

IBM System Storage N series



Linux Host Utilities 6.2 Installation and Setup Guide

Contents

Preface	8
Supported features	8
Websites	8
Getting information, help, and service	8
Before you call	9
Using the documentation	9
Hardware service and support	9
Firmware updates	9
How to send your comments	10
Changes to this document: July 2013	11
The Linux Host Utilities	12
Overview of the supported Linux environments	13
How to find instructions for your Linux environment	15
Quick start checklist for planning the Host Utilities installation	17
Setup procedures for Intel NIC, QLogic, Brocade, and Emulex	
HBAs	20
(FC) What you need to know about Fibre Channel HBAs and drivers	20
(iSCSI) Configuring Broadcom 10 Gb iSCSI offload	21
Installing the Linux Host Utilities software	24
(Red Hat) Configuring space reclamation on hosts running Red Hat Linux 6.2	25
(Oracle VM) Configuring the O2CB_HEARTBEAT_THRESHOLD	26
(iSCSI) How to configure iSCSI for Linux	27
(iSCSI) Getting the initiator node name	27
(iSCSI) Setting the timeout values to enable multipathing	29
(iSCSI) Setting up CHAP for Red Hat Linux 5, 6 and SUSE Linux 10, 11	30
(iSCSI) Setting up CHAP for Red Hat Enterprise Linux 4 series	31
(iSCSI) Starting the iSCSI service	33
(iSCSI) Methods for setting up target discovery with software initiators	33
(iSCSI) Discovering the iSCSI target by using the iscsiadm utility on	
Red Hat 5, 6, SUSE 10, 11	34
(iSCSI) Setting up target discovery on Red Hat Enterprise Linux 4 series	
.....	34

(iSCSI) Discovering targets by using YaST2 on SUSE 10, 11	35
(iSCSI) Configuring the iSCSI service to start automatically	36
(iSCSI) Configuring manual or automatic node login	36
Configuring the storage system	38
DM-Multipath configuration	40
Verifying the required multipathing packages	40
Editing the DM-Multipath configuration file	41
Starting DM-Multipath	43
Configuring DM-Multipath to start automatically while booting	44
Verifying the DM-Multipath configuration	45
Stopping DM-Multipath	48
Veritas Dynamic Multipath configuration	49
(Veritas) VxDMP restore daemon and LUN retries tunable configuration	49
(Veritas) Setting the restore daemon and LUN retry tunable values	50
(Veritas, Red Hat) Configuring Red Hat 6 to support Veritas Storage Foundation	51
(Veritas, SUSE) Configuring SUSE Linux 11 to support Veritas Storage Foundation	52
(Veritas) The Array Support Library and Array Policy Module	53
(Veritas) What the ASL is	53
(Veritas) What the APM is	54
(Veritas) Installing the ASL and APM software	54
(Veritas) Removing the ASL and APM	55
(Veritas) Information about ASL error messages	56
Methods for working with LUNs in native Linux environments	57
(FC, Hardware iSCSI) Discovering new LUNs	57
(Software iSCSI, multipathing) Discovering new LUNs on Red Hat 5, 6, SUSE 10, 11	58
(Software iSCSI, multipathing) Discovering new LUNs on Red Hat 4	59
Viewing a list of LUNs	59
Examples of sanlun, iscsiadm, iscsi output when used to view LUNs	60
(Native multipathing) Using sanlun to display DM-Multipath information	66
(Native multipathing) Examples of sanlun output containing DM- Multipath information	67
Enabling device persistence for newly discovered LUNs	70
Removing an unmapped LUN	71

(Veritas) LUN access when using VxDMP	73
(Veritas, FC) Discovering new LUNs	73
(Veritas, Software iSCSI) Discovering new LUNs for Red Hat 5, 6, or SUSE 10, 11	74
(Veritas) Viewing LUNs mapped to the host	75
(Veritas) Examples of sanlun, iscsiadm, iscsi output when used to view LUNs	75
(Veritas) Displaying multipathing information for VxDMP	80
(Veritas) Examples of sanlun output for VxDMP	80
(Veritas) Removing a LUN	81
(Veritas) Displaying available paths using VxVM	82
(FC) Setting up a SAN boot LUN on Red Hat Enterprise Linux	84
(FC) Setting up a SAN boot LUN on SUSE Linux Enterprise Server ...	85
(FC) Configuring the root partition with DM-Multipath on SUSE Linux Enterprise Server	85
(iSCSI) SAN boot configuration for iSCSI hardware, software initiators	87
(Hardware iSCSI) Configuring SAN boot on Red Hat Enterprise Linux	87
(Software iSCSI) Configuring SAN boot on Red Hat Enterprise Linux	88
(Software iSCSI) Configuring SAN boot on SUSE Linux Enterprise Server	89
(Software iSCSI) Configuring multipathing for a SAN boot LUN using SUSE Linux Enterprise Server	91
(Veritas) Configuring SAN boot in a Veritas environment	93
Support for host virtualization	95
Hypervisor VHD requires alignment for best performance	96
Supported Linux and Data ONTAP features	97
Protocols and configurations supported by Host Utilities	97
The FC protocol	97
The FCoE protocol	98
The iSCSI protocol	98
SAN booting	99
Support for Linux Device Mapper Multipathing	99
Volume management and multipathing with Veritas Storage Foundation ...	99
Linux configurations that support ALUA	100
Troubleshooting	101
LUNs not recognized by DM-Multipath	101

(FC) LUNs not visible on the Linux host	101
FC troubleshooting	102
(FC) Warnings displayed during Host Utilities installation	102
(FC) Linux hosts with QLogic initiators fail to detect LUNs	103
(FC) The SAN booted root on a DM-Multipath host freezes during FC path faults	104
(FC) Poor performance due to misalignment of partitioned LUNs	104
(FC) sanlun command displays error when HBA libHBAAPI.so is missing	104
(FC) LUNs are unavailable after a storage system failover or giveback	105
(FC) SCSI rescan issues on SUSE Linux Enterprise Server hosts	105
Warning messages displayed when using LVM with multipathing software	105
iSCSI troubleshooting	106
(iSCSI) LVM devices are not automatically mounted during system boot on SUSE Linux Enterprise Server 11	106
(iSCSI) LVM devices are not automatically mounted during system boot on SUSE Linux Enterprise Server 10	106
Multipathd occasionally fails to update the path status of DM-Multipath devices	106
(iSCSI) Poor performance due to misalignment of partitioned LUNs	107
iSCSI: LUNs not visible on the Linux host	107
Veritas Storage Foundation troubleshooting	108
(Veritas) Error while stopping the fencing driver	108
(Veritas) Secondary paths identified incorrectly	109
(Veritas) Enclosure-based naming not reflected on N series storage	109
(Veritas) Setting up a LUN for SAN booting fails	109
(Veritas) Encapsulation of the root disk fails	109
Installing the nSANity data collection program	109
Sample configuration files for Red Hat Enterprise Linux 6	111
Sample configuration file for Red Hat Enterprise Linux 5	119
Sample configuration file for Red Hat Enterprise Linux 4	124
Sample configuration file for Oracle Linux	126
Sample configuration file for Red Hat Enterprise Virtualization Hypervisor	127
Sample configuration file for Oracle VM 3.0.1, 3.1.1, and 3.2.1	131

Sample configuration files for SUSE Linux Enterprise Server 11	
series	137
Sample configuration file for SUSE Linux Enterprise Server 10	141
Sample Configuration for Citrix XenServer	144
Where to find more information	149
Copyright and trademark information	152
Trademark information	153
Index	156

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 8) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Changes to this document: July 2013

This section contains information on the changes made to this guide since it was originally released.

The 6.2 version of the Linux Host Utilities adds support for more configurations and features. This document contains information about those features. In addition, this document is updated when information about new configurations and features becomes available.

Any time this document is updated, a note is added to the Release Notes. It is a good practice to check the online Release Notes regularly to determine whether there is new information about using the Linux Host Utilities or changes to this guide. The most current versions of the Release Notes and this guide are posted on the N series support website (accessed and navigated as described in [Websites](#) on page 8).

Note: For the most current information on what the Linux Host Utilities support, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

July 2013 update

This document was updated in July 2013 to support the Linux Host Utilities 6.2. The following information was added to this document:

- Support for Citrix XenServer and the sample `multipath.conf` file has been added.
For information, see [Sample Configuration for Citrix XenServer](#) on page 144.
For information about iSCSI timeout settings, see [Setting the timeout values to enable multipathing](#) on page 29.
- Support for SUSE Linux Enterprise Server 11 SP2 KVM and XEN and the sample `multipath.conf` file has been added.
For information, see [Sample configuration file for SUSE Linux Enterprise Server 11 series](#) on page 137.
- Support for Red Hat Enterprise Virtualization Hypervisor 6.3 and the sample `multipath.conf` file has been added.
For information, see [Sample configuration file for Red Hat Enterprise Virtualization Hypervisor](#) on page 127.
- Support for Oracle VM 3.2.1 and the sample `multipath.conf` file has been added.
For information, see [Sample configuration file for Oracle VM](#) on page 131.
- Support for 5.9 and 6.4 versions of Red Hat Linux Enterprise Server.
For information, see [Sample configuration file for Red Hat Enterprise Linux Series 6](#) on page 111.
- Support for Intel NIC card and Brocade HBA/CNA.
For information, see [Setup procedures for Intel NIC, QLogic, Brocade, and Emulex HBAs](#) on page 20.

The Linux Host Utilities

The Host Utilities provide software programs and documentation that you can use to connect your Linux host to IBM N series storage systems running Data ONTAP. The software is available as either a 32-bit or a 64-bit .rpm file.

The Host Utilities include the following components:

- The SAN Toolkit

The toolkit is installed automatically when you install the Host Utilities. This kit provides the following key tools:

Note: This toolkit is common across all configurations and protocols of the Host Utilities. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance.

- The `sanlun` utility, which helps you to manage LUNs and HBAs.
- The `san_version` command, which displays the versions of the Host Utilities.

Note: Previous versions of the Host Utilities also included diagnostics programs. These programs have been replaced by the nSANity Diagnostic and Configuration Data Collector, and are no longer installed with the Host Utilities. The nSANity program is not part of the Host Utilities. You should download, install, and execute it only when requested to do so by technical support.

See the man pages for these commands for details on using them.

- Documentation

The documentation provides information about installing, setting up, using, and troubleshooting the Host Utilities. The documentation consists of the following:

- This installation and setup guide
- *Quick Command Reference*
- *Host Settings Affected by Linux Host Utilities*
- *Quick Start Guide*
- *Release Notes*

Note: The *Release Notes* are updated whenever new information about the Host Utilities is available. You should check the *Release Notes* before installing the Host Utilities to see if there is new information about installing and working with the Host Utilities.

Overview of the supported Linux environments

The Host Utilities support several Linux environments. These environments consist of different versions of the Linux operating system using protocols such as Fibre Channel (FC), FCoE, or iSCSI. Some environments also use Veritas Storage Foundation.

For details about the environments, including the supported operating system versions and Veritas Storage Foundation versions, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

The following table summarizes key aspects of the main environments:

Linux Environment	Notes
<p>Red Hat Enterprise Linux, Oracle Linux, SUSE Linux Enterprise Server KVM and Xen, Citrix XenServer, Red Hat Enterprise Linux KVM, SUSE Linux Enterprise Server, and Red Hat Enterprise Virtualization Hypervisor</p>	<ul style="list-style-type: none"> • This environment works with features provided by the Linux operating system. • Multipathing: DM-Multipath • Volume management: Native Linux volume managers • Protocols: FC, FCoE, and iSCSI • ALUA (Asymmetric Logical Unit Access): For systems using clustered Data ONTAP, ALUA is supported with FC, FCoE, and iSCSI. For systems using Data ONTAP operating in 7-Mode, ALUA is supported with FC and FCoE only. • Setup issues: <ul style="list-style-type: none"> • If you are using HBAs, set them up before you install the Host Utilities. • If you are using multipathing, set up DM-Multipath. • (Red Hat Enterprise Linux 6 update 2 or later) If you want to use space reclamation, you must use this version of the operating system. • Configuration issues: <ul style="list-style-type: none"> • If you are using a Hypervisor Virtual hard Disk (HVD), make sure the LUNs are properly aligned for best performance. <p>Oracle Linux:</p> <p>This guide provides instructions and examples using Red Hat Enterprise Linux and SUSE Linux Enterprise Server. In most cases, Oracle Linux uses the same setup procedures as Red Hat Enterprise Linux. To simplify this guide, it uses "Red Hat" to refer to both systems using Red Hat Enterprise Linux and systems using Oracle Linux. If Oracle Linux requires a different procedure, that procedure is included.</p> <p>Note: Ensure that the kernel and dm-multipath versions are as per interoperability matrix tool. If not, install the versions from Oracle ULN (Oracle Unbreakable Linux Network).</p> <p>Note: The examples in this documentation apply to current Linux distributions. If you have an earlier version of Linux, you must check the documentation for previous versions of the Linux Host Utilities. For example, if you use SUSE Linux Enterprise Server 9, see the <i>Linux Host Utilities Installation and Setup Guide</i> for those instructions and examples.</p>

Linux Environment	Notes
Veritas Storage Foundation	<ul style="list-style-type: none"> • This environment works with features provided by Veritas Storage Foundation • Multipathing: Veritas Dynamic Multipathing (VxDMP) • Linux operating systems: Red Hat Enterprise Linux, Oracle Linux, or SUSE Linux Enterprise Server • Volume management: Veritas Volume Manager (VxVM) • Protocols: Fibre Channel (FC) and, with certain versions of Veritas Storage Foundation, iSCSI • ALUA: Supported with Veritas 5.1 and FC • Setup issues: <ul style="list-style-type: none"> • Set up the HBAs before you install the Host Utilities. • If you have Veritas Storage Foundation 5.0 or earlier, you must install the Array Support Library (ASL) and the Array Policy Module (APM).

How to find instructions for your Linux environment

Many instructions in this guide apply to all the environments that the Linux Host Utilities support. Sometimes, though, a command or step applies to only one environment. In those cases, this guide uses a qualifier to identify the environment.

The qualifier lets you quickly determine whether the information applies to your environment. If the qualifier does not apply to your environment, then you do not need read the section. For example, a title that starts with **(Software iSCSI)** indicates a section that applies only to environments using software iSCSI initiators.

If there is no qualifier, then the information applies to all environments.

Qualifier	The section that follows applies to
(Veritas)	Environments using Veritas VxDMP as the multipathing solution.
(Red Hat)	Environments that use Red Hat Enterprise Linux. Red Hat Enterprise Linux and Oracle Linux use many of the same setup procedures. Unless otherwise specified, the information for Red Hat environments also applies to Oracle Linux environments.
(Oracle VM)	Environments that use Oracle VM.
(SUSE)	Environments that use SUSE Linux Enterprise Server.
(Native multipathing)	Environments using DM-Multipath.

Qualifier	The section that follows applies to
(FC)	Environments using the FC or FCoE protocol.
(iSCSI)	Environments using iSCSI.
(Software iSCSI)	Environments using the iSCSI software initiator.
(Hardware iSCSI)	Environments using iSCSI with a hardware HBA.

Quick start checklist for planning the Host Utilities installation

Installing the Host Utilities and setting up your system involves numerous tasks that are performed on both the storage system and the host. The following checklist provides a high-level overview of these tasks.

If you are an experienced Linux user, this checklist can serve as a quick start guide to installing and setting up the Host Utilities.

Regardless of your experience level, planning how you will install and configure the Host Utilities is a good practice.

The detailed steps for each of the tasks presented in the checklist are provided later in this guide.

Note: Occasionally there are known problems that can affect your system setup. Review the *Host Utilities Release Notes* before you install the Host Utilities. The *Release Notes* are updated whenever an issue is found, and may contain information about the Host Utilities that was discovered after this guide was produced.

Task 1: Make sure the prerequisites for installing and setting up the Host Utilities have been met

1. Verify that your system setup is correct. Check the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8) for the most current information about system requirements. This includes:

- Host operating system version and patches.
- HBAs and drivers, or software initiators, model, and version.

Note: This guide contains information on QLogic and Emulex HBAs. Other HBAs are also supported. Check the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8) for complete information on supported HBAs.

- **(Veritas)** Veritas Storage Foundation
Set up the Veritas Volume Manager (VxVM), and then set up the Array Support Library (ASL) and Array Policy Module (APM), if this software was not installed with your version of Veritas Storage Foundation.
- Volume management and multipathing.
- Optional setup based on your operating system
 - **(Red Hat Enterprise Linux 6 or later)** This version of the Red Hat Enterprise Linux operating system supports space reclamation, but you must perform some setup steps to use space reclamation successfully.
 - **(Oracle VM)** If you are using Oracle VM, you need to set the `O2CB_HEARTBEAT_THRESHOLD` to 65.

2. Verify that your storage system is:

- Running a supported version of Data ONTAP.

Note: If you are using Red Hat Enterprise Linux 6.0 or later, you should run Data ONTAP 8.0.1 or higher on the storage systems. That way you can use the Block Limits VPD page (0xb0) information. Using the Block Limits VPD page helps maintain block alignment, which results in enhanced I/O performance on the N series LUN.

- Running the correct, licensed protocol for your environment.
- Set up to work with the host and the initiator HBAs or software initiators, as needed by your protocol.
- Set up to work with ALUA, if it is supported with your version of Data ONTAP and your protocol.
For clustered Data ONTAP, ALUA works with the FC, FCoE and iSCSI protocols and **must** be enabled. For Data ONTAP operating in 7-Mode, ALUA works with the FC and FCoE protocols.
- Set up with working volumes.

3. (FC) If you are using a switch, verify that it is:

- Set up correctly.
- Zoned.
- Cabled correctly.

Note: For information about supported topologies and cabling, see the *Data ONTAP SAN Configuration Guide for 7-Mode* for your version of Data ONTAP.

4. Confirm that the host and the storage system can communicate.

Task 2: Install the Host Utilities

1. If you are upgrading the Host Utilities from an earlier version, remove the currently installed version of the Host Utilities.

If you are uninstalling the current version of the Host Utilities or Linux Host Utilities 5.3, you can use the `rpm -e` command. Prior to the Linux Host Utilities 5.3, the software for the Host Utilities was not provided in .rpm format. If you have one of those versions of the Host Utilities installed, you must use the `/uninstall` command.

2. Get a copy of the .rpm file containing the Host Utilities software for your host architecture.

Note: The Linux Host Utilities provide both a 32-bit software installation package and a 64-bit software installation package.

3. Install the Host Utilities software using the `rpm -ivh package_name` command.

Note: Starting with Linux Host Utilities 5.3, if you are upgrading the Host utilities from Linux Host Utilities, you can use the `rpm -Uvh package_name` command.

(iSCSI) Task 3: Configure the iSCSI protocol

1. Record the host's iSCSI node name.
2. Start the iSCSI service.

3. Set up target discovery.
4. **(Optional)** Set up CHAP.
5. Configure the iSCSI service to start automatically.
6. Configure automatic node login.

Task 4: Set up multipathing

1. Supply the correct parameters based on your multipathing system.
If you are using the native DM-Multipath, you must configure the multipath.conf file.
If you are using Veritas Dynamic Multipathing (VxDMP) you must set the appropriate parameters.
2. Verify the configuration.
3. **(Native multipathing only)** Configure DM-Multipath to start automatically during system boots.

Task 5: Set up access between the host and the LUNs on the storage system

1. Discover new LUNs.
2. View the LUNs.
3. Display multipathing information about the LUNs.
4. Enable device persistence for the LUNs.
5. Remove LUNs.

Setup procedures for Intel NIC, QLogic, Brocade, and Emulex HBAs

Linux Host Utilities environments using FC or hardware iSCSI require that you set up your HBAs before you install the Host Utilities. The setup involves making sure that the HBAs and their drivers use the correct parameter values. If you use Intel NIC, QLogic, Brocade, or Emulex HBAs, you will need to change some of these values from the default settings.

(FC) What you need to know about Fibre Channel HBAs and drivers

If your Linux Host Utilities environment uses HBAs, you must install the HBAs before you install the Host Utilities software. There are certain things you need to know about working with the HBAs.

- Make sure you are using a supported HBA and driver.
If your driver firmware is not supported, uninstall it and get a supported version of the firmware. If the HBA is not supported, remove it and get a supported HBA.
Check the interoperability matrix for the most current information about supported HBAs and drivers.
For information on installing and uninstalling the HBAs and drivers, see the instructions provided by the vendor.
- If you have multiple HBAs, make sure they are all the same brand.
You should not use HBAs from different vendors on the same Linux host to access LUNs.

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

- Make sure you are using the correct values for the HBA and driver parameters.
Note: The HBA and driver parameter values recommended in this guide apply only to environments using the FC protocol or a hardware iSCSI HBA, not to environments using the FCoE protocol. In some cases, the values you supply for the parameters differ based on whether the driver was bundled with the Linux operating system distribution or downloaded separately from an OEM.
- For QLogic and Emulex HBAs you can use the default values for the HBA and driver parameters.
- You must install the HBA management packages provided by the vendors on their websites. The management software enables the `sanlun` command to gather information about the FC HBAs, such as their WWPNNs.

Make sure you download the correct package for your host architecture. For `sanlun fcp show adapter` to work, ensure that the following packages are installed:

QLogic HBA

- QConvergeConsole CLI

Emulex HBA

- OneCommand Manager core application (CLI) package for your host

Brocade HBA/CNA

- Brocade Command Utility (BCU), which provides the APIs.

Note: The BCU is part of the Linux Adapter Software Installer.

- libhbaapi
- libhbalinux

Intel NIC

- libhbaapi
- libhbalinux
- If you did not record the WWPN of the HBAs, you can use the `sanlun fcp show adapter` command to list the HBAs.

(iSCSI) Configuring Broadcom 10 Gb iSCSI offload

You can use a Broadcom 10 Gb adapter in iSCSI environments that use iSCSI offload.

Before you begin

Get the Broadcom 10 Gb adapter. For detailed instructions, see the Broadcom documentation.

About this task

iSCSI Offload uses the TCP Offload Engine (TOE) technology in network interface cards (NICs) to offload the processing of the TCP/IP stack to a network controller. Normally, this TOE is used with high-speed network interfaces such as gigabit Ethernet and 10 Gb Ethernet, where the overhead of processing a network stack can become significant.

Steps

1. Install the Linux operating system.

If you are installing the operating system on your machine, follow the standard installation procedure. If you are using a SAN boot LUN, see the instructions in this guide for setting up a SAN boot LUN.

