

iSCSI SUSE® Linux® Host Utilities 3.0

Setup Guide

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

Part number 215-02057_A0
October 2006

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, Data ONTAP, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. FlexVol, NOW, NetApp on the Web, RAID-DP, SnapValidator, and vFiler are trademarks of Network Appliance, Inc. in the United States and other countries.

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	v
Chapter 1	Setup Overview	1
Chapter 2	Installing software on the host	7
	Installing the Host Utilities	8
	Verifying or installing the required RPMs	10
	Editing the host's /etc/multipath.conf file	11
	Recording or changing the initiator node name on the host	14
	Editing the host's /etc/iscsi.conf file	17
Chapter 3	Configuring the storage system for iSCSI	21
Chapter 4	Accessing LUNs from the host	23
	Starting the iSCSI service on the host	24
	Configuring the iSCSI service to start automatically.	25
	Starting the multipath service.	26
	Configuring the multipath service to start automatically.	27
	Accessing LUNs with SLES 8 SP4.	28
	Creating a partition	30
	Creating a file system on a partition of a LUN.	32
	Labeling and mounting the file system.	33
	Accessing LUNs as raw devices	35
	Accessing LUNs with SLES 8 SP3 and SLES 9	36
	Creating a file system on a LUN	43
	Accessing LUNs as raw devices	46
	Accessing LUNs with dm-multipath support.	47
	Verifying the dm-multipath policy and mapping	48
	Viewing a list of LUNs.	49
	Creating and mounting a file system on a LUN	51
	Accessing LUNs as raw devices	53
	Tuning the dm-multipath configuration	54

Stopping dm-multipath 55

Appendix A **Troubleshooting 57**

Preface

About this guide

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

See the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, available on the NetApp on the Web™ (NOW) site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ for information on the specific Linux versions supported.

Audience

This guide is for administrators of SUSE Linux Enterprise Server hosts and NetApp storage systems.

Terminology

This guide uses the following terms:

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

Command conventions

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

Formatting conventions

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

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

Keyboard conventions

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

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

Special messages

This guide contains special messages that are described as follows:

Note

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

Caution

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

Understanding how Data ONTAP® software implements iSCSI

The iSCSI protocol is a licensed service on the storage system that enables you to transfer block data to hosts using the SCSI protocol over TCP/IP. The iSCSI protocol standard is defined by RFC 3720 (www.ietf.org).

In an iSCSI network, storage systems are targets that have storage target devices, which are referred to as LUNs (Logical Units). A Linux host running the iSCSI initiator software uses the iSCSI protocol to access 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 *Data ONTAP Block Access Management Guide* for your version of Data ONTAP.

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) 1.2.1 (o2cb the cluster service as part of 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 Oracle documentation Web site. For details see “[Where to go for more information](#)” on page 4.

Understanding device mapper multipathing

Linux device mapper multipathing (dm-multipath) is supported on NetApp storage systems with SLES 9 SP3. This support enables you to configure multiple network 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 SLES 9 SP3 binary kernel RPM. You may have to install additional userspace RPMs 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/disk/by-name/` for each LUN that represents all the paths. You should create file systems and mount the LUN using the device in `/dev/disk/by-name/`. You do not need to create partitions or labels when using LUNs with dm-multipath support; the device in `/dev/disk/by-name/` creates a persistent association with a LUN.

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

Note

Automounting and autounmounting of file systems on multipath devices do not work on SLES9 SP3. See [“Automounting and autounmounting of file system on multipath device do not work on SLES 9 SP3”](#) on page 57 for details

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 software on the host, including: ◆ Installing the iSCSI Linux host utilities software. ◆ Installing required RPMs on the host and recording or changing the host's iSCSI node name. ◆ Optionally configuring dm-multipath support ◆ Setting initiator parameters, including the IP address of the target on the storage system ◆ Optionally configuring CHAP See Chapter 2, “Installing software on the host,” on page 7.

