iscsi.conf</code>, <code>iscsi-ls</code>, and <code>iscsid</code> man pages on the host. ◆ Red Hat Enterprise Linux: ❖ Documentation on Red Hat at: http://www.redhat.com

If you need information about...	See this source...
Information about installing and configuring SnapDrive® for UNIX®	The <i>SnapDrive for UNIX Installation and Administration Guide</i> for your version of SnapDrive.
Information about configuring multipath support on Red Hat Enterprise Linux	The multipath-usage.txt file included with the device-mapper-multipath RPM for Red Hat Enterprise Linux 4 Update 3 in /usr/share/doc/device-mapper-multipath-<version>. <version> is the latest version number shipped with the Red Hat Enterprise Linux release.
Information about Red Hat Cluster and GFS	http://www.redhat.com
Information about configuring multipath support on SUSE Linux Enterprise Server	The Novell Knowledgebase article <i>How to set up / use multipathing on SUSE Linux Enterprise Server</i> at: http://support.novell.com
Information about managing the iSCSI initiator on SUSE Linux Enterprise Server 9 series	◆ The README file in /usr/share/doc/packages/linux-iscsi on the host. ◆ The iscsi.conf, iscsi-ls, and iscsid manual (man) pages on the host. ◆ Documentation on the SUSE Web site at: http://www.novell.com
Information about configuring and managing the iSCSI initiator on SUSE Linux Enterprise Server 10 series	◆ The README file in /usr/share/doc/packages/open-iscsi on the host. ◆ The iscsid and iscsiadm man pages on the host. ◆ SUSE Linux Enterprise Server 10 series documentation on the Novell Web site at: http://www.novell.com

Installing the HBAs for FC

Configuration overview

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

Topics in this chapter

This chapter includes the following topics:

- ◆ [“HBA Installation”](#) on page 22
- ◆ [“Removing the previous external HBA driver”](#) on page 23
- ◆ [“Downloading and installing the external HBA driver”](#) on page 24
- ◆ [“Installing HBA management software”](#) on page 25
- ◆ [“Setting the required HBA and bundled driver parameters”](#) on page 26
- ◆ [“Setting the required HBA and external driver parameters”](#) on page 31
- ◆ [“Loading the bundled HBA driver on the host”](#) on page 34
- ◆ [“Loading the external HBA driver on the host”](#) on page 35
- ◆ [“Verifying the HBA settings”](#) on page 36
- ◆ [“Recording the WWPNs of the HBAs”](#) on page 38
- ◆ [“Unloading the bundled or external HBA driver”](#) on page 40

Requirement privileges login

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

HBA Installation

Verify the HBA model is supported

Verify that the Linux host supports your HBA model. For information about the supported HBA models, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Do not mix HBA models

If you install more than one HBA, ensure that all HBAs are the same brand and bus type (PCI-X or PCI Express). Mixing HBA brands might lead to inconsistent performance and behavior.

Hardware installation

If the HBAs are not already installed, install them on the Linux host. For instructions, see the HBA and host hardware vendors' documentation.

Removing the previous external HBA driver

Removing previous external QLogic HBA drivers

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

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

Removing previous external Emulex HBA drivers

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

Step	Action
1	Log in to the Linux host as root.
2	Go to the directory where the driver installer files are located.
3	Uninstall the previous Emulex driver on the host. Ensure that the driver is not in use (that is, the usage count of lpfc is 0). To uninstall the driver, run the following command: ./lpfc-install --uninstall

Downloading and installing the external HBA driver

Verify required driver version

Download the appropriate HBA driver from the vendor Web site. Ensure that you verify the required driver version by using the support matrix before downloading. For details about the supported driver versions, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

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

Download the external HBA driver

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

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

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

Install the driver software

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

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 download and install the HBA management packages from the respective HBA vendor Web site. These packages contains software components used to display HBA information (also used by the `sanlun` utility). You should also download the rescan script from the HBA vendor Web site.

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

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

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

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

Setting the required HBA and bundled driver parameters

Setting the required HBA and bundled driver parameters

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

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

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

Step	Action	
7	Depending on your Linux version, complete the following steps.	
	If you are using...	Then...
	Red Hat Enterprise Linux 4	a. In the <code>/etc/modprobe.conf</code> file, append the following line: <pre>options qla2xxx qlport_down_retry=30 ql2xsuspendcount=6</pre> b. Re-create the <code>initrd</code> by running the following command: <pre>mkinitrd -f /boot/initrd-"`uname -r`.img "`uname -r` -- with=qioctlmod</pre> Note _____ Backup of the current <code>initrd</code> before running the <code>mkinitrd</code> command. It is best to take a backup as a precaution. c. Reboot the host. d. Verify that the values of the <code>qlport_down_retry</code> and <code>ql2xsuspendcount</code> parameters in the corresponding <code>qla2xxx</code> directory in either <code>/sys/module/qla2xxx/</code> or <code>/sys/module/qla2xxx/parameters/</code> matches the values in <code>modprobe.conf</code> file.

Step	Action
	Example: You can check the value of the ql2xsuspendcount on a Red Hat Enterprise Linux host by running the following command: <pre>cat /sys/module/qla2xxx/ql2xsuspendcount</pre>
	Red Hat Enterprise Linux 5 - In the /etc/modprobe.conf file, append the following line: <pre>options qla2xxx qlport_down_retry=30</pre> - Re-create the initrd by running the following command: <pre>mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`</pre> Note _____ Backup of the current initrd before running the mkinitrd command. It is best take a backup as a precaution. - Reboot the host. - Verify that the value of the qlport_down_retry parameter in the corresponding qla2xxx directory in /sys/module/qla2xxx/parameters/ matches the values in modprobe.conf.

Step	Action	
		Example: You can check the value of the qlport_down_retry on a Red Hat Enterprise Linux host by running the following command: <pre>cat /sys/module/qla2xxx/parameters/qlport_down_retry</pre> Note _____ There are no qla2300/qla2400/qla2500 firmware modules in Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module itself.
	SUSE Linux Enterprise Server 10	a. In the /etc/modprobe.conf.local file, append the following line: <pre>options qla2xxx qlport_down_retry=30</pre> b. Re-create the initrd by running the mkinitrd command. Note _____ Take a backup of the current initrd before running the mkinitrd command. It is best to take backup as a precaution. c. Reboot the host.

Step	Action
	d. Verify that the value of qlport_down_retry parameter in the corresponding qla2xxx directory in /sys/module/qla2xxx/parameters/ matches the values in modprobe.conf.local. Example: You can check the value of the qlport_down_retry on a SUSE Enterprise Linux Server host by running the following command: <pre data-bbox="892 715 1217 802">cat /sys/module/qla2xxx/parameters/qlport_down_retry</pre>

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

Setting the required HBA and external driver parameters

For QLogic HBAs

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

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

Step	Action	
1	Run the following command on the host: scli	
2	From the main menu, select the Configure HBA Settings option.	
3	Select the first HBA port that is listed.	
4	Configure the QLogic HBA port depending on the topology that you are using.	
	If you storage system is operating in...	Then set ConnectionOptions to...
	SSI, Standby, or Partner CF modes	Point to Point Only
	Dual Fabric or Mixed CF modes	Loop Only
	Direct-attached	Point to Point Only
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.	

Step	Action
9	Repeat the settings for each HBA port that is 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	Depending on your HBA model, unload the external QLogic HBA driver by entering the following command: <pre>./qlinstall -ul <driver model></pre> Where, <i><driver model></i> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx
2	Append the <code>options qla2xxx</code> statement with the settings in the <code>modprobe.conf</code> file on a single line. <pre>ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> Note Only one <code>options qla2xxx</code> statement is allowed.
3	Save the file and exit.

Step	Action
4	Depending on your HBA model, unload the external QLogic HBA driver by entering the following command: <pre>./qlinstall -l <driver model></pre> Where, <i><driver model></i> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx
5	Depending on your HBA model, unload the external QLogic HBA driver by entering the following command: <pre>./qlinstall -br -in <driver model></pre> Where, <i><driver model></i> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx
6	Reboot the Linux host by using the updated image.
7	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 and a SUSE Linux Enterprise Server hosts by running the following command: <pre>cat /sys/module/qla2xxx/ql2xsuspendcount</pre> Note _____ You cannot verify the same in SUSE Linux Enterprise Server 9 as there is no qla2xxx directory available in sysfs. _____

For Emulex HBAs

No changes required.

Loading the bundled HBA driver on the host

Starting FC

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

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

Step	Action
1	Depending on your HBA model, unload the bundled QLogic HBA driver by entering the following command: <pre>modprobe -v <driver model></pre> Where, <driver model> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx

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

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

Loading the external HBA driver on the host

Starting FC

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

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

Step	Action
1	Depending on your HBA model, unload the bundled QLogic HBA driver by entering the following command: <pre>modprobe -v <driver model></pre> Where, <driver model> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx

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

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

Verifying the HBA settings

Verifying the QLogic HBA settings

To verify the PortDownRetryCount and LinkDownTimeout parameters, complete the following step.

Step	Action
1	To verify PortDownRetryCount and LinkDownTimeout along with Connection Options and Data Rate (which need to be manually set instead of the default autonegotiate), run the <code>-c</code> option of the <code>scli</code> command: <pre>/opt/QLogic_Corporation/SANsurferCLI/scli -c</pre> Example: The following is a sample output of the <code>scli</code> command: <pre>----- 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</pre>

Step	Action
	<pre> 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 </pre>

Verifying the Emulex HBA settings

The default settings work fine in the case of Emulex.

Recording the WWPNs of the HBAs

About WWPNs

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

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

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

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

Obtaining the WWPn for QLogic HBAs

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

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

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

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

Obtaining the WWPN for Emulex HBAs

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

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

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

Unloading the bundled or external HBA driver

Before unloading the HBA driver

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

Unloading the HBA driver

To unload the HBA driver, complete the following steps.

Step	Action
1	QLogic driver: Depending on your HBA model, unload the QLogic HBA driver by entering the following command: <code>modprobe -r <driver model></code> Where <i><driver model></i> is one of the following: ◆ qla2500 ◆ qla2400 ◆ qla2300 ◆ qla2xxx
	Emulex driver: Enter the following command: <code>modprobe -r lpfc</code>

Installing the HBAs for Hardware iSCSI

Configuration overview

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

Topics in this chapter

This chapter includes the following topics:

- ◆ [“HBAs Installation”](#) on page 42
- ◆ [“Install the HBA management software”](#) on page 43
- ◆ [“Setting the required HBA and bundled driver parameters”](#) on page 44
- ◆ [“Loading the bundled HBA driver on the host”](#) on page 45
- ◆ [“Verifying the HBA settings”](#) on page 46
- ◆ [“Recording the initiator nodename of the HBAs”](#) on page 47
- ◆ [“Adding targets to the host”](#) on page 48
- ◆ [“Unloading the bundled HBA driver”](#) on page 49

HBAs Installation

Verify the HBA model is supported

Verify that the Linux host supports your HBA model. For information about the supported HBA models, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Do not mix HBA models

If you install more than one HBA, ensure 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

If the HBAs are not already installed, install them on the Linux host. For instructions, see the HBA and host hardware vendors' documentation.

Install the HBA management software

You should download and install the HBA management packages from the respective HBA vendor Web site. These packages contains software components used to display HBA information (also used by the `sanlun` utility). You should also download the rescan script from the HBA vendor Web site.

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

- ❖ SANsurfer iSCSI HBA Manager
- ❖ SANsurfer iSCLI

Setting the required HBA and bundled driver parameters

For QLogic HBAs

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

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

Step	Action
1	Run the following command on the host: <code>iscli</code>
2	From the main menu, select the Configure HBA Settings option.
3	Set the following HBA parameters: ◆ <code>IP_ARP_Redirect=on</code> ◆ <code>KeepAliveTO=30</code>
4	Select Commit Changes for this HBA so that the changes are applied to the HBA ports.
5	Repeat the settings for each HBA port that is listed.
6	Return to the main menu and select Quit.

Note

You can also use `iscli -n` command to set the HBA parameters directly.

Example:

```
iscli -n <port inst no> IP_ARP_Redirect on
iscli -n <port inst no> KeepAliveTO 30
```

`<port inst no>` is the port instance number which can be obtained through `iscli -i` command.

Loading the bundled HBA driver on the host

Starting iSCSI

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

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

Step	Action
1	Enter the following command on the Linux host for Red Hat Enterprise Linux 5 Update 3: <code>modprobe -v qla4xxx</code>

Verifying the HBA settings

Verifying the QLogic HBA settings

To verify the ql4xkeepalive parameter, complete the following steps.

Step	Action
1	Check the value for the parameter in the corresponding qla4xxx directory in the sysfs, /sys/module/qla4xxx/ or /sys/module/qla4xxx/parameters/. Example: You can check the value of ql4xkeepalive on a Red Hat Enterprise Linux host by running the following command: <pre>cat /sys/module/qla4xxx/parameters/ql4xkeepalive</pre>
2	To verify IP_ARP_Redirect (which need to be manually set), run the -c option of the <code>iscli</code> command: <pre>/opt/QLogic_Corporation/SANsurferCLI/iscli -c</pre>

Recording the initiator nodename of the HBAs

**About initiator
nodename**

The initiator node name is required to create igroups on the storage system. You map igroups to specific LUNs. Only the hosts in the igroup can discover the LUNs as local devices.

**Obtaining the
initiator nodename
for QLogic HBAs**

To obtain the initiator nodenames of QLogic HBAs by using the SANsurfer iSCSI HBA CLI client, complete the following steps.

Step	Action
1	Enter the following command on the Linux host console: iscli -i
2	Record the initiator nodename for each port listed.

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

Adding targets to the host

To add target details to the host, complete the following steps.

Step	Action
1	Run the following command on the host: iscsi
2	Add target details which includes the following: ◆ Target nodename ◆ Target IP address
3	Select Commit Changes for this HBA so that the changes are applied to the HBA ports.
4	Repeat the settings for each HBA port that is listed.

Unloading the bundled HBA driver

Before unloading the HBA driver

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

Unloading the HBA driver

To unload the HBA driver, complete the following steps.

Step	Action
1	Enter the appropriate command for QLogic HBA driver model: <code>modprobe -r qla4xxx</code>

Configuration overview

This chapter explains how to complete the installation and initial configuration of the iSCSI initiator on the Linux host. Complete the final configuration of the initiator after you have configured the storage system. To configure initiator settings, you require root privileges.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Configuring Red Hat Enterprise Linux”](#) on page 52
- ◆ [“Configuring SUSE Linux Enterprise Server”](#) on page 62

Configuring Red Hat Enterprise Linux

Verifying or installing the required iSCSI package

You need to verify that the correct iSCSI package is installed on your Linux host. To verify the iSCSI package, complete the following steps.

Step	Action
1	Enter the following command: <code>rpm -q iscsi-initiator-utils</code>
	Result: The <code>rpm -q</code> command returns the name and version of the iSCSI package. Install the correct version of the packages that are shipped along with the operating system.
2	If you plan to use DM-Multipath, see “Configuring DM-Multipath” on page 89.

Recording or changing the initiator node name on the host

The initiator node name is required to create igroups on the storage system. You map igroups to specific LUNs. Only the hosts in the igroup can discover the LUNs as local devices. The initiator node name is generated the first time the iSCSI service on the host is started. If you have never used the iSCSI service on the Linux host, create the initiator node name manually.

The default initiator node name uses the following format:

- ◆ In Red Hat Enterprise Linux 5 series:
`iqn.2005-03.com.RedHat:RandomNumber`
- ◆ In Red Hat Enterprise Linux 4 series Update 3 and later:
`iqn.1987-05.com.cisco:RandomNumber`

It is best to set the initiator node name as follows:

- ◆ For Red Hat Enterprise Linux 5 series:
`iqn.2005-03.com.RedHat: host_name`
- ◆ For Red Hat Enterprise Linux 4 series Update 3 and later:
`iqn.1987-05.com.cisco: host_name`

host_name is the host name of your Linux host.

You can change the initiator node name or use the default node name.

To view or change the default node name, complete the following steps.

Red Hat Enterprise Linux 5 series:

Step	Action	
1	With a text editor, open the host's /etc/iscsi/initiatorname.iscsi file.	
2	If...	Then...
	A default initiator name has not been generated	a. Replace the text in the host's /etc/iscsi/initiatorname.iscsi file with the following line: InitiatorName=iqn.2005-03.com.RedHat:RandomNumber b. Record the node name and proceed to “Editing the host's iSCSI configuration file” on page 56.
	You want to use the default initiator node name	Record the node name and proceed to “Editing the host's iSCSI configuration file” on page 56.

Step	Action	
	You want to change the initiator node name	a. Modify the RandomNumber part of the initiator node name. The following line shows an example node name: <pre>iqn.2005-03.com.RedHat:linux-host1</pre> b. Record the node name and proceed to “Editing the host’s iSCSI configuration file” on page 56. Note _____ The initiator node name should be unique. _____

Red Hat Enterprise Linux 4 series Update 3 and later:

Step	Action	
1	With a text editor, open the host's /etc/initiatorname.iscsi file.	
2	If...	Then...
	A default initiator name has not been generated	a. Replace the text in the host's /etc/initiatorname.iscsi file with the following line: <pre>iqn.1987-05.com.cisco:RandomNumber</pre> b. Record the node name and proceed to “Editing the host's iSCSI configuration file” on page 56.
	You want to use the default initiator node name	Record the node name and proceed to “ Editing the host's iSCSI configuration file ” on page 56.
	You want to change the initiator node name	a. Modify the <i>RandomNumber</i> part of the initiator node name. The following line shows an example node name. <pre>iqn.1987-05.com.cisco:linux-host1</pre> b. Record the node name and proceed to “Editing the host's iSCSI configuration file” on page 56. Note _____ The initiator node name should be unique. _____

Node name rules: If you change the host's initiator node name, ensure that the new name follows these rules:

- ◆ A node name can be up to 223 bytes.
- ◆ Uppercase characters are always mapped to lowercase characters.
- ◆ A node name can contain alphabetic characters (a to z), numbers (0 to 9) and three special characters:
 - ❖ Period (".")
 - ❖ Hyphen ("-")
 - ❖ Colon (":")
- ◆ The underscore character ("_") is *not* supported.

Editing the host's iSCSI configuration file

For configuring iSCSI, you should edit host's iSCSI configuration file.

Configuring iSCSI to use DM-Multipath:

To configure iSCSI to use DM-Multipath for Red Hat Enterprise Linux 5 series, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi/iscsid.conf</code> file.
2	If you are using DM-Multipath, set the value of <code>node.session.timeo.replacement_timeout</code> to 5.