For more information, check the *Broadcom NetXtreme II™ Network Adapter User Guide*.

Note: Do not install the `iscsi-util` package that comes with the operating system distribution. This package is part of the Broadcom driver package and is installed when you install the drivers.

2. Install the Broadcom iSCSI package.
3. Install the appropriate Broadcom drivers. For instructions, see the Broadcom README file that comes with the drivers.
4. Configure the initiator and connect it to the target using the offload path by performing the following steps:
 - a) Add a new target or node by entering the following command:

```
iscsiadm -m node -p <target ipaddr:[port],TPGT> -T <iqn.targetname> -o new
```

Example

The following command line provides an example of adding a new target:

```
iscsiadm -m node -p 192.168.1.250:3260,1 -T iqn.2005-03.broadcom:jho-target -o new
```

- b) Get the Broadcom NX2 transport name by entering the following command:

```
dmes | grep "bnx2i: netif"
```

Example

This command line returns the following type of output on a system that has two NetXtreme II devices. The transport names are shown in bold:

```
bnx2i: netif=eth0, iscsi=bcm570x-030000
bnx2i: netif=eth1, iscsi=bcm570x-030001
```

- c) Bind the iSCSI target to the Broadcom NX2 transport using the following command:

```
iscsiadm -m node -p <target ipaddr:[port]> -T <iqn.targetname> --o=update --name=iface.transport_name --value=$XPORT_NAME
```

Example

The following command line binds the iSCSI target to the Broadcom NX2 transport:

```
iscsiadm -m node -p 192.168.1.250:3260 -T iqn.2005-03.broadcom:jho-target --op=update --name=iface.transport_name --value=bcm570x-030000
```

- d) Log in to the target using the command:

```
iscsiadm -m node -p <target ipaddr:[port]> -T <iqn.targetname> --login
```

Example

The following command line provides an example of logging in to the target:

```
iscsiadm -m node -p 192.168.1.250:3260 -T iqn.2005-03.broadcom:jho-target --login
```

Verifying your configuration

Before you offload the iSCSI sessions, the information has the following format:

```
tcp: [1] 192.168.200.6:3260,1005 iqn.1992-08.com.app:sn.118045890  
tcp: [2] 192.168.200.5:3260,1004 iqn.1992-08.com.app:sn.118045890
```

After you offload the iSCSI sessions, the information has the following format:

```
bcm570x-1a0000: [6] 192.168.200.6:3260,1005 iqn.1992-08.com.app:sn.  
118045890  
bcm570x-1a0000: [5] 192.168.200.5:3260,1004 iqn.1992-08.com.app:sn.  
118045890
```

After you finish

You must get the host to discover the LUNs.

Installing the Linux Host Utilities software

The Linux Host Utilities software is packaged in both a 32-bit and a 64-bit .rpm file. You must get a copy of the appropriate .rpm file either from the IBM website or the physical media. After you have the correct .rpm file for your system, you can install the software.

Before you begin

It is a good practice to confirm that you have the correct version of the Linux Host Utilities for your version of the Linux host. For the most current information about the correct version, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

(Red Hat Enterprise Linux 6) If you are using Red Hat Enterprise Linux 6.0 or later, it is recommended that you run Data ONTAP 8.0.1 or higher on the storage systems connected to the host. That way you can use the Block Limits VPD page (0xb0) information. Using this page helps maintain block alignment, which leads to enhanced I/O performance on the N series LUN.

Steps

1. If you have a version of the Linux Host Utilities currently installed, you must remove it.

If you are uninstalling the Linux Host Utilities 5.3 or higher, you can use the `rpm -e` command.

If you have an earlier version of the Host Utilities installed, go to the directory where the Host Utilities software is installed (the default is `/opt/ontap/santools`) enter the `/uninstall` command.

Note: Prior to the Linux Host Utilities 5.3, the software for the Host Utilities was not provided in .rpm format.

2. Obtain a copy of the .rpm file containing the Host Utilities software. This software is available on physical media and by downloading it.

The Linux Host Utilities provide two versions of the software package:

- A 32-bit version: `ontap_linux_host_utilities-6-2.i386.rpm`
- A 64-bit version: `ontap_linux_host_utilities-6-2.x86_64.rpm`

To download the software, complete the following steps:

- a) Go to the N series support website (accessed and navigated as described in [Websites](#) on page 8).
- b) Get a copy of the .rpm file containing the Host Utilities software.
- c) Sign in with your IBM ID and password.

If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration, click the Register link, follow the online instructions, and then sign in.

- d) Select the N series software you want to download, and then select the Download view.
 - e) Use the “Software Packages ...” link on the Web page presented and follow the online instructions to download the software.
3. Check the publication matrix page for important alerts, news, interoperability details, and other information about the product before beginning the installation.
 4. Obtain the product software from the physical media kit or from the version you downloaded.
 5. Launch the software installation program from the CD-ROM or from where you downloaded the software, and then follow the prompts.

(Red Hat) Configuring space reclamation on hosts running Red Hat Linux 6.2

If you are running Red Hat Enterprise Linux 6 update 2 on your host and Data ONTAP 8.1 or later on the storage systems, you can set up your system to use space reclamation.

Before you begin

The storage systems must be running Data ONTAP 8.1 or later.

About this task

The following steps take you through the process of configuring space reclamation. For more details, see the Red Hat Enterprise Linux 6 *Storage Administration Guide* information on tuning considerations for Solid-State Disk Deployment. At the time this guide was created, this information was located at [Chapter 20. Solid-State Disk Deployment Guidelines](#)

Steps

1. On the storage controller, enable the LUN to report provisioning threshold events by entering the following command on the LUN:
`-e space_alloc`
2. If you have not already done so, map the LUN to the host running Red Hat Enterprise Linux 6 update 2.
3. Configure DM-Multipath on the LUN so that it shows up as a DM-Multipath device.
4. Mount the DM-Multipath device using the `discard` option of the `ext4` command.

Example

The following is an example of mounting the DM-Multipath device:

```
mount -t ext4 -o discard /dev/mapper/  
360a98000572d4461524a64793868696d /mnt
```

(Oracle VM) Configuring the O2CB_HEARTBEAT_THRESHOLD

When you run Oracle VM with multipathing in a clustered host environment, you need to configure the O2CB_HEARTBEAT_THRESHOLD to 65.

About this task

This timeout value works with the I/O layer and sets the maximum amount of time in seconds that a node waits for an I/O operation. In the heartbeat system file, the node writes to the block every two seconds. The block offset is equal to its global node number. This means that node 0 writes to the first block, node 1 to the second block, and so on. All the nodes also read the heartbeat sysfile every two seconds. As long as the timestamp keeps changing, that node is considered "alive."

Steps

1. On each hypervisor, run the command:

```
service o2cb configure
```

2. Set O2CB_HEARTBEAT_THRESHOLD to 65 by entering:

```
O2CB_HEARTBEAT_THRESHOLD = 65
```

Heartbeat settings and multipath settings should be exactly same on all the hypervisors in a server pool.

3. Either restart the O2CB service or reboot your system.

(iSCSI) How to configure iSCSI for Linux

The iSCSI initiator software comes with the operating system distribution. Its components include a kernel module, which is already compiled into the Linux kernel, and user space packages. For it to work properly, you must perform configuration tasks.

These tasks include the following:

- Record the host's initiator node name.
You must supply the node name when you set up igroups on the storage system. Once the igroup is created, you map LUNs to it so that all the hosts associated with that igroup can discover the LUNs.
It is a good practice to change the initiator node name from the default name to one that includes the host name. That way the name is easier to use.
- If you want to use multipathing, you must edit the iSCSI configuration file to set it up.
- If you want to use CHAP, you must edit the iSCSI configuration file to set it up.
- You must set up target discovery so that the host can access LUNs on the storage system.
- Configure the initiator with the IP address for each storage system using either static, ISNS, or dynamic discovery.

The following sections explain how to perform these tasks.

(iSCSI) Getting the initiator node name

You must supply the iSCSI initiator node name when you set up igroups. To make this process go smoothly, it is a good practice to write down the node name now. Before you record the node name, you might want to change it. By default, a node name ends with a string of random numbers. Changing this part of the name to something such as the host name can make the node name easier to use.

About this task

Each node name must be unique. Do not change multiple node names to the same name.

If you decide to change the node name, you can only change the *RandomNumber* portion of the name that appears at the end. The first part of a node name starts with `iqn.` and uses a standard format that you cannot change. For example, if your node name is `iqn.2005-03.com.RedHat:012345`, you could change the last six numbers to a word or phrase so that the name becomes `iqn.2005-03.com.RedHat:Toaster`.

Steps

1. Use a text editor to open the file containing the node names.

If you are using...	Open the following file...
Red Hat Enterprise Linux 6 series and 5 series, SUSE Linux Enterprise Server 10 series, or SUSE Linux Enterprise Server 11	<code>/etc/iscsi/initiatorname.iscsi</code>
Red Hat Enterprise Linux 4 series	<code>/etc/initiatorname.iscsi</code>

2. **(Optional)** If you want to change the default name, edit the line in the file containing the node name.

You can only replace the *RandomNumber* portion of the name. When you replace this, make sure you follow the node name rules:

- A node name can be up to 223 characters.
- 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 (-)

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

- Colon (:)

Note: If the node name does not exist, you can create one by adding a line to the file containing node names. Use the same format shown below for modifying a node name.

If you are using...	Then...
Red Hat Enterprise Linux 6 series and 5 series, Oracle VM, and Red Hat Enterprise Virtualization	<p>Locate the node name (<code>InitiatorName=iqn.2005-03.com.RedHat:RandomNumber</code>) you want to modify and change the <i>RandomNumber</i> part of the name to a unique value.</p> <p>The following line shows an example of a modified node name:</p> <p><code>InitiatorName=iqn.2005-03.com.RedHat:linux-host1</code></p>
Red Hat Enterprise Linux 4 series	<p>Locate the node name (<code>InitiatorName=iqn.1987-05.com.cisco:RandomNumber</code>) you want to modify and change the <i>RandomNumber</i> part of the name to a unique value.</p> <p>The following line shows an example of a modified node name:</p> <p><code>InitiatorName=iqn.1987-05.com.cisco:linux-host1</code></p>

If you are using...	Then...
SUSE Linux Enterprise Server 10 series or Server 11	<p>Locate the node name (<code>InitiatorName=iqn.1996-04.de.suse:RandomNumber</code>) you want to modify and change the <i>RandomNumber</i> part of the name to a unique value.</p> <p>The following line shows an example of a modified node name:</p> <pre>InitiatorName=iqn.1996-04.de.suse:linux-host1</pre>
Citrix XenServer 6.0 and 6.0.2	<p>The <code>initiatorname.iscsi</code> file in <code>/etc/iscsi</code> is not present in the earlier version of Citrix XenServer releases. To set the IQN, use the following command:</p> <pre>xe host-param-set uuid=UUID_of_the_host other-config:iscsi_iqn=IQN_In_the_Standard_format.</pre> <p>The following line shows an example of a modified IQN:</p> <pre>xe host-param-set uuid=163c53be-8de5-4035-8770-fbe012ab1f56 other-config:iscsi_iqn=iqn.2013-06.com.example:RanjanTest123</pre>

3. Write down the node name so that you can easily enter it when you configure the storage system.
4. If you modified the file, make sure you save it before you close it.

(iSCSI) Setting the timeout values to enable multipathing

If you use multipathing, you need to edit the timeout value in the iSCSI configuration file.

Step

1. To use multipathing, edit the following file to provide the correct timeout value for your Host Utilities environment (DM-Multipath or Veritas Storage Foundation):

If you are using...	Enter the following values...
Red Hat Enterprise Linux 6 series or 5 series, SUSE Linux Enterprise Server 10 series or 11 series, SUSE Linux Enterprise Server KVM and XEN, and Red Hat Enterprise Linux KVM.	<p>Edit <code>/etc/iscsi/iscsid.conf</code> to include the following value:</p> <ul style="list-style-type: none"> • DM-Multipath environments: <pre>node.session.timeo.replacement_timeout = 5</pre> • Veritas Storage Foundation environments: <pre>node.session.timeo.replacement_timeout = 120</pre>
Oracle VM 3.0.1 and 3.1.1	<p>Edit <code>/etc/iscsi/iscsid.conf</code> to include the following value:</p> <pre>node.session.timeo.replacement_timeout = 5</pre>

If you are using...	Enter the following values...
Red Hat Enterprise Linux 4 series	<p>Edit <code>/etc/iscsi.conf</code> to include the following DM-Multipath value:</p> <ol style="list-style-type: none">Remove the comment indicator from the <code>ConnFailTimeout</code> line in the Session Timeout Settings section.Set the value of <code>ConnFailTimeout</code> to 5
Citrix XenServer 6.0 and 6.0.2	<p>For existing Storage Repositories:</p> <ol style="list-style-type: none">Edit <code>/etc/iscsi/iscsid.conf</code> to include the following value:<ul style="list-style-type: none">DM-Multipath environments: <code>node.session.timeo.replacement_timeout = 5</code>Detach and then reattach SRs from XenCenter or by using XenServer <code>xe</code> CLI. This will reflect the new iSCSI timeout for the existing SRs. <p>For new Storage Repositories:</p> <ol style="list-style-type: none">Edit <code>/etc/iscsi/iscsid.conf</code> to include the following value:<ul style="list-style-type: none">DM-Multipath environments: <code>node.session.timeo.replacement_timeout = 5</code>Create the new SR. New as well as existing SRs will be updated with new iSCSI timeout settings.
Red Hat Enterprise Virtualization Hypervisor 6.2 and 6.3	<ol style="list-style-type: none">Edit <code>/etc/iscsi/iscsid.conf</code> to include the following value:<ul style="list-style-type: none">DM-Multipath environments: <code>node.session.timeo.replacement_timeout = 5</code>Execute the following command to ensure your <code>iscsid.conf</code> file is persistent across reboots: <code>persist /etc/iscsi/iscsid.conf</code>

(iSCSI) Setting up CHAP for Red Hat Linux 5, 6 and SUSE Linux 10, 11

You can use the CHAP protocol on hosts running Red Hat Enterprise Linux 5 and 6 series and SUSE Linux Enterprise Server 10 and 11 series to provide enhanced security. To set up CHAP, you must add CHAP user names and passwords to the `/etc/iscsi/iscsid.conf` file. To complete the

setup, you must use the `iscsi security` command to set up the same user names and passwords on the storage system.

Steps

1. Open the `/etc/iscsi/iscsid.conf` file with a text editor.
2. Enable CHAP authentication by setting `node.session.auth.authmethod` to CHAP
The default is `None`.
3. Provide a CHAP user name and password for the target to use when authenticating the initiator. You must remove the comment indicators and supply values for the options `username` and `password` in the following configuration entries:

```
node.session.auth.username = username
node.session.auth.password = password
```

4. Provide a CHAP user name and password for the initiator to use when authenticating the target. You must remove the comment indicators and supply values for the options `username_in` and `password_in` in the following configuration entries:

```
node.session.auth.username_in = username_in
node.session.auth.password_in = password_in
```

5. For a successful session discovery, enable discovery CHAP authentication by supplying the passwords in the `discovery.sendtargets.auth.options`.

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 that you used when you set up CHAP on the storage system with the `iscsi security` command.

```
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
```

(iSCSI) Setting up CHAP for Red Hat Enterprise Linux 4 series

You can use the CHAP protocol on hosts running Red Hat Enterprise Linux 4 series to provide enhanced security. To set up CHAP, you need edit the `/etc/iscsi.conf` file to add CHAP user

names and passwords. To complete the setup, you must use the `iscsi security` command to set up the same user names and passwords on the storage system.

Steps

1. Open the `/etc/iscsi.conf` file with a text editor.
2. Add CHAP user names and passwords to the storage system's `DiscoveryAddress` section. Use a white space or tab to indent the CHAP settings.

You can set up CHAP as either unidirectional authentication or bidirectional authentication.

- 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 (*iname* and *inpassword*).
- For bidirectional authentication, you should define both sets of user names/passwords: outgoing and incoming.
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`.

Example

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

3. 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: This example specifies two targets. You now need to edit the sections under the targets to add the user names and passwords.

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


(iSCSI) Starting the iSCSI service

After you edit the iSCSI configuration file, you must start the iSCSI service so that the host can discover the targets and access LUNs. If the iSCSI service is running, you must restart it.

Step

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

If you are using...	Enter...
Red Hat Enterprise Linux 6 series	<code>service iscsid force-start</code>
Note: You must execute this command the first time you start the iSCSI service on a host running Red Hat Enterprise Linux 6 series and Oracle Linux 6 series. If you execute <code>/etc/init.d/iscsi start</code> without previously executing <code>service iscsid force-start</code> , you will get an error message.	
Red Hat Enterprise Linux 5 or 4 series and Oracle Linux 5 series, Red Hat Enterprise Virtualization, Oracle Linux, and Oracle VM	<code>/etc/init.d/iscsi start</code>
SUSE Linux Enterprise Server 10 or 11 series	<code>/etc/init.d/open-iscsi start</code>

Citrix discourages the use of `iscsiadm` tool. The native XAPI stack accomplishes the tasks of starting and stopping the `iscsi` service, automatic login on boot, and other iSCSI operations.

(iSCSI) Methods for setting up target discovery with software initiators

You need to configure the iSCSI initiator to discover the target so that the host can access LUNs on the target. The method you use to do this depends on your version of the operating system.

- Red Hat Enterprise Linux 5 and 6 series: use the `iscsiadm` utility.
- Red Hat Enterprise Linux 4 series: modify the `/etc/iscsi.conf` file.
- SUSE Linux Enterprise Server 10 and 11 series: use either the `iscsiadm` utility or YaST2.

Note: Use Management GUI for setting up target discovery for RHEV, Oracle VM, and Citrix XenServer.

The following sections provide instructions for setting up targets on Red Hat Enterprise Linux 5 and 6 series, Red Hat Enterprise Linux 4 series, and SUSE Linux Enterprise Server 10 and 11 series.

Note: If you are using SUSE Linux Enterprise Server 9 series, see the *Linux Host Utilities Installation and Setup Guide* at www.ibm.com/storage/support/nseries/.

(iSCSI) Discovering the iSCSI target by using the `iscsiadm` utility on Red Hat 5, 6, SUSE 10, 11

You can use the `iscsiadm` utility to manage (update, delete, insert, and query) the persistent database on Red Hat Enterprise 5 or 6 series, SUSE Linux Enterprise Server 10 or 11 series. 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 iSCSI target:

```
iscsiadm --mode discovery --op update --type sendtargets --portal  
targetIP
```

targetIP is the IP address of the target.

When you enter the `iscsiadm` command, the host discovers the target specified by the *targetip* variable. The `iscsiadm` utility displays each target it discovers on a separate line. It stores the values associated with the target in an internal persistent database.

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

```
iscsiadm --mode node -l all
```

The initiator logs into 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 of the output this command produces:

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

022e71 is the record ID.

(iSCSI) Setting up target discovery on Red Hat Enterprise Linux 4 series

When you are using Red Hat Enterprise Linux 4 series, you can set up target discovery by editing the `/etc/iscsi.conf` file and adding the IP addresses of the storage systems you want to use as targets.

Steps

1. Open the `/etc/iscsi.conf` file with a text editor.
2. Configure the storage system as a target by adding the following line for any one iSCSI-enabled interface on each storage system that you used for iSCSI LUNs:

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

The following lines set up the storage systems with the IP addresses 192.168.10.100 and 192.168.10.20 as targets:

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

(iSCSI) Discovering targets by using YaST2 on SUSE 10, 11

If you are running SUSE Linux Enterprise Server 10 series or SUSE Linux Enterprise Server 11, you can use YaST2 to discover and configure iSCSI connections. By 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. If that is not the port number, then change the port number to 3260.
4. Enter the IP address of the iSCSI target.
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 required for using 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 website.

Related information

[Novell website](#)

(iSCSI) 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	<code>chkconfig iscsi on</code>
SUSE Linux Enterprise Server 10 or 11 series	<code>chkconfig open-iscsi on</code>

(iSCSI) Configuring manual or automatic node login

When you are running Red Hat Enterprise Linux 6 or 5 series or SUSE Linux Enterprise Server 11 or 10 series, you can specify whether the system automatically logs in to an iSCSI node at startup or whether you must manually log it in to the node.

Before you begin

Make sure the iSCSI service is running when the logins are supposed to occur.

About this task

If you set your login mode to manual, you will need to log in to the nodes manually the next time the system starts up. If you set your login mode to automatic, the system will log in to the nodes automatically when it starts up.

Note: (Red Hat Enterprise Linux 4 series) When you are running Red Hat Enterprise Linux 4 series, all sessions are logged in automatically when you start the iSCSI service.

Setting the login mode affects only nodes that are discovered after the value is set.

Step

1. To set the login mode for a specific portal on a target or for all the portals on a target, use the `iscsiadm` command.

To set a login mode for all targets and their ports, edit the `/etc/iscsi/iscsid.conf` file.

The `iscsiadm` command has more than one format. Use the format that applies to your system. You must supply information such as the target name as well as whether the login will be `manual` or `automatic`.

Note: For more information on the `iscsiadm` options, see the man page.

To set the login mode for...	Do the following...
A specific port on a target	<p>Enter the appropriate command for your system:</p> <pre>iscsiadm --mode node -T <i>targetname</i> -p <i>ip:port</i> -o update -n node.startup -v manual automatic</pre> <pre>iscsiadm --mode node -T <i>targetname</i> -p <i>ip:port</i> -o update -n node.conn[0].startup -v manual automatic</pre>
All the ports on a target	<p>Enter the appropriate command for your system:</p> <pre>iscsiadm --mode node -T <i>targetname</i> -o update -n node.startup -v manual automatic</pre> <pre>iscsiadm --mode node -T <i>targetname</i> -o update -n node.conn[0].startup -v manual automatic</pre>
All the targets	<ol style="list-style-type: none"> <li data-bbox="400 835 1201 895">a. Modify the <code>/etc/iscsi/iscsid.conf</code> file to add the following line. You must specify either <code>manual</code> or <code>automatic</code>: <pre>node.startup = manual automatic</pre> <li data-bbox="400 956 709 979">b. Rediscover the iSCSI target. <li data-bbox="400 1003 682 1025">c. Restart the iSCSI service.

Configuring the storage system

You need to configure the storage system so that the protocol you are using can access it. For example, if your environment uses FC, you must supply the host's WWPN and make sure the storage system is using a supported cfmode. For environments using the iSCSI protocol, you must supply the names of the initiator nodes and, if you are using CHAP, set up the user names and passwords on the storage system.

