iSCSI Red Hat® Enterprise Linux® Host Utilities 3.2
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 215-04008_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 **Setup Overview** 1

Chapter 2 **Installing software on the host** 7

Installing the Host Utilities 8

Removing the previous initiator 10

Verifying or installing the required RPMs 11

Editing the /etc/multipath.conf file 13

Starting the multipath service. 18

Configuring the multipath service to start automatically 19

Recording or changing the initiator node name on the host 20

Editing the host's configuration file 24

Chapter 3 **Configuring the storage system for iSCSI** 29

Chapter 4 **Accessing LUNs from the host** 35

Starting the iSCSI service on the host 36

Configuring the iSCSI service to start automatically. 37

Accessing LUNs using dm-multipath 38

- Viewing a list of LUNs. 39
- Viewing the dm-multipath configuration and multipath devices 43
- Creating and mounting a file system on a multipath device 44
- Accessing LUNs as raw devices 46
- Tuning the dm-multipath configuration 47

Accessing LUNs without dm-multipath 48

- Creating a partition 51
- Creating a file system 53
- Labeling and mounting the file system. 54
- Accessing LUNs as raw devices 56
- Viewing LUN information 57

Appendix A **Troubleshooting** 63

Preface

About this guide

This guide describes how to configure the iSCSI initiator in Red Hat® Enterprise Linux® 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 support.

Audience

This guide is for administrators of Red Hat Enterprise Linux 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
Italic type	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the `fctest` *adaptername* command, you enter the characters "fctest" 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.
Enter	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
type	Used to mean pressing one or more keys on the keyboard.
enter	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.

Setup Overview 1

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 Logical Units. A Linux host running the iSCSI initiator software uses the iSCSI protocol to access Logical Unit Numbers (LUNs) on a storage system running Data ONTAP® software. The 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 the Clustering suite on Linux hosts.

See "Where to go for more information" on page 4 for details.

Understanding device mapper multipathing

Linux device mapper multipathing (dm-multipath) is supported on NetApp storage systems with RHEL 5 series, RHEL 4 series Update 3 and later. 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 RHEL 5 series, RHEL 4 series Update 3 and later 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/sd*x* 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 should create file systems and mount the LUN using the device 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 throughout this guide.

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.
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 29.

Step	Action
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 35.

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 files, known problems, and limitations	The *Release Notes* for this Host Utilities.
The most current system requirements for the Windows host	*NetApp iSCSI Support Matrix* 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 http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190
The supported storage system models for Data ONTAP licensed with iSCSI	◆ *NetApp iSCSI Support Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/iscsi_support_matrix.shtml ◆ *System Configuration Guide for Data ONTAP* 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 *Block Access Management Guide* ◆ Data ONTAP *Release Notes*
Information about managing the iSCSI initiator	For RHEL 5 series: ◆ The iscsid.conf, iscsiadm and iscsid man pages on the host. For RHEL 4 series Update 3 and later: ◆ The iscsi.conf, iscsi-ls, and iscsid man pages on the host. Red Hat Enterprise Linux Documentation: ◆ Documentation on the Red Hat at http://www.redhat.com/docs/manuals/enterprise/

If you want...	Go to...
Information about configuring multipathing on RHEL	The Multipath example files included with the device-mapper-multipath RPM are in /usr/share/doc/device-mapper-multipath-<version> where <version> is the latest version number shipped with the RHEL release.
Information about installing and configuring SnapDrive for UNIX	The *SnapDrive® for UNIX® Installation and Administration Guide* for your version of SnapDrive.
Information on Red Hat Cluster and GFS	http://www.redhat.com/docs/manuals/csgfs/
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://www.netapp.com/library/tr/3423.pdf http://www.netapp.com/library/tr/3369.pdf

Installing software on the host 2

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.

Red Hat Enterprise 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 iscsi initiator 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 device-mapper-multipath RPMs.

Note

Do not start the iSCSI service on the Linux host until instructed to in "Starting the iSCSI service on the host" on page 36.

Topics in this chapter

This chapter includes the following topics:

- "Installing the Host Utilities" on page 8
- "Removing the previous initiator" on page 10
- "Verifying or installing the required RPMs" on page 11.
- "Editing the /etc/multipath.conf file" on page 13
- "Starting the multipath service" on page 18
- "Configuring the multipath service to start automatically" on page 19
- "Recording or changing the initiator node name on the host" on page 20.
- "Editing the host's configuration file" on page 24.

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 release for your version of Linux. See the *NetApp iSCSI Support Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Downloading the host utilities software

Download the iSCSI Linux host utilities file (netapp_linux_host_utilities_3_2.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 you downloaded, complete the following steps:

Step	Action
1	Remove any previous version of the Host Utilities (formerly, Support Kit). 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 downloaded the host utilities file.
3	Enter the following command to uncompress the file: `gunzip netapp_linux_host_utilities_3_2.tar.gz`
4	Enter the following command to extract the files: `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 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.