To configure DM-Multipath for Red Hat Enterprise Linux 4 series Update 3 and later, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi.conf</code> file.
2	If you are using DM-Multipath: ◆ Remove the comment of the <code>ConnFailTimeout</code> line in the Session Timeout Settings section ◆ Set the value of <code>ConnFailTimeout</code> to 5

Setting up CHAP authentication:

Note

Setting up CHAP authentication is optional.

To set up CHAP for Red Hat Enterprise Linux 5 series, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi/iscsid.conf</code> file.
2	Add CHAP user names and passwords, as shown in the following steps: a. To enable CHAP authentication, set <code>node.session.auth.authmethod</code> to CHAP. The default is None. <code>node.session.auth.authmethod = CHAP</code> b. To set a CHAP user name and password for initiator authentication by the target or targets, remove the comment from the following lines: <code>node.session.auth.username = username</code> <code>node.session.auth.password = password</code> c. To set a CHAP user name and password for target authentication by the initiator, remove the comment from the following lines: <code>node.session.auth.username_in = username_in</code> <code>node.session.auth.password_in = password_in</code>

Step	Action
	d. For a successful session discovery, you need to enable the discovery CHAP authentication. <pre>discovery.sendtargets.auth.authmethod = CHAP discovery.sendtargets.auth.username = username discovery.sendtargets.auth.password = password discovery.sendtargets.auth.username_in = username_in discovery.sendtargets.auth.password_in = password_in</pre> Note _____ The username and password should be same for both session and discovery. Ensure that you use the same user names and passwords when you set up CHAP on the storage system with the <code>iscsi security</code> command.

To set up CHAP for Red Hat Enterprise Linux 4 series Update 3 and later, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi.conf</code> file.
2	Add CHAP user names and passwords, as shown in the following example. CHAP settings should be indented below the Discovery Address of the storage system they apply to. ◆ For bidirectional CHAP, you should define outgoing and incoming user names and passwords. ◆ If you do not want to use bidirectional CHAP, define only the <code>OutgoingUsername</code> and <code>OutgoingPassword</code>.

Step	Action
	Example: <pre>DiscoveryAddress=192.168.10.20 OutgoingUsername=username_out OutgoingPassword=password_out IncomingUsername=username_in IncomingPassword=password_in</pre> Ensure that you use the same user names and passwords when you set up CHAP on the storage system with the <code>iscsi security</code> command. ◆ Use the <code>OutgoingUsername</code> and <code>OutgoingPassword</code> for the storage system's inbound user name and password (<i>inname</i> and <i>inpassword</i>). ◆ If you are using bidirectional authentication, use <code>IncomingUsername</code> and <code>IncomingPassword</code> for the storage system's outbound user name and password (<i>outname</i> and <i>outpassword</i>) for incoming settings.

Note

If you want to configure global CHAP, that is the same user name and password for all the targets, ensure that the CHAP settings are mentioned before the Discovery Address.

Setting up target discovery:

To set up target discovery for Red Hat Enterprise Linux 5 series, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi/iscsid.conf</code> file.
2	To discover the target, use the following command: <pre>iscsiadm -m discovery -t st -p <ip_address></pre> <code><ip_address></code> is the target address.
3	To set up the session start-up option of iSCSI as manual or automatic, use the following command: <pre>iscsiadm -m node -T <target_iqn> -p <ip_address> :3260 --op update -n node.startup -v <mode></pre> <pre>iscsiadm -m node -T <target_iqn> -p <ip_address> :3260 -- op update -n node.conn[0].startup -v <mode></pre> <code><target_iqn></code> is the iqn number of the target. <code><ip_address></code> is the target address. <code><mode></code> is manual or automatic. Alternatively, You can set up the session start-up by editing the <code>/etc/iscsi/iscsid.conf</code> file before target discovery. The default setting is <code>node.startup=automatic</code>.

To set up target discovery for Red Hat Enterprise Linux 4 series Update 3 and later, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi.conf</code> file.

Step	Action
2	Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you used for iSCSI LUNs: <pre>DiscoveryAddress=<storage_system_IPaddress></pre> <storage_system_IPaddress> is the IP address of an Ethernet interface on the storage system. Specify an interface that is used for iSCSI communication. It is best to use Gigabit Ethernet interfaces. Example: <pre>DiscoveryAddress=192.168.10.100</pre> <pre>DiscoveryAddress=192.168.10.20</pre>

Configuring SUSE Linux Enterprise Server

Verifying or installing the required iSCSI packages

You need to verify that the correct iSCSI package are installed on your Linux host.

To verify the iSCSI package, complete the following steps.

Step	Action	
1	Enter the following command:	
	If you are using...	Then run...
	SUSE Linux Enterprise Server 9 series	<code>rpm -q linux-iscsi</code>
	SUSE Linux Enterprise Server 10 series	<code>rpm -q open-iscsi</code>
	Result: The <code>rpm -q</code> command returns the name and version of the iSCSI package. Install the correct version of the packages that are shipped along with the operating system.	
2	If you plan to use DM-Multipath, see “ Configuring DM-Multipath ” on page 89.	

Recording or changing the initiator node name on the host

Initiator node name: The initiator node name is required to create initiator groups (igroups) on the storage system. You map LUNs to specific igroups. Only the hosts in the igroup can discover the LUNs as local devices.

The default initiator node name for SUSE Linux Enterprise Server 9 series and SUSE Linux Enterprise Server 10 series are as follows:

- ◆ SUSE Linux Enterprise Server 9 series:
`iqn.1987-05.com.cisco:RandomNumber`
- ◆ SUSE Linux Enterprise Server 10 series:
`iqn.1996-04.de.suse:RandomNumber`

You should set the initiator node name as follows:

- ◆ SUSE Linux Enterprise Server 9 series:
`iqn.1987-05.com.cisco:host_name`
- ◆ SUSE Linux Enterprise Server 10 series:
`iqn.1996-04.de.suse:host_name`
host_name is the name of your Linux host.

You can change the initiator node name or use the default node name.

Viewing or changing the initiator node name: To view or change the default initiator node name, complete the following steps.

Step	Action	
1	With a text editor, open the host's <code>/etc/initiatorname.iscsi</code> file.	
2	If...	Then...
	A default initiator name has not been generated for SUSE Linux Enterprise Server 9 series	a. Replace the text in the host's <code>/etc/initiatorname.iscsi</code> file with the initiator name in the following format: <code>InitiatorName=iqn.1987-05.com.cisco:RandomNumber</code> b. Record the initiator node name and proceed to “Editing the host's iSCSI configuration file” on page 56.
	A default initiator name has not been generated for SUSE Linux Enterprise Server 10 series	a. Replace the text in the host's <code>/etc/iscsi/initiatorname.iscsi</code> file with the initiator name in the following format: <code>InitiatorName=iqn.1996-04.de.suse:RandomNumber</code> b. Record the initiator node name and proceed to “Editing the host's iSCSI configuration file” on page 56.

Step	Action	
	You want to use the default initiator node name	Record the initiator node name and proceed to “ Editing the host’s iSCSI configuration file ” on page 56.
	You want to change the initiator node name for SUSE Linux Enterprise Server 9 series	a. Modify the <i>RandomNumber</i> part of the initiator node name. The following is a sample node name: iqn.1987-05.com.cisco:01.82ea2060f487 b. Record the initiator node name and proceed to “Editing the host’s iSCSI configuration file” on page 56.
	You want to change the initiator node name for SUSE Linux Enterprise Server 10 series	a. Modify the <i>RandomNumber</i> part of the initiator node name. The following is a sample node name: iqn.1996-04.de.suse:lnx.118 b. Record the initiator node name and proceed to “Editing the host’s iSCSI configuration file” on page 56.

Configuring the iSCSI initiator

SUSE Linux Enterprise Server 9 series: SUSE Linux includes the iSCSI initiator software in the software distribution. The initiator's components include a kernel module, which is already compiled into the Linux kernel, and the user space packages.

Note **iSCSI session timeout values**

`ConnFailTimeout` specifies the length of time to wait for session reestablishment before failing SCSI commands to a target device. If the value is set to a default value 0, then there is no timeout.

Example: `ConnFailTimeout=0`

If multipath is enabled, then you can set `ConnFailTimeout` to a lower value, such as 10 or 30. In other cases, the default value can be retained.

`MinDiskCommandTimeout` specifies the minimum length of time to wait before failing SCSI commands to a target device. The default value is -1, which disables the command timeout, allowing commands to be retried forever.

Example: `MinDiskCommandTimeout=-1`

`MaxDiskCommandTimeout` specifies the maximum length of time to wait before failing SCSI commands to a target device.

To configure the iSCSI initiator on SUSE Linux Enterprise Server 9 series, edit the `/etc/iscsi.conf` file.

You can edit the file to configure the following:

- ◆ Required settings
- ◆ The storage system's IP address
- ◆ User names and passwords for CHAP
- ◆ Multipath support

To edit the `/etc/iscsi.conf` file, complete the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi.conf</code> file.
2	At the beginning of the file, add the following configuration entries. Ensure that you add these configuration entries before any other configuration settings. <pre>Continuous=no HeaderDigest=never DataDigest=never</pre> Attention _____ These configuration entries are required for proper operation. If you do not add these entries at the beginning of the file, you might experience problems with connectivity and performance. _____
3	If you are using DM-Multipath, perform the following actions: ◆ Remove the comment to enable <code>Multipath=portal</code> in the Multipath Settings section. ◆ Remove the comment from <code>ConnFailTimeout</code> in the Session Timeout Settings section and set the value to <code>ConnFailTimeout=5</code>.
4	To discover the storage systems, add the IP address of the storage system by using the following format: <pre>DiscoveryAddress=storage_system_IPaddress</pre> <i>storage_system_IPaddress</i> is the IP address of an Ethernet interface on the storage system, preferably a Gigabit Ethernet interface. Specify an IP address that is used for iSCSI communication. Example: The following are sample <code>DiscoveryAddress</code> entries that is using the IP addresses of the storage systems: <pre>DiscoveryAddress=192.168.10.100 DiscoveryAddress=10.61.208.200</pre>

Step	Action
5	Add CHAP user names and passwords, as shown in the following example. CHAP settings should be indented below the Discovery Address of the storage system they apply to. ◆ For bidirectional CHAP, you should define outgoing and incoming user names and passwords. ◆ If you do not want to use bidirectional CHAP, define only the outgoing user name and outgoing password. When adding the global CHAP entries in <code>/etc/iscsi.conf</code> file, the sequence should be as follows: ◆ Outgoing CHAP settings ◆ Discovery Addresses ◆ Incoming CHAP settings (in case of bidirectional CHAP) Example: <pre>DiscoveryAddress=192.168.10.20 OutgoingUsername=username_out OutgoingPassword=password_out IncomingUsername=username_in IncomingPassword=password_in</pre> Ensure that you use the same user names and passwords when you set up CHAP on the storage system with the <code>iscsi security</code> command. ◆ Use the outgoing user name and outgoing password for the storage system's inbound user name and password (<i>inname</i> and <i>inpassword</i>). ◆ If you are using bidirectional authentication, use incoming user name and incoming password for the storage system's outbound user name and password (<i>outname</i> and <i>outpassword</i>).

SUSE Linux Enterprise Server 10 series: SUSE Linux includes the iSCSI initiator software in the software distribution. The initiator's components include a kernel module, which is already compiled into the Linux kernel, and the user space packages.

The configuration of the iSCSI initiator on SUSE Linux Enterprise Server 10 series involves two steps:

- ◆ Discovering the available iSCSI targets
- ◆ Setting up an iSCSI session

Both the discovery and configuration of iSCSI connections require a running `iscsid` daemon. During first-time discovery, the internal database of the iSCSI initiator is created in the `/etc/iscsi/send_targets/` and `/etc/iscsi/nodes/` directories.

The default location of the iSCSI database files is `/etc/iscsi`.

iSCSI parameter settings: To set the iSCSI parameter, complete the following steps.

Step	Action
1	Open <code>/etc/iscsi/iscsid.conf</code> file with a text editor.
2	Change the following iSCSI parameter for SUSE Linux Enterprise Server 10 series: <code>node.session.timeo.replacement_timeout = 5</code>

You should use the YaST2 and `iscsiadm` utilities for the discovery and the configuration of iSCSI connections.

For more information about using YaST2, see the SUSE Linux Enterprise Server 10 series documentation available at:

<http://www.novell.com>

Configuring the iSCSI initiator using YaST2: YaST2 has three tabs—Services, Connected Targets, and Discovered Targets. You can use the Services tab to enable the iSCSI initiator at boot time. Connected Targets provides you with an overview of the iSCSI targets that are currently connected. Connected Targets and Discovered Targets enable you to add new targets to the system. Discovered Targets also enables you to discover the iSCSI targets in the network.

To configure the iSCSI initiator using YaST2, complete the following steps.

Step	Action
1	Enter the following command: yast2
2	Click Network Services.
3	Click iSCSI Initiator.
4	Select the Discovered Targets tab.
5	Click Discovery.
6	Enter the IP address of the iSCSI target, and change the port if necessary.
7	If you have an authentication mechanism in place, enter the credentials.
8	Click Next to start the discovery.
9	After a successful discovery, use Login to activate the target.
10	Enter the authentication credentials to use the selected iSCSI target.
11	Click Next to finish the configuration. The target now appears in Connected Targets.
12	Click the Connected Targets tab. Change the start-up option to Manual or Automatic, depending on your requirement, by using the Toggle Start-Up button for all the discovered targets.

Configuring the iSCSI initiator using iscsiadm: The `iscsiadm` utility is a command-line tool that is used to manage (update, delete, insert, and query) the persistent database. The utility enables you to perform a set of operations on iSCSI nodes, sessions, connections, and discovery records.

For more information, see the `iscsiadm` manual (`man`) pages.

To discover the available iSCSI targets, complete the following steps.

Step	Action
1	Enter the following command to initiate discovery: <pre>iscsiadm --mode discovery --op update --type sendtargets --portal <i>targetip</i></pre> <i>targetip</i> is the IP address of the target. Result: The target <i>targetip</i> is discovered. For each target defined on the iSCSI target, a separate line is displayed. All received values are stored in an internal persistent database.
2	Enter the following command to create all needed devices: <pre>iscsiadm -m node -l all</pre> Result: The initiator logs in to all the discovered nodes as maintained in the iSCSI database.
3	If your discovery is password-protected, you need to provide the authentication information in the iscsid.conf file. For a successful session discovery, you need to enable the configuration entries of the discovery CHAP authentication. <pre>discovery.sendtargets.auth.authmethod = CHAP discovery.sendtargets.auth.username = <i>username</i> discovery.sendtargets.auth.password = <i>password</i> discovery.sendtargets.auth.username_in = <i>username_in</i> discovery.sendtargets.auth.password_in = <i>password_in</i></pre>
4	Enter the following command to display all the active iSCSI sessions: <pre>iscsiadm --mode session</pre> Result: All the established iSCSI sessions are displayed. Example: The following is a sample output: <pre>tcp [022e71] 10.72.200.37:3260,1000 iqn.1992-08.com.netapp:sn.101176399</pre> The ID of the target is 022e71. This ID is used to access the target.

Step	Action
5	To automate logging in to a node, use the following commands with the record ID of the node discovered: <pre>iscsiadm -m node -r recordId -o update -n node.startup -v automatic</pre> Alternatively, to set the <code>node.startup</code> attribute to <code>startup</code> as default for all sessions, add the following to <code>/etc/iscsi/iscsid.conf</code> and perform a rediscovery: <pre>node.startup = automatic</pre>

If you want to use CHAP, add CHAP user names and passwords, as described in the following steps.

Step	Action
1	With a text editor, open the host's <code>/etc/iscsi/iscsid.conf</code> file.
2	To enable CHAP authentication, set <code>node.session.auth.authmethod</code> to CHAP. The default is None. <pre>node.session.auth.authmethod = CHAP</pre>
3	To set a CHAP user name and password for initiator authentication by the target or targets, remove the comment from the following configuration entries: <pre>node.session.auth.username = <i>username</i> node.session.auth.password = <i>password</i></pre>
4	To set a CHAP user name and password for target or targets authentication by the initiator, remove the comment from the following configuration entries: <pre>node.session.auth.username_in = <i>username_in</i> node.session.auth.password_in = <i>password_in</i></pre>

Note

Ensure that you use the same user names and passwords when you set up CHAP on the storage system with the `iscsi` security command.

Login mechanism: You can log in manually or set up automatic login.

To manually log in to the iSCSI service, complete the following steps.

If you want to set the login mode to this...	Then do this...
Manual on a single portal on a target	Run the following command: iscsiadm -m node -T targetname -p ip:port --op update -n node.startup -v manual
Manual on all portals on a target	Run the following command: iscsiadm -m node -T targetname --op update -n node.startup -v manual
Manual mode for all the targets	Add the following to the /etc/iscsi/iscsid.conf file: node.startup = manual

Setting the login mode to manual affects only nodes that are discovered after setting the value. To log in to all the manual nodes, restart the iSCSI service:

```
/etc/init.d/open-iscsi restart
```

During the next start-up, you should manually log in to the nodes.

To automatically log in to the iSCSI service, complete the following steps.

If want to set the login mode to this...	Then do this...
Automatic on a single portal on a target	Run the following command: iscsiadm -m node -T targetname -p ip:port --op update -n node.startup -v automatic
Automatic on all portals on a target	Run the following command: iscsiadm -m node -T targetname --op update -n node.startup -v automatic

If want to set the login mode to this...	Then do this...
Automatic mode for all the targets	Add the following to the /etc/iscsi/iscsid.conf file: node.startup = automatic

Setting the login mode to automatic affects only nodes that are discovered after setting the value. To log in to all the manual nodes, restart the iSCSI service:

```
/etc/init.d/open-iscsi restart
```

During the next start-up, you are logged in to the nodes automatically.

Configuring the Storage System

Before you begin

Before you configure the storage system, verify that you have the following:

- ◆ **FC:** You need the host's WWPNs that you recorded in "[Installing the HBAs for FC](#)" on page 21.

You can use the WWPNs when creating igroups on the storage system.

- ◆ **Software iSCSI:** You need the host's software iSCSI initiator node name that you recorded in "[Configuring Red Hat Enterprise Linux](#)" on page 52 and "[Configuring SUSE Linux Enterprise Server](#)" on page 62.

You can use the initiator node name when creating igroups on the storage system.

- ◆ **Hardware iSCSI:** You need the host's initiator name that you recorded in "[Installing the HBAs for Hardware iSCSI](#)" on page 41.

You can use the initiator names when creating igroups on the storage system.