Before you begin

Make sure you have the host's WWPNs that you recorded when you installed the FC HBAs or the iSCSI initiator nodes that you recorded.

Note: If you are using Red Hat Enterprise Linux 6.0 or later, it is recommended that you run Data ONTAP 8.0.1 or higher on the storage systems connected to the host. That way you can use the Block Limits VPD page (0xb0) information. Using the Block Limits VPD page helps maintain block alignment, which leads to enhanced I/O performance on the N series LUN.

About this task

This checklist steps you through the requirements for making sure your storage system is correctly set up.

Steps

1. Ensure that the protocol you are using (FC or iSCSI) is licensed and the service is running.
2. **(iSCSI)** If you want to use CHAP authentication, use the `iscsi security` command or the FilerView interface to configure the CHAP user name and password on the storage system. Make sure you use the same user names and passwords that you supplied when you set up CHAP on the host.
3. Create an igroup that includes the Linux host.

Note: You can use the `sanlun fcp show adapter -c` command to get the information necessary to create an igroup on the controller. The information supplied by this command can only be used to create an igroup if the controller is running Data ONTAP operating in 7-Mode.

4. Create and map the LUNs to the igroup.

Make sure that you specify the LUN operating system type and igroup type as `linux`.

5. Optionally, enable ALUA.

For clustered Data ONTAP, ALUA works with the FC, FCoE and iSCSI protocols. For Data ONTAP operating in 7-Mode, ALUA works with the FC and FCoE protocols. For information on which versions of Linux support ALUA, see the section *Linux configurations that support ALUA*.

Note: ALUA must be enabled if you are using clustered Data ONTAP.

6. If you are using clustered Linux hosts, ensure that the igroup contains either the WWPNs or 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 for use in multipathing in environments that use native Linux solutions. With DM-Multipath, you can configure multiple I/O paths between a host and storage controllers into a single device. If one path fails, DM-Multipath reroutes I/Os to the remaining paths.

Note: If you are running Veritas Storage Foundation, you need to use VxDMP as your multipathing solution.

When you have multiple paths to a LUN, Linux creates a SCSI device for each path. This means that a single LUN might appear as `/dev/sdd` and `/dev/sdf` if there are two paths to it. To make it easy to keep track of the LUNs, DM-Multipath creates a single device in `/dev/mapper/` for each LUN that includes all the paths. For example, `/dev/mapper/360a9800043346852563444717a513571` is the multipath device that is created on top of `/dev/sdd` and `/dev/sdf`.

When you are using DM-Multipath, you should create a file system for each LUN and then mount the LUN using the device in `/dev/mapper/`.

Note: To create a file system on a LUN, use `/dev/mapper/device` on a Linux host console.

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

You also use the DM-Multipath's configuration file to specify whether ALUA is being used and if the hardware handler should be enabled for ALUA.

When DM-Multipath is running, it automatically checks the paths. As a result, if you look at the output of a command such as `lun stats -o`, you see a small amount of FC partner path traffic listed under the operations per second. On average, this is usually about 4 kb of operations per path per LUN every 20 seconds, which is the default time period. This is expected behavior for DM-Multipath.

Verifying the required multipathing packages

It is a good practice to verify that you have the multipathing packages that enable you to use DM-Multipath. These packages are part of your Linux distribution.

Steps

1. Use the `rpm -q` command to display information about the name and version of the DM-Multipath package that you have installed.

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

2. If you do not have the required packages, get a copy of your operating system RPM and install the multipathing package.

Related information

[Red Hat website](#)

[Novell website](#)

Editing the DM-Multipath configuration file

For DM-Multipath to function properly, you must edit the `/etc/multipath.conf` configuration file.

Steps

1. If the `/etc/multipath.conf` file exists, edit it to include the sections needed for your system.
2. If you do not have the `/etc/multipath.conf` file, copy the sample configuration file for your operating system. The following sections provide sample configuration files for several versions of Red Hat Enterprise Linux and SUSE Linux Enterprise Server.
3. After you have the file, you can specify which devices you want to exclude (blacklist).

You should exclude all the devices that do not correspond to LUNs configured on the storage system that are mapped to your Linux host. That is, exclude those devices that are not displayed by the `sanlun lun show` command.

Note: You **must** supply information that is specific to your system. Otherwise, you will encounter problems.

The following steps explain what you need to do to exclude a device:

- a) In the `blacklist` section of the configuration file, enter the WWID of all non-IBM N series SCSI devices installed on your host.

You can get the WWID by running the `scsi_id` command on a device.

Example

For example, assume that `/dev/sda` is a local SCSI drive. To obtain the WWID on systems running Red Hat Enterprise Linux 6 series or SUSE Linux Enterprise Server 11, enter:

```
/lib/udev/scsi_id -gud /dev/sda
```

To obtain the WWID on systems running other Linux operating systems, enter:

```
scsi_id -gus /block/sda
```

In both cases, the output looks similar to the following:

```
SIBM-ESXSMAW3073NC_FDAR9P66067W
```

To exclude that device, enter SIBM-ESXSMAW3073NC_FDAR9P66067W in the `blacklist` section of the configuration file:

```
blacklist
{
    wwid IBM-ESXSMAW3073NC_FDAR9P66067W
    devnode "^hd[a-z]"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
```

- b) Exclude other devices by using the devnode, such as:

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

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

```
devnode "^cciss.*"
```

Example

On Red Hat Enterprise Linux 4 hosts, the `blacklist` section might appear as the following:

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

4. Make sure you use the correct settings based on whether you are using ALUA.

If you are using ALUA, you must specify the ALUA callout program. If you are not using ALUA, you must specify the Data ONTAP callout program. The following table provides information the values that you must supply.

Note: If you are using clustered Data ONTAP, you must have ALUA enabled.

If you are running...	With ALUA, use the value	Without ALUA, use the value
SUSE Linux Enterprise Server 10 SP2 and later, SUSE Linux Enterprise Server 11 and later, or Red Hat Enterprise Linux 6.0 and later	Set <code>prio</code> to: "alua"	Set <code>prio</code> to: "ontap"
Any other Red Hat Enterprise Linux operating system	Set <code>prio_callout</code> to: /sbin/mpath_prio_alua	Set <code>prio_callout</code> to: /sbin/mpath_prio_ontap
SUSE Linux Enterprise Server 10 SP1 and SP2	~	Set <code>prio_callout</code> to: /sbin/mpath_prio_ontap
All supported Linux operating systems that support ALUA	Set <code>hardware_handler</code> to: "1 alua"	Set <code>hardware_handler</code> to: "0"

Note: ALUA is supported in Red Hat Enterprise Linux 5 Update 1 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later.

5. Save the changes.

Related references

[Sample configuration files for Red Hat Enterprise Linux 6](#) on page 111
[Sample configuration file for Red Hat Enterprise Linux 5](#) on page 119
[Sample configuration file for Red Hat Enterprise Linux 4](#) on page 124
[Sample configuration file for Red Hat Enterprise Virtualization Hypervisor](#) on page 127
[Sample configuration file for Oracle VM 3.0.1, 3.1.1, and 3.2.1](#) on page 131
[Sample configuration files for SUSE Linux Enterprise Server 11 series](#) on page 137
[Sample configuration file for SUSE Linux Enterprise Server 10](#) on page 141

Starting DM-Multipath

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

Steps

1. To start the multipath daemon, enter the `start` command line:

If you are using...	Enter the following command...
Red Hat Enterprise Linux, Red Hat Enterprise Virtualization, and Oracle VM	<code># /etc/init.d/multipathd start</code>
SUSE Linux Enterprise Server	<code># /etc/init.d/boot.multipath start</code> <code># /etc/init.d/multipathd start</code>

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

```
# multipath
```

3. Perform the following steps to enable multipath using Xen center or Xen CLI:

- a) Enter maintenance mode:

```
xe host-disable uuid=<host UUID>
```

- b) Enable multipath:

```
xe host-param-set other-config:multipathing=true uuid=host_uuid
```

- c) xe host-param-set other-config:

```
multipathhandle=dmp uuid=host_uuid
```

- d) Release maintenance mode:

```
# xe host-enable uuid=<host UUID>
```

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 following action on the Linux host console:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	<code>chkconfig --add multipathd</code> <code>chkconfig multipathd on</code> Note: You should reboot the host if you are configuring a SAN boot LUN on the host.
SUSE Linux Enterprise Server	<code>chkconfig --add boot.multipath</code> <code>chkconfig --add multipathd</code> <code>chkconfig boot.multipath on</code> <code>chkconfig multipathd on</code>

Note: Hypervisors Oracle VM, Red Hat Enterprise Virtualization, and Citrix XenServer management stack ensures automatic start of multipath service.

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
```

Note: To determine if `multipathd` is working correctly on your system, enter the `multipathd show config` command. This command displays the values currently being used for the `multipath.conf` file. You can then confirm that `multipathd` is using the values you specified.

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

```
multipath -ll
```

Example

(FC) The output looks similar to the following for FC:

```
# multipath -ll
3360a98000486e5372635a44646a505643Ontap,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

(FC and clustered Data ONTAP) If the storage system is running clustered Data ONTAP and using iSCSI, the output looks similar to the following:

```
# multipath -ll
360a98000486e2f66426f2f7a32745338 dm-0 Ontap,LUN C-Mode
```

```
size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-- policy='round-robin 0' prio=50 status=active
|   |-- 3:0:1:1 sdd 8:48 active ready running
|   |-- 3:0:4:1 sdh 8:112 active ready running
|   |-- 2:0:0:1 sdj 8:144 active ready running
|   |-- 2:0:2:1 sdn 8:208 active ready running
|-- policy='round-robin 0' prio=10 status=enabled
|   |-- 3:0:0:1 sdb 8:16 active ready running
|   |-- 3:0:2:1 sdf 8:80 active ready running
|   |-- 2:0:1:1 sdl 8:176 active ready running
|   |-- 2:0:3:1 sdq 8:240 active ready running
```

Example

(FC on Red Hat Enterprise Linux 6, SUSE Linux Enterprise Server 11 SP1 or later) The output looks similar to the following when you have ALUA enabled and are using either Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 with SP1 or later:

```
# multipath -ll
360a98000486e2f66426f2f7a32745338 dm-0 Ontap,LUN
size=1.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-- policy='round-robin 0' prio=50 status=active
|   |-- 3:0:1:1 sdd 8:48 active ready running
|   |-- 3:0:4:1 sdh 8:112 active ready running
|   |-- 2:0:0:1 sdj 8:144 active ready running
|   |-- 2:0:2:1 sdn 8:208 active ready running
|-- policy='round-robin 0' prio=10 status=enabled
|   |-- 3:0:0:1 sdb 8:16 active ready running
|   |-- 3:0:2:1 sdf 8:80 active ready running
|   |-- 2:0:1:1 sdl 8:176 active ready running
|   |-- 2:0:3:1 sdq 8:240 active ready running
```

Example

(FC on Red Hat Enterprise Virtualization Hypervisor 6.2) The output looks similar to the following when you have ALUA enabled and are using Red Hat Enterprise Virtualization Hypervisor 6.2:

```
# multipath -ll
360a98000316b5a776b3f2d7035505a6f dm-0 Ontap,LUN
size=60G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-- policy='round-robin 0' prio=50 status=active
|   |-- 4:0:0:0 sda 8:0 active ready running
|   |-- 5:0:0:0 sdc 8:32 active ready running
|-- policy='round-robin 0' prio=10 status=enabled
|   |-- 4:0:1:0 sdb 8:16 active ready running
|   |-- 5:0:1:0 sdd 8:48 active ready running
```

Example

(iSCSI and clustered Data ONTAP) If the storage system is running clustered Data ONTAP and using iSCSI, the output looks similar to the following:

```
# multipath -ll
3600a09803246676c703f2d6c36506138 dm-0 Ontap,LUN C-Mode
size=25G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
|  `-- 8:0:0:0   sdc 8:32 active ready running
|+- policy='round-robin 0' prio=10 status=enabled
|  `-- 9:0:0:0   sdd 8:48 active ready running
|  `-- 7:0:0:0   sdb 8:16 active ready running
|  `-- 10:0:0:0  sde 8:64 active ready running
```

Example

(iSCSI on Red Hat Enterprise Linux 6, SUSE Linux Enterprise Server 11 SP1 or later) The output looks similar to the following when you are using iSCSI with either Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11 with SP1 or later:

```
# multipath -ll
360a9800050334b50694a625563326f37 dm-3  Ontap,LUN
size=9.0G features='1 queue_if_no_path' hwhandler='0' wp=rw
`+- policy='round-robin 0' prio=2 status=active
|  `-- 4:0:0:1   sdc 8:32  active ready running
|  `-- 3:0:0:1   sdb 8:16  active ready running
```

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	<code>/etc/init.d/multipathd stop</code>
SUSE Linux Enterprise Server	<code>/etc/init.d/multipathd stop</code> <code>/etc/init.d/boot.multipath stop</code>

Veritas Dynamic Multipath configuration

When you are using Veritas Dynamic Multipathing (VxDMP), you need to perform some configuration tasks to enable multipathing.

There are several tunable parameters that you might need to set, depending on your version of Linux. These parameters might include:

- `vxmpadm start restore interval`
- `vxmpadm settune dmp_lun_retry_timeout`
- `vxmpadm settune dmp_path_age`

In addition, if you are running either Red Hat Enterprise Linux 6 series or SUSE Linux Enterprise Server 11 series, you must modify the `/etc/udev/rules.d/40-rport.rules` file.

You must have the Array Support Library (ASL) and the Array Policy Module (APM) that Symantec provides for storage systems installed. The amount of work you need to do to set up the ASL and APM depends on your version of Veritas Storage Foundation.

The following sections describe the configuration work necessary to set up VxDMP.

(Veritas) VxDMP restore daemon and LUN retries tunable configuration

It is a good practice to set the Veritas restore daemon values for the restore policy and the polling interval and the VxDMP tunable `dmp_lun_retry_timeout` to the Host Utilities recommended values.

The restore policy and the polling interval settings determine how frequently the Veritas daemon checks paths between the host and the storage system. At the time this document was produced, the Host Utilities recommended the following restore daemon settings:

- A restore policy of "disabled"
- A polling interval of "60"

Check the *Release Notes* to see if these recommendations have changed.

The tunable `dmp_lun_retry_timeout` tells the VxDMP to continue retrying I/O requests to a LUN where all the paths to the disk have failed. Setting this value provides for more robust handling of temporary path failures that are quickly restored. The recommended value for this tunable is 300, which means that the VxDMP continues to retry paths to the LUN until either the I/O succeeds or 300 seconds have elapsed.

(Veritas) Setting the restore daemon and LUN retry tunable values

To access LUNs by using VxDMP, you should configure the restore daemon and verify that the daemon value has been successfully configured. If you have Veritas Storage Foundation 5.1, you should also set the VxDMP tunable `dmp_lun_retry_timeout` command.

Steps

1. Set the value for the VxDMP restore daemon to an interval of 60 by entering the following command:

```
vxdmptune dmp_restore_interval=60
```

On reboot, this value takes 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 takes effect immediately; however, it is not persistent on reboot. You must reset the value each time you reboot the system.

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.

3. **(Veritas Storage Foundation 5.1 till 6.0)** Set the value of the `dmp_lun_retry_timeout` to an interval of 300 by entering the following command:

```
vxdmpadm settune dmp_lun_retry_timeout=300
```

The tunable value changes immediately.

(Veritas Storage Foundation 6.0.x series and above) Set the value of the `dmp_lun_retry_timeout` to an interval of 60 by entering the following command:

```
vxdmpadm settune dmp_lun_retry_timeout=60
```

The tunable value changes immediately.

4. **(Veritas Storage Foundation 5.1 SP1 and later)** Set the value of the `dmp_path_age` to an interval of 120 by entering the following command:

```
vxdmpadm settune dmp_path_age=120
```

The tunable value changes immediately.

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

```
vxdmpadm gettune
```

This command displays the current values for all the Veritas DMP tunables.

(Veritas, Red Hat) Configuring Red Hat 6 to support Veritas Storage Foundation

If you are using Red Hat Enterprise Linux 6 or later, you must configure it to support Veritas Storage Foundation.

About this task

Note: If you are running Red Hat Enterprise Linux 6 series, you must perform some additional steps and make sure the value of the IOFENCE timeout parameter is correct.

Steps

1. Log in to the host as the root user.
2. Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add", RUN
+="/bin/sh -c 'echo 20 > /sys/class/fc_remote_ports/%k/
fast_io_fail_tmo;echo 864000 >/sys/class/fc_remote_ports/%k/
dev_loss_tmo'"
```

3. **(Red Hat Enterprise Linux 6 or later only)** Check the value of the IOFENCE timeout parameter is make sure it is set to 30000.

The IOFENCE timeout parameter specifies the amount of time in milliseconds that it takes clients to respond to an IOFENCE message before the system halts. When clients receive an IOFENCE message, they must unregister from the GAB driver within the number of milliseconds specified by the IOFENCE timeout parameter. If they do not unregister within that time, the system halts. The default value for this parameter is 15000 milliseconds or 15 seconds.

Example

To check the value of this parameter, enter the command `gabconfig -l` on the host. The following is an example of the type of output this command produces.

```
# gabconfig -l
GAB Driver Configuration
Driver state           : Configured
Partition arbitration: Disabled
Control port seed      : Enabled
Halt on process death: Disabled
Missed heartbeat halt: Disabled
Halt on rejoin         : Disabled
Keep on killing         : Disabled
Quorum flag            : Disabled
Restart                : Enabled
Node count              : 2
Send queue limit        : 128
```

```
Recv queue limit      : 128
IOFENCE timeout (ms)  : 15000
Stable timeout (ms)   : 5000
```

4. If the value of the IOFENCE timeout parameter is not 30000, change it by entering the command:

```
gabconfig -f 30000
```

This value is not persistent across reboots, so you must check it each time you boot the host and reset it if necessary.

(Veritas, SUSE) Configuring SUSE Linux 11 to support Veritas Storage Foundation

If you are using SUSE Linux Enterprise Server 11 or later, you must configure it to support Veritas Storage Foundation.

Before you begin

Check Symantec TechNote 124725 for the latest update on using SUSE Linux Enterprise Server 11 with Veritas Storage Foundation.

Steps

1. Log in to the host as the root user.
2. Install SUSE Linux Enterprise Server 11 with kernel version 2.6.27.45-0.1.1 or later from Novell.
3. Create the file `/etc/udev/rules.d/40-rport.rules` with the following content line:

```
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add",RUN
+="/bin/sh -c 'echo 20 > /sys/class/fc_remote_ports/%k/
fast_io_fail_tmo;echo 864000 >/sys/class/fc_remote_ports/%k/
dev_loss_tmo'"
```

4. Reboot the host.

Related information

[Novell website](#)

[Symantec TechNote for setting up SUSE Linux Enterprise Server 11](#)

(Veritas) The Array Support Library and Array Policy Module

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

Note: These are Symantec products; therefore, Symantec provides technical support if you encounter a problem using them.

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

(Veritas) What the ASL is

The ASL is a Data ONTAP-qualified library that provides information about storage array attributes and multipathing configurations to the Device Discovery Layer (DDL) and Veritas 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, disk array, or Vserver. 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.

(Veritas) ASL array type

The ASL reports information about the multipathing configuration to the DDL as an Active/Active (A/A), 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.
- **ALUA**
A LUN in an ALUA-enabled array can be accessed through both controllers, by using optimized and non-optimized paths. The array notifies 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*.

(Veritas) 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 APM is customized to optimize I/O error handling and failover path selection for the N series environment. After the ASL discovers the storage array as an N series array, the ASL instructs VxDMP to use the APM to handle I/O error processing and path failures for the N series storage array.

(Veritas) Installing the 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.
- Download 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 website. See the appropriate interoperability matrix at N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8). You should download this software before you start the installation.

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

About this task

Only one version of the 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 storage configured as JBOD in your VxVM configuration, remove the JBOD support for storage by entering the following command:

```
# vxddladm rmjbod vid=IBM
```

3. Install the ASL and APM according to the instructions provided by Symantec.
4. If your host is connected to a storage system, verify the installation by following these steps:
 - a. Run the following command:

```
# vxdlmpadm listenclosure all
```

The output shows the model name of the storage device if you are using enclosure- based naming with VxVM.

(Veritas) 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. Use the `rpm` command to remove the ASL package. This command has the format: `rpm -ev asl_rpm_name`

Example

The following command line removes a previous version of the ASL:

If you are using...	Enter the following command...
Red Hat Enterprise Linux	<code>rpm -ev VRTSNTAPas1-5.0-3.0_RHEL5</code>
SUSE Linux Enterprise Server	<code>rpm -ev VRTSNTAPas1-5.0-3.0_SLES10</code>

4. Use the `rpm` command to remove the APM package. This command has the format: `rpm -ev apm_rpm_name`.

Example

The following command line removes a previous version of the APM:

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

(Veritas) 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 VxDMP 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 for working with LUNs in native Linux environments

The method you use when working with LUNs often varies depending on your environment--whether you are using multipathing, which protocol you are using, whether you have an HBA, hardware iSCSI initiator, or a software iSCSI initiator, and whether you are using Veritas Storage Foundation. In some cases, the method also varies depending on your version of the Linux operating system.

Note: If you are using Veritas Storage Foundation, see the section on accessing LUNs with VxDMP. The sections here focus on working with LUNs in a native Linux environment.

The sections that follow provide information about the tools you need to use to work LUNs as well as what actions you should take when working with LUNs in your environment. For example, if you do not have multipathing enabled, it is a good practice to provide persistent identification for the LUNs. Or, if you are using the iSCSI software initiator, you can use either the `sanlun` or `iscsiadm` command to view LUNs.

As you work with LUNs, remember that the host cannot distinguish multiple LUNs from multiple paths to the same LUN without multipathing software. As a result:

- If you have more than one path from the host to a LUN, you should use DM-Multipath.
- If you are not using multipathing software, you should limit each LUN to a single path.

For information about the supported configurations for DM-Multipath, see the appropriate interoperability matrix for your N series product.

(FC, Hardware iSCSI) Discovering new LUNs

When you are using an FC or hardware iSCSI environment, you can use the `rescan` script to discover the LUNs you have created and mapped to the Linux host.

Before you begin

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4 or later, Red Hat Enterprise Linux 6 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later. For earlier versions, use the vendor-specific `rescan` scripts, which are available on their websites. See HBA vendor-specific documentation.

Step

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

```
/usr/bin/rescan-scsi-bus.sh
```

(Software iSCSI, multipathing) Discovering new LUNs on Red Hat 5, 6, SUSE 10, 11

When you are running Red Hat Enterprise Linux 5 or 6 series or SUSE Linux Enterprise Server 10 or 11 series with DM-Multipath and the software iSCSI initiator, you can discover new LUNs by rescanning the iSCSI service on the host. Rescanning the service displays all the newly created LUNs that have been mapped to the host.

Before you begin

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4, Red Hat Enterprise Linux 6 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later.

About this task

You cannot view new LUNs until after the operating system discovers them. Once the LUNs have been discovered, they are automatically added to the DM-Multipath configuration.

Note: You can use the `rescan` script to discover new LUNs in both those environments that are using multipathing and those that are not using multipathing.

Steps

1. To discover a new LUN on a system running DM-Multipath, enter one of the following commands:

If you want to...	Enter the following command...
Obtain the list of all the current sessions	<code>iscsiadm -m session</code>
Rescan a specific session	<code>iscsiadm -m session --sid=N --rescan</code> <i>N</i> is the specific session ID.
Rescan all the sessions	<code>iscsiadm -m session --rescan</code>
Rescan using the SCSI <code>rescan</code> script	<code>/usr/bin/rescan-scsi-bus.sh</code>

2. To verify that the new LUNs have been discovered, use the `sanlun` command or the `iscsiadm` command.

(Software iSCSI, multipathing) Discovering new LUNs on Red Hat 4

When you are running Red Hat Enterprise Linux 4 series with DM-Multipath and the software iSCSI initiator, you can discover new LUNs by reloading the iSCSI service on the host. Reloading the service displays all the newly created LUNs that have been mapped to the host.

About this task

You cannot view new LUNs until after the operating system discovers them. Once the LUNs have been discovered, they are automatically added to the DM-Multipath configuration.

Steps

1. Enter the following command on the Linux host to reload the iSCSI service:
`/etc/init.d/iscsi reload`
2. Use the `sanlun` or `iscsi-ls` command to verify that the new LUNs have been discovered.

Viewing a list of LUNs

Whether your environment uses multipathing or not, you have several options for viewing LUNs that are mapped to the host. The `sanlun lun show all` command works for all environments and protocols. If you are using an iSCSI software initiator, you also have a choice between using the `iscsiadm` command or the `iscsi` command, depending on your version of Linux.

Step

1. To view the list of LUNs mapped to your host, run the appropriate command for your system environment.

The following table summarizes the commands and the environments that support them. For more information on the commands and their options, see the man pages.

Note: You can use the `sanlun` command to display LUN information and the `iscsiadm` command to view iSCSI information.

If you are using...	Enter one of the following commands...
All environments	<code>sanlun lun show all</code>
(iSCSI software initiator) Red Hat Enterprise Linux 5 or 6 series	<code>iscsiadm --mode session --sid=N -P 3</code>

If you are using...	Enter one of the following commands...
(iSCSI software initiator) Red Hat Enterprise Linux 4 series	<code>iscsi-ls -l</code>
(iSCSI software initiator) SUSE Linux Enterprise Server 10 or 11 series	<ul style="list-style-type: none">• (SUSE Linux Enterprise Server 10 SP2): <code>iscsiadm --mode session --sid=N -P 3</code>• (SUSE Linux Enterprise Server 10 SP1): <code>iscsiadm --mode session --sid=N -i</code>

The sections that follow contain examples of the type of output these commands produce with different protocols and operating systems.

Examples of sanlun, iscsiadm, iscsi output when used to view LUNs

You can use either the `sanlun` command, the `iscsiadm` command, or the `iscsi` command to view the LUNs configured on your Linux host. The examples in this section show the type of output you would see if you ran one of these commands on your Linux operating system in an environment with DM-Multipath enabled or one with it disabled.

The tool you use depends on your version of Linux and what you want to view as well as whether you have DM-Multipath enabled or not. The `sanlun` command displays the host device names and the LUNs to which they are mapped. The `iscsiadm` command lists the available storage systems and LUNs. The `iscsi-ls -l` command lists storage system node names, IP addresses, and available LUNs.

The following sections provide examples of the type of output you would see if you ran one of these commands in a specific environment; for example with iSCSI and DM-Multipath on Red Hat Enterprise Linux 5 series.

- FC with DM-Multipath running `sanlun`
- iSCSI with multipathing running `sanlun`
- (Red Hat Linux) Software iSCSI with DM-Multipath running `iscsiadm`
- (Red Hat 5, 6) Software iSCSI without multipathing running `iscsiadm`
- (SUSE Linux 10, 11) Software iSCSI with DM-Multipath running `iscsiadm`
- (SUSE Linux 10, 11) Software iSCSI without multipathing running `iscsiadm`
- (Red Hat 4) Software iSCSI without multipathing running `iscsi-ls`

FC with DM-Multipath example of using sanlun to view LUNs

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the FC protocol with DM-Multipath and the storage system running clustered Data ONTAP.

```
# sanlun lun show all
controller(7mode)/vs
server(Cmode) lun-pathname    device      host      lun
                               filename    adapter   protocol  size      mode
```

```

-----
vs_data28_2    /vol/vol1/lun1    /dev/sdgf    host1    FCP    3g    C
vs_data28_2    /vol/vol2/lun2    /dev/sdge    host1    FCP    3g    C
vs_data28_2    /vol/vol3/lun3    /dev/sdgd    host1    FCP    3g    C

```

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the FC protocol with DM-Multipath and the storage system running Data ONTAP operating in 7-Mode.

```

# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname    device
                               filename    host
                               adapter    protocol    lun
                               size    mode
-----
fas3040-201-25  /vol/vol1/lun1    /dev/sdu    host1    FCP    5g    7
fas3040-201-25  /vol/vol2/lun2    /dev/sdt    host1    FCP    5g    7
fas3040-201-25  /vol/vol3/lun3    /dev/sds    host1    FCP    5g    7

```

iSCSI with DM-Multipath example of using `sanlun` to view LUNs

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the iSCSI protocol with DM-Multipath and the storage system running clustered Data ONTAP.

```

# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname    device
                               filename    host
                               adapter    protocol    lun
                               size    mode
-----
vs_data78_1    /vol/vol1/lun1    /dev/sdcx    host29    iSCSI    3g    C
vs_data78_0    /vol/vol2/lun2    /dev/sdcw    host20    iSCSI    3g    C
vs_data79_1    /vol/vol3/lun3    /dev/sdck    host14    iSCSI    3g    C

```

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the iSCSI protocol with DM-Multipath and the storage system running Data ONTAP operating in 7-Mode.

```

# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname    device
                               filename    host
                               adapter    protocol    lun
                               size    mode
-----
fas3040-201-24  /vol/vol1/lun1    /dev/sdb    host6    iSCSI    10m    7
fas3040-201-24  /vol/vol2/lun2    /dev/sdb    host6    iSCSI    10m    7
fas3040-201-24  /vol/vol3/lun3    /dev/sdb    host6    iSCSI    10m    7

```

(Red Hat Enterprise Linux 5, 6) Software iSCSI with DM-Multipath example of using `iscsiadm` to view LUNs

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol and DM-Multipath on a Red Hat Enterprise Linux 5 or 6 series system.

Note: This 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.

```

# iscsiadm -m session -P 3 -r 2
Target: ign.1992-08.com.ibm: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: ign.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: LOGGED_IN
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

```

(Red Hat Enterprise Linux 5, 6) Software iSCSI without multipathing example of using iscsiadm to view LUNs

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol without multipathing on a Red Hat Enterprise Linux 5 or 6 series system. This example uses a specific session ID.

Note: This 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.

```
# iscsiadm -m session -P 3 -r 2
Target: iqn.1992-08.com.ontap: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: LOGGED_IN
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

```

(SUSE Linux 10, 11) Software iSCSI with DM-Multipath example of using iscsiadm command

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol and DM-Multipath on a SUSE Linux Enterprise Server 10 or 11 series.

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

```

# iscsiadm -m session --sid=0 -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.ibm: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: LOGGED_IN
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 10, 11) Software iSCSI without multipathing example of using iscsiadm command

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol without multipathing on a SUSE Linux Enterprise Server 10 or 11 series.

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

```

# iscsiadm -m session --sid=N -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.ibm: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: LOGGED_IN
        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
```

(Red Hat 4) Software iSCSI without multipathing example of using `iscsi-ls` command

This example shows sample output from the `iscsi-ls -l` command when it is issued in a Host Utilities environment that is running the iSCSI protocol without multipathing on a Red Hat Enterprise Linux 4 series system.

```
# /sbin/iscsi-ls -l
*****
SFNet iSCSI Driver Version ... 3.6.2 (27-Sep-2004 )
*****
TARGET NAME           : ign.1992-08.com.ibm: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: IBM Model: Nseries Rev: 0.2
  Type:   Direct-Access          ANSI SCSI revision: 04
  page83 type3: 60a980004f6443745359763759367733
  page83 type1: 4e45544150502020204c554e204f644374535976375936773300000000000000
  page80: 4f6443745359763759367733
  Device: /dev/sdb
LUN ID : 1
  Vendor: IBM Model: Nseries Rev: 0.2
  Type:   Direct-Access          ANSI SCSI revision: 04
  page83 type3: 60a980004f644374535976426253674b
  page83 type1: 4e45544150502020204c554e204f644374535976426253674b00000000000000
  page80: 4f644374535976426253674b
  Device: /dev/sdc
*****
```

(Native multipathing) Using `sanlun` to display DM-Multipath information

When you are using DM-Multipath, you can use the `sanlun` command to confirm that DM-Multipath is set up.

Step

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

```
sanlun lun show -p
```

You can also use the `sanlun lun show all` command to display more information about your LUN setup, such as whether you are using LUNs mapped with clustered Data ONTAP or Data ONTAP operating in 7-Mode.

Note: Check the interoperability matrix to determine if clustered Data ONTAP is supported with your Host Utilities environment.

(Native multipathing) Examples of sanlun output containing DM-Multipath information

When you use the `sanlun` command to confirm that DM-Multipath is set up, the command output varies depending on the protocol you are using and whether you are using clustered Data ONTAP or Data ONTAP operating in 7-Mode.

The following sections provide examples of the type of output produced when you run the `sanlun` command in the following environments:

- Clustered Data ONTAP with FC
- Data ONTAP operating in 7-Mode with FC
- Clustered Data ONTAP with iSCSI
- Data ONTAP operating in 7-Mode iSCSI

Clustered Data ONTAP with FC: Example of using sanlun to display DM-Multipath information

The following examples show the output from the `sanlun lun show -p` command and the `sanlun lun show all` command in a Host Utilities environment that is running clustered Data ONTAP with FC and DM-Multipath.

The first example uses the `sanlun lun show -p` command. The output from the command shows that DM-Multipath (Multipath Provider: Native) is configured.

```
# sanlun lun show -p
                        ONTAP Path: vs_data28_2:/vol/vol1/lun1
                        LUN: 2
                        LUN Size: 3g
                        Mode: C
                        Host Device: 3600a09803246664c422b2d51674f7470
                        Multipath Policy: round-robin 0
                        Multipath Provider: Native
```

host path state	vserver path type	/dev/ node	host adapter	vserver LIF
up	primary	sdfo	host0	lif1
up	primary	sdfk	host0	lif2
up	primary	sdga	host1	lif3
up	primary	sdge	host1	lif4
up	secondary	sdgm	host1	lif5
up	secondary	sdgj	host0	lif6

up	secondary	sdfw	host0	lif7
up	secondary	sdgq	host1	lif8

This example uses the `sanlun lun show all` command. The output shows that the LUNs are mapped to clustered Data ONTAP operating an environment using FC.

```
# sanlun lun show all
controller(7mode)/
vserver(Cmode) lun-pathname      device
                                filename      host
                                adapter      protocol      lun
                                size      mode
-----
vs_data28_2      /vol/vol1/lun1      /dev/sdgf      host1      FCP      3g      C
vs_data28_2      /vol/vol2/lun2      /dev/sdgc      host1      FCP      3g      C
vs_data28_2      /vol/vol3/lun3      /dev/sdgd      host1      FCP      3g      C
```

Data ONTAP operating in 7-Mode with FC: Example of using sanlun to display DM-Multipath information

The following examples show the output from the `sanlun lun show -p` command and the `sanlun lun show all` in a Host Utilities environment that is running Data ONTAP operating in 7-Mode with FC and DM-Multipath.

Note: With the Linux Host Utilities 6.0 release, the output format of the `sanlun` utility has changed. The format no longer maintains backward compatibility when using LUNs mapped for Data ONTAP operating in 7-Mode.

The first example uses the `sanlun lun show -p` command. The output from the command shows that DM-Multipath (Multipath Provider: Native) is configured.

```
# sanlun lun show -p
ONTAP Path: fas3040-201-25:/vol/vol1/lun1
LUN: 27
LUN Size: 5g
Controller CF State: Cluster Enabled
Controller Partner: fas3040-201-24
Mode: 7
Host Device: 360a98000486e2f65686f647876573362
Multipath Policy: round-robin 0
Multipath Provider: Native
```

host path state	controller path type	/dev/ node	host adapter	controller target port
up	primary	sdbq	host0	2a
up	primary	sdeu	host0	1a
up	primary	sdgj	host1	2a
up	primary	sdjn	host1	1a
up	secondary	sdab	host0	2a
up	secondary	sddf	host0	1a
up	secondary	sdhy	host1	2a
up	secondary	sdlc	host1	1a

This example uses the `sanlun lun show all` command. The output shows that the LUNs are mapped to Data ONTAP operating in 7-Mode in an environment using FC.

```
# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname      device      host      protocol  lun   mode
                  filename      adapter
-----
fas3040-201-25  /vol/vol1/lun1  /dev/sdu    host1     FCP        5g    7
fas3040-201-25  /vol/vol2/lun2  /dev/sdt    host1     FCP        5g    7
fas3040-201-25  /vol/vol3/lun3  /dev/sds    host1     FCP        5g    7
```

Clustered Data ONTAP with iSCSI: Example of using sanlun to display DM-Multipath information

The following examples show the output from the `sanlun lun show -p` command and the `sanlun lun show all` in a Host Utilities environment that is running clustered Data ONTAP with iSCSI and DM-Multipath. The output is the same regardless of whether you are using a software iSCSI initiator or hardware iSCSI initiator.

The first example uses the `sanlun lun show -p` command. The output from the command shows that DM-Multipath (Multipath Provider: Native) is configured.

```
# sanlun lun show -p
ONTAP Path: vs_data28_2:/vol/vol1/lun1
LUN: 2
LUN Size: 3g
Mode: C
Host Device: 3600a09803246664c422b2d51674f7470
Multipath Policy: round-robin 0
Multipath Provider: Native
```

host path state	vserver path type	/dev/ node	host adapter	vserver LIF
up	primary	sdfo	host0	lif1
up	primary	sdfk	host0	lif2
up	primary	sdga	host1	lif3
up	primary	sdge	host1	lif4
up	secondary	sdgm	host1	lif5
up	secondary	sdgj	host0	lif6
up	secondary	sdfw	host0	lif7
up	secondary	sdgg	host1	lif8

This example uses the `sanlun lun show all` command. The output shows that the LUNs are mapped to clustered Data ONTAP in an environment using iSCSI.

```
# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname      device      host      protocol  lun   mode
                  filename      adapter
-----
vs_data78_1     /vol/vol1/lun1  /dev/sdcx    host29    iSCSI      3g    C
vs_data78_0     /vol/vol2/lun2  /dev/sdcw    host20    iSCSI      3g    C
vs_data79_1     /vol/vol3/lun3  /dev/sdck    host14    iSCSI      3g    C
```

Data ONTAP operating in 7-Mode iSCSI: Example of using sanlun to display DM-Multipath information

The following examples show the output from the `sanlun lun show -p` command and the `sanlun lun show all` in a Host Utilities environment that is running Data ONTAP operating in 7-Mode with iSCSI and DM-Multipath. The output is the same regardless of whether you are using a software iSCSI initiator or hardware iSCSI initiator.

The first example uses the `sanlun lun show -p` command. The output from the command shows that DM-Multipath (Multipath Provider: Native) is configured.

```
# sanlun lun show -p
ONTAP Path: f3170-201-37:/vol/vol1/lun1
LUN: 5
LUN Size: 6g
Controller CF State: Cluster Enabled
Controller Partner: f3170-201-36
Mode: 7
Host Device: 360a98000572d5a74526f6374342f5658
Multipath Policy: round-robin 0
Multipath Provider: Native

-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
up        iscsi     sdh       host8     10.72.201.37
up        iscsi     sdp       host10    192.168.100.37
```

This example uses the `sanlun lun show all` command. The output shows that the LUNs are mapped to Data ONTAP operating in 7-Mode in an environment using iSCSI.

```
# sanlun lun show all
controller(7mode)/      device      host      lun
vserver(Cmode)  lun-pathname filename  adapter  protocol  size  mode
-----
fas3040-201-24  /vol/vol1/lun1 /dev/sdb  host6    iSCSI     10m   7
fas3040-201-24  /vol/vol2/lun2 /dev/sdb  host6    iSCSI     10m   7
fas3040-201-24  /vol/vol3/lun3 /dev/sdb  host6    iSCSI     10m   7
```

Enabling device persistence for newly discovered LUNs

The Linux Host sees LUNs on the storage system as SCSI devices. When you have multiple LUNs, you must be able to persistently identify each LUN across system reboots. This means you should make sure each LUN has a unique file system label and then mount the file system using that label.

Before you begin

Discover the LUNs.

About this task

The issue of persistent identification occurs because the Linux operating system assigns a unique device name, such as `/dev/sda`, to each LUN as it discovers the LUN. If you reboot or restart the iSCSI service, these names might change because the order in which the operating system discovers LUNs cannot be predicted.

When you use multipathing, DM-Multipath automatically creates persistent devices for each LUN in the `/dev/mapper/` directory on the Linux host.

Step

1. Mount the new file system by adding an entry to `/etc/fstab`.

If you are running...	'Enter the following line...
DM-Multipath with Red Hat Enterprise Linux	<code>device mount_point type _netdev,defaults 0 0</code>
DM-Multipath with SUSE Linux Enterprise Server	<code>device mount_point type _netdev,defaults nofail,rw 0 0</code>

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

`mount_point` is the mount point you created for the file system.

`type` is the file system type, such as `ext2` or `ext3`.

`_netdev` is used for any network-dependent devices such as iSCSI. It is only used in iSCSI environments and lets you add iSCSI mount point devices to `/etc/fstab`.

Removing an unmapped LUN

You can use the SCSI `rescan` script to remove an unmapped LUN from the host.

Before you begin

Make sure that the LUN is not in use and that it has been unmapped.

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4 or later and Red Hat Enterprise Linux 6 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later. For earlier versions, use the vendor-specific `rescan` scripts, which are available on their websites. See the documentation for your HBA.

Step

1. Enter the following command to remove the LUN:

```
/usr/bin/rescan-scsi-bus.sh -r
```


(Veritas) LUN access when using VxDMP

If you configured your system correctly for Veritas Storage Foundation, you can use VxDMP for multipathing and VxVM to manage the LUNs. You can still use commands such as `sanlun` to display information about the LUN paths.

There are some things you should check before you start working with LUNs in a Veritas Storage Foundation environment:

- If you are using iSCSI, make sure you are using a version of Veritas Storage Foundation that supports that protocol.
See the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).
- Make sure you have set the HBA driver parameters correctly for your system setup.
Having the correct values for these parameters ensures that the multipathing and storage system failover work correctly.
- If you configured VxDMP, multipath devices are created for all the LUNs that are discovered by the HBA driver.
Each time an HBA driver is started, it scans the storage system and discovers all mapped LUNs.
- Make sure you set the VxDMP restore daemon to the correct values.
These values ensure that Veritas Storage Foundation works efficiently and correctly.
- When you use Veritas Storage Foundation, the VxVM manages the LUNs.
This means that, in addition to using tools such as `sanlun` and `iscsadm` to display information about the LUNs, you can also use the VxVM interface to display information about the VxVM devices.

(Veritas, FC) Discovering new LUNs

After you create a new LUN and map it to the Linux host, you can discover the LUN by using the SCSI rescan script.

Before you begin

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4 or later, Red Hat Enterprise Linux 6 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later. For earlier versions, use the vendor-specific `rescan` scripts, which are available on their websites. See HBA vendor-specific documentation.

Steps

1. Enter the following command to rescan the operating system:
`/usr/bin/rescan-scsi-bus.sh`

2. Initiate a rescan of the operating system device tree from the Veritas Volume Manager by entering the following command:

```
vxdisk scandisks
```

(Veritas, Software iSCSI) Discovering new LUNs for Red Hat 5, 6, or SUSE 10, 11

When you are running Red Hat Enterprise Linux 5 or 6 series or SUSE Linux Enterprise Server 10 or 11 series and the software iSCSI initiator, you can discover new LUNs by rescanning the iSCSI service on the host. Rescanning the service displays all the newly created LUNs that have been mapped to the host.

Before you begin

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4, Red Hat Linux 6 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later. For earlier versions, use the vendor-specific `rescan` scripts, which are available on their websites. See HBA vendor-specific documentation.

About this task

You cannot view new LUNs until after the operating system discovers them. Once the LUNs have been discovered, they are automatically added to the VxDMP configuration.

Steps

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	/usr/bin/rescan-scsi-bus.sh
Rescan Veritas devices	vxdisk scandisks

2. To verify that the new LUNs have been discovered, use the `sanlun` command or the `iscsiadm` command.

(Veritas) Viewing LUNs mapped to the host

Regardless of the protocol your environment uses, you have several options for viewing LUNs that are mapped to the host. The `sanlun` command works for all protocols. If you are using an iSCSI software initiator, you also have a choice between using the `iscsiadm` command or the `iscsi` command, depending on your version of Linux. You can also use `vxdisk list` to see the LUNs on VxVM disks.

Steps

1. To view a list of LUNs mapped to your host, run the appropriate command for your system environment.

The following table summarizes the commands and the environments that support them. For more information on the commands and their options, see the man pages.

If you are using...	Enter one of the following commands...
(FC, iSCSI hardware initiator) All versions of Linux	<code>sanlun lun show all</code>
(iSCSI software initiator) Red Hat Enterprise Linux 5 or 6 series	<ul style="list-style-type: none"> • <code>sanlun lun show all</code> • <code>iscsiadm --mode session --sid=N -P 3</code>
(iSCSI software initiator) SUSE Linux Enterprise Server 10 or 11 series	<ul style="list-style-type: none"> • <code>sanlun lun show all</code> • (SUSE Linux Enterprise Server 10 SP2): <code>iscsiadm --mode session --sid=N -P 3</code> • (SUSE Linux Enterprise Server 10 SP1): <code>iscsiadm --mode session --sid=N -i</code>

The sections that follow contain examples of the type of output these commands produce with different protocols and operating systems.

2. To view the LUNs on the VxVM disks, enter the `vxdisk list` command.

(Veritas) Examples of `sanlun`, `iscsiadm`, `iscsi` output when used to view LUNs

You can use either the `sanlun` command, the `iscsiadm` command, or the `iscsi` command to view the LUNs configured on your Linux host. The examples in this section show the type of output you would see if you ran one of these commands on your Linux operating system in an environment running VxDMP.

The tool you use depends on your version of Linux and what you would like to view. The `sanlun` command displays the host device names and the LUNs to which they are mapped. The `iscsiadm` command lists the available storage systems and LUNs.

The following sections provide examples of the type of output you would see if you ran one of these commands in a specific environment; for example with iSCSI and DM-Multipath on Red Hat Enterprise Linux 5 series.

- FC running `sanlun`
- FC running `vxdisk`
- (Red Hat Linux) Software iSCSI running `sanlun`
- (SUSE Linux 10, 11) Software iSCSI running `sanlun`
- (Red Hat Linux) Software iSCSI running `iscsiadm`
- (SUSE Linux 10, 11) Software iSCSI running `iscsiadm`

Note: The output in the examples below has been modified to better fit the screen.

Data ONTAP operating in 7-Mode with FC example of using `sanlun` to view LUNs

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the Data ONTAP operating in 7-Mode with FC and Veritas Storage Foundation.

Note: With the Linux Host Utilities 6.0 release, the output format of the `sanlun` utility has changed. The format no longer maintains backward compatibility when using LUNs mapped for Data ONTAP operating in 7-Mode.

```
# sanlun lun show all
controller(7mode)/
vsriver(Cmode)  lun-pathname  device      host      protocol  lun  mode
                  filename      adapter
-----
fas3040-201-25  /vol/vol1/lun1  /dev/sdu    host1     FCP        5g   7
fas3040-201-25  /vol/vol2/lun2  /dev/sdt    host1     FCP        5g   7
fas3040-201-25  /vol/vol3/lun3  /dev/sds    host1     FCP        5g   7
```

If you executed the `sanlun lun show all` command in a Data ONTAP operating in 7-Mode FC environment, you would get the following output:

```
# sanlun lun show -p

ONTAP Path: fas6030-201-71:/vol/vol1/lun1
LUN: 12
LUN Size: 7g
Controller_CF_State: Cluster Enabled
Controller Partner: fas6030-201-70
Mode: 7
DMP NODE: fas60300_9
Multipath Provider: Veritas
```

host path state	controller path type	/dev/ node	host adapter	controller target port
up	secondary	sdn	host3	2a
up	primary	sdbg	host3	2a
up	primary	sdes	host4	2b
up	secondary	sdcz	host4	2b

FC example of using vxdisk to view LUNs

This example shows sample output from the `vxdisk list` command when it is issued in a Host Utilities environment that is running the FC protocol. The `vxdisk list` displays the LUNs on the VxVM disks:

```
# vxdisk list
DEVICE      TYPE      DISK      GROUP     STATUS
s20200_0    auto:cdsdisk data_dg01 data_dg    online thinrclm shared
s20200_1    auto:cdsdisk data_dg02 data_dg    online thinrclm shared
s20200_2    auto:cdsdisk data_dg113 data_dg    online thinrclm shared
s20200_3    auto:cdsdisk data_dg180 data_dg    online thinrclm shared
```

Data ONTAP operating in 7-Mode with iSCSI example of using sanlun to view LUNs

This example shows sample output from the `sanlun lun show all` command when it is issued in a Host Utilities environment that is running the Data ONTAP operating in 7-Mode with iSCSI and Veritas Storage Foundation.

Note: With the Linux Host Utilities 6.0 release, the output format of the `sanlun` utility has changed. The format no longer maintains backward compatibility when using LUNs mapped for Data ONTAP operating in 7-Mode.

```
# sanlun lun show all
controller(7mode)/
vserver(Cmode)  lun-pathname  device      host      protocol  lun      mode
                  filename      adapter
-----
fas3040-201-24  /vol/vol1/lun1  /dev/sdb    host6     iSCSI     10m      7
fas3040-201-24  /vol/vol2/lun2  /dev/sdb    host6     iSCSI     10m      7
fas3040-201-24  /vol/vol3/lun3  /dev/sdb    host6     iSCSI     10m      7
```

If you executed the `sanlun lun show all` command in a Data ONTAP operating in 7-Mode iSCSI environment, you would get the following output:

```
# ./sanlun lun show -p

ONTAP Path: fas6030-201-70:/vol/vol1/lun1
LUN: 21
LUN Size: 7g
Controller_CF_State: Cluster Enabled
Controller Partner: fas6030-201-71
Mode: 7
DMP NODE: fas60300_44
Multipath Provider: Veritas
```

host path state	controller path type	/dev/ node	host adapter	controller target port
up	iscsi	sdaq	host6	10.72.201.24
up	iscsi	sdbp	host7	10.72.201.25

(Red Hat Enterprise Linux 5) Software iSCSI example of using iscsiadm to view LUNs

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol and Veritas Storage Foundation on a Red Hat Enterprise Linux 5 series system.

Note: This 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.

```
# iscsiadm -m session -P 3 -r 2
Target: ign.1992-08.com.ontap: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: ign.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 sdq State: running
scsi4 Channel 00 Id 0 Lun: 7
    Attached scsi disk sdq State: running
```

(SUSE Linux 10, 11) Software iSCSI example of using iscsiadm command

This example shows sample output from the `iscsiadm` command when it is issued in a Host Utilities environment that is running the iSCSI protocol and Veritas Storage Foundation on a SUSE Linux Enterprise Server 10 or 11 system.

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

```
# iscsiadm -m session --sid=N -P 3
iSCSI Transport Class version 2.0-724
iscsiadm version 2.0-868
Target: iqn.1992-08.com.ibm: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
```

(Veritas) Displaying multipathing information for VxDMP

You can use the `sanlun` command to display multipathing information for VxDMP.

Step

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

```
sanlun lun show -p
```

(Veritas) Examples of sanlun output for VxDMP

When you use the `sanlun` command to display multipathing information for VxDMP, the command output varies depending on the protocol you are using.

The following sections provide examples of the type of output produced when you run the `sanlun` command in the following VxDMP environments:

- FC
- Software iSCSI

FC example of using sanlun to display VxDMP information

In the following example, the `sanlun lun show -p` command is issued in a Host Utilities environment that is running the FC protocol with Veritas Storage Foundation and Data ONTAP operating in 7-Mode. The output from the command shows that VxDMP is configured.

```
# sanlun lun show -p
ONTAP Path: fas6030-201-71:/vol/vol1/lun1
LUN: 12
LUN Size: 7g
Controller_CF_State: Cluster Enabled
Controller Partner: fas6030-201-70
Mode: 7
DMP NODE: fas60300_9
Multipath Provider: Veritas
-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
-----
up        secondary sdn        host3     2a
up        primary   sdbg       host3     2a
up        primary   sdes       host4     2b
up        secondary sdcz       host4     2b
```


Software iSCSI example of using sanlun to display VxDMP information

In the following example, the `sanlun lun show -p` command is issued in a Host Utilities environment that is running the iSCSI protocol with a software initiator Veritas Storage Foundation and Data ONTAP operating in 7-Mode. The output from the command shows that VxDMP is configured.

```
# sanlun lun show -p

ONTAP Path: fas6030-201-70:/vol/vol1/lun1
  LUN: 21
  LUN Size: 7g
Controller_CF_State: Cluster Enabled
Controller Partner: fas6030-201-71
  Mode: 7
  DMP NODE: fas60300_44
Multipath Provider: Veritas
-----
host      controller
path      path      /dev/      host      controller
state     type      node      adapter   target
-----
up        iscsi      sdaq      host6      10.72.201.24
up        iscsi      sdbp      host7      10.72.201.25
```

(Veritas) Removing a LUN

After you unmap the LUN from the Linux host, you can remove the LUN from the Veritas VxVM by running the SCSI rescan script.

Before you begin

You must ensure that the LUN is not in use.

If you do not have a copy of the `rescan` script, get one. The `rescan` script is available with the `sg3_utils` package. In addition, the `rescan` script is available with Red Hat Enterprise Linux 5 Update 4, Red Hat Enterprise Linux 6.0 or later, SUSE Linux Enterprise Server 10 SP2 or later, and SUSE Linux Enterprise Server 11 or later. For earlier versions, use the vendor-specific `rescan` scripts, which are available on their websites. See HBA vendor-specific documentation.

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 s30200_1
```

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

```
/usr/bin/rescan-scsi-bus.sh -r
```
4. Initiate a rescan of the operating system device tree from Veritas Volume Manager by entering the following command:

```
vxdisk scandisks
```

(Veritas) Displaying available paths using VxVM

You can use the VxVM management interface to display information about the VxVM devices and see which devices are managed by the VxVM. The `vxdisk` command displays information the device, type, disk, group, and status.

Steps

1. Enter the following command to view all VxVM devices:

```
vxdisk list
```

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.

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
s31700_0    auto:cdsdisk data_dg01 data_dg    online thinrclm shared
s31700_1    auto:cdsdisk data_dg02 data_dg    online thinrclm shared
s31700_2    auto:cdsdisk data_dg08 data_dg    online thinrclm shared
s31700_3    auto:cdsdisk data_dg09 data_dg    online thinrclm shared
s31700_4    auto:cdsdisk data_dg10 data_dg    online thinrclm shared
s31700_5    auto:cdsdisk data_dg11 data_dg    online thinrclm shared
s31700_6    auto:cdsdisk data_dg12 data_dg    online thinrclm shared
s31700_7    auto:cdsdisk data_dg13 data_dg    online thinrclm shared
```

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=s31700_1
NAME  STATE[A]    PATH-TYPE[M]  CTRL-NAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
=====
sdb   ENABLED(A)  -             c7         S3170      s31700      -
sdq   ENABLED(A)  -             c8         S3170      s31700      -
```

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

```
vxdmpadm getsubpaths ctrl=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 ctrl=controller_name` command is similar to the following:

```
#vxdmpadm getsubpaths ctrl=c7
NAME  STATE[A]    PATH-TYPE[M]  DMPNODENAME  ENCLR-TYPE  ENCLR-NAME  ATTRS
=====
sdj   ENABLED(A)  -             s31700_0     S3170      s31700      -
sdb   ENABLED(A)  -             s31700_1     S3170      s31700      -
sd1   ENABLED(A)  -             s31700_10    S3170      s31700      -
sdm   ENABLED(A)  -             s31700_11    S3170      s31700      -
sdn   ENABLED(A)  -             s31700_12    S3170      s31700      -
```

(FC) Setting up a SAN boot LUN on Red Hat Enterprise Linux

You can set up a SAN boot LUN to work in a Red Hat Enterprise Linux environment that is using the FC protocol.

Before you begin

Verify that your system setup supports SAN boot LUNs. See the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

Steps

1. Create a LUN on the storage system and map it to the host. This LUN will be the SAN boot LUN.

You should ensure the following:

- The SAN boot LUN is mapped to the host.
- Multiple paths to the LUN are available.
- The LUN is visible to the host during the boot process.

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

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

3. Configure the paths to the HBA boot BIOS as primary, secondary, tertiary, and so on, on the boot device.

For more information, see your vendor-specific documentation.

4. Save and exit.
5. Reboot the host.
6. Install the operating system on the SAN boot LUN.

Note: You must 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.

7. Install the Host Utilities.
8. Configure DM-Multipath.

(FC) Setting up a SAN boot LUN on SUSE Linux Enterprise Server

You can set up a SAN boot LUN to work in a SUSE Linux Enterprise Server environment that is using the FC protocol.

Before you begin

Verify that your system setup supports SAN boot LUNs. See the appropriate interoperability matrix for your N series product, available on the IBM support website.

Steps

1. Create a LUN on the storage system and map it to the host. This LUN will be the SAN boot LUN.

You should ensure the following:

- The SAN boot LUN is mapped to the host.

Note: If you are running a SAN boot installation using a version of SUSE Linux Enterprise Server prior to 11 SP2, then only one primary path to the LUN is available. SUSE Linux Enterprise Server 11 SP2 supports multiple paths to the LUN.

- The LUN is visible to the host during the boot process.

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

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

3. Install the operating system on the SAN boot LUN.
4. Install the Host Utilities.
5. Configure DM-Multipath.

(FC) Configuring the root partition with DM-Multipath on SUSE Linux Enterprise Server

You can configure the 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	<code>mkinitrd -f multipath</code>
SUSE Linux Enterprise Server 10 SP2	<code>mkinitrd -f mpath</code>

Note: You must keep a backup of the initrd 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.

(iSCSI) SAN boot configuration for iSCSI hardware, software initiators

When you set up a SAN boot device, you use a SAN-attached disk, such as a LUN, as a root device for a host. A SAN boot LUN 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 a SAN boot device, you can have the root device on an iSCSI LUN, and you can use any of the following options to load the kernel:

- A host's locally attached disk (for storing kernel and initrd images)
- A Preboot Execution Environment (PXE) Server

- **Hardware iSCSI**

If the SAN boot LUN uses an iSCSI HBA, then, because the protocol stack runs on the HBA, it is ready to communicate with the storage system and discover a LUN when it starts up.

You can have both the boot device and root device on an iSCSI LUN.

Note: Not all operating systems supported by the Host Utilities work with iSCSI SAN boot LUNs. For example, Oracle VM does not support creating a SAN boot LUN that uses software iSCSI.

(Hardware iSCSI) Configuring SAN boot on Red Hat Enterprise Linux

When you are running Red Hat Enterprise Linux, you can configure SAN boot LUN to use an iSCSI hardware initiator.

Steps

1. Create a LUN on the storage system and map it to the host. This will be the SAN boot LUN.

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.

4. After making changes to the HBA BIOS, save and exit.
5. Reboot the host.

6. Install the operating system on the boot LUN 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.

(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.

Before you begin

Check the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

Steps

1. When you initiate the installation, specify the Boot Option as `linux mpath` and press Enter.
2. Continue with the installation until you reach the storage configuration page. Click **Advanced storage configuration**.
3. Select **Add iSCSI target** and click **Add drive**.
4. Enter the Target IP address and the iSCSI initiator name.

Note: You should ensure that you associate this IQN with the correct 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. At the storage devices selection screen, 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 mount point 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

You can configure a SAN boot LUN on SUSE Linux Enterprise Server. Doing this requires multiple steps and close attention to the requirements in each step.

Before you begin

Verify that your system setup supports SAN boot LUNs. See the appropriate interoperability matrix for your N series product.

Steps

1. Log in to the storage system console or the Web interface of the storage system.
2. When you initiate the installation, specify Boot Option as follows: `linux withiscsi=1 netsetup=1`
3. In the **iSCSI Initiator Overview** page, select the **Service** tab and enter the Target IP address and the **iSCSI initiator** name.

Note: You should ensure that you associate this IQN with the appropriate privileges on the storage controller.

4. On the storage controller, create an igroup with the initiator name that you provided in the previous step.
5. Create a LUN on the storage system on which you can create the root partition, and map it to the igroup.
6. Return to the host screen. Select the **Connected Targets** tab and click **Add**.
7. On 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) Specify 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 one. You might also have to do this for the authentication credentials also.

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 `onboot` 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 mount point as `/`.

17. Click the **Fstab Options** button.

18. Select the `Mount by Device ID` option.

19. Ensure that you have the `_netdev`, `nofail` 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 **Boot** 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`

28. Click **OK.**

29. Click the **Boot Loader Installation tab.**

30. In the **Boot Loader Location pane, select the `Boot from Master Boot Record` option.**

Click **Finish**. Doing this returns you to the Installation Settings page.

31. Review the configuration settings and click **Accept. The **Confirm Installation** page is displayed.**

32. Click **Install and follow the prompts to complete the installation.**

(Software iSCSI) Configuring multipathing for a SAN boot LUN using SUSE Linux Enterprise Server

When you set up a SAN boot LUN for SUSE Linux Enterprise Server, you install the operating system on a SCSI device. After you do that, you should set up multipathing. Multipathing must be configured on the root partition of the SAN boot LUN.

Before you begin

Verify that your system setup supports SAN boot LUNs. See the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

Steps

1. Enable additional paths to the LUN (root LUN).
2. Use YaST2 to change the start-up mode for all the iSCSI sessions.

If you are using...	Set the start-up mode to...
SUSE Linux Enterprise Server 11	Onboot
SUSE Linux Enterprise Server 10 SP2	automatic

3. Modify the `/etc/sysconfig/network/ifcfg x` file so that the value of `STARTMODE` now reads `nfsroot`.
4. Change the value of the session re-establishment timeout for iSCSI sessions fetching the SAN boot LUN by entering the following command:

```
iscsiadm -m node -T targetname -p ip:port -o update -n  
node.session.timeo.replacement_timeout -v 5
```

5. Create a new `initrd` image with the root partition on multipath enabled.

If you are using...	Enter the command...
SUSE Linux Enterprise Server 11	mkinitrd -f "iscsi multipath"
SUSE Linux Enterprise Server 10 series	mkinitrd -f "iscsi mpath"

6. Change the start-up mode of the iSCSI sessions fetching the SAN boot 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
```

7. Reboot the host.

The system boots with multipathing 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.

(Veritas) Configuring SAN boot in a Veritas environment

You can set up a SAN boot LUN to work in a Veritas Storage Foundation environment.

Before you begin

Verify that SAN boot LUNs are supported with your version of Veritas Storage Foundation and your Linux operating system. See the appropriate interoperability matrix for your N series product, available on the IBM support website.

About this task

When you are working in a Veritas Storage Foundation environment, the steps you need to perform to set up a SAN boot LUN are essentially the same for both Red Hat Enterprise Linux and SUSE Linux Enterprise Server.

Steps

1. Create a LUN on the storage system. This LUN will be the SAN boot LUN.
2. Map the LUN to the host.
3. Ensure that only one primary path is available to the LUN.
4. Ensure that only one SAN boot LUN is available to the host.
5. 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 HBA vendor-specific documentation.

6. 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) The Red Hat Enterprise Linux 4 Update 4 distribution does not include 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 appropriate interoperability matrix for your N series product, available on the IBM support website.

Note: (SUSE) When you install the SUSE Linux Enterprise Server operating system, you should ensure that GRUB is installed in the Master Boot Record. You can do this from the Expert tab in the software package selection screen during installation.

7. After installing the operating system, reboot the host.

The host boots from the SAN boot LUN on the storage system through a primary path.

8. Install the Linux Host Utilities.
9. If you are using HBA drivers acquired from an OEM, install the supported versions of the drivers.
10. Verify the HBA settings.
11. Install Veritas Storage Foundation and any appropriate patches or fixes for it.
12. Configure the `vxdump restore` daemon by setting it to an interval of 60:

```
vxdmptune dmp_restore_interval 60
```

On reboot, this value takes effect and remains persistent across system reboots.

13. **(Veritas Storage Foundation 5.1 and later)** Set the Veritas DMP LUN retries tunable to a value of 300:

```
vxdmppadm settune dmp_lun_retry_timeout=300
```

The new value takes effect immediately.

14. **(Veritas Storage Foundation 5.1 SP1 and later)** Set the value of the `dmp_path_age` to an interval of 120 by entering the following command:

```
vxdmppadm settune dmp_path_age=120
```

The new value takes effect immediately.

15. Enable persistence by entering the following command:

```
vxddladm set namingscheme=osn persistence=yes
```

You must enable persistence before you can encapsulate the root disk.

16. Encapsulate the root disk for use in VxVM by entering 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.

17. Reboot the host after encapsulation.

18. Verify the encapsulation by entering the following command:

```
vxprint
```

This command displays the rootvol and swapvol volumes under the corresponding disk group.

19. 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.

Support for host virtualization

The Host Utilities support virtualization, including products based out of technologies such as; Kernel-based Virtual Machine (KVM) and Xen. The products derived from KVM are Red Hat Enterprise Linux KVM, SUSE Linux Enterprise Server KVM, and RHEV. Products derived from Xen are Citrix XenServer, SUSE Linux Enterprise Server Xen, and Oracle VM.

Server virtualization is a method of dividing computer resources into multiple, isolated environments. In a virtual ecosystem, a host operating system runs one or more guest virtual machines in a simulated environment. Each guest virtual machine (VM) has access to all of the host's hardware. You can configure guest VMs in full and paravirtualized modes.

Full virtualization includes a virtualization translation layer (VMM) that ensures all guest system calls get translated to native hypervisor's format. In this approach all guest operating systems run completely unmodified. Commercial products based out of Full Virtualization concept are RHEL-KVM, RHEV, Oracle VirtualBox. KVM is a full virtualization solution for Linux on x86 hardware that contains virtualization extensions (Intel VT or AMD-V). It consists of a loadable kernel module (kvm.ko) that provides the core virtualization infrastructure and a processor-specific module (kvm-intel.ko or kvm-amd.ko). KVM is supported with Red Hat Enterprise Linux 5 Update 4 or later and Red Hat Enterprise Linux 6.0 or later.

Note: KVM leverages the Linux kernel's infrastructure and extends it to serve as a full-fledged hypervisor. Thus, any limitations that apply to the Linux kernel also apply to the KVM hypervisor.

Note: If you are using Oracle VM hypervisor, see the sections for configuring FC, iSCSI, and multipathing. In this guide, sections that refer to Red Hat Enterprise Linux also apply to Oracle Linux. Both operating systems use the same instructions for the tasks featured in this guide.

Paravirtualization support involves providing a paravirtualized network driver, a paravirtualized block I/O device (disk) driver, and a balloon driver to affect the operation of the guest virtual memory manager and provide CPU optimization for Linux guests. Both KVM and Oracle VM provide limited paravirtualization support for guest VMs. This approach to virtualization includes a slim hypervisor kernel running on the bare metal host hardware. Each guest OS needs to be ported to the hypervisor's kernel interface. The most renowned project utilizing the paravirtualization concept is Xen. Different commercial products based on Xen include: OVM, Citrix XenServer and SLES XEN.

Oracle VM is a server virtualization solution that consists of Oracle VM Server (OVS). This is a self-contained virtualization environment designed to provide a lightweight, secure, server-based platform for running virtual machines. OVS is based on an updated version of the underlying Xen hypervisor technology. Oracle VM Manager provides the user interface to manage OVS and guests.

Hypervisor VHD requires alignment for best performance

A virtual hard disk (VHD) is partitioned with a master boot record that is used by a Linux virtual machine. It must be aligned with the underlying LUN for best performance.

If the data block boundaries of a disk partition do not align with the block boundaries of the underlying LUN, the storage system often has to complete two block reads or writes for every operating system block read or write. The additional block reads and writes caused by the misalignment can lead to serious performance problems.

The misalignment is caused by the location of the starting sector for each partition defined by the master boot record. Partitions created by Linux usually are not aligned with underlying IBM N series LUNs.

See technical report TR-3747 for best practices on configuring the aligned LUNs.

Note: This technical report contains information about NetApp products that IBM licenses and in some cases customizes. Technical reports might contain information about models and features that are not supported by IBM.

Supported Linux and Data ONTAP features

The Host Utilities support a number of features and configurations for Linux hosts and storage systems running Data ONTAP.

Some of the supported features include

- SAN booting
- The Linux Device Mapper Multipathing
- Volume management and multipathing with Veritas Storage Foundation
- Host and storage virtualization
- ALUA

Note: Your specific environment can affect what the Host Utilities support.

Protocols and configurations supported by Host Utilities

The 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.

These protocols enable the host to access data on storage systems. The storage systems are targets that have storage target devices called LUNs.

Protocols enable hosts to access LUNs to store and retrieve data.

For more information about using the protocols with your storage system, see the *Data ONTAP SAN Administration Guide for 7-Mode* for your version of Data ONTAP.

For more details on supported topologies, including diagrams, see the *Data ONTAP SAN Configuration Guide for 7-Mode* for your version of Data ONTAP.

The sections that follow provide high-level information about these protocols.

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 the FC protocol with fabric-attached SAN and direct-attached configurations:

- 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. Like the traditional FC protocol, FCoE maintains existing FC management and controls, but it uses a 10-gigabit Ethernet network as the hardware transport.

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 the same way you use 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.

You need to make a note of the iSCSI node name so that you can supply it when you create an igroup.

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 booting uses a SAN-attached disk, such as a LUN configured on a storage controller, as a boot device for a host.

Note: SAN booting is not supported with Oracle VM.

A SAN boot LUN 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 configurations that are supported for SAN boot LUNs, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 8).

Support for Linux Device Mapper Multipathing

You can use Linux Device Mapper Multipathing (DM-Multipath) to provide I/O failover and path load sharing for multipathed block devices.

The 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.

For information on setting up DM-Multipath, see section *DM-Multipath configuration*.

Volume management and multipathing with Veritas Storage Foundation

If your environment uses Veritas Storage Foundation, you can use the tools it supplies to work with LUNs and provide multipathing. These tools include Veritas Volume Manager (VxVM) and Dynamic Multipathing (VxDMP).

With VxVM, you can manage the LUNs that you create on the storage systems. You can use VxDMP to provide the multipathing solution for your system.

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

To determine which versions of Veritas are supported, see the interoperability matrix.

Linux configurations that support ALUA

The Linux Host Utilities support ALUA (asymmetric logical unit access) on hosts running Red Hat Enterprise Linux or SUSE Linux Enterprise Server and a version of Data ONTAP that supports ALUA.

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

ALUA defines a standard set of SCSI commands for discovering path priorities to LUNs on SANs. When you have the host and storage controller configured to use ALUA, it automatically determines which target ports provide optimized and unoptimized access to LUNs.

Note: If you are using clustered DATA ONTAP, ALUA is supported with the FC, FCoE, and iSCSI protocols and you must use it. If you are using DATA ONTAP operating in 7-Mode, ALUA is supported only with the FC and FCoE FC protocols.

ALUA is automatically enabled when you set up your storage for FC.

The following configurations support ALUA:

Host Utilities Version	Host requirements	Data ONTAP versions
Host Utilities 4.0 and later	<ul style="list-style-type: none"> Red Hat Enterprise Linux 6 or later Red Hat Enterprise Linux 5 Update 1 or later SUSE Linux Enterprise Server 10 SP2 or later SUSE Linux Enterprise Server 11 or later <p>Note: Veritas Storage Foundation 5.1 or later support ALUA with the FC protocol.</p>	7.2.4 or later

Troubleshooting

If you encounter problems while running the Host Utilities on FC, iSCSI, or Veritas Storage Foundation, you can check the sections that follow for troubleshooting tips.

LUNs not recognized by DM-Multipath

LUNs are mapped to the Linux host and are correctly reported by the `sanlun lun show all` command. However, the LUNs are not reported by the `multipath -ll` command. Check to make sure the LUNs are not blacklisted (excluded) from the DM-Multipath support configuration in the `/etc/multipath.conf` file.

Enter the following command on the Linux host to display the devices that are blacklisted:

```
multipath -v3 -d | grep blacklist
```

If any devices show up as being 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
```

(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. Refer to the appropriate interoperability matrix for your N series product, available on the IBM support website.
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.

Configuration settings	What you should do
FC service status	Verify that the FC service is licensed and started on the storage system. For more information, see the <i>Data ONTAP SAN Administration Guide for 7-Mode</i> .
HBA World Wide Port Names	Verify that you are using the correct initiator WWPNs in the igroup configuration. On the storage system, use the <code>igroup show</code> command to display the WWPNs of the initiators in the storage system's igroups. On the host, use the HBA tools to display the initiator WWPN. The initiator WWPNs configured in the igroup and on the host should match.
Initiator login	Verify that the initiator is logged in to the storage system by entering the <code>fcp show initiator</code> command on the storage system console. If the initiator is configured and logged in to the storage system, the storage system console displays the initiator WWPN.
LUN mappings	Verify that the LUNs are mapped to an igroup. On the storage system, use one of the following commands: <ul style="list-style-type: none"> <code>lun show -m</code> displays all the LUNs and the igroups to which they are mapped. <code>lun show -g igroup-name</code> displays the LUNs mapped to a specific igroup.
System requirements	Verify that the components of your configuration are supported. Verify that you have the correct host Operating System (OS) Service Pack level, initiator version, Data ONTAP version, and other system requirements. Refer to the appropriate interoperability matrix for your N series product, available on the IBM support website.

FC troubleshooting

The troubleshooting sections that follow provide tips for dealing with problems that might occur when you are running the FC protocol with the Host Utilities.

(FC) Warnings displayed during Host Utilities installation

Occasionally, you might see warnings during the Host Utilities installation.

These warnings might be similar to the following:

- Warning: libnl.so (32-bit) library not found, some sanlun commands may not work. Refer Linux Host Utilities Installation and Setup Guide for more details
- Warning: libHBAAPI.so (32-bit) library not found, some sanlun commands may not work. Refer Linux Host Utilities Installation and Setup Guide for more details
- Warning: libdevmapper.so (32-bit) library not found, some sanlun commands may not work. Refer Linux Host Utilities Installation and Setup Guide for more details

To avoid these warnings, make sure you install the packages that provide the libraries before you install the Host Utilities software. For more information, see the information on installing and configuring QLogic and Emulex HBAs.

(FC) Linux hosts with QLogic initiators fail to detect LUNs

A Linux host with QLogic initiators fails to detect LUNs mapped to it when configured to autonegotiate port topology and speed. This can cause the host to miss a few LUNs 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 set the following settings for port topology and speed using the QConvergeConsole CLI 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/QConvergeConsoleCLI/quacli
```
 - 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 the QLogic port speed to the highest speed possible, depending on its maximum and the maximum of the switch or target port to which it is connected.
 - a) Run the following command:


```
/opt/QLogic_Corporation/QConvergeConsoleCLI/quacli
```
 - 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) 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 appropriate errata multipathing package from the Red Hat Network website. For more information, see appropriate interoperability matrix for your N series product, available on the IBM support website.

(FC) Poor performance due to misalignment of partitioned LUNs

If you experience poor performance on some or all FC LUNs that have partitions, it might be due to an alignment mismatch between the LUN block boundaries and the underlying WAFL boundaries.

This problem occurs only with partitions and only with certain combinations of disk geometry parameters.

You can resolve the issue by using the `fdisk` command to manually enter the disk geometry parameters when creating partitions. For specific steps, go to the interoperability matrix and search for “S1002716”.

(FC) sanlun command displays error when HBA libHBAAPI.so is missing

The `sanlun` command uses the HBAAPI software provided by both the host operating system and the HBA vendor's API plugins to gather information about the HBAs, including their WWPNs. If this software is not available, the `sanlun` command returns an error message.

Error messages similar to the following indicate that the operating system's `libHBAAPI` library is not installed:

- Unable to locate /usr/lib/libHBAAPI.so library
- Unable to locate /usr/lib64/libHBAAPI.so library

Error messages similar to the following indicate that the HBA vendor API plug-in is not installed:

- No supported adapters are present
- Unable to load HBA control library

To avoid this problem, make sure you have the correct management software package for your HBA and host architecture installed:

- For QLogic HBAs, install the QLogic QConvergeConsole CLI package.
- For Emulex HBAs, install the Emulex OneCommand Manager core application (CLI) package.

Note: If you are using a different HBA that is supported by the Host Utilities, you must install the management software for that HBA.

These packages are available from the HBA vendor.

(FC) LUNs are unavailable after a storage system failover or giveback

The host requires specific driver, HBA, and FC switch zoning settings to maintain LUN access across a storage system takeover, giveback, and reboot. You should confirm that 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`

(FC) 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.

Warning messages displayed when using LVM with multipathing software

When you use LVM in a multipathing environment, it displays a warning message.

In a default configuration, LVM scans all attached disks and identifies the disks that contain physical volumes. In environments that use multipathing software, each path to a particular LUN is registered as a different SCSI device. LVM detects this and chooses one of the devices. LVM then displays a warning message.

To avoid this problem, you should modify the paths to the `preferred_names` parameter in the `/etc/lvm/lvm.conf` file.

The following is an example of how the `preferred_names` parameter line should look:

```
preferred_names = [ "^/dev/mapper/*" ]
```

After you make the change, perform a rescan (`pvscan` and `vgscan`) to ensure all devices are properly displayed.

iSCSI troubleshooting

Sometimes you might encounter a problem while running iSCSI. The sections that follow provide tips for resolving any issues that might occur.

(iSCSI) LVM devices are not automatically mounted 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
```

(iSCSI) LVM devices are not automatically mounted during system boot on SUSE Linux Enterprise Server 10

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 for each logical volumes, the logical volumes can be automatically mounted by adding entries in `/etc/fstab`.

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) Poor performance due to misalignment of partitioned LUNs

If you experience poor performance on some or all 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 IBM N series storage system and search for “S1002716”.

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/sd*|", "r|/dev/disk/.*/|", "r|/dev/block/.*/|", "a/.*/"]`

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.
Network connectivity	Verify that there is TCP/IP connectivity between the host and the storage system. From the storage system command line: <ul style="list-style-type: none"> • 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 <i>Data ONTAP SAN Administration Guide for 7-Mode</i> .
Initiator login	Verify that the initiator is logged in to the storage system by entering the <code>iscsi show initiator</code> command on the storage system console. If the initiator is configured and logged in to the storage system, the storage system console displays the initiator node name and the target portal group to which it is connected. If the command output shows 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.

Configuration setting	What you should do
iSCSI node names	<p>Verify that you are using the correct initiator node names in the igroup configuration.</p> <p>On the storage system, use the <code>igroup show</code> command to display the node name of the initiators in the storage system's igroups. On the host, use the initiator tools and commands to display the initiator node name. The initiator node names configured in the igroup and on the host should match.</p>
LUN mappings	<p>Verify that the LUNs are mapped to an igroup.</p> <p>On the storage system, use one of the following commands:</p> <ul style="list-style-type: none"> • <code>lun show -m</code> displays all the LUNs and the igroups to which they are mapped. • <code>lun show -g igroup-name</code> displays the LUNs mapped to a specific igroup.
System requirements	<p>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. Refer to the appropriate interoperability matrix for your N series product, available on the IBM support website.</p>
Jumbo frames	<p>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.</p>
Firewall settings	<p>Verify that the iSCSI port (3260) is open in the firewall rule.</p>

Veritas Storage Foundation troubleshooting

Sometimes you might encounter a problem while running Veritas Storage Foundation with the Host Utilities. The sections that follow provide tips for resolving any issues that might occur.

(Veritas) 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.

(Veritas) Secondary paths identified incorrectly

Secondary paths are identified incorrectly as primary paths while configuring VxDMP. In such a scenario, you should verify that the ASL for the storage system was installed correctly and that no error messages occurred during installation.

(Veritas) Enclosure-based naming not reflected on N series storage

The VxVM enclosure-based naming feature is enabled, but the VxDMP devices do not contain a storage model string. In such a scenario, verify that the ASL for the storage system was installed correctly and that no error messages occurred during installation.

(Veritas) Setting up a LUN for SAN booting fails

The process to set up a LUN for SAN booting fails if you have multiple paths configured to the boot LUN before installing the operating system. You should verify that the host can see only one path to the LUN.

(Veritas) Encapsulation of the root disk fails

Encapsulation of the root disk fails after SAN boot installation. Ensure that you follow the recommended partitioning scheme.

For more information on partitioning scheme, see the *Veritas Volume Manager Administrator's Guide*.

Installing the nSANity data collection program

Obtain and install the nSANity Diagnostic and Configuration Data Collector program when instructed to do so by your technical support representative.

About this task

Contact your technical support representative to obtain the nSANity Diagnostic and Configuration Data Collector program. The nSANity program replaces the diagnostic programs included in previous versions of the Host Utilities. The nSANity program runs on a Windows or Linux system with network connectivity to the component from which you want to collect data.

Steps

1. Obtain the Windows zip or Linux tgz version of the nSANity program from your technical support representative and copy it to the workstation or server that you want to run it on.
2. Change to the directory to which you downloaded the zip or tgz file.
3. Extract all of the files and follow the instructions in the `README.txt` file. Also be sure to review the `RELEASE_NOTES.txt` file for any warnings and notices.

After you finish

Run the specific nSANity commands specified by your technical support representative.

Sample configuration files for Red Hat Enterprise Linux 6

All versions of Red Hat Enterprise Linux 6 series use a DM-Multipath configuration file, but there might be slight variations in the file based on which Red Hat update you have installed.

You can use the sample Red Hat Enterprise Linux 6 series configuration files shown here to create your own multipath.conf file. When you create your file, keep the following in mind:

Red Hat Enterprise Linux 6 series notes	Explanation
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
Blacklisting the local device	If you want to blacklist the local disk, you must put the WWID of the local disk in the blacklist section. You do not need to add other devnode lines because DM-Multipath adds them by default.
SAN boot LUNs and the <code>user_friendly_names</code> parameter	It is recommended that you set the <code>user_friendly_names</code> parameter to <code>no</code> . There have been reports of problems when this parameter is set to <code>yes</code> .

Red Hat Enterprise Linux 6 series, SAN boot LUNs, and `user_friendly_names` parameter

When you create a SAN boot LUN, the installer sets the `user_friendly_names` parameter to `yes` by default. It is recommended that you set the `user_friendly_names` parameter to `no`. There have been reports of problems when this parameter is set to `yes`.

If you create a SAN boot LUN and the installer sets the `user_friendly_names` parameter to `yes`, you must perform the following steps.

1. Change the `user_friendly_names` parameter to `no`.

Note: For **Red Hat Enterprise Linux 6.4 series**, replace first step with the following and continue with steps 2 to 5:

- Create an empty multipath.conf and all the settings for both RHEL 6.4 with and without ALUA are automatically updated by default.

2. Make a backup of initrd-image.

3. Re-create the initrd-image using the command `mkinitrd`.

Red Hat Enterprise Linux 6 series and later use either the command

```
mkinitrd -f /boot/ initramfs-`uname -r`.img `uname -r`
```

or the command

```
dracut -f
```

4. Change the root dm-multipath device name to the WWID-based device name in all the locations that refer to the device, such as `/etc/fstab` and `/boot/grub/device.map`.

For example, suppose the name of the root device is `/dev/mapper/mpatha` and the WWID of the device is `360a98000486e2f66426f583133796572`. You must recreate the initrd-image. Then you must change the device name to `/dev/mapper/360a98000486e2f66426f583133796572` in `/etc/fstab` and `/boot/grub/device.map` and any other place that that refers to device `/dev/mapper/mpatha`. After that, reboot the host.

5. Append the following parameter value to the kernel for ALUA and non-ALUA to work:

```
rdloaddriver=scsi_dh_alua
```

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/  
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/  
lv_root LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD  
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc  
KEYTABLE=us rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

6. Reboot the host.

Red Hat Enterprise Linux 6 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 with ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {  
    user_friendly_names    no  
    max_fds                max  
    flush_on_last_del      yes  
    queue_without_daemon   no  
}  
# All data under blacklist must be specific to your system.  
blacklist {  
    devnode "^hd[a-z]"  
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"  
    devnode "^cciss.*"  
}  
devices {  
    device {  
        vendor  
            "IBM"
```



```

        product                "LUN"
        path_grouping_policy    group_by_prio
        features                "3 queue_if_no_path pg_init_retries 50"
        prio                    "alua"
        path_checker             tur
        failback                 immediate
        path_selector            "round-robin 0"
        hardware_handler        "1 alua"
        rr_weight                uniform
        rr_min_io                128
        getuid_callout           "/lib/udev/scsi_id -g -u -d /dev/%n"
    }
}

```

Red Hat Enterprise Linux 6 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults {
    user_friendly_names    no
    max_fds                 max
    flush_on_last_del      yes
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw/loop/fd/md/dm-|sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio                   "ontap"
        path_checker           tur
        failback                immediate
        path_selector           "round-robin 0"
        hardware_handler        "0"
        rr_weight               uniform
        rr_min_io               128
        getuid_callout          "/lib/udev/scsi_id -g -u -d /dev/%n"
    }
}

```

Red Hat Enterprise Linux 6 update 1 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 update 1 with ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
    queue_without_daemon   no
    dev_loss_tmo           2147483647
    fast_io_fail_tmo       5
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode  "^hd[a-z]"
    devnode  "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode  "^cciss.*"
}
devices {
    device {
        vendor            "IBM"
        product           "LUN"
        path_grouping_policy group_by_prio
        features           "3 queue_if_no_path pg_init_retries 50"
        prio              "alua"
        path_checker       tur
        failback           immediate
        path_selector      "round-robin 0"
        hardware_handler   "1 alua"
        rr_weight          uniform
        rr_min_io          128
        getuid_callout     "/lib/udev/scsi_id -g -u -d /dev/%n"
    }
}
```

Red Hat Enterprise Linux 6 update 1 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 update 1 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    max_fds                max
```

```

        flush_on_last_del      yes
        queue_without_daemon   no
        dev_loss_tmo           2147483647
        fast_io_fail_tmo       5
    }
    # All data under blacklist must be specific to your system.
    blacklist {
        devnode "^hd[a-z]"
        devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
        devnode "^cciss.*"
    }
    devices {
        device {
            vendor              "IBM"
            product              "LUN"
            path_grouping_policy group_by_prio
            features              "3 queue_if_no_path pg_init_retries 50"
            prio                  "ontap"
            path_checker          tur
            failback               immediate
            path_selector          "round-robin 0"
            hardware_handler       "0"
            rr_weight              uniform
            rr_min_io              128
            getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/%n"
        }
    }
}

```

Red Hat Enterprise Linux 6 update 2 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 update 2 with ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults {
    user_friendly_names      no
    max_fds                   max
    flush_on_last_del        yes
    queue_without_daemon     no
    dev_loss_tmo              infinity
    fast_io_fail_tmo         5
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor              "IBM"
        product              "LUN"
        path_grouping_policy group_by_prio
    }
}

```

```

        features                "3 queue_if_no_path pg_init_retries 50"
        prio                    "alua"
        path_checker             tur
        failback                 immediate
        path_selector            "round-robin 0"
        hardware_handler        "1 alua"
        rr_weight                uniform
        rr_min_io               128
        getuid_callout           "/lib/udev/scsi_id -g -u -d /dev/%n"
    }
}

```

Red Hat Enterprise Linux 6 update 2 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 update 2 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
    queue_without_daemon   no
    dev_loss_tmo           infinity
    fast_io_fail_tmo       5
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio                  "ontap"
        path_checker           tur
        failback               immediate
        path_selector          "round-robin 0"
        hardware_handler       "0"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/%n"
    }
}

```

Red Hat Enterprise Linux 6 update 3 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 6 update 3 with ALUA enabled. By default, the hardware table sets the rest of the parameters.

```
devices {
    device {
        vendor            "IBM"
        product           "LUN.*"
        prio              "alua"
        hardware_handler  "1 alua"
    }
}
```

Red Hat Enterprise Linux 6 update 3 without ALUA enabled sample configuration file

When you are not using ALUA, you only need to list any devices that must be blacklisted. All the other parameter values are set by the hardware table.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

If you are using a SAN boot LUN and must blacklist the local disk, you must supply the WWID of the local disk. You do not need to add other devnode information. DM-Multipath adds that information by default.

Remember: When you use the blacklist section, you must **replace** the sample information with information for your system.

```
# All data under blacklist must be specific to your system.
blacklist {
    wwid < wwid_of_the_local_disk>
}
```

Red Hat Enterprise Linux 6 update 4 with and without ALUA enabled sample configuration file

Refer to the steps listed in section *Red Hat Enterprise Linux 6 series, SAN boot LUNs, and user_friendly_names parameter* for setting the parameters and then continue to the steps 1 and 2.

Remember: When you use the blacklist section, you must **replace** the sample information with information for your system.

```
# All data under blacklist must be specific to your system.
blacklist {
    wwid < wwid_of_the_local_disk>
}
```

1. Append the following parameter value to the kernel for ALUA and non-ALUA to work and then reboot: `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/  
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/  
lv_root LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD  
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc  
KEYTABLE=us rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

2. Verify the output of `cat /proc/cmdline` command to ensure that the setting is complete.

Sample configuration file for Red Hat Enterprise Linux 5

All versions of Red Hat Enterprise Linux 5 series use a DM-Multipath configuration file, but there might be slight variations in the file based on which Red Hat update you have installed. You can replace your current file with this sample file and change the values to ones that are appropriate for your system.

You can use the sample Red Hat Enterprise Linux 5 series configuration files shown here to create your own multipath.conf file. When you create your file, keep the following in mind:

Red Hat Enterprise Linux 5 series notes	Explanation
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
SAN boot LUNs and the <code>user_friendly_names</code> parameter	It is recommended that you set the <code>user_friendly_names</code> parameter to <code>no</code> . There have been reports of problems when this parameter is set to <code>yes</code> .
Red Hat Enterprise Linux 5 prior to Update 6	If you are using a version of Red Hat Enterprise Linux 5 series prior to update 6, check the <i>Recommended Host Settings for Linux Host Utilities</i> to see if there are any parameter values specific to that version.

Red Hat Enterprise Linux 5 series, SAN boot LUNs, and `user_friendly_names` parameter

When you create a SAN boot LUN, the installer sets the `user_friendly_names` parameter to `yes` by default. It is recommended that you set the `user_friendly_names` parameter to `no`. There have been reports of problems when this parameter is set to `yes`.

If you create a SAN boot LUN and the installer sets the `user_friendly_names` parameter to `yes`, you must perform the following steps.

1. Change the `user_friendly_names` parameter to `no`.
2. Make a backup of `initrd-image`.
3. Re-create the `initrd-image` using the command `mkinitrd`.

You could use the following command line:

```
mkinitrd -f /boot/initrd-"`uname -r`".img `uname -r`
```

4. Change the root dm-multipath device name to the WWID-based device name in all the locations that refer to the device, such as `/etc/fstab` and `/boot/grub/device.map`.
5. Reboot the host.

For example, suppose the name of the root device is `/dev/mapper/mpatha` and the WWID of the device is `360a98000486e2f66426f583133796572`. You must recreate the `initrd-image`. Then you must change the device name to `/dev/mapper/360a98000486e2f66426f583133796572` in `/etc/fstab` and `/boot/grub/device.map` and any other place that that refers to device `/dev/mapper/mpatha`. After that, reboot the host.

Red Hat Enterprise Linux 5 update 9, 8 or update 7 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 5 with update 8 or update 7 and has ALUA enabled.

Note: Both Red Hat Enterprise Linux 5 update 8 and Red Hat Enterprise Linux 5 update 7 use the same values in the DM-Multipath configuration file, so this file can apply to either version.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    queue_without_daemon   no
    flush_on_last_del      yes
    max_fds                 max
    pg_prio_calc            avg
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio_callout           "/sbin/mpath_prio_alua /dev/%n"
        path_checker           tur
        path_selector           "round-robin 0"
        failback               immediate
        hardware_handler       "1 alua"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
    }
}
```


Red Hat Enterprise Linux 5 update 9, 8, or update 7 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 5 with update 9, 8, or update 7 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    queue_without_daemon   no
    flush_on_last_del      yes
    max_fds                 max
    pg_prio_calc            avg
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw/loop/fd/md/dm-|sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor            "IBM"
        product           "LUN"
        path_grouping_policy group_by_prio
        features           "3 queue_if_no_path pg_init_retries 50"
        prio_callout       "/sbin/mpath_prio_ontap /dev/%n"
        path_checker       tur
        path_selector      "round-robin 0"
        failback           immediate
        hardware_handler   "0"
        rr_weight           uniform
        rr_min_io          128
        getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
    }
}
```

Red Hat Enterprise Linux 5 update 6 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 5 with update 6 and has ALUA enabled:

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
```

```

        queue_without_daemon    no
        flush_on_last_del       yes
        max_fds                 max
        pg_prio_calc            avg
    }
    # All data under blacklist must be specific to your system.
    blacklist {
        devnode "^hd[a-z]"
        devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
        devnode "^cciss.*"
    }
    devices {
    device {
        vendor                    "IBM"
        product                   "LUN"
        path_grouping_policy      group_by_prio
        features                   "1 queue_if_no_path"
        prio_callout              "/sbin/mpath_prio_alua /dev/%n"
        path_checker              directio
        path_selector             "round-robin 0"
        failback                  immediate
        hardware_handler          "1 alua"
        rr_weight                 uniform
        rr_min_io                 128
        getuid_callout            "/sbin/scsi_id -g -u -s /block/%n"
    }
    }
}

```

Red Hat Enterprise Linux 5 update 6 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 5 with update 6 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults {
    user_friendly_names    no
    queue_without_daemon    no
    flush_on_last_del       yes
    max_fds                 max
    pg_prio_calc            avg
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
device {

```

```
vendor          "IBM"
product         "LUN"
path_grouping_policy group_by_prio
features        "1 queue_if_no_path"
prio_callout    "/sbin/mpath_prio_ontap /dev/%n"
path_checker    directio
path_selector   "round-robin 0"
failback        immediate
hardware_handler "0"
rr_weight        uniform
rr_min_io        128
getuid_callout  "/sbin/scsi_id -g -u -s /block/%n"
}
}
```

Sample configuration file for Red Hat Enterprise Linux 4

All versions of Red Hat Enterprise Linux 4 series use a DM-Multipath configuration file, but there might be slight variations in the file based on which Red Hat update you have installed. You can replace your current file with this sample file and change the values to ones that are appropriate for your system.

Configuration Notes for DM-Multipath

Please review the following configuration notes before you set up your configuration file to enable DM-Multipath on Red Hat Enterprise Linux 4 series:

Red Hat Enterprise Linux 4 series version	Parameter Notes
4 Update 7 or later	<ul style="list-style-type: none">Use the options <code>flush_on_last_del</code> and <code>max_fds</code>Set <code>pathchecker</code> to <code>directio</code>
(iSCSI only) 4 Update 7 and earlier	In environments running iSCSI, set <code>path_grouping_policy</code> to <code>multibus</code>
4 Update 7 and earlier	Place <code>rr_min_io</code> in the default section, not the device section, of the <code>multipath.conf</code> file and set its value to 128.
4 Update 6 and earlier	Set <code>pathchecker</code> to <code>readsector0</code> .

Red Hat Enterprise Linux 4 Update 9 sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Linux 4 Update 9.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults
{
    user_friendly_names    no
    queue_without_daemon   no
    max_fds                 max
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
devnode_blacklist
{
```

```

devnode "^hd[a-z]"
devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
devnode "^cciss.*"
}
devices
{
    device
    {
        vendor            "IBM"
        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
        path_selector       "round-robin 0"
        failback           immediate
        rr_weight           uniform
        rr_min_io          128
        path_checker        directio
    }
}

```

Sample configuration file for Oracle Linux

All versions of Oracle Linux use a DM-Multipath configuration file, but there might be slight variations in the file based on which Oracle update you have installed. You can replace your current file with sample file and change the values to ones that are appropriate for your system.

In most cases, Oracle Linux uses the same setup procedures as Red Hat Enterprise Linux. To simplify this guide, it uses "Red Hat" to refer to both systems using Red Hat Enterprise Linux and systems using Oracle Linux. If Oracle Linux requires a different procedure, that procedure is included.

Note: Ensure that the kernel and dm-multipath versions are as per interoperability matrix tool. If not, install the versions from Oracle ULN (Oracle Unbreakable Linux Network).

Note: Ensure to use the RHEL based setting when using Red Hat Compatible Kernel and multipath packages.

Sample configuration file for Red Hat Enterprise Virtualization Hypervisor

All versions of Red Hat Enterprise Virtualization Hypervisor (RHEV) use a DM-Multipath configuration file, but there might be slight variations in the file. You can replace your current file with this sample file and change the values to ones that are appropriate for your system.

You can use the sample Red Hat Enterprise Virtualization Hypervisor configuration files shown here to create your own `multipath.conf` file. When you create your file, keep the following in mind:

Red Hat Enterprise Virtualization Hypervisor	Explanation
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
wwid <Dev ID>	Use the actual device ID. Run the <code>/lib/udev/scsi_id -g -u -d /dev/<sd_device></code> command on the SCSI drive to be blacklisted.
Revision and private tags	The first and second lines of the file should be the revision and private tags: <pre># RHEV REVISION # RHEV PRIVATE</pre> These tags tell the VDSM not to overwrite the configuration file. If there is no tag then configuration file will be overwritten with the defaults after server reboot/update.
multipath.conf and persistence	You must set the <code>multipath.conf</code> file to persist across reboots.

DM-Multipath configuration file requires command to persist across reboots

By default, the DM-Multipath configuration file for Red Hat Enterprise Virtualization Hypervisor does not persist across reboots.

You must enter the following command to ensure that it persists:

```
persist /etc/multipath.conf
```

Red Hat Enterprise Virtualization Hypervisor 6.2 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Virtualization Hypervisor 6.2 with ALUA enabled:

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
# RHEV REVISION 0.7
# RHEV PRIVATE
defaults {
    user_friendly_names    no
    max_fds                 max
    flush_on_last_del      yes
    queue_without_daemon   no
    dev_loss_tmo            infinity
    fast_io_fail_tmo       5
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid "<wwid_of_the_local_disk>"
    devnode "^(ram/raw/loop/fd/md/dm-|sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio                   "alua"
        path_checker           tur
        failback               immediate
        path_selector          "round-robin 0"
        hardware_handler       "1 alua"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/
%n"
    }
}
```

Red Hat Enterprise Virtualization Hypervisor 6.2 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Virtualization Hypervisor 6.2 and does not have ALUA enabled:

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
# RHEV REVISION 0.7
# RHEV PRIVATE
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
    queue_without_daemon   no
    dev_loss_tmo            infinity
    fast_io_fail_tmo       5
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid <wwid_of_the_local_disk>
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio                   "ontap"
        path_checker           tur
        failback               immediate
        path_selector           "round-robin 0"
        hardware_handler       "0"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -g -u -d /dev/
%n"
    }
}
}
```

Red Hat Enterprise Virtualization Hypervisor 6.3 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Virtualization Hypervisor 6.3 with ALUA enabled:

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
# RHEV REVISION 0.8
# RHEV PRIVATE
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
```

```

        queue_without_daemon    no
        dev_loss_tmo            infinity
        fast_io_fail_tmo        5
    }
    # All data under blacklist must be specific to your system.
    blacklist {
        wwid <wwid_of_the_local_disk>
    }
    devices {
        device {
            vendor                "IBM"
            product               "LUN.*"
            prio                  "alua"
            hardware_handler       "1 alua"
        }
    }
}

```

Red Hat Enterprise Virtualization Hypervisor 6.3 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Red Hat Enterprise Virtualization Hypervisor 6.3 without ALUA enabled:

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

# RHEV REVISION 0.8
# RHEV PRIVATE
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del       yes
    queue_without_daemon    no
    dev_loss_tmo            infinity
    fast_io_fail_tmo        5
}
# All data under blacklist must be specific to your system.
blacklist {
    wwid <wwid_of_the_local_disk>
}
devices {
    device {
        vendor                "IBM"
        product               "LUN.*"
        prio                  "ontap"
        hardware_handler       "0"
    }
}
}

```

Sample configuration file for Oracle VM 3.0.1, 3.1.1, and 3.2.1

All versions of Oracle VM use a DM-Multipath configuration file, but there might be slight variations in the file based on which Oracle update you have installed. You can replace your current file with this sample file and change the values to ones that are appropriate for your system.

You can use the sample Oracle VM configuration files shown here to create your own `multipath.conf` file.

When you create your file, keep the following in mind:

Oracle VM 3.0.1, 3.1.1, and 3.2.1	Parameter notes
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
wwid <Dev ID>	Use the actual device ID. Run the <code>/lib/udev/scsi_id -gus /block/%n</code> command on the SCSI drive to be blacklisted.
Multipath settings and heartbeat settings	Multipath settings and heartbeat settings must be the same on all hypervisors in a server pool.
SAN boot support	SAN booting is supported only with Oracle VM 3.1.1 and the FC protocol.
SAN boot LUNs and the <code>user_friendly_names</code> parameter	It is recommended that you set the <code>user_friendly_names</code> parameter to <code>no</code> . There have been reports of problems when this parameter is set to <code>yes</code> .

Oracle VM, SAN boot LUNs, and `user_friendly_names` parameter

The recommended value for the `user_friendly_names` parameter is `no`. There have been reports of problems when this parameter is set to `yes`.

If you create a SAN boot LUN and the installer sets the `user_friendly_names` parameter to `yes`, you must perform the following steps.

1. Change the `user_friendly_names` parameter to `no`.
2. Make a backup of `initrd-image`.
3. Re-create the `initrd-image` using the command `mkinitrd`.
You could use the following command line:

```
mkinitrd -f /boot/initrd-`uname -r`.img `uname -r`
```

4. Change the root dm-multipath device name to the WWID-based device name in all the locations that refer to the device, such as `/etc/fstab` and `/boot/grub/device.map`.
5. Reboot the host.

Oracle VM 3.0.1 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM 3.0.1 with ALUA enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid "<wwid_of_the_local_disk>"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "1 queue_if_no_path"
        prio                   "alua"
        path_checker           directio
        no_path_retry          "queue"
        failback               immediate
        hardware_handler       "1 alua"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -gus /block/%n"
    }
}
```

Oracle VM 3.0.1 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM 3.0.1 and ALUA is not enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26. Also, unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid "<wwid_of_the_local_disk>"
    devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "1 queue_if_no_path"
        prio                   "ontap"
        path_checker           directio
        no_path_retry          "queue"
        failback               immediate
        hardware_handler       "0"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -gus /block/%n"
    }
}
```

Oracle VM 3.1.1 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM 3.1.1 with ALUA enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26.

```
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
```

```

        queue_without_daemon    yes
    }
    # All data under blacklist must be specific to your system.
    blacklist {
        devnode    "^hd[a-z]"
        wwid    "<wwid_of_the_local_disk>"
        devnode    "^(ram/raw|loop/fd/md/dm-/sr/scd/st)[0-9]*"
        devnode    "^cciss.*"
    }
    devices {
        device {
            vendor            "IBM"
            product           "LUN.*"
            path_grouping_policy    group_by_prio
            features          "3 queue_if_no_path pg_init_retries 50"
            prio              "alua"
            path_checker      tur
            no_path_retry     "queue"
            failback          immediate
            hardware_handler  "1 alua"
            rr_weight         uniform
            rr_min_io         128
            getuid_callout    "/lib/udev/scsi_id -gus /block/%n"
        }
    }
}

```

Oracle VM 3.1.1 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM and ALUA is not enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26. Also, unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

```

defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon    yes
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode    "^hd[a-z]"
    wwid    "<wwid_of_the_local_disk>"
    devnode    "^(ram/raw|loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode    "^cciss.*"
}
devices {
    device {
        vendor            "IBM"
        product           "LUN.*"
        path_grouping_policy    group_by_prio
    }
}

```

```

        features                "3 queue_if_no_path pg_init_retries 50"
        prio                    "ontap"
        path_checker             tur
        no_path_retry            "queue"
        failback                 immediate
        hardware_handler         "0"
        rr_weight                uniform
        rr_min_io               128
        getuid_callout           "/lib/udev/scsi_id -gus /block/%n"
    }
}

```

Oracle VM 3.2.1 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM and ALUA is enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26. Also, unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

```

defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid "<wwid_of_the_local_disk>"
    devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN.*"
        path_grouping_policy  group_by_prio
        features               "3 queue_if_no_path pg_init_retries 50"
        prio                  "alua"
        path_checker           tur
        no_path_retry          "queue"
        failback               immediate
        hardware_handler       "1 alua"
        rr_weight              uniform
        rr_min_io              128
        getuid_callout         "/lib/udev/scsi_id -gus /block/%n"
    }
}

```

Oracle VM 3.2.1 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Oracle VM and ALUA is not enabled.

Note: In addition to providing a DM-Multipath configuration file, you must also set the `O2CB_HEARTBEAT_THRESHOLD` timeout. For more information see, [\(Oracle VM\) Configuring the O2CB_HEARTBEAT_THRESHOLD](#) on page 26. Also, unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

```
defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid "<wwid_of_the_local_disk>"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
devices {
    device {
        vendor            "IBM"
        product           "LUN.*"
        path_grouping_policy group_by_prio
        features           "3 queue_if_no_path pg_init_retries 50"
        prio              "ontap"
        path_checker       tur
        no_path_retry      "queue"
        failback           immediate
        hardware_handler   "0"
        rr_weight          uniform
        rr_min_io          128
        getuid_callout     "/lib/udev/scsi_id -gus /block/%n"
    }
}
```


Sample configuration files for SUSE Linux Enterprise Server 11 series

All versions of SUSE Linux Enterprise Server use a DM-Multipath configuration file, but there might be slight variations in the file based on which version and update of SUSE Linux Enterprise server you have. If you are running SUSE Linux Enterprise Server 11, 11 SP1, or 11 SP2, you can replace your current file with one of the sample configuration files and then change the values to ones that are appropriate for your system.

You can use the sample SUSE Linux Enterprise Server 11 series configuration files shown here to create your own `multipath.conf` file. When you create your file, keep the following in mind:

SUSE Linux Enterprise Server 11 series notes	Explanation
Blacklist section	<p>You must provide information of your system in the blacklist section.</p> <p>Any names shown in the sample files are examples and will not work with your system.</p>

SUSE Linux Enterprise Server 11, SP1 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running either SUSE Linux Enterprise Server 11 or 11 SP1 with ALUA.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults
{
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode                "^hd[a-z]"
    wwid                    "<wwid_of_the_local_disk>"
    devnode                "^(/ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode                 "^cciss.*"
}
devices
{
    device
    {
        vendor              "IBM"
        product             "LUN"
```

```

        getuid_callout      "/lib/udev/scsi_id -g -u -d /dev/%n"
        prio               "alua"
        features           "1 queue_if_no_path"
        hardware_handler    "1 alua"
        path_grouping_policy group_by_prio
        path_selector       "round-robin 0"
        failback            immediate
        rr_weight           uniform
        rr_min_io           128
        path_checker        tur
    }
}

```

SUSE Linux Enterprise Server 11, 11 SP1 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running SUSE Linux Enterprise Server 11 or 11 SP1 and ALUA is not enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults
{
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode      "^hd[a-z]"
    wwid         "<wwid_of_the_local_disk>"
    devnode      "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode      "^cciss.*"
}
devices
{
    device
    {
        vendor      "IBM"
        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
        path_selector "round-robin 0"
        failback     immediate
    }
}

```

```

        rr_weight          uniform
        rr_min_io          128
        path_checker       tur
    }
}

```

SUSE Linux Enterprise Server 11 SP2 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running SUSE Linux Enterprise Server 11 SP2 KVM and Xen with ALUA.

Note: This configuration file applies to both SLES 11 SP2 KVM and Xen also.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults
{
    user_friendly_names    no
    max_fds                max
    queue_without_daemon   no
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode      "^hd[a-z]"
    wwid         "<wwid_of_the_local_disk>"
    devnode      "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode      "^cciss.*"
}
devices
{
    device
    {
        vendor          "IBM"
        product         "LUN"
        getuid_callout   "/lib/udev/scsi_id -g -u -d /dev/%n"
        prio            "alua"
        features        "3 queue_if_no_path pg_init_retries 50"
        hardware_handler "1 alua"
        path_grouping_policy group_by_prio
        failback         immediate
        rr_weight        uniform
        rr_min_io        128
        path_checker     tur
    }
}

```

SUSE Linux Enterprise Server 11 SP2 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running SUSE Linux Enterprise Server 11 SP2 KVM and Xen when ALUA is not enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Note: This configuration file applies to both SLES 11 SP2 KVM and Xen also.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults
{
    user_friendly_names    no
    max_fds                 max
    queue_without_daemon   no
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode    "^hd[a-z]"
    wwid       "<wwid_of_the_local_disk>"
    devnode    "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode    "^cciss.*"
}
devices
{
    device
    {
        vendor            "IBM"
        product           "LUN"
        getuid_callout    "/lib/udev/scsi_id -g -u -d /dev/%n"
        prio              "ontap"
        features          "3 queue_if_no_path pg_init_retries 50"
        hardware_handler  "0"
        path_grouping_policy group_by_prio
        failback          immediate
        rr_weight         uniform
        rr_min_io         128
        path_checker       tur
    }
}
```

Sample configuration file for SUSE Linux Enterprise Server 10

All versions of SUSE Linux Enterprise Server use a DM-Multipath configuration file, but there might be slight variations in the file based on which version and update of SUSE Linux Enterprise server you have. If you are running SUSE Linux Enterprise Server 10, you can replace your current file with this sample configuration files and then change the values to ones that are appropriate for your system.

Review the following configuration notes before you set up your configuration file to enable DM-Multipath on SUSE Linux Enterprise Server 10:

SUSE Linux Enterprise Server 10 series notes	Explanation
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
SUSE Linux Enterprise Server 10 SP3 and earlier	If you are using a version of SUSE Linux Enterprise Server 10 series prior to SP4, check the <i>Recommended Host Settings for Linux Host Utilities</i> to see if there are any parameter values specific to that version.

SUSE Linux Enterprise Server 10 SP4 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running SUSE Linux Enterprise Server 10 SP4 and has ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```
defaults
{
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode "^hd[a-z]"
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^cciss.*"
}
```

```

devices
{
    device
    {
        vendor            "IBM"
        product           "LUN"
        getuid_callout     "/sbin/scsi_id -g -u -s /block/%n"
        prio              "alua"
        features           "1 queue_if_no_path"
        hardware_handler   "1 alua"
        path_grouping_policy group_by_prio
        path_selector      "round-robin 0"
        failback           immediate
        rr_weight          uniform
        rr_min_io          128
        path_checker       tur
    }
}

```

SUSE Linux Enterprise Server 10 SP4 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running SUSE Linux Enterprise Server 10 SP4 and does not have ALUA enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults
{
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      yes
}
# All data under blacklist must be specific to your system.
blacklist
{
    devnode "^hd[a-z]"
    devnode "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode "^cciss.*"
}
devices
{
    device
    {
        vendor            "IBM"
        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  
path_selector         "round-robin 0"  
failback              immediate  
rr_weight             uniform  
rr_min_io             128  
path_checker          tur  
}  
{
```

Sample Configuration for Citrix XenServer

All versions of Citrix XenServer use a DM-Multipath configuration file, but there might be slight variations in the file based on which Citrix update you have installed. You can replace your current file with sample file and change the values to the ones that are appropriate for your system.

You can use the sample Citrix XenServer configuration files shown here to create your own multipath.conf file.

When you create your file, keep the following in mind:

Citrix XenServer 6.0 and 6.0.2	Parameter notes
Blacklist section	You must provide information of your system in the blacklist section. Any names shown in the sample files are examples and will not work with your system.
wwid <Dev ID>	Use the actual device ID. Run the <code>/sbin/scsi_id -g -u -s /block/%n</code> command on the SCSI drive to be blacklisted.
SAN boot support	SAN booting is supported only with Citrix XenServer 6.0, 6.0.2 and the FC/FCoE protocol.
SAN boot LUNs and the <code>user_friendly_names</code> parameter	It is recommended that you set the <code>user_friendly_names</code> parameter to <code>no</code> . There have been reports of problems when this parameter is set to <code>yes</code> .

Citrix XenServer, SAN boot LUNs, and `user_friendly_names` parameter

The recommended value for the `user_friendly_names` parameter is `no`. There have been reports of problems when this parameter is set to `yes`.

If you create a SAN boot LUN and the installer sets the `user_friendly_names` parameter to `yes`, you must perform the following steps.

1. Change the `user_friendly_names` parameter to `no`.
2. Make a backup of `initrd-image`.
3. Place the above `multipath.conf` file in the `/etc` directory.
4. Run the following command:

```
service multipathd reload  
multipath -v3.
```


Note: `multipath -ll` should reflect the changes on the SAN Boot LUN as per the changed `multipath.conf` file.

5. Re-create the `initrd`-image using the command `mkinitrd`.

You could use the following command line:

```
mkinitrd -f /boot/initrd-"`uname -r`.img`" `uname -r`
```

6. Reboot the host.

Note: Currently SAN Boot LUN does not reflect `queue_if_no_path` feature. To overcome this limitation, perform the steps mentioned below:

- Run `multipath` command after boot.

Note: To avoid running `multipath` command on every boot, add `/sbin/multipath` to `/etc/rc.d/rc.local` file.

- Running `multipath` forces the unused LUNs (on which there is no Storage Repository) to get added to the XAPI Layer.
- To remove the `scsi_device` ID of those unused LUNs from `mpathutil` use the following command: `# /opt/xensource/sm/mpathutil.py remove <scsi_device-id>`. Skip this step if you have no unused LUNs (where no SR is created).

Citrix XenServer 6.0 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Citrix XenServer 6.0 with ALUA enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with the information from your system.

```
defaults{
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode    "^hd[a-z]"
    wwid       "<wwid_of_the_local_disk>"
    devnode    "^(ram/raw/loop/fd/md/dm-|sr/scd/st)[0-9]*"
    devnode    "^cciss.*"
}
devices {
    device {
        vendor            "IBM"
        product           "LUN"
        path_grouping_policy group_by_prio
        features           "1 queue_if_no_path"
        prio_callout       "/sbin/mpath_prio_alua /dev/%n"
        path_checker        directio
```

```

        failback                immediate
        hardware_handler        "1 alua"
        rr_weight                uniform
        rr_min_io                128
        getuid_callout           "/sbin/scsi_id -g -u -s /block/%n"
    }
}

```

Citrix XenServer 6.0 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Citrix XenServer 6.0 and ALUA is not enabled.

Remember: If you use the blacklist section, you must **replace** the sample information with information for your system.

```

defaults {
    user_friendly_names    no
    max_fds                 max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode "^hd[a-z]"
    wwid     "<wwid_of_the_local_disk>"
    devnode  "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode  "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy   group_by_prio
        features               "1 queue_if_no_path"
        prio_callout           "/sbin/mpath_prio_ontap /dev/%n"
        path_checker            directio
        failback                immediate
        hardware_handler        "0"
        rr_weight                uniform
        rr_min_io                128
        getuid_callout          "/sbin/scsi_id -g -u -s /block/%n"
    }
}

```

Citrix XenServer 6.0.2 with ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Citrix XenServer 6.0.2 with ALUA enabled.

```

defaults {
    user_friendly_names    no

```

```

        max_fds                max
        flush_on_last_del      no
        queue_without_daemon   no
    }
    # All data under blacklist must be specific to your system.
    blacklist {
        devnode  "^hd[a-z]"
        wwid     "<wwid_of_the_local_disk>"
        devnode  "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
        devnode  "^cciss.*"
    }
    devices {
        device {
            vendor                "IBM"
            product               "LUN"
            path_grouping_policy  group_by_prio
            features               "1 queue_if_no_path"
            prio_callout           "/sbin/mpath_prio_alua /dev/%n"
            path_checker           tur
            failback               immediate
            hardware_handler       "1 alua"
            rr_weight              uniform
            rr_min_io              128
            getuid_callout         "/sbin/scsi_id -g -u -s /block/%n"
        }
    }
}

```

Citrix XenServer 6.0.2 without ALUA enabled sample configuration file

The following file provides an example of the values you need to supply when your host is running Citrix XenServer and ALUA is not enabled.

Note: Unless you are running the iSCSI protocol and Data ONTAP operating in 7-Mode, you should have ALUA enabled.

```

defaults {
    user_friendly_names    no
    max_fds                max
    flush_on_last_del      no
    queue_without_daemon   no
}
# All data under blacklist must be specific to your system.
blacklist {
    devnode  "^hd[a-z]"
    wwid     "<wwid_of_the_local_disk>"
    devnode  "^(ram/raw/loop/fd/md/dm-/sr/scd/st)[0-9]*"
    devnode  "^cciss.*"
}
devices {
    device {
        vendor                "IBM"
        product               "LUN"
        path_grouping_policy  group_by_prio
        features               "1 queue_if_no_path"
    }
}

```

```
    prio_callout          "/sbin/mpath_prio_ontap /dev/%n"
    path_checker          tur
    failback              immediate
    hardware_handler      "0"
    rr_weight             uniform
    rr_min_io             128
    getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
    }
}
```

Where to find more information

This table contains pointers to documents and websites that provide more information about host and storage system requirements, supported configurations, best practices, your operating system, and troubleshooting.

If you need more information about...	See this source...
Known issues, system requirements, and last minute updates	<p>The latest Host Utilities <i>Release Notes</i></p> <p>Note: Generally, the Release Notes are updated more frequently than other documentation. It is a good practice to check the release notes before installing the Host Utilities to see if there have been any changes to the installation or setup process since this document was prepared. After that, you should check the Release Notes periodically to see if there is new information on using the Host Utilities. A summary of what has been updated and when is on the Release Notes index page.</p>
The latest supported configurations and requirements	The N series interoperability matrix website
Configuring the storage system and managing SAN storage on it	<ul style="list-style-type: none"> • Data ONTAP documentation Index • <i>Best Practices for Reliability: New System Installation</i> • <i>Data ONTAP Software Setup Guide for 7-Mode</i> • <i>Data ONTAP SAN Administration Guide for 7-Mode</i> • <i>Data ONTAP Release Notes for 7-Mode</i> • <i>Data ONTAP Commands: Manual Page Reference for 7-Mode, Volume 1</i> and <i>Data ONTAP Commands: Manual Page Reference for 7-Mode, Volume 2</i>
Supported SAN topologies	The <i>Data ONTAP SAN Administration Guide for 7-Mode</i> for your version of Data ONTAP software
Verifying compatibility of a storage system with environmental requirements	<i>IBM System Storage N series Introduction and Planning Guide</i>
Upgrading Data ONTAP	<i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>
Installing and configuring the HBA on the host	HBA vendor documentation

If you need more information about...	See this source...
Enabling ALUA	<i>Data ONTAP SAN Administration Guide for 7-Mode</i> Note: In particular, see the section “Enabling ALUA for a Fibre Channel igroup”.
Host Utilities examples using SUSE Linux Enterprise Server 9	<i>Linux Host Utilities Installation and Setup Guide</i> Note: There were no changes to the information about using SUSE Linux Enterprise Server 9 between the Host Utilities 5.2 release and 5.3 release.
Working with QLogic	Refer to the QLogic documentation
Working with Emulex	Refer to the Emulex documentation
Additional supported HBAs	Check the N series interoperability matrix website (accessed and navigated as described in Websites on page 8).
Broadcom 10GB iSCSI	Check vendor documentation, including the README and the <i>Broadcom NetXtreme II™ Network Adapter User Guide</i>
Information about Oracle VM	<ul style="list-style-type: none"> • Oracle® VM Release Notes • Oracle® VM Installation and Upgrade Guide • Oracle® VM User's Guide Release 3.0 for x86
Information about Oracle Cluster File System	<i>Oracle Cluster File System (OCFS2) User's Guide and OCFS2 - Frequently Asked Questions</i>
Information about Oracle Real Application Cluster (RAC)	Oracle RAC installation information from the Oracle website
Technical reports for Oracle RAC	<ul style="list-style-type: none"> • Technical report: IBM System Storage N series and Oracle Database 11g • Best Practice Guidelines for Oracle
Information about managing the iSCSI initiator	OS vendor documentation
Information about installing and configuring SnapDrive for UNIX	The <i>SnapDrive for UNIX Installation and Administration Guide</i> for your version of SnapDrive
Veritas Storage Foundation and its features	Refer to the Veritas documentation
Information about Oracle VM	Refer to the Oracle documentation

If you need more information about...	See this source...
Information about KVM	Refer to the Linux documentation
Information about configuring multipath support on Red Hat Enterprise Linux	<p>The <code>Multipath-usage.txt</code> file included with the <code>device-mapper-multipath</code> RPM for Red Hat Enterprise Linux 4 Update 3 and later in <code>/usr/share/doc/device-mapper-multipath-<i>version</i></code>.</p> <p><i>version</i> is the latest version number shipped with the Red Hat Enterprise Linux release.</p>

Related information

[*Oracle website*](#)

[*Red Hat website*](#)

[*Novell website*](#)

Copyright and trademark information

Copyright ©1994 - 2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2013 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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

IBM, the IBM logo, and [ibm.com](http://www.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>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP,

ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

- ## A
- alignment
 - VHD partition [96](#)
 - ALUA
 - configuring in `/etc/multipath.conf` [41](#)
 - supported environments [100](#)
 - APM
 - failover path [54](#)
 - I/O error handling [54](#)
 - I/O error processing [54](#)
 - path failures [54](#)
 - Array Policy Module
 - installing [54](#)
 - Array Support Library
 - installing [54](#)
- ## B
- block alignment
 - VHD partition [96](#)
 - Broadcom
 - configuring 10 Gb iSCSI offload [21](#)
- ## C
- Citrix Xen Server
 - sample multipath-conf file [144](#)
 - SAN booting not supported [144](#)
 - command, iSCSI service
 - `chkconfig` [36](#)
 - configure DM-Multipath, command
 - `chkconfig` [44](#)
- ## D
- Device Mapper Multipathing
 - DM-Multipath [99](#)
 - Red Hat Enterprise Linux [99](#)
 - SUSE Linux Enterprise Server [99](#)
 - discovering LUNs
 - rescan script [57](#)
 - SCSI rescan [57](#)
 - `sg3_utils` [57](#)
 - discovering LUNs, verification
 - `iscsi-ls` [59](#)
 - `sanlun` [59](#)
- ## DM-Multipath
- `/etc/multipath.conf` [41](#)
 - Citrix Xen Server s sample multipath-conf file [144](#)
 - configuring ALUA [41](#)
 - displaying information using `sanlun lun show -p` [66](#), [67](#)
 - editing configuration file [41](#)
 - I/O failover [99](#)
 - Oracle VM s sample multipath-conf file [131](#)
 - path load sharing [99](#)
 - Red Hat Enterprise Linux 4 sample multipath-conf file [124](#)
 - Red Hat Enterprise Linux 5 sample multipath-conf file [119](#)
 - Red Hat Enterprise Linux 6 sample multipath-conf file [111](#)
 - Red Hat Enterprise Virtualization Hypervisor sample multipath-conf file [127](#)
 - SCSI device [99](#)
 - SUSE Linux Enterprise Server 10 sample multipath-conf file [141](#)
 - SUSE Linux Enterprise Server 11 sample multipath-conf files [137](#)
 - `user_friendly_names` parameter and SAN boot LUNs [119](#)
 - verifying [40](#)
 - viewing with `multipath -ll` command [45](#)
- ## DM-Multipath, configuration, command
- `# multipath` [43](#)
- ## E
- environment
 - finding instructions [15](#)
 - environments
 - specifying [15](#)
 - Examples
 - `sanlun FC`, iSCSI output [67](#)
- ## F
- FC protocol
 - discovering LUNS [57](#)
 - host bus adapters

HBA port [97](#)
 sanlun output [67](#)

FCoE

converged network adapters [98](#)
 data center bridging [98](#)
 Ethernet switch [98](#)
 traditional FC [98](#)

H

HBA

general notes [20](#)
 getting WWPN [20](#)

HBA driver

bundled with Linux distribution [20](#)
 downloaded from OEM [20](#)

Host Utilities

contents [12](#)
 defined [12](#)
 drivers [20](#)
 HBA [20](#)
 host configuration [20](#)
 installing [24](#)
 Linux environments [13](#)
 planning installation [17](#)

Hypervisor

align partitions for best performance [96](#)

I

initiator group

igroup [97](#)

initiator node name

See node name

installation

overview [17](#)
 planning [17](#)

installing

Host Utilities software [24](#)

iSCSI

setting multipathing timeout values [29](#)

iSCSI protocol

automatic node login [36](#)
 changing initiator node name [27](#)
 configuring a SAN boot LUN [87](#)
 configuring Broadcom [21](#)
 configuring Red Hat SAN boot LUN [87](#)
 configuring software iSCSI SAN boot LUN [88](#), [89](#)
 discovering LUNs in hardware iSCSI environments [57](#)

HBA [98](#)

initiator node name [27](#)

iSCSI target HBA [98](#)

manual node login [36](#)

methods of target discovery with software initiators [33](#)

preboot execution environment (PXE) [87](#)

sanlun output [67](#)

setting multipathing timeout values [29](#)

standard Ethernet interfaces [98](#)

starting automatically [36](#)

starting the iSCSI service [33](#)

iSCSI protocol, configure SAN boot

network interface card (NIC) [87](#)

software iSCSI

locally attached disk [87](#)

preboot execution environment (PXE) [87](#)

K

Kernel-based Virtual Machine (KVM)

supported in Linux Host Utilities [95](#)

L

Linux

Host Utilities [12](#)

Linux configurations

ALUA support

automatically enabled [100](#)

asymmetric logical unit access

Target Port Group Support [100](#)

Linux Host Utilities

direct-attached configurations

high-availability [97](#)

single-zoned [97](#)

stand-alone [97](#)

fabric-attached SAN

dual-port [97](#)

single-port [97](#)

Linux versions

Oracle Linux [13](#)

Red Hat Enterprise Linux [13](#)

SUSE Linux Enterprise Server [13](#)

Veritas Storage Foundation [13](#)

LUN retry tunable values

configuring for Veritas Storage Foundation [50](#)

LUNs

discovering in FC environments [57](#)

discovering in hardware iSCSI environments [57](#)

- discovering in iSCSI environments [58](#)
- discovering in Veritas, iSCSI environments [74](#)
- iSCSI protocol
 - discovering LUNs [58](#)
- native Linux environments [57](#)
- removing [71](#)
- removing in Veritas environments [81](#)
- viewing [59](#)

M

- man pages
 - installed by Host Utilities [12](#)
- misalignment
 - VHD partition [96](#)
- mounting, new file
 - device mount_point type _netdev 0 0 [59](#)
- multipath -ll command [45](#)
- multipath.conf
 - Citrix Xen Server sample file [144](#)
 - Oracle VM sample file [131](#)
 - Red Hat Enterprise Linux 4 sample file [124](#)
 - Red Hat Enterprise Linux 5 sample file [119](#)
 - Red Hat Enterprise Linux 6 sample file [111](#)
 - Red Hat Enterprise Virtualization Hypervisor sample file [127](#)
 - SUSE Linux Enterprise Server 10 sample file [141](#)
 - SUSE Linux Enterprise Server 11 sample files [137](#)
 - user_friendly_names parameter and SAN boot LUNs [119](#)
- multipathing
 - displaying with sanlun lun show -p [80](#)
 - setting iSCSI timeout values [29](#)
- multipathing packages
 - DM-Multipath [40](#)
 - verifying [40](#)

N

- node name
 - changing [27](#)
 - recording [27](#)
- nSANity
 - diagnostic utility [12](#)
 - installing [109](#)

O

- O2CB_HEARTBEAT_THRESHOLD
 - set to 65 [26](#)

- Oracle Linux
 - Host Utilities environment [13](#)
- Oracle Linux, Sample configuration [126](#)
- Oracle VM
 - sample multipath-conf file [131](#)
 - SAN booting not supported [131](#)
 - setting O2CB_HEARTBEAT_THRESHOLD to 65 [26](#)
 - setting up [26](#)
 - supported in Linux Host Utilities [95](#)

P

- performance
 - align VHD partitions [96](#)
- protocols
 - Fibre Channel (FC) [97](#)
 - Fibre Channel over Ethernet (FCoE) [97](#)
 - iSCSI [97](#)

Q

- qualifiers
 - specifying environments [15](#)

R

- Red Hat Enterprise Linux
 - configuring Red Hat Linux 6, 5 to work with Veritas Storage Foundation [51](#)
 - Host Utilities environment [13](#)
 - setting iSCSI timeout values [29](#)
 - setting the IOFENCE timeout value [51](#)
 - setting up space reclamation [25](#)
- Red Hat Enterprise Linux 4
 - sample multipath-conf file [124](#)
- Red Hat Enterprise Linux 5
 - sample multipath-conf file [119](#)
- Red Hat Enterprise Linux 6
 - sample multipath-conf file [111](#)
- Red Hat Enterprise Virtualization Hypervisor
 - sample multipath-conf file [127](#)
- reloading iSCSI, command
 - /etc/init.d/iscsi reload [59](#)
- requirements
 - configurations
 - more information [149](#)
 - more information [149](#)
- rescan script
 - discovering LUNs [73](#), [74](#)

- LUNs
 - discovering in Veritas environments [73](#)
 - removing LUNs [71](#)

S

- SAN boot
 - centralized storage [99](#)
 - disaster recovery [99](#)
 - maintenance [99](#)
 - not supported on Citrix Xen Server [144](#)
 - not supported on Oracle VM [131](#)
 - reliability [99](#)
 - server swaps [99](#)
 - user_friendly_names parameter and SAN boot LUNs [119](#)
- SAN boot LUN
 - configuring iSCSI hardware initiator [87](#)
 - configuring software iSCSI on Red Hat [88](#)
 - configuring software iSCSI on SUSE Linux [89](#)
 - configuring with iSCSI [87](#)
- SAN boot, FC
 - BIOS [84](#)
 - configuring root partition, SUSE Linux Enterprise Server [85](#)
 - DM-Multipath [85](#)
 - HBA BIOS [85](#)
 - initird [85](#)
 - installing, Red Hat Enterprise Linux [84](#)
 - installing, SUSE Linux Enterprise Server [85](#)
 - linux mpath [84](#)
 - SAN boot LUN [84](#), [85](#)
 - spinup delay [85](#)
- SAN boot, hardware iSCSI
 - Boot Option [87](#)
 - linux mpath [87](#)
 - Red Hat Enterprise Linux, SAN boot [87](#)
 - SAN boot LUN [87](#)
 - spinup delay [87](#)
- SAN boot, root partition
 - iSCSI sessions [91](#)
 - mount [91](#)
 - Start-Up mode [91](#)
 - SUSE Linux Enterprise Server, SAN boot [91](#)
- SAN boot, software iSCSI
 - /boot [89](#)
 - Boot Option [88](#)
 - initiator name [89](#)
 - IQN [89](#)
 - iSCSI initiator name [88](#)
 - linux mpath [88](#)
 - Red Hat Enterprise Linux, SAN boot [88](#)
 - root file system [88](#)
 - SUSE Linux Enterprise Server, SAN boot [89](#)
 - SWAP partition [88](#), [89](#)
 - target IP address [88](#)
- san_version
 - installed by Host Utilities [12](#)
- sanlun
 - displaying Veritas DMP multipathing information [80](#)
 - examples of FC, iSCSI output [67](#)
 - iscsi
 - viewing LUNs [59](#)
 - iscsiadm
 - viewing LUNs [59](#)
 - sanlun lun show -p output [67](#)
 - sanlun lun show all output [67](#)
 - viewing LUNs [59](#)
- sanlun utility
 - installed by Host Utilities [12](#)
- server virtualization
 - KVM [95](#)
 - Oracle VM [95](#)
 - supported in Linux Host Utilities [95](#)
- space reclamation
 - setting up [25](#)
- starting DM-Multipath, command
 - /etc/init.d/boot.multipath start [43](#)
 - /etc/init.d/multipathd start [43](#)
- stopping DM-Multipath, command
 - /etc/init.d/boot.multipath stop [48](#)
 - /etc/init.d/multipathd stop [48](#)
- storage system
 - direct-attached [97](#)
 - Ethernet network [97](#)
 - fabric-attached [97](#)
- SUSE Linux Enterprise Server
 - configuring SUSE Linux 11 to work with Veritas Storage Foundation [52](#)
 - Host Utilities environment [13](#)
 - setting iSCSI timeout values [29](#)
- SUSE Linux Enterprise Server 10
 - sample multipath-conf files [141](#)
- SUSE Linux Enterprise Server 11 sample multipath-conf file [137](#)

T

- timeout values
 - iSCSI multipathing timeout values [29](#)

V

Veritas

Dynamic Multipathing

VxDMP [99](#)

using VxDMP [73](#)

using VxVM [73](#)

Veritas Volume Manager

VxVM [99](#)

Veritas DMP

sanlun lun show -p [80](#)

Veritas Dynamic Multipathing

VxDMP [99](#)

Veritas restore daemon

configuring [50](#)

Veritas Storage Foundation

Array Policy Module, [54](#)

Array Support Library [54](#)

boot BIOS [93](#)

configure Red Hat 6, 5 [51](#)

configure SAN boot, Red Hat Enterprise Linux [93](#)

configuring LUN retry tunable values [50](#)

configuring the restore daemon [50](#)

HBA BIOS [93](#)

HBA port [93](#)

Host Utilities environment [13](#)

SUSE Linux 11 must be configured [52](#)

viewing a DM-Multipath configuration [45](#)

viewing LUNs, command

sanlun lun show all [75](#)

virtual hard disk (VHD)

align partitions for best performance [96](#)

VxDMP

prerequisites [73](#)

VxVM

managing LUNs [73](#)

VxVM, command

vxdisk list [82](#)

vxdkmpadm getsubpaths ctrl [82](#)

vxdkmpadm getsubpