Linux® Host Utilities 5.2
Installation and Setup Guide

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

Part number: 215-05185_A0
March 2010

Contents

Copyright information **9**
Trademark information **11**
About this guide **13**
Audience 13
Where to enter commands 13
Keyboard and formatting conventions 14
Special messages 15
How to send your comments 15
Changes to this document: March 2010 ..... **17**
Introduction to the Host Utilities **19**
Supported Linux versions 19
What the Linux Host Utilities contain 20
Protocols and configurations supported by the Linux Host Utilities 20
Notes about the supported protocols 20
The FC protocol 21
The FCoE protocol 21
The iSCSI protocol 22
Supported Linux and Data ONTAP features 22
SAN booting 22
Support for Linux Device Mapper Multipathing 23
What Device Mapper Multipathing does 23
Linux configurations that support ALUA 23
Device persistence 24
Support for Veritas Dynamic Multipathing 24
Where to find more information about Linux 24
HBA installation **27**
Host and storage system communication using FC 27
Verification of the HBA model 28
About uninstalling external HBA drivers 28
Download of the external HBA driver software for FC 28
Installation of the external HBA driver software for FC 28
Installation of the HBA management software 29

Setting the QLogic HBA parameters for FC 29
Setting the QLogic bundled driver parameters for FC 30
Setting the QLogic external driver parameters for FC 32
Setting the QLogic HBA parameters for hardware iSCSI 33
Emulex HBA parameters for FC 33
Driver settings for Red Hat Enterprise Linux 34
Driver settings for SUSE Linux Enterprise Server 34
Loading a bundled or an external HBA driver on a Linux host using FC 35
Loading a bundled HBA driver on a Linux host using hardware iSCSI 36
Verifying the QLogic HBA settings for FC 36
Emulex HBA settings for FC 37
Verifying the QLogic HBA settings for hardware iSCSI 37
Obtaining WWPNs of HBAs for FC 38
Obtaining the WWPN of QLogic HBAs using FC 39
Obtaining the WWPN for Emulex HBAs using FC 39
Adding targets to the host using hardware iSCSI 39
Obtaining the initiator node name of QLogic HBAs for hardware iSCSI 40
Unloading the bundled or external HBA driver for FC 40
Unloading the bundled HBA driver for hardware iSCSI 41
The Host Utilities installation 43
Downloading the Host Utilities software 43
Installing the Host Utilities software 43
Configuration of the software iSCSI initiator 45
Verifying the required iSCSI package on your Linux host 45
Why you need the iSCSI initiator node name on the Linux host 46
Node name rules 46
Recording or editing the initiator node name on your Linux host 46
Configuration of software iSCSI on Red Hat Enterprise Linux 48
Configuration of the iSCSI initiator on Red Hat Enterprise Linux 5 series 49
Configuration of the iSCSI initiator on Red Hat Enterprise Linux 4 series 51
Configuration of software iSCSI on SUSE Linux Enterprise Server 53
Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 10 series and 11 53

Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 9 series 58
Configuring the iSCSI service to start automatically 60
Storage system configuration 61
Configuring the storage system for FC 61
Configuring the storage system for hardware and software iSCSI 62
DM-Multipath configuration 65
Verifying the required multipathing packages 65
Editing the DM-Multipath configuration file 66
Starting DM-Multipath 71
Configuring DM-Multipath to start automatically while booting 71
Verifying the DM-Multipath configuration 72
Stopping DM-Multipath 73
The Veritas Array Support Library and Array Policy Module 75
What the ASL is 75
ASL array type 75
What the APM is 76
Installing the downloaded ASL and APM software 76
Removing the ASL and APM 77
Information about ASL error messages 78
Methods to access LUNs 81
Accessing LUNs with DM-Multipath 81
FC: Discovering new LUNs 82
FC: Viewing the list of LUNs 82
FC: Displaying multipath information 82
FC: Removing the LUN 83
Software iSCSI: Discovering new LUNs 84
Software iSCSI: Viewing the list of LUNs 87
Software iSCSI: Displaying multipath information 92
Software iSCSI: Removing the LUN 92
Hardware iSCSI: Discovering new LUNs 93
Hardware iSCSI: Viewing the list of LUNs 93
Hardware iSCSI: Displaying multipath information 93
Hardware iSCSI: Removing the LUN 94
Accessing LUNs without multipathing software 94
Persistent identification of LUNs without multipathing software 95

FC: Viewing information about LUNs 95
Software iSCSI: Viewing information about LUNs 95
Software iSCSI: Automatic mounting of the file system 101
Accessing LUNs with Veritas DMP 103
HBA and driver parameter settings on Veritas Storage Foundation 103
Accessing LUNs using Veritas DMP 103
LUN access using Veritas DMP with FC 104
Veritas: Discovering new LUNs with FC 104
Veritas: Viewing the list of LUNs with FC 105
Veritas: Displaying multipath information with FC 105
Veritas: Removing the LUN with FC 106
LUN access using Veritas DMP with software iSCSI 107
Veritas: Viewing the list of LUNs with software iSCSI 107
Veritas: Discovering new LUNs with iSCSI 114
Veritas: Displaying multipath information with software iSCSI 116
Veritas: Removing the LUN with software iSCSI 117
Veritas DMP restore daemon values for Red Hat Enterprise Linux 118
Veritas DMP restore daemon values for SUSE Linux Enterprise Server 118
Displaying available paths using VxVM 118
Obtaining multipathing information by using the sanlun command 120
FC: SAN boot configuration 121
FC: Configuring SAN boot on Red Hat Enterprise Linux 121
FC: Installing the Red Hat Enterprise Linux operating system 121
FC: Installing and configuring the external or bundled HBA driver on Red Hat Enterprise Linux 122
FC: Configuring SAN boot on a SUSE Linux Enterprise Server 124
FC: Installing the SUSE Linux Enterprise Server operating system 125
FC: Installing and configuring the external or bundled HBA driver on SUSE Linux Enterprise Server 125
FC: Configuring root partition with DM-Multipath on SUSE Linux Enterprise Server 127
iSCSI: SAN boot configuration 129
Software iSCSI: Configuring SAN boot on Red Hat Enterprise Linux 129
Software iSCSI: Configuring SAN boot on SUSE Linux Enterprise Server 131
Software iSCSI: Configuring the root partition on multipath on the SUSE Linux Enterprise Server 133

Hardware iSCSI: Configuring SAN boot on Red Hat Enterprise Linux 134
Veritas Storage Foundation: SAN boot configuration 135
Veritas Storage Foundation: Configuring SAN boot on Red Hat Enterprise Linux 135
Veritas Storage Foundation: Configuring SAN boot on SUSE Linux Enterprise Server 137
Troubleshooting 139
FC troubleshooting 139
FC: Unloading of the QLogic module fails on a Linux host 139
FC: Linux hosts with QLogic initiators fail to detect LUNs 140
FC: multipath -l displays an undef LUN status 140
FC: LUNs not recognized by DM-Multipath 141
The SAN booted root on a DM-Multipath host freezes during FC path faults 141
FC: Poor performance due to misalignment of partitioned LUNs 141
Error message from the sanlun fcp show adapter command 142
LUNs unavailable after storage system failover or giveback 142
sanlun commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts .. 142
SCSI rescan issues on SUSE Linux Enterprise Server hosts 143
FC: LUNs not visible on the Linux host 143
iSCSI troubleshooting 144
OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 10 series and 9 series 144
iSCSI: multipath -l displays an undef LUN status 145
Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 11 145
Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 10 series 145
Automatically mounting and unmounting the file system on multipath device do not work on SUSE Linux Enterprise Server 9 series 146
Multipathd occasionally fails to update the path status of DM-Multipath devices 147
iSCSI: LUNs not recognized by DM-Multipath 147
Storage system console displays continuous ISID Rule messages on SUSE Linux Enterprise Server 9 series 147
iSCSI: Poor performance due to misalignment of partitioned LUNs 148

iSCSI: LUNs not visible on the Linux host 148
Veritas Storage Foundation troubleshooting 150
Error while stopping the fencing driver 150
Secondary paths identified incorrectly 151
Enclosure-based naming not reflected on NetApp storage 151
SAN boot installation to LUN configured on a storage controller fails 151
Encapsulation of the root disk fails 151
Index 153

Copyright information

Copyright © 1994–2010 NetApp, 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 NetApp material is subject to the following license and disclaimer:

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

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

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

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

Trademark information

NetApp; the NetApp logo; the Network Appliance logo; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW (NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, 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 StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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

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

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

About this guide

You can use your product more effectively when you understand this document's intended audience and the conventions that this document uses to present information.

This guide describes how to install the Linux Host Utilities 5.2 for Native OS and Veritas on a Linux host and how to set up the host to access the storage object described by a logical unit number (LUN). It also explains how to configure the storage system to work with the host.

This guide describes how to use both the iSCSI and Fibre Channel (FC) protocols. It also describes using both protocols to connect to the same LUN.

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

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

Next topics

Audience on page 13
Where to enter commands on page 13
Keyboard and formatting conventions on page 14
Special messages on page 15
How to send your comments on page 15

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability

Audience

This document is written with certain assumptions about your technical knowledge and experience.

This guide is for administrators of Red Hat Enterprise Linux, SUSE Linux Enterprise Server, Oracle Enterprise Linux, and NetApp storage systems.

Where to enter commands

You can use your product more effectively when you understand how this document uses command conventions to present information.

You can perform common administrator tasks in one or more of the following ways:

- You can enter commands either at the system console or from any client computer that can obtain access to the storage system using a Telnet or Secure Shell (SSH) session.

In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

- You can enter Windows, ESX, HP-UX, AIX, Linux, and Solaris commands at the applicable client console.

  In examples that illustrate command execution, the command syntax and output shown might differ from what you enter or see displayed, depending on your version of the operating system.

Keyboard and formatting conventions

You can use your product more effectively when you understand how this document uses keyboard and formatting conventions to present information.

Keyboard conventions

Convention	What it means
The NOW site	Refers to *NetApp On the Web at http://now.netapp.com/*.
Enter, *enter*	• Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. • Used to mean pressing one or more keys on the keyboard and then pressing the Enter key, or clicking in a field in a graphical interface and then typing information into the field.
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
type	Used to mean pressing one or more keys on the keyboard.

Formatting conventions

Convention	What it means
Italic font	• Words or characters that require special attention. • Placeholders for information that you must supply. For example, if the guide says to enter the `arp -d` *`hostname`* command, you enter the characters "arp -d" followed by the actual name of the host. • Book titles in cross-references.

Convention	What it means
`Monospaced` font	• Command names, option names, keywords, and daemon names. • Information displayed on the system console or other computer monitors. • Contents of files. • File, path, and directory names.
`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.

Special messages

This document might contain the following types of messages to alert you to conditions that you need to be aware of.

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

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

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to *doccomments@netapp.com*. To help us direct your comments to the correct division, include in the subject line the name of your product and the applicable operating system. For example, *FAS6070—Data ONTAP 7.3*, or *Host Utilities—Solaris*, or *Operations Manager 3.8—Windows*.

Changes to this document: March 2010

Changes have been made to this document since it was published for the Linux Host Utilities 5.1 release.

This document has been updated to add the following information:

- SUSE Linux Enterprise Server 11 support for software iSCSI.
- Veritas Storage Foundation 5.1 support for the software iSCSI initiator.
- Starting from this release the following diagnostic scripts are not bundled along with the toolkit
  - `linux_info`: Collects configuration information about the Linux host, including Veritas stack information
  - `filer_info`: Collects information about the storage system (for Data ONTAP versions 7.1 and earlier)
  - `controller_info`: Collects information about the storage system
  - `qlogic_info`: Collects information about QLogic switches installed in the network
  - `brocade_info`: Collects information about Brocade switches installed in the network
  - `cisco_info`: Collects information about Cisco switches installed in the network
  - `mcdata_info`: Collects information about McData switches installed in the network
  - `SHsupport.pm`: Support library for all the above info scripts
  - `Telnet.pm`: Support library for all the above info scripts
  - `FTP.pm`: Support library for all the above info scripts
  - `NaElement.pm`: Support library for all the above info scripts
  - `NaErrno.pm`: Support library for all the above info scripts
  - `NaServer.pm`: Support library for all the above info scripts

Introduction to the Host Utilities

The Host Utilities are a set of software programs and documentation that enable you to connect Linux hosts to storage systems running Data ONTAP software.

When you install the Host Utilities, the software installs the SAN Toolkit, which contains:

- The `sanlun` utility, which helps you manage LUNs and the host HBA.
- The support files which help in collecting and troubleshooting information.

The documentation that comes with the Host Utilities includes:

- This *Installation and Setup Guide*
- *Release Notes*
- *Quick Command Reference*

Next topics

Supported Linux versions on page 19
What the Linux Host Utilities contain on page 20
Protocols and configurations supported by the Linux Host Utilities on page 20
Supported Linux and Data ONTAP features on page 22
Where to find more information about Linux on page 24

Supported Linux versions

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

For information about the supported operating system version and Veritas Storage Foundation version, see the Interoperability Matrix.

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

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

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/

What the Linux Host Utilities contain

The Host Utilities bundle the software tools into a SAN toolkit.

The toolkit contains the following components:

- The `sanlun` utility that displays information about LUNs on the storage systems that are available to the Linux host
- The `install` and `uninstall` command scripts to add and remove the Linux Host Utilities components
- `san_version`: Displays the version of the Host Utilities

Note: For the detailed information about running the `sanlun` command, see the man page in the `/opt/netapp/man/man1` directory.

You can obtain the NOTICES.PDF in the `/opt/netapp/santools` directory.

Protocols and configurations supported by the Linux Host Utilities

The Linux Host Utilities provide support for Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI connections to the storage system using direct-attached, fabric-attached, and Ethernet network configurations.

Next topics

Notes about the supported protocols on page 20
The FC protocol on page 21
The FCoE protocol on page 21
The iSCSI protocol on page 22

Notes about the supported protocols

The FC, FCoE, and iSCSI protocols enable the host to access data on storage systems.

The storage systems are targets that have storage target devices called LUNs. The protocol enables the host to access the LUNs to store and retrieve data.

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

The FC protocol

The FC protocol requires one or more supported host bus adapters (HBAs) in the host. Each HBA port is an initiator that uses FC to access the LUNs on the storage system. The HBA port is identified by a worldwide port name (WWPN).

You need to make a note of the WWPN so that you can supply it when you create an initiator group (igroup). To enable the host to access the LUNs on the storage system using the FC protocol, you must create an igroup on the storage system and provide the WWPN as an identifier. Then, when you create the LUN, you map it to that igroup. This mapping enables the host to access that specific LUN.

The Linux Host Utilities support fabric-attached SAN and direct-attached configurations. The following configurations are supported:

- Fabric-attached SAN

  The Host Utilities support two variations of fabric-attached SANs:

  - A single-port FC connection from the HBA to the storage system through a single switch

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

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

  A single host with a direct FC connection from the host HBA to stand-alone or high-availability storage system configurations.

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

The FCoE protocol

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. FCoE is similar to traditional Fibre Channel (FC), as it maintains existing FC management and controls, but the hardware transport is a 10-gigabit Ethernet network.

Setting up an FCoE connection requires one or more supported converged network adapters (CNAs) in the host, connected to a supported data center bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

In general, you can configure and use FCoE connections just like traditional FC connections.

The iSCSI protocol

The iSCSI protocol is implemented on both the host and the storage system.

On the host, the iSCSI protocol is implemented over either the host's standard Ethernet interfaces or an HBA.

On the storage system, the iSCSI protocol can be implemented over the storage system's standard Ethernet interface using one of the following:

- A software driver that is integrated into Data ONTAP.
- (Data ONTAP 7.1 and later) An iSCSI target HBA or an iSCSI TCP/IP offload engine (TOE) adapter.

The connection between the initiator and target uses a standard TCP/IP network. The storage system listens for iSCSI connections on TCP port 3260.

Supported Linux and Data ONTAP features

The Host Utilities work with both Linux and Data ONTAP features.

Next topics

SAN booting on page 22
Support for Linux Device Mapper Multipathing on page 23
What Device Mapper Multipathing does on page 23
Linux configurations that support ALUA on page 23
Device persistence on page 24
Support for Veritas Dynamic Multipathing on page 24

SAN booting

SAN booting is the general term for booting a Linux host from a storage system LUN instead of an internal hard disk. SAN boot uses a SAN-attached disk, such as a LUN configured on a storage controller, as a boot device for a host.

SAN boot provides the following advantages:

- No maintenance and servicing costs for hard drives
  You can remove the hard drives from your servers and use the SAN for booting needs.
- Consolidated and centralized storage, because the host uses the SAN
- Lower cost
  The hardware and operating costs are lowered.
- Greater reliability
  Systems without the disks are less prone to failure.
- Quick server swaps

If a server fails, systems without the disks can be swapped.

- Better disaster recovery
  Site duplication is simplified.

For information about the supported operating system version and Veritas Storage Foundation version, see the Interoperability Matrix.

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/

Support for Linux Device Mapper Multipathing

Linux Device Mapper Multipathing (DM-Multipath) is supported on storage systems connected to hosts running on Red Hat Enterprise Linux and SUSE Linux Enterprise Server.

The required software module for DM-Multipath is included in the standard Red Hat Enterprise Linux and SUSE Linux Enterprise Server distributions. You should install additional user space packages as part of the setup process.

What Device Mapper Multipathing does

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

When you have multiple paths to a LUN, Linux creates a SCSI device for each path. For example, a single LUN might appear as`/dev/sdd and /dev/sdf` if there are two paths to the LUN. DM-Multipath creates a single device in `/dev/mapper/` for each LUN that represents all the paths.

For example, `/dev/mapper/360a98000433468525634447 17a513571` is the multipath device created on top of `/dev/sdd` and `/dev/sdf`. You should create file systems and mount the LUN using the device in `/dev/mapper/`.

Linux configurations that support ALUA

The Linux Host Utilities support asymmetric logical unit access (ALUA) on hosts running Red Hat Enterprise Linux or SUSE Linux Enterprise Server. ALUA is also known as Target Port Group Support (TPGS).

Veritas Dynamic Multipathing and DM-Multipath work with ALUA to determine primary paths and secondary or partner paths, which are used in the event of a failover.

ALUA is automatically enabled for Linux operating system. The following configurations support ALUA:

Host Utilities Version	Host requirements	Data ONTAP versions
Host Utilities 4.0 and later	• Red Hat Enterprise Linux 5 Update 1 and later • SUSE Linux Enterprise Server 10 SP1 and later	7.2.4 and later

Note: The Host Utilities do not support ALUA iSCSI.

Device persistence

The DM-Multipath code creates persistent devices for each LUN in the `/dev/mapper/` directory on the Linux host. You should use these entries for multipath LUNs.

You can create a file system directly on a multipath device in `/dev/mapper/`. You do not have to create a partition or label on the multipath device.

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

Support for Veritas Dynamic Multipathing

Veritas Dynamic Multipathing (VxDMP) is supported on storage systems connected to hosts running on Red Hat Enterprise Linux and SUSE Linux Enterprise Server. Veritas Volume Manager (VxVM) uses Dynamic Multipathing (DMP) to provide multipathing.

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

Where to find more information about Linux

For more information about host and storage system requirements, supported configurations, operating system, and troubleshooting, see the documents listed in the following table.

If you need more information about...	See this source...
Known issues, system requirements, and last minute updates	The latest Host Utilities *Release Notes*
The latest supported configurations	• The Interoperability Matrix • *System Configuration Guide*
Configuring the storage system	• The *Data ONTAP Software Setup Guide* • The *Data ONTAP Block Access Management Guide for iSCSI and FC*

If you need more information about...	**See this source...**
Supported SAN topologies	The *FC and iSCSI Configuration Guide* for your version of Data ONTAP software
Installing and configuring the HBA on the host	HBA vendor documentation
Managing SAN on the storage system	• The Command Reference • The *Data ONTAP Block Access Management Guide for iSCSI and FC* • The FilerView online Help • The Best Practices for Reliability: Ongoing Maintenance
Enabling ALUA	*Data ONTAP Block Access Management Guide for iSCSI and FC* **Note:** In particular, see the section “Enabling ALUA for a Fibre Channel igroup”.
The most current system requirements for the Linux host	The NetApp Interoperability Matrix
Best practices or configuration issues	The NetApp Knowledge Base
Information about how to configure and manage the FC and iSCSI service on the storage system	The following documentation for your release of Data ONTAP: • *Data ONTAP Block Access Management Guide for iSCSI and FC* • *Data ONTAP Release Notes*
Information about Oracle Cluster File System	*Oracle Cluster File System (OCFS2) User’s Guide and OCFS2* - Frequently Asked Questions
Information about Oracle Real Application Cluster (RAC)	Oracle RAC installation information from the Oracle Web site
Information about managing the iSCSI initiator	OS vendor documentation.
Information about installing and configuring SnapDrive for UNIX	The *SnapDrive for UNIX Installation and Administration Guide* for your version of SnapDrive

If you need more information about...	**See this source...**
Information about configuring multipath support on Red Hat Enterprise Linux	The `Multipath-usage.txt` file included with the device-mapper-multipath RPM for Red Hat Enterprise Linux 4 Update 3 and later in `/usr/share/doc/device-mapper-multipath-`*`version`*. *`version`* is the latest version number shipped with the Red Hat Enterprise Linux release.
Information about Red Hat Cluster and GFS	The Red Hat Web site

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/

FC and iSCSI Configuration Guide - http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

Data ONTAP Block Access Management Guide - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml/

FilerView® online Help - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml/

Best Practices for Reliability: Ongoing Maintenance - http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm

NetApp Knowledge Base - https://now.netapp.com/eservice/kbbrowse

Oracle Web site - http://oss.oracle.com

Knowledge base for Oracle RAC - http://www.netapp.com/library/tr/3423.pdf

Knowledge base for Oracle RAC - http://www.netapp.com/library/tr/3369.pdf

Red Hat Web site - http://www.redhat.com

Novell Web site - http://www.novell.com

HBA installation

You need to first install the HBAs and drivers and complete the host configuration after you have configured the storage system. You need to install the HBAs before you install the Host Utilities.