Topics in this chapter

This chapter includes the following topics:

- ◆ "[Configuring the storage system with FC](#)" on page 76
- ◆ "[Configuring the storage system with hardware and software iSCSI](#)" on page 83

Configuring the storage system with FC

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

- ◆ Ensure that the FC service is licensed and started.
- ◆ For storage system configurations, verify that the storage system is running a supported cfmode by running the `fcv show cfmode` command. The support matrix lists the supported cfmode values for your configuration. If you need to change the cfmode, see “[Understanding single_image mode](#)” on page 78. For standby and mixed cfmode, special FC switch zoning is required. See “[Zoning with standby or mixed cfmode](#)” on page 86.
- ◆ 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. Ensure that you specify the LUN type and FC igroup type as `linux`.
- ◆ **Optional:** Enable ALUA, if required.
- ◆ If you are using clustered Linux hosts, the igroup should contain the WWPNs of all the hosts in the cluster that need access to the mapped LUN.

For the detailed configuration steps

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

Understanding cfmode

What cfmode is: Cluster failover mode (cfmode) is a storage system FC failover mode setting which controls how the storage appliance handles FC before and after cluster failover events. The cfmode is not available and not relevant on nonclustered 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 storage controller affects the following items:

- ◆ Switch configuration and cabling

In mixed and dual fabric modes, the target ports log in to the fabric in loop mode. In `single_image` mode, the target ports log in to 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 continues 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.

For the detailed information about the cabling configurations supported, and examples for systems with onboard FC ports, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

- ◆ Zoning

Updating the storage system to `single_image` cfmode changes the WWNN of the storage system and the WWPNNs of the target ports on each controller. If the storage system connects to a switch that is zoned by WWPNN, you should check the zoning setup and make the required updates to your switch. Switches that are zoned by physical port are not affected.

- ◆ 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 should reconfigure every host that is connected to the storage system to discover the new paths. The LUNs might be accessible but should not be used until you reconfigure the host.

- ◆ Multipathing policy

Changing the cfmode setting affects the number of paths available to the host. You should reconfigure the multipath devices, so that the new changes take effect.

Prerequisites for changing cfmode: Complete the following tasks before you change the cfmode.

- ◆ **Schedule downtime for the storage system:** Hosts cannot access LUNs when you are changing cfmode.

You should allow enough downtime to change the cfmode setting, and reconfigure the host. Enter commands that quiesce host I/O and take host applications offline.

- ◆ **Identify affected hosts and initiators:** You should identify any hosts and initiators that access LUNs on the systems you are upgrading.

Enter the following command:

```
igroup show
```

- ◆ **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 should 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.

Understanding single_image mode

While configuring storage system for FC, you can migrate to single_image mode while installing the new 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 FC cables to the Linux host HBAs.

What single_image mode is: In single_image mode, both controllers in the storage system configuration share the same WWNN. single_image mode provides the following advantages:

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

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

- ◆ To increase the number of paths to the storage system
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, ensure that the following mapping conflicts do not occur:

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

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

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

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

1. Prepare the host
2. Upgrade the storage system to `single_image` mode
3. Validate zoning and cabling
4. Reconfigure the host

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

Step	Action
1	Stop host I/O and 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 FC HBAs.
3	Proceed to “Migrating to single_image mode” on page 80.

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

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

Step	Action
8	If you are using switch zoning by WWPN, you need the WWPNs of the target ports to reconfigure zones on your switch. To obtain the WWPNs, enter the following command on each controller in the active/active configuration: <pre>fcp config</pre> For more information about zoning, see “Validating zoning and connectivity” on page 81.
9	Return to administrative privileges. On each storage system, enter the following command: <pre>priv set admin</pre>

Validating zoning and connectivity:

- ◆ Changes to WWNNs and WWPNs

After you change the cfmode setting to single_image, both controllers share a single WWNN. (One controller in the storage system 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 should 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 best to change the cfmode setting, ensure that each HBA has one primary and at least one secondary path to each LUN. For more information, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

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

Configuring the storage system with hardware and software iSCSI

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

- ◆ Ensure that the iSCSI service is licensed and started.
- ◆ A Linux host running without multipathing support can access more than one iSCSI portal (typically an Ethernet interface) on the storage system.
The host creates multiple sessions to the storage system and accesses each LUN as one SCSI device for each path. For more information, see [“Restricting access to a single iSCSI portal”](#) on page 83.
- ◆ If you want to use CHAP authentication, use the `iscsi security` command or the FilerView® interface to configure a CHAP user name and password.
- ◆ Create LUNs, create an igroup using the host’s initiator names, 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. Ensure that you specify the LUN type and iSCSI igroup type as `linux`.
- ◆ If you are using clustered Linux hosts, the igroup should contain the initiator names of all hosts in the cluster that need access to the mapped LUN.

For the detailed configuration steps

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

Restricting access to a single iSCSI portal

The iSCSI initiator creates a separate iSCSI session to each iSCSI portal that it discovers on the storage system. An iSCSI portal is a physical or logical network device on the storage system, such as an Ethernet interface or an entity VLAN which is enabled for iSCSI.

If you do not want to use DM-Multipath, you have several options for restricting access, depending on whether other iSCSI hosts need to access the storage system:

- ◆ Enable iSCSI on only one storage system interface.
- ◆ Create vFiler™ units, each with only one iSCSI-enabled interface.
Each vFiler unit acts as a separate iSCSI target.
- ◆ Enable iSCSI login from the host to only one portal on the controller.

The details of each option are as follows:

- ◆ Enable iSCSI on only one interface

If the storage system has only a few iSCSI hosts, a simple solution is to disable iSCSI traffic on all but one interface. Use the `iscsi interface disable` command (`iswt interface disable` for versions earlier than Data ONTAP 7.1) or the Web-based FilerView interface to enable and disable interfaces.

- ◆ Create vFiler units

For more information, see *Data ONTAP MultiStore Management Guide*.

- ◆ Enable iSCSI login on the controller

If you do not want to disturb the storage system configuration and want to restrict access through only one path, enable iSCSI login from the host to only one portal on the controller. This can be achieved by setting one path to automatic and others to manual.

Configuring FC Cabling and Switches

Fabric or direct connections

You can connect your host to the storage system by using either FC 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.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Switched fabric configurations](#)” on page 86
- ◆ “[Direct-connected configurations](#)” on page 87

Switched fabric configurations

High availability

To avoid single points of failure, use two physically independent FC switch fabrics to connect your Linux host to a storage system configuration.

Configuration and zoning details

For details about supported FC configurations and switch zoning, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Use default HBA and target topology

For a fabric-connected configuration, the default autonegotiated HBA and storage system target FC topology works correctly. Ensure that you have set the HBA and driver parameters to the values specified in “[Setting the required HBA and external driver parameters](#)” on page 31.

Single initiator zoning

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

Zoning with standby or mixed cfmode

If a storage system configuration is using the standby or mixed cfmode, your switches should 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 the zones, so the host can access them.

Connect cables

Connect the storage system cables to the FC switches.

Direct-connected configurations

High availability

To avoid single points of failure, use two FC cables to connect your Linux host to both the controllers of a storage system configuration.

Use default HBA and target topology

For a direct-connected configuration, the default autonegotiated HBA and storage system target FC topology works correctly. Ensure that you have set the HBA and driver parameters to the values specified in “[Setting the required HBA and external driver parameters](#)” on page 31.

Single_image cfmode required

A direct FC connection to a storage system requires that the storage system should be using single_image cfmode.

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

```
fcp show cfmode
```

For information about changing to single_image cfmode, see “[Understanding single_image mode](#)” on page 78.

Connect cables

Connect the FC cables from the host to the target ports on the storage system.

This chapter explains how to configure DM-Multipath.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Verifying the required multipathing packages](#)” on page 90
- ◆ “[Editing the configuration file](#)” on page 91
- ◆ “[Starting DM-Multipath](#)” on page 98
- ◆ “[Verifying the configuration](#)” on page 100
- ◆ “[Stopping DM-Multipath](#)” on page 102

Verifying the required multipathing packages

To verify that you have installed the correct multipathing package, complete the following steps.

Step	Action	
1	To verify DM-Multipath packages, enter one of the following commands:	
	If you are using...	Then run the following commands...
	SUSE Linux Enterprise Server	<code>rpm -q device-mapper</code> <code>rpm -q multipath-tools</code>
	Red Hat Enterprise Linux	<code>rpm -q device-mapper</code> <code>rpm -q device-mapper-multipath</code>
	Result: The <code>rpm -q</code> commands return the name and version of the installed packages.	
2	If you do not have the required packages, install the appropriate packages based on your version of Linux. For more information about obtaining RPMs, see “ Obtaining RPMs ”.	

Obtaining RPMs: You can get the RPMs either from the Red Hat Enterprise Linux media or from the Red Hat Network at <http://www.redhat.com/software/rhn/> (subscription required).

Copying the file: The `multipath.conf` file contains configuration settings for multipath. A sample file is copied to your host when you install the multipath-tools packages:
`/usr/share/doc/packages/multipath-tools/multipath.conf.annotated`.

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

Editing the configuration file

To edit DM-Multipath configuration file, complete the following steps.

Step	Action	
1	Check if the /etc/multipath.conf file exists.	
	If...	Then...
	The file exists	a. Open the file. b. Comment out all the lines. c. Append the sample specified in “Red Hat Enterprise Linux 5 Update 3” on page 93, or see the note on “page 94”, “Red Hat Enterprise Linux 4 Update 7” on page 94, “SUSE Linux Enterprise Server 10 SP2 and later” on page 96, or “SUSE Linux Enterprise Server 9 SP3” on page 97, depending on your version of the operating system.
	The file does not exist	a. Create a new file. b. Append the sample specified in “Red Hat Enterprise Linux 5 Update 3” on page 93, or see the note on “page 94”, “Red Hat Enterprise Linux 4 Update 7” on page 94, “SUSE Linux Enterprise Server 10 SP2 and later” on page 96, or “SUSE Linux Enterprise Server 9 SP3” on page 97, depending on your version of the operating system.

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

Step	Action	
3	If you are using ALUA, specify the ALUA callout program, that is:	
	If you are using...	Then...
	Red Hat Enterprise Linux	/sbin/mpath_prio_alua program, instead of /sbin/mpath_prio_ontap
	SUSE Linux Enterprise Server	prio “alua” program, instead of prio “ontap”
	ALUA callout programs are available in the Red Hat Enterprise Linux distribution. Note _____ ALUA is supported from SUSE Linux Enterprise Server 10 SP2, Red Hat Enterprise Linux 5 Update 2 and later. _____	
4	Save the changes.	

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

Red Hat Enterprise Linux 5 Update 3:

```
defaults
{
    user_friendly_names    yes
    max_fds                max
    queue_without_daemon   no
}
blacklist
{
    wwid <DevId>
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]]*"
}
devices
{
    device
    {
        vendor            "NETAPP"
        product            "LUN"
        getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
```

```

        prio_callout          "/sbin/mpath_prio_ontap /dev/%n"
        features              "1 queue_if_no_path"
        hardware_handler      "0"
        path_grouping_policy  group_by_prio
        fallback              immediate
        rr_weight              uniform
        rr_min_io             128
        path_checker          directio
        flush_on_last_del     yes
    }
}

```

<DevID> is the unique ID. To generate WWID, see [Step 2](#) of “[Editing the configuration file](#)” on page 91.

Note

If you are configuring DM-Multipath on an iSCSI host, set `path_grouping_policy` to `multibus`.

`flush_on_last_del` is available with Red Hat Enterprise Linux 5 Update 3. For Red Hat Enterprise Linux 5 Update 2 and earlier, `flush_on_last_del` is not supported.

`directio` pathchecker is available with Red Hat Enterprise Linux 5 Update 1 and later. Set the `max_fds` value to 4096 on Red Hat Enterprise Linux 5 Update 2 host. For Red Hat Enterprise Linux 5.0, use `readsector0` pathchecker.

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

Red Hat Enterprise Linux 4 Update 7:

```

defaults
{
    user_friendly_names    yes
    rr_min_io              128
    max_fds                4096
}
devnode_blacklist
{
    wwid <DevId>
    devnode "^ (ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] * "
    devnode "^hd[a-z] "
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*] "
}
devices

```

```

{
    device
    {
        vendor                "NETAPP"
        product                "LUN"
        flush_on_last_del      yes
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout           "/sbin/mpath_prio_ontap /dev/%n"
        features               "1 queue_if_no_path"
        hardware_handler        "0"
        path_grouping_policy    group_by_prio
        failback                immediate
        rr_weight               uniform
        path_checker            directio
    }
}

```

<DevID> is the unique ID. To generate WWID, see [Step 2](#) of “[Editing the configuration file](#)” on page 91.

Note

If you are configuring DM-Multipath on an iSCSI host, set `path_grouping_policy` to `multibus`.

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

`flush_on_last_del` is available with Red Hat Enterprise Linux 4 Update 7 and later. For Red Hat Enterprise Linux 4 Update 6 and earlier, `flush_on_last_del` is not supported.

`max_fds` is available from Red Hat Enterprise Linux 4 Update 7. For Red Hat Enterprise Linux 4 Update 6 and earlier, `max_fds` is not supported.

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

SUSE Linux Enterprise Server 10 SP2 and later:

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

<DevID> is the unique ID. To generate WWID, see [Step 2](#) of “[Editing the configuration file](#)” on page 91.

Note

If you are configuring DM-Multipath on an iSCSI host, set `path_grouping_policy` to `multibus`.

You should use `prio_callout "/sbin/mpath_prio_ontap /dev/%n"` for SUSE Linux Enterprise Server 10 SP1, instead of `prio "ontap"` in the above example.

SUSE Linux Enterprise Server 9 SP3:

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

<DevID> is the unique ID. To generate WWID, see [Step 2](#) of “[Editing the configuration file](#)” on page 91.

Starting DM-Multipath

You have to start the DM-Multipath to configure LUNs with DM-Multipath.

To start the DM-Multipath on the Linux host, complete the following steps.

Step	Action	
1.	To start the multipath daemon, enter the following command:	
	If you are using...	Then...
	Red Hat Enterprise Linux	# /etc/init.d/multipathd start
	SUSE Linux Enterprise Server	# /etc/init.d/boot.multipath start # /etc/init.d/multipathd start
2.	To configure the DM-Multipath devices, run the following command: # multipath	

Configuring the DM-Multipath to start automatically while booting

To configure DM-Multipath to start automatically while booting, complete the following step.

Step	Action	
1.	To add the multipath service to the boot sequence, enter the following commands on the Linux console:	
	If you are using...	Then...
	Red Hat Enterprise Linux	<code>chkconfig --add multipathd</code> <code>chkconfig multipathd on</code>
	SUSE Linux Enterprise Server	<code>insserv boot.multipath multipathd</code> <code>chkconfig --add boot.multipath</code> <code>chkconfig --add multipathd</code> <code>chkconfig boot.multipath on</code> <code>chkconfig multipathd on</code>

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, you can enter **`multipath -h`**.

To view all the configuration details, complete the following steps.

Step	Action
1	Enter the following command on the Linux host: <pre>multipath -v3 -d -ll</pre> Note _____ The <code>-d</code> (“dry run”) parameter prevents the command from actually updating the multipath maps. _____
2	Verify that the <code>multipathd</code> is running by issuing the following command: <pre>/etc/init.d/multipathd status</pre>
3	To view a list of multipath devices, including which <code>/dev/sdx</code> devices are used, enter the following command: <pre>multipath -ll</pre> Example: The output looks similar to the following for FC: <pre>3360a98000486e5372635a44646a505643NETAPP, 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]</pre>

Step	Action
	Example: The output looks similar to the following for software and hardware iSCSI: <pre># multipath -ll 360a9800043346773435a48514d624367 dm-0 NETAPP,LUN [size=7.0G][features=1 queue_if_no_path][hwhandler=0][rw] _ round-robin 0 [prio=4][active] _ 4:0:0:6 sdc 8:32 [active][ready] _ 3:0:0:6 sdd 8:48 [active][ready]</pre>
4	To view the /dev/mapper devices, enter the following command: <pre>ls -l /dev/mapper/</pre> Example: The following example shows the sample output of <code>ls -l dev/mapper</code> command. <pre>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</pre>

Stopping DM-Multipath

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

To stop DM-Multipath on the Linux host, complete the following steps.

Step	Action	
1	Unmount all file systems on the LUNs.	
2	Flush the DM-Multipath devices by running the following command: multipath -F	
3	Stop the multipath daemon by entering the following command:	
	If you are using...	Then...
	Red Hat Enterprise Linux	/etc/init.d/multipathd stop
	SUSE Linux Enterprise Server	/etc/init.d/multipathd stop /etc/init.d/boot.multipath stop

About this chapter

If you are using DM-Multipath, 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 multipathing software.

For information about the supported configuration, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Note

You should 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:

- ◆ “[Accessing LUNs with DM-Multipath](#)” on page 104
- ◆ “[Accessing LUNs without multipathing software](#)” on page 109

Accessing LUNs with DM-Multipath

To access LUNs with DM-Multipath, perform the following tasks.

1. View the list of LUNs. For more information, see “[Viewing the list of LUNs](#)” on page 105.
2. Create and mount the file system. For more information, see “[Creating and mounting a file system on a multipath device](#)” on page 107.

Accessing LUNs with DM-Multipath

Viewing the list of LUNs

Viewing list of LUNs To view information about LUNs, you can use the `sanlun` command.

