

iSCSI SUSE® Linux® Host Utilities 3.2 for Native OS **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 215-03971_A0
June 2008

Copyright and trademark information

Copyright information

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

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

Software derived from copyrighted material of NetApp, Inc. is subject to the following license and disclaimer:

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

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

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

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

Trademark information

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

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

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

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

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

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

Table of Contents

	Preface	vii
Chapter 1	Overview	1
Chapter 2	Installing Host Utilities (SUSE Linux Enterprise Server 9 SP3)	7
	Installing the Host Utilities software	8
	Verifying or installing the required multipathing packages	10
	Configuring dm-multipath support	11
	Recording or changing the initiator node name on the host	14
	Configuring the iSCSI initiator	17
Chapter 3	Configuring the Storage System (SUSE Linux Enterprise Server 9 SP3) .	21
Chapter 4	Accessing LUNs from the Host (SUSE Linux Enterprise Server 9 SP3) .	23
	Starting the iSCSI service	24
	Starting the multipath service.	25
	Accessing LUNs without dm-multipath support	26
	Identifying and viewing information about LUNs	27
	Creating a partition	33
	Creating a file system on a partition of a LUN.	34
	Mounting the file system	35
	Accessing LUNs with dm-multipath support.	37
	Verifying the dm-multipath configuration	38
	Identifying and viewing a list of LUNs	39
	Creating and mounting a file system on a LUN	43
	Tuning the dm-multipath configuration	45
	Stopping dm-multipath	46
Chapter 5	Installing Host Utilities (SUSE Linux Enterprise Server 10 SP1 and SP2)	47
	Installing the Host Utilities software	48
	Verifying or installing the required packages	50

	Configuring dm-multipath support	51
	Recording or changing the initiator node name on the host	55
	Configuring the iSCSI initiator	58
Chapter 6 SP2)	Configuring the Storage System (SUSE Linux Enterprise Server 10 SP1 and 65	
Chapter 7 SP2)	Accessing LUNs from the Host (SUSE Linux Enterprise Server 10 SP1 and 67	
	Starting the iSCSI service	68
	Starting the multipath service.	69
	Accessing LUNs without dm-multipath support	70
	Identifying and viewing information about LUNs	71
	Creating a partition	78
	Creating a file system on a partition of a LUN.	79
	Mounting the file system	80
	Accessing LUNs with dm-multipath support.	82
	Verifying the dm-multipath configuration	83
	Identifying and viewing a list of LUNs	84
	Creating and mounting a file system on a LUN	88
	Tuning the dm-multipath configuration	90
	Stopping dm-multipath	91
Chapter 8 93	Implementing SAN Boot (SUSE Linux Enterprise Server 10 SP1 and SP2) .	
	SAN boot overview.	94
	Setting up the host for SAN boot	95
	Configuring root partition on multipath	102
Appendix A	Troubleshooting	105

Preface

About this guide

This guide describes how to configure the iSCSI initiator in SUSE® Linux® Enterprise Server to access a logical unit of storage identified by numbers (LUNs) on a NetApp® storage system. It also explains how to configure the storage system to work with the initiator.

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

Audience

This guide is for NetApp storage systems administrators who configure the iSCSI initiator in SUSE Linux Enterprise Server host to access NetApp storage system.

Terminology

This guide uses the following terms:

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

Command conventions

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

Formatting conventions

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

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcctest <i>adaptername</i></code> command, you enter the characters “fcctest” 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.

Overview

About this chapter

This chapter provides an overview of the following:

- ◆ How Data ONTAP software implements iSCSI
- ◆ The main features that Host Utilities supports
- ◆ The Host Utilities setup process

This chapter also helps you to find additional information about device mapper multipath support.

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. The RFC 3720 defines the iSCSI protocol standard, published by the Internet Engineering Task Force. You can get a copy of the standard from 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 the logical units of storage identified by numbers (LUNs) on a storage system running Data ONTAP® software. The host does not have a hardware iSCSI host bus adapter (HBA). 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 regular public network. The storage system listens for iSCSI connections on IP port 3260.

For more information on using iSCSI with your storage system, see the *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 SP1 and SP2 come with an Open-iSCSI software initiator. iSCSI SUSE Linux Host Utilities 3.2 (Host Utilities) supports the Open-iSCSI software initiator on SUSE Linux Enterprise Server 10 SP1 and SP2 hosts. See “[Where to go for more information](#)” on page 5 for links to Novell® documentation.

Root partition on multipath (SAN boot) support: This Host Utilities version supports implementation of root partition on multipath (SAN boot) on hosts running SUSE Linux Enterprise Server 10 SP1 and SP2. See [“Implementing SAN Boot \(SUSE Linux Enterprise Server 10 SP1 and SP2\)”](#) on page 93 for details.

Linux clustering support: This Host Utilities version supports the usage of cluster file systems and clustering suite on Linux hosts.

The following file system is supported:

- ◆ Oracle® Clustered File System (OCFS2)

The following Cluster suite is supported:

- ◆ Oracle Clusterware and Oracle Real Application Cluster (RAC) 10g R2

For information on OCFS2, Oracle Clusterware, and Oracle RAC, see the Oracle documentation Web site. For details see [“Where to go for more information”](#) on page 5.

Understanding device mapper multipathing

Linux device mapper multipathing (dm-multipath) is supported on NetApp storage systems with SUSE Linux Enterprise Server 9 SP3 and SUSE Linux Enterprise Server 10 SP1 and SP2. It provides IO failover and path load sharing for multipathed block devices. This support enables you to configure multiple iSCSI paths between the Linux host and storage system. If one path fails, iSCSI traffic continues on the remaining paths.

The required kernel modules for dm-multipath are included in the SUSE Linux Enterprise Server 9 SP3 and SUSE Linux Enterprise Server 10 SP1 and SP2 binary kernel packages. You might have to install additional user space packages as part of the setup process.

When you have multiple paths to a LUN, Linux creates a /dev/sdx 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. The dm-multipath support creates a single device in /dev/mapper for each LUN that represents all the paths. You will need to create file systems and mount the LUN using the device entries in /dev/mapper. You do not need to create partitions or labels when using LUNs with dm-multipath support; the device in /dev/mapper creates a persistent association with a LUN.

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

Note

Automatically mounting and unmounting the file systems on multipath devices do not work on SUSE Linux Enterprise Server 9 SP3. For details see, [“Automatically mounting and unmounting the file system on multipath device do not work on SUSE Linux Enterprise Server 9 SP3”](#) on page 106.

Setup procedure

The procedure for setting up the iSCSI protocol on a host and storage system follows the same basic sequence for all host types.

Step	Action
1	Install and configure Host Utilities software on the host, including the following: ◆ Installing required packages on the host, and recording or changing the host's iSCSI node name ◆ Optional: Configuring dm-multipath support ◆ Setting the initiator parameters, including the IP address of the target on the storage system ◆ Optional: Configuring Challenge Handshake Authentication Protocol (CHAP) See Chapter 2, “Installing Host Utilities (SUSE Linux Enterprise Server 9 SP3),” on page 7, or Chapter 5, “Installing Host Utilities (SUSE Linux Enterprise Server 10 SP1 and SP2),” on page 47 depending on the version of the operating system you are using.
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 Chapter 3, “Configuring the Storage System (SUSE Linux Enterprise Server 9 SP3),” on page 21, or Chapter 6, “Configuring the Storage System (SUSE Linux Enterprise Server 10 SP1 and SP2),” on page 65 depending on the version of the operating system you are using.

Step	Action
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 Chapter 4, “Accessing LUNs from the Host (SUSE Linux Enterprise Server 9 SP3),” on page 23, or Chapter 7, “Accessing LUNs from the Host (SUSE Linux Enterprise Server 10 SP1 and SP2),” on page 67 depending on the version of the operating system you are using.

Where to go for more information

The following table describes where to find additional information about using iSCSI on your storage system.