Removing the previous initiator

Removing an iSCSI initiator downloaded from SourceForge

You need to remove any previous iSCSI initiator downloaded from SourceForge before upgrading to any versions later to Red Hat 3.0 Update 4. To remove the iSCSI initiator:

Step	Action
1	Log in to the Linux host as root.
2	Stop the iSCSI service by running the `/etc/init.d/iscsi stop` command.
3	Change to the directory from which the previous iSCSI initiator was installed. The directory name is usually linux-iscsi-*version*.
4	Run the `make remove` command.
5	Change to the parent directory using the `cd ..` command.
6	Remove the old iSCSI initiator directory using the `rm -rf linux-iscsi-`*`version`* command.

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: **`rpm -q iscsi-initiator-utils`** **Result:** The `rpm -q` command returns the name and version of the iSCSI RPM. See the *Release Notes* for the correct RPM version for your specific version of Linux.
2	If you plan to use dm-multipath support, enter the following commands: **`rpm -q device-mapper`** **`rpm -q device-mapper-multipath`** **Result:** The `rpm -q` commands return the name and version of the installed RPMs. See the *Release Notes* for the correct RPM version for your specific version of Linux.

Step	Action	
3	**If...**	**Then...**
	You have the correct RPMs for your version of Linux.	Proceed to "Editing the /etc/multipath.conf file" on page 13.
	You do not have the correct RPMs for your version of Linux.	Install the iscsi-initiator-utils RPM. You can get the RPM either from the Red Hat Enterprise Linux media or from the Red Hat Network at http://www.redhat.com/software/rhn/ (subscription required). **Example:** In the following procedure example, you use the Red Hat Network. 1. Using an X terminal or local console, log in to your Linux host as root. 2. Enter `rpm -Uvh iscsi-initiator-utils-<version>.rpm` to install or upgrade the rpm package. 3. Enter the following command to verify the RPM installation: `rpm -q iscsi-initiator-utils` Proceed to "Editing the /etc/multipath.conf file" on page 13

Editing the /etc/multipath.conf file

Note

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

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:

For RHEL 5 series:

Step	Action
1	Open /etc/multipath.conf with a text editor.
2	Comment out the following line: `blacklist {devnode "*" }`
3	Delete the comment characters from the remaining blacklist command.
4	The default command excludes IDE hard drives /dev/hdx devices from /dev/hda through /dev/hdz. If your local drives are IDE drives in this range, no further changes to the blacklist section are needed.

Step	Action

5	If your local drives are SCSI drives, add the wwid within the `blacklist` command to exclude the specific /dev/sd*x* devices. Be sure to blacklist only your local drives. Each path to an iSCSI LUN is also a /dev/sd*x* device that should be configured for multipathing.
	
	To obtain the WWID of a device, run thefollowing `scsi_id` command on the concerned device:
	
	`scsi_id -g -u -s /block/sda`
	
	The output may look like:
	`SIBM-ESXSMAW3073NC_FDAR9P66067WJ`
	
	As the WWID is unique and constant for each device, this would work fine across reboots as well.

6	Create a device-specific section at the end of the file for the storage system as follows. These settings apply only to NetApp storage systems. **`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 directio`** **`failback immediate`** **`}`** **`}`**
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.

```
defaults {
        user_friendly_names yes
```

```
}
blacklist {
    wwid     SIBM-ESXSMAT3073NC_FAAR9P5A0ADG6
    wwid     SIBM-ESXSMAT3073NC_FAAR9P5A0ADKP
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
          }
devices {
    device {
    vendor                  "NETAPP"
    product                 "LUN"
    path_grouping_policy    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            directio
    failback                immediate
           }

}
```

For RHEL 4 series Update 3 and later:

Step	Action
1	Open /etc/multipath.conf with a text editor.
2	Comment out the following line: `devnode_blacklist {devnode "*" }`
3	Delete the comment characters from the remaining `devnode_blacklist` command.
4	The default command excludes IDE hard drives /dev/hd*x* devices from /dev/hda through /dev/hdz. If your local drives are IDE drives in this range, no further changes to the blacklist section are needed.

Step	Action

5	If your local drives are SCSI drives, add the wwid within the `blacklist` command to exclude the specific /dev/sd*x* devices. Be sure to blacklist only your local drives. Each path to an iSCSI LUN is also a /dev/sd*x* device that should be configured for multipathing.
	
	To obtain the WWID of a device, run thefollowing `scsi_id` command on the concerned device:
	
	`scsi_id -g -u -s /block/sda`
	
	The output may look like:
	`SIBM-ESXSMAW3073NC_FDAR9P66067WJ`
	
	As the WWID is unique and constant for each device, this would work fine across reboots as well.