To view a list of LUNs that are mapped to the Linux host, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show all</code>

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

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

To display multipath information for native DM-Multipath, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show -p</code>

Example: The following output shows when DM-Multipath is configured.

```
# 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

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, use the multipath device in the /dev/mapper/ directory.

Creating the file system

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

Step	Action
1	Enter the following command on the Linux host console: <pre>mkfs -t type /dev/mapper/device</pre> <i>type</i> is the file system type, such as <i>ext3</i>. <i>device</i> is the multipath device name of the LUN in the /dev/mpath/ directory. Example: The following command creates an ext3 file system on a LUN: <pre>mkfs -t ext3 /dev/mapper/360a980004334686267344374474d4f75</pre>

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.

Note Automatically mounting and unmounting the file systems on multipath devices are not supported in SUSE Linux Enterprise Server 9 series.

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

To mount the file system, complete the following steps.

Step	Action	
1	Create a mountpoint using the following command: mkdir mount_point <i>mount_point</i> is the name of the mountpoint you are creating.	
2	Open the /etc/fstab file with a text editor.	
3	Add an entry to /etc/fstab as follows: device mount_point type defaults 0 0 <i>device</i> is the name of the device in the /dev/mpath/ directory. <i>mount_point</i> is the mountpoint you created for the file system. <i>type</i> is the file system type, such as ext2 or ext3. Example: /dev/mapper/360a980004334686267344374474d4f75 /mnt/ssl/lun_0 ext3 defaults 0 0	
4	To mount the file system, enter the following command:	
	If you are using...	Then...
	Red Hat Enterprise Linux	mount mountpoint Example: mount /mnt/ssl/lun_0
	SUSE Linux Enterprise Server	mount -a
5	Do either of the following: ◆ Enter df to display the mounted file systems. Verify that the new file system is listed. ◆ Change to the mountpoint directory and try writing and reading a file to verify that the file system was mounted successfully.	

Accessing LUNs without multipathing software

Overview

Without multipathing software, the host cannot distinguish multiple LUNs from multiple paths to the same LUN.

If you are not using multipathing software, you should limit each LUN to a single path. You should use the standby cfmode to avoid having a secondary path visible on the storage system.

If you are using a switched FC fabric, you should zone your switch to prevent the host from having multiple paths to any one LUN. See “[Zoning with standby or mixed cfmode](#)” on page 86. 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 multipathing software.

To set the parameters for Emulex FC HBA, complete the following steps.

Step	Action	
1	Append <code>options lpfc lpfc_nodev_tmo=180</code> to the following file:	
	If you are using...	Then...
	Red Hat Enterprise Linux	/etc/modprobe.conf
	SUSE Linux Enterprise Server	/etc/modprobe.conf.local
2	Re-create initrd image using the following command and reboot the host:	
	If you are using...	Then...
	Red Hat Enterprise Linux	<code>mkinitrd -f /boot/initrd-"`uname -r`".img `uname -r`</code>
	SUSE Linux Enterprise Server	<code>mkinitrd</code>

Configurations without multipathing software can have additional restrictions on host OS versions and HBA models. See the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

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 should complete the following tasks.

1. Identify and view the list of LUNs. For more information, see “[Identifying and viewing information about LUNs](#)” on page 111.
2. Create a file system. For more information, see “[Creating a file system](#)” on page 113.
3. Create a label for the file system and add an entry for the file system’s label in `/etc/fstab`.

Then mount the file system. For more information, see “[Labeling and mounting the file system](#)” on page 114.

Identifying and viewing information about LUNs

Persistently identifying the 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 a 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 using the following commands:

- ◆ `sanlun`

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

- ◆ `lsscsi`

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

Using the sanlun command: The sanlun program is installed from Linux Host Utilities.

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

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

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

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

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

Creating a file system

About file systems

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

Using mkfs to create a file system

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

Step	Action
1	On the host console, enter the following command: <code>mkfs -t filesystem_type device</code> <i>filesystem_type</i> is the type of file system that you want to create. <i>device</i> is the partition.

Accessing LUNs without multipathing software

Labeling and mounting the file system

About labels

A file system label creates a persistent mapping between a LUN and a file system. For more information, see “[Viewing list of LUNs](#)” on page 105.

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

Using e2label to create a file system label

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

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

Mounting the file system

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

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

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

Overview of accessing LUNs

When the iSCSI service is started, it scans the storage systems and discovers all the mapped LUNs. If you configured the optional DM-Multipath, multipath devices are created for all the LUNs that are discovered.

Accessing multipath LUNs

If you are using DM-Multipath, you can access LUNs differently than if you are not using multipath. This chapter has two main sections, one that describes accessing LUNs with DM-Multipath and another that describes accessing LUNs without multipathing software.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Starting the iSCSI service on the host”](#) on page 118
- ◆ [“Configuring the iSCSI service to start automatically”](#) on page 119
- ◆ [“Accessing LUNs with DM-Multipath on Red Hat Enterprise Linux”](#) on page 120
- ◆ [“Accessing LUNs with DM-Multipath on SUSE Linux Enterprise Server”](#) on page 128
- ◆ [“Accessing LUNs without multipathing software on Red Hat Enterprise Linux”](#) on page 139
- ◆ [“Accessing LUNs without multipathing software on SUSE Linux Enterprise Server”](#) on page 150

Starting the iSCSI service on the host

Starting iSCSI

To start the iSCSI service, complete the following step.

Step	Action	
1	Start the iSCSI service by entering the following command at the Linux host command prompt:	
	If you are using...	Then...
	Red Hat Enterprise Linux	<code>/etc/init.d/iscsi start</code>
	SUSE Linux Enterprise Server 9 series host	<code>/etc/init.d/iscsi start</code>
	SUSE Linux Enterprise Server 10 series	<code>/etc/init.d/open-iscsi start</code>

Configuring the iSCSI service to start automatically

If you want to configure the iSCSI service to start automatically at the system boot, complete the following step.

Step	Action	
1	Perform the following steps to configure the iSCSI service:	
	If you are using...	Then...
	Red Hat Enterprise Linux	<code>chkconfig iscsi on</code>
	SUSE Linux Enterprise Server 9	<code>chkconfig iscsi on</code>
	SUSE Linux Enterprise Server 10	<code>chkconfig open-iscsi on</code>

Note—
You can also use the Service tab in YaST2 to enable the iSCSI initiator at the system boot for SUSE Linux Enterprise Server.

Accessing LUNs with DM-Multipath on Red Hat Enterprise Linux

Understanding automatic mounting

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

For detailed information

The following sections discuss how you can access LUNs with DM-Multipath:

- ◆ [“Viewing the list of LUNs”](#) on page 121
- ◆ [“Creating and mounting a file system on a multipath device”](#) on page 126

Accessing LUNs with DM-Multipath

Viewing the list of LUNs

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

Using the `sanlun` command

To view a list of LUNs that are mapped to the Linux host, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show all</code>

Example: The output looks similar to the following:

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

To display multipath information for native DM-Multipath, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show -p</code>

Example: The output looks similar to the following.

```
# sanlun lun show -p
f3020-236-39:/vol/vol1/lun1 (LUN 6)          Lun state: GOOD
Lun Size:          7g (7516192768)          Filer_CF_State: Cluster
Enabled
Protocol: iSCSI          Filer Partner: f3020-236-40
DM-MP DevName:      360a9800043346773435a48514d624367
dm-0
Multipath-provider: NATIVE
```

sanlun	Filer			Primary	Partner
path	Path	/dev/	Host	Filer	Filer
state	type	node	HBA	port	port
UP	iscsi	sdc	host4	10.72.236.69	--
UP	iscsi	sdd	host3	10.72.236.39	--

Using the iscsi-ls command for Red Hat Enterprise Linux 4

The `iscsi-ls -l` command also lists one entry for each path to each LUN.

Example: The output looks similar to the following:

```
TARGET NAME           : iqn.1992-08.com.netapp:sn.101175345
TARGET ALIAS          :
HOST ID               : 21
BUS ID                : 0
TARGET ID             : 0
TARGET ADDRESS        : 192.168.2.5:3260,3
SESSION STATUS        : ESTABLISHED AT Fri Apr 14 11:44:04 EDT 2006
SESSION ID            : ISID 00023d000002 TSIH 91d
```

DEVICE DETAILS:

LUN ID : 0

Vendor: NETAPP Model: LUN Rev: 0.2

Type: Direct-Access ANSI SCSI revision: 04

page83 type3: 60a9800043346536534a344a794e6d59

page80: 43346536534a344a794e6d590a

Device: /dev/sde

Using the iscsiadm command Red Hat Enterprise Linux 5

The `iscsiadm` command lists the LUNs that are mapped to the Linux host.

Example: The output looks similar to the following:

```
# iscsiadm -m session -P 3 -r 2
Target: iqn.1992-08.com.netapp:sn.101183016
    Current Portal: 10.72.199.71:3260,1001
    Persistent Portal: 10.72.199.71:3260,1001
    *****
    Interface:
    *****
    Iface Name: default
    Iface Transport: tcp
    Iface Initiatorname: iqn.1994-
05.com.redhat:5e3e11e0104d
    Iface IPaddress: 10.72.199.119
    Iface HWaddress: default
    Iface Netdev: default
    SID: 2
    iSCSI Connection State: LOGGED IN
    iSCSI Session State: Unknown
    Internal iscsid Session State: NO CHANGE
    *****
    Negotiated iSCSI params:
    *****

    HeaderDigest: None
    DataDigest: None
    MaxRecvDataSegmentLength: 131072
    MaxXmitDataSegmentLength: 65536
    FirstBurstLength: 65536
    MaxBurstLength: 65536
    ImmediateData: Yes
    InitialR2T: No
    MaxOutstandingR2T: 1
    *****
    Attached SCSI devices:
    *****

    Host Number: 4   State: running
    scsi4 Channel 00 Id 0 Lun: 0
        Attached scsi disk sdc           State: running
```

```

scsi4 Channel 00 Id 0 Lun: 1
    Attached scsi disk sde      State: running
scsi4 Channel 00 Id 0 Lun: 2
    Attached scsi disk sdg      State: running
scsi4 Channel 00 Id 0 Lun: 3
    Attached scsi disk sdi      State: running
scsi4 Channel 00 Id 0 Lun: 4
    Attached scsi disk sdk      State: running
scsi4 Channel 00 Id 0 Lun: 5
    Attached scsi disk sdm      State: running
scsi4 Channel 00 Id 0 Lun: 6
    Attached scsi disk sdp      State: running
scsi4 Channel 00 Id 0 Lun: 7
    Attached scsi disk sdq      State: running

```

The preceding example lists the available storage systems and LUNs for a session with a specific session ID. To view the details of all the sessions, use the `iscsiadm -m session -P 3` command.

Using the `lsscsi` command

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

Example: The output looks similar to the following:

```

[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

```

Discovering new LUNs

In Red Hat Enterprise Linux 5 series: If you create a new LUN and map it to the Linux host, you can discover the LUN by rescanning the iSCSI service on the host.

- ◆ Use the following command to get the list of all the current sessions:
`iscsiadm -m session`
- ◆ Use the following command to rescan a specific session:

```
iscsiadm -m session --sid=N --rescan
```

N is the specific session ID.

- ◆ Use the following command to rescan all the sessions:

```
iscsiadm -m session --rescan
```

You can use the `sanlun` command in Red Hat Enterprise Linux 5 series to verify that the new LUNs were discovered.

After the LUNs are discovered, they are automatically added to the DM-Multipath configuration.

In Red Hat Enterprise Linux 4 series: If you create a new LUN and map it to the Linux host, you can discover the LUN by reloading the iSCSI service on the host. To reload the iSCSI service, enter the following command on the Linux host console:

```
/etc/init.d/iscsi reload
```

Use the `sanlun` or `iscsi-ls` command to verify that the new LUNs were discovered.

After the LUNs are discovered, they are automatically added to the DM-Multipath configuration.

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, use the DM-Multipath device in the `/dev/mapper/` directory.

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

Step	Action
1	Enter the following command on the Linux host console: <pre>mkfs -t type /dev/mapper/ device</pre> <i>type</i> is the file system type, such as <code>ext3</code> . <i>device</i> is the multipath device name of the LUN in the <code>/dev/mpath/</code> directory. Example: The following command creates an <code>ext3</code> file system on a LUN: <pre>mkfs -t ext3 /dev/mapper/mpath0</pre>

Mounting the file system

To mount the new file system, 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 directory by using the <code>mkdir</code> command.
2	Open the <code>/etc/fstab</code> file with an editor.
3	Add an entry to <code>/etc/fstab</code> as follows: <pre>device mount_point type _netdev 0 0</pre> <i>device</i> is the name of the device in the <code>/dev/mapper/</code> directory. <i>mount_point</i> is the mount point you created for the file system. <i>type</i> is the file system type, such as <code>ext2</code> or <code>ext3</code>. <i>_netdev</i> is used for any network-dependent devices like iSCSI. Example: <pre>/dev/mapper/mpath0 /mnt/ss1/lun_0 ext3 _netdev 0 0</pre>
4	To mount the file system, enter the following command: <pre>mount mountpoint</pre> Example: <pre>mount /mnt/ss1/lun_0</pre>
5	Enter <code>df</code> to display the mounted file systems. Verify that the new file system is listed. Another option is to change to the mount point directory and try writing and reading a file to verify that the file system was mounted successfully.

Accessing LUNs with DM-Multipath on SUSE Linux Enterprise Server

Overview

To access LUNs with DM-Multipath, you need to perform the following tasks:

1. Identify and view the list of LUNs on which the partition is to be created. Depending on the series you are using, for more information see:
 - ❖ [“Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 9 series”](#) on page 129
 - ❖ [“Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 10 series”](#) on page 132
2. Create and mount the file system. Depending on the series you are using, for more information, see:
 - ❖ [“Creating and mounting a file system on a multipath device”](#) on page 126
 - ❖ [“Creating and mounting a file system on a LUN on SUSE Linux Enterprise Server 10 series”](#) on page 137

Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 9 series

To identify LUNs, you can view information about LUNs in the following ways:

- ◆ The `sanlun` command is included in Host Utilities. The command displays the host device names and the storage system LUNs to which they are mapped.
- ◆ The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.
- ◆ The `lsscsi` command lists all the SCSI devices and their attributes.

Note

You can use the `sanlun`, `iscsi-ls`, or `lsscsi` command to verify that the new LUNs were discovered.

Using the `sanlun` command to view LUNs

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

To view a list of LUNs that are mapped to the Linux host, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show all</code>

Example: The output looks similar to the following:

```
# sanlun lun show
filer:      lun-pathname      device filename  adapter protocol   lun size      lun state
f3020-43:   /vol/fvol1/lun122  /dev/sdb        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1    /dev/sdc        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol1/lun122  /dev/sdd        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1    /dev/sde        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol1/lun122  /dev/sdf        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1    /dev/sdg        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol1/lun122  /dev/sdh        host4   iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1    /dev/sdi        host4   iSCSI    1g (1073741824)  GOOD
```

Using the `iscsi-ls` command

The `iscsi-ls -l` command also lists one entry for each path to each LUN.

Example: The output looks similar to the following:

```
DEVICE DETAILS :
-----
LUN ID : 0
  Vendor: NETAPP      Model: LUN              Rev: 0.2
  Type:   Direct-Access      ANSI SCSI revision: 04
*****
TARGET NAME           : iqn.1992-08.com.netapp:sn.101175385
TARGET ALIAS          :
HOST NO               : 5
BUS NO                : 1
TARGET ID             : 0
TARGET ADDRESS        : 10.61.208.5:3260
SESSION STATUS        : ESTABLISHED AT Tue Mar 28 16:03:15 2006
NO. OF PORTALS        : 1
PORTAL ADDRESS 1      : 10.61.208.5:3260,2
SESSION ID            : ISID 00023d000002 TSID 544
```

Using the `lsscsi` command

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

Example: The output looks similar to the following:

[0:0:0:0]	cd/dvd	HL-DT-ST	RW/DVD	GCC-4246N	OY03	/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

Viewing newly created LUNs

To view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration.

Discovering new LUNs: If you create a new LUN and map it to the Linux host, you can discover the LUN by reloading the iSCSI service on the host. To reload the iSCSI service, enter the following command on the Linux host console:

```
/etc/init.d/iscsi reload
```

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

```
multipath
```

The output of the multipath command displays the configuration of the new LUN.

Identifying and viewing a list of LUNs on SUSE Linux Enterprise Server 10 series

To identify LUNs, you can view information about LUNs in the following ways:

- ◆ The `sanlun` command is included in Host Utilities. The command displays the host device names and the storage system LUNs to which they are mapped.
- ◆ The `iscsiadm -m session --sid=N -P 3` command lists the available storage systems and LUNs for a session with a specific session ID, *N*.
- ◆ The `lsscsi` command lists all the SCSI devices including the iSCSI devices and their attributes.

Using the `sanlun` command to view LUNs

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

To view a list of LUNs that are mapped to the Linux host, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show all</code>

Example: The output looks similar to the following:

```
# sanlun lun show all
filer:  lun-pathname    device filename  adapter  protocol  lun size      lun state
f3020-94: /vol/vol1/lun2    /dev/sdg        host6    iSCSI     4g (4294967296)  GOOD
f3020-94: /vol/vol1/lun1    /dev/sdf        host6    iSCSI     4g (4294967296)  GOOD
f3020-95: /vol/vol1/lun95_3 /dev/sde        host5    iSCSI     4g (4294967296)  GOOD
f3020-95: /vol/vol1/lun95_4 /dev/sdd        host5    iSCSI     4g (4294967296)  GOOD
f3020-95: /vol/vol1/lun95_2 /dev/sdc        host5    iSCSI     4g (4294967296)  GOOD
f3020-95: /vol/vol1/lun95_1 /dev/sdb        host5    iSCSI     4g (4294967296)  GOOD
```

Using the iscsiadm command

The `iscsiadm -m session --sid=N -P 3` command lists the available storage systems and LUNs for a session with a specific session ID, *N*.

Example: The output looks similar to the following:

```
# iscsiadm -m session --sid=0 -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.netapp:sn.101180456
    Current Portal: 10.72.199.144:3260,1001
    Persistent Portal: 10.72.199.144:3260,1001
    *****
    Interface:
    *****
    Iface Name: default
    Iface Transport: tcp
    Iface Initiatorname: iqn.1996-
04.de.suse:lnx.200.109
    Iface IPaddress: 10.72.200.109
    Iface HWaddress: default
    Iface Netdev: default
    SID: 1
    iSCSI Connection State: LOGGED IN
    iSCSI Session State: Unknown
Internal iscsid Session State: NO CHANGE
    *****
    Negotiated iSCSI params:
    *****
    HeaderDigest: CRC32C
    DataDigest: None
    MaxRecvDataSegmentLength: 131072
    MaxXmitDataSegmentLength: 65536
    FirstBurstLength: 65536
    MaxBurstLength: 65536
    ImmediateData: Yes
    InitialR2T: No
    MaxOutstandingR2T: 1
    *****
    Attached SCSI devices:
    *****
```

```

Host Number: 47 State: running
scsi47 Channel 00 Id 0 Lun: 2
    Attached scsi disk sdj          State: running
scsi47 Channel 00 Id 0 Lun: 1
    Attached scsi disk sdf          State: running
scsi47 Channel 00 Id 0 Lun: 0
    Attached scsi disk sdb          State: running

```

Using the lsscsi command

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

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

```

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

```

Viewing newly discovered LUNs

To view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration.

Discovering new LUNs: If you create a new LUN and map it to the Linux host, you can discover the LUN by rescanning the iSCSI service on the host.

- ◆ Use the following command to get the list of all the current sessions:

```
iscsiadm -m session
```

- ◆ Use the following command to rescan a specific session:

```
iscsiadm -m session --sid=N --rescan
```

N is the specific session ID.

- ◆ Use the following command to rescan all the sessions:

```
iscsiadm -m session --rescan
```

Note

To verify that the new LUNs were discovered, you can use the `sanlun`, `iscsiadm`, or `lscsi` command.

Creating and mounting a file system on a LUN on SUSE Linux Enterprise Server 9 series

After a LUN is added to the DM-Multipath configuration, you can create a file system.

When creating a file system, 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

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

To mount the file system, complete the following steps.

Step	Action
1	Create a mountpoint directory by using the <code>mkdir</code> command.
2	Open the <code>/etc/fstab.iscsi</code> file with an editor.

Step	Action
3	Add an entry to <code>/etc/fstab.iscsi</code> as follows: <pre>device mount_point type defaults 0 0</pre> <i>device</i> is the name in the <code>/dev/mapper</code> directory. <i>mount_point</i> is the mountpoint you created for the file system. <i>type</i> is the file system type, such as <code>ext2</code>, <code>ext3</code>, or <code>reiserfs</code>. Example: <pre>/dev/mapper/360a9800043346536684a337a7553424f /mnt/ssl/lun_0 ext3 defaults 0 0</pre>
4	To mount the file system, enter the following command: <pre>/sbin/iscsi-mountall</pre>
5	Perform one of the following. ◆ Enter df to display the mounted file systems. Verify that the new file system is listed. ◆ Change to the mountpoint directory and try writing and reading a file to verify that the file system was mounted successfully.

Creating and mounting a file system on a LUN on SUSE Linux Enterprise Server 10 series

After a LUN is added to the DM-Multipath configuration, you can create a file system.

When creating a file system, 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
```

Automatically mounting the file system

To automatically mount the file system on SUSE Linux Enterprise Server 10 series, complete the following steps.