Step	Action
2	Configure the storage system, including: ◆ Licensing and starting the iSCSI service ◆ Optionally 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 for iSCSI,” on page 21.
3	Access the LUNs from the host, including: ◆ Starting the iSCSI service ◆ Creating file systems on the LUNs and mounting them, or configuring the LUNs as raw devices ◆ Creating persistent mappings of LUNs to file systems See Chapter 4, “Accessing LUNs from the host,” on page 23.

Note

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

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...	Go to...
Changes in this release of the Host Utilities, a list of Host Utilities known problems, and limitations	The <i>Release Notes</i> for this Host Utilities.
The most current system requirements for the Windows host	<i>NetApp iSCSI Support Matrix</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml
Best practices or configuration issues	NetApp Knowledge Base at https://now.netapp.com/eservice/kbbrowse
The supported storage system models for Data ONTAP licensed with iSCSI	◆ <i>NetApp iSCSI Support Matrix</i> at: http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml ◆ <i>System Configuration Guide for Data ONTAP</i> online version at: http://now.netapp.com/NOW/knowledge/docs
Information about how to configure and manage the iSCSI service on the storage system	The following documentation for your release of Data ONTAP: ◆ <i>Data ONTAP Block Access Management Guide</i> ◆ <i>Data ONTAP Release Notes</i>
Information about managing the iSCSI initiator	◆ The README file in /usr/share/doc/packages/linux-iscsi on the host. ◆ The iscsi.conf, iscsi-ls, and iscsid man pages on the host. ◆ Documentation on the SUSE Web site at http://www.novell.com/linux/suse/portal/
Information about configuring multipathing on SLES	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

If you want...	Go to...
Information about configuring multipathing on SLES	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 on OCFS2	◆ Oracle Cluster File System (OCFS2) User's Guide http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf ◆ OCFS2 - Frequently Asked Questions http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html (The FAQs on OCFS2)
Information on Oracle RAC	http://www.oracle.com/technology/pub/articles/smiley_rac10g_install.html
NetApp Knowledge Base for Oracle RAC	http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190 http://www.netapp.com/library/tr/3369.pdf http://www.netapp.com/library/tr/3423.pdf

Configuration overview

This chapter explains how to install the host utilities and how to complete the installation and initial configuration of the iSCSI initiator and multipathing software on the Linux host. You will complete the final configuration of the initiator after you have configured the storage system.

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 linux-iscsi RPM.

If you want to use the Linux device mapper multipathing (dm-multipath) support to create a highly-available connection between the Linux host and storage system, you also need the device-mapper and multipath-tools RPMs.

Note

Do not start the iSCSI or multipath services on the Linux host until instructed to in [Chapter 4](#) on page 23.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Installing the Host Utilities](#)” on page 8
- ◆ “[Verifying or installing the required RPMs](#)” on page 10.
- ◆ “[Editing the host's /etc/multipath.conf file](#)” on page 11
- ◆ “[Recording or changing the initiator node name on the host](#)” on page 14.
- ◆ “[Editing the host's /etc/iscsi.conf file](#)” on page 17.

Requirement for root privileges

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

Installing the Host Utilities

Verifying the correct version

Verify that you have the correct host utilities for your version of Linux. See the *NetApp iSCSI Support Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

Getting the host utilities software

Download the iSCSI Linux host utilities file (netapp_linux_tools_3_0.tar.gz) from the NOW Web site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on your Linux host.

Installing the host utilities software

To install the host utilities software, complete the following steps:

Step	Action
1	Remove any previous version of the host utilities. 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	Enter the following command to uncompress the file: <code>gunzip netapp_linux_tools_3_0.tar.gz</code>
4	Enter the following command to extract the files: <code>tar -xvf netapp_linux_tools_3_0.tar</code>
5	Change to the netapp_linux_tools_3_0 directory. By default, this directory is a subdirectory of the working directory in which you extracted the host utilities files in the previous step.
6	Enter the following command: <code>./install</code>

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

To verify that the correct iSCSI and multipathing RPMs are installed on your Linux host, 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 RPM. See the <i>Release Notes</i> for the correct RPM version for your specific version of Linux.	
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 RPMs. See the <i>Release Notes</i> for the correct RPM version for your specific version of Linux.	
3	If...	Then...
	You have the correct RPMs for your version of SLES.	Proceed to “ Editing the host’s /etc/multipath.conf file ” on page 11.
	You do not have the correct RPMs for your version of SLES.	Install the required RPM versions from the media for your version of SLES. You can use the <code>rpm</code> or <code>yast</code> command.