6	Create a device-specific section at the end of the file for the storage system as follows. These settings apply only to NetApp storage systems. **`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`** **`}`** **`}`**
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.

```
defaults {
        user_friendly_names no
```

```
}
blacklist {
    wwid     SIBM-ESXSMAT3073NC_FAAR9P5A0ADG6
    wwid     SIBM-ESXSMAT3073NC_FAAR9P5A0ADKP
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
    devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
          }
devices
{
 device {
  vendor                "NETAPP"
  product               "LUN"
  path_grouping_policy  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
        }
}
```

Starting the multipath service

Note

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

Once you have edited the `/etc/multipath.conf` file, start the multipath service on the Linux host. To start the multipath service, enter the following command on the Linux host console:

Step	Action
1	Start the multipath daemon by entering the following command: **`/etc/init.d/multipathd start`**

Configuring the multipath service to start automatically

Note

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

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

Step	Action

1	Add the multipath service to the boot sequence by entering the following commands on the Linux console:
	`chkconfig --add multipathd`
	`chkconfig multipathd on`

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:

In RHEL 5 series:

iqn.2005-03.com.RedHat:*RandomNumber*

In RHEL 4 series Update 3 and later:

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

Note

The initiator node name is generated the first time the iSCSI service on the host is started. However, it is important that you have a LUN mapped to the Linux host before you start the iSCSI service. If you have never used the iSCSI service on the Linux host, create the initiator node name manually.

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

For RHEL 5 series:

`iqn.2005-03.com.RedHat:` *host-name*

For RHEL 4 series Update 3 and later:

`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.

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

For RHEL 5 series:

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

Step	Action	
2	**If...**	**Then...**
	If a default initiator name has not been generated	**1.** Replace the text in the host's /etc/iscsi/initiatorname.iscsi file with the following line: `InitiatorName=iqn.2005-03.com.RedHat:`***`RandomNumber`*** The recommended value for *RandomNumber* is the host name of the Linux host. **2.** Record the node name and proceed to "Editing the host's configuration file" on page 24.
	You want to use the default initiator node name	Record the node name and proceed to "Editing the host's configuration file" on page 24.
	You want to change the initiator node name	**1.** Modify the *RandomNumber* part of the initiator node name. The following line shows an example node name. iqn.2005-03.com.RedHat:linux-host1 **2.** Record the node name and proceed to "Editing the host's configuration file" on page 24. **Note** The *RandomNumber* 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.

For RHEL 4 series Update 3 and later:

Step	Action	
1	With a text editor, open the host's /etc/initiatorname.iscsi file.	
2	**If...**	**Then...**
	If a default initiator name has not been generated	**1.** Replace the text in the host's /etc/initiatorname.iscsi file with the following line: **`iqn.1987-05.com.cisco:`*`RandomNumber`*** The recommended value for *RandomNumber* is the host name of the Linux host. **2.** Record the node name and proceed to "Editing the host's configuration file" on page 24.
	You want to use the default initiator node name	Record the node name and proceed to "Editing the host's configuration file" on page 24.
	You want to change the initiator node name	**1.** Modify the *RandomNumber* part of the initiator node name. The following line shows an example node name. iqn.1987-05.com.cisco:linux-host1 **2.** Record the node name and proceed to "Editing the host's configuration file" on page 24. **Note** The *RandomNumber* 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.

Multipathing support:

For RHEL 5 series, complete the following steps:

Step	Action
1	With a text editor, open the host's /etc/iscsi/iscsid.conf file.
2	If you are using dm-multipath support: ◆ Set the value of `node.session.timeo.replacement_timeout` to `5`

For RHEL 4 series Update 3 and later, complete the following steps:

Step	Action
1	With a text editor, open the host's /etc/iscsi.conf file.
2	If you are using dm-multipath support: ◆ Remove the comment of the `ConnFailTimeout` line in the Session Timeout Settings section and set the value to `ConnFailTimeout=5`.

CHAP authentication:

For RHEL 5 series, complete the following steps:

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

Step	Action

2	If you want to use CHAP, add CHAP user names and passwords, as shown below.
	
	To enable CHAP authentication set node.session.auth.authmethod to CHAP. The default is None.
	
	`node.session.auth.authmethod = CHAP`
	
	To set a CHAP username and password for initiator authentication by the target(s), remove the comment from the following lines:
	
	`node.session.auth.username = username`
	
	`node.session.auth.password = password`
	
	To set a CHAP username and password for target(s) authentication by the initiator, remove the comment from the following lines:
	
	`node.session.auth.username_in = username_in`
	
	`node.session.auth.password_in = password_in`
	
	Make sure you use the same user names and passwords when you set up CHAP on the storage system with the iscsi security command.