Next topics

Host and storage system communication using FC on page 27
Verification of the HBA model on page 28
About uninstalling external HBA drivers on page 28
Download of the external HBA driver software for FC on page 28
Installation of the external HBA driver software for FC on page 28
Installation of the HBA management software on page 29
Setting the QLogic HBA parameters for FC on page 29
Setting the QLogic HBA parameters for hardware iSCSI on page 33
Emulex HBA parameters for FC on page 33
Driver settings for Red Hat Enterprise Linux on page 34
Driver settings for SUSE Linux Enterprise Server on page 34
Loading a bundled or an external HBA driver on a Linux host using FC on page 35
Loading a bundled HBA driver on a Linux host using hardware iSCSI on page 36
Verifying the QLogic HBA settings for FC on page 36
Emulex HBA settings for FC on page 37
Verifying the QLogic HBA settings for hardware iSCSI on page 37
Obtaining WWPNs of HBAs for FC on page 38
Adding targets to the host using hardware iSCSI on page 39
Obtaining the initiator node name of QLogic HBAs for hardware iSCSI on page 40
Unloading the bundled or external HBA driver for FC on page 40
Unloading the bundled HBA driver for hardware iSCSI on page 41

Host and storage system communication using FC

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

After installing and configuring FC HBA, the Linux host starts using the FC protocol to access data on storage systems.

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

Verification of the HBA model

You should verify that the Linux host supports your HBA model.

If you install more than one HBA, ensure that all HBAs are the same brand and bus type (PCI-X or PCI-E). You should not install HBAs from different vendors on the same host because this might lead to inconsistent behavior and performance of the host.

> **Note:** For hardware iSCSI, enabling or initializing an additional network driver or an additional iSCSI driver for the second port while accessing iSCSI storage might result in inconsistent behavior of the host.

Ensure HBAs are installed on the Linux host. For installation instructions, see the HBA and host hardware vendor documentation.

For information about the supported HBA models, see the NetApp Interoperability Matrix.

Related information

NetApp Interoperability Matrix-http://now.netapp.com/NOW/products/interoperability

About uninstalling external HBA drivers

You might need to uninstall the QLogic or Emulex external HBA drivers on the Linux host. For more information about how to uninstall your external HBA driver, follow the instructions on the respective vendor's Web site or see the README included with the driver.

Download of the external HBA driver software for FC

You can download the external HBA driver software from the vendor's Web site. Before downloading, verify that you are using the appropriate version by checking the NetApp Interoperability Matrix.

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/

Installation of the external HBA driver software for FC

To install the driver software on the Linux host, you should follow the instructions on the vendor’s Web site or the README included with the driver.

Installation of the HBA management software

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

You should download and install the following software components depending on the HBA from the respective vendor Web site:

- QLogic HBA for FC
  - SANsurfer FC HBA Manager
  - SANsurfer CLI
- Emulex HBA for FC
  - Applications kit
- QLogic HBA for hardware iSCSI
  - SANsurfer iSCSI HBA Manager
  - SANsurfer iSCLI

Setting the QLogic HBA parameters for FC

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

Steps

1. Run the following command on the host:

   `/opt/QLogic_Corporation/SANsurferCLI/scli`

2. Set the following HBA parameters:
   - Port Down Retry Count to `30`
   - Link Down Timeout to `20`
3. Select **Commit Changes** for this HBA so that the changes are applied to the HBA port.
4. Repeat the settings for each HBA port that is listed.
5. Depending on your Linux version, configure the driver with the required values.

Next topics

Setting the QLogic bundled driver parameters for FC on page 30
Setting the QLogic external driver parameters for FC on page 32

Related references

Driver settings for Red Hat Enterprise Linux on page 34
Driver settings for SUSE Linux Enterprise Server on page 34

Setting the QLogic bundled driver parameters for FC

You should set the QLogic bundled driver parameters to ensure that the multipathing and storage system failover work correctly.

Step

1. Depending on your Linux version, configure the driver by performing the following step.

If you are using...	Then...
Red Hat Enterprise Linux 5 series	**a.** In the `/etc/modprobe.conf` file, append the following line: **`options qla2xxx qlport_down_retry=30`** **b.** Re-create the initrd by running the following command: **`mkinitrd -f /boot/initrd-"`uname -r`".img `uname - r``** **Note:** It is best to back up the current initrd before running the `mkinitrd` command. **c.** Reboot the host. **d.** Verify that the value of the qlport_down_retry parameter in the corresponding `qla2xxx` directory in `/sys/module/qla2xxx/parameters/` matches the values in the `modprobe.conf` file. Example: You can check the value of the qlport_down_retry parameter on a Red Hat Enterprise Linux host by running the following command: **`cat /sys/module/qla2xxx/parameters/qlport_down_retry`** **Note:** There are no qla2300/qla2400/qla2500 firmware modules in Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module.

If you are using...	Then...
Red Hat Enterprise Linux 4 series	**a.** In the `/etc/modprobe.conf` file, append the following line: **`options qla2xxx qlport_down_retry=30 ql2xsuspendcount=6`** **b.** Re-create the initrd by running the following command: **`mkinitrd -f /boot/initrd-"`uname - r`".img `uname -r` --with=qioctlmod`** **Note:** It is best to back up the current initrd before running the `mkinitrd` command. **c.** Reboot the host. **d.** Verify that the values of the qlport_down_retry and ql2xsuspendcount parameters in the corresponding `qla2xxx` directory in either `/sys/module/qla2xxx/` or `/sys/module/qla2xxx/parameters/` matches the values in the `modprobe.conf` file. Example: You can check the value of the ql2xsuspendcount parameter on a Red Hat Enterprise Linux host by running the following command: **`cat /sys/module/qla2xxx/ql2xsuspendcount`**
SUSE Linux Enterprise Server 10 series	**a.** In the `/etc/modprobe.conf.local` file, append the following line: **`options qla2xxx qlport_down_retry=30`** **b.** Re-create the initrd by running the following command: **`mkinitrd`** **Note:** It is best to back up the current initrd before running the `mkinitrd` command. **c.** Reboot the host. **d.** Verify that the value of the qlport_down_retry parameter in the corresponding `qla2xxx` directory in `/sys/module/qla2xxx/parameters/` matches the values in the `modprobe.conf.local` file. Example: You can check the value of the qlport_down_retry parameter on a SUSE Enterprise Linux Server host by running the following command: **`cat /sys/module/qla2xxx/parameters/qlport_down_retry`**
SUSE Linux Enterprise Server 11	No changes are required for QLogic driver parameters. **Note:** QLogic FC HBA Link Down Timeout should be set to `20` from the default value of `30`.

Setting the QLogic external driver parameters for FC

You can set the QLogic driver parameters for multipathing and storage system failover by using the `qlinstall` command.

Before you begin

You need to edit either the `/etc/modprobe.conf` file (Red Hat Enterprise Linux) or the `/etc/modprobe.conf.local` file (SUSE Linux Enterprise Server).

Steps

1. Depending on your HBA model, unload the external QLogic HBA driver by entering the following command:

   `./qlinstall -ul` *`driver module`*

   Example

   `./qlinstall -ul qla2xxx`

2. Append the `options qla2xxx` statement with the settings in the `modprobe.conf` file on a single line:

   `ql2xsuspendcount=6 MaxRetriesPerPath=10 MaxRetriesPerIo=10 ql2xfailover=0`

3. Save the file and exit.

4. Depending on your HBA model, load the external QLogic HBA driver by entering the following command:

   `./qlinstall -l` *`driver module`*

   Example

   `./qlinstall -l qla2xxx`

5. Depending on your HBA model to make the changes persistent, enter the following command:

   `./qlinstall -br -in` *`driver module`*

   Example

   `./qlinstall -br -in qla2xxx`

6. Verify the values for the external driver parameters in the corresponding `qla2xxx` directory in `sysfs`.

   The value of the external driver parameter can be found in `/sys/module/qla2xxx/` or `/sys/module/qla2xxx/parameters/` directory, depending on the driver version.

   Example
   You can check the value of ql2xsuspendcount on a Red Hat Enterprise Linux and a SUSE Linux Enterprise Server hosts by running the following command:

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

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

Setting the QLogic HBA parameters for hardware iSCSI

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

Steps

1. Run the following command on the host:

   ```
   /opt/QLogic_Corporation/SANsurferCLI/iscli
   ```

   Note: For more information about options, see the `iscli` utility help.

2. Set the following HBA parameters:
   - IP_ARP_Redirect=`on`
   - KeepAliveTO=`30`
3. Select **Commit Changes** for this HBA so that the changes are applied to the HBA ports.
4. Repeat the settings for each HBA port that is listed.
5. Return to the main menu and select **Quit**.

   Note: Alternatively, you can use the `iscli -n` command to set the HBA parameters directly.

   Example:

   - `iscli -n` *`port inst no`*`IP_ARP_Redirect on`
   - `iscli -n` *`port inst no`* `KeepAliveTO 30`

   `port inst no` is the port instance number that you can obtain by using the `iscli -i` command.

Emulex HBA parameters for FC

You do not need to set the parameters if you are using Emulex HBAs.

Driver settings for Red Hat Enterprise Linux

You should configure the driver settings to use the required values.

The default versus required settings for Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5 are as follows:

Component	**Parameter**	**Red Hat Enterprise Linux 4**		**Red Hat Enterprise Linux 5**	
		Default setting	**Required setting**	**Default setting**	**Required setting**
QLogic external driver for FC	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6
QLogic bundled driver for FC	ql2xsuspendcount	10	6	Not applicable	Not applicable
QLogic FC HBA	Link Down Timeout	30	20	30	20
QLogic iSCSI HBA	IP_ARP_Redirect	Not applicable	Not applicable	off	on

Note: For Emulex, no changes are required to the default settings.

Driver settings for SUSE Linux Enterprise Server

You can configure the driver settings to use the required values.

The default versus required settings for SUSE Linux Enterprise Server 9, SUSE Linux Enterprise Server 10, and SUSE Linux Enterprise Server 11 are as follows:

Component	**Parameter**	**SUSE Linux Enterprise Server 9**		**SUSE Linux Enterprise Server 10 and SUSE Linux Enterprise Server 11**	
		Default setting	**Required setting**	**Default setting**	**Required setting**
QLogic external driver for FC	MaxRetriesPerIo	0	10	0	10
	MaxRetriesPerPath	0	10	0	10
	ql2xsuspendcount	10	6	10	6

Component	Parameter	SUSE Linux Enterprise Server 9		SUSE Linux Enterprise Server 10 and SUSE Linux Enterprise Server 11	
		Default setting	**Required setting**	**Default setting**	**Required setting**
QLogic bundled driver for FC	ql2xsuspendcount	10	6	Not available	Not available
	MaxRetriesPerIo	0	10	Not available	Not available
	MaxRetriesPerPath	0	10	Not available	Not available
QLogic FC HBA	Link Down Timeout	30	20	30	20

Note: For Emulex, no changes are required to the default settings.

Loading a bundled or an external HBA driver on a Linux host using FC

You can load a bundled or an external HBA driver on your Linux host by using the `modprobe` command.

Step

1. Depending on your HBA model, load the bundled or external HBA driver by performing the following step:

If you are using a...	Enter the following command...
QLogic HBA driver	**`modprobe -v`** ***`driver model`*** Example: **`modprobe -v qla2xxx`**
Emulex HBA driver	**`modprobe -v lpfc`**

Loading a bundled HBA driver on a Linux host using hardware iSCSI

You can load a bundled QLogic HBA driver on your Linux host by using the `modprobe` command.

Step

1. Enter the following command on the Linux host:

   `modprobe -v qla4xxx`

Verifying the QLogic HBA settings for FC

You can verify the QLogic HBA settings for the PortDownRetryCount and LinkDownTimeout parameters for FC.

Step

1. To verify the PortDownRetryCount and LinkDownTimeout parameters, run the `scli` command with the `-c` option:

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

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

```
-------------------------------------------------
HBA 0: QLE2462 Port 1 WWPN 21-00-00-E0-8B-80-72-A1 Port
ID 03-05-00
-------------------------------------------------------
Connection Options             : 1 - Point-to-Point Only.
Data Rate                      : 4 Gbps
Frame Size                     : 2048
Hard Loop ID                   : 0
Loop Reset Delay (seconds)     : 5
Enable Host HBA BIOS           : Disabled
Enable Hard Loop ID            : Disabled
Enable FC Tape Support         : Enabled
Operation Mode                 : 0 - Interrupt for every I/O completion
Interrupt Delay Timer (100ms)  : 0
Execution Throttle             : 16
Login Retry Count              : 8
Port Down Retry Count          : 30
Enable LIP Full Login          : Enabled
Link Down Timeout (seconds)    : 20
Enable Target Reset            : Enabled
LUNs Per Target                : 128
-------------------------------------------------------
HBA 1: QLE2462 Port 2 WWPN 21-01-00-E0-8B-A0-72-A1 Port
```

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

Emulex HBA settings for FC

If you are using an Emulex HBA, no changes are required to the default settings.

Verifying the QLogic HBA settings for hardware iSCSI

You can verify the QLogic HBA settings for the ql4xkeepalive parameter for hardware iSCSI.

Steps

1. Check the value for the ql4xkeepalive parameter in the corresponding `qla4xxx` directory.

   The values of the external driver parameter can be found in `sysfs`, `/sys/module/qla2xxx/`, or `/sys/module/qla2xxx/parameters/`, depending on the driver version.

   Example
   You can check the value of ql4xkeepalive on a host by running the following command:

   `cat /sys/module/qla4xxx/parameters/ql4xkeepalive`

2. To check for IP_ARP_Redirect value, run the `iscli` command with the `-c` option:

   `/opt/QLogic_Corporation/SANsurferCLI/iscli -c`

   Example

```
*** Displaying Advanced Settings inst=0 ***
FW_Marker                           :  on(*)
FW_Stat_Alarm                       :  off(*)
FW_Accept_AEN                       :  off(*)
FW_Access_Control                   :  off(*)
FW_Session_Mode                     :  on(*)
```

```
FW_Initiator_Mode                    :  on(*)
FW_Target_Mode                       :  off(*)
FW_Fast_Posting                      :  off(*)
FW_Sense_Buffer_Desc                 :  off(*)
FW_ZIO_Enable_Mode                   :  off
AFW_Device_Timeouts                  :  on
AFW_Delayed_Ack                      :  on
AFW_AutoConnect                      :  on
DevType                              :  0(*)
ExeThrottle                          :  64
FirstBurstLen                        :  256
IP_Fragmentation                     :  on(*)
IP_ARP_Redirect                      :  on
VLAN_Enable                          :  off
VLAN_User_Priority                   :  0
VLAN_ID                              :  0
IPv4_TOS_ENABLE                      :  off
Force_Negotiate_Main_iSCSI_Keys      :  off
iSCSI_Send_Markers                   :  off(*)
iSCSI_Header_Digests                 :  off
iSCSI_Data_Digests                   :  off
iSCSI_Immediate_Data                 :  on
iSCSI_Initial_R2T                    :  off
iSCSI_Data_Seq_In_Order              :  on(*)
iSCSI_Data_PDU_In_Order              :  on(*)
iSCSI_CHAP_Auth                      :  off(*)
iSCSI_Bidi_CHAP_Auth                 :  off(*)
iSCSI_Error_Recovery_Level           :  0(*)
KeepAliveTO                          :  30
DefaultTime2Retain                   :  20(*)
DefaultTime2Wait                     :  2(*)
MaxBurstLen                          :  512
MaxOutstandingR2T                    :  1
MaxRxDataSegmentLen                  :  128(*)
```

Obtaining WWPNs of HBAs for FC

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

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

Each physical port has its own WWPN. A two-port HBA has two unique WWPNs.

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

Next topics

Obtaining the WWPN of QLogic HBAs using FC on page 39
Obtaining the WWPN for Emulex HBAs using FC on page 39

Obtaining the WWPN of QLogic HBAs using FC

You can obtain the WWPN of QLogic HBAs by using the SANsurfer FC HBA CLI client.

Step

1. Enter the following command on the Linux host console:

   `/opt/QLogic_Corporation/SANsurferCLI/scli -i`

Obtaining the WWPN for Emulex HBAs using FC

You can obtain the WWPN of Emulex HBA by using the Emulex HBAnyware utility.

Step

1. Enter the following command on the Linux host console:

   `/usr/sbin/hbanyware/hbacmd ListHBAs`

Adding targets to the host using hardware iSCSI

You should add targets to the host by using the `iscli` command.

Steps

1. Run the following command on the host:

   `/opt/QLogic_Corporation/SANsurferCLI/iscli`

2. Add the target details, which include the following:
   - Target node name
   - Target IP address
3. Select `Commit Changes` for this HBA so that the changes are applied to the HBA port.
4. Repeat the settings for each HBA port that is listed.

Obtaining the initiator node name of QLogic HBAs for hardware iSCSI

The initiator node name is required to create igroups on the storage system. You map LUNs to specific igroups for the hosts in the igroup to discover the LUNs.

Step

1. To obtain the initiator node name of QLogic HBA, enter the following command on the Linux host console:

   `/opt/QLogic_Corporation/SANsurferCLI/iscli -i`

Unloading the bundled or external HBA driver for FC

You can unload the HBA driver for FC by using the `modprobe` command.

Before you begin

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

Step

1. Depending on your HBA model, unload the QLogic HBA driver by performing the following step:

If you are using a...	Enter the following command...
QLogic HBA driver	**`modprobe -r`** ***`driver module`*** Example: **`modprobe -r qla2xxx`**
Emulex HBA driver	**`modprobe -r lpfc`**

Unloading the bundled HBA driver for hardware iSCSI

You can unload the HBA driver for hardware iSCSI by using the `modprobe` command.

Before you begin

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

Step

1. Enter the following command for the QLogic HBA driver model:

   ```
   modprobe -r qla4xxx
   ```

The Host Utilities installation

The Linux Host Utilities installation process involves several tasks. You must download the appropriate Host Utilities installation file and install the software.

Next topics

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43

Downloading the Host Utilities software

You can download the Host Utilities software package from the NOW site.

Before you begin

Verify that you have the correct version of the Linux Host Utilities for your version of the Linux host. For information about the correct version, see the Interoperability Matrix.

Step

1. Download the Linux Host Utilities file (`netapp_linux_host_utilities_5_2.tar.gz`) from the NOW site to your Linux host.

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability/
NOW site - http://now.netapp.com/NOW/cgi-bin/software/

Installing the Host Utilities software

After you download the Host Utilities software by extracting the files from the tar file, you can install the software on your Linux host.

Before you begin

Ensure that the following packages are installed on the host before installing Linux Host Utilities 5.2:

- perl-HTML-Tagset
- perl-HTML-Parser
- perl-XML-Parser
- perl-libwww-perl

Steps

1. To remove any previous version of the Linux Host Utilities, go to the directory where the previous version is installed (the default is `/opt/netapp/santools`) and enter the following command:

   ```
   ./uninstall
   ```

2. Go to the directory to which you downloaded the latest Host Utilities file.
3. To extract the files, enter the following command:

   ```
   tar -zxvf netapp_linux_host_utilities_5_2.tar.gz
   ```

4. Change to the `netapp_linux_host_utilities_5_2` directory.

   By default, this is a subdirectory of the working directory in which you extracted the Host Utilities file in the previous step.

5. Enter the following command:

   ```
   ./install
   ```

Configuration of the software iSCSI initiator

You need to configure the iSCSI initiator on the Linux host before you can discover the target and access the LUNs on the target through the iSCSI protocol. You need to configure the iSCSI initiator on the Linux host to configure the host on the storage system. To configure the initiator settings, you require root privileges.

Note: Both Red Hat Enterprise Linux and SUSE Linux Enterprise Server contain the software iSCSI initiator package.

Next topics

Verifying the required iSCSI package on your Linux host on page 45
Why you need the iSCSI initiator node name on the Linux host on page 46
Node name rules on page 46
Recording or editing the initiator node name on your Linux host on page 46
Configuration of software iSCSI on Red Hat Enterprise Linux on page 48
Configuration of software iSCSI on SUSE Linux Enterprise Server on page 53
Configuring the iSCSI service to start automatically on page 60

Verifying the required iSCSI package on your Linux host

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

Step

1. Depending on your Linux host, perform the appropriate action:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	`rpm -q iscsi-initiator-utils`
SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11	`rpm –q open-iscsi`
SUSE Linux Enterprise Server 9 series	`rpm -q linux-iscsi`

The `rpm -q` command returns the name and version of the iSCSI package.

Why you need the iSCSI initiator node name on the Linux host

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

When you start the iSCSI service for the first time, the initiator node name is generated.

For ease of use, you should change the random number of the iSCSI initiator node name to the host name.

For example, if the random number of the iSCSI initiator node name is `iqn.2005-03.com.RedHat:012345`, you should change it to `iqn.2005-03.com.RedHat:`*`host_name`*.

`host_name` is the name of your Linux host.

Node name rules

You should remember certain rules when you change the host's initiator node name.

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

- A node name can be up to 223 bytes.
- Uppercase characters are always mapped to lowercase characters.
- A node name can contain alphabetic characters (a to z), numbers (0 to 9), and the following three special characters:
  - Period (.)
  - Hyphen (-)
  - Colon (:)

  Note: The underscore character (_) is not supported.

Recording or editing the initiator node name on your Linux host

You need to record the initiator node name on your Linux host to use them later while configuring the storage system. You can also edit the initiator node name if required.

Steps

1. With a text editor, open the following file:

If you are using...	Open the following file...
Red Hat Enterprise Linux 5 series, SUSE Linux Enterprise Server 10 series, and SUSE Linux Enterprise Server 11	`/etc/iscsi/initiatorname.iscsi`
Red Hat Enterprise Linux 4 series and SUSE Linux Enterprise Server 9 series	`/etc/initiatorname.iscsi`

2. To view or change the default initiator name, perform the appropriate action:

If...	Then...
A default initiator name has not been generated	**Red Hat Enterprise Linux 5 series** Replace the text in the host's `/etc/iscsi/initiatorname.iscsi` file with the following line: **`InitiatorName=iqn.2005-03.com.RedHat:`*`RandomNumber`*** **Red Hat Enterprise Linux 4 series** Replace the text in the host's `/etc/initiatorname.iscsi` file with the following line: **`InitiatorName=iqn.1987-05.com.cisco:`*`RandomNumber`*** **SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11** Replace the text in the host's `/etc/iscsi/initiatorname.iscsi` file with the initiator name in the following format: **`InitiatorName=iqn.1996-04.de.suse:`*`RandomNumber`*** **SUSE Linux Enterprise Server 9 series** Replace the text in the host's `/etc/initiatorname.iscsi` file with the initiator name in the following format: **`InitiatorName=iqn.1987-05.com.cisco:`*`RandomNumber`***
You want to use the default initiator node name	Record the node name and proceed to editing the host's iSCSI configuration file.

If...	Then...
You want to change the initiator node name	**Red Hat Enterprise Linux 5 series** Modify the *RandomNumber* part of the initiator node name. The following line shows an example node name: **`InitiatorName=iqn.2005-03.com.RedHat:linux-host1`** **Red Hat Enterprise Linux 4 series** Modify the *RandomNumber* part of the initiator node name. The following line shows an example node name: **`InitiatorName=iqn.1987-05.com.cisco:linux-host1`** **SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11** Modify the *RandomNumber* part of the initiator node name. The following is a sample node name: **`InitiatorName=iqn.1996-04.de.suse:lnx.118`** **SUSE Linux Enterprise Server 9 series** Modify the *RandomNumber* part of the initiator node name. The following is a sample node name: **`InitiatorName=iqn.1987-05.com.cisco:01.82ea2060f487`**

Note: The initiator node name should be unique.

3. Record the node name and proceed to editing the host's iSCSI configuration file.

Related concepts

DM-Multipath configuration on page 65
Configuration of software iSCSI on SUSE Linux Enterprise Server on page 53
Configuration of software iSCSI on Red Hat Enterprise Linux on page 48

Related tasks

Editing the DM-Multipath configuration file on page 66

Configuration of software iSCSI on Red Hat Enterprise Linux

You need to edit your Linux host's configuration file to configure iSCSI to use DM-Multipath, to set up CHAP (optional), and to configure target discovery.

Next topics

Configuration of the iSCSI initiator on Red Hat Enterprise Linux 5 series on page 49
Configuration of the iSCSI initiator on Red Hat Enterprise Linux 4 series on page 51

Configuration of the iSCSI initiator on Red Hat Enterprise Linux 5 series

Red Hat Enterprise Linux includes the iSCSI initiator software in the operating system distribution. The initiator's components include a kernel module, which is already compiled into the Linux kernel, and user space packages.

Next topics

Editing the iSCSI configuration file on Red Hat Enterprise Linux 5 series on page 49
Starting the iSCSI service on Red Hat Enterprise Linux 5 series on page 50
Setting up target discovery on Red Hat Enterprise Linux 5 series on page 50
Configuring manual or automatic login on Red Hat Enterprise Linux 5 series on page 50

Editing the iSCSI configuration file on Red Hat Enterprise Linux 5 series

You should edit the `/etc/iscsi/iscsid.conf` file to set the iSCSI parameter on Red Hat Enterprise Linux 5 series.

Steps

1. Open the `/etc/iscsi/iscsid.conf` file with the text editor.
2. If you are using DM-Multipath, change the following iSCSI parameter:
   - For native:
     `node.session.timeo.replacement_timeout` to `5`
   - For Veritas Storage Foundation:
     `node.session.timeo.replacement_timeout` to `120`
3. Add CHAP user names and passwords, as shown in the following steps:

   Note: You can set up CHAP authentication for enhanced security on Red Hat Enterprise Linux 5 series, if required.

   a. To enable CHAP authentication, set `node.session.auth.authmethod` to `CHAP`. The default is None.
   ```
   node.session.auth.authmethod = CHAP
   ```
   b. To set a CHAP user name and password for initiator authentication by the target or targets, remove the comment and provide a user name and password from the following lines:
   ```
   node.session.auth.username = username
   node.session.auth.password = password
   ```
   c. To set a CHAP user name and password for target authentication by the initiator, remove the comment and provide `username_in` and `password_in` from the following lines:
   ```
   node.session.auth.username_in = username_in
   node.session.auth.password_in = password_in
   ```

d. For a successful session discovery, enable 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
```

Note: The user name and password must match for both session and discovery on the host. Ensure that you use the same user names and passwords when you set up CHAP on the storage system with the `iscsi security` command.

Starting the iSCSI service on Red Hat Enterprise Linux 5 series

You should start the iSCSI service on the Red Hat Enterprise Linux 5 series to discover the targets and access the LUN.

Step

1. To start the iSCSI service, enter the following command at the Linux host command prompt:

   `/etc/init.d/iscsi start`

Setting up target discovery on Red Hat Enterprise Linux 5 series

You can set up target discovery to access the LUNs on the host by using the `iscsiadm` utility.

Step

1. To discover the target, use the following command:

   `iscsiadm -m discovery -t st -p` ***`ip_address`***

   `ip_address` is the target address.

Configuring manual or automatic login on Red Hat Enterprise Linux 5 series

You can log in to the nodes either automatically or manually. To log in automatically to the nodes, you should start the iSCSI service. To manually log in, you should use the `iscsiadm` utility.

Step

1. To set the start-up mode of iSCSI sessions as manual or automatic, perform the appropriate action:

If you want to set the login mode to...	Then...
Manual or automatic on a single portal on a target	Run the following command: **`iscsiadm -m node -T targetname -p ip:port -o update -n node.startup -v manual|automatic`** **`iscsiadm -m node -T targetname -p ip:port -o update -n node.conn[0].startup -v manual|automatic`**
Manual or automatic on all of the portals on a target	Run the following command: **`iscsiadm -m node -T targetname -o update -n node.startup -v manual|automatic`** **`iscsiadm -m node -T targetname -o update -n node.conn[0].startup -v manual|automatic`**
Manual or automatic mode for all the targets	**a.** Add the following to the `/etc/iscsi/iscsid.conf` file: **`node.startup = manual|automatic`** **b.** Rediscover the iSCSI target. **c.** Restart the iSCSI service

Note: If you set up your login mode to manual, you should manually log in to the nodes during the next start-up.

If you set up your login mode to automatic, you are logged in to the nodes automatically during the next start-up.

Setting the login mode to manual or automatic affects only nodes that are discovered after setting the value.

Configuration of the iSCSI initiator on Red Hat Enterprise Linux 4 series

Red Hat Enterprise Linux 4 series includes the iSCSI initiator software in the operating system distribution. The initiator's components include a kernel module, which is already compiled into the Linux kernel, and user space packages.

Next topics

Editing the iSCSI configuration file on Red Hat Enterprise Linux 4 series on page 52
Starting the iSCSI service on Red Hat Enterprise Linux 4 series on page 53

Editing the iSCSI configuration file on Red Hat Enterprise Linux 4 series

You should edit the `/etc/iscsi.conf` file to set the iSCSI parameter on Red Hat Enterprise Linux 4 series.

Steps

1. Open the `/etc/iscsi.conf` file with a text editor.
2. If you are using DM-Multipath, perform the following:
   a. Remove the comment of the `ConnFailTimeout` line in the Session Timeout Settings section.
   b. Set the value of `ConnFailTimeout` to `5`.
3. Add CHAP user names and passwords.

   Example

   ```
   DiscoveryAddress=192.168.10.20
         OutgoingUsername=username_out
         OutgoingPassword=password_out
         IncomingUsername=username_in
         IncomingPassword=password_in
   ```

   Note: You can set up CHAP authentication for enhanced security on Red Hat Enterprise Linux 4 series, if required.

   CHAP settings should be indented by using white space or tab, below the `DiscoveryAddress` of the storage system they apply to.

   - For unidirectional authentication, you should define only the `OutgoingUsername` and `OutgoingPassword`.

     Use the `OutgoingUsername` and `OutgoingPassword` for the storage system's inbound user name and password (*`inname`* and *`inpassword`*).
   - For bidirectional authentication, you should define both outgoing and incoming user names and passwords.

     Use `IncomingUsername` and `IncomingPassword` of the host as the storage system's outbound user name and password (*`outname`* and *`outpassword`*).

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

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

4. Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you used for iSCSI LUNs:

   `DiscoveryAddress=`*`storage_system_IPaddress`*

storage_system_IPaddress is the IP address of an Ethernet interface on the storage system. You should specify an interface that is used for iSCSI communication.

Example:

```
DiscoveryAddress=192.168.10.100
DiscoveryAddress=192.168.10.20
```

Starting the iSCSI service on Red Hat Enterprise Linux 4 series

You should start the iSCSI service on Red Hat Enterprise Linux 4 series to access the LUN.

Step

1. To start the iSCSI service, enter the following command at the Linux host command prompt:

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

Configuration of software iSCSI on SUSE Linux Enterprise Server

You should perform certain tasks to configure software iSCSI on a SUSE Linux Enterprise Server host.

Next topics

Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 10 series and 11 on page 53
Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 9 series on page 58

Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 10 series and 11

SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11 include the iSCSI initiator software in the operating system distribution. The initiator's components include a kernel module, which is already compiled into the Linux kernel, and user space packages.

Next topics

Editing the iSCSI configuration file on SUSE Linux Enterprise Server 10 series and 11 on page 54
Starting the iSCSI service on SUSE Linux Enterprise Server 10 series and 11 on page 55
Setting up target discovery on SUSE Linux Enterprise Server 10 series and 11 on page 55
Configuring manual or automatic login on SUSE Linux Enterprise Server 10 series and 11 on page 57

Editing the iSCSI configuration file on SUSE Linux Enterprise Server 10 series and 11

You should edit the `/etc/iscsi/iscsid.conf` file to set the iSCSI parameter on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11.

Steps

1. Open the `/etc/iscsi/iscsid.conf` file with a text editor.
2. If you are using DM-Multipath, change the following iSCSI parameter for SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11:
   - For native:

     `node.session.timeo.replacement_timeout = 5`
   - For Veritas Storage Foundation:

     `node.session.timeo.replacement_timeout = 120`
3. Add CHAP user names and passwords, as shown in the following steps:
   a. To enable CHAP authentication, set `node.session.auth.authmethod` to `CHAP`.
      The default is None.
      `node.session.auth.authmethod = CHAP`
   b. To set a CHAP user name and password for initiator authentication by the target or targets, remove the comment and provide *`username`* and *`password`* from the following configuration entries:

      `node.session.auth.username = `*`username`*

      `node.session.auth.password = `*`password`*
   c. To set a CHAP user name and password for target or targets authentication by the initiator, remove the comment and provide *`username_in`* and *`password_in`* from the following configuration entries:

      `node.session.auth.username_in = `*`username_in`*

      `node.session.auth.password_in = `*`password_in`*

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

Starting the iSCSI service on SUSE Linux Enterprise Server 10 series and 11

You should start the iSCSI service on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11 to discover the targets and access the LUN.

Step

1. To start the iSCSI service, enter the following command at the Linux host command prompt:

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

Setting up target discovery on SUSE Linux Enterprise Server 10 series and 11

You should configure the iSCSI initiator to discover the target and to set the iSCSI session on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11. You can configure the iSCSI initiator by using the YaST2 or `iscsiadm` utility.

Next topics

Discovering the iSCSI target by using YaST2 on page 55
Discovering the iSCSI target by using the iscsiadm utility on page 56

Discovering the iSCSI target by using YaST2

You can use YaST2 for the discovery and configuration of iSCSI connections on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11. Using YaST2, you can enable the iSCSI initiator at boot time, add new targets to the system, and discover iSCSI targets in the network. You can also view the iSCSI targets that are currently connected.

Steps

1. Enter the following command:

   ```
   yast2
   ```

2. Click **Network Services > iSCSI Initiator > Discovered Targets > Discovery** in the **YaST2** window.
3. Ensure the port number is 3260.
4. Enter the IP address of the iSCSI target, and change the port if necessary.
5. If you have an authentication mechanism in place, enter the credentials.
6. Click **Next** to start the discovery.
7. After discovering the targets, use **Login** to activate the target.
8. Enter the authentication credentials to use the selected iSCSI target.
9. Click **Next** to finish the configuration.

The target now appears in **Connected Targets**.

10. Click **Connected Targets**.
11. Change the start-up option to Manual or Automatic, depending on your requirement, by using the Toggle Start-Up button for all the discovered targets.

    For more information about using YaST2, see the SUSE Linux Enterprise Server 10 series documentation available at the Novell Web site.

Related information

Novell Web site - http://www.novell.com

Discovering the iSCSI target by using the iscsiadm utility

You can use the `iscsiadm` utility to manage (update, delete, insert, and query) the persistent database on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11. The utility enables you to perform a set of operations on iSCSI nodes, sessions, connections, and discovery records.

Steps

1. Enter the following command to discover the available iSCSI targets:

   `iscsiadm --mode discovery --op update --type sendtargets --–portal targetip`

   `targetip` is the IP address of the target.

   The target *`targetip`* is discovered. For each target defined on the iSCSI target, a separate line is displayed. All received values are stored in an internal persistent database.

2. Enter the following command to create all needed devices:

   `iscsiadm -m node -l all`

   The initiator logs in to all the discovered nodes that are maintained in the iSCSI database.

3. Enter the following command to see all the active iSCSI sessions:

   `iscsiadm --mode session`

   Example
   The following is a sample output:

   ```
   tcp [022e71] 10.72.200.37:3260,1000 iqn.1992-08.com.netapp:sn.101176399
   ```

   022e71 is the record ID.

   All the established iSCSI sessions are displayed.

Configuring manual or automatic login on SUSE Linux Enterprise Server 10 series and 11

You can log in to the nodes either automatically or manually on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11. To log in to the nodes automatically, you should start the iSCSI service. To log in manually, you should use the `iscsiadm` utility.

Step

1. To set the startup mode of iSCSI sessions as manual or automatic, perform the appropriate action:

If you want to set the login mode to...	Do this...
Manual or automatic on a single portal on a target	Run the following command: **`iscsiadm -m node -T targetname -p ip:port -o update -n node.startup -v manual|automatic`** **`iscsiadm -m node -T targetname -p ip:port -o update -n node.conn[0].startup -v manual|automatic`**
Manual or automatic on all of the portals on a target	Run the following command: **`iscsiadm -m node -T targetname -o update -n node.startup -v manual|automatic`** **`iscsiadm -m node -T targetname -o update -n node.conn[0].startup -v manual|automatic`**
Manual or automatic mode for all the targets	**a.** Add the following to the `/etc/iscsi/iscsid.conf` file: **`node.startup = manual|automatic`** **b.** Rediscover the iSCSI target. **c.** Restart the iSCSI service.

Note: If you set up your login mode to manual, you should manually log in to the nodes during the next start-up.

If you set up your login mode to automatic, you are logged in to the nodes automatically during the next start-up.

Setting the login mode to manual or automatic affects only nodes that are discovered after setting the value.

Configuration of the iSCSI initiator on SUSE Linux Enterprise Server 9 series

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

Next topics

Editing the iSCSI configuration file on SUSE Linux Enterprise Server 9 series on page 58
Starting the iSCSI service on SUSE Linux Enterprise Server 9 series on page 59

Editing the iSCSI configuration file on SUSE Linux Enterprise Server 9 series

You should edit the `/etc/iscsi.conf` file to configure the iSCSI initiator.

Steps

1. Edit the `/etc/iscsi.conf` file on the Linux host console.
2. At the beginning of the file, add the following configuration entries before any other configuration setting:

   `Continuous=no`

   `HeaderDigest=never`

   `DataDigest=never`

   Attention: These configuration entries are required for iSCSI to work properly. If you do not add these entries at the beginning of the file, you might experience problems with connectivity and iSCSI performance.

3. If you are using DM-Multipath, perform the following steps:
   a. Remove the comment to enable `Multipath=portal` in the Multipath Settings section.
   b. Remove the comment from `ConnFailTimeout` in the Session Timeout Settings section and set the value to `ConnFailTimeout=5`.
4. To discover the storage system, add the IP address of the storage system by using the following format:

   `DiscoveryAddress=`*`storage_system_IPaddress`*

   `storage_system_IPaddress` is the IP address of an Ethernet interface on the storage system that is used for iSCSI communication.

   Example
   The following are sample `DiscoveryAddress` entries that are using the IP addresses of the storage systems:

```
DiscoveryAddress=192.168.10.100
DiscoveryAddress=10.61.208.200
```

5. Add the CHAP user names and passwords, as shown in the following example.

   CHAP settings should be indented below the DiscoveryAddress of the storage system they apply to.

   - For unidirectional authentication, you should define only the outgoing user name and outgoing password.
     Use the `OutgoingUsername` and `OutgoingPassword` for the storage system's inbound user name and password (*`inname`* and *`inpassword`*).
   - For bidirectional authentication, you should define both outgoing and incoming user names and passwords.
     Use `IncomingUsername` and `IncomingPassword` of the host as the storage system's outbound user name and password (*`outname`* and *`outpassword`*).

   When adding the global CHAP entries in the `/etc/iscsi.conf` file, the sequence should be as follows:

   - Outgoing CHAP settings
   - Discovery Addresses
   - Incoming CHAP settings (in case of bidirectional CHAP)

   Example

```
DiscoveryAddress=192.168.10.20
      OutgoingUsername=username_out
      OutgoingPassword=password_out
      IncomingUsername=username_in
      IncomingPassword=password_in
```

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

Starting the iSCSI service on SUSE Linux Enterprise Server 9 series

You should start the iSCSI service on the SUSE Linux Enterprise Server 9 series to access the LUN.

Step

1. To start the iSCSI service, enter the following command at the Linux host command prompt:

   `/etc/init.d/iscsi start`

Configuring the iSCSI service to start automatically

You can use the `chkconfig` command to configure the iSCSI service to start automatically at system boot.

Step

1. Enter the following command at the Linux host command prompt:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	**`chkconfig iscsi on`**
SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11	**`chkconfig open-iscsi on`**
SUSE Linux Enterprise Server 9 series	**`chkconfig iscsi on`**

Storage system configuration

You should configure your storage system and prepare it for initiator access by using FC, software iSCSI, or hardware iSCSI.

When creating igroups on the storage system, you need the following:

- FC
  The host's WWPNs that you recorded when installing the HBAs.
- Software iSCSI
  The host's software iSCSI initiator node name that you recorded when configuring software iSCSI on your Linux host.
- Hardware iSCSI
  The host's initiator name that you recorded when installing the HBAs.

Next topics

Configuring the storage system for FC on page 61
Configuring the storage system for hardware and software iSCSI on page 62

Related concepts

HBA installation on page 27
Obtaining WWPNs of HBAs for FC on page 38
Configuration of the software iSCSI initiator on page 45

Related tasks

Obtaining the initiator node name of QLogic HBAs for hardware iSCSI on page 40

Configuring the storage system for FC

You need to configure the storage system and prepare it for initiator access.

Before you begin

You need the host's WWPNs that you recorded when installing the HBAs. Ensure that you have added WWPNs either when you create igroups on the storage systems or after creating the igroups.

About this task

You need to know the different cluster failover modes that are supported by NetApp storage systems. The cfmode setting controls how the storage system handles FC before and after cluster failover events. The cfmode is not available on nonclustered storage controllers.

The following cfmodes are supported:

- Single System Image (single_image) cfmode
- Dual Fabric (dual_fabric) cfmode
- Partner cfmode
- Standby cfmode
- Mixed cfmode

Steps

1. Ensure that the FC service is licensed and started.
2. Verify that the storage system is running a supported cfmode by running the `fcp show cfmode` command.

   The support matrix lists the supported cfmode values for your configuration. For standby and mixed cfmode, special FC switch zoning is required.

   Note: In a switched fabric use single initiator zoning for each of the initiator ports, that is, each zone should have only a single initiator port.

3. Create and map the LUNs to the igroup for the Linux host.

   You can manage LUNs by using the Web-based FilerView interface or the Data ONTAP command-line interface. Ensure that you specify the LUN operating system type and igroup type as `linux`.

4. Optionally, enable ALUA.
5. If you are using clustered Linux hosts, ensure that the igroup contains the WWPNs of all the hosts in the cluster that need access to the mapped LUN.

Related concepts

Obtaining WWPNs of HBAs for FC on page 38

Configuring the storage system for hardware and software iSCSI

You can configure your storage system for hardware and software iSCSI.

Before you begin

Ensure that you have added WWPNs either when you create igroups on the storage systems or after creating the igroups.

Steps

1. Ensure that the iSCSI service is licensed and started.
2. If you want to use CHAP authentication, use the `iscsi security` command or the FilerView interface to configure a CHAP user name and password.
3. Create and map the LUNs to the igroup for the Linux host.

   You can manage the LUNs by using the Web-based FilerView interface or the Data ONTAP command-line interface. Ensure that you specify the LUN operating system type and igroup type as `linux`.
4. If you are using clustered Linux hosts, ensure that the igroup contains the initiator names of all the hosts in the cluster that need access to the mapped LUN.

DM-Multipath configuration

You can configure DM-Multipath on your host. DM-Multipath enables you to configure multiple I/O paths between the host and storage controllers into a single device. If one path fails, DM-Multipath reroutes I/Os to the remaining paths.

Next topics

Verifying the required multipathing packages on page 65
Editing the DM-Multipath configuration file on page 66
Starting DM-Multipath on page 71
Verifying the DM-Multipath configuration on page 72
Stopping DM-Multipath on page 73

Verifying the required multipathing packages

You need to verify that you have installed the multipathing packages.

Steps

1. Enter one of the following commands:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	`rpm -q device-mapper` `rpm -q device-mapper-multipath`
SUSE Linux Enterprise Server	`rpm -q device-mapper` `rpm -q multipath-tools`

   The `rpm -q` commands return the name and version of the installed packages.

2. If you do not have the required packages, install the appropriate packages based on your version of Linux.

   - Obtaining RPMs
     You can get the RPMs either from the Red Hat Enterprise Linux media or from the Red Hat Network at the Red Hat Web site (subscription required).
   - Copying the file
     The `multipath.conf` file contains configuration settings for multipath. A sample file is copied to your host when you install the multipath-tools package `/usr/share/doc/packages/multipath-tools/multipath.conf.annotated`.

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

Related information

Red Hat Web site - http://www.redhat.com/software/rhn/

Editing the DM-Multipath configuration file

For DM-Multipath to function properly, you should edit the configuration file.

Before you begin

If the `/etc/multipath.conf` file exists, remove the content or, alternatively, create a new file.

Steps

1. Open the `/etc/multipath.conf` file and remove the content.
2. Depending on your version of the operating system, copy the respective example.

   The following examples show the required changes to the `/etc/multipath.conf` file.

   Example

   Red Hat Enterprise Linux 5 series

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