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

Note_____

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

Copying the file

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

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

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

Editing the file

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

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

Step	Action
1	Open <code>/etc/multipath.conf</code> with a text editor.
2	Locate the section labeled <code>## name : blacklist</code>
3	Delete the comment characters (<code>#</code>) to enable the code in the blacklist section.
4	If the command in the blacklist section is <code>blacklist</code> , rename the command <code>devnode_blacklist</code> .

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

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

```
##
```

```

## name      : blacklist
## scope     : multipath & multipathd
## desc      : list of device names to discard as not multipath
               candidates
## default   : cciss, fd, hd, md, dm, sr, scd, st, ram, raw, loop
##
devnode_blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z][0-9]*"
    devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"
    devnode "sd[a-b]$"
}
#
##

```

Recording or changing the initiator node name on the host

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

The default initiator node name uses the following format:

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

It is recommended that you set the initiator node name to:

`iqn.1987-05.com.cisco:` *host-name*

where *host-name* is the host name of your Linux host.

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

Note

Testing of SLES 8 on i386 systems found that the same default node name is generated on all hosts. Be sure to set a new node name if you have more than one host running SLES 8.

To view or change the default 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	1. Replace the text in the host's <code>/etc/initiatorname.iscsi</code> file with the following line: <code>iqn.1987-05.com.cisco:RandomNumber</code> The recommended value for <i>RandomNumber</i> is the host name of the Linux host. 2. Record the node name and proceed to “Editing the host's /etc/iscsi.conf file” on page 17.
	You want to use the default initiator node name	Record the node name and proceed to “Editing the host's /etc/iscsi.conf file” on page 17. If you are using SLES 8, do not use the default node name.
	You want to change the initiator node name	1. Modify the <i>RandomNumber</i> part of the initiator node name. The following line shows an example node name. <code>iqn.1987-05.com.cisco:linux-host1</code> 2. Record the node name and proceed to “Editing the host's /etc/iscsi.conf file” on page 17. Note_____ The <i>RandomNumber</i> is the only component of the node name you modify. You do not modify the other components of the node name. It is recommended that you 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 these rules:

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

Editing the host's /etc/iscsi.conf file

You edit the /etc/iscsi.conf file to configure:

- ◆ Required iSCSI parameter 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.
2	At the beginning of the file, add the following lines. Make sure you add these lines before any other settings: Caution_____ These lines are required for proper operation. If you do not add these lines at the beginning of the file, you will experience problems with connectivity and performance. _____ Continuous=no HeaderDigest=never DataDigest=never
3	If you have SLES 8 and you want to enable ImmediateData support for improved write performance, add the following line to the file: ImmediateData=yes Note_____ For SLES 8 SP 4 and SLES 9, ImmediateData is set to yes by default. _____
4	If you are using dm-multipath support: ◆ Remove the comment to enable the Multipath=portal line in the Multipath Settings section. ◆ Remove the comment from the ConnFailTimeout line in the Session Timeout Settings section and set the value to ConnFailTimeout=5.

Step	Action
5	Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you will use for iSCSI LUNs: <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 interface 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
6	If you want to use Challenge-Handshake Authentication Protocol (CHAP), add CHAP user names and passwords, as shown in the example below. 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 OutgoingUsername and OutgoingPassword. Example: <pre>OutgoingUsername=linuxhostout OutgoingPassword=aow9857f1 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 OutgoingUsername and OutgoingPassword for the storage system's inbound user name and password (<i>inname</i> and <i>inpassword</i>). ◆ If you are using bidirectional authentication, use IncomingUsername and IncomingPassword for the storage system's outbound user name and password (<i>outname</i> and <i>outpassword</i>) for incoming settings.

Before you begin

You need the host's iSCSI initiator node name that you recorded in [Chapter 2](#) 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:

- ◆ Ensure that the iSCSI service is licensed and started.
- ◆ If you want to use CHAP authentication, use the `iscsi security` command or the FilerView interface to configure a CHAP user name and password.
- ◆ Create LUNs, create an igroup using the host's initiator 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 `linux`.

For detailed configuration steps

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

Data ONTAP Upgrade note

If you upgrade the Data ONTAP software running on the storage system from version 6.4.x to version 6.5 or later, the CHAP configuration on the storage system is not saved. The format of the CHAP configuration file changed in Data ONTAP 6.5. Be sure to run the `iscsi security add` command for each iSCSI initiator, even if you had previously configured the CHAP settings.

Accessing LUNs from the host

Starting the iSCSI service

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.

Overview of accessing LUNs

The steps for accessing LUNs depend on which version of SUSE Linux 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 the host”](#) on page 24
- ◆ [“Configuring the iSCSI service to start automatically”](#) on page 25
- ◆ [“Starting the multipath service”](#) on page 26
- ◆ [“Configuring the multipath service to start automatically”](#) on page 27
- ◆ [“Accessing LUNs with SLES 8 SP4”](#) on page 28
- ◆ [“Accessing LUNs with SLES 8 SP3 and SLES 9”](#) on page 36
- ◆ [“Accessing LUNs with dm-multipath support”](#) on page 47

Starting the iSCSI service on the host

To start the iSCSI service, complete the following step.

Step	Action
1	Start the iSCSI service by entering the following command at the Linux host command prompt: <code>/etc/init.d/iscsi start</code>

Configuring the iSCSI service to start automatically

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 iscsi	
2	If...	Then...
	The <code>chkconfig</code> command indicates the iSCSI service is enabled, as follows: <code>iscsi on</code>	No action needed.
	The <code>chkconfig</code> command indicates that the <code>iscsi</code> service is not enabled as follows: <code>iscsi off</code>	Start the iSCSI service by entering the following command: chkconfig iscsi on

Starting the multipath service

Note_____

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

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

Step	Action
1	Start the scan for multipath LUNs by entering the following command: <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

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

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

Step	Action
1	Add the multipath service to the boot sequence by entering the following command: <code>insserv boot.multipath multipathd</code>

Accessing LUNs with SLES 8 SP4

Understanding device persistence

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

If you have more than one LUN, you need to have a persistent way of identifying each LUN. The iSCSI initiator included in SLES 8 SP3, SLES 9, and SLES 9 SP1 automatically generates persistent symbolic links for the LUNs.

With the initiator included in SLES 8 SP4, those symbolic links are not created. For SLES 8 SP4, you can write a file system label to the LUN and mount the file system by that label to provide a persistent identifier.

Understanding automounting

The Linux system mounts all of the file systems listed in the `/etc/fstab.iscsi` file when it first boots up. When you create file systems labels, you must add them to the `/etc/fstab.iscsi` file so that the file systems will be automounted.

Viewing information about LUNs

You can view information about available LUNs in the following ways:

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

Tasks for accessing LUNs

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

- ◆ If you are creating a file system on the LUN:
 - ❖ Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN.
 - ❖ Create a file system on the partition.

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

For detailed information

For detailed information, see the following sections:

- ◆ [“Creating a partition”](#) on page 30
- ◆ [“Creating a file system on a partition of a LUN”](#) on page 32
- ◆ [“Labeling and mounting the file system”](#) on page 33
- ◆ [“Accessing LUNs as raw devices”](#) on page 35

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 still need to create a partition.

Before you begin

You should have one or more LUNs available to the host. Log into the host as root and run the following command to verify the LUNs are available:

```
iscsi-ls -l
```

You should see information about the iSCSI driver, including a description of available LUNs. The following example shows one LUN:

```
DEVICE DETAILS :
-----
LUN ID   : 0
Vendor:  NETAPP   Model:  LUN               Rev:  0.2
Type:    Direct-Access      ANSI SCSI revision: 04
page83  type3:  60a980004f6443745359763759367733
page83  type1:
4e45544150502020204c554e204f644374535976375936773300000000000000
page80:  4f6443745359763759367733
Device:  /dev/sda
```

In this case, the LUN with ID 0 is mapped to the device /dev/sda.

Using fdisk to create a partition

To create a partition on a LUN:

Step	Action
1	Enter /sbin/fdisk device <i>device</i> is the device name assigned to the LUN, such as /dev/sda.
2	Enter the command n to create a new partition.
3	Enter p to create a primary partition
4	Enter a partition number. For the first partition, enter 1 .

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

Creating a file system on a partition of a LUN

About file systems A file system enables users and applications to read and write files on a LUN. Linux supports several different file systems. These instructions show how to create a journaling ext3 file system, but you can create any type you like.

If your application needs a raw device, do not create a file system. Instead, see [“Accessing LUNs as raw devices”](#) on page 35.

Creating a file system To create a file system on the partition you created, complete the following step:

Step	Action	
1	If...	Then...
	You want to create and label a Reiser file system	Enter the following command: mkreiserfs --label <i>text</i> <i>device</i> <i>text</i> is the file system label used to create a persistent identifier for the file system. <i>device</i> is the partition. For example /dev/sda1 is the first partition on /dev/sda.
	You want to create an ext3 file system	Enter the following command: mkfs -t ext3 <i>device</i> <i>device</i> is the partition. For example /dev/sda1 is the first partition on /dev/sda.

Accessing LUNs with SLES 8 SP4

Labeling and mounting the file system

About labels

A file system label creates a persistent mapping between a LUN and a file system. See “[Understanding device persistence](#)” on page 28 for more information.

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

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:

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

Note _____
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.

Mounting the file system

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

Step	Action
1	Create a mount point for the file system. Example: mkdir /mnt/ss0_lun0
2	Open the file /etc/fstab.iscsi using a text editor.

Step	Action
3	Add an entry to /etc/fstab.iscsi as follows, separating each entry with a Tab (except the final two zeros, which are separated by a space): <pre>LABEL=<i>label</i> <i>mount_point</i> <i>type</i> defaults 0 0</pre> <i>label</i> is the file system label you created above <i>mount_point</i> is the mount point you created for the file system <i>type</i> is the file system type, such as ext2, ext3, or reiserfs Example: <pre>LABEL=lun0_label /mnt/ss0_lun0 ext3 defaults 0 0</pre>
4	Save your changes to /etc/fstab.iscsi and exit the editor.
5	Enter /usr/sbin/iscsi-mountall to mount the file system.
6	Enter df to display the mounted file systems. Verify the new file system is listed. Another option is to change to the mount point directory and try writing and reading a file to verify the file system mounted successfully.

Accessing LUNs as raw devices

To access LUNs as raw devices, complete the following step.

Step	Action
1	On the host console, map a raw device to an iSCSI device in the /dev/sdx directory by entering the following command: <pre>raw /dev/raw/rawN iscsi_device</pre> <i>N</i> is a number identifying the raw device. <i>iscsi_device</i> is the device name assigned to the iSCSI LUN. Example: <pre>raw /dev/raw/raw1 /dev/sdb</pre>

Accessing LUNs with SLES 8 SP3 and SLES 9

How LUNs are integrated into the Linux file system

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 might be inconsistent each time the host iSCSI service is started, and SCSI commands might be sent to the wrong storage system or LUN. If you have applications that access storage system LUNs, you must ensure that these applications use the symbolic links instead of the standard SCSI device nodes.

SLES 8 SP3: For SLES 8 SP3, the initiator software creates a tree of directories and symbolic links under the host's `/dev/iscsi` directory. This directory has subdirectories for each SCSI bus number, target ID on that bus, and LUN ID. The LUN subdirectory contains symbolic links that map the LUN to a SCSI disk device.