For RHEL 4 series Update 3 and later, complete the following steps:

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

Step	Action

2	If you want to use 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:
	
	`OutgoingUsername=linuxhostout`
	
	`OutgoingPassword=aow9857fl`
	
	`IncomingUsername=linuxhostin`
	
	`IncomingPassword=e09fj30`
	
	Make sure you use the same user names and passwords when you set up CHAP on the storage system with the `iscsi security` command.
	
	Use the `OutgoingUsername` and `OutgoingPassword` for the storage system's inbound user name and password (*inname* and *inpassword*).
	
	If you are using bidirectional authentication, use `IncomingUsername` and `IncomingPassword` for the storage system's outbound user name and password (*outname* and *outpassword*) for incoming settings.

Note

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

Target Discovery:

For RHEL 5 series complete the following steps:

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

2	For successful session discovery, you need to enable the discovery CHAP authentication.
	`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`

3	To discover target use the following command: `iscsiadm -m discovery -t st -p <ip_address>` where <ip_address> is the target address
4	To setup the session startup option of iSCSI as manual or automatic, use the following command: `iscsiadm -m node -T <target_iqn> -p <ip_address>:3260 --op update -n node.startup -v <mode>` where <target_iqn> is the iqn number of the target, <ip_address> is the target address and <mode> is manual or automatic. Alternatively, You can setup the session startup by editing the /etc/iscsi/iscsid.conf file before target discovery. The default setting is node.startup=automatic

For RHEL 4 series Update 3 and later, complete the following steps:

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

Step	Action

2	Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you will use for iSCSI LUNs:
	`DiscoveryAddress=storage_system_IPaddress`
	storage_system_IPaddress 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:
	`DiscoveryAddress=192.168.10.100`
	`DiscoveryAddress=10.61.208.200`

Configuring the storage system for iSCSI 3

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 when creating igroups on the storage system.

Tasks for configuring the storage system

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

- Ensure that the iSCSI service is licensed and started.
- For RHEL 5 series, RHEL 4 series Update 2 and later versions without dm-multipath, restrict the host's access to a single iSCSI portal on the storage system.

Note

The Linux host running RHEL 5 series, RHEL 4 series Update 2 and later versions without dm-multipath, can access more than one iSCSI portal (typically an Ethernet interface) on the storage system. It will create multiple sessions to the storage system and access each LUN as more than one SCSI device for each path. This can lead to file system inconsistencies. See "Restricting access to a single iSCSI portal (RHEL 5 series, RHEL 4 series Update 2 and later)" for more information.

- 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`, and be sure at least one LUN is mapped as LUN 0.

Restricting access to a single iSCSI portal (RHEL 5 series, RHEL 4 series Update 2 and later)

The iSCSI initiator included with RHEL 5 series, RHEL 4 series Update 2 and later versions creates a separate iSCSI session to each iSCSI portal it discovers on the storage system. An iSCSI portal is a physical or logical network device on the storage system, such as an Ethernet interface or VLAN, that is enabled for iSCSI. A network portal has a unique IP address.

When the Linux host detects multiple iSCSI portals, it creates multiple iSCSI sessions. If the same LUNs appear in each session, the Linux host treats the duplicate LUNs as if they were unique devices. This new behavior is intended to support multipathing. Without multipathing, this new behavior can cause problems on the Linux host.

For example, if a storage system has two interfaces enabled for iSCSI, and a single LUN mapped to a host running RHEL 5 series, the host would create two sessions to the storage system and would treat the LUN as two separate SCSI devices.

See the following output of the `iscsiadm -m session -P 3 -r 2` command:

```
[root@199-119 ~]# iscsiadm -m session  -P 3 -r 2
Target: iqn.1992-08.com.netapp:sn.101183016
        Current Portal: 10.72.199.71:3260,1001
        Persistent Portal: 10.72.199.71:3260,1001
                **********
                Interface:
                **********
                Iface Name: default
                Iface Transport: tcp
                Iface Initiatorname: iqn.1994-
05.com.redhat:5e3e11e0104d
                Iface IPaddress: 10.72.199.119
                Iface HWaddress: default
                Iface Netdev: default
                SID: 2
                iSCSI Connection State: LOGGED IN
                iSCSI Session State: Unknown
                Internal iscsid Session State: NO CHANGE
                ************************
                Negotiated iSCSI params:
                ************************
                HeaderDigest: None
                DataDigest: None
                MaxRecvDataSegmentLength: 131072
                MaxXmitDataSegmentLength: 65536
                FirstBurstLength: 65536
                MaxBurstLength: 65536
                ImmediateData: Yes
                InitialR2T: No
                MaxOutstandingR2T: 1
                ************************
                Attached SCSI devices:
                ************************
                Host Number: 4  State: running
                scsi4 Channel 00 Id 0 Lun: 0
```

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

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

For RHEL 4 series Update 3 and later:

The following output from the `iscsi-ls -l` command shows this situation. Note that the values for TARGET NAME and LUN ID are the same, but the SESSION ID values are different. The same LUN shows up as both `/dev/sdc` and `/dev/sdd`.

```
[root@storagesystem]# iscsi-ls -l
*******************************************************************
SFNet iSCSI Driver Version ...4:0.1.11(12-Jan-2005)
*******************************************************************
TARGET NAME           : iqn.1992-08.com.netapp:sn.33604646
TARGET ALIAS          :
HOST ID               : 26
BUS ID                : 0
TARGET ID             : 0
TARGET ADDRESS        : 10.61.208.8:3260,2
SESSION STATUS        : ESTABLISHED AT Thu Nov 17 16:50:52 EST 2005
SESSION ID            : ISID 00023d00001f TSIH 505

DEVICE DETAILS:
---------------
LUN ID : 0
  Vendor: NETAPP Model: LUN Rev: 0.2
  Type:   Direct-Access                    ANSI SCSI revision: 04
  page83 type3: 60a980004f644374535a31524e495a77
  page80: 4f644374535a31524e495a770a
  Device: /dev/sdc
```

```
**************************************************************
TARGET NAME           : iqn.1992-08.com.netapp:sn.33604646
TARGET ALIAS          :
HOST ID               : 27
BUS ID                : 0
TARGET ID             : 0
TARGET ADDRESS        : 10.61.208.9:3260,3
SESSION STATUS        : ESTABLISHED AT Thu Nov 17 16:50:52 EST 2005
SESSION ID            : ISID 00023d00001f TSIH 902

DEVICE DETAILS:
---------------
LUN ID : 0
  Vendor: NETAPP Model: LUN Rev: 0.2
  Type:   Direct-Access                    ANSI SCSI revision: 04
  page83 type3: 60a980004f644374535a31524e495a77
  page80: 4f644374535a31524e495a770a
  Device: /dev/sdd
**************************************************************
```

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

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

The details of each option are as follows.

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

Create vFiler units. If you have a MultiStore® software licence on your storage system, you can create several virtual storage systems, called vFiler units. Each vFiler unit has its own iSCSI target. Be sure the vFiler unit used by the Linux host has only one iSCSI-enabled interface. This option also enables you to use other interfaces for other iSCSI hosts. For more information, see the *Data ONTAP MultiStore Management Guide*.

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

To setup the session startup option of iSCSI as manual or automatic, use the following command:

```
iscsiadm -m node -T <target_iqn> -p <ip_address>:3260 --op update -
n node.startup -v <mode>
```

where `<target_iqn>` is the iqn number of the target, `<ip_address>` is the target address and `<mode>` is manual or automatic.

For detailed configuration steps

For detailed storage system configuration 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 4

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.

Accessing multipath LUNs

If you are using the dm-multipath support with RHEL 5 series and RHEL 4 series Update 3 and later, you access LUNs differently than if you are not using multipathing. This chapter has two main sections, one for accessing LUNs with dm-multipath, and one for accessing LUNs without dm-multipath.

For detailed information

For detailed information, see the following sections:

- "Starting the iSCSI service on the host" on page 36
- "Configuring the iSCSI service to start automatically" on page 37
- "Accessing LUNs using dm-multipath" on page 38
- "Accessing LUNs without dm-multipath" on page 48

Starting the iSCSI service on the host

Before you start iSCSI

Be sure that you have at least one LUN mapped to the Linux host as LUN 0 before starting the iSCSI service.

Starting iSCSI

To start the iSCSI service, complete the following step.

If you are using dm-multipath, be sure you have configured multipathing before starting iSCSI.

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

Configuring the iSCSI service to start automatically

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

Step	Action	
1	Verify the status of the iSCSI service by entering the following command: **`chkconfig --list iscsi`**	
2	**If...**	**Then...**
	The `chkconfig` command indicates the iSCSI service is enabled, as follows: `iscsi on`	No action needed.
	The `chkconfig` command indicates that the iscsi service is not enabled as follows: `iscsi off`	Enable the iSCSI service by entering the following command: **`chkconfig iscsi on`**

Accessing LUNs using dm-multipath

RHEL 5 series, RHEL 4 series Update 3 and later

Support for dm-multipath is currently on RHEL 5 series, RHEL 4 series Update 3 and later. However, new configurations might be qualified between support kit / 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

Understanding device persistence

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

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

Understanding automounting

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

For detailed information

For detailed information, see the following sections:

- "Viewing a list of LUNs" on page 39
- "Viewing the dm-multipath configuration and multipath devices" on page 43
- "Creating and mounting a file system on a multipath device" on page 44
- "Accessing LUNs as raw devices" on page 46
- "Tuning the dm-multipath configuration" on page 47