```
        }
}
```

DevID is the unique ID.

Note:

- If you are configuring DM-Multipath on an iSCSI host, set path_grouping_policy to `multibus`.
- flush_on_last_del is available from Red Hat Enterprise Linux 5 Update 3 and later.
- directio pathchecker is available in Red Hat Enterprise Linux 5 Update 1 and later. For Red Hat Enterprise Linux 5, use readsector0 pathchecker.
- max_fds is available from Red Hat Enterprise Linux 5 Update 2 and later. Set the max_fds value to `4096` on Red Hat Enterprise Linux 5 Update 2 host.

Example

Red Hat Enterprise Linux 4 series

```
defaults
{
   user_friendly_names   yes
   max_fds               4096
}
devnode_blacklist
{
   wwid DevId
   devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
   devnode "^hd[a-z]"
   devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
   device
   {
     vendor                  "NETAPP"
     product                 "LUN"
     flush_on_last_del       yes
     getuid_callout          "/sbin/scsi_id -g -u -s /block/%n"
     prio_callout            "/sbin/mpath_prio_ontap /dev/%n"
     features                "1 queue_if_no_path"
     hardware_handler        "0"
     path_grouping_policy    group_by_prio
     failback                immediate
     rr_weight               uniform
     rr_min_io               128
     path_checker            directio
     }
}
```

DevID is the unique ID.

Note:

- If you are configuring DM-Multipath on an iSCSI host, set path_grouping_policy to `multibus` for Red Hat Enterprise Linux 4 Update 7 and earlier.
- `directio` path_checker is available from Red Hat Enterprise Linux 4 Update 7 and later. For Red Hat Enterprise Linux 4 Update 6 and earlier, use `readsector0`path_checker.
- flush_on_last_del is available from Red Hat Enterprise Linux 4 Update 7 and later.
- max_fds is available in Red Hat Enterprise Linux 4 Update 7.
- For Red Hat Enterprise Linux 4 Update 7 and earlier, rr_min_io should be set to `128` in the default section and not in the device section.

Example

SUSE Linux Enterprise Server 11

```
defaults
{
  user_friendly_names    no
  max_fds                8192
}
blacklist
{
  wwid       DevId
  devnode    "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
  devnode    "^hd[a-z]"
  devnode    "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
   device
    {
      vendor                  "NETAPP"
      product                 "LUN"
      getuid_callout          "/lib/udev/scsi_id -g -u -d /dev/%n"
      prio                    "ontap"
      features                "1 queue_if_no_path"
      hardware_handler        "0"
      path_grouping_policy    group_by_prio
      failback                immediate
      rr_weight               uniform
      rr_min_io               128
      path_checker            directio
    }
}
```

`DevID` is the unique ID.

Example

SUSE Linux Enterprise Server 10 series

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

```
   devnode "^hd[a-z]"
   devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"
}
devices
{
    device
     {
      vendor                "NETAPP"
      product               "LUN"
      getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
      prio                  "ontap"
      features              "1 queue_if_no_path"
      hardware_handler      "0"
      path_grouping_policy  group_by_prio
      failback              immediate
      rr_weight             uniform
      rr_min_io             128
      path_checker          directio
      }
}
```

DevID is the unique ID.

Note:

- If you are configuring DM-Multipath on an iSCSI host, set path_grouping_policy to `multibus` for SUSE Linux Enterprise Server 10 SP2 and earlier.
- You should use prio_callout `"/sbin/mpath_prio_ontap /dev/%n"` for SUSE Linux Enterprise Server 10 SP1, instead of prio `"ontap"` in the above example.

Example

SUSE Linux Enterprise Server 9 series

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

```
        rr_min_io              128
        path_checker           readsector0
        }
}
```

DevID is the unique ID.

3. Exclude (blacklist) all the devices that do not correspond to any LUNs configured on the storage controller and which are mapped to your Linux host (that is, exclude those devices that are not displayed by the `sanlun lun show` command).

   a. Blacklist all non-NetApp SCSI devices installed on your host by using the WWID method. To obtain the WWID of a device, run the `scsi_id` command on the respective device.

      Example
      Assume `/dev/sda` is a local SCSI drive. To obtain the WWID, run the following command:

      - SUSE Linux Enterprise Server 11:

        `/lib/udev/scsi_id -gud /dev/sda`
      - For any other Linux operating system:

        `scsi_id -gus /block/sda`

      The output looks similar to the following:
      `SIBM-ESXSMAW3073NC_FDAR9P66067W`

      Add this entry for blacklisting.

   b. Blacklist other devices by using the devnode method by adding the following:

      `devnode "^hd[a-z]"`

      `devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"`

      `devnode "^cciss!c[0-9]d[0-9]*"`

      Example
      On Red Hat Enterprise Linux 4 and SUSE Linux Enterprise Server 9 hosts:

```
devnode_blacklist
{
devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
devnode "^hd[a-z][0-9]*"
devnode "^cciss!c[0-9]d[0-9]*p[0-9]*"
wwid SIBM-ESXSMAW3073NC_FDAR9P66067WJ
}
```

4. If you are using ALUA, specify the ALUA callout program:

If you are using...	Perform the following...
SUSE Linux Enterprise Server 10 SP2 and later	Instead of `"ontap"` for prio, enter the following: **`"alua"`**

If you are using...	Perform the following...
For any other Linux operating system	Instead of `/sbin/mpath_prio_ontap` for prio_callout, enter the following: **`/sbin/mpath_prio_alua`**

Note: ALUA is supported from Red Hat Enterprise Linux 5 Update 2 and SUSE Linux Enterprise Server 10 SP2 and later.

5. Save the changes.

Starting DM-Multipath

You can start DM-Multipath manually to configure LUNs with DM-Multipath.

Steps

1. To start the multipath daemon, perform the appropriate action:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	**`# /etc/init.d/multipathd start`**
SUSE Linux Enterprise Server	**`# /etc/init.d/boot.multipath start`** **`# /etc/init.d/multipathd start`**

2. To configure the DM-Multipath devices, run the following command:

   `# multipath`

Configuring DM-Multipath to start automatically while booting

You can configure DM-Multipath to start automatically while booting.

Step

1. To add the multipath service to the boot sequence, perform the appropriate action on the Linux host console:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	**`chkconfig --add multipathd`** **`chkconfig multipathd on`** **Note:** You should reboot the host if you are configuring the SAN boot on your Linux host.
SUSE Linux Enterprise Server	**`chkconfig --add boot.multipath`** **`chkconfig --add multipathd`** **`chkconfig boot.multipath on`** **`chkconfig multipathd on`**

Verifying the DM-Multipath configuration

You can use the `multipath` command on the Linux host to view the DM-Multipath configuration. You can change the amount of configuration detail that is displayed by using the `-v` parameter.

Steps

1. Enter the following command on the Linux host:

   `multipath -v3 -d -ll`

   Note: The `-d` (dry run) parameter prevents the command from updating the multipath maps.

2. Verify that the multipathd is running by entering the following command:

   `/etc/init.d/multipathd status`

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

   `multipath -ll`

   Example
   The output looks similar to the following for FC:

```
# multipath -ll
3360a98000486e5372635a44646a505643NETAPP,LUN
[size=10G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=8][active]
 \_ 6:0:0:0 sda 8:0   [active][ready]
 \_ 7:0:1:0 sdg 8:96  [active][ready]
\_ round-robin 0 [prio=2][enabled]
```

```
\_ 6:0:1:0 sdc 8:32  [active][ready]
 \_ 7:0:0:0 sde 8:64  [active][ready]
```

Example

The output looks similar to the following for software and hardware iSCSI:

```
# multipath -ll
360a9800043346773435a48514d624367 dm-0 NETAPP,LUN
[size=7.0G][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=4][active]
 \_ 4:0:0:6 sdc 8:32  [active][ready]
 \_ 3:0:0:6 sdd 8:48  [active][ready]
```

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

   `ls -l /dev/mapper/`

   Example

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

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

Stopping DM-Multipath

When you want to stop DM-Multipath on the Linux host, you should stop the affected services.

Steps

1. Unmount all the file systems on the LUNs.
2. Flush the DM-Multipath devices by running the following command:

   `multipath -F`

3. To stop the multipath daemon, perform the appropriate action:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	**`/etc/init.d/multipathd stop`**
SUSE Linux Enterprise Server	**`/etc/init.d/multipathd stop`** **`/etc/init.d/boot.multipath stop`**

The Veritas Array Support Library and Array Policy Module

Symantec provides both the Array Support Library (ASL) and the Array Policy Module (APM) for NetApp storage systems.

Note: Starting with Veritas Storage Foundation 5.1, ALUA is supported on FC.

To determine which version of the ASL and APM you need for this version of the Host Utilities, check the NetApp Interoperability Matrix.

Next topics

What the ASL is on page 75
What the APM is on page 76
Installing the downloaded ASL and APM software on page 76
Removing the ASL and APM on page 77
Information about ASL error messages on page 78

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability

What the ASL is

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

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

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

ASL array type

The ASL reports information about the multipathing configuration to the DDL as an Active/Active (A/A), Asymmetric Logical Unit Access (ALUA) or an Active/Passive Concurrent (A/P-C) disk array type.

- Active/Active (A/A)

There are multiple active paths to a storage system, and simultaneous I/O is supported on each path. If a path fails, I/O is distributed across the remaining paths.

- Active/Passive Concurrent (A/P-C)

  An A/P-C array is a variant of the A/P array type that supports concurrent I/O and load balancing by having multiple primary paths to LUNs. Failover to the secondary (passive) path occurs only if all the active primary paths fail.
- Asymmetric Logical Unit Access (ALUA)

  A LUN in an ALUA-enabled array can be accessed through both controllers, using optimized and non-optimized paths. The array informs the host of path options, their current state, and state changes. Using this information, the host can determine which paths are optimized. Failover to the non-optimized path occurs only if all the optimized paths fail.

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

What the APM is

The APM is a kernel module that defines I/O error handling, failover path selection, and other failover behavior for a specific array.

The NetApp APM is customized to optimize I/O error handling and failover path selection for the NetApp environment. After the ASL discovers the storage array as a array, the ASL instructs VxDMP to use the APM to handle I/O error processing and path failures for the storage array.

Installing the downloaded ASL and APM software

If you are using Veritas Storage Foundation for multipathing, you should install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM) for storage systems.

Before you begin

- Verify that your configuration meets the system requirements.

  For more information, see the NetApp Interoperability Matrix.
- Obtain the ASL and APM software.

  The ASL and APM are not distributed with the Host Utilities software. You can obtain the ASL and APM from the Symantec Web site. For a direct link to the ASL and APM on the Symantec Web site, see the NetApp Interoperability Matrix. You should download this software before you start the installation.

Note: In Veritas Storage Foundation 5.1, the NetApp ASL and APM is included in the Veritas Storage Foundation product.

About this task

- Only one version of the NetApp ASL and APM package can be installed on the host at any given time.

Steps

1. Log in to the Linux host.
2. If you already have the NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering the following command:

   `vxddladm rmjbod vid=NETAPP`
3. Install the ASL and APM according to the instructions provided by Symantec.
4. If your host is connected to a NetApp storage system, verify the installation by following these steps:

   a. Run the following command:

   `vxdmpadm listenclosure all`

   b. Locate NetApp under ARRAY_TYPE in the command output.
   The output shows the model name of the storage device if you are using enclosure- based naming with VxVM.

   Example
   The `vxdmpadm listenclosure all` command shows the Enclosure Type as FAS2020 in this example.

```
#vxdmpadm listenclosure all

ENCLR_NAME  ENCLR_TYPE  ENCLR_SNO  STATUS       ARRAY_TYPE     LUN_COUNT
=====================================================================
=
disk        Disk        DISKS      CONNECTED    Disk              1
fas31700    FAS3170     80010081   CONNECTED    A/A-NETAPP        15
fas31701    FAS3170     80010082   CONNECTED    A/A-NETAPP        15
```

Related information

NetApp Interoperability Matrix-http://now.netapp.com/NOW/products/interoperability

Removing the ASL and APM

If you do not require the ASL and APM, you can remove them from the host.

About this task

You do not need to stop any volumes created on unaffected disk arrays, such as disk arrays from other vendors, before removing ASL and APM support. This is also true for arrays or disks in the

OTHER_DISKS category. The OTHER_DISKS category includes local non-FC attached disks. Volumes created on these arrays remain accessible because they do not have multipath support.

Steps

1. Log in to the host as the root user.
2. Stop all I/O to LUNs configured on storage controllers.

   Note: In a Storage Foundation RAC cluster, you should also stop clustering on a node before you remove the ASL and APM.

3. Enter the following command to remove the ASL package:

   `rpm -ev asl_rpm_name`

   Example

   The following command removes a previous version of the ASL:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	`rpm -ev VRTSNTAPasl-5.0-3.0_RHEL5`
SUSE Linux Enterprise Server	`rpm -ev VRTSNTAPasl-5.0-3.0_SLES10`

4. Enter the following command to remove the APM package:

   `rpm -ev apm_rpm_name`

   Example

   The following command removes a previous version of the APM:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	`rpm -ev VRTSNTAPapm-5.0-1.0_RHEL5`
SUSE Linux Enterprise Server	`rpm -ev VRTSNTAPapm-5.0-1.0_SLES10`

Information about ASL error messages

The ASL works silently and seamlessly with the VxVM DDL. If an error, misconfiguration, or malfunction occurs, messages from the library are logged to the console by using the host's logging facility.

The following table lists the importance and severity of these messages.

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

Methods to access LUNs

The method you use to access LUNs depends on whether you are using DM-Multipath or not.

Next topics

Accessing LUNs with DM-Multipath on page 81
Accessing LUNs without multipathing software on page 94

Accessing LUNs with DM-Multipath

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

For information about the supported configuration for DM-Multipath, see the NetApp Interoperability Matrix.

Note: You should use DM-Multipath if there is more than one physical path from the host to a LUN.

Next topics

FC: Discovering new LUNs on page 82
FC: Viewing the list of LUNs on page 82
FC: Displaying multipath information on page 82
FC: Removing the LUN on page 83
Software iSCSI: Discovering new LUNs on page 84
Software iSCSI: Viewing the list of LUNs on page 87
Software iSCSI: Displaying multipath information on page 92
Software iSCSI: Removing the LUN on page 92
Hardware iSCSI: Discovering new LUNs on page 93
Hardware iSCSI: Viewing the list of LUNs on page 93
Hardware iSCSI: Displaying multipath information on page 93
Hardware iSCSI: Removing the LUN on page 94

Related information

NetApp Interoperability Matrix-http://now.netapp.com/NOW/products/interoperability

FC: Discovering new LUNs

After you create a new LUN and map it to the Linux host, you can discover the LUN using the SCSI rescan script with DM-Multipath. The rescan script is available with the sg3_utils package.

Step

1. Enter the following command to discover a new LUN:

   `rescan-scsi-bus.sh`

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4, SUSE Linux Enterprise Server 10 SP2, and SUSE Linux Enterprise Server 11. For earlier versions, see the HBA vendor-specific documentation.

FC: Viewing the list of LUNs

You can use the `sanlun` command to view a list of LUNs with DM-Multipath that are mapped to the Linux host.

Step

1. Enter the following command on the Linux host console:

   `sanlun lun show all`

   Example

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

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

   Note: The output is modified to fit the screen.

FC: Displaying multipath information

You can use the `sanlun` command to display multipath information for native DM-Multipath.

Step

1. Enter the following command on the Linux host console:

   `sanlun lun show -p`

Example

The following output shows that DM-Multipath is configured.

```
# sanlun lun show -p
f3020-92:/vol/vol1/lun2 (LUN 1)         Lun state: GOOD
Lun Size:    3.6g (3879731200)   Controller_CF_State: Cluster Enabled
Protocol: FCP           Controller Partner: f3020-93
DM-MP DevName: mpath3   (360a9800043346852563444717a513571) dm-1
Multipath-provider: NATIVE
------ ---------- ------- -------- ----------- ----------
sanlun Controller                   Primary     Partner
path    Path     /dev/   Host       Controller  Controller
state   type     node    HBA        port        port
------ ---------- ------- -------- ----------- ----------
 UP    primary     sde    host17      1b          --
 UP    primary     sdg    host18      1a          --
 UP    secondary   sdc    host17      --          1a
 UP    secondary   sdi    host18      --          1b
```

Note: You should use `/dev/mapper/`*`device`* on a Linux host console to create the file system on a LUN.

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

FC: Removing the LUN

After you unmap the LUN from the controller to the Linux host, you can remove the LUN by running the SCSI rescan script. The rescan script is available with the sg3_utils package.

Before you begin

Before unmapping the LUN, ensure that the LUN is not in use.

Step

1. Enter the following command to remove the LUN:

   `rescan-scsi-bus.sh –r`

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4, SUSE Linux Enterprise Server 10 SP2, and SUSE Linux Enterprise Server 11. For earlier versions, see the HBA vendor-specific documentation.

Software iSCSI: Discovering new LUNs

After you create a new LUN and map it to the Linux host, you can discover the LUN using the software iSCSI or SCSI rescan script with DM-Multipath. The rescan script is available with the sg3_utils package.

Next topics

Red Hat Enterprise Linux 5 series: Discovering new LUNs on page 84
Red Hat Enterprise Linux 4 series: Discovering new LUNs on page 85
SUSE Linux Enterprise Server 10 series and 11: Discovering new LUNs on page 85
SUSE Linux Enterprise Server 9 series: Discovering new LUNs on page 86

Red Hat Enterprise Linux 5 series: Discovering new LUNs

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

Step

1. To discover a new LUN, enter the following command:

If you want to...	Enter the following command...
Obtain the list of all the current sessions	**`iscsiadm -m session`**
Rescan a specific session	**`iscsiadm -m session --sid=N --rescan`** N is the specific session ID.
Rescan all the sessions	**`iscsiadm -m session --rescan`**
Rescan using the SCSI rescan script **Note:** The SCSI rescan script is available from Red Hat Enterprise Linux 5 Update 4 and later.	**`rescan-scsi-bus.sh`**

You can use the `sanlun` command in Red Hat Enterprise Linux 5 series to verify that the new LUNs are discovered. After the LUNs are discovered, they are automatically added to the DM-Multipath configuration. To view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration.

Attention: You can mount the new file system by adding an entry to `/etc/fstab` as follows:

***device mount_point type* _netdev 0 0**

`device` is the name of the device in the `/dev/mapper/` directory.

`mount_point` is the mountpoint 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 such as iSCSI.

Red Hat Enterprise Linux 4 series: 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 using software iSCSI on Red Hat Enterprise Linux 4 series.

Steps

1. Enter the following command on the Linux host console to reload the iSCSI service:

   /etc/init.d/iscsi reload

2. Use the sanlun or iscsi-ls command to verify that the new LUNs are discovered.

   After the LUNs are discovered, they are automatically added to the DM-Multipath configuration. To view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration.

   Attention: You can mount the new file system by adding an entry to /etc/fstab as follows:

   ***device mount_point type* _netdev 0 0**

   device is the name of the device in the /dev/mapper/ directory.

   mount_point is the mountpoint 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 such as iSCSI.

SUSE Linux Enterprise Server 10 series and 11: Discovering new LUNs

If you create a new LUN and map it to the Linux host, you can discover the LUN by rescanning the iSCSI service on the host. To view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration using software iSCSI on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11.

Step

1. To discover a new LUN, perform the appropriate action:

If you want to...	Enter the following command...
Obtain the list of all the current sessions	**iscsiadm -m session**

If you want to...	Enter the following command...
Rescan a specific session	**`iscsiadm -m session --sid=N --rescan`** N is the specific session ID.
Rescan all the sessions	**`iscsiadm -m session --rescan`**
Rescan using the SCSI rescan script **Note:** The SCSI rescan script is available in SUSE Linux Enterprise Server 10 SP2 and SUSE Linux Enterprise Server 11.	**`rescan-scsi-bus.sh`**

You can verify that the new LUNs are discovered by using the `sanlun` or `iscsiadm` command.

Note:

- Before mounting the file system, modify the network configuration file and change the value of STARTMODE to `nfsroot`. This value is required for the successful unmounting of the iSCSI devices during system halts.
- You can mount the new file system by adding an entry to `/etc/fstab` for SUSE Linux Enterprise Server 11 as follows:

  ***device mount_point type* _netdev,nofail,rw 0 0**

  `device` is the name in the `/dev/mapper` directory.

  `mount_point` is the mountpoint you created for the file system.

  `type` is the file system type
- You can mount the new file system by adding an entry to `/etc/fstab` for SUSE Linux Enterprise Server 10 series as follows:

  ***device mount_point type* hotplug 0 0**

  `device` is the name in the `/dev/mapper` directory.

  `mount_point` is the mountpoint you created for the file system.

  `type` is the file system type.

SUSE Linux Enterprise Server 9 series: 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 view newly created LUNs, you need to first discover them and then add them to the DM-Multipath configuration using software iSCSI on SUSE Linux Enterprise Server 9 series.

Step

1. To reload the iSCSI service, enter the following command on the Linux host console:

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

Note: You can mount the new file system by adding an entry to `/etc/fstab.iscsi` as follows:

```
device mount_point type defaults 0 0
```

`device` is the name in the `/dev/mapper` directory.

`mount_point` is the mountpoint you created for the file system.

`type` is the file system type.

Software iSCSI: Adding new LUNs to the DM-Multipath configuration

After you have discovered new LUNs, you need to add them to the DM-Multipath configuration using software iSCSI on SUSE Linux Enterprise Server 9 series.

Step

1. To add the discovered LUNs to the DM-Multipath configuration, enter the following command on the Linux host console:

   ```
   multipath
   ```

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

Software iSCSI: Viewing the list of LUNs

You can view the list of LUNs with DM-Multipath that are mapped to the Linux host by using various commands, depending on the Linux version of your host.

Step

1. To view the list of LUNs, perform the appropriate action:

If you are using...	Enter one of the following commands...
Red Hat Enterprise Linux 5 series	• `sanlun lun show all` • `iscsiadm -m session --sid=N -P 3`
Red Hat Enterprise Linux 4 series	• `sanlun lun show all` • `iscsi-ls -l`
SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11	• `sanlun lun show all` • `iscsiadm -m session --sid=N -P 3`

If you are using...	Enter one of the following commands...
SUSE Linux Enterprise Server 9 series	• `sanlun lun show all` • `iscsi-ls -l`

Next topics

Red Hat Enterprise Linux: Commands to view LUNs on page 88
SUSE Linux Enterprise Server 10 series and 11: Commands to view LUNs on page 90
SUSE Linux Enterprise Server 9 series: Commands to view LUNs on page 91

Red Hat Enterprise Linux: Commands to view LUNs

The commands you use to view LUN information differ depending on your Red Hat Enterprise Linux version. The output enables you to view the list of LUNs configured on your Linux host.

The sanlun command (Red Hat Enterprise Linux 4 series and 5 series)

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

The output of the `sanlun` command looks similar to the following:

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

Note: The output is modified to fit the screen.

The iscsi-ls command (Red Hat Enterprise Linux 4 series)

The `iscsi-ls` command lists storage system node names, IP addresses, and available LUNs.

The output of the `iscsi-ls` command looks similar to the following:

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

DEVICE DETAILS:
---------------
LUN ID : 0
  Vendor: NETAPP Model: LUN Rev: 0.2
```

```
  Type:   Direct-Access                    ANSI SCSI revision: 04
  page83 type3: 60a9800043346536534a344a794e6d59
  page80: 43346536534a344a794e6d590a
  Device: /dev/sde
```

The iscsiadm command (Red Hat Enterprise Linux 5)

The `iscsiadm` command lists the available storage systems and LUNs for a session with a specific session ID, *N*.

The output of the `iscsiadm` command looks similar to the following:

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

```
                        Attached scsi disk sdm      State: running
                scsi4 Channel 00 Id 0 Lun: 6
                        Attached scsi disk sdp      State: running
                scsi4 Channel 00 Id 0 Lun: 7
                        Attached scsi disk sdq      State: running
```

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

SUSE Linux Enterprise Server 10 series and 11: Commands to view LUNs

The commands you use to view LUN information differ depending on the version of your SUSE Linux Enterprise Server. The output enables you to view the list of LUNs configured on your Linux host.

The sanlun command

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

The output of the `sanlun` command looks similar to the following:

```
# sanlun lun show all
controller: lun-pathname device filename adapter protocol lun size lun
state
f3020-94: /vol/vol1/lun2    /dev/sdg     host6  iSCSI      4g      GOOD
f3020-95: /vol/vol1/lun95_3 /dev/sde     host5  iSCSI      4g      GOOD
```

Note: The output is modified to fit the screen.

The iscsiadm command

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

The output of the `iscsiadm` command looks similar to the following:

```
# iscsiadm -m session --sid=0 -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.netapp:sn.101180456
        Current Portal: 10.72.199.144:3260,1001
        Persistent Portal: 10.72.199.144:3260,1001
                **********
                Interface:
                **********
                Iface Name: default
                Iface Transport: tcp
                Iface Initiatorname: iqn.1996-04.de.suse:lnx.200.109
                Iface IPaddress: 10.72.200.109
                Iface HWaddress: default
                Iface Netdev: default
                SID: 1
                iSCSI Connection State: LOGGED IN
```

```
                iSCSI Session State: Unknown
Internal iscsid Session State: NO CHANGE
                ************************
                Negotiated iSCSI params:
                ************************
                HeaderDigest: CRC32C
                DataDigest: None
                MaxRecvDataSegmentLength: 131072
                MaxXmitDataSegmentLength: 65536
                FirstBurstLength: 65536
                MaxBurstLength: 65536
                ImmediateData: Yes
                InitialR2T: No
                MaxOutstandingR2T: 1
                ************************
                Attached SCSI devices:
                ************************
                Host Number: 47 State: running
                scsi47 Channel 00 Id 0 Lun: 2
                   Attached scsi disk sdj       State: running
                scsi47 Channel 00 Id 0 Lun: 1
                   Attached scsi disk sdf       State: running
                scsi47 Channel 00 Id 0 Lun: 0
                   Attached scsi disk sdb       State: running
```

SUSE Linux Enterprise Server 9 series: Commands to view LUNs

The commands you use to view LUN information differ depending on your SUSE Linux Enterprise Server version. The output enables you to view the list of LUNs configured on your Linux host.

The sanlun command

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

The output of the `sanlun` command looks similar to the following:

```
# sanlun lun show
controller: lun-pathname  device filename adapter protocol lun size lun
state
f3020-43: /vol/fvol1/lun122  /dev/sdb       host4   iSCSI   1g       GOOD
f3020-43: /vol/fvol6/lun1    /dev/sdc       host4   iSCSI   1g       GOOD
f3020-43: /vol/fvol1/lun122  /dev/sdd       host4   iSCSI   1g       GOOD
```

The iscsi-ls command

The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The output of the `iscsi-ls -l` command looks similar to the following:

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

```
****************************************************************
TARGET NAME             : iqn.1992-08.com.netapp:sn.101175385
TARGET ALIAS            :
HOST NO                 : 5
BUS NO                  : 1
TARGET ID               : 0
TARGET ADDRESS          : 10.61.208.5:3260
SESSION STATUS          : ESTABLISHED AT Tue Mar 28 16:03:15 2006
NO. OF PORTALS          : 1
PORTAL ADDRESS 1        : 10.61.208.5:3260,2
SESSION ID              : ISID 00023d000002 TSID 544
```

Software iSCSI: Displaying multipath information

You can use the `sanlun` command to display multipath information with DM-Multipath using software iSCSI.

Step

1. Enter the following command on the Linux host console:

   `sanlun lun show -p`

   Example

   The following output shows that DM-Multipath is configured.

```
# sanlun lun show -p
f3020-236-39:/vol/vol1/lun1 (LUN 6) Lun state: GOOD
Lun Size: 7g (7516192768) Controller_CF_State: Cluster
Enabled
Protocol: iSCSI Controller Partner: f3020-236-40
DM-MP DevName: 360a9800043346773435a48514d624367
dm-0
Multipath-provider: NATIVE
------ ---------- ------- -------- ----------- ----------
sanlun Controller                   Primary     Partner
path    Path       /dev/    Host    Controller  Controller
state   type       node     HBA     port        port
------ ---------- ------- -------- ------------ ---------
UP      iscsi      sdc     host4   10.72.236.69   --
UP      iscsi      sdd     host3   10.72.236.39   --
```

Software iSCSI: Removing the LUN

After you unmap the LUN from the controller to the Linux host, you can remove the LUN by running the SCSI rescan script. The rescan script is available with the sg3_utils package.

Before you begin

Before unmapping the LUN, ensure that the LUN is not in use.

Step

1. Enter the following command to remove the LUN:

   rescan-scsi-bus.sh –r

   Note: The SCSI rescan script is available only from Red Hat Enterprise Linux 5 Update 4 and SUSE Linux Enterprise Server 10 SP2 and later.

Hardware iSCSI: Discovering new LUNs

After you create a new LUN and map it to the Linux host, you can discover the LUN using SCSI rescan script with DM-Multipath. The rescan script is available with the sg3_utils package.

Step

1. Enter the following command to add a new LUN:

   rescan-scsi-bus.sh

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4 and later.

Hardware iSCSI: Viewing the list of LUNs

You can use the `sanlun` command to view a list of LUNs with DM-Multipath that are mapped to the Linux host.

Step

1. Enter the following command on the Linux host console:

   sanlun lun show all

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

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

Hardware iSCSI: Displaying multipath information

You can use the `sanlun` command to display multipath information with DM-Multipath using hardware iSCSI.

Step

1. Enter the following command on the Linux host console:

`sanlun lun show -p`

Example

The following output shows that DM-Multipath is configured.

```
# sanlun lun show -p
f3020-94:/vol/vol7/lun1 (LUN 0)            Lun state: GOOD
Lun Size:     9g (9663676416)   Controller_CF_State: Cluster Enabled
Protocol: iSCSI          Controller Partner: f3020-95
DM-MP DevName:            360a9800043346851436f4f48457a4e31    dm-0
Multipath-provider: NATIVE
------ ---------- ------- -------- -------------- ----------
sanlun Controller                    Primary       Partner
path   Path       /dev/    Host     Controller    Controller
state  type       node     HBA      port          port
------ ---------- ------- -------- -------------- ----------
UP       iscsi    sdb      host4    10.72.199.71    --
UP       iscsi    sda      host4    10.72.199.94    --
```

Hardware iSCSI: Removing the LUN

After you unmap the LUN from the controller to the Linux host, you can remove the LUN by running the SCSI rescan script. The rescan script is available with the sg3_utils package.

Before you begin

Before unmapping the LUN, ensure that the LUN is not in use.

Step

1. Enter the following command to remove the LUN:

   `rescan-scsi-bus.sh –r`

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4 later.

Accessing LUNs without multipathing software

Without multipathing software, the host cannot distinguish multiple LUNs from multiple paths to the same LUN. If you are not using multipathing software, you should limit each LUN to a single path. You should also use the standby cfmode to avoid having a secondary path visible on the storage system.

Next topics

Persistent identification of LUNs without multipathing software on page 95
FC: Viewing information about LUNs on page 95
Software iSCSI: Viewing information about LUNs on page 95

Software iSCSI: Automatic mounting of the file system on page 101

Persistent identification of LUNs without multipathing software

LUNs on the storage system are seen by the Linux host as SCSI devices. If you have more than one LUN, you need to have a persistent way of identifying each LUN. You can write a file system label to the LUN and mount the file system by that label.

Linux assigns a unique device name, such as `/dev/sda` or `/dev/sdb`, to each LUN. The device names are assigned as the LUNs are discovered. Because of variable network delays and other issues, the order in which LUNs are discovered cannot be predicted. Therefore, a LUN can be mapped to a different device name after a Linux reboot or restarting the iSCSI service.

FC: Viewing information about LUNs

You can view information about LUNs without multipathing software that are mapped to the Linux host by using the `sanlun` command.

Step

1. Enter the `sanlun` command to view the list of LUNs.

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

   #### Example

   The `sanlun` command displays information about the FC devices and the storage system LUNs they are mapped to, including LUN size and state.

   The output of the `sanlun` command looks similar to the following:

```
# sanlun lun show all
controller: lun-pathname device filename adapter protocol lun size lun
state
f3020-94: /vol/vol_1/lun_1 /dev/sdb1 host0       FCP        4g      GOOD
f3020-95: /vol/vol_1/lun_2 /dev/sdc1 host0       FCP        4g      GOOD
```

   Note: The output is modified to fit the screen.

Software iSCSI: Viewing information about LUNs

For software iSCSI, you can view information about LUNs without multipathing software that are mapped to the Linux host by using different commands.

Step

1. To view the list of LUNs, perform the appropriate action:

If you are using...	Enter one of the following commands...
Red Hat Enterprise Linux	• **`iscsiadm -m session -P 3`** • **`iscsi-ls -l`** • **`sanlun`**
SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11	• **`sanlun`** • **`iscsiadm -m session --sid=N -P 3`**
SUSE Linux Enterprise Server 9 series	• **`sanlun`** • **`iscsi-ls -l`**

Next topics

Red Hat Enterprise Linux: Commands to view LUNs on page 96
SUSE Linux Enterprise Server 10 series and 11: Commands to view LUNs on page 98
SUSE Linux Enterprise Server 9 series: Commands to view LUNs on page 100

Red Hat Enterprise Linux: Commands to view LUNs

You can use a list of commands to view LUN information on Red Hat Enterprise Linux using software iSCSI.

The sanlun command (Red Hat Enterprise Linux 4 series and 5 series)

The `sanlun` command is included in the iSCSI Linux Host Utilities package.

```
# /opt/netapp/santools/sanlun lun show all
controller: lun-pathname device filename adapter protocol lun size lun
state
f3020-94: /vol/vol_1/lun_1  /dev/sdb1  <unknown>0  iSCSI    4g      GOOD
f3020-95: /vol/vol_1/lun_2  /dev/sdc1  <unknown>0  iSCSI    4g      GOOD
```

Note: The output is modified to fit the screen.

The iscsiadm command (Red Hat Enterprise Linux 5 series)

The `iscsiadm` command displays the storage system and LUN details.

The console lists the available storage systems and LUNs for a session with a specific session ID, as shown by the following example:

```
# 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
```

To view the details of all the sessions, use the `iscsiadm -m session -P 3` command.

The iscsi-ls command (Red Hat Enterprise Linux 4 series)

The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The console lists the available storage systems and LUNs, as shown by the following example:

```
# /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
*************************************************************************
****
```

SUSE Linux Enterprise Server 10 series and 11: Commands to view LUNs

You can use a list of commands to view LUN information on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11 using software iSCSI without multipathing software.

The sanlun command

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

The `sanlun` command displays information about the iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. The output looks similar to the following:

```
# /opt/netapp/santools/sanlun lun show all
controller: lun-pathname device filename adapter protocol lun size lun
state
```

```
f3020-94: /vol/vol1/lun2    /dev/sdg      host6   iSCSI      4g      GOOD
f3020-95: /vol/vol1/lun95_3 /dev/sde      host5   iSCSI      4g      GOOD
```

The iscsiadm command

The `iscsiadm` command lists the available storage systems and LUNs for a session with a specific session ID, *N*.

- SUSE Linux Enterprise Server 10 SP1
  The `iscsiadm -m session --sid=N -i` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.
- SUSE Linux Enterprise Server 10 SP2
  The `iscsiadm -m session --sid=N -P 3` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.

The output looks similar to the following:

```
# iscsiadm -m session --sid=N -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.netapp:sn.101180456
        Current Portal: 10.72.199.144:3260,1001
        Persistent Portal: 10.72.199.144:3260,1001
                **********
                Interface:
                **********
                Iface Name: default
                Iface Transport: tcp
                Iface Initiatorname:iqn.1996-04.de.suse:lnx.200.109
                Iface IPaddress: 10.72.200.109
                Iface HWaddress: default
                Iface Netdev: default
                SID: 1
                iSCSI Connection State: LOGGED IN
                iSCSI Session State: Unknown
                Internal iscsid Session State: NO CHANGE
                ************************
                Negotiated iSCSI params:
                ************************
                HeaderDigest: CRC32C
                DataDigest: None
                MaxRecvDataSegmentLength: 131072
                MaxXmitDataSegmentLength: 65536
                FirstBurstLength: 65536
                MaxBurstLength: 65536
                ImmediateData: Yes
                InitialR2T: No
                MaxOutstandingR2T: 1
                ************************
                Attached SCSI devices:
                ************************
                Host Number: 47 State: running
                scsi47 Channel 00 Id 0 Lun: 2
                   Attached scsi disk sdj       State: running
                scsi47 Channel 00 Id 0 Lun: 1
```

```
                   Attached scsi disk sdf       State: running
                scsi47 Channel 00 Id 0 Lun: 0
                   Attached scsi disk sdb       State: running
```

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

SUSE Linux Enterprise Server 9 series: Commands to view LUNs

You can use a list of commands to view LUN information on SUSE Linux Enterprise Server 9 series using software iSCSI without multipathing software.

The iscsi-ls -l command

The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The console lists the available storage systems and LUNs. The following example shows the output for SUSE Linux Enterprise Server 9 SP3.

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

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

The sanlun command

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

The output looks similar to the following:

```
# /opt/netapp/santools/sanlun lun show all
controller: lun-pathname  device filename adapter protocol lun size lun
state
```

```
f3020-43: /vol/fvol1/lun122  /dev/sdb      host4    iSCSI    1g       GOOD
f3020-43: /vol/fvol6/lun1    /dev/sdc      host4    iSCSI    1g       GOOD
```

The iscsi-ls command

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

The output looks similar to the following:

```
*****************************************************************
        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
```

Software iSCSI: Automatic mounting of the file system

The Linux host using Red Hat Enterprise Linux without multipathing software cannot mount file systems on iSCSI LUNs when it first boots up, because the networking and iSCSI driver resources are not yet available.

You can 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/mpt1_lun0 ext3  _netdev    0        0
```

Accessing LUNs with Veritas DMP

You can configure the host to access LUNs by using Veritas DMP with both the iSCSI and FC protocols.

Next topics

HBA and driver parameter settings on Veritas Storage Foundation on page 103
Accessing LUNs using Veritas DMP on page 103
LUN access using Veritas DMP with FC on page 104
LUN access using Veritas DMP with software iSCSI on page 107
Veritas DMP restore daemon values for Red Hat Enterprise Linux on page 118
Veritas DMP restore daemon values for SUSE Linux Enterprise Server on page 118
Displaying available paths using VxVM on page 118
Obtaining multipathing information by using the sanlun command on page 120

HBA and driver parameter settings on Veritas Storage Foundation

You need to set either bundled or external HBA driver parameters for Veritas Storage Foundation to ensure that the multipathing and storage system failover work correctly.

When the HBA driver is started, it scans the storage systems and discovers all mapped LUNs. If you configure Veritas DMP, multipath devices are created for all the LUNs that are discovered.

Related tasks

Setting the QLogic HBA parameters for FC on page 29
Loading a bundled or an external HBA driver on a Linux host using FC on page 35

Accessing LUNs using Veritas DMP

To access LUNs by using Veritas DMP, you should configure the restore daemon and verify that the daemon value has been successfully configured.

Steps

1. Configure the vxdmp restore daemon with an interval of 60 by entering the following command:

   ```
   vxdmptune dmp_restore_interval 60
   ```

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

Alternatively, you can use the following command to configure the daemon value.

```
vxdmpadm stop restore
```

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

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

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

```
vxdmpadm stat restored
```

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

LUN access using Veritas DMP with FC

You can configure the bundled HBA or external driver and the parameters, and access LUNs on your host by using Veritas DMP.

Next topics

Veritas: Discovering new LUNs with FC on page 104
Veritas: Viewing the list of LUNs with FC on page 105
Veritas: Displaying multipath information with FC on page 105
Veritas: Removing the LUN with FC on page 106

Veritas: Discovering new LUNs with FC

After you create a new LUN and map it to the Linux host, you can discover the LUN using the SCSI rescan script. The rescan script is available with the sg3_utils package

Step

1. Enter the following command to discover a new LUN:

```
rescan-scsi-bus.sh
```

```
vxdisk scandisks
```

Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4, SUSE Linux Enterprise Server 10 SP2, and SUSE Linux Enterprise Server 11. For earlier versions, see the HBA vendor-specific documentation.

Veritas: Viewing the list of LUNs with FC

You can use the `sanlun` command to view a list of LUNs with Veritas DMP that are mapped to the Linux host.

Steps

1. Enter the following command on the Linux host console:

   `sanlun lun show all`

   Example

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

   ```
   # sanlun lun show all
   controller:     lun              device    adapter protocol lun   lun
                   -pathname        filename                   size  state
   fas2020-201-42: /vol/vol1/lun0 /dev/sdix  host3   FCP       500m  GOOD
   fas2020-201-42: /vol/vol1/lun1 /dev/sdiy  host3   FCP       500m  GOOD
   fas2020-201-42: /vol/vol1/lun2 /dev/sdiz  host3   FCP       500m  GOOD
   fas2020-201-42: /vol/vol1/lun3 /dev/sdja  host3   FCP       500m  GOOD
   ```

   Note: The output is modified to fit the screen.

2. Enter the `vxdisk list` command to view a list of LUNs.

   ```
   # vxdisk list
   DEVICE      TYPE         DISK       GROUP   STATUS
   fas20200_0  auto:cdsdisk data_dg01  data_dg online thinrclm shared
   fas20200_1  auto:cdsdisk data_dg02  data_dg online thinrclm shared
   fas20200_2  auto:cdsdisk data_dg113 data_dg online thinrclm shared
   fas20200_3  auto:cdsdisk data_dg180 data_dg online thinrclm shared
   ```

Veritas: Displaying multipath information with FC

You can use the `sanlun` command to display multipath information for Veritas DMP.

Step

1. Enter the following command on the Linux host console:

   `sanlun lun show -p`

   Example

   The following output shows that Veritas DMP is configured.

   ```
   #  sanlun lun show -p
   fas2020-201-42:/vol/vol1/lun2  (LUN 2)            Lun state: GOOD
   Lun Size:    500m (524288000)    Controller_CF_State: Cluster Enabled
   Protocol: FCP           Controller Partner: fas2020-201-43
   DMP Nodename: fas20200_0
   Multipath-provider: VERITAS
   ```

```
--------  --------- ---------- ------- -------- ---------- -----------
----------
Veritas   sanlun    Controller                  Primary     Partner
Veritas
path       path      Path       /dev/    Host   Controller  Controller
IO
state     state      type       node     HBA    port        port
Policy
--------  --------- ---------- ------- -------- ---------- -----------
----------
ENABLED(A) GOOD      PRIMARY    sdiz     host3   0a           --
MinimumQ
ENABLED    GOOD      SECONDARY  sdd      host3   --           0a
MinimumQ
ENABLED(A) GOOD      PRIMARY    sdnk     host4   0b           --
MinimumQ
ENABLED    GOOD      SECONDARY  sdacr    host4   --           0b
MinimumQ
```

Veritas: Removing the LUN with FC

After you unmap the LUN from the controller to the Linux host, you can remove the LUN by running the SCSI rescan script. The rescan script is available with the sg3_utils package.

Before you begin

Before unmapping the LUN, ensure that the LUN is not in use.

Steps

1. Remove the devices from the control of Veritas Volume Manager by entering the following command:

   `vxdisk rm device`

   `device` is the name of the LUN that is being removed.

   Example

   ```
   # vxdisk rm fas30200_1
   ```

2. Unmap the LUN from the controller to the Linux host.
3. Update the logical unit configuration of the host by entering the following command:

   `rescan-scsi-bus.sh -r`

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4, SUSE Linux Enterprise Server 10 SP2, and SUSE Linux Enterprise Server 11. For earlier versions, see the HBA vendor-specific documentation.

4. Initiate a rescan of the OS device tree by Veritas Volume Manager:

   `vxdisk scandisks`

LUN access using Veritas DMP with software iSCSI

You can configure software iSCSI to access LUNs on your host by using Veritas DMP.

Next topics

Veritas: Viewing the list of LUNs with software iSCSI on page 107
Veritas: Discovering new LUNs with iSCSI on page 114
Veritas: Displaying multipath information with software iSCSI on page 116
Veritas: Removing the LUN with software iSCSI on page 117

Related concepts

Configuration of the software iSCSI initiator on page 45

Veritas: Viewing the list of LUNs with software iSCSI

You can view the list of LUNs with Veritas DMP that are mapped to the Linux host by using various commands, depending on the Linux version of your host.

Step

1. To view the list of LUNs, perform the appropriate action:

If you are using...	Enter one of the following commands...
Red Hat Enterprise Linux 5 series	• `sanlun lun show all` • `iscsiadm -m session --sid=N -P 3` • `vxdisk list`
Red Hat Enterprise Linux 4 series	• `sanlun lun show all` • `iscsi-ls -l` • `vxdisk list`
SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11	• `sanlun lun show all` • `iscsiadm -m session --sid=N -P 3` • `vxdisk list`
SUSE Linux Enterprise Server 9 series	• `sanlun lun show all` • `iscsi-ls -l` • `vxdisk list`

Next topics

Veritas: Commands to view LUNs for Red Hat Enterprise Linux with software iSCSI on page 108
Veritas: Commands to view LUNs for SUSE Linux Enterprise Server 10 series and 11 with software iSCSI on page 110
Veritas: Commands to view LUNs for SUSE Linux Enterprise Server 9 series with software iSCSI on page 112

Veritas: Commands to view LUNs for Red Hat Enterprise Linux with software iSCSI

You can use a list of commands to view LUN information on Red Hat Enterprise Linux using software iSCSI.

The sanlun command (Red Hat Enterprise Linux 4 series and 5 series)

The `sanlun` command is included in the iSCSI Linux Host Utilities package.

```
# /opt/netapp/santools/sanlun lun show all
controller:     lun-pathname   device filename  adapter  protocol   lun
size   lun state
fas6030-201-71: /vol/vol1/lun1 /dev/sdawh       host3    iSCSI
200m     GOOD
fas6030-201-71: /vol/vol1/lun2 /dev/sdawd       host3    iSCSI
200m     GOOD
fas6030-201-71: /vol/vol1/lun3 /dev/sdawf       host3    iSCSI
200m     GOOD
fas6030-201-71: /vol/vol1/lun4 /dev/sdawg       host3    iSCSI
200m     GOOD
```

Note: The output is modified to fit the screen.

The iscsiadm command (Red Hat Enterprise Linux 5 series)

The `iscsiadm` command displays the storage system and LUN details.

The console lists the available storage systems and LUNs for a session with a specific session ID, as shown by the following example:

```
# 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
```

To view the details of all the sessions, use the `iscsiadm -m session -P 3` command.

The iscsi-ls command (Red Hat Enterprise Linux 4 series)

The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The console lists the available storage systems and LUNs, as shown by the following example:

```
# /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
type1:4e45544150502020204c554e204f6443745359763759367733000000000000
  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
*************************************************************************
****
```

Veritas: Commands to view LUNs for SUSE Linux Enterprise Server 10 series and 11 with software iSCSI

You can use a list of commands to view LUN information on SUSE Linux Enterprise Server 10 series and SUSE Linux Enterprise Server 11 using software iSCSI without multipathing software.

The sanlun command

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

The `sanlun` command displays information about the iSCSI devices and the storage system LUNs they are mapped to, including LUN size and state. The output looks similar to the following:

```
# /opt/netapp/santools/sanlun lun show all
controller:     lun-pathname    device filename  adapter  protocol lun
size  lun state
fas6030-201-71: /vol/vol1/lun1  /dev/sdawh       host3    iSCSI
200m       GOOD
fas6030-201-71: /vol/vol1/lun2  /dev/sdawd       host3    iSCSI
200m       GOOD
fas6030-201-71: /vol/vol1/lun3  /dev/sdawf       host3    iSCSI
200m       GOOD
```

```
fas6030-201-71: /vol/vol1/lun4  /dev/sdawg       host3    iSCSI
200m       GOOD
```

The iscsiadm command

The `iscsiadm` command lists the available storage systems and LUNs for a session with a specific session ID, *N*.

- SUSE Linux Enterprise Server 10 SP1
  The `iscsiadm -m session --sid=N -i` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.
- SUSE Linux Enterprise Server 10 SP2
  The `iscsiadm -m session --sid=N -P 3` command lists information about the available storage systems and LUNs for a session with a specific session ID, *N*.

The output looks similar to the following:

```
# iscsiadm -m session --sid=N -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.netapp:sn.101180456
        Current Portal: 10.72.199.144:3260,1001
        Persistent Portal: 10.72.199.144:3260,1001
                **********
                Interface:
                **********
                Iface Name: default
                Iface Transport: tcp
                Iface Initiatorname:iqn.1996-04.de.suse:lnx.200.109
                Iface IPaddress: 10.72.200.109
                Iface HWaddress: default
                Iface Netdev: default
                SID: 1
                iSCSI Connection State: LOGGED IN
                iSCSI Session State: Unknown
                Internal iscsid Session State: NO CHANGE
                ************************
                Negotiated iSCSI params:
                ************************
                HeaderDigest: CRC32C
                DataDigest: None
                MaxRecvDataSegmentLength: 131072
                MaxXmitDataSegmentLength: 65536
                FirstBurstLength: 65536
                MaxBurstLength: 65536
                ImmediateData: Yes
                InitialR2T: No
                MaxOutstandingR2T: 1
                ************************
                Attached SCSI devices:
                ************************
                Host Number: 47 State: running
                scsi47 Channel 00 Id 0 Lun: 2
                   Attached scsi disk sdj       State: running
                scsi47 Channel 00 Id 0 Lun: 1
```

```
                   Attached scsi disk sdf       State: running
                scsi47 Channel 00 Id 0 Lun: 0
                   Attached scsi disk sdb       State: running
```

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

Veritas: Commands to view LUNs for SUSE Linux Enterprise Server 9 series with software iSCSI

You can use a list of commands to view LUN information on SUSE Linux Enterprise Server 9 series using software iSCSI without multipathing software.

The iscsi-ls -l command

The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The console lists the available storage systems and LUNs. The following example shows the output for SUSE Linux Enterprise Server 9 SP3.

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

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

The sanlun command

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

The output looks similar to the following:

```
# /opt/netapp/santools/sanlun lun show all
controller:    lun-pathname    device filename  adapter  protocol lun size
```

```
lun state
fas6030-201-71: /vol/vol1/lun1  /dev/sdawh       host3    iSCSI
200m       GOOD
fas6030-201-71: /vol/vol1/lun2  /dev/sdawd       host3    iSCSI
200m       GOOD
fas6030-201-71: /vol/vol1/lun3  /dev/sdawf       host3    iSCSI
200m       GOOD
fas6030-201-71: /vol/vol1/lun3  /dev/sdawg       host3    iSCSI
200m       GOOD
```

The iscsi-ls command

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

The output looks similar to the following:

```
*****************************************************************
        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
```

Veritas: Discovering new LUNs with iSCSI

After you create a new LUN and map it to the Linux host, you can discover the LUN using the software iSCSI or SCSI rescan script. The rescan script is available with the sg3_utils package.

Next topics

Veritas: Discovering new LUNs for Red Hat Enterprise Linux 5 series with software iSCSI on page 114

Veritas: Discovering new LUNs for Red Hat Enterprise Linux 4 series with software iSCSI on page 115

Veritas: Discovering new LUNs for SUSE Linux Enterprise Server 10 series and 11 with software iSCSI on page 115

Veritas: Discovering new LUNs for SUSE Linux Enterprise Server 9 series with software iSCSI on page 116

Veritas: Discovering new LUNs for Red Hat Enterprise Linux 5 series with software iSCSI

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 using software iSCSI on Red Hat Enterprise Linux 5 series.

Step

1. To discover a new LUN, enter the following command:

If you want to...	Enter the following command...
Obtain the list of all the current sessions	**`iscsiadm -m session`**
Rescan a specific session	**`iscsiadm -m session --sid=N --rescan`** N is the specific session ID.
Rescan all the sessions	**`iscsiadm -m session --rescan`**
Rescan using the SCSI rescan script **Note:** The SCSI rescan script is available from Red Hat Enterprise Linux 5 Update 4 and later.	**`rescan-scsi-bus.sh`**
Rescan Veritas devices	**`vxdisk scandisks`**

You can use the `sanlun` command in Red Hat Enterprise Linux 5 series to verify that the new LUNs are discovered. After the LUNs are discovered, they are automatically added to the Veritas DMP configuration.

Veritas: Discovering new LUNs for Red Hat Enterprise Linux 4 series with software iSCSI

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 using software iSCSI on Red Hat Enterprise Linux 4 series.

Steps

1. Enter the following command on the Linux host console to reload the iSCSI service:

   `/etc/init.d/iscsi reload`

2. Rescan the Veritas devices using the following command:

   `vxdisk scandisks`

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

   After the LUNs are discovered, they are automatically added to the Veritas DMP configuration.

Veritas: Discovering new LUNs for SUSE Linux Enterprise Server 10 series and 11 with software iSCSI

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. To view newly created LUNs, you need to first discover them and then add them to the Veritas DMP configuration using software iSCSI on SUSE Linux Enterprise Server 10 series and 11.

Step

1. To discover a new LUN, perform the following step:

If you want to...	Enter the following command...
Obtain the list of all the current sessions	**`iscsiadm -m session`**
Rescan a specific session	**`iscsiadm -m session --sid=N --rescan`** N is the specific session ID.
Rescan all the sessions	**`iscsiadm -m session --rescan`**
Rescan using the SCSI rescan script **Note:** The SCSI rescan script is available from SUSE Linux Enterprise Server 10 SP2 and later.	**`rescan-scsi-bus.sh`**
Rescan Veritas device	**`vxdisk scandisks`**

You can verify that the new LUNs are discovered by using the `sanlun` or `iscsiadm` command.

Note:

- Before mounting the file system, modify the network configuration file and change the value of STARTMODE to `nfsroot`. This value is required for the successful unmounting of the iSCSI devices during system halts.
  For SUSE Linux Enterprise Server 10 series, modify `/etc/sysconfig/network/ifcfg-eth-id`*`xxxxxxxxx`*.

Veritas: Discovering new LUNs for SUSE Linux Enterprise Server 9 series with software iSCSI

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 view newly created LUNs, you need to first discover them and then add them to the Veritas DMP configuration using software iSCSI on SUSE Linux Enterprise Server 9 series.

Steps

1. To reload the iSCSI service, enter the following command on the Linux host console:

   `/etc/init.d/iscsi reload`

2. Rescan the Veritas devices using the following command:

   `vxdisk scandisks`

Veritas: Displaying multipath information with software iSCSI

You can use the `sanlun` command to display multipath information with Veritas DMP using software iSCSI.

Step

1. Enter the following command on the Linux host console:

   `sanlun lun show -p`

   Example

   The following output shows that Veritas DMP is configured.

```
# sanlun lun show -p
fas6030-201-71:/vol/vol1/lun92 (LUN 92)           Lun state: GOOD
Lun Size:    200m (209715200)    Controller_CF_State: Cluster Enabled
Protocol: iSCSI         Controller Partner: fas6030-201-70
DMP Nodename: fas60300_597
Multipath-provider: VERITAS
--------- ------- ---------- ------- ------- -------------- ---------
----------
Veritas   sanlun  Controller                  Primary
```

```
Partner     Veritas
path      path    Path        /dev/   Host    Controller
Controller  IO
state     state   type        node    HBA     port
port        Policy
--------- ------- ---------- ------- ------- --------------- ----------
----------
ENABLED(A) GOOD    -          sdne    host5   192.168.200.71
--        MinimumQ
ENABLED(A) GOOD    -          sdmi    host4   192.168.200.96
--        MinimumQ
ENABLED(A) GOOD    -          sdaop   host3   10.72.201.71
--        MinimumQ
------------------------------------------------------------------------
----------
```

Veritas: Removing the LUN with software iSCSI

After you unmap the LUN from the controller to the Linux host, you can remove the LUN by running the SCSI rescan script. The rescan script is available with the sg3_utils package.

Before you begin

Before unmapping the LUN, ensure that the LUN is not in use.

Steps

1. Remove the devices from the control of Veritas Volume Manager by entering the following command:

   `vxdisk rm device`

   `device` is the name of the LUN that is being removed.

   Example

   ```
   # vxdisk rm fas30200_1
   ```

2. Unmap the LUN from the controller to the Linux host.
3. Update the logical unit configuration of the host by entering the following command:

   `rescan-scsi-bus.sh -r`

   Note: The rescan script is available from Red Hat Enterprise Linux 5 Update 4, SUSE Linux Enterprise Server 10 SP2, and SUSE Linux Enterprise Server 11. For earlier versions, see the HBA vendor-specific documentation.

4. Initiate a rescan of the OS device tree by Veritas Volume Manager:

   `vxdisk scandisks`

Veritas DMP restore daemon values for Red Hat Enterprise Linux

You should configure the driver settings to use the required values.

The default versus required settings for Red Hat Enterprise Linux 4 and Red Hat Enterprise Linux 5 are as follows:

Component	**Parameter**	**Red Hat Enterprise Linux 4**		**Red Hat Enterprise Linux 5**	
		Default setting	**Required setting**	**Default setting**	**Required setting**
Veritas DMP	dmp_restore_interval	300	60	300	60

Veritas DMP restore daemon values for SUSE Linux Enterprise Server

You can configure the driver settings to use the required values.

The default versus required settings for SUSE Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10 are as follows:

Component	**Parameter**	**SUSE Linux Enterprise Server 9**		**SUSE Linux Enterprise Server 10**	
		Default setting	**Required setting**	**Default setting**	**Required setting**
Veritas DMP	dmp_restore_interval	300	60	300	60

Displaying available paths using VxVM

The VxVM management interface displays the vxdisk device, type, disk, group, and status.

Steps

1. Enter the following command to view all devices:

   ```
   vxdisk list
   ```

Example
The output of the `vxdisk list` command is similar to the following:

```
# vxdisk list
DEVICE         TYPE            DISK         GROUP             STATUS
disk_0       auto:none       -            -        online invalid
fas31700_0   auto:cdsdisk    data_dg01    data_dg  online thinrclm
shared
fas31700_1   auto:cdsdisk    data_dg02    data_dg  online thinrclm
shared
fas31700_2   auto:cdsdisk    data_dg08    data_dg  online thinrclm
shared
fas31700_3   auto:cdsdisk    data_dg09    data_dg  online thinrclm
shared
fas31700_4   auto:cdsdisk    data_dg10    data_dg  online thinrclm
shared
fas31700_5   auto:cdsdisk    data_dg11    data_dg  online thinrclm
shared
fas31700_6   auto:cdsdisk    data_dg12    data_dg  online thinrclm
shared
fas31700_7   auto:cdsdisk    data_dg13    data_dg  online thinrclm
shared
```

Note: For Veritas Storage Foundation 5.0 MP1 and MP2, the ASL displays the enclosure-based naming disk objects in uppercase.

For Veritas Storage Foundation 5.0 MP3 and Veritas Storage Foundation 5.1, the default behavior of the ASL is to display the enclosure-based naming disk objects in lowercase.

You can change the enclosure names to uppercase by using the `vxddladm set namingscheme=ebn lowercase=no` command.

2. On the host console, enter the following command to display path information for the device you want:

   `vxdmpadm getsubpaths dmpnodename=`*`device`*

   `device` is the name listed under the output of the `vxdisk list` command.

   Example
   The output of the `vxdmpadm getsubpaths dmpnodename=`*`device`* command is similar to the following:

```
vxdmpadm getsubpaths dmpnodename=fas31700_1
NAME   STATE[A]   PATH-TYPE[M] CTLR-NAME ENCLR-TYPE  ENCLR-NAME  ATTRS
=====================================================================
sdb    ENABLED(A)    -          c7        FAS3170     fas31700       -
sdq    ENABLED(A)    -          c8        FAS3170     fas31700       -
```

3. To obtain path information for a host HBA, enter the following command:

   `vxdmpadm getsubpaths ctlr=`*`controller_name`*

   `controller_name` is the controller displayed under “CTRL-NAME” in the output of the `vxdmpadm getsubpaths dmpnodename` command.

The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to.

Example

The output of the `vxdmpadm getsubpaths ctlr=` *`controller_name`* command is similar to the following:

```
#vxdmpadm getsubpaths ctlr=c7

NAME  STATE[A]   PATH-TYPE[M] DMPNODENAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
==========================================================================
sdj   ENABLED(A)    -          fas31700_0   FAS3170    fas31700    -
sdb   ENABLED(A)    -          fas31700_1   FAS3170    fas31700    -
sdl   ENABLED(A)    -          fas31700_10  FAS3170    fas31700    -
sdm   ENABLED(A)    -          fas31700_11  FAS3170    fas31700    -
sdn   ENABLED(A)    -          fas31700_12  FAS3170    fas31700    -
```

Obtaining multipathing information by using the sanlun command

You can obtain multipathing information to check the configuration settings by using the `sanlun` command.

Step

1. On the host command line, enter the following command:

   `sanlun lun show -p`

   `sanlun` displays path information for each LUN.

   Example

   The `sanlun lun show -p` command output is similar to the following:

```
# sanlun lun show -p
fas3170-201-78:/vol/nate_4_1_IBMx336_200_144/
nate_IBMx336-200-144_linux_10
LUN 10) Lun state: GOOD
Lun Size:      8g (8589934592)   Controller_CF_State: Cluster Enabled
Protocol: iSCSI         Controller Partner: fas3170-201-79
DMP Nodename: fas31700_10
Multipath-provider: VERITAS
------- ------- --------- ---------- ------- ------------ ------------
Veritas  sanlun  Controller          Primary      Partner     Veritas
path     path    Path    /dev/ Host  Controller   Controller  IO
state    state   type    node  HBA   port         port        Policy
------- ------- --------- ---------- ------- ------------ -------------
ENABLED(A) GOOD    -     sdaa  host8 10.72.201.78     --      MinimumQ
ENABLED(A) GOOD    -     sdl   host7 192.168.200.78   --      MinimumQ
```

FC: SAN boot configuration

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

Next topics

FC: Configuring SAN boot on Red Hat Enterprise Linux on page 121
FC: Configuring SAN boot on a SUSE Linux Enterprise Server on page 124

FC: Configuring SAN boot on Red Hat Enterprise Linux

You can configure SAN boot on your Red Hat Enterprise Linux host.

Steps

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

Next topics

FC: Installing the Red Hat Enterprise Linux operating system on page 121
FC: Installing and configuring the external or bundled HBA driver on Red Hat Enterprise Linux on page 122

Related tasks

Installing the Host Utilities software on page 43

FC: Installing the Red Hat Enterprise Linux operating system

You can install the Red Hat Enterprise Linux operating system on the SAN boot LUN configured on the storage controller.

Steps

1. On the storage system, select a cfmode setting.
2. Create a LUN on the storage system that acts as the SAN boot LUN, and map it to the host.

   You should ensure that the SAN boot LUN is mapped, and multiple paths to the LUN are available. You should also ensure that the SAN boot LUN is visible to the host during the boot process.

3. Enable the BIOS of the HBA port to which the SAN boot LUN is mapped.

   For information about how to enable the HBA BIOS, see your HBA vendor-specific documentation.

   Note: It is best to enable the spinup delay option for individual HBA ports in the HBA BIOS.

4. Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device.
5. Save and exit.

   For more information, see your vendor-specific documentation.

6. Reboot the host.
7. Install the operating system on the LUN configured on the storage controller.

   Note: You should specify Boot Option as `linux mpath` during the operating system installation. When you specify `linux mpath`, you can see the multipath devices (`/dev/mapper/mpathx`) as installation devices.

Related tasks

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43
Setting the QLogic HBA parameters for FC on page 29

FC: Installing and configuring the external or bundled HBA driver on Red Hat Enterprise Linux

You need to set the required HBA and driver parameters to install and configure the external or bundled HBA driver on your Linux host.

Before you begin

Ensure that you download the required HBA driver from the specific vendor Web site.

Steps

1. To install the external driver on Red Hat Enterprise Linux, perform the appropriate action:

If you are using...	Perform the following...
A QLogic external HBA driver	**a.** Run the `qlinstall` script with the `-up` option. **b.** Reboot the host to load the external QLogic driver.
An Emulex external HBA driver	Run the `lpfc-install` script to install the external Emulex driver on the SAN booted host.

2. Set the required HBA driver parameters.

If you are using...	Perform the following...
A QLogic external HBA driver	**a.** Edit the `/etc/modprobe.conf` file. **b.** Append the `options qla2xxx` statement with settings in the `modprobe.conf` file on a single line. **Note:** Only one `options qla2xxx` statement is allowed. **`ql2xsuspendcount=6 MaxRetriesPerPath=10`** **`MaxRetriesPerIo=10 ql2xfailover=0`** **c.** Save the file and exit. **d.** Enter the appropriate command for your HBA model to rebuild the initial ramdisk (initrd) image with the new parameter values: **`./qlinstall -br -in driver module`** Example **`./qlinstall -br qla2xxx`** **e.** Reboot the host.
A QLogic bundled HBA driver	**a.** In the `/etc/modprobe.conf` file, append the following line: **`options qla2xxx qlport_down_retry=30`** **b.** Re-create the initrd by running the following command: **`mkinitrd -f /boot/initrd-"`uname -r`".img `uname -r``** **Note:** It is best to back up the current initrd before running the `mkinitrd` command. **c.** Reboot the host. **d.** Verify that the value of the qlport_down_retry parameter in the corresponding `qla2xxx` directory in `/sys/module/qla2xxx/parameters/` matches the values in `modprobe.conf` file. Example: You can check the value of the qlport_down_retry on a Red Hat Enterprise Linux host by running the following command: **`cat /sys/module/qla2xxx/parameters/qlport_down_retry`** **Note:** There are no qla2300/qla2400/qla2500 firmware modules in Red Hat Enterprise Linux 5 series. These modules have been integrated into the base qla2xxx module.

If you are using...	Perform the following...
An Emulex external or bundled HBA driver	No changes are required to the HBA driver parameters.

Related concepts

DM-Multipath configuration on page 65

Related tasks

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43
Setting the QLogic HBA parameters for FC on page 29

FC: Configuring SAN boot on a SUSE Linux Enterprise Server

You can configure SAN boot with multipath on your SUSE Linux Enterprise Server host.

Steps

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

Next topics

FC: Installing the SUSE Linux Enterprise Server operating system on page 125
FC: Installing and configuring the external or bundled HBA driver on SUSE Linux Enterprise Server on page 125
FC: Configuring root partition with DM-Multipath on SUSE Linux Enterprise Server on page 127

Related tasks

Installing the Host Utilities software on page 43
Setting the QLogic bundled driver parameters for FC on page 30
Setting the QLogic HBA parameters for hardware iSCSI on page 33

FC: Installing the SUSE Linux Enterprise Server operating system

You can install the operating system on a LUN that is configured on the storage controller.

Steps

1. Create a LUN on the storage system that acts as the SAN boot LUN, and map it to the host.

   Note: You should ensure that the SAN boot LUN is mapped, and only a single primary path to the LUN is available. You should also ensure that the SAN boot LUN is visible to the host during boot process.

2. Enable the BIOS of the HBA port to which the SAN boot LUN is mapped.

   For information about how to enable the HBA BIOS, see your HBA vendor-specific documentation.

   Note: It is best to enable the `spinup delay` option for individual HBA ports in the HBA BIOS.

3. Install the operating system on the SAN boot LUN.

FC: Installing and configuring the external or bundled HBA driver on SUSE Linux Enterprise Server

You need to set the required HBA and external driver parameters to install and configure the external or bundled HBA driver on your Linux host.

Before you begin

Ensure that you download the required HBA driver from the specific vendor Web site.

Steps

1. To install the external driver on SUSE Linux Enterprise Server, perform the appropriate action:

If you are using...	Perform the following...
A QLogic external HBA driver	**a.** Run the `qlinstall` script with the `-up` option. **b.** Reboot the host to load the external QLogic driver.
An Emulex external HBA driver	Run the `lpfc-install` script to install the external Emulex driver on the SAN booted host.

2. Set the required HBA driver parameters.

If you are using...	Perform the following...
A QLogic external HBA driver	**a.** Edit the `/etc/modprobe.conf.local` file. **b.** Append the `options qla2xxx` statement with settings in the `modprobe.conf` file on a single line. **Note:** Only the `options qla2xxx` statement is allowed. **`ql2xsuspendcount=6 MaxRetriesPerPath=10`** **`MaxRetriesPerIo=10 ql2xfailover=0`** **c.** Save the file and exit.
A QLogic bundled HBA driver	**a.** In the `/etc/modprobe.conf.local` file, append the following line: **`options qla2xxx qlport_down_retry=30`** **b.** Re-create the initrd by running the `mkinitrd` command. **Note:** It is best to back up the current initrd before running the `mkinitrd` command. **c.** Reboot the host. **d.** Verify that the value of qlport_down_retry parameter in the corresponding `qla2xxx` directory in `/sys/module/qla2xxx/parameters/` matches the values in `modprobe.conf.local` file. Example: You can check the value of the qlport_down_retry on a SUSE Enterprise Linux Server host by running the following command: **`cat /sys/module/qla2xxx/parameters/qlport_down_retry`**
An Emulex external or bundled HBA driver	No changes are required to the HBA driver parameters.

Related tasks

Download of the external HBA driver software for FC on page 28
Installation of the external HBA driver software for FC on page 28

FC: Configuring root partition with DM-Multipath on SUSE Linux Enterprise Server

You can configure root partition with DM-Multipath on your SUSE Linux Enterprise Server host.

Steps

1. Configure DM-Multipath support.
2. Install the Linux Host Utilities.
3. Enable all paths to the LUN on the storage controller.
4. Re-create the initrd by running the following command:

If you are using...	Enter the following command...
SUSE Linux Enterprise Server 11	`mkinitrd -f multipath`
SUSE Linux Enterprise Server 10 SP2	`mkinitrd -f mpath`

Note: You must keep a backup of the initird image.

5. Reboot the host.

Note: Because all paths to the SAN boot LUN are now available on the host, you should configure all paths in the HBA boot BIOS as primary boot device, secondary boot device, and so on.

Related concepts

DM-Multipath configuration on page 65

Related tasks

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43

iSCSI: SAN boot configuration

SAN boot uses a SAN-attached disk, such as a LUN, as a root device for a host. SAN boot can be implemented either by an iSCSI HBA, or a network interface card (NIC) and software iSCSI stack.

- Software iSCSI
  For a software initiator to implement SAN boot, you can have the root device on an iSCSI LUN, and you can use any of the following options to load the kernel:
  - Using a host's locally attached disk (for storing kernel and initrd images)
  - Using a preboot execution environment (PXE) server
- Hardware iSCSI
  If you use the iSCSI HBA, because the protocol stack runs on the HBA, it is ready to communicate with the storage system and discover a LUN. You can have both the boot device and root device on an iSCSI LUN.

Next topics

Software iSCSI: Configuring SAN boot on Red Hat Enterprise Linux on page 129
Software iSCSI: Configuring SAN boot on SUSE Linux Enterprise Server on page 131
Hardware iSCSI: Configuring SAN boot on Red Hat Enterprise Linux on page 134

Software iSCSI: Configuring SAN boot on Red Hat Enterprise Linux

You can configure SAN boot on a Red Hat Enterprise Linux host using software iSCSI.

Steps

1. While initiating the installation, specify Boot Option as `linux mpath` and press Enter.
2. Proceed with the installation until you reach the storage configuration page and click **Advanced storage configuration**.
3. Select **Add iSCSI target** and click **Add drive**.
4. Enter the Target IP address and the proper iSCSI initiator name.

   Note: You should ensure that you associate this IQN with proper privileges on the storage controller.
5. On the storage controller, create an igroup with the initiator name that you provided in Step 4.
6. Create a LUN on the storage system on which you intend to create root partition, and map it to the igroup.

7. Return to the host screen.
8. Click **Add Target** in the **Configure iSCSI Parameters** window.

   When you add the target, the target portal is discovered.

   Note: You should ensure that multiple target portals are discovered, because the Red Hat installer does not identify the iSCSI device as a multipathed device unless it has more than one path.

9. To discover more target portals, repeat Step 2 through Step 8.

   You should now see a multipathed iSCSI device listed in the drives section.

   Note: If the iSCSI multipathed device is not listed, you should check the configuration.

10. Select a partitioning layout as `Create custom layout` and Click **Next**.

    You can now proceed with the installation process and enter choices until you reach the **Installation Summary** page.

11. In the **Partitioner** page, select the iSCSI multipathed device from the list of allowable drives where you want to install the root file system.
12. Create the root file system on the selected device and select the mountpoint as `/`.
13. Create a SWAP partition.

    Note: You can create a SWAP partition on the same LUN that contains the root partition or on a different LUN.

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

14. Create the /boot partition.

    You can create a /boot partition on a locally attached disk or use a PXE server to load the kernel boot image.

15. Click **Next** and follow the installation prompts to complete the installation.

Software iSCSI: Configuring SAN boot on SUSE Linux Enterprise Server

Implementing SAN boot (root partition on multipath) requires significant configuration and the process needs to be carefully managed. You can configure SAN boot on SUSE Linux Enterprise Server.

Before you begin

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

Steps

1. While initiating the installation, specify Boot Option as follows:

   ```
   linux withiscsi=1 netsetup=1
   ```

2. In the **iSCSI Initiator Overview** page, select the **Service** tab and enter the Target IP address and the proper iSCSI initiator name.

   Note: You should ensure that you associate this IQN with proper privileges on the storage controller.

3. On the storage controller, create an igroup with the initiator name that you provided in Step 2.
4. Create a LUN on the storage system on which you intend to create the root partition, and map it to the igroup.
5. Return to the host screen.
6. Select the **Connected Targets** tab, click **Add**.
7. In the **iSCSI Initiator Discovery** page, perform the following steps:
   a. Specify the IP address of the storage system.
   b. Specify the port. The default is 3260.
   c. Enter the credentials if you are using an authentication mechanism.
   d. Click **Next**.
8. In the list of storage systems that are discovered, click **Connect** for each, including authentication credentials if necessary.

   Note: During the installation, you should enable only one path to the root LUN.

   Click **Next**.
9. Verify that the value for Connected is true for all the targets and click **Next**.

The Connected Targets pane lists all the targets.

10. Set the Start-up mode to automatic by using the **Toggle Start-up** button, and click **Finish**.
11. In the **Installation Settings** page, select the **Expert** tab.
12. Click **Partitioning**.
13. Select the `Create Custom Partition Setup` option.

    Note: You can view the list of local disks and LUNs.

14. Select the `Custom Partitioning (for experts)` option.
15. In the **Expert Partitioner** page, select the LUN where you want to install the root file system.
16. Create the root file system on the selected LUN and select the mountpoint as `/`.
17. Click the **Fstab Options** button.
18. Select the `Mount by Device ID` option.
19. Ensure that you have the hotplug keyword in the Arbitrary Option Value text box, and click **OK**.
20. In the **Create a Primary Partition** page, click **OK**.
21. Create a SWAP partition.

    Note: You can create a SWAP partition on the same LUN that contains the root partition or on a different LUN.

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

22. Create the /boot partition.

    You can create a /boot partition on a locally attached disk or use a PXE server to load the kernel boot image.

23. After you return to the **Expert Partitioner** page, review the configuration. Click **Finish**.
24. In the **Installation Settings** page, click the **Booting** tab.
25. Select **Default Label**.
26. Click **Edit**.
27. For the Optional Kernel Command Line Parameter, ensure that all references to installer arguments are removed.

    The parameter should look similar to the following:
    `resume=/dev/sda1 splash=silent showopts`

    Click **OK**.

28. Click the **Boot Loader Installation** tab.

29. In the **Boot Loader Location** pane, select the `Boot from Master Boot Record` option.

    Click **Finish**. You are returned to the **Installation Settings** page.

30. Review the configuration settings.

    Click **Accept**. The **Confirm Installation** page is displayed.

31. Click **Install** and follow the installation prompts to complete the installation.

Related concepts

DM-Multipath configuration on page 65
Configuration of the software iSCSI initiator on page 45

Software iSCSI: Configuring the root partition on multipath on the SUSE Linux Enterprise Server

After SAN boot is configured on the host, you can configure the root partition on multipath for your SUSE Linux Enterprise Server host.

Steps

1. Enable additional paths to the NetApp LUN (root LUN) and change the start-up mode of all the iSCSI sessions to automatic by using YaST2.
2. Modify the `/etc/sysconfig/network/ifcfg` *x* file and change the value of STARTMODE to `nfsroot`.
3. Change the session re-establishment timeout of iSCSI sessions fetching root LUN by entering the following command:

   ```
   iscsiadm -m node -T targetname -p ip:port -o update -n
   node.session.timeo.replacement_timeout -v 5
   ```

4. To create a new initrd image with the root partition on multipath enabled, run the following command:

   ```
   mkinitrd -f “iscsi mpath"
   ```

5. Change the start-up mode of the iSCSI sessions fetching the root LUN by entering the following command:

   ```
   iscsiadm -m node -T targetname -p ip:port -o update -n node.startup -v
   onboot
   ```

   ```
   iscsiadm -m node -T targetname -p ip:port -o update -n
   node.conn[0].startup -v onboot
   ```

6. Reboot the host.

   Note: The system now boots with multipath enabled on the root device. You can verify this by running the `mount` command and ensuring that the root partition is on a DM-Multipath device.

Hardware iSCSI: Configuring SAN boot on Red Hat Enterprise Linux

You can configure SAN boot on Red Hat Enterprise Linux using hardware iSCSI.

Steps

1. Create a LUN on the storage system that acts as the SAN boot LUN, and map it to the host.

   You should ensure that the SAN boot LUN is mapped, and multiple paths to the SAN boot LUN are available on the host. You should also ensure that the SAN boot LUN is visible to the host during the boot process.

2. Set the Initiator IP Settings and Initiator iSCSI Name in **Host Adapter Settings**.

3. Set the Primary and Alternate Target IP and iSCSI Name and Adapter Boot Mode to `Manual` in **iSCSI Boot Settings**.

   For information, see your HBA vendor-specific documentation.

   Note: It is best to enable the `spinup delay` option for the individual HBA ports in the HBA BIOS.

4. After making changes to the HBA BIOS, save and exit.

5. Reboot the host.

6. Install the operating system on the LUN configured on the storage controller and follow the installation prompts to complete the installation.

   Note: You should specify Boot Option as `linux mpath` during the operating system installation. When you specify `linux mpath`, you can see the multipath devices (`/dev/mapper/mpathx`) as installation devices.

Veritas Storage Foundation: SAN boot configuration

You can set up a SAN boot LUN to work in the Veritas Storage Foundation environment.

Next topics

Veritas Storage Foundation: Configuring SAN boot on Red Hat Enterprise Linux on page 135
Veritas Storage Foundation: Configuring SAN boot on SUSE Linux Enterprise Server on page 137

Veritas Storage Foundation: Configuring SAN boot on Red Hat Enterprise Linux

You can configure SAN boot on the host using a storage system LUN.

Steps

1. On the storage system, select a supported cfmode.
2. Create a LUN on the storage system, which acts as the SAN boot LUN.
3. Map the LUN to the host.
4. Ensure that only one primary path is available to the LUN.
5. Ensure that only one SAN boot LUN is available to the host.
6. Enable the boot BIOS of the HBA port to which the SAN boot LUN is mapped.

   Note: It is best to enable the `spinup delay` option for the HBA port.

   For information about how to enable the boot BIOS, see the respective HBA vendor-specific documentation.
7. After performing the appropriate changes to the HBA BIOS and ensuring that the SAN boot LUN is visible, install the operating system on the SAN boot LUN.

   Before installing the operating system, see the section on rootability in the *Veritas Volume Manager Administrator's Guide* for Linux that is shipped along with the software for partitioning information.

   Note: Red Hat Enterprise Linux 4 Update 4 does not have bundled HBA drivers for 4-Gb and 8-Gb QLogic cards; therefore, you should use the device driver kit provided by QLogic. For more information about the supported drivers, see the NetApp Interoperability Matrix.
8. After installing the operating system, reboot the host.

The host boots from the SAN boot LUN on the storage system through a primary path.

9. Install the Linux Host Utilities.
10. If you are using external HBA drivers, install the supported versions of the drivers.
11. Verify the HBA settings.
12. Install Veritas Storage Foundation and relevant patches or fixes.
13. Configure the `vxdmp restore` daemon by setting it to an interval of 60:

    `vxdmptune dmp_restore_interval 60`

    On reboot, this value comes into effect and remains persistent.
14. Encapsulate the root disk for use in VxVM by running the following command:

    `vxdiskadm`

    For the detailed steps, see the section on encapsulating the disk in the *Veritas Volume Manager Administrator's Guide* for Linux that is shipped along with the software.
15. Reboot the host after encapsulation.
16. Verify the encapsulation by running the following command:

    `vxprint`

    This command displays the rootvol and swapvol volumes under the corresponding disk group.
17. Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device.

    For more information, see the respective HBA vendor-specific documentation.

Related tasks

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43
Download of the external HBA driver software for FC on page 28
Installation of the external HBA driver software for FC on page 28
Verifying the QLogic HBA settings for FC on page 36

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability

Veritas Storage Foundation: Configuring SAN boot on SUSE Linux Enterprise Server

You can configure SAN boot on the host using a storage system LUN.

Steps

1. On the storage system, select a supported cfmode.
2. Create a LUN on the storage system, which acts as the SAN boot LUN.
3. Map the LUN to the host.
4. Ensure that only one primary path is available to the LUN.
5. Ensure that only one SAN boot LUN is available to the host and that LUN.
6. Enable the boot BIOS of the HBA port to which the SAN boot LUN is mapped.

   Note: It is best to enable the `spinup delay` option for the HBA port.

   For information about how to enable the boot BIOS, see the respective HBA vendor-specific documentation.
7. After performing the appropriate changes to the HBA BIOS and ensuring that the SAN boot LUN is visible, install the operating system on the SAN boot LUN.

   Before installing the operating system, see the section on rootability in the *Veritas Volume Manager Administrator's Guide* for Linux that is shipped along with the software for partitioning information.

   Note: While installing the SUSE Linux Enterprise Server operating system, you should ensure that GRUB is installed in Master Boot Record (MBR). You can do this from the Expert tab in the software package selection screen during installation.
8. After the operating system installation, reboot the host.

   The host boots from the SAN boot LUN on the storage system through a primary path.
9. Install the Linux Host Utilities.
10. Install the supported version of the drivers, if you are using external HBA drivers.
11. Verify the HBA settings.
12. Install Veritas Storage Foundation and relevant patches or fixes.
13. Configure the `vxdmp restore` daemon by setting it to an interval of 60:

    ```
    vxdmptune dmp_restore_interval 60
    ```

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

14. Encapsulate the root disk for use in VxVM by running the following command:

    ```
    vxdiskadm
    ```

    For the detailed steps, see the section on encapsulating the disk in *Veritas Volume Manager Administrator's Guide* for Linux that is shipped along with the software.

15. Reboot the host after encapsulation.

16. Verify the encapsulation by entering the following command:

    ```
    vxprint
    ```

    This command displays the rootvol and swapvol volumes under the corresponding disk group.

17. Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device.

    For more information, see the respective HBA vendor-specific documentation.

Related tasks

Downloading the Host Utilities software on page 43
Installing the Host Utilities software on page 43
Download of the external HBA driver software for FC on page 28
Installation of the external HBA driver software for FC on page 28
Verifying the QLogic HBA settings for FC on page 36

Troubleshooting

If you encounter some problems while running the Host Utilities on FC, iSCSI, or Veritas Storage Foundation, you need to know some troubleshooting tips.

Next topics

FC troubleshooting on page 139
iSCSI troubleshooting on page 144
Veritas Storage Foundation troubleshooting on page 150

FC troubleshooting

If you encounter a problem while running FC on your Host Utilities, you need to know some troubleshooting tips.

Next topics

FC: Unloading of the QLogic module fails on a Linux host on page 139
FC: Linux hosts with QLogic initiators fail to detect LUNs on page 140
FC: multipath -l displays an undef LUN status on page 140
FC: LUNs not recognized by DM-Multipath on page 141
The SAN booted root on a DM-Multipath host freezes during FC path faults on page 141
FC: Poor performance due to misalignment of partitioned LUNs on page 141
Error message from the sanlun fcp show adapter command on page 142
LUNs unavailable after storage system failover or giveback on page 142
sanlun commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts on page 142
SCSI rescan issues on SUSE Linux Enterprise Server hosts on page 143
FC: LUNs not visible on the Linux host on page 143

FC: Unloading of the QLogic module fails on a Linux host

If you connect Linux FC host initiators to multiple initiator FC zones, the HBA module usage count on the initiator host that you specify always shows a higher value than the actual value, even when you do not use the QLogic HBA module.

You should connect each FC host initiator to a separate individual initiator zone instead of a multiple initiator zone.

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

```
qla2400 207360 0
```

```
qla2xxx 3654906 5 qla2400
```

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

FC: Linux hosts with QLogic initiators fail to detect LUNs

Linux hosts with QLogic initiators fail to detect LUNs mapped to it when configured to autonegotiate port topology and speed. This can cause the host to miss a few LUNs that are mapped to it, during the HBA scan. This behavior persists even after rebooting the host.

About this task

In hosts using QLogic initiators, the host initiator fails to establish a consistent topology and speed, even though both are set to autonegotiate. When using QLogic initiators on a Linux host, ensure that you configure the settings provided in this section for port topology and speed using the SANsurfer scli package.

Steps

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

      `/opt/QLogic_Corporation/SANsurferCLI/scli`

   b. Select **Main Menu** > **Configure HBA Settings**.
   c. For each of the WWPNs listed, select **ConnectionOptions** and set it to Point to Point Only or Loop Only, as required.
2. Set the value of QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch or target port it is connected to.
   a. Run the following command:

      `/opt/QLogic_Corporation/SANsurferCLI/scli`

   b. Select **Main Menu** > **Configure HBA Settings**.
   c. For each of the WWPNs listed, select the **Data Rate** option and set it to the specified speed.

FC: multipath -l displays an undef LUN status

With DM-Multipath configured, the `multipath -l` output displays the LUN status as `undef`. To resolve this issue, you can use the `multipath-ll` command.

When you run the `multipath-ll` command, the LUN status is displayed correctly:
`ready`

The `undef` status is displayed because of a bug in the command. It does not indicate an error and can be safely ignored.

FC: LUNs not recognized by DM-Multipath

LUNs are mapped to the Linux host and are correctly reported by `sanlun lun show all` commands. However, the LUNs are not reported by the `multipath -ll` command. Ensure that the LUNs are not blacklisted (excluded) from the DM-Multipath support configuration.

Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v3 -d | grep blacklist
```

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

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

```
multipath -v3
```

For more information, see bug number 201302 on Bugs Online.

Related information

Bugs Online - http://now.netapp.com/NOW/cgi-bin/bol

The SAN booted root on a DM-Multipath host freezes during FC path faults

The SAN booted root on a DM-Multipath host triggers a freeze during FC path faults on Red Hat Enterprise Linux 5 Update 1.

For SAN boot support on Red Hat Enterprise Linux 5 Update 1, you have to download the respective errata multipathing package from the Red Hat Network Web site. For more information, see NetApp Interoperability Matrix.