If you want to see...	Go to...
Changes in this release of Host Utilities, a list of known problems, and limitations of Host Utilities	The <i>Release Notes</i> for this Host Utilities.
The most current system requirements for the Linux host	The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability , or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml .
Best practices or configuration issues	The NetApp Knowledgebase at ◆ https://now.netapp.com/eservice/kbbrowse ◆ http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190
The supported storage system models for Data ONTAP licensed with iSCSI	◆ The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml. ◆ The <i>System Configuration Guide 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: ◆ Data ONTAP <i>Block Access Management Guide</i> ◆ Data ONTAP <i>Release Notes</i>
Information about managing the iSCSI initiator on SUSE Linux Enterprise Server 9 SP3	◆ 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/linux/suse/portal/.

If you want to see...	Go to...
Information about configuring multipathing on SUSE Linux Enterprise Server	The Novell Knowledgebase article <i>How to setup / use multipathing on SLES</i> at http://support.novell.com/techcenter/sdb/en/2005/04/sles_multipathing.html .
Information about configuring and managing the iSCSI initiator on SUSE Linux Enterprise Server 10 SP1 and SP2	◆ 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 SP1 and SP2 documentation on the Novell Web site at http://www.novell.com/documentation/sles10/index.html.
Information about Oracle Cluster File System	◆ <i>Oracle Cluster File System (OCFS2) User's Guide</i> at http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf. ◆ <i>OCFS2 - Frequently Asked Questions</i> at http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html.
Information about Oracle RAC	<i>Installing the Oracle RAC</i> at http://www.oracle.com/technology/pub/articles/smiley_rac10g_install.html .
The NetApp Knowledgebase for Oracle RAC	http://www.netapp.com/library/tr/3369.pdf . http://www.netapp.com/library/tr/3423.pdf .

About this chapter

This chapter explains how to install Host Utilities, and how to complete the installation and initial configuration of the iSCSI initiator and multipathing software on the Linux host.

Note

You will complete the final configuration of the initiator after you have configured the storage system.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the Host Utilities software”](#) on page 8
- ◆ [“Verifying or installing the required multipathing packages”](#) on page 10
- ◆ [“Configuring dm-multipath support”](#) on page 11
- ◆ [“Recording or changing the initiator node name on the host”](#) on page 14
- ◆ [“Configuring the iSCSI initiator”](#) on page 17

Requirement for root privileges

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

Installing the Host Utilities software

Verifying the correct version

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

Getting the Host Utilities software

Download the Host Utilities file (netapp_linux_host_utilities_3_2.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.

Installing the Host Utilities software

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

Step	Action
1	Remove any previous version of Host Utilities installed on your system. Change to the directory where the previous version is installed (default is /opt/netapp/santools) and enter the following command: <code>./uninstall</code>
2	Change to the working directory to which you saved the Host Utilities file.
3	To uncompress the file, enter the following command: <code>gunzip netapp_linux_host_utilities_3_2.tar.gz</code>
4	To extract the files, enter the following command: <code>tar -xvf netapp_linux_host_utilities_3_2.tar</code>
5	Change to the netapp_linux_host_utilities_3_2 directory. By default, this directory is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.
6	Enter the following command: <code>./install</code>

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

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

Verifying or installing the required multipathing packages

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

Step	Action	
1	Enter the following command: <pre>rpm -q linux-iscsi</pre> 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 support, enter the following commands: <pre>rpm -q device-mapper</pre> <pre>rpm -q multipath-tools</pre> Result: The <code>rpm -q</code> commands return the name and version of the installed packages. Install the correct version of the packages that are shipped along with the operating system.	
3	If...	Then...
	You have the correct packages for your version of SUSE Linux Enterprise Server.	Proceed to “ Configuring dm-multipath support ” on page 11.
	You do not have the correct packages for your version of SUSE Linux Enterprise Server.	Install the required package versions from the media for your version of SUSE Linux Enterprise Server. You can use the <code>rpm</code> or <code>yast</code> command.

Configuring dm-multipath support

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

Copying the file The multipath.conf file contains configuration settings for multipathing. A sample 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 multipath.conf.

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

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

Step	Action
1	Open /etc/multipath.conf with a text editor.
2	Add the following syntax for blacklisting devices: <code>devnode_blacklist {}</code>

Step	Action
3	Blacklist all devices that do not correspond to any NetApp LUN and are mapped to your Linux host (that is, blacklist those devices that are not displayed in the <code>sanlun lun show</code> command). - 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 that device. Example: Assume that <code>/dev/sda</code> is a local SCSI drive. To obtain its WWID, you will run the following command: <pre>scsi_id -gus /block/sda</pre> The output be as follows: <pre>SIBM-ESXSMAW3073NC_FDAR9P66067WJ</pre> Add this entry for blacklisting. - Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]" devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*" devnode "^cciss!c[0-9]d[0-9]*"</pre> Example: The following is an example for a SUSE Linux Enterprise Server 9 SP3 host: <pre>devnode_blacklist { devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ }</pre>
4	Save the changes.

Example: The following example shows the required changes to the `/etc/multipath.conf` file for a system with one local SCSI drive.

```
##
## name      : blacklist
## scope     : multipath & multipathd
## desc      : list of device names to discard as not multipath
               candidates
## default   : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
```

```

devnode_blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z][0-9]*"
    devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"
    wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ
}
#
##
devices {
    device {
        vendor                "NETAPP"
        product                "LUN"
        path_grouping_policy   multibus
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout           "/sbin/mpath_prio_ontap /dev/%n"
        features                "1 queue_if_no_path"
        path_checker            readsector0
        failback                immediate
    }
}

```

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 uses the following format on SUSE Linux Enterprise Server 9 SP3 hosts:

`iqn.1987-05.com.cisco:RandomNumber`

You must set the initiator node name to `iqn.1987-05.com.cisco:host.name` on SUSE Linux Enterprise Server 9 SP3 hosts, where *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.

Step	Action	
2	If...	Then...
	If a default initiator name has not been generated	a. Replace the text in the host's <code>/etc/initiatorname.iscsi</code> file with the initiator name in the following format: <pre>InitiatorName=ign.1987-05.com.cisco:RandomNumber</pre> The recommended value for <i>RandomNumber</i> is the host name of the Linux host. b. Record the initiator node name and proceed to “Configuring the iSCSI initiator” on page 17.
	You want to use the default initiator node name	Record the initiator node name and proceed to “ Configuring the iSCSI initiator ” on page 17.
	You want to change the initiator node name	a. Modify the <i>RandomNumber</i> part of the initiator node name. The following is a sample node name: <pre>ign.1987-05.com.cisco:01.82ea2060f487</pre> b. Record the initiator node name and proceed to “Configuring the iSCSI initiator” on page 17. Note _____ <i>RandomNumber</i> is the only component of the initiator node name you modify. You do not modify the other components of the initiator node name. You must use the host name in place of the default random number.

Node name rules

If you change the host's initiator node name, be sure the new name follows all of the following rules:

- ◆ An initiator node name can be up to 223 bytes.
- ◆ Upper-case characters are always mapped to lower-case characters.
- ◆ An initiator node name can contain alphabetic characters (a to z), numbers (0 to 9), and the following three special characters:
 - ❖ Period (.)
 - ❖ Hyphen (-)
 - ❖ Colon (:)
- ◆ The underscore character (_) is *not* supported.

Configuring the iSCSI initiator

SUSE Linux includes the iSCSI initiator software in the software distribution. The initiator’s components include a kernel module that 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 SP3, edit the /etc/iscsi.conf file. You edit the file to configure the following:

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

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

Step	Action
1	With a text editor, open the host’s /etc/iscsi.conf file.

Step	Action
2	At the beginning of the file, add the following configuration entries. Make sure 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 will experience problems with connectivity and performance.
3	If you are using dm-multipath support, perform the following actions: ◆ Remove the comment to enable the Multipath=portal in the Multipath Settings section. ◆ Remove the comment from the ConnFailTimeout in the Session Timeout Settings section and set the value to ConnFailTimeout=5.
4	To discover the NetApp storage systems, add the IP address of the storage system in 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. Specify an IP address that will be used for iSCSI communication. Gigabit Ethernet interfaces are strongly recommended. Example: The following are sample DiscoveryAddress entries using storage system IP addresses: <pre>DiscoveryAddress=192.168.10.100 DiscoveryAddress=10.61.208.200</pre>

Step	Action
5	If you want to use CHAP, add CHAP user names and passwords, as shown in the following example. CHAP settings must be indented below the Discovery Address of the storage system they apply to. For bidirectional CHAP, you must define outgoing and incoming user names and passwords. If you do not want to use bidirectional CHAP, you define only the outgoing user name and outgoing password. While adding the global CHAP entries in <code>/etc/iscsi.conf</code> file, the sequence must be as follows: ◆ Outgoing CHAP settings ◆ Discovery Addresses ◆ Incoming CHAP settings (in case of bidirectional CHAP) Example: <pre>OutgoingUsername=linuxhostout OutgoingPassword=aow9857f1 DiscoveryAddress=10.72.199.109 IncomingUsername=linuxhostin IncomingPassword=e09fj30</pre> Make sure 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>).

Before you begin

You need the host’s iSCSI initiator node name that you recorded in [Chapter 2](#), before configuring the storage system. You use the initiator node name to create igroups on the storage system.

Tasks for configuring the storage system

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

Step	Action
1	Ensure that the iSCSI service is licensed and started.
2	If you want to use CHAP authentication, use the <code>iscsi security</code> command or the FilerView® interface to configure a CHAP user name and password.
3	Create LUNs, create an igroup using the host’s initiator node name, and then map the LUNs to the igroup for the Linux host. You can manage LUNs using the Web-based FilerView interface or using the Data ONTAP command line. Be sure to specify the LUN type and iSCSI igroup type as <code>linux</code> .

For detailed configuration steps

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

Overview of accessing LUNs

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

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Starting the iSCSI service”](#) on page 24
- ◆ [“Starting the multipath service”](#) on page 25
- ◆ [“Accessing LUNs without dm-multipath support”](#) on page 26
- ◆ [“Accessing LUNs with dm-multipath support”](#) on page 37
- ◆ [“Stopping dm-multipath”](#) on page 46

Starting the iSCSI service

To start the iSCSI service on SUSE Linux Enterprise Server 9 SP3 hosts, enter the following command at the Linux host command prompt:

```
/etc/init.d/iscsi start
```

Configuring the iSCSI service to start automatically after reboot: If you want to configure the iSCSI service to start automatically at system bootup, complete the following steps.

Step	Action	
1	Verify the status of the iSCSI service by entering the following command: <code>chkconfig iscsi</code>	
2	If...	Then...
	The <code>chkconfig</code> command indicates the iSCSI service is enabled	No action is required.
	The <code>chkconfig</code> command indicates that the iSCSI service is not enabled	Start the iSCSI service by entering the following command: <code>chkconfig iscsi on</code>

Starting the multipath service

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

After you have started the iSCSI service, start the multipath service on the Linux host. Complete the following steps.

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

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

To integrate the multipath setup into the boot sequence, enter the following command on the Linux host:

```
chkconfig boot.multipath on
chkconfig multipathd on
```

Accessing LUNs without dm-multipath support

Tasks for accessing LUNs

After the initiator is configured, you need to set up the host to access the LUNs you created on the storage system. The following is an overview of the tasks you must complete to access LUNs.

1. Identify the LUN on which the partition is to be created. For more information, see “[Identifying and viewing information about LUNs](#)” on page 27.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN. For more information, see “[Creating a partition](#)” on page 33.
3. Create a file system on the partition. For more information, see “[Creating a file system on a partition of a LUN](#)” on page 34.
4. Create a label for the file system. For more information, see “[Before mounting the file system](#)” on page 35.
5. Mount the file system. For more information, see “[Mounting the file system](#)” on page 35.

Identifying and viewing information about LUNs

Persistently identifying each LUN

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

The Linux host sees LUNs on the storage system as SCSI devices. Linux assigns an identifier, such as /dev/sda or /dev/sdb, to each LUN. The identifiers are assigned as the LUNs are discovered. Because of network latencies, the order in which LUNs are discovered cannot be predicted. Therefore, a LUN may be mapped to a different device name after a Linux reboot or a reload of the iSCSI service.

If you have more than one LUN, you need to have a persistent way of identifying each LUN. To identify the persistence of a device you can use the `scsi_id` command. The unique string in the output will be common across the multiple SCSI devices belonging to the same LUN. See the following example:

```
scsi_id -gus /block/sdb
[root@linux-119 ~]# scsi_id -gus /block/sdb
360a98000687046754c3030364b736c74

scsi_id -gus /block/sdc
[root@linux-119 ~]# scsi_id -gus /block/sdc
360a98000687046754c3030364b736c74
```

The output for `scsi_id -gus /block/sdb` and `scsi_id -gus /block/sdc` command are the same, which shows that the iSCSI devices are two paths to the same LUN.

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

- ◆ `iscsi-ls -l` - The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.
- ◆ `lsscsi` - The `lsscsi` command lists all SCSI devices including iSCSI devices and their attributes.
- ◆ `sanlun` - The `sanlun` command is included in Host Utilities. It displays the host device names and the storage system LUNs to which they are mapped.
- ◆ `ls -l` - The `ls -l` command lists the symbolic links the initiator created under the /dev/disk/by-id directory. For more information on symbolic links, see [“Viewing symbolic links”](#).

Viewing storage system and LUN details

Using the `iscsi-ls` command: To view storage system and LUN details, enter the following command on the host:

```
iscsi-ls -l
```

Result: The console lists available storage systems and LUNs.

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

```
*****
                Cisco iSCSI Driver Version ... 4.0.188.26 ( 21-May-2004 )
*****
TARGET NAME           : iqn.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
*****
```

Viewing the mapping of host devices to LUNs

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 /opt/netapp/santools directory.
2	Enter the following command: sanlun lun show all Result: The sanlun command displays information about 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 SP3.

```
[root@host_1 /]# /opt/netapp/santools/sanlun lun show all
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
```

Viewing information about target configuration

Using iscsi-ls command: To view information about target configuration, enter the following command on the host:

```
iscsi-ls -c
```

Result: The console lists available storage systems and LUNs, and iSCSI session parameters.

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

```
*****
          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
```

Viewing information about SCSI devices

Using the `lsscsi` command: To view information about all the SCSI devices and their attributes, enter the following command at the host console:

```
lsscsi
```

Result: Lists information about all the SCSI devices and their attributes.

Example: The following example shows the output of `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
```

Viewing symbolic links

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

For SUSE Linux Enterprise Server 9 SP3, the iSCSI devices are under the `/dev/disk/by-id/` directory in the following format:

```
iscsi-<target_name>-LUN_ID > symbolic_link_to_disk
```

Example: The following example shows three iSCSI devices in the `/dev/disk/by-id/` directory. For the second device listed, “-0p1” appended to the node name indicates the disk partition.

```
lrwxrwxrwx  1 root root 9 Aug 23 18:48 iscsi-iqn.1992-08.com.netapp:sn.12345678-0 -> ../../sda
lrwxrwxrwx  1 root root 10 Aug 23 18:48 iscsi-iqn.1992-08.com.netapp:sn.12345678-0p1 -> ../../sda1
lrwxrwxrwx  1 root root 9 Aug 23 18:48 iscsi-iqn.1992-08.com.netapp:sn.12345678-1 -> ../../sdb
```

Using the ls -l command: To view symbolic links and the storage system node name to which they are mapped, enter the following command at the host console:

```
# ls -l /dev/disk/by-path
```

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

```
# ls -l /dev/disk/by-path
total 0
drwxr-xr-x 2 root root 304 2004-08-24 14:57 .
drwxr-xr-x 4 root root 96 2004-08-18 14:26 ..
lrwxrwxrwx 1 root root 9 2004-08-24 13:32 iscsi-iqn.1992-
08.com.netapp:sn.12345678.ss1-0 -> ../../sde
lrwxrwxrwx 1 root root 10 2004-08-24 13:32 iscsi-iqn.1992-
08.com.netapp:sn.12345678.ss1-0p1 -> ../../sde1
lrwxrwxrwx 1 root root 9 2004-08-24 14:57 iscsi-iqn.1992-
08.com.netapp:sn.12345678.ss1-1 -> ../../sdg
```

Creating a partition

About partitions

You can divide a LUN into one or more partitions. Each partition can contain a file system. If you want to use the entire LUN for a single file system, You must create a partition.

Before you begin

You should have one or more LUNs available to the host. Log in to the host as root and verify that the LUNs are available.

Using fdisk to create a partition

To create a partition of a LUN perform the following steps.

Step	Action
1	On the host console, enter the following command: <code>/sbin/fdisk <i>iscsi_device</i></code> <i>iscsi_device</i> is the iSCSI LUN.
2	Enter n to create a new partition.
3	Enter p to create a primary partition.
4	Enter a partition number. For the first partition, enter 1 .
5	Press Enter twice to accept the default cylinder values.
6	Enter w to save the changes and exit fdisk.

Creating a file system on a partition of a LUN

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

On the host console, enter the following command:

```
mkfs -V -t filesystem_type partition
```

partition is the name of the partition you just created.

filesystem_type is the type of file system you want to create.

Example: The following command creates an ext3 file system on the first partition of LUN 0.