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/sde        host21  iSCSI     7g (7516192768)      GOOD
ss1:    /vol/vol1/linux_1 /dev/sdd        host20  iSCSI     7g (7516192768)      GOOD
ss2:    /vol/vol1/linux_2 /dev/sdc        host19  iSCSI     5g (5368709120)      GOOD
ss2:    /vol/vol1/linux_2 /dev/sdb        host18  iSCSI     5g (5368709120)      GOOD
```

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

Using iscsi-ls

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

Note

In RHEL 5 series, the list option iscsi-ls -l is not available.

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/sde device on the host.

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

DEVICE DETAILS:
---------------
LUN ID : 0
  Vendor: NETAPP Model: LUN Rev: 0.2
  Type:   Direct-Access                    ANSI SCSI revision: 04
  page83 type3: 60a9800043346536534a344a794e6d59
  page80: 43346536534a344a794e6d590a
  Device: /dev/sde
```

Using iscsiadm command

The `iscsiadm` command lists the LUNs that are mapped to the Linux host. The following example shows the output of the `iscsiadm -m session --sid=0 -P 3 -r 2` command:

```
[root@199-119 ~]# iscsiadm -m session  -P 3 -r 2
Target: iqn.1992-08.com.netapp:sn.101183016
        Current Portal: 10.72.199.71:3260,1001
        Persistent Portal: 10.72.199.71:3260,1001
                **********
                Interface:
                **********
                Iface Name: default
                Iface Transport: tcp
                Iface Initiatorname: iqn.1994-05.com.redhat:5e3e11e0104d
                Iface IPaddress: 10.72.199.119
                Iface HWaddress: default
                Iface Netdev: default
                SID: 2
                iSCSI Connection State: LOGGED IN
                iSCSI Session State: Unknown
                Internal iscsid Session State: NO CHANGE
                ************************
                Negotiated iSCSI params:
                ************************
```

```
HeaderDigest: None
DataDigest: None
MaxRecvDataSegmentLength: 131072
MaxXmitDataSegmentLength: 65536
FirstBurstLength: 65536
MaxBurstLength: 65536
ImmediateData: Yes
InitialR2T: No
MaxOutstandingR2T: 1
************************
Attached SCSI devices:
************************
Host Number: 4  State: running
scsi4 Channel 00 Id 0 Lun: 0
        Attached scsi disk sdc      State: running
scsi4 Channel 00 Id 0 Lun: 1
        Attached scsi disk sde      State: running
scsi4 Channel 00 Id 0 Lun: 2
        Attached scsi disk sdg      State: running
scsi4 Channel 00 Id 0 Lun: 3
        Attached scsi disk sdi      State: running
scsi4 Channel 00 Id 0 Lun: 4
        Attached scsi disk sdk      State: running
scsi4 Channel 00 Id 0 Lun: 5
        Attached scsi disk sdm      State: running
scsi4 Channel 00 Id 0 Lun: 6
        Attached scsi disk sdp      State: running
scsi4 Channel 00 Id 0 Lun: 7
        Attached scsi disk sdq      State: running
```

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

Discovering new LUNs

In RHEL 5 series:

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

Use the following command to get the list of all the current sessions

`iscsiadm -m session`

Use the following command to rescan a specific session

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

where N is the specific session ID.

Use the following command to rescan all the sessions

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

You can use the `sanlun` command in RHEL 5 series to verify that the new LUNs were discovered.

Once discovered, the LUNs are automatically added to the dm-multipath configuration.

In RHEL 4 series Update 3 and later:

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

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

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

Once discovered, the LUNs are automatically added to the dm-multipath configuration.

Viewing the dm-multipath configuration and multipath devices

Viewing the configuration

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

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

`multipath -v3 -d -ll`

Viewing the multipath devices

To view a list of multipath devices, including which /dev/sd*x* devices are used, enter the following command.

`multipath -d -ll`

Example: The following output in RHEL 5 series , shows two LUNs, with two paths each.

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

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

Note that the mpath*x* values are persistent across reboots, but the /dev/sd*x* devices are not. After a reboot, or a restart of the iscsi service, you might find that different /dev/sd*x* devices make up a given mpath*x* device. The mpath*x* device, however, will always correspond to the same LUN.

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

Example: The following example in RHEL 5 series, shows two devices, mpath0 and mpath1.

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

Creating and mounting a file system on a multipath device

Overview

You can optionally create a file system on the multipath device that represents a LUN. If you are using the LUN as a raw device, skip this section.

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

Creating the file system

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

Enter the following command on the Linux host console:

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

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

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

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

```
mkfs -t ext3 /dev/mapper/mpath0
```

Mounting the file system

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

To mount the file system, complete the following steps:

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

Step	Action
3	Add an entry to /etc/fstab as follows, separating each entry with a Tab (except the final two zeros, which are separated by a space): *device* *mount_point* *type* `_netdev 0 0` *device* is the name of the device in the /dev/mpath/ directory *mount_point* is the mount point you created for the file system *type* is the file system type, such as `ext2 or ext3` *_netdev* is used for any network dependent devices like iscsi **Example:** `/dev/mapper/mpath0 /mnt/ss1/lun_0 ext3 _netdev 0 0`
4	Enter **`mount mountpoint`** to mount the file system. **Example:** `mount /mnt/ss1/lun_0`
5	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

Overview

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

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

For the Linux 2.6 kernel used by RHEL 5 series, RHEL 4 series and later, do not use the raw interface. Instead open the device using the open call with the O_DIRECT flag.

Tuning the dm-multipath configuration

Tuning multipath parameters

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

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

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

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

Accessing LUNs without dm-multipath

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. iSCSI initiator version 3.4.x automatically generates persistent symbolic links for the LUNs. With initiator versions 3.6.x and 4.0.x, those symbolic links are no longer created.

For initiator versions 3.6.x and 4.0.x, you can write a file system label to the LUN and mount the file system by that label. The readme file for the iSCSI initiator (usr/share/doc/iscsi-initiator-utils-<version>/README) discusses other methods of creating device persistence.

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 both `scsi_id` command are the same, which shows that the iscsi devices are two paths to the same LUN.

Understanding automounting

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

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

```
#device            mount             FS    mount   backup     fsck
```

```
#to mount          point              type  options frequency  pass
LABEL=lun0_label  /mnt/filer0_lun0  ext3  _netdev    0          0
```

The SourceForge iSCSI initiator automount script `/sbin/iscsi-mountall` is not included with the iSCSI initiator in RHEL5 series, RHEL 4 series, RHEL 3 Update 4 and later. Use the automounting method described above instead.

Tasks for accessing LUNs

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

- If you are creating a file system on the LUN:
  - Create one or more partitions of the LUN. A file system needs a partition, even if the one partition uses all of the space in the LUN.
  - Create a file system on the partition.
  - Create a label for the file system and add an entry for the file system's label in /etc/fstab.
- Configure the LUN as a raw device.

For detailed information

For detailed information, see the following sections:

- "Creating a partition" on page 51
- "Creating a file system" on page 53
- "Labeling and mounting the file system" on page 54
- "Accessing LUNs as raw devices" on page 56
- "Viewing LUN information" on page 57

Creating a partition

About partitions

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

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

Before you begin

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

`sanlun lun show all`

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

Using fdisk to create a partition

To create a partition of a LUN:

Step	Action
1	Enter **`/sbin/fdisk`** ***`device`*** *device* 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`**.
5	Press Enter twice to accept the default cylinder values.

Step	Action

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`
	
	`  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

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 56.

Using mkfs to create a file system

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

Step	Action
1	Enter **`mkfs -t ext3 device`** Where *device* is the partition. For example /dev/sda1 is the first partition on /dev/sda.

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 48 for more information.

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

Using e2label to create a file system label

To create a file system label, run the `e2label` command as follows:

Step	Action
1	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`

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/filer0_lun0`
2	Open the file /etc/fstab using a text editor.

Step	Action

3	Add an entry to /etc/fstab as follows, separating each entry with a Tab (except the final two zeros, which are separated by a space):
	`LABEL=`*label* *mount_point* `auto _netdev 0 0`
	label is the file system label you created above
	mount_point is the mount point you created for the file system
	Example:
	`LABEL=lun0_label /mnt/filer0_lun0 auto _netdev 0 0`

4	Save your changes to /etc/fstab and exit the editor.
5	Enter **`mount LABEL=`*`label`*** to mount the file system. The file system will be automatically mounted whenever the Linux host is rebooted.
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

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

For the Linux 2.6 kernel used by RHEL 5 series, RHEL 4 series and later, do not use the raw interface. Instead open the device using the open call with the O_DIRECT flag.

Viewing LUN information

Viewing information about LUNs

You can view information about LUNs in the following ways:

- The `iscsiadm` command displays the storage system and LUN details.
- The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

Note

The iscsi-ls command is not available in RHEL 5 series.

- 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.