For example, the local disk device that corresponds to LUN 0 on the storage system is `/dev/iscsi/bus0/target0/lun0/disk`. The symbolic link `/dev/iscsi/bus0/target0/lun0/generic` maps to the SCSI generic interface to the disk. The symbolic links `/part1` through `/part15` represent the disk partitions.

SLES 9: For SLES 9, the iSCSI devices are under the `/dev/disk/by-id/` directory in the following format:

`iscsi-<target_name>-LUN_ID > symbolic_link_to_disk`

The following shows example shows three iSCSI devices in the `/dev/disk/by-id/` directory. For the second device listed, the “-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
```

Viewing information about LUNs

You can view information about LUNs in the following ways:

- ◆ Symbolic links—The `ls -l` command lists the symbolic links the initiator creates under the `/dev/iscsi` directory (SLES 8) or under the `/dev/disk/by-id` directory (SLES 9).
- ◆ `iscsi-ls -l` command—The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.
- ◆ `sanlun` command—The `sanlun` command is included in the iSCSI Linux Host Utilities. It displays the host device names and the storage system LUNs to which they are mapped.

SLES 8 SP3: Viewing symbolic links. To view symbolic links in the `/dev/iscsi` directory, complete the following step.

Step	Action
1	On the SLES 8 host console, enter the following command: <pre>ls -l /dev/iscsi/busID/targetID/lunID</pre> Example: The following sample symbolic links show that LUN 0 is mapped to the first SCSI device <code>/dev/sda</code>: <pre># ls -l /dev/iscsi/bus0/target0/lun0 /dev/iscsi/bus0/target0/lun0: total 0 lrwxrwxrwx 1 root root 8 Jun 9 15:45 disk -> /dev/sda lrwxrwxrwx 1 root root 8 Jun 9 15:45 generic -> /dev/sg0 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part1 -> /dev/sda1 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part10 -> /dev/sda10 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part11 -> /dev/sda11 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part12 -> /dev/sda12 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part13 -> /dev/sda13 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part14 -> /dev/sda14 lrwxrwxrwx 1 root root 10 Jun 9 15:45 part15 -> /dev/sda15 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part2 -> /dev/sda2 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part3 -> /dev/sda3 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part4 -> /dev/sda4 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part5 -> /dev/sda5 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part6 -> /dev/sda6 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part7 -> /dev/sda7 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part8 -> /dev/sda8 lrwxrwxrwx 1 root root 9 Jun 9 15:45 part9 -> /dev/sda9</pre>

SLES 9: Viewing symbolic links: To view symbolic links and the storage system node name to which they are mapped, complete the following step:

Step	Action
1	On the SLES 9 host console, enter the following command: <pre># ls -l /dev/disk/by-id</pre> Example: <pre># ls -l /dev/disk/by-id 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</pre>

Viewing storage system and LUN details: To view storage system and LUN details, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi-ls -l</pre> Result: The console lists available storage systems and LUNs. Example: The following example shows output for SLES 8. The iSCSI Driver Version depends on whether you have SLES 8 or SLES 9. <pre>***** Cisco iSCSI Driver Version ... 3.4.0.3 (22-Aug-2003) ***** TARGET NAME : iqn.1992-08.com.netapp:sn.12345678 TARGET ALIAS : HOST NO : 0 BUS NO : 0 TARGET ID : 1 DISCOVERY ADDRESS : 10.10.10.10:3260 SESSION ESTABLISHED AT : Mon Aug 22 15:44:58 2003 NO. OF PORTALS : 2 TARGET ADDRESS 1 : 10.10.10.10:3260,2 TARGET ADDRESS 2 : 10.10.10.11:3260,1 SESSION ID : ISID 00023d000001 TSID 239 DEVICE DETAILS : ----- LUN ID : 0 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 03 LUN ID : 1 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 03</pre>

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

Step	Action
1	On the host, change to the /opt/netapp/santools directory.
2	Enter the following command: <pre>sanlun lun show all</pre> Result: The <code>sanlun</code> command displays information about iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. Example: The following lines show sample <code>sanlun</code> command output. The adapter is displayed as <code><unknown></code> because the host is using a software initiator instead of a Host Bus Adapter (HBA). <pre>[root@host_1 ~]# /opt/netapp/santools/sanlun lun show all filer: lun-pathname device filename adapter protocol lun size lun state ss1: /vol/luns/iscsilun0 /dev/sda <unknown> iSCSI 10g (10737418240) GOOD ss1: /vol/luns/iscsilun1 /dev/sdb <unknown> iSCSI 10g (10737418240) GOOD ss1: /vol/luns/iscsilun2 /dev/sdc <unknown> iSCSI 10g (10737418240) GOOD ss1: /vol/luns/iscsilun3 /dev/sdd <unknown> iSCSI 10g (10737418240) GOOD</pre>

Viewing information about target configuration To view information about target configuration, complete the following step.

Step	Action
1	Enter the following command: <pre>iscsi-ls -c</pre> Result: The console lists available storage systems and LUNs, and iSCSI session parameters Example: The following lines show example output for SLES 9. The iSCSI Driver Version depends on whether you are running SLES 8 or SLES 9. <pre>***** 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</pre>

Accessing LUNs

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

- ◆ Creating a file system on the LUN—For details, see “[Creating a file system on a LUN](#)” on page 43.
- ◆ Accessing the LUN as a raw device—For details, see “[Accessing LUNs as raw devices](#)” on page 35.

Creating a file system on a LUN

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

Note Reiser file systems (ReiserFS) are not supported for SLES 9. They are supported for SLES 8 SP3 and SLES 9 SP1.

Using fdisk to create a partition

To create a partition of a LUN:

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

Creating a file system

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

Step	Action
1	On the host console, enter the following command: <pre>mkfs -V -t <i>filesystem_type</i> <i>partition</i></pre> <i>partition</i> is the name of the partition you just created. The partition name is different for SLES 8 and SLES 9. For SLES 8 SP3, the partition is the symbolic link to the LUN with <i>partn</i>. For SLES 9, the partition name is the name of the link that maps the iSCSI LUN on the storage system to a local disk device concatenated with the letter p and the partition number. <i>filesystem_type</i> is the type of file system you want to create. SLES 8 example: The following example creates an ext3 file system on the first partition of LUN 0. <pre>mkfs -V -t ext3 /dev/iscsi/bus0/target0/lun0/part1</pre> SLES 9 example: The following example creates an ext3 file system on the first partition of LUN 0. <pre>mkfs -V -t ext3 /dev/disk/by-id/iscsi-ign.1992-08.com.netapp:sn.12345678.ss_hq-0p1</pre>

Automount the file system

To automount the file system, complete the following steps:

Step	Action
1	On the host, create a mount point for the LUNs.
2	With a text editor, open the host's <code>/etc/fstab.iscsi</code> file.
3	Add an entry for the file system you created in “Creating a file system on a partition of a LUN” on page 32, SLES 8 example: <pre>/dev/iscsi/bus0/target0/lun0/part1 /mnt_pt ext3 defaults 0 0</pre> SLES 9 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 mount point 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 SLES 8 SP3 and SLES 9

Accessing LUNs as raw devices

SLES 8 SP3

To access LUNs as raw devices on SLES 8, complete the following step.

Step	Action
1	On the host console, map a raw device to an iSCSI device in the /dev/iscsi directory by entering the following command: <pre>raw /dev/raw/rawN iscsi_device</pre> <i>N</i> is a number identifying the raw device. <i>iscsi_device</i> is the symbolic link that maps the iSCSI LUN on the storage system to a local disk device. Example: <pre>raw /dev/raw/raw1 /dev/iscsi/bus1/target1/lun1/disk</pre>

SLES 9

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

Accessing LUNs with dm-multipath support

SLES 9 SP3 only

Support for dm-multipath is currently on SLES 9 SP3 only. However, new configurations might be qualified between host utilities releases. For the latest information on supported configurations, see the *NetApp iSCSI Support Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml.

Overview

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

- ◆ “[Verifying the dm-multipath policy and mapping](#)” on page 48
- ◆ “[Viewing a list of LUNs](#)” on page 49
- ◆ “[Creating and mounting a file system on a LUN](#)” on page 51
- ◆ “[Accessing LUNs as raw devices](#)” on page 53
- ◆ “[Tuning the dm-multipath configuration](#)” on page 54
- ◆ “[Stopping dm-multipath](#)” on page 55

Verifying the dm-multipath policy and mapping

Overview

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

Verifying the configuration

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

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

Updating the configuration

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