Attention

Before mounting the file system, modify the network configuration file and change the value of `STARTMODE` to `nfsroot`. This value is required for the successful unmounting of the iSCSI devices during system halts.

For SUSE Linux Enterprise Server 10 series, modify `/etc/sysconfig/network/ifcfg-eth-idxxxxxxxxx`.

Step	Action
1	On the host, create a mountpoint for the LUNs.
2	With a text editor, open the host's <code>/etc/fstab</code> file.

Step	Action
3	Add an entry for the file system you created in “Creating a file system” on page 158. Example: <pre>/dev/mapper/360a98044334685f433e6774722-part1 /mnt/data ext3 hotplug,rw 0 0</pre> where, /dev/mapper/360a98044334685f433e6774722-part1 is the iSCSI device. /mnt/data is the mountpoint. ext3 is the file system.
4	Save the changes and exit the /etc/fstab file.
5	On the host console, enter the following command to mount the file systems you added: <pre>/bin/mount -a</pre>

Accessing LUNs without multipathing software on Red Hat Enterprise Linux

Persistently identifying each 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 a different device name after a Linux reboot or restart iSCSI service.

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.

Understanding automatic mounting

The Linux system cannot mount file systems on iSCSI LUNs when it first boots up, because the networking and iSCSI driver resources are not yet available. The Linux `netfs` `init` script mounts these file systems after the networking and iSCSI systems are available.

To take advantage of the `netfs` `init` script, specify the mount options for the file system in `/etc/fstab` as `_netdev`. For example:

#device	mount	FS	mount	backup	fsck
#to mount	point	type	options	frequency	pass
LABEL=lun0_label	/mnt/filer0_lun0	ext3	_netdev	0	0

The SourceForge iSCSI initiator automount script `/sbin/iscsi-mountall` is not included with the iSCSI initiator in Red Hat Enterprise Linux 5 series, Red Hat Enterprise Linux 4 series, or Red Hat Enterprise Linux 3 Update 4 and later. Use the automounting method described earlier instead.

Tasks for accessing LUNs

After the initiator is configured, you need to set up the host to access the LUNs that you created on the storage system. To access the LUNs, you should complete the following task:

- ◆ If you are creating a file system on the LUN:
 - a. Create a file system.
 - b. Create a label for the file system and add an entry for the file system's label in `/etc/fstab`.

For detailed information

The following sections discuss how you can access LUNs without multipathing software:

- ◆ [“Creating a file system”](#) on page 141
- ◆ [“Labeling and mounting the file system”](#) on page 142
- ◆ [“Viewing LUN information”](#) on page 144

Creating a file system

About file systems A file system enables users and applications to read and write files on a LUN.

Linux supports different file systems. You can create any type of file system. This section explains how to create a journaling ext3 file system.

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

Step	Action
1	On the host console, enter the following command: <code>mkfs -t filesystem_type device</code> <i>filesystem_type</i> is the type of file system that you want to create. <i>device</i> is the partition.

Accessing LUNs without multipathing software

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 each LUN](#)” on page 139 for more information.

After you create a label, you can 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 the following command: <code>/sbin/e2label device label</code> <i>device</i> is the partition where you created the file system. <i>label</i> is a string (maximum 16 characters) used to identify the file system. Example: <code>/sbin/e2label /dev/sda1 lun0_label</code>

Mounting the file system

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

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

Step	Action
3	Add an entry to /etc/fstab as follows: <pre><code>LABEL=<i>label</i> <i>mount_point</i> auto _netdev 0 0</code></pre> <i>label</i> is the file system label that you created earlier. <i>mount_point</i> is the mountpoint that you created for the file system. Example: <pre><code>LABEL=lun0_label /mnt/filer0_lun0 auto _netdev 0 0</code></pre>
4	Save your changes to /etc/fstab and exit the editor.
5	To mount the file system, enter the following command: <pre><code>mount LABEL=<i>label</i></code></pre> The file system is automatically mounted whenever the Linux host is rebooted.
6	Enter df to display the mounted file systems. Verify that the new file system is listed. Another option is to change to the mount point directory and try writing and reading a file to verify that the file system was mounted successfully.

Viewing LUN information

You can view information about LUNs in the following ways:

- ◆ The `iscsiadm` command displays the storage system and LUN details.
- ◆ The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.
- ◆ The `sanlun` command, which is included in the iSCSI Linux Host Utilities package.
- ◆ The `lsscsi` command lists information about all the SCSI devices and their attributes.

The command displays the host device names and the storage system LUNs to which they are mapped.

Using the iscsiadm command for Red Hat Enterprise Linux 5

To view the storage system and LUN details using the iscsiadm command, complete the following step.

Step	Action
1	Enter the following command at the host’s shell prompt: <code>iscsiadm -m session -P 3 -r 2</code> Result: The console lists the available storage systems and LUNs for a session with a specific session ID, as shown by the following example. <pre># iscsiadm -m session -P 3 -r 2 Target: iqn.1992-08.com.netapp:sn.101183016 Current Portal: 10.72.199.71:3260,1001 Persistent Portal: 10.72.199.71:3260,1001 ***** Interface: ***** Iface Name: default Iface Transport: tcp Iface Initiatorname: iqn.1994-05.com.redhat:5e3e11e0104d Iface IPaddress: 10.72.199.119 Iface HWaddress: default Iface Netdev: default SID: 2 iSCSI Connection State: LOGGED IN iSCSI Session State: Unknown Internal iscsid Session State: NO CHANGE ***** Negotiated iSCSI params: ***** HeaderDigest: None DataDigest: None MaxRecvDataSegmentLength: 131072 MaxXmitDataSegmentLength: 65536 FirstBurstLength: 65536 MaxBurstLength: 65536 ImmediateData: Yes InitialR2T: No MaxOutstandingR2T: 1 ***** Attached SCSI devices: *****</pre>

Step	Action
	<pre> Host Number: 4 State: running scsi4 Channel 00 Id 0 Lun: 0 Attached scsi disk sdc State: running scsi4 Channel 00 Id 0 Lun: 1 Attached scsi disk sde State: running scsi4 Channel 00 Id 0 Lun: 2 Attached scsi disk sdg State: running scsi4 Channel 00 Id 0 Lun: 3 Attached scsi disk sdi State: running scsi4 Channel 00 Id 0 Lun: 4 Attached scsi disk sdk State: running scsi4 Channel 00 Id 0 Lun: 5 Attached scsi disk sdm State: running scsi4 Channel 00 Id 0 Lun: 6 Attached scsi disk sdq State: running scsi4 Channel 00 Id 0 Lun: 7 Attached scsi disk sdq State: running </pre> To view the details of all the sessions, use the <code>iscsiadm -m session -P 3</code> command.

Using the iscsi-ls command for Red Hat Enterprise Linux 4

To view the storage system and LUN details using the `iscsi-ls` command, complete the following step.

Step	Action
1	Enter the following command at the host's shell prompt: <pre> /sbin/iscsi-ls -l </pre> Result: The console lists the available storage systems and LUNs, as shown by the following example: <pre> # /sbin/iscsi-ls -l ***** SFNet iSCSI Driver Version ... 3.6.2 (27-Sep-2004) ***** TARGET NAME : iqn.1992-08.com.netapp:sn.33604646 TARGET ALIAS : HOST NO : 0 BUS NO : 0 TARGET ID : 0 TARGET ADDRESS : 10.60.128.100:3260 SESSION STATUS : ESTABLISHED AT Mon Jan 3 10:05:14 2005 NO. OF PORTALS : 1 PORTAL ADDRESS 1 : 10.60.128.100:3260,1 SESSION ID : ISID 00023d000001 TSID 103 DEVICE DETAILS : ----- LUN ID : 0 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 04 page83 type3: 60a980004f6443745359763759367733 page83 type1:4e45544150502020204c554e204f644374535976375936773300000000000000 page80: 4f6443745359763759367733 Device: /dev/sdb LUN ID : 1 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 04 page83 type3: 60a980004f644374535976426253674b page83 type1:4e45544150502020204c554e204f644374535976426253674b00000000000000 page80: 4f644374535976426253674b Device: /dev/sdc ***** </pre>

Using the sanlun command: The sanlun program is installed from iSCSI Host Utilities on Red Hat Enterprise Linux.

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

Step	Action
1	Enter the following command at the host's shell prompt: <pre>/opt/netapp/santools/sanlun lun show all</pre> Result: The sanlun command displays information about the iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The following lines show a sample sanlun command output: <pre># /opt/netapp/santools/sanlun lun show all filer: lun-pathname device filename adapter protocol lun size lun state filer1: /vol/vol_1/lun_1 /dev/sdb1 <unknown>0 iSCSI 4g (4294967296) GOOD filer1: /vol/vol_1/lun_2 /dev/sdc1 <unknown>0 iSCSI 4g (4294967296) GOOD</pre>

Using the lsscsi command

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

Example: The output is similar to the following:

```
Internal iscsid Session State: NO CHANGE
*****
Negotiated iSCSI params:
*****
HeaderDigest: CRC32C
DataDigest: None
MaxRecvDataSegmentLength: 131072
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 65536
ImmediateData: Yes
InitialR2T: No
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 47 State: running
```

```
scsi47 Channel 00 Id 0 Lun: 2
Attached scsi disk sdj State: running
scsi47 Channel 00 Id 0 Lun: 1
Attached scsi disk sdf State: running
scsi47 Channel 00 Id 0 Lun: 0
Attached scsi disk sdb State: running
[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
```

Accessing LUNs without multipathing software on SUSE Linux Enterprise Server

Persistently identifying each 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 a different device name after a Linux reboot or restart iSCSI service.

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

Tasks for accessing LUNs

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

1. Identify the LUN on which the partition is to be created. For more information, see [“Viewing information about LUNs”](#) on page 151.
2. Create a file system. For more information, see [“Creating a file system”](#) on page 158.
3. Mount the file system. For more information, see [“Mounting the file system”](#) on page 159.

Viewing information about LUNs

Viewing information about LUNs

You can view information about LUNs in the following ways:

SUSE Linux Enterprise Server 9 series: To identify LUNs, you can view information about LUNs using the following commands:

- ◆ The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.
- ◆ The `lsscsi` command lists all the SCSI devices including the iSCSI devices and their attributes.
- ◆ The `sanlun` command is included in Host Utilities. The command displays the host device names and the storage system LUNs to which they are mapped.

Using the `iscsi-ls` command: To view the storage system and LUN details, complete the following step.

Step	Action
1	Enter the following command on the host: <code>iscsi-ls</code> Result: The console lists the available storage systems and LUNs.

Example: The following example shows the output for SUSE Linux Enterprise Server 9 SP3.

```
*****
Cisco iSCSI Driver Version ... 4.0.188.26 ( 21-May-2004 )
*****
TARGET NAME           : ign.1992-08.com.netapp:sn.84209894
TARGET ALIAS          :
HOST NO               : 5
BUS NO                : 0
TARGET ID             : 2
TARGET ADDRESS        : 10.72.199.38:3260
SESSION STATUS        : ESTABLISHED AT Sun Jun 24 01:47:46 2007
NO. OF PORTALS        : 1
PORTAL ADDRESS 1      : 10.72.199.38:3260,1000
SESSION ID            : ISID 00023d000001 TSID 72

DEVICE DETAILS :
-----
LUN ID : 0
  Vendor: NETAPP   Model: LUN                      Rev: 0.2
  Type:   Direct-Access                      ANSI SCSI revision: 04
LUN ID : 1
  Vendor: NETAPP   Model: LUN                      Rev: 0.2
  Type:   Direct-Access                      ANSI SCSI revision: 04
LUN ID : 2
  Vendor: NETAPP   Model: LUN                      Rev: 0.2
  Type:   Direct-Access                      ANSI SCSI revision: 04
*****
```

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

Step	Action
1	On the host, change to the /opt/netapp/santools directory.
2	Enter the following command: <pre>sanlun lun show all</pre> Result: The sanlun command displays information about the iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state.

Example: The following example shows the `sanlun` command output for SUSE Linux Enterprise Server 9 series.

```
# /opt/netappontap/santools/sanlun lun show all
filer:      lun-pathname      device filename adapter  protocol lun size      lun state
f3020-43:   /vol/fvol11/lun122 /dev/sdb      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1     /dev/sdc      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol11/lun122 /dev/sdd      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1     /dev/sde      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol11/lun122 /dev/sdf      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1     /dev/sdg      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol11/lun122 /dev/sdh      host4    iSCSI    1g (1073741824)  GOOD
f3020-43:   /vol/fvol6/lun1     /dev/sdi      host4    iSCSI    1g (1073741824)  GOOD
```

Using the `iscsi-ls` command: To view information about target configuration, complete the following step.

Step	Action
1	Enter the following command on the host: <code>iscsi-ls -c</code> Result: The console lists the available storage systems and LUNs, and the iSCSI session parameters.

Example: The following example shows the output for SUSE Linux Enterprise Server 9 series.

```
*****
        Cisco iSCSI Driver Version ... 4.0.188.13 ( 21-May-2004 )
*****
TARGET NAME           : iqn.1992-08.com.netapp:sn.12345678
TARGET ALIAS          :
HOST NO               : 23
BUS NO                : 0
TARGET ID             : 0
TARGET ADDRESS        : 192.123.10.50:3260
SESSION STATUS        : ESTABLISHED AT Mon Aug 23 18:48:31 2004
NO. OF PORTALS        : 1
PORTAL ADDRESS 1      : 192.123.10.50:3260,2
SESSION ID            : ISID 00023d000001 TSID 66f

TARGET CONFIGURATION KEYS :
-----
INITIAL R2T           : Yes
IMMEDIATE DATA       : Yes
HEADERDIGEST          : NONE
DATA DIGEST           : NONE
FIRSTBURSTLENGTH      : 65536
MAXBURSTLENGTH        : 262144
MAXRECVDATASEGMENTLENGTH: initiator - 131072
                        target      - 65536
LOGIN TIMEOUT         : 15
AUTH TIMEOUT          : 45
ACTIVE TIMEOUT        : 5
IDLE TIMEOUT          : 60
PING TIMEOUT          : 5
```

Using the lsscsi command: To view information about all the SCSI devices and their attributes, complete the following step.

Step	Action
1	Enter the following command on the host console: <pre>lsscsi</pre> Result: The <code>lsscsi</code> command lists information about all the SCSI devices and their attributes.

For SUSE Linux Enterprise Server 10 series: To identify the LUN, you can view information about LUNs in the following ways:

- ◆ **sanlun:** The `sanlun` command is included in Host Utilities. The command displays the host device names and the storage system LUNs to which they are mapped.
- ◆ **iscsiadm**
 - ❖ For SUSE Linux Enterprise Server 10 SP1: The `iscsiadm -m session --sid=N -i` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.
 - ❖ For SUSE Linux Enterprise Server 10 SP2: The `iscsiadm -m session --sid=N -P 3` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.
- ◆ **lsscsi:** The `lsscsi` command lists all the SCSI devices including the iSCSI devices and their attributes.

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, change to the <code>/opt/netapp/santools</code> directory.
2	Enter the following command: <pre>sanlun lun show all</pre> Result: The <code>sanlun</code> command displays information about the iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The sample output of the <code>sanlun</code> command is as follows: <pre># /opt/netappontap/santools/sanlun lun show all filer: lun-pathname device filename adapter protocol lun size lunstate f3020-94: /vol/vol1/lun2 /dev/sdg host6 iSCSI 4g (4294967296) GOOD f3020-94: /vol/vol1/lun1 /dev/sdf host6 iSCSI 4g (4294967296) GOOD f3020-95: /vol/vol1/lun95_3 /dev/sde host5 iSCSI 4g (4294967296) GOOD f3020-95: /vol/vol1/lun95_4 /dev/sdd host5 iSCSI 4g (4294967296) GOOD f3020-95: /vol/vol1/lun95_2 /dev/sdc host5 iSCSI 4g (4294967296) GOOD f3020-95: /vol/vol1/lun95_1 /dev/sdb host5 iSCSI 4g (4294967296) GOOD</pre>

Using the `iscsiadm -m session --sid=N` command: To view information about the available storage systems and LUNs for a session with a specific session ID, *N*, complete the following step.

Step	Action
1	Enter the following command on the host console: <pre>iscsiadm -m session --sid=N -P 3</pre> <i>N</i> is the specific session ID. Result: The available storage systems and LUNs for a session with a specific session ID, <i>N</i>, are listed.

Example: The following example shows the output on SUSE Linux Enterprise Server 10 series.

```
# iscsiadm -m session --sid=N -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: ign.1992-08.com.netapp:sn.101180456
Current Portal: 10.72.199.144:3260,1001
Persistent Portal: 10.72.199.144:3260,1001
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface Initiatorname:ign.1996-04.de.suse:lnx.200.109
Iface IPaddress: 10.72.200.109
Iface HWaddress: default
Iface Netdev: default
SID: 1
iSCSI Connection State: LOGGED IN
iSCSI Session State: Unknown
Internal iscsid Session State: NO CHANGE
*****
Negotiated iSCSI params:
*****
HeaderDigest: CRC32C
DataDigest: None
MaxRecvDataSegmentLength: 131072
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
```

```

MaxBurstLength: 65536
ImmediateData: Yes
InitialR2T: No
MaxOutstandingR2T: 1
*****
Attached SCSI devices:
*****
Host Number: 47 State: running
scsi47 Channel 00 Id 0 Lun: 2
    Attached scsi disk sdj          State: running
scsi47 Channel 00 Id 0 Lun: 1
    Attached scsi disk sdf          State: running
scsi47 Channel 00 Id 0 Lun: 0
    Attached scsi disk sdb          State: running

```

To list all the active iSCSI sessions, use the `iscsiadm -m session` command.

Using the `lsscsi` command: To view information about all the SCSI devices and their attributes, complete the following step.

Step	Action
1	Enter the following command on the host console: <code>lsscsi</code> Result: The <code>lsscsi</code> command lists information about all the SCSI devices and their attributes.

Creating a file system

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

Step	Action
1	On the host console, enter the following command: <pre>mkfs -t <i>filesystem_type</i> <i>device</i></pre> <i>filesystem_type</i> is the type of file system that you want to create. <i>device</i> is the partition. Example for SUSE Linux Enterprise Server 9 series: The following command creates an ext3 file system on the device. <pre>mkfs -t ext3 /dev/disk/by-id/iscsi-ign.1992-08.com.netapp:sn.12345678.ss_hq-0</pre> Example for SUSE Linux Enterprise Server 10 series: The following command creates an ext3 file system on the device. <pre>mkfs -t ext3 /dev/disk/by-id/scsi-360a980004334654c6d4a39436c686a32</pre>

Mounting the file system

Before mounting the file system

Before mounting the file system, you need to create a label for the file system.

About labels: A file system label creates persistent mapping between a LUN and a file system. For more information, see “[Accessing LUNs without multipathing software on SUSE Linux Enterprise Server](#)” on page 150.

Using e2label to create a file system label: To create a file system label for an ext2 or ext3 file system, run the **e2label** command as follows:

```
/sbin/e2label device label
```

device is the partition where you created the file system.

label is a string (maximum 16 characters) that is used to identify the file system.

Example: `/sbin/e2label /dev/sda1 lun0_label`

For Reiser file systems, create the label using the `--label` option when you create the file system. You should not use the `e2label` command because it is specific only to ext2/ext3 file systems and not to Reiser file systems..

Note

You should mount devices using *by-id* and not using labels in SUSE Linux Enterprise Server 10 series.

Automatically mounting the file system for SUSE Linux Enterprise Server 9 series

The Linux system mounts all the file systems listed in the `/etc/fstab.iscsi` file when it first boots up. To automatically mount the file system, 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 <code>/etc/fstab.iscsi</code> file.
3	Add an entry for the device on which you have created the file system. Example: <pre>/dev/disk/by-id/iscsi-iqn.1992-08.com.netapp:sn.12345678.ss_hq-0p1 /mnt_pt ext3 defaults 0 0</pre> <i>mnt_pt</i> is the mountpoint you created in Step 1. Result: The <code>iscsi-mountall</code> script reads the <code>/etc/fstab.iscsi</code> file when the iSCSI service is started, and then mounts the file systems that are listed.
4	Save your changes and exit the <code>/etc/fstab.iscsi</code> file.
5	On the host console, enter the following command to mount the file systems you added: <pre>/sbin/iscsi-mountall</pre>

Automatically mounting the file system - SUSE Linux Enterprise Server 10 series

To automatically mount the file system, complete the following steps.

Attention

Before mounting the file system, modify the network configuration file and change the value of `STARTMODE` to *nfsroot*. This value is required for the successful unmounting of the iSCSI devices during system halts.

For SUSE Linux Enterprise Server 10 series, modify `/etc/sysconfig/network/ifcfg-eth-idxxxxxxxxxx`.

Step	Action
1	On the host, create a mountpoint for the LUNs.
2	With a text editor, open the host's <code>/etc/fstab</code> file.

Step	Action
3	Add an entry for the device on which you have created the file system. Example: ◆ For SUSE Linux Enterprise Server 10 SP2 /dev/disk/by-id/scsi-360a98044334685f433e6774722-part1 /mnt/data ext3 hotplug,rw 0 0 where, /dev/disk/by-id/scsi-360a98044334685f433e6774722-part1 is the iSCSI device and it needs to be listed by its ID. /mnt/data is the mountpoint. ext3 is the file system.
4	Save your changes, and exit the /etc/fstab file.
5	On the host console, enter the following command to mount the file systems you added: /bin/mount -a

Accessing LUNs using Hardware iSCSI

About this chapter

If you are using the DM-Multipath, the method you use to access LUNs is different from when you are not using multipathing.

Topic in this chapter

This chapter includes the following topic:

- ◆ [“Accessing LUNs with DM-Multipath”](#) on page 164

Accessing LUNs with DM-Multipath

To access LUNs with DM-Multipath, perform the following tasks.

1. View the list of LUNs. For more information, see “[Viewing the list of LUNs](#)” on page 165.
2. Create and mount the file system. For more information, see “[Creating and mounting a file system on a multipath device](#)” on page 167.

Viewing the list of LUNs

Viewing list of LUNs To view information about LUNs, you can use the `sanlun` command.

To view a list of LUNs that are mapped to the Linux host, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show all</code>

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

```
# sanlun lun show all
filer:   lun-pathname      device filename adapter  protocol lun size lun state
f3020-94: /vol/vol7/lun1  /dev/sdb host4   iSCSI  9g      (9663676416)  GOOD
f3020-94: /vol/vol7/lun1  /dev/sda host4   iSCSI  9g      (9663676416)  GOOD
```

To display multipath information for native device-mapper multipath, complete the following step.

Step	Action
1	Enter the following command on the Linux host console: <code>sanlun lun show -p</code>

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

```
# sanlun lun show -p
f3020-94:/vol/vol7/lun1 (LUN 0)          Lun state: GOOD
Lun Size:      9g (9663676416)          Filer_CF_State: Cluster Enabled
Protocol: iSCSI          Filer Partner: f3020-95
DM-MP DevName:          360a9800043346851436f4f48457a4e31    dm-0
Multipath-provider: NATIVE
```

sanlun	Filer			Primary	Partner
path	Path	/dev/	Host	Filer	Filer
state	type	node	HBA	port	port
UP	iscsi	sdb	host4	10.72.199.71	--
UP	iscsi	sda	host4	10.72.199.94	--

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, use the multipath device in the /dev/mapper/ directory.

Creating the file system

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

Step	Action
1	Enter the following command on the Linux host console: <pre>mkfs -t <i>type</i> /dev/mapper/<i>device</i></pre> <i>type</i> is the file system type, such as ext3 . <i>device</i> is the multipath device name of the LUN in the /dev/mpath/ directory. Example: The following command creates an ext3 file system on a LUN: <pre>mkfs -t ext3 /dev/mapper/mpath1</pre>

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, create a mountpoint and add an entry to the /etc/fstab file.

To mount the file system, complete the following steps.

Step	Action
1	Create a mountpoint using the following command: mkdir <i>mount_point</i> <i>mount_point</i> is the name of the mountpoint you are creating.
2	Open the /etc/fstab file with a text editor.
3	Add an entry to /etc/fstab as follows, separating each entry by pressing the Tab key (except the final two zeros, which are separated by a space): <i>device</i> <i>mount_point</i> <i>type</i> defaults 0 0 <i>device</i> is the name of the device in the /dev/mapper/ directory. <i>mount_point</i> is the mountpoint you created for the file system. <i>type</i> is the file system type, such as ext2 or ext3. Example: /dev/mapper/mpath1 /mnt/ss1/lun_0 ext3 defaults 0 0
4	To mount the file system for Red Hat Enterprise Linux, enter the following command: mount <i>mountpoint</i> Example: mount /mnt/ss1/lun_0
5	Do either of the following: ◆ Enter df to display the mounted file systems. Verify that the new file system is listed. ◆ Change to the mountpoint directory and try writing and reading a file to verify that the file system was mounted successfully.

About this chapter

SAN boot uses a SAN-attached disk, such as a LUN, as a boot device for a host.

Note

SAN boot is currently supported on Red Hat Enterprise Linux 5 Update 1 and later, and SUSE Linux Enterprise Server 10 SP2.

For information about the supported configuration, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Configuring SAN boot on a Red Hat Enterprise Server](#)” on page 170
- ◆ “[Configuring SAN boot on a SUSE Linux Enterprise Server](#)” on page 174

Configuring SAN boot on a Red Hat Enterprise Server

Tasks for configuring SAN boot

To configure SAN boot, complete the following tasks.

Step	Action
1	Install the operating system on the LUN configured on the storage controller.
2	Install Host Utilities.
3	Install and configure the external HBA driver or the bundled HBA driver.
4	Configure DM-Multipath.

Installing the operating system

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

Step	Action
1	On the storage system, select a cfmode setting.
2	Create a LUN on the storage system that acts as the SAN boot LUN, and map it to the host. Note_____ Ensure that the SAN boot LUN is mapped, and multiple paths to the LUN are available. Ensure that the SAN boot LUN is visible to the host during the boot process. _____

Step	Action
3	Enable the BIOS of the HBA port to which the SAN boot LUN is mapped. For information about how to enable the HBA BIOS, see your HBA vendor-specific documentation. Note _____ It is best to enable the spinup delay option for the individual HBA ports in the HBA BIOS.
4	Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device. For more information, see your vendor-specific documentation.
5	After making the relevant changes to the HBA BIOS, save and exit.
6	Reboot the host.
7	Install the operating system on the LUN configured on the storage controller. Note _____ Specify <i>linux mpath</i> in the Anaconda installer during the operating system installation. When <i>linux mpath</i> is specified, you can see the multipath devices (/dev/mapper/mpathx) as installation devices.

Installing Host Utilities

Install Linux Host Utilities 5.0. For information about installing the Linux Host Utilities 5.0, see “[Downloading and Installing the Host Utilities Software](#)” on page 4.

Installing and configuring the bundled HBA driver

To install and configure the bundled HBA driver, see “[Setting the required HBA and bundled driver parameters](#)” on page 26.

Installing and configuring the external HBA driver

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

Step	Action
1	Download and install the HBA driver. For more information on downloading, see “Download the external HBA driver” on page 24. QLogic: To install the external QLogic driver on the SAN booted host, complete the following steps. a. Run the <code>qlinstall</code> script with the <code>-up</code> option. b. Reboot the host to load the external QLogic driver.
2	Set the required HBA driver parameters. For more information about setting the QLogic HBA parameters, see “QLogic HBA parameters” on page 31. To set the HBA driver parameters, perform the following steps. a. Edit the <code>/etc/modprobe.conf</code>. b. Append the <code>options qla2xxx</code> statement with settings in the <code>modprobe.conf</code> file on a single line; only one <code>options qla2xxx</code> statement is allowed. <code>ql2xsuspendcount=6 MaxRetriesPerPath=10</code> <code>MaxRetriesPerIo=10 ql2xfailover=0</code> c. Save the file and exit.

Step	Action
	d. Enter the appropriate command for your HBA model to rebuild the initial ramdisk (initrd) image with the new parameter values by entering the following command: <pre>./qlinstall -br -in <driver model></pre> Where, <i><driver model></i> is one of the following: ❖ qla2500 ❖ qla2400 ❖ qla2300 ❖ qla2xxx e. Reboot the host to load the external QLogic driver. Emulex: a. Run the <code>lpfc-install</code> script to install the external Emulex driver on the SAN booted host. Note _____ No changes are required to the HBA and driver parameters. _____
3	To configure DM-Multipath, see “ Configuring DM-Multipath ” on page 89.

For SAN boot support for Red Hat Enterprise Linux 5 Update 1, you have to download the respective errata multipathing package from the Red Hat Network Web site (<http://www.redhat.com/software/rhn/>). For more information about the errata package, see the Notes section of the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Configuring SAN boot on a SUSE Linux Enterprise Server

Tasks for configuring SAN boot

To configure SAN boot with multipath, complete the following tasks.

Step	Action
1	Install the operating system on the LUN configured on a storage controller with root that is on a single SCSI device.
2	Install Host Utilities.
3	Configure the external or bundled HBA driver.
4	Configure root with DM-Multipath.

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

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

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

Installing Host Utilities

Install Linux Host Utilities 5.0. For information about installing the Linux Host Utilities 5.0, see “[Downloading and Installing the Host Utilities Software](#)” on page 4.

Installing and configuring the bundled HBA driver

To install and configure the bundled HBA driver, see “[Setting the required HBA and bundled driver parameters](#)” on page 26.

Installing and configuring the external HBA driver

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

Step	Action
1	Download and install the external HBA driver and applications. For more information, see “Downloading and installing the external HBA driver” on page 24. QLogic: To install the external QLogic driver on the SAN booted host, complete the following the steps. a. Run the qlinstall script with the -up option. b. Reboot the host to load the external QLogic driver.

Step	Action
2	Set the required HBA and driver parameters. For more information about setting the QLogic HBA parameters, see “QLogic HBA parameters” on page 31. To set the external HBA driver parameters, perform the following steps. - Edit the <code>/etc/modprobe.conf.local</code>. - Append the <code>options qla2xxx</code> statement with settings in the <code>modprobe.conf</code> file on a single line; only the <code>options qla2xxx</code> statement is allowed. <pre>ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0</pre> - Save the file and exit. Emulex: - Run the <code>lpfc-install</code> script to install the external Emulex driver on the SAN booted host. Note _____ No changes are required to the HBA and driver parameters.

Configuring root with DM-Multipath

To configure root with DM-Multipath, complete the following steps.

Step	Action
1	You have to configure DM-Multipath support. For more information about configuring DM-Multipath support, see “Configuring DM-Multipath” on page 89.
2	Install Linux Host Utilities 5.0. For information about installing Linux Host Utilities 5.0, see “Downloading and Installing the Host Utilities Software” on page 4.

Step	Action
3	Enable all paths to the LUN on the storage controller.
4	Re-create the initrd by running the following command: <pre>mkinitrd -f mpath</pre> Note_____ Keep a backup of the initrd image and make necessary modifications in the bootloader. _____
5	Reboot the host. Note_____ As all paths to the SAN boot LUN are now available on the host, configure all paths in the HBA boot BIOS as primary boot device, secondary boot device, and so on. _____

About this chapter This chapter provides instructions for implementing SAN boot on Linux hosts.

Note

SAN boot is currently supported on Red Hat Enterprise Linux 5 Update 2 and later, SUSE Linux Enterprise Server 10 SP1 and later.

Topics in this chapter

This chapter has the following topics:

- ◆ [“Configuring SAN boot on Red Hat Enterprise Linux”](#) on page 180
- ◆ [“Configuring SAN boot on SUSE Linux Enterprise Server”](#) on page 185

Configuring SAN boot on Red Hat Enterprise Linux

Implementing SAN boot (root partition on multipath) requires significant configuration and the process needs to be carefully managed.

You should log in to the storage system console or the Web interface of the storage system as you proceed with the configuration steps.

This section provides you information about the prerequisites for setting up root partition on multipath. The process of setting up root partition on multipath includes the following steps.

- 1. The initial setup tasks. For more information, see “[Initial setup tasks](#)” on page 180.
- 2. Selecting the partition. For more information, see “[Selecting the partition](#)” on page 182.
- 3. Configuration settings. For more information, see “[Configuring the system](#)” on page 184.

Root partition on multipath prerequisites

The prerequisites for root partition on multipath on Red Hat Enterprise Linux 5 Update 2 and later are as follows:

- ◆ You should implement the Operating System boot (kernel) from a local disk or PXE server. For more information, see [Step 8](#) of “[Selecting the partition](#)” on page 182.
- ◆ You should ensure that you have configured your storage system as per iSCSI requirements.

Initial setup tasks

Run the installation of Red Hat Enterprise Linux.

After the installer starts, complete the following steps.

Step	Action
1	In the initial installation page, select Installation.

2	Specify Boot Option as follows: linux mpath Press Enter.
3	Select the Language of your choice, and select OK.
4	Select your keyboard type and select OK.
5	Select your network device.
6	Select your network device for installation and select OK.
7	The Red Hat Enterprise Linux 5 GUI screen appears. Click Next.
8	The storage configuration page appears. Enter Red Hat Enterprise Linux installation number if any. Click OK.
9	Select a partitioning layout as Create custom layout.
10	Click Advanced storage configuration.
11	Select Add iSCSI target and click Add drive.
12	Enter Target IP address. Enter the proper iSCSI initiator name. Note _____ Ensure that you associate this iqn number with proper privileges on the storage controller. _____
13	On the storage controller, create an igroup with the initiator name used in Step 12 .
14	Create a LUN on the storage system on which you intend to create root partition, and map it to the igroup.
15	Return to the host screen.

16	Click Add Target in Configure iSCSI Parameters window. Adding the target discovers the target portal. Note_____ Ensure that multiple target portals are discovered, because the Red Hat installer does not identify the iSCSI device as a multipathed device unless it has more than one paths. _____
17	To discover more target portals, repeat Step 8 to Step 16 .
18	You should now see a multipathed iSCSI device listed in the drives section. Note_____ If the iSCSI multipathed device is not listed, check the configuration. _____ Click Next. Note_____ After Step 18, you can proceed with the installation process and enter choices until you reach the Installation Summary page. See Red Hat Enterprise Linux documentation to guide you through the rest of the installation procedure. _____

Selecting the partition

To select the partition, complete the following steps.

Step	Action
1	In the Partitioner page, you can view the list of local disks and multipathed device.
2	Create root partition.

3	In the Partitioner page, select the iSCSI multipathed device where you want to install the root file system. Click New. Note_____ Ensure that the LUNs are listed as multipathed devices in the Partitioner page. _____
4	The Add Partition window is displayed. Select the appropriate multipathed device from Allowable Drives.
5	Select the Mount Point as follows: /
6	Select the File System Type and partition size. Click OK.
7	Create a SWAP partition. You might create a SWAP partition on the same LUN that contains the root partition or on a different LUN. To create the partition, follow Step 3 to Step 6 and specify the File system value = SWAP and do not select Mount Point. Note_____ If you are using the software suspend functionality, ensure that the SWAP partition is on a local disk. _____
8	Create the /boot partition. You can create a /boot partition on a locally attached disk or use a PXE server to load the kernel boot image. To create the partition, follow Step 3 to Step 6 and specify the Mount Point as /boot. Note_____ It is best to install the /boot on a locally attached disk. _____
9	Click Next.

Configuring the system

To configure the system before beginning the installation process, complete the following steps.

Step	Action
1	Boot Loader installation: Ensure that the grub boot loader is installed on the master boot record (MBR) of the local disk on which /boot is installed. Click Next.
2	Network Device Configuration: Ensure that you have at least one network device configured which can talk to the storage controller. Click Next.
3	Select a region. Click Next.
4	Enter root user password. Click Next.
5	Select the required packages and start installation. Note _____You should select the iSCSI package (iscsi-initiator-utils-xxx) and the multipath package (device-mapper-multipath-xxx) manually. _____
6	After installation, click Reboot.
7	To configure DM-Multipath, see “ Configuring DM-Multipath ” on page 89.

Note_____If you rediscover the targets after root partition on multipath configuration, ensure that all the logged in connections are automatic.

Configuring SAN boot on SUSE Linux Enterprise Server

Implementing SAN boot (root partition on multipath) requires significant configuration and the process needs to be carefully managed.

You should log in to the storage system console or the Web interface of the storage system as you proceed with the configuration steps.

This section provides you information about the prerequisites for setting up root partition on multipath. The process of setting up root partition on multipath includes the following steps.

1. The initial setup tasks. For more information, see [“Initial setup tasks”](#) on page 180.
2. Selecting the partition. For more information, see [“Selecting the partition”](#) on page 187.
3. Setting up the bootloader. For more information, see [“Setting up the bootloader”](#) on page 189.
4. Configuring root partition on multipath. For more information, see [“Configuring root partition on multipath”](#) on page 190.

Root partition on multipath prerequisites

The prerequisites for root partition on multipath on SUSE Linux Enterprise Server 10 SP1 and SUSE Linux Enterprise Server 10 SP2 are as follows:

- ◆ You should implement the OS boot (kernel) from a local disk or PXE server. For more information, see [Step 15](#) of [“Selecting the partition”](#) on page 187.
- ◆ You should ensure that you have configured your storage system as per iSCSI requirements.

Initial setup tasks

Run the installation of SUSE Linux Enterprise Server.

After the installer starts, complete the following steps.

Step	Action
1	In the initial installation page, select Installation.

Step	Action
2	Specify Boot Option as follows: <code>linux withiscsi=1 netsetup=1</code>
3	Press Enter. You can view the list of the available network devices.
4	Select your network device and configure the network on it.
5	Select the Language of your choice, and click Next.
6	The License Agreement appears. Click Next. The iSCSI Initiator Overview page appears.
7	Select the Service tab and set the iqn number.
8	On the storage controller, create an igroup with the initiator name used in Step 7 .
9	Create a LUN on the storage system on which you intend to create root partition, and map it to the igroup.
10	Return to the host screen.
11	Select the Connected Targets tab. Click Add.
12	Perform the following (on the iSCSI Initiator Discovery page): a. Specify the IP address of the storage system. b. Specify the port. The default is 3260. c. Enter the credentials if you are using an authentication mechanism. d. Click Next.

Step	Action
13	In the list of storage systems that are discovered, click Connect for each, including authentication credentials if necessary. Note_____ During installation, only one path to the root LUN should be enabled. _____ Click Next.
14	Verify that the value for Connected is <i>true</i> for all the targets and click Next. The Connected Targets pane lists all the targets.
15	Set the Start-up mode to automatic by using Toggle Start-up button:
16	Click Finish. Note_____ After Step 16, you can proceed with the installation process and enter choices until you reach the Installation Summary page. See SUSE Linux Enterprise Server documentation to guide you through the rest of the installation procedure. _____

Selecting the partition

To select the partition, complete the following steps.

Step	Action
1	In the Installation Settings page, select the Expert tab.
2	Click Partitioning.
3	Select the Create Custom Partition Setup option. Note_____ You can view the list of local disks and LUNs. _____
4	Select the Custom Partitioning (for experts) option.

5	In the Expert Partitioner page, select the LUN where you want to install the root file system. Click Create. Note the corresponding device name. Note_____ Ensure that the LUNs are listed as SCSI devices in the Expert Partitioner page. _____
6	The Disk to Partition window is displayed. Select the appropriate device (device name corresponding to the LUN you selected), and click OK.
7	The Partition Type window is displayed. Select Primary Partition.
8	Select the File system, Size, and other applicable options.
9	Select the Mount Point as follows: /
10	Click the Fstab Options button.
11	Select the Mount by Device ID option.
12	Ensure that you have the <i>hotplug</i> keyword in the Arbitrary Option Value text box, click OK.
13	The Create a Primary Partition page appears. Click OK.
14	Create a SWAP partition. You might create a SWAP partition on the same LUN that contains the root partition or on a different LUN. To do this, follow Step 1 to Step 7 in the section “Selecting the partition”, and specify the File system value = SWAP. Note_____ If you are using the software suspend functionality, ensure that the SWAP partition is on a local disk. _____

15	Create the /boot partition. You can create a /boot partition on a locally attached disk or use a PXE server to load the kernel boot image. To do this, follow Step 1 to Step 7 in the section “Selecting the partition”, and specify the Mount Point in Step 9 as /boot. Note_____ You should install the /boot partition on a locally attached disk. _____
16	After you return to the Expert Partitioner page, review the configuration. Click Finish. The Installation Settings page is displayed.

Setting up the bootloader

To set up the bootloader, complete the following steps.

Note_____

Complete the following steps if you have the kernel image installation (/boot) on a local disk. In the case of PXE, ensure that the proper kernel images are set on the PXE server.

Step	Action
1	Click the Booting tab.
2	Select Default Label.
3	Click Edit.
4	In Optional Kernel Command Line Parameter, ensure that all references to installer arguments are removed. Your final line should look something similar to <code>resume=/dev/sda1 splash=silent showopts</code>. Click OK.
5	Click the Boot Loader Installation tab.

6	In the Boot Loader Location pane, select the Boot from Master Boot Record option. Click Finish. You return to the Installation Settings page.
7	Review the configuration settings. Click Accept. The Confirm Installation page is displayed.
8	Click Install.

Configuring root partition on multipath

After SAN boot is configured on the host, you can configure root partition on multipath.

To configure SAN boot using root partition on multipath, complete the following steps.

Step	Action
1	Configure the DM-Multipath. For details on configuring the DM-Multipath, see “ Configuring DM-Multipath ” on page 89.
2	Configure iSCSI on the host. For details on configuring iSCSI on host, see “ Configuring the iSCSI initiator ” on page 49.
3	Enable the additional paths to the NetApp LUN (root LUN) and change the Start-Up mode to automatic by using YaST2.
4	Before mounting the file system, modify the network configuration file and change the value of STARTMODE to <i>nfsroot</i> . Modify /etc/sysconfig/network/ifcfg-eth-idxxxxxxxx file.

Step	Action
5	To create a new initrd image with the root partition on multipath enabled, run the following command: <pre>mkinitrd -f "iscsi mpath"</pre> Note_____ Ensure that you set Start-Up mode to automatic before re-creating the initrd image. _____
6	Change the Start-Up mode of iSCSI session fetching the root LUN to onboot by using following commands: <pre>iscsiadm -m node -T targetname -p ip:port -o update -n node.startup -v onboot</pre>
7	Reboot the host. Note_____ The system now boots with the multipath enabled on the root device. You can verify this by running the <code>mount</code> command and ensuring that the root partition is on a DM-Multipath device. _____

About this chapter

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

Note

SAN boot is currently supported on Red Hat Enterprise Linux 5 Update 3.

Topic in this chapter

This chapter includes the following topic:

- ◆ [“Configuring SAN boot on Red Hat Enterprise Server”](#) on page 194

Configuring SAN boot on Red Hat Enterprise Server

Tasks for configuring SAN boot

To configure SAN boot, complete the following tasks.

Step	Action
1	Install the operating system on the LUN configured on the storage controller.
2	Install Host Utilities.
3	Install and configure the bundled HBA driver.
4	Configure DM-Multipath.

Installing the operating system

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

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

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

Installing Host Utilities

Install Linux Host Utilities 5.0. For information about installing the Linux Host Utilities 5.0, see “[Downloading and Installing the Host Utilities Software](#)” on page 4.

Installing and configuring the bundled HBA driver

To install and configure the bundled HBA driver, see “[Setting the required HBA and bundled driver parameters](#)” on page 44.

Configuring DM-Multipath

To configure root with DM-Multipath, complete the following steps.

Step	Action
1	Configure DM-Multipath. For more information about configuring DM-Multipath, see “ Editing the configuration file ” on page 91.
2	To activate multipathing during the boot process, enter the following command on the SAN booted host: chkconfig multipathd on

About this chapter

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

To determine which version of Veritas Storage Foundation and ASL to use, check the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Overview of the ASL and APM](#)” on page 198
- ◆ “[Installing the ASL and APM](#)” on page 199
- ◆ “[Removing the ASL and APM](#)” on page 201
- ◆ “[Understanding ASL error message types](#)” on page 203

Overview of the ASL and APM

What the ASL is

The 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 APM is

The 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 the system requirements
See the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

- ◆ Obtaining the ASL and APM
- ◆ Installing the latest version of the ASL and APM

Note

ALUA is not supported with Veritas Storage Foundation.

Obtaining the ASL and APM software

The ASL and APM are not distributed with the Host Utilities software. You can obtain the ASL and APM from the Symantec Web site. For a direct link to the ASL and APM on the Symantec Web site, see the NetApp Interoperability Matrix at:

<http://now.netapp.com/NOW/products/interoperability>

Download this software before you start the installation.

Installing the downloaded ASL and APM software

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

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

Step	Action
4	If your host is connected to a NetApp storage system, verify the installation by following these steps: a. Run the following command: <code>vxddmpadm listenclosure all</code> b. Locate NetApp under ARRAY_TYPE in the command output. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. Example: The <code>vxddmpadm listenclosure all</code> command shows the Enclosure Type as FAS2020 in this example. <pre># vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE LUN_COUNT ===== fas20200 FAS2020 30003571 CONNECTED A/P-C-NETAPP 4</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 category includes local non-FC attached disks. Volumes created on these arrays remain accessible because they do not have multipath support.

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

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

Step	Action						
1	Log in to the host as the root user.						
2	Stop all I/O to LUNs configured on storage controllers. Note _____In a Storage Foundation RAC cluster, you should also stop clustering on a node before you remove the ASL and APM. _____						
3	Enter the following command to remove the ASL package: <code>rpm -ev asl_rpm_name</code> Examples: The following command removes a previous version of the ASL:						
	<table><tr><th>If you are using...</th><th>Then...</th></tr><tr><td>Red Hat Enterprise Linux host</td><td><code>rpm -ev VRTSNTAPas1-5.0-3.0_RHEL5</code></td></tr><tr><td>SUSE Linux Enterprise Server host</td><td><code>rpm -ev VRTSNTAPas1-5.0-3.0_SLES10</code></td></tr></table>	If you are using...	Then...	Red Hat Enterprise Linux host	<code>rpm -ev VRTSNTAPas1-5.0-3.0_RHEL5</code>	SUSE Linux Enterprise Server host	<code>rpm -ev VRTSNTAPas1-5.0-3.0_SLES10</code>
	If you are using...	Then...					
	Red Hat Enterprise Linux host	<code>rpm -ev VRTSNTAPas1-5.0-3.0_RHEL5</code>					
SUSE Linux Enterprise Server host	<code>rpm -ev VRTSNTAPas1-5.0-3.0_SLES10</code>						

Step	Action	
4	Enter the following command to remove the APM package: <code>rpm -ev apm_rpm_name</code> Example: The following command removes a previous version of the APM:	
	If you are using...	Then...
	Red Hat Enterprise Linux host	<code>rpm -ev VRTSNTAPapm-5.0-1.0_RHEL5</code>
	SUSE Linux Enterprise host	<code>rpm -ev VRTSNTAPapm-5.0-1.0_SLES10</code>

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 the 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.	

About this Chapter

This chapter applies to Red Hat Enterprise Linux and SUSE Linux Enterprise Server hosts.

Topics in this chapter

This chapter discusses the following sections:

- ◆ [“Setting the required HBA and bundled driver parameters”](#) on page 205
- ◆ [“Loading the bundled HBA driver on the host”](#) on page 205
- ◆ [“Loading the external HBA driver on the host”](#) on page 205
- ◆ [“Accessing LUNs with Veritas DMP”](#) on page 206
- ◆ [“Obtaining multipathing information”](#) on page 206
- ◆ [“Unloading the external HBA driver and shutting down”](#) on page 209

Setting the required HBA and bundled driver parameters

You need to set both HBA and bundled driver parameters for Veritas to ensure that the multipathing and storage system failover work correctly. For more information on setting the HBA and bundled driver parameters, see [“Setting the required HBA and bundled driver parameters”](#) on page 26.

Loading the bundled HBA driver on the host

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

For more information about loading external HBA driver, see [“Loading the bundled HBA driver on the host”](#) on page 34.

Loading the external HBA driver on the host

After you have at least one LUN mapped to the Linux host, you can load the external HBA driver. When the external HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configure Veritas DMP, multipath devices are created for all the LUNs that are discovered. For more information about loading external HBA driver, see [“Loading the external HBA driver on the host”](#) on page 35.

Accessing LUNs with Veritas DMP

For details about configuring Veritas Volume Manager and accessing LUNs, see the Veritas Volume Manager Administrator's Guide for Linux that is shipped along with the software.

The vxdump restore daemon should be configured with an interval of 60. To configure the daemon, run the following command:

```
vxdump tune dmp_restore_interval 60
```

On reboot, this value comes into effect and remains persistent.

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

```
vxdumpadm stop restore
```

```
vxdumpadm start restore interval=60
```

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

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

```
vxdumpadm stat restored
```

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 a storage system, and simultaneous I/O is supported on each. 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 more information about system management, see the *Veritas Volume Manager Administrator's Guide*.

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

Step	Action
1	Enter the following command to view all devices: <pre>vxdisk list</pre> Result: The VxVM management interface displays the vxdisk device, type, disk, group, and status. It also shows which disks are managed by VxVM. Example: The following is a sample output of the vxdisk list command: <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS disk_0 auto:none - - online invalid disk_1 auto - - error fas30200_0 auto - - error fas30200_1 auto:cdsdisk fas30200_1 sharedg online shared fas30200_2 auto:cdsdisk fas30200_79 sharedg1_2 online shared fas30200_3 auto:cdsdisk fas30200_80 sharedg1_2 online shared fas30200_4 auto:cdsdisk fas30200_2 sharedg online shared fas30200_5 auto:cdsdisk fas30200_3 sharedg online shared fas30200_6 auto:cdsdisk fas30200_101 sharedg3_4 online shared fas30200_7 auto:cdsdisk fas30200_102 sharedg3_4 online shared</pre> Note_____ For Veritas Storage Foundation 5.0 MP1 and MP2, the ASL displays the enclosure-based naming disk objects in uppercase. For Veritas Storage Foundation 5.0 MP3, the default behavior of the ASL is to display the enclosure-based naming disk objects in lowercase. To change the enclosure names to uppercase, run the following command: <pre>vxddladm set namingscheme=ebn lowercase=no</pre>

Step	Action
2	On the host console, enter the following command to display path information for the device you want: <pre>vxddmpadm getsubpaths dmpnodename=device</pre> <i>device</i> is the name listed under the output of the <code>vxdisk list</code> command. Example: The following is a sample output: <pre># vxddmpadm getsubpaths dmpnodename=fas20200_1 NAME STATE[A] PATH-TYPE[M] CTRL-NAME ENCLR-TYPE ENCLR-NAME ATTRS ===== sdb ENABLED SECONDARY c1 FAS2020 fas20200 - sdf ENABLED(A) PRIMARY c1 FAS2020 fas20200 - sdj ENABLED SECONDARY c2 FAS2020 fas20200 - sdn ENABLED(A) PRIMARY c2 FAS2020 fas20200 -</pre>
3	To obtain path information for a host HBA, enter the following command: <pre>vxddmpadm getsubpaths ctrlr=controller_name</pre> <i>controller_name</i> is the controller displayed under “CTRL-NAME” in the output of the <code>vxddmpadm getsubpaths dmpnodename</code> command. Result: The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to. Example: The following is a sample output: <pre># vxddmpadm getsubpaths ctrlr=c1 NAME STATE[A] PATH-TYPE[M] DMPNODENAME ENCLR-TYPE ENCLR-NAME ATTRS ===== sdb ENABLED SECONDARY fas20200_1 FAS2020 fas20200 - sdf ENABLED(A) PRIMARY fas20200_1 FAS2020 fas20200 -</pre>

Using the sanlun command to obtain multipathing information: To obtain multipathing information by using the `sanlun` command, complete the following step.

Step	Action																																								
1	On the host command line, enter the following command: <pre>sanlun lun show -p</pre> Result: sanlun displays path information for each LUN. Example: The following is a sample output of the sanlun lun show -p command: <pre># sanlun lun show -p f2020-200-85:/vol/vol_c/c1 (LUN 1) Lun state: GOOD Lun Size: 100m (104857600) Filer_CF_State: Cluster Enabled Protocol: FCP Filer Partner: f2020-200-86 DMP Nodename: fas20200_1 Multipath-provider: VERITAS</pre> <table><tr><th>Veritas path state</th><th>sanlun path state</th><th>Filer Path type</th><th>/dev/ node</th><th>Host HBA</th><th>Primary Filer port</th><th>Partner Filer port</th><th>Veritas IO Policy</th></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdf</td><td>host1</td><td>0a</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDARY</td><td>sdb</td><td>host1</td><td>--</td><td>0a</td><td>Round-Robin</td></tr><tr><td>ENABLED(A)</td><td>GOOD</td><td>PRIMARY</td><td>sdn</td><td>host2</td><td>0b</td><td>--</td><td>Round-Robin</td></tr><tr><td>ENABLED</td><td>GOOD</td><td>SECONDARY</td><td>sdj</td><td>host2</td><td>--</td><td>0b</td><td>Round-Robin</td></tr></table>	Veritas path state	sanlun path state	Filer Path type	/dev/ node	Host HBA	Primary Filer port	Partner Filer port	Veritas IO Policy	ENABLED(A)	GOOD	PRIMARY	sdf	host1	0a	--	Round-Robin	ENABLED	GOOD	SECONDARY	sdb	host1	--	0a	Round-Robin	ENABLED(A)	GOOD	PRIMARY	sdn	host2	0b	--	Round-Robin	ENABLED	GOOD	SECONDARY	sdj	host2	--	0b	Round-Robin
Veritas path state	sanlun path state	Filer Path type	/dev/ node	Host HBA	Primary Filer port	Partner Filer port	Veritas IO Policy																																		
ENABLED(A)	GOOD	PRIMARY	sdf	host1	0a	--	Round-Robin																																		
ENABLED	GOOD	SECONDARY	sdb	host1	--	0a	Round-Robin																																		
ENABLED(A)	GOOD	PRIMARY	sdn	host2	0b	--	Round-Robin																																		
ENABLED	GOOD	SECONDARY	sdj	host2	--	0b	Round-Robin																																		

Unloading the external HBA driver and shutting down

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

About this chapter This chapter provides instructions for implementing SAN boot on Linux hosts.

Topics in this chapter

This chapter has the following topics:

- ◆ [“Configuring SAN boot on Red Hat Enterprise Server”](#) on page 212
- ◆ [“Configuring SAN boot on SUSE Linux Enterprise Server”](#) on page 215

Configuring SAN boot on Red Hat Enterprise Server

Configuring SAN boot on the host

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

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

Step	Action
8	After installing the operating system, reboot the host. As a result, the host boots from the SAN boot LUN on the storage system through a nonproxy path.
9	Install Linux Host Utilities 5.0. For information about installing the Linux Host Utilities 5.0, see “Downloading and Installing the Host Utilities Software” on page 4.
10	Install the supported version of the drivers, if you are using external HBA drivers. For more information about the installation and configuration of the drivers, see “Downloading and installing the external HBA driver” on page 24.
11	Verify the HBA settings. For more information, see “Verifying the HBA settings” on page 36.
12	Install Veritas Storage Foundation, any relevant patches or fixes, and supported ASL and APM.
13	Configure the vxdump restore daemon by setting it to an interval of 60. To configure the daemon, run the following command: <pre>vxdmptune dmp_restore_interval 60</pre> On reboot this value comes into effect and remains persistent.
14	Encapsulate the root disk for use in VxVM. To encapsulate the disk, run the following command: <pre>vxdiskadm</pre> For the detailed steps, see the section on encapsulating the disk in <i>Verita Volume Manager Administrator’s Guide</i> for Linux that is shipped along with the software.
15	Reboot the host after encapsulation.
16	Verify the encapsulation by issuing the following command: <pre>vxprint</pre> This command displays the rootvol and swapvol under the corresponding disk group.

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

Configuring SAN boot on SUSE Linux Enterprise Server

Configuring SAN boot on the host

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

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

Step	Action
9	Install Linux Host Utilities 5.0. For information about installing the Linux Host Utilities 5.0, see “Downloading and Installing the Host Utilities Software” on page 4.
10	Install the supported version of the drivers, if you are using external HBA drivers. For more information about the installation and configuration of the drivers, see “Downloading and installing the external HBA driver” on page 24.
11	Verify the HBA settings. For more information, see “Verifying the HBA settings” on page 36.
12	Install Veritas Storage Foundation, any relevant patches or fixes, and the respective ASL and APM.
13	Configure the <code>vxddmp restore</code> daemon by setting it to an interval of 60. To configure the daemon, run the following command: <pre>vxddmptune dmp_restore_interval 60</pre> On reboot this value comes into effect and remains persistent.
14	Encapsulate the root disk for use in VxVM. To encapsulate the disk, run the following command: <pre>vxdiskadm</pre> For the detailed steps, see the section on encapsulating the disk in <i>Veritas Volume Manager Administrator's Guide</i> for Linux that is shipped along with the software.
15	Reboot the host after encapsulation.
16	Verify the encapsulation by issuing the following command: <pre>vxprint</pre> This command displays the rootvolume and swapvolume under the corresponding disk group.

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

Troubleshooting

For a list of known problems and limitations, check the *Release Notes* for Host Utilities 5.0.

Topics in this appendix

This appendix includes the following topics:

- ◆ [“Troubleshooting—FC”](#) on page 220
- ◆ [“Troubleshooting—iSCSI”](#) on page 224
- ◆ [“Troubleshooting—Veritas Storage Foundation”](#) on page 229
- ◆ [“LUNs not visible on the host”](#) on page 231

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

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

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

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

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

Linux hosts with QLogic initiators fail to detect LUNs

Problem: Linux 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, even though both are set 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 host, ensure that the following settings are done for port topology and speed using the SANsurfer scli package:

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

To do this, perform the following steps:

- a. Run the following command:

```
/opt/QLogic_Corporation/SANsurferCLI/scli
```

- b. Select Configure HBA Settings from the Main Menu.
- c. For each of the WWPNs listed here, select ConnectionOptions and set it to Point to Point Only or Loop Only as required.

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

To do this, perform the following steps:

- a. Run the following command:

```
/opt/QLogic_Corporation/SANsurferCLI/scli
```

- b. Select the Configure HBA Settings from the Main Menu.

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

LUNs not recognized by DM-Multipath

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

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

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

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

For information about blacklisting devices on a Red Hat Enterprise Linux host, see [“Accessing LUNs with DM-Multipath”](#) on page 104.

For more information about blacklisting devices on a SUSE Linux Enterprise Server host, see “[Accessing LUNs with DM-Multipath](#)” on page 104.

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

```
multipath -v3
```

For more information, see bug number 228744 on Bugs Online at:

<http://now.netapp.com/cgi-bin/bol>

FC performance is poor

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

Workaround: Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. For specific steps, see 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 find /usr/lib/libHBAAPI.so library  
Ensure the package installing the library is installed & loaded
```

Workaround: The `libHBAAPI.so` file is included with the QLogic SANsurfer Pro package and the Emulex applications package. Install the appropriate package for your HBA model. See “[Downloading and installing the external HBA driver](#)” on page 24.

LUNs unavailable after storage system failover or giveback

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

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

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

```
/etc/init.d/multipathd start  
multipath -v3
```

sanlun commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts

Problem: The `sanlun` commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts.

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

Download the 32-bit compatibility Red Hat Package Manager (RPM) for `device-mapper-multipath` from the Novell Web site at:

<http://support.novell.com/techcenter/psdb/8ac2a27b3b63ac95fd776298026596eb.html>

Install the RPM which provides the 32-bit `libdevmapper.so` compatibility library. When `sanlun` is run, it links to the 32-bit library.

SCSI rescan issues on SUSE Linux Enterprise Server hosts

Problem: On a SUSE Linux Enterprise Server host connected to a storage controller, the presence of dummy LUNs causes issues during a SCSI rescan. Dummy LUNs are those which show up as "UNKNOWN" in `/proc/scsi/scsi/` even when no LUNs are mapped to the host. They have a generic device node (`/dev/sg*` entry), but no corresponding disk node (`/dev/sd*` entry). Due to these dummy LUNs, subsequent SCSI rescans on the host fail to recognize newly added LUNs, unless preceded by a write to the `sysfs` delete attribute of the dummy LUN.

Workaround: This issue is not yet resolved in SUSE Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10, but is fixed from the SUSE Linux Enterprise Server 10 SP2 updated kernel v2.6.16.60-0.31.

OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 10 series

Problem: During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, and none of the nodes in the OCFS2 cluster is able to write or read the heartbeat values. As a result, the nodes might fence by panicking.

Workaround: Set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and the storage system, a reasonable value usually is 181. This value corresponds to a total of 360 seconds before a node self-fences.

Depending on the load in the iSCSI network, the value should be raised. It is best to use a value in the range of 61 to 201 for a configuration with OCFS2 on multipath, and iSCSI devices.

For more information, see the following:

http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT

"multipath -l" displays the LUN status as "undef" on SUSE Linux Enterprise Server 10 series

Problem: With DM-Multipath configured, the `multipath -l` output displays the LUN status as `undef`.

This status is displayed because of a bug in the command.

Workaround: The `undef` status does not indicate an error and can be safely ignored. Use the `multipath-ll` command which displays the LUN status correctly as `ready`.

Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 10 series

Problem: Currently, the volume groups created on iSCSI devices are not automatically scanned when iSCSI LUNs are discovered. Therefore during the system boot, the volume groups that were created on the iSCSI devices are unavailable.

Workaround: To overcome this problem, a helper script is provided in `/usr/share/doc/packages/lvm2/lvm-vg-to-udev-rules.sh`. This script can be used to generate udev rules for iSCSI Logical Volume Manager (LVM) devices, which can be automatically mounted during system boot.

Example:

- ◆ `/bin/bash /usr/share/doc/packages/lvm2/lvm-vg-to-udev-rules.sh/dev/vgname/lvname`. Here, `lvname` is the name of the logical volume that is created on top of volume group `vgname`.
- ◆ Add the generated rules to the `/etc/udev/rules.d/85-mount-fstab-lvm.rules` file.

After completing the preceding steps after each logical volumes, the logical volumes can be automatically mounted by adding the entries in `/etc/fstab`.

Automatically mounting and unmounting the file system on multipath device do not work on SUSE Linux Enterprise Server 9 series

Problem: Automated mounting of multipath-based file systems at boot time using entries in `/etc/fstab.conf` or `/etc/fstab.iscsi` does not work. This is because the placement of the multipath bootup and shutdown scripts in the `rc` script list in SUSE Linux Enterprise Server 9 series is out of order with respect to the mount and unmount scripts for iSCSI and non-iSCSI devices.

The `multipathd` and `iscsid` startup scripts are at the same serial number for a run level in `/etc/init.d/rc3.d`. Because alphabetically *i* comes before *m*, `iscsid` starts before `multipathd`.

As a result, the multipath bootup and shutdown sequence is not synchronized with the iSCSI bootup and shutdown sequence. Multipath devices are not available when the `mount` command is run, and the `mount` command fails to mount these devices.

Similarly in the shutdown path, by the time the unmount script runs, the multipath utility would have already been shut down.

Workaround: Manual mounting and unmounting is required for multipath-based file systems on SUSE Linux Enterprise Server 9 series. After the multipath devices are available on bootup, the file systems can be mounted.

For unmounting, ensure that any multipath-based file systems are unmounted before the multipath utility is shut down.

To ensure the proper shutdown of multipath-based iSCSI file systems, complete the following steps.

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

For more information, see bug number 196997 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 9 series

Problem: During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, and none of the nodes in the OCFS2 cluster is able to write or read the heartbeat values. As a result, the nodes might fence by panicking.

Workaround: Set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and the storage system, a reasonable value usually is 61. This value corresponds to a total of 120 seconds before a node self-fences.

Depending on the load in the iSCSI network, the value should be raised. It is best to use a value in the range of 61 to 201 that would be ideal for a configuration with OCFS2 on mutipath, and iSCSI devices.

For more information, see the following:

http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT

See also, bug 210340 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

Multipathd occasionally fails to update path status of DM-Multipath devices

Problem: Sometime multipathd does not update the path status of dm-multipath devices after running I/O faults. The reason is, the failure of multipathd to start "event checker" for some of the DM-Multipath devices. Because of this failure multipathd is not keeping track of path up or down status for those devices.

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

LUNs not recognized by DM-Multipath on SUSE Linux Enterprise Server

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

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

```
multipath -v3 -d | grep blacklist
```

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

If the iSCSI 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 number 201302 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>.

Storage system console displays continuous ISID Rule messages (SUSE Linux Enterprise Server 9 series)

Problem: The storage system console displays continuous ISID Rule messages, similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0

Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1

Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10

Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0

Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10

Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

Workaround: Verify that you added the following line to the `/etc/iscsi.conf` file. Ensure that this line is at the beginning of the file, before any “DiscoveryAddress” entries.

```
Continuous=no
```

iSCSI performance is poor on SUSE Linux Enterprise Server

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

Workaround: Manually, enter the disk geometry parameters when creating partitions with the `fdisk` command. For more details, see bug number 156121 on Bugs Online at:

<http://now.netapp.com/NOW/cgi-bin/bol>

Workaround: In a default setup, LVM scans all attached disks and determines which of them contain physical volumes. When using multipathing, multiple devices can contain identical physical volumes. If LVM detects the SCSI device first, the underneath device is shown as SCSI device.

Modify the filter in `/etc/lvm/lvm.conf` file to only scan the multipath device and not the SCSI device underneath. `filter = [ "rl/dev/.*by-path/.*|", "rl/dev/.*by-id/.*|", "rl/dev/sda2|", "rl/dev/sd*|", "a./.*" ]`

Troubleshooting—Veritas Storage Foundation

Error when stopping the fencing driver

If you get an error when 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 the fencing driver does not stop, remove the name of the coordinator diskgroup from <code>/etc/vxfendg</code> .
4	In the <code>/etc/vxfenmode</code> file, make the following change: <code>vxfen_mode=disabled</code>
5	Stop the fencing driver by using the following command: <code>/etc/init.d/vxfen stop</code>

Note

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

Secondary paths identified incorrectly

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

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

Enclosure-based naming not reflected on NetApp storage

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

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

**SAN boot
installation to LUN
configured on a
storage controller
fails**

Problem: SAN boot fails if you have multiple paths configured to boot LUN before OS installation.

Workaround: Verify that the host can only see one path to the LUN configured on a storage controller during installation.

**Encapsulation of
root disk fails**

Problem: Encapsulation of root disk fails after SAN boot installation

Workaround: Ensure that you follow the partitioning scheme given in the *Veritas Volume Manager Administrator's Guide* for Linux.

LUNs not visible on the host

iSCSI LUNs appear as local disks to the host. If the storage system LUNs are not available as disks on the hosts, verify the following configuration settings:

Configuration setting	What to do
Driver version	Verify that the external HBA driver version you are running is supported. You can check the most up-to-date system requirements in the FC Host Compatibility Matrix at the following URL: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Cabling	Verify that the cables between the host and the storage system are properly connected.
Zoning	Verify that the zoning is correct. For information about zoning see “ Switched fabric configurations ” on page 86.
Network connectivity	Verify that there is TCP/IP connectivity between the host and the storage system. From the storage system command line: ◆ Ping the host interfaces that are used for iSCSI. ◆ Ping the storage system interfaces that are used for iSCSI.
HBA Setting	Verify the recommended HBA settings. For information about verifying the HBA settings, see “ Verifying the HBA settings ” on page 36.
FC service status	Verify that the FC service is licensed and started on the storage system according to the procedure described in the <i>Data ONTAP Block Access Management Guide</i> .
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system according to the procedure described in the <i>Data ONTAP Block Access Management Guide</i> .

Configuration setting	What to do
Initiator login for FC	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. If the initiator is configured and logged in to the storage system, the storage system console displays the initiator WWPN.
HBA World Wide Port Names	Verify that you are using the correct initiator WWPNs in the igroup configuration. On the storage system, use the <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 should match.
Initiator login for iSCSI	Verify that the initiator is logged in to the storage system by entering the <code>iscsi show initiator</code> command on the storage system console. If the initiator is configured and logged in to the storage system, the storage system console displays the initiator node name and the target portal group to which it is connected. If the command output shows that no initiators are logged in, check the initiator configuration on the host. Verify that the storage system is configured as a target of the initiator. For the detailed information, see “Editing the host’s iSCSI configuration file” on page 56.

Configuration setting	What to do
iSCSI node names	Verify that you are using the correct initiator node names in the igroup configuration. On the storage system, use the <code>igroup show</code> command to display the node name of the initiators in the storage system's igroups. On the host, use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host should match.
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system, use one of the following commands: ◆ <code>lun show -m</code> displays all the LUNs and the igroups to which they are mapped. ◆ <code>lun show -g igroup-name</code> displays the LUNs mapped to a specific igroup.
System requirements	Verify that the components of your configuration are supported. Verify that you have the correct host Operating System (OS) Service Pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up-to-date system requirements in the NetApp Interoperability Matrix at: http://now.netapp.com/NOW/products/interoperability
Jumbo frames	If you are using jumbo frames in your configuration, ensure that the jumbo frames are enabled on all the devices in the network path: the host Ethernet NIC, the storage system, and any switches.
Firewall settings	Verify that the iSCSI port(3260) is open in the firewall rule.

Default and Recommended Settings

Default versus recommended settings

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

Component	Parameter	Red Hat Enterprise Linux 4		Red Hat Enterprise Linux 5	
		Default value	NetApp recommended value	Default value	NetApp recommended value
Qlogic external driver for FC	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6
Qlogic bundled driver for FC	qlport_down_retry	0	30	30	30
	ql2xsuspendcount	10	6	Not applicable	Not applicable
Qlogic FC HBA	Link Down Timeout	30	20	30	20
QLogic iSCSI HBA	IP_ARP_Redirect	Not applicable	Not applicable	off	on
	KeepAliveTO	Not applicable	Not applicable	30	30
Multipath.conf	user_friendly_names	no	yes	no	yes
Veritas DMP	dmp_restore_interval	300	60	300	60

Note

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

On Red Hat Enterprise Linux 5, the `user_friendly_names` feature is disabled by default. However, it is enabled by default in the Red Hat Enterprise Linux 5 Anaconda installer for enabling the SAN boot root on the DM-Multipath feature.

For Emulex, the default settings work fine, there are no recommended settings for Emulex.

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

Component	Parameter	SUSE Linux Enterprise Server 9		SUSE Linux Enterprise Server 10	
		Default value	NetApp recommended value	Default value	NetApp recommended value
QLogic external driver for FC	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6
QLogic bundled driver for FC	qlport_down_retry	0	30	0	30
	ql2xsuspendcount	10	6	Not applicable	Not applicable
	MaxRetriesPerIo	0	10	Not applicable	Not applicable
	MaxRetriesPerPath	0	10	Not applicable	Not applicable
QLogic FC HBA	Link Down Timeout	30	20	30	20
Multipath.conf	<code>user_friendly_names</code>	no	no	no	no
Veritas DMP	<code>dmp_restore_interval</code>	300	60	300	60

About this information

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

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

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

Copyrights and licenses

The following components are subject to the Perl Artistic License:

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

Perl Artistic License:

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

Definitions:

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

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

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

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

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

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

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

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

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

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

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

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

Portions copyright © 2001 Microsoft Corporation. All rights reserved.

SNIA Public License 1.1:

1. Definitions.

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

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

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

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

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

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

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

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

1.9 “License” means this document.

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

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

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

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

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

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

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

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

2. Source Code License.

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

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

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

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

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

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

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

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

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

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

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

3. Distribution Obligations.

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

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

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

3.4 Intellectual Property Matters.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

8. **TERMINATION.**

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

8.2 If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor upon whom You file

such action is referred to as “Participant”) alleging that:

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

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

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

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

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

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

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

EXHIBIT A—The SNIA Public License.

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

This file itself is the license.

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

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Contributor(s): _____.

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

NetSSLeay License:

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

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

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

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

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

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

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

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. The licence and distribution terms for any publicly available version or derivative of this code cannot be changed; i.e., this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.

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

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

GNU General Public License 2.0:

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

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

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

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

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

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

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

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

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION, AND MODIFICATION

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NO WARRANTY

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

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

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

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

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

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

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

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

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

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

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

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

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

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

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

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

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

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

Robin's Nest Software License:

This Software Provided By Robin's Nest Software Inc.

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

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

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

Copyright © Microsoft Corporation. All rights reserved.

Common Public License:

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

1. DEFINITIONS

"Contribution" means:

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

b) in the case of each subsequent Contributor:

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

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

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

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

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

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

2. GRANT OF RIGHTS

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

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

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

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

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

3. REQUIREMENTS

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

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

When the Program is made available in source code form:

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

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

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

4. COMMERCIAL DISTRIBUTION

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

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

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

5. NO WARRANTY

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

6. DISCLAIMER OF LIABILITY

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

SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

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

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

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

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

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

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