Using the iscsiadm command

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

Step	Action

1	Enter the following command at the host's shell prompt:
	`iscsiadm -m session -P 3 -r 2`
	Result: The console lists the available stroage systems and LUNs for a session with a specific session ID, as shown by the following example.
	`[root@199-119 ~]# iscsiadm -m session  -P 3 -r 2`
	`Target: iqn.1992-08.com.netapp:sn.101183016`
	`Current Portal: 10.72.199.71:3260,1001`
	`Persistent Portal: 10.72.199.71:3260,1001`
	`**********`
	`Interface:`
	`**********`
	`Iface Name: default`
	`Iface Transport: tcp`
	`Iface Initiatorname: iqn.1994-05.com.redhat:5e3e11e0104d`
	`Iface IPaddress: 10.72.199.119`
	`Iface HWaddress: default`
	`Iface Netdev: default`
	`SID: 2`
	`iSCSI Connection State: LOGGED IN`
	`iSCSI Session State: Unknown`
	`Internal iscsid Session State: NO CHANGE`
	`************************`
	`Negotiated iSCSI params:`
	`************************`
	`HeaderDigest: None`
	`DataDigest: None`
	`MaxRecvDataSegmentLength: 131072`
	`MaxXmitDataSegmentLength: 65536`
	`FirstBurstLength: 65536`
	`MaxBurstLength: 65536`
	`ImmediateData: Yes`
	`InitialR2T: No`
	`MaxOutstandingR2T: 1`
	`************************`
	`Attached SCSI devices:`
	`************************`

Step	Action

	`Host Number: 4  State: running`
	`scsi4 Channel 00 Id 0 Lun: 0`
	`  Attached scsi disk sdc  State: running`
	`scsi4 Channel 00 Id 0 Lun: 1`
	`  Attached scsi disk sde  State: running`
	`scsi4 Channel 00 Id 0 Lun: 2`
	`  Attached scsi disk sdg  State: running`
	`scsi4 Channel 00 Id 0 Lun: 3`
	`  Attached scsi disk sdi  State: running`
	`scsi4 Channel 00 Id 0 Lun: 4`
	`  Attached scsi disk sdk  State: running`
	`scsi4 Channel 00 Id 0 Lun: 5`
	`  Attached scsi disk sdm  State: running`
	`scsi4 Channel 00 Id 0 Lun: 6`
	`  Attached scsi disk sdp  State: running`
	`scsi4 Channel 00 Id 0 Lun: 7`
	`  Attached scsi disk sdq  State: running`
	
	The example above lists the available stroage systems and LUNs for a session with a specific session ID. To view the details of all the sessions, use `iscsiadm -m session -P 3` command.

Using the iscsi-ls command

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

Note

The iscsi-ls command is not available in RHEL 5 series.

Step	Action
1	Enter the following command at the host's shell prompt: **/sbin/iscsi-ls -l** **Result:** The console lists available storage systems and LUNs, as shown by the following example: [root@host /]# /sbin/iscsi-ls -l ********** SFNet iSCSI Driver Version ... 3.6.2 (27-Sep-2004) ********** TARGET NAME : iqn.1992-08.com.netapp:sn.33604646 TARGET ALIAS : HOST NO : 0 BUS NO : 0 TARGET ID : 0 TARGET ADDRESS : 10.60.128.100:3260 SESSION STATUS : ESTABLISHED AT Mon Jan 3 10:05:14 2005 NO. OF PORTALS : 1 PORTAL ADDRESS 1 : 10.60.128.100:3260,1 SESSION ID : ISID 00023d000001 TSID 103 DEVICE DETAILS : ---------- LUN ID : 0 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 04 page83 type3: 60a980004f6443745359763759367733 page83 type1:4e45544150502020204c554e204f6443745359763759367733000000000000000 page80: 4f6443745359763759367733 Device: /dev/sdb LUN ID : 1 Vendor: NETAPP Model: LUN Rev: 0.2 Type: Direct-Access ANSI SCSI revision: 04 page83 type3: 60a980004f644374535976426253674b page83 type1:4e45544150502020204c554e204f644374535976426253674b00000000000000 page80: 4f644374535976426253674b Device: /dev/sdc **********

Using the sanlun command

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

The sanlun program is installed from the iSCSI Red Hat Enterprise Linux Host Utilities.

Step	Action

1	Enter the following command at the host's shell prompt:
	`/opt/netapp/santools/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 lines show sample `sanlun` command output:
	`[root@host /]# /opt/netapp/santools/sanlun lun show all`
	`filer: lun-pathname device filename adapter protocol lun size lun state`
	`filer1: /vol/vol_1/lun_1 /dev/sdb1 <unknown>0 iSCSI 4g (4294967296) GOOD`
	`filer1: /vol/vol_1/lun_2 /dev/sdc1 <unknown>0 iSCSI 4g (4294967296) GOOD`

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.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system according to the procedure described in the Data ONTAP *Block Access Management Guide*.
Initiator login	Verify that the initiator is logged in to the storage system by entering the `iscsi show initiator` 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 configuration file" on page 24.

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 `igroup show` 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: `lun show -m` displays all LUNs and the igroups to which they are mapped. `lun show -g` *`igroup-name`* displays the LUNs mapped to a specific igroup.
System requirements	Verify that the components of your configuration are supported. Verify that you have the correct host Operating System (OS) service pack level, initiator version, Data ONTAP version, and other system requirements. You can check the most up-to-date system requirements in the NetApp 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.
Firewall settings	Verify that the iSCSI port(3260) is open in the firewall rule.

Note

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