```
multipath -v2
```

Viewing a list of LUNs

Using sanlun

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

```
sanlun lun show all
```

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

```
host:/ # sanlun lun show all
filer:  lun-pathname      device filename  adapter  protocol  lun size    lun state
ss1:    /vol/vol1/linux_1  /dev/sdc        <unknown>0  iSCSI     5g (5368709120)  GOOD
ss2:    /vol/vol1/linux_2  /dev/sdd        <unknown>0  iSCSI     6g (6442450944)  GOOD
ss1:    /vol/vol1/linux_1  /dev/sde        <unknown>0  iSCSI     5g (5368709120)  GOOD
ss2:    /vol/vol1/linux_2  /dev/sdf        <unknown>0  iSCSI     6g (6442450944)  GOOD
```

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

Using iscsi-ls

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 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
```

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
```

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

Adding new LUNs to the dm-multipath configuration

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

```
multipath
```

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

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

Accessing LUNs with dm-multipath support

Creating and mounting a file system on a LUN

Overview

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

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

Creating the file system

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

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

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

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

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

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

Mounting the file system

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

To mount the file system, complete the following steps:

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

Step	Action
3	Add an entry to <code>/etc/fstab.iscsi</code> as follows, separating each entry with a Tab (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/disk/by-name/</code> directory <i>mount_point</i> is the mount point you created for the file system <i>type</i> is the file system type, such as <code>ext2</code>, <code>ext3</code>, or <code>reiserfs</code> Example: <pre>/dev/disk/by-name/360a9800043346536684a337a7553424f /mnt/ss1/lun_0 ext3 defaults 0 0</pre>
4	Enter <code>/usr/sbin/iscsi-mountall</code> to mount the file system.
5	Enter <code>df</code> to display the mounted file systems. Verify the new file system is listed. Another option is to change to the mount point directory and try writing and reading a file to verify the file system mounted successfully.

Accessing LUNs as raw devices

Overview

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

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

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

Tuning the dm-multipath configuration

Tuning multipath parameters

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

rr_min_io: The `rr_min_io` setting for dm-multipath specifies the number of I/O operations sent through a path before switching to the next path. Lowering this value from the default value of 1000 has been seen to dramatically affect overall throughput for dm-multipath, especially for large I/O workloads (64 KB or more) 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

Overview

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

Stopping the services

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

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>

Note

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

Automounting and autounmounting of file system on multipath device do not work on SLES 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 SLES9 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 SLES9 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>

Step	Action
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 SLES 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 iSCSI link between Linux hosts and NetApp filer, 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:

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.
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> .

Configuration setting	What to do
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 “Editing the host’s /etc/iscsi.conf file” on page 17.
iSCSI node names	Verify that you are using the correct initiator node names in the <code>igroup</code> 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 <code>igroups</code>. On the host, use the initiator tools and commands to display the initiator node name. The initiator node names configured in the <code>igroup</code> and on the host must match.
LUN mappings	Verify that the LUNs are mapped to an <code>igroup</code>. On the storage system, use one of the following commands: <code>lun show -m</code> displays all LUNs and the <code>igroups</code> to which they are mapped. <code>lun show -g <i>igroup-name</i></code> displays the LUNs mapped to a specific <code>igroup</code>.

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

If the 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

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>