```
mkfs -V -t ext3 /dev/disk/by-id/iscsi-ign.1992-  
08.com.netapp:sn.12345678.ss_hq-0p1
```

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. See “[Tasks for accessing LUNs](#)” on page 26 for more information.

When you create file systems’ labels, you must add them to the `/etc/fstab.iscsi` file so that the file systems will be automatically mounted.

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) used to identify the file system.

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

For Reiser file systems, you create the label using the `--label` option when you create the file system. You do not use the `e2label` command.

Automatically mounting the file system

The Linux system mounts all of 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 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>

Accessing LUNs with dm-multipath support

Overview

Note

Support for dm-multipath is on SUSE Linux Enterprise Server 9 SP3.

To access LUNs with dm-multipath support, the following is an overview of the tasks you need to perform:

1. Verify the dm-multipath configuration. For more information, see “[Verifying the dm-multipath configuration](#)” on page 38.
2. Identify the LUN on which the partition is to be created. For more information, see “[Identifying and viewing a list of LUNs](#)” on page 39.
3. Create and mount the file system. For more information, see “[Creating and mounting a file system on a LUN](#)” on page 43.
4. Tune the dm-multipath configuration. For more information, see “[Tuning the dm-multipath configuration](#)” on page 45.

Verifying the dm-multipath configuration

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

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

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

Identifying and viewing a list of LUNs

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

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

The `sanlun` command is installed by Host Utilities in the `/opt/netapp/santools` directory. To view a list of LUNs that are mapped to the Linux host, enter the following command on the Linux host console:

```
sanlun lun show all
```

Example: The following output shows one LUN mapped to the Linux host from each of two storage systems. There are two paths from the host to each storage system, which is why each LUN appears twice.

```

lnx200-165:~ # 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

```

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

Using the `iscsi-ls` command

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

Example: The following example shows the part of the `iscsi-ls -l` command output for the storage system ss1 LUN `/vol/vol1/linux_1`. This path maps to the `/dev/sdc` device on the host.

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 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 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: After 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 -ll`

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

Example: The following example output shows one new LUN with two paths:

```
360a9800043346461524a43456e623179
[size=1 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [prio=8][active]
  \_ 4:0:0:1 sdc 8:32 [active][ready]
  \_ 4:1:0:1 sde 8:64 [active][ready]
  \_ 4:2:0:1 sdg 8:96 [active][ready]
  \_ 4:3:0:1 sdi 8:128 [active][ready]

360a9800043346461524a427837582f42
[size=1 GB][features="1 queue_if_no_path"][hwhandler="0"]
\_ round-robin 0 [prio=8][active]
  \_ 4:0:0:0 sdb 8:16 [active][ready]
  \_ 4:1:0:0 sdd 8:48 [active][ready]
  \_ 4:2:0:0 sdf 8:80 [active][ready]
  \_ 4:3:0:0 sdh 8:112 [active][ready]
```

Accessing LUNs with dm-multipath support

Creating and mounting a file system on a LUN

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

Tuning the dm-multipath configuration

Tuning the dm-multipath parameters

Advanced administrators who are familiar with Linux performance tuning might want to modify one of the dm-multipath parameters for best performance.

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

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

```
defaults {  
    rr_min_io      128  
}
```

Stopping dm-multipath

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

To stop dm-multipath and the iSCSI service, complete the following steps.

Step	Action
1	Unmount all iSCSI devices.
2	Flush the multipath maps by entering the following command: <code>multipath -F</code>
3	Stop the dm-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>

Configuration overview

This chapter explains how to install Host Utilities, and how to complete the installation and initial configuration of the iSCSI initiator and multipathing software on the Linux host.

Note

You will complete the final configuration of the initiator after you have configured the storage system.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the Host Utilities software”](#) on page 48
- ◆ [“Verifying or installing the required packages”](#) on page 50
- ◆ [“Configuring dm-multipath support”](#) on page 51
- ◆ [“Recording or changing the initiator node name on the host”](#) on page 55
- ◆ [“Configuring the iSCSI initiator”](#) on page 58

Requirement for root privileges

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

Installing the Host Utilities software

Verifying the correct version

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

Getting the Host Utilities software

Download the Host Utilities file (netapp_linux_host_utilities_3_2.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.

Installing the Host Utilities software

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

Step	Action
1	Remove any previous version of Host Utilities installed on your system. Change to the directory where the previous version is installed (default is /opt/netapp/santools) and enter the following command: ./uninstall
2	Change to the working directory to which you saved the Host Utilities file.
3	To uncompress the file, enter the following command: gunzip netapp_linux_host_utilities_3_2.tar.gz
4	To extract the files, enter the following command: tar -xvf netapp_linux_host_utilities_3_2.tar
5	Change to the netapp_linux_host_utilities_3_2 directory. By default, this directory is a subdirectory of the working directory in which you extracted the Host Utilities files in the previous step.

Step	Action
6	Enter the following command: ./install

The diagnostic scripts are installed on the /opt/netapp/santools directory. Note that this directory is different from the directory used by the previous version of Host Utilities.

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

Verifying or installing the required packages

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

Step	Action	
1	For SUSE Linux Enterprise Server 10 SP1 and SP2, enter the following command: <pre>rpm -q open-iscsi</pre> 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 support, enter the following commands: <pre>rpm -q device-mapper</pre> <pre>rpm -q multipath-tools</pre> Result: The <code>rpm -q</code> commands return the name and version of the installed packages. Install the correct version of the packages that are shipped along with the operating system.	
3	If...	Then...
	You have the correct packages for your version of SUSE Linux Enterprise Server.	Proceed to “Configuring dm-multipath support” on page 51.
	You do not have the correct packages for your version of SUSE Linux Enterprise Server.	Install the required package versions from the media for your version of SUSE Linux Enterprise Server. You can use the <code>rpm</code> or <code>yast</code> command.

Configuring dm-multipath support

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

Copying the file The multipath.conf file contains configuration settings for multipathing. A sample 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 multipath.conf.

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

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

Step	Action
1	Open /etc/multipath.conf with a text editor.
2	Add the following syntax for blacklisting devices: blacklist {}

Step	Action
3	Blacklist all devices that do not correspond to any NetApp LUN and are mapped to your Linux host (that is, blacklist those devices that are not displayed in the <code>sanlun lun show</code> command). - 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 that device. Example: assume that <code>/dev/sda</code> is a local SCSI drive. To obtain its WWID, you will run the following command: <pre>scsi_id -gus /block/sda</pre> The output be as follows: <pre>SIBM-ESXSMW3073NC_FDAR9P66067WJ</pre> Add this entry for blacklisting. - Blacklist other devices using the devnode method by adding the following: <pre>devnode "^hd[a-z]" devnode "^(ram raw loop fd md dm- sr scd st)[0-9]*" devnode "^cciss!c[0-9]d[0-9]*"</pre> Example: The following is an example for a SUSE Linux Enterprise Server 10 SP1 and SP2 host: <pre>blacklist { devnode "^(ram raw loop fd md dm-ls scd st)[0-9]*" devnode "^hd[a-z][0-9]*" devnode "^cciss!c[0-9]d[0-9]*p[0-9]*" wwid SIBM-ESXSMW3073NC_FDAR9P66067WJ }</pre>
4	Save the changes.

Example: The following example shows the required changes to the `/etc/multipath.conf` file for a system with one local SCSI drive for SUSE Linux Enterprise Server 10 SP1.

```

##
## name      : blacklist
## scope     : multipath & multipathd
## desc      : list of device names to discard as not multipath
candidates
## default   : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
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
}
#
##
devices {
    device {
        vendor            "NETAPP"
        product           "LUN"
        path_grouping_policy  multibus
        getuid_callout      "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout        "/sbin/mpath_prio_ontap /dev/%n"
        features            "1 queue_if_no_path"
        path_checker        readsector0
        failback            immediate
    }
}

```

Example: The following example shows the required changes to the `/etc/multipath.conf` file for a system with one local SCSI drive for SUSE Linux Enterprise Server 10 SP2.

```

##
## name      : blacklist
## scope     : multipath & multipathd
## desc      : list of device names to discard as not multipath
candidates
## default   : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
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
}
#
##
devices {

```

```

device {
    vendor                "NETAPP"
    product               "LUN"
    path_grouping_policy  multibus
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    prio                  "ontap"
    features              "1 queue_if_no_path"
    path_checker           directio
    failback              immediate
}
}

```

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 uses the following format on SUSE Linux Enterprise Server 10 SP1 and SP2:

`iqn.1996-04.de.suse:RandomNumber`

You must set the initiator node name to `iqn.1996-04.de.suse:host.name` on SUSE Linux Enterprise Server 10 SP1 and SP2 hosts, where *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 default 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/iscsi/initiatorname.iscsi</code> file.

Step	Action	
2	If...	Then...
	A default initiator name has not been generated	a. Replace the text in the host's <code>/etc/iscsi/initiatorname.iscsi</code> file with the initiator name in the following format: <pre>InitiatorName=iqn.1996-04.de.suse:RandomNumber</pre> The recommended value for <i>RandomNumber</i> is the host name of the Linux host. b. Record the initiator node name and proceed to “Configuring the iSCSI initiator” on page 58.
	You want to use the default initiator node name	Record the initiator node name and proceed to “ Configuring the iSCSI initiator ” on page 58.
	You want to change the initiator node name	a. Modify the <i>RandomNumber</i> part of the initiator node name. The following is a sample node name: <pre>iqn.1996-04.de.suse:lnx.118</pre> b. Record the initiator node name and proceed to “Configuring the iSCSI initiator” on page 58. Note _____ <i>RandomNumber</i> is the only component of the initiator node name you modify. You do not modify the other components of the initiator node name. You must use the host name in place of the default random number. _____

Node name rules

If you change the host's initiator node name, be sure the new name follows all of the following rules:

- ◆ An initiator node name can be up to 223 bytes.
- ◆ Upper-case characters are always mapped to lower-case characters.
- ◆ An initiator node name can contain alphabetic characters (a to z), numbers (0 to 9), and the following three special characters:
 - ❖ Period (.)
 - ❖ Hyphen (-)
 - ❖ Colon (:)
- ◆ The underscore character (_) is *not* supported.

Example: `iqn.1996-04.de.suse:lnx.118`

Configuring the iSCSI initiator

SUSE Linux includes the iSCSI initiator software in the software distribution. The initiator's components include a kernel module that is already compiled into the Linux kernel, and the user space packages.

The configuration of the iSCSI initiator on SUSE Linux Enterprise Server 10 SP1 and SP2 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. When discovery is run for the first time, 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 are `/etc/iscsi`.

iSCSI parameter settings

The default value for `replacement_timeout` is 120 seconds. `replacement_timeout` can be set to a lower value like 5 to 15 seconds, because the `multipath queue_if_no_path` and `no_path_retry` settings can prevent IO errors from being propagated to the application. If a lower value is set, then the pending IO is quickly sent to a new path and is executed, while the iSCSI layer attempts reestablishment of the session. If all paths fail, then the multipath and the device mapper layer internally queue the IO based on the `multipath.conf` settings, instead of the iSCSI layer.

Example: `node.session.timeo.replacement_timeout = 5`

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

For more information about using YaST2, see the SUSE Linux Enterprise Server 10 SP1 and SP2 documentation available at http://www.novell.com/documentation/sles10/sles_admin/data/sles_admin.html.

Note

NetApp recommends you to configure the iSCSI initiator using YaST2.

Configuring the iSCSI initiator using YaST2

The configuration is divided into three tabs. The Service tab can be used to enable the iSCSI initiator at boot time. The Connected Targets tab gives an overview of the currently connected iSCSI targets. Similar to the Discovered Targets tab, it provides the option to add new targets to the system. The Discovered Targets tab provides the possibility of discovering iSCSI targets in the network.

You must do the following tasks to configure the iSCSI initiator using YaST2.

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 startup option to Manual or Automatic, depending on your requirement, 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 used to manage (update, delete, insert, and query) the persistent database. The utility presents a set of operations that you can perform on iSCSI nodes, sessions, connections, and discovery records.

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

Target Discovery: Complete the following steps.

Step	Action
1	Enter the following command to perform discovery: <pre>iscsiadm --mode discovery --op update --type sendtargets --portal targetip</pre> <i>targetip</i> is the IP address of the target. Result: The target targetip 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 <code>iscsid.conf</code> file. For successful session discovery, you need to enable the discovery CHAP authentication configuration entries. <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>

Step	Action
4	Enter the following command to display all the active iscsi sessions: <pre>iscsiadm --mode session</pre> Result: All established iSCSI sessions are displayed. Example: The following is 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.
5	To automate login to a node, use the following command with the record ID of the node discovered: <pre>iscsiadm -m node -r recordId -o update -n node.conn[0].startup -v automatic</pre> Alternatively, to set the node.startup attribute to startup as default for all sessions add the following to /etc/iscsi/iscsid.conf and perform rediscovery: <pre>node.startup = automatic</pre> Note _____ The internal database does not exist during the first discovery, therefore it cannot be used at this time.

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

Step	Action
1	With a text editor, open the host's /etc/iscsi/iscsid.conf file.
2	To enable CHAP authentication set node.session.auth.authmethod to CHAP. The default is None. <pre>node.session.auth.authmethod = CHAP</pre>

Step	Action
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

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

Login mechanism:

There are two modes of login for iSCSI:

- ◆ Manual login
- ◆ Automatic login

Manual login:

If...	Then...
You want to set the login mode to manual on a single portal on a target	Run the following command: <pre>iscsiadm -m node -T targetname -p ip:port --op update -n node.startup -v manual</pre>
You want to set the login mode to manual on all portals on a target	Run the following command: <pre>iscsiadm -m node -T targetname --op update -n node.startup -v manual</pre>

If...	Then...
You want to set the login mode to 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 in iscsid.conf does not affect existing nodes. It will only affects nodes that are discovered after setting the value. To login to all the manual nodes, simply restart the iSCSI service:

Example: /etc/init.d/open-iscsi restart

During the next startup you must manually log in to the nodes.

Automatic login:

If...	Then...
You want to set the login mode to 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
You want to set the login mode to automatic on all portals on a target	Run the following command: iscsiadm -m node -T targetname --op update -n node.startup -v automatic
You want to set the login mode to automatic mode for all the targets	Add the following to the /etc/iscsi/iscsid.conf: node.startup = automatic

Setting the automatic login mode in iscsid.conf does not affect existing nodes. It will only affects nodes that are discovered after setting the value. To login to all the manual nodes, simply restart the iSCSI service:

Example: /etc/init.d/open-iscsi restart

During the next startup the nodes are logged into automatically.

Before you begin

You need the host’s iSCSI initiator node name that you recorded in [Chapter 5](#), before you can configure the storage system. You use the initiator node name to create igroups on the storage system.

Tasks for configuring the storage system

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

Step	Action
1	Ensure that the iSCSI service is licensed and started.
2	If you want to use CHAP authentication, use the <code>iscsi security</code> command or the FilerView® interface to configure a CHAP user name and password.
3	Create LUNs, create an igroup using the host’s initiator node name, and then map the LUNs to the igroup for the Linux host. You can manage LUNs using the Web-based FilerView interface or using the Data ONTAP command line. Be sure to specify the LUN type and iSCSI igroup type as <code>linux</code> .

For detailed configuration steps

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

Overview of accessing LUNs

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

The steps for accessing LUNs depend on which version of SUSE Linux Enterprise Server you are using on your Linux host and whether you are using dm-multipath support. Be sure to follow the procedures for your version.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Starting the iSCSI service](#)” on page 68
- ◆ “[Starting the multipath service](#)” on page 69
- ◆ “[Accessing LUNs without dm-multipath support](#)” on page 70
- ◆ “[Accessing LUNs with dm-multipath support](#)” on page 82
- ◆ “[Stopping dm-multipath](#)” on page 91

Starting the iSCSI service

To start the iSCSI service on SUSE Linux Enterprise Server 10 SP1 and SP2 hosts, enter the following command at the Linux host command prompt:

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

Configuring the iSCSI service to start automatically after reboot: If you want to configure the iSCSI service to start automatically at system bootup, complete the following steps.

Step	Action	
1	Verify the status of the iSCSI service by entering the following command: chkconfig open-iscsi	
2	If...	Then...
	The <code>chkconfig</code> command indicates the iSCSI service is enabled	No action is required.
	The <code>chkconfig</code> command indicates that the iSCSI service is not enabled	Start the iSCSI service by entering the following command: chkconfig open-iscsi on

Also, the Service tab in YaST2 can be used to enable the iSCSI initiator at boot time.

Starting the multipath service

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

After you have started the iSCSI service, start the multipath service on the Linux host. Complete the following steps.

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

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

To integrate the multipath setup into the boot sequence, enter the following commands on the Linux host console:

```
chkconfig boot.multipath on
chkconfig multipathd on
```

Accessing LUNs without dm-multipath support

Tasks for accessing LUNs

After the initiator is configured, you need to set up the host to access the LUNs you created on the storage system. The following is an overview of the tasks you must complete to access LUNs.

1. Identify the LUN on which the partition is to be created. For more information, see “[Identifying and viewing information about LUNs](#)” on page 71.
2. Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN. For more information, see “[Creating a partition](#)” on page 78.
3. Create a file system on the partition. For more information, see “[Creating a file system on a partition of a LUN](#)” on page 79.
4. Create a label for the file system. For more information, see “[Before mounting the file system](#)” on page 80.
5. Mount the file system. For more information, see “[Mounting the file system](#)” on page 80.

Identifying and viewing information about LUNs

Persistently identifying each LUN

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

The Linux host sees LUNs on the storage system as SCSI devices. Linux assigns an identifier, such as `/dev/sda` or `/dev/sdb`, to each LUN. The identifiers are assigned as the LUNs are discovered. Because of network latencies, the order in which LUNs are discovered cannot be predicted. Therefore, a LUN may be mapped to a different device name after a Linux reboot or a restart of the iSCSI service.

If you have more than one LUN, you need to have a persistent way of identifying each LUN. To identify the persistence of a device you can use the `scsi_id` command. The unique string in the output will be common across the multiple SCSI devices belonging to the same LUN. See the following example:

```
scsi_id -gus /block/sdb
[root@linux-119 ~]# scsi_id -gus /block/sdb
360a98000687046754c3030364b736c74

scsi_id -gus /block/sdc
[root@linux-119 ~]# scsi_id -gus /block/sdc
360a98000687046754c3030364b736c74
```

The output for `scsi_id -gus /block/sdb` and `scsi_id -gus /block/sdc` command are the same, which shows that the iSCSI devices are two paths to the same LUN.

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

- ◆ **sanlun:** The `sanlun` command is included in Host Utilities. It 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 the 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 the information about the available storage systems and LUNs for a session with a specific session ID, *N*.
- ◆ **lsscsi:** The `lsscsi` command lists all SCSI devices including iSCSI devices and their attributes.

- ◆ `ls -l`: The `ls -l` command lists the symbolic links the initiator created under the `/dev/disk/by-id` directory. For more information about symbolic links, see “[Viewing symbolic links](#)”.

Viewing the mapping of host devices to LUNs

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 /opt/netapp/santools directory.
2	Enter the following command: sanlun lun show all Result: The sanlun command displays information about iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The following shows a sample output of the sanlun command. <pre>[root@host_1 /]# /opt/netapp/santools/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</pre>

Viewing information about sessions and LUNs

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.

Enter the following command at the host console:

```
iscsiadm -m session --sid=N -P 3
```

N is the specific session ID.

Result: The available storage systems and LUNs for a session with a specific session ID, *N* are listed.

Example: The following example shows the output on SUSE Linux Enterprise Server 10 SP1 and SP2.

```
lnx200-109:~ # iscsiadm -m session
tcp: [1] 10.72.199.144:3260,1001 ign.1992-08.com.netapp:sn.101180456
tcp: [2] 10.72.199.146:3260,1002 ign.1992-08.com.netapp:sn.101180456
tcp: [3] 10.72.199.142:3260,1000 ign.1992-08.com.netapp:sn.101180456
tcp: [4] 10.72.199.148:3260,1003 ign.1992-08.com.netapp:sn.101180456
```

```
lnx200-109:~ # 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 active iSCSI sessions, use the `iscsiadm -m session` command.

Viewing information about SCSI devices

Using the `lsscsi` command: To view information about all the SCSI devices and their attributes, enter the following command at the host console:

```
lsscsi
```

Result: Lists information about all the SCSI devices and their attributes.

Example: The following example shows the output of `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 symbolic links

The initiator uses symbolic links to provide host applications with reliable mappings between device nodes and LUNs. Linux assigns SCSI device nodes dynamically when iSCSI targets or LUNs are detected. This means that the mapping of SCSI devices to LUNs cannot be persistent each time the host iSCSI service is started; SCSI commands can be sent to the wrong storage system or LUN. If you have applications that access storage system LUNs, you must ensure that these applications use the symbolic links instead of the standard SCSI device nodes.

For SUSE Linux Enterprise Server 10 SP1 and SP2, the iSCSI devices are under the `/dev/disk/by-id/` directory in the following format:

```
scsi-diskID > symbolic_link_to_disk
```

Example: The following example shows three iSCSI devices in the `/dev/disk/by-id/` directory. For the first device listed, “-part1” appended to the node name indicates the disk partition.

```
lrwxrwxrwx 1 root root 9 Dec 27 21:22 scsi-
360a980004334654c6d4a39436c686a32 -> ../../sdk

lrwxrwxrwx 1 root root 10 Dec 28 12:45 scsi-
360a980004334654c6d4a39436c686a32-part1 -> ../../sdk1

lrwxrwxrwx 1 root root 9 Dec 27 21:22 scsi-
360a980004334654c6d4a39436c69394d -> ../../sd1

lrwxrwxrwx 1 root root 9 Dec 27 21:22 scsi-
360a980004334654c6d4a39436c695553 -> ../../sdm
```

Using the `ls -l` command: To view symbolic links and the storage system node name to which they are mapped, enter the following command at the host console:

```
# ls -l /dev/disk/by-path
```

Example: The following is sample output of the `ls -l /dev/disk/by-path` command.

```
lnx200-109:~ # ls -l /dev/disk/by-path/
total 0

lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.142:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-0 -> ../../sde
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.142:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-1 -> ../../sdh
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.142:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-2 -> ../../sdl
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.144:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-0 -> ../../sdb
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.144:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-1 -> ../../sdf
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.144:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-2 -> ../../sdj
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.146:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-0 -> ../../sdd
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.146:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-1 -> ../../sdg
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.146:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-2 -> ../../sdk
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.148:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-0 -> ../../sdc
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.148:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-1 -> ../../sdi
lrwxrwxrwx 1 root root  9 Apr 22 15:57 ip-10.72.199.148:3260-iscsi-iqn.1992-08.com.netapp:sn.101180456-lun-2 -> ../../sdm
```

Creating a partition

About partitions

You can divide a LUN into one or more partitions. Each partition can contain a file system. If you want to use the entire LUN for a single file system, you must create a partition.

Before you begin

You should have one or more LUNs available to the host. Log in to the host as root and verify that the LUNs are available.

Using fdisk to create a partition

To create a partition of a LUN perform the following steps.

Step	Action
1	On the host console, enter the following command: <code>/sbin/fdisk <i>iscsi_device</i></code> <i>iscsi_device</i> is the iSCSI LUN.
2	Enter n to create a new partition.
3	Enter p to create a primary partition.
4	Enter a partition number. For the first partition, enter 1 .
5	Press Enter twice to accept the default cylinder values.
6	Enter w to save the changes and exit fdisk.

Creating a file system on a partition of a LUN

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

On the host console, enter the following command:

```
mkfs -V -t filesystem_type partition
```

partition is the name of the partition you just created.

filesystem_type is the type of file system you want to create.

Example: The following command creates a reiserfs file system on the first partition of the specified iSCSI disk.

```
mkfs -V -t reiserfs /dev/disk/by-id/scsi-  
360a980004334654c6d4a39436c686a32-part1
```

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. See “[Tasks for accessing LUNs](#)” on page 70 for more information.

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:

Enter `/sbin/e2label device label`

device is the partition where you created the file system.

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

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

For Reiser file systems, you create the label using the `--label` option when you create the file system. You do not use the `e2label` command.

Note

You should mount devices using *by-id* and not using labels in SLES 10 SP1 and SP2.

Automatically mounting the file system

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

Attention — Before mounting the file system, modify the `/etc/sysconfig/network/ifcfg-eth-id-xxxxxxx` and change the value of `STARTMODE` to `nfsroot`. This is required for the successful unmounting of the iscsi devices during system halts.

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.
3	Add an entry for the device on which you have created the file system. Example: <pre>/dev/disk/by-id/scsi-360a98044334685f433e6774722-part1 /mnt/data ext3 hotplug,rw 0 0</pre> <code>/dev/disk/by-id/scsi-360a98044334685f433e6774722-part1</code> is the iSCSI device and it needs to be listed by its ID. <code>/mnt/data</code> is the mountpoint. <code>ext3</code> is the file system. <code>hotplug</code> is the option which instructs the kernel that the device is a hot-pluggable device. <code>hotplug</code> option is required to automatically mount the iSCSI device.
4	Save your changes, and exit the <code>/etc/fstab</code> file.
5	On the host console, enter the following command to mount the file systems you added: <pre>/bin/mount -a</pre>

Accessing LUNs with dm-multipath support

Overview

To access LUNs with dm-multipath support, the following is an overview of the tasks you need to perform:

1. Verify the dm-multipath configuration. For more information, see “[Verifying the dm-multipath configuration](#)” on page 83.
2. Identify the LUN on which the partition is to be created. For more information, see “[Identifying and viewing a list of LUNs](#)” on page 84.
3. Create and mount the file system. For more information, see “[Creating and mounting a file system on a LUN](#)” on page 88.
4. Tune the dm-multipath configuration. For more information, see “[Tuning the dm-multipath configuration](#)” on page 90.

Verifying the dm-multipath configuration

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

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

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

iSCSI timeout values behaviour

The default value for replacement_timeout is 120 seconds. replacement_timeout can be set to a lower value like 5 to 15 seconds, because the multipath queue_if_no_path and no_path_retry settings can prevent IO errors from being propagated to the application. If a lower value is set, then the pending IO is quickly sent to a new path and is executed, while the iSCSI layer attempts reestablishment of the session. If all paths fail, then the multipath and the device mapper layer internally queue the IO based on the multipath.conf settings, instead of the iSCSI layer.

Identifying and viewing a list of LUNs

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

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

Note

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

Using the `sanlun` command to view LUNs

The `sanlun` command is installed by Host Utilities in the `/opt/netapp/santools` directory. To view a list of LUNs that are mapped to the Linux host, enter the following command on the Linux host console:

```
sanlun lun show all
```

Example: The following output shows one LUN mapped to the Linux host from each of two storage systems. There are two paths from the host to each storage system, which is why each LUN appears twice.

```

host:/ # sanlun lun show all
filer:      lun-pathname  device filename  adapter  protocol      lun size
lun state
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

```

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

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 following example lists the available storage systems and LUNs for a session with the ID 0.

```

lnx200-109:~ # iscsiadm -m session
tcp: [1] 10.72.199.144:3260,1001 iqn.1992-08.com.netapp:sn.101180456
tcp: [2] 10.72.199.146:3260,1002 iqn.1992-08.com.netapp:sn.101180456
tcp: [3] 10.72.199.142:3260,1000 iqn.1992-08.com.netapp:sn.101180456
tcp: [4] 10.72.199.148:3260,1003 iqn.1992-08.com.netapp:sn.101180456

lnx200-109:~ # iscsiadm -m session --sid=1 -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

```

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 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
```

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

Accessing LUNs with dm-multipath support

Creating and mounting a file system on a LUN

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 SP1 and SP2, complete the following steps.

Attention

Before mounting the file system, modify the /etc/sysconfig/network/ifcfg-eth-id-xxxxxxx and change the value of STARTMODE to *nfsroot* . This is required for the successful unmounting of the iscsi devices during system halts.

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

Step	Action
3	Add an entry for the file system you created in “Creating a file system on a partition of a LUN” on page 32. Example: <pre>/dev/mapper/360a98044334685f433e6774722-part1 /mnt/data ext3 hotplug,rw 0 0</pre> /dev/mapper/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. hotplug is the option which instructs the kernel that the device is a hot-pluggable device. hotplug option is required to automatically mount the iSCSI device.
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>

Tuning the dm-multipath configuration

Tuning the dm-multipath parameters

Advanced administrators who are familiar with Linux performance tuning might want to modify one of the dm-multipath parameters for best performance.

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

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

```
defaults {  
    rr_min_io      128  
}
```

Stopping dm-multipath

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

To stop dm-multipath and the iSCSI service, complete the following steps.

Step	Action
1	Unmount all iSCSI devices.
2	Flush the multipath maps by entering the following command: <code>multipath -F</code>
3	Stop the dm-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/open-iscsi stop</code>

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

Topics in this chapter This chapter has the following topics:

- ◆ [“SAN boot overview”](#)
- ◆ [“Setting up the host for SAN boot”](#)
- ◆ [“Configuring root partition on multipath”](#)

Note

SAN boot is currently supported on SUSE Linux Enterprise Server 10 SP1 and SP2 only.

SAN boot overview

About SAN boot

SAN boot can be implemented either by an iSCSI HBA, or a network interface card (NIC) and software iSCSI stack.

If you use the iSCSI HBA, since the protocol stack runs on the HBA, it is ready to communicate with the storage system and discover a LUN.

For a software initiator to implement SAN boot, any of the following options can be used to load the kernel:

- ◆ Using a host's locally attached disk (for storing kernel and initrd images)
- ◆ Using a preboot execution environment (PXE) server

You should use a locally attached disk for SAN boot to store the kernel and initrd images.

Note

Setting up a PXE server is beyond the scope of this guide.

Advantages of SAN boot

SAN boot uses a SAN-attached disk, such as a NetApp LUN, as a root(/) device for a host. SAN boot provides the following advantages:

- ◆ You can remove the hard drives from your servers and use the SAN for your booting needs, eliminating the cost associated with maintaining and servicing hard drives
- ◆ The host uses the SAN, consolidating and centralizing storage
- ◆ Lower cost: The hardware and operating costs are lowered.
- ◆ Greater reliability: Systems without disks are less prone to failure.
- ◆ Quick server swaps: In the event of a server failure, systems without disks can be swapped.
- ◆ Better disaster recovery: Site duplication is simplified.

Setting up the host for SAN boot

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

You are required to 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”](#).
2. Selecting the partition. For more information, see [“Partition selection”](#).
3. Setting up the bootloader. For more information, see [“Setting up the Bootloader”](#).
4. Configuring root partition on multipath. For more information, see [“Configuring root partition on multipath”](#).

Root partition on multipath prerequisites

Following are the prerequisites for root partition on multipath on SUSE Linux Enterprise Server 10 SP1 and SP2:

- ◆ You must implement the OS boot (kernel) from a local disk or PXE server. For more information see [“Creating /boot partition”](#).
- ◆ You must ensure that you have configured your storage system as per iSCSI requirements.

Initial setup tasks

Run the installation of SUSE Linux Enterprise Server from a CD-ROM. This section describes the steps you need to follow after the host boots from the CD-ROM and the installer runs.

Step	Action
1	In the initial installation page, select Installation.
2	Specify Boot Option as follows: <code>withiscsi=1 netsetup=1</code>

3	Press Enter. You can view the list of the network devices available.
4	Select your network device and configure network on it.
5	Select the Language of your choice, and click Next.
6	The Licence Agreement appears. Click Next. The iSCSI Initiator Overview page appears.
7	Select the Service tab and set the iqn number.
8	Log in to the Web interface of the NetApp storage system, and click FilerView.
9	Click LUNs.
10	Click Initiator group.
11	Click Manage. The Manage Initiator Groups page is displayed. Click the relevant group.
12	In the initiator text box, add the iqn number to the initiator groups of the NetApp storage system. Note _____ If you have more than one storage system, add the iqn to that storage system also. _____ Click Apply. Ensure that the operation succeeds.
13	Map a LUN to the igroup to which you added the iqn number of the initiator.
14	Return to the host screen.
15	Select the Connected Targets tab. Click Add.

16	Do the following: (On iSCSI Initiator Discovery page) • Specify the IP address of the NetApp storage system. • Specify the port. Default is 3260. • Enter the credentials if you are using an authentication mechanism. • Click Next. You can view the list of NetApp storage system which are successfully discovered.
17	Click Connect for each discovered NetApp storage system. Specify the authentication credentials if you have an authentication mechanism. Click Next. Note _____ During installation, only one path to the root LUN must be enabled.
18	Verify that the value for Connected is <i>true</i> for all the targets. Click Next. The Connected Targets pane will list all the targets.
19	Set the Start up mode to Automatic using the Toggle Start-up button.
20	Click Finish. Note _____ After this step, 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 procedures.

Partition selection

To select the partition, perform 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 NetApp 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 ————— Make sure that the NetApp 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: / Click Fstab Options button.
10	Select the Mount by Device ID radio button.
11	Make sure that you have the <i>hotplug</i> keyword in the Arbitrary Option Value text box. Click OK.
12	The Create a Primary Partition page. Click OK.

Creating a SWAP Partition: You may 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 8](#) in the section “[Partition selection](#)”, and specify the File system value = SWAP.

Note

If you are using the software suspend functionality, make sure that the SWAP partition is on a local disk.

Creating /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 8](#) in the section “[Partition selection](#)”, and specify the Mount Point in [Step 9](#) as /boot.

Note

You must install the /boot partition on a locally attached disk.

After you return to the Expert Partitioner page, review the configuration. Then, 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 Booting.
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> .
5	Click OK.
6	Click the Boot Loader Installation tab.
7	In the Boot Loader Location pane, select the Boot from Master Boot Record option.
8	Click Finish. You return to the Installation Settings page.
9	Review the configuration settings. Click Accept. The Confirm Installation page is displayed.

10	Click Install.
-----------	----------------

Configuring root partition on multipath

After SAN boot is configured on the host, root partition on multipath can be configured. To configure SAN boot using root partition on multipath, perform the following steps.

Step	Action
1	Configure the multipath. For details on configuring the dm-multipath, see “ Configuring dm-multipath support ” on page 51.
2	Start the <code>multipathd</code> daemon by running the following command: <code>service multipathd start</code>
3	Enable the additional paths to the remote NetApp LUN (root LUN) and set it to automatic (use YaST2). Note _____ Make sure that you re-create the <code>initrd</code> image as given in Step 6 . Make sure that you set Start-Up Mode to automatic before re-creating the <code>initrd</code> image.
4	To activate multipathing during booting up, run the following commands on the SAN booted host: <code>chkconfig boot.multipath on</code> <code>chkconfig multipathd on</code> Note _____ Make sure you do not configure the <code>open-iscsi</code> service to automatically the start at system boot.
5	Modify <code>/etc/sysconfig/network/ifcfg-eth-idxxxxxxxxx</code> and change the value of <code>STARTMODE</code> to <i>nfsroot</i> .
6	To create a new <code>initrd</code> image with the root partition on multipath enabled, run the following command: <code>mkinitrd -f "iscsi mpath"</code>

Step	Action
7	Reboot the host. Note _____ The system will now boot with the multipath enabled on the root device. This can be verified by running the <code>mount</code> command and ensuring that the root partition is on a dm-multipath device.

Note_____

If you rediscover the targets after root partition on multipath configuration, make sure that all the logged in connections are automatic, and re-create the `initrd` image. You must re-create the `initrd` image if the iSCSI or the network configuration are changed in any way.

Note

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

OCFS2 self fencing during cluster failover on SUSE Linux Enterprise Server 10 SP1 and SP2

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 can become unavailable, and none of the nodes in the OCFS2 cluster will be able to write or read the heartbeat values. As a result, the nodes may fence by panicking.

Workaround: The workaround for this problem is to set the `O2CB_HEARTBEAT_THRESHOLD` value in `/etc/sysconfig/o2cb` to a value large enough to ensure that the nodes will not self-fence.

NetApp recommends that, with a dedicated TCP/IP network for the iSCSI link between Linux hosts and the NetApp storage system, a reasonable value in most cases is 181. This 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. NetApp recommends that a value in the range of 61 to 201 would be ideal 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 SP1 and SP2**

Problem: With dm-multipath configured, `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 actual error and can be safely ignored. Use `multipath-ll` command which displays the LUN status correctly as *ready*.

**Automatically
mounting LVM
devices during
system boot on
SUSE Linux
Enterprise Server
10 SP1 and SP2**

Problem: Currently, Volume Groups created on iSCSI devices do not get automatically scanned when iSCSI LUNs are discovered. Therefore during system boot, the volume groups that were created on 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 vgname`
Here, `vgname` is the name of the Volume Group created on the iSCSI device.
- ◆ Add the generated rules to the `/etc/udev/rules.d/85-mount-fstab-lvm.rules` file.

After completing the preceding steps once for each Volume Group, 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 SP3**

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 multipath bootup and shutdown scripts in the `rc` script list in SUSE Linux Enterprise Server 9 SP3 is out of order with respect to the mount and unmount scripts for iSCSI and non-iSCSI devices.

`multipathd` and `iscsid` startup scripts are at the same serial number for a run level in `/etc/init.d/rc3.d`. Since alphabetically *i* comes before *m*, `iscsid` starts before `multipathd`.

As a result, the multipath bootup and shutdown sequence is not synchronized with iSCSI bootup and shutdown sequence. Multipath devices will not be 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 shutdown.

Workaround: Manual mounting and unmounting is required for multipath-based file systems on SUSE Linux Enterprise Server 9 SP3. After multipath devices are available on bootup, the file systems can be mounted.

For unmounting, make sure that any multipath-based file systems are unmounted before the multipath utility is shutdown.

The following sequence of actions should be performed to ensure proper shutdown of multipath-based iSCSI file systems:

Step	Action
1	Unmount all file systems on 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 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 SP3

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 can become unavailable, and none of the nodes in the OCFS2 cluster will be able to write or read the heartbeat values. As a result, the nodes may fence by panicking.

Workaround: The workaround for this problem is to set the O2CB_HEARTBEAT_THRESHOLD value in /etc/sysconfig/o2cb to a value large enough to ensure that the nodes will not self-fence.

NetApp recommends that, with a dedicated TCP/IP network for the iSCSI link between Linux hosts and the NetApp storage system, a reasonable value in most cases is 61. This 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. NetApp recommends that a value in the range of 61 to 201 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>.

LUNs are 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
Cabling	Verify that the cables between the host and the storage system are properly connected.
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 being used for iSCSI. ◆ From the host command line, ping the storage system interfaces that are being used for iSCSI.

Configuration setting	What to do
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 iSCSI Support Matrix at the following URL: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml
Jumbo frames	If you are using jumbo frames in your configuration, ensure that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
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>.
Initiator login	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 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 detailed information, see “Configuring the iSCSI initiator” on page 17 for SUSE Linux Enterprise Server 9 SP3, or “Configuring the iSCSI initiator” on page 58 for SUSE Linux Enterprise Server 10 SP1 and SP2.

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 must 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 LUNs and the igroups to which they are mapped. <code>lun show -g igroup-name</code> displays the LUNs mapped to a specific igroup.

LUNs not recognized by dm-multipath

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

```
multipath -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 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 SP3)

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. Make sure this line is at the beginning of the file, before any “DiscoveryAddress” entries.

```
Continuous=no
```

iSCSI performance is poor

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 may 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>.