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability
Red Hat Network Web site - http://www.redhat.com/software/rhn/

FC: Poor performance due to misalignment of partitioned LUNs

If you experience poor performance on some or all FC LUNs that have partitions, the alignment between the LUN block boundaries and the underlying WAFL boundaries might not match.

This problem occurs only with partitions and only with certain combinations of disk geometry parameters. Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. For specific steps, see bug number 156121 on Bugs Online.

Related information

Bugs Online - http://now.netapp.com/NOW/cgi-bin/bol

Error message from the sanlun fcp show adapter command

When you use the `sanlun fcp show adapter` command, it returns an error message.

You see an error message similar to the following:

```
Unable to find /usr/lib/libHBAAPI.so library
Ensure the package installing the library is installed & loaded
```

You should install the QLogic SANsurfer package or the Emulex applications package to obtain the `libHBAAPI.so` library. Install the appropriate package for your HBA model.

Related tasks

Download of the external HBA driver software for FC on page 28
Installation of the external HBA driver software for FC on page 28

LUNs unavailable after storage system failover or giveback

The host requires specific driver, HBA, and FC switch zoning settings to maintain LUN access across storage system takeover, giveback, and reboot. You should check if you have set the required driver and HBA settings. Also, verify that you have correctly zoned your FC switches.

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

- **`/etc/init.d/multipathd restart`**
- **`multipath -v3`**

sanlun commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts

The `sanlun` commands fail on SUSE Linux Enterprise Server 9 x86_64 hosts as there is no 32-bit compatibility library available in the installation media for libdevmapper. Because `sanlun` is a 32-bit application, it can link to only 32-bit libraries.

When `sanlun` is run on a SUSE Linux Enterprise Server 9x86_64 host, it exits with the error message that it cannot find the libdevmapper.so library.

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

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

Related information

Novell Web site - http://support.novell.com/techcenter/psdb/8ae2a27b3b63ac95fd776298026596eb.html

SCSI rescan issues on SUSE Linux Enterprise Server hosts

On a SUSE Linux Enterprise Server host connected to a storage controller, the presence of dummy LUNs causes issues during a SCSI rescan. This issue is not yet resolved in SUSE Linux Enterprise Server 9 and SUSE Linux Enterprise Server 10.

Dummy LUNs are those that show up as "UNKNOWN" in `/proc/scsi/scsi/` even when no LUNs are mapped to the host. These LUNs have a generic device node (/dev/sg* entry), but no corresponding disk node (/dev/sd* entry). Due to these dummy LUNs, subsequent SCSI rescans on the host fail to recognize newly added LUNs, unless preceded by a write to the sysfs delete attribute of the dummy LUN.

FC: LUNs not visible on the Linux host

FC LUNs appear as local disks to the host. If the storage system LUNs are not available as disks on the hosts, try rescanning the HBAs by reloading the drivers. If the LUNs still do not appear, verify the configuration settings.

Configuration settings	**What you should do**
Driver version	Verify that the version of the external HBA driver you are running is supported. You can check the up-to-date system requirements in the NetApp Interoperability Matrix..
Cabling	Verify that the cables between the host and the storage system are properly connected.
Zoning	Verify that the zoning is correct.
HBA Setting	Verify the recommended HBA settings.
FC service status	Verify that the FC service is licensed and started on the storage system. For more information, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.
HBA World Wide Port Names	Verify that you are using the correct initiator WWPNs in the igroup configuration. On the storage system, use the `igroup show` command to display the WWPNs of the initiators in the storage system's igroups. On the host, use the HBA tools to display the initiator WWPN. The initiator WWPNs configured in the igroup and on the host should match.
Initiator login	Verify that the initiator is logged in to the storage system by entering the `fcp 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 WWPN.

Configuration settings	What you should do
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 the 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 up-to-date system requirements in the NetApp Interoperability Matrix.

iSCSI troubleshooting

If you encounter a problem while running iSCSI on your Host Utilities, you need to know some troubleshooting tips.

Next topics

OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 10 series and 9 series on page 144
iSCSI: multipath -l displays an undef LUN status on page 145
Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 11 on page 145
Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 10 series on page 145
Automatically mounting and unmounting the file system on multipath device do not work on SUSE Linux Enterprise Server 9 series on page 146
Multipathd occasionally fails to update the path status of DM-Multipath devices on page 147
iSCSI: LUNs not recognized by DM-Multipath on page 147
Storage system console displays continuous ISID Rule messages on SUSE Linux Enterprise Server 9 series on page 147
iSCSI: Poor performance due to misalignment of partitioned LUNs on page 148
iSCSI: LUNs not visible on the Linux host on page 148

OCFS2 self-fencing during cluster failover on SUSE Linux Enterprise Server 10 series and 9 series

During a clustered failover or network disruption, all paths to the storage system can become unavailable. As a result, the nodes might fence by panicking. Set the

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

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

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

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

Depending on the load in the iSCSI network, you should raise the value. It is best to use a value in the range of 91 to 201 for a configuration with OCFS2 on multipath, and iSCSI devices.

Related information

Oracle Web site-http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT

iSCSI: multipath -l displays an undef LUN status

With DM-Multipath configured, the `multipath -l` output displays the LUN status as `undef`. To resolve this issue, you can use the `multipath-ll` command.

When you run the `multipath-ll` command, the LUN status is displayed correctly:
`ready`

The `undef` status is displayed because of a bug in the command. It does not indicate an error and can be safely ignored.

Automatically mounting LVM devices during system boot on SUSE Linux Enterprise Server 11

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

To overcome this problem, manually mount the logical volumes by using the following command:

```
/sbin/mount -a
```

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

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

iSCSI devices are unavailable. To overcome this problem, a helper script is provided in `/usr/share/doc/packages/lvm2/ lvm-vg-to-udev-rules.sh`.

You can use this script to generate `udev` rules for iSCSI Logical Volume Manager (LVM) devices, which can be automatically mounted during system boot.

Example:

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

After completing the preceding steps after each logical volumes, the logical volumes can be automatically mounted by adding the entries in `/etc/fstab`.

Automatically mounting and unmounting the file system on multipath device do not work on SUSE Linux Enterprise Server 9 series

Automated mounting of multipath-based file systems at boot time using entries in `/etc/fstab.conf` or `/etc/fstab.iscsi` does not work. You need to manually mount and unmount the multipath-based file systems on SUSE Linux Enterprise Server 9 series.

Automated mounting of multipath-based file systems does not work because the placement of the multipath bootup and shutdown scripts in the rc script list in SUSE Linux Enterprise Server 9 series is out of order with respect to the mount and unmount scripts for iSCSI and non-iSCSI devices.

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

As a result, the multipath bootup and shutdown sequence is not synchronized with the iSCSI bootup and shutdown sequence. Multipath devices are not available when the `mount` command is run, and the `mount` command fails to mount these devices.

Similarly in the shutdown path, by the time the `unmount` script runs, the multipath utility would have already shut down.

After the multipath devices are available on bootup, the file systems can be mounted. For unmounting, ensure that any multipath-based file systems are unmounted before the multipath utility is shut down.

1. Unmount all the file systems on the iSCSI LUNs.
2. Flush the multipath maps by entering the following command:

   `multipath -F`

3. Stop the multipath daemon by entering the following command:

   `/etc/init.d/multipathd stop`

4. Stop the iSCSI service by entering the following command:

```
/etc/init.d/iscsi stop
```

For more information, see bug number 196997 on Bugs Online.

Related information

Bugs Online - http://now.netapp.com/NOW/cgi-bin/bol

Multipathd occasionally fails to update the path status of DM-Multipath devices

Occasionally, multipathd does not update the path status of DM-Multipath devices after running I/O faults. Set both HOTPLUG_USE_HWSCAN and HOTPLUG_USE_SUBFS to `no` in the `/etc/sysconfig/hotplug` file.

Multipathd fails occasionally because it fails to start "event checker" for some of the DM-Multipath devices. Because of this failure, multipathd is unable to keep track of the path up or down status for those devices.

iSCSI: LUNs not recognized by DM-Multipath

LUNs are mapped to the Linux host and are correctly reported by `sanlun lun show all` commands. However, the LUNs are not reported by the `multipath -ll` command. Ensure that the LUNs are not blacklisted (excluded) from the DM-Multipath support configuration.

Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v3 -d | grep blacklist
```

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

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

```
multipath -v3
```

For more information, see bug number 201302 on Bugs Online.

Related information

Bugs Online - http://now.netapp.com/NOW/cgi-bin/bol

Storage system console displays continuous ISID Rule messages on SUSE Linux Enterprise Server 9 series

The storage system console displays continuous ISID Rule messages. Verify that you have added `Continuous=no` to the `/etc/iscsi.conf` file. Ensure that this line is at the beginning of the file, before any “DiscoveryAddress” entries.

ISID Rule messages look 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 iqn.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 iqn.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
```

iSCSI: Poor performance due to misalignment of partitioned LUNs

If you experience poor performance on some or all iSCSI LUNs that have partitions, the alignment between the LUN block boundaries and the underlying WAFL boundaries might not match. Manually, enter the disk geometry parameters when creating partitions with the `fdisk` command.

This problem occurs only with partitions and certain combinations of disk geometry parameters. For more details, see bug number 156121 on Bugs Online.

In a default setup, LVM scans all attached disks and determines which of them contain physical volumes. When using multipathing, multiple devices can contain identical physical volumes. If LVM detects the SCSI device first, the device is treated as a separate physical volume.

Modify the filter in the `/etc/lvm/lvm.conf` file to scan only the multipath device and not the SCSI device underneath: `filter = [ "r|/dev/.*/by-path/.*|", "r|/dev/.*/by-id/.*|", "r|/dev/sda2|", "r|/dev/sd*|", "a/.*/" ]`

Related information

Bugs Online - http://now.netapp.com/NOW/cgi-bin/bol

iSCSI: LUNs not visible on the Linux 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 configuration settings.

Configuration setting	What you should do
Cabling	Verify that the cables between the host and the storage system are properly connected.

Configuration setting	**What you should do**
Network connectivity	Verify that there is TCP/IP connectivity between the host and the storage system. From the storage system command line: • Ping the host interfaces that are used for iSCSI. • Ping the storage system interfaces that are used for iSCSI.
iSCSI service status	Verify that the iSCSI service is licensed and started on the storage system. For more information, see the *Data ONTAP Block Access Management Guide for iSCSI and FC*.
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 that no initiators are logged in, check the initiator configuration on the host. Verify that the storage system is configured as a target of the initiator.
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 should 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 the 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 up-to-date system requirements in the NetApp Interoperability Matrix.

Configuration setting	What you should do
Jumbo frames	If you are using jumbo frames in your configuration, ensure that the jumbo frames are enabled on all the devices in the network path: the host Ethernet NIC, the storage system, and any switches.
Firewall settings	Verify that the iSCSI port (3260) is open in the firewall rule.

Related concepts

Configuration of software iSCSI on Red Hat Enterprise Linux on page 48

Related tasks

Verifying the QLogic HBA settings for FC on page 36
Verifying the QLogic HBA settings for hardware iSCSI on page 37

Related information

NetApp Interoperability Matrix - http://now.netapp.com/NOW/products/interoperability

Veritas Storage Foundation troubleshooting

If you encounter a problem while running Veritas Storage Foundation on your Host Utilities, you need to know some troubleshooting tips.

Next topics

Error while stopping the fencing driver on page 150
Secondary paths identified incorrectly on page 151
Enclosure-based naming not reflected on NetApp storage on page 151
SAN boot installation to LUN configured on a storage controller fails on page 151
Encapsulation of the root disk fails on page 151

Error while stopping the fencing driver

You might get an error when you stop the fencing driver.

Steps

1. Stop the Veritas cluster service by using `hastop` on all the nodes.
2. Try stopping the fencing driver again.
3. If the fencing driver does not stop, remove the name of the coordinator diskgroup from `/etc/vxfendg`.
4. In the `/etc/vxfenmode` file, make the following change:

```
vxfen_mode=disabled
```

5. Stop the fencing driver by using the following command:

```
/etc/init.d/vxfen stop
```

Note: It is best to halt all the nodes except the last one in the cluster.

Secondary paths identified incorrectly

Secondary paths are identified incorrectly as primary paths while configuring DMP. In such a scenario, you should verify that the ASL for the storage system is installed correctly without any error messages during installation.

Enclosure-based naming not reflected on NetApp storage

The VxVM enclosure-based naming feature is enabled, but the DMP devices do not contain a storage model string. In such a scenario, verify that the ASL for the storage system is installed correctly without any error messages during installation.

SAN boot installation to LUN configured on a storage controller fails

SAN boot fails if you have multiple paths configured to boot the LUN before the operating system installation. You should verify that the host can see only one path to the LUN that is configured on a storage controller during installation.

Encapsulation of the root disk fails

Encapsulation of the root disk fails after SAN boot installation. Ensure that you follow the partitioning scheme.

For more information on partitioning scheme, see the *Veritas Volume Manager Administrator's Guide*.

Index

A

add targets
    target IP address 39
    target node name 39
APM
    failover path 76
    I/O error handling 76
    I/O error processing 76
    path failures 76

C

CHAP, authentication
    password_in 49
    passwords 49
    user names 49
    username_in 49
command
    ./qlinstall -br -in 32
    ./qlinstall -l 32
    ./qlinstall -ul 32
    fcp show cfmode 61
    iscsi security 62
    lsmod 40
    modprobe 35
    scli 36
command, hardware iSCSI
    modprobe 36
command, iSCSI initiator
    iscsiadm 55
    yast2 55
command, iSCSI package
    rpm -q iscsi-initiator-utils 45
    rpm -q linux-iscsi 45
    rpm –q open-iscsi 45
command, iSCSI service
    chkconfig 60
commands
    mkinitrd 30
    modprobe.conf 30
configuration file
    /etc/iscsi/iscsid.conf 49
configure DM-Multipath, command
    chkconfig 71
current sessions
    iscsiadm -m session 85, 115

D

Device Mapper Multipathing
    DM-Multipath 23
    Red Hat Enterprise Linux 23
    SUSE Linux Enterprise Server 23
device persistence
    multipath device 24
    persistent devices 24
discovering LUNs
    iscsiadm 85, 115
    rescan script 82, 104
    sanlun 85, 115
    SCSI rescan 82, 104
    sg3_utils 82, 104
discovering LUNs, command
    current session
        iscsiadm -m session 84, 114
    rescan, all sessions
        iscsiadm -m session --rescan 84, 114
    rescan, specific session
        iscsiadm -m session --sid=N --rescan 84, 114
    SCSI rescan script
        rescan-scsi-bus.sh 84, 114
discovering LUNs, software iSCSI
    SCSI rescan script 84, 114
    sg3_utils 84, 114
discovering LUNs, verification
    iscsi-ls 85, 115
    sanlun 85, 115
displaying multipath, command
    sanlun lun show -p 82, 105
DM-Multipath
    I/O failover 23
    methods to access LUNs 81
    path load sharing 23
    SCSI device 23
DM-Multipath configuration file
    /etc/multipath.conf 66
DM-Multipath configuration, command
    ls -l /dev/mapper/ 72
    multipath -ll 72
    multipath -v3 -d -ll 72
DM-Multipath, adding new LUN

multipath 87
DM-Multipath, configuration command
multipath 71
DM-Multipath, viewing LUNs, commands
iscsi-ls -l 87, 107
iscsiadm -m session --sid=N -P 3 87, 107
sanlun lun show all 87, 107
driver parameters, QLogic
ql2xsuspendcount 30
qlport_down_retry 30
driver parameters, QLogic external
MaxRetriesPerIo 32
MaxRetriesPerPath 32
ql2xsuspendcount 32
driver settings, Red Hat Enterprise Linux
default setting 34
required setting 34
driver settings, SUSE Linux Enterprise Server
default setting 34
required setting 34

F

FC protocol
host bus adapters
HBA port 21
FCoE
converged network adapters 21
data center bridging 21
Ethernet switch 21
traditional FC 21

H

HBA
management software 29
software components
Applications kit 29
SANsurfer CLI 29
SANsurfer FC HBA Manager 29
SANsurfer iSCLI 29
SANsurfer iSCSI HBA Manager 29
HBA installation
bus type 28
HBA parameters, QLogic hardware iSCSI
IP_ARP_Redirect 33
KeepAliveTO 33
HBA setting, FC QLogic
LinkDownTimeout 36
PortDownRetryCount 36
Host Utilities
diagnostic scripts 19
download software 43
drivers 27
HBA 27
host configuration 27
SAN Toolkit 19
software version 43
support files 19

I

initiator group
igroup 21
initiator node name
create igroup 40
igroup 40
installing
latest Host Utilities 43
Linux software 43
iSCSI protocol
HBA 22
iSCSI target HBA 22
standard Ethernet interfaces 22
iSCSI service reload, command
/etc/init.d/iscsi reload 86, 116
iSCSI, configure SAN boot
hardware iSCSI
protocol stack 129
network interface card (NIC) 129
software iSCSI
locally attached disk 129
preboot execution environment (PXE) 129

L

Linux configurations
ALUA support
automatically enabled 23
asymmetric logical unit access
Target Port Group Support 23
Linux Host Utilities
direct-attached configurations
high-availability 21
single-zoned 21
stand-alone 21
fabric-attached SAN
dual-port 21
single-port 21
Linux software packages

perl-HTML-Parser 43
perl-HTML-Tagset 43
perl-libwww-perl 43
perl-XML-Parser 43
Linux versions
Oracle Enterprise Linux 19
Red Hat Enterprise Linux 19
SUSE Linux Enterprise Server 19
Veritas Storage Foundation 19
list of LUNs, commands
iscsi-ls -l 95
iscsiadm -m session --sid=N -P 3 95
iscsiadm -m session -P 3 95
sanlun 95
login mode
automatic 50
manual 50
LUN access, Veritas DMP
restore daemon 103
vxdmp 103

M

mounting, new file
device mount_point type _netdev 0 0 85, 115
multipathing command
rpm -q 65
multipathing packages
device-mapper 65
device-mapper-multipath 65
multipath-tools 65
multipathing software
persistent identification 95

P

parameter, hardware iSCSI
IP_ARP_Redirect 37
ql4xkeepalive 37
protocols
Fibre Channel (FC) 20
Fibre Channel over Ethernet (FCoE) 20
iSCSI 20

Q

QLogic FC HBA
HBA parameters
Link Down Timeout 29
Port Down Retry Count 29

R

reloading iSCSI, command
/etc/init.d/iscsi reload 85, 115
removing LUN
rescan script 83, 106
SCSI rescan 83, 106
removing LUN, command
rescan-scsi-bus.sh –r 83, 106
rescan, all sessions
iscsiadm -m session --rescan 85, 115
rescan, SCSI rescan script
rescan-scsi-bus.sh 85, 115
restore daemon, command
vxdmpadm start restore 103
vxdmpadm stat restored 103
vxdmpadm stop restore 103
vxdmptune dmp_restore_interval 103

S

SAN boot
centralized storage 22
disaster recovery 22
maintenance 22
reliability 22
server swaps 22
SAN boot, FC
BIOS 121
bundled HBA driver 121, 124
configuring root partition, SUSE Linux Enterprise Server 127
DM-Multipath 121, 124, 127
Emulex external HBA driver 122
external HBA driver 121, 124
firmware modules 122
HBA BIOS 125
initird 127
installing, Red Hat Enterprise Linux 121
installing, SUSE Linux Enterprise Server 125
linux mpath 121
QLogic external HBA driver 122
Red Hat Enterprise Linux, SAN boot 121
SAN boot LUN 121, 127
SCSI device 124
spinup delay 125
SUSE Linux Enterprise Server, SAN boot 124
SAN boot, hardware iSCSI

Boot Option 134
linux mpath 134
Red Hat Enterprise Linux, SAN boot 134
SAN boot LUN 134
spinup delay 134
SAN boot, root partition
iSCSI sessions 133
mount 133
Start-Up mode 133
SUSE Linux Enterprise Server, SAN boot 133
SAN boot, software iSCSI
/boot 131
Boot Option 129
initiator name 131
IQN 131
iSCSI initiator name 129
linux mpath 129
Red Hat Enterprise Linux, SAN boot 129
root file system 129
SUSE Linux Enterprise Server, SAN boot 131
SWAP partition 129, 131
target IP address 129
sanlun command, Veritas
sanlun lun show -p 120
software iSCSI initiator
initiator setting 45
root previleges 45
specific session, rescan
iscsiadm -m session --sid=N --rescan 85, 115
starting DM-Multipath, command
/etc/init.d/boot.multipath start 71
/etc/init.d/multipathd start 71
stopping DM-Multipath, command
/etc/init.d/boot.multipath stop 73
/etc/init.d/multipathd stop 73
storage system
direct-attached 20
Ethernet network 20
fabric-attached 20
storage system configuration
FC, WWPN 61
software, hardware - iSCSI initiator node name 61
storage system configuration, FC
cfmode setting 61
creating igroup 61
managing LUN 61
supported protocols
FC 20
FCoE 20
iSCSI 20

U

unloading HBA driver, bundled or external
LVM volume 40
multipath map 40
unmounting 40

V

Veritas
Dynamic Multipathing
DMP 24
Veritas Volume Manager
VxVM 24
Veritas DMP
bundled HBA 103
external driver 103
restore daemon 118
Veritas Dynamic Multipathing
VxDMP 24
Veritas Storage Foundation
Array Policy Module, APM 76
Array Support Library, ASL 76
boot BIOS 135
configure SAN boot, Red Hat Enterprise Linux 135
configure SAN boot, SUSE Linux Enterprise
Server 137
GRUB 137
HBA BIOS 135
HBA port 135
Master Boot Record 137
primary path 137
spinup delay 137
swapvolume 137
viewing LUN information
sanlun 95
viewing LUNs, command
sanlun lun show all 82, 105
VxVM, command
vxddladm rmjbod vid=NETAPP 76
vxdisk list 118
vxdmpadm getsubpaths ctlr 118
vxdmpadm getsubpaths dmpnodename 118
vxdmpadm listenclosure all 76

W

WWPN
creating igroup 38
Emulex HBAnyware 39

HBA ports 38
initiator access 38
SANsurfer FC HBA CLI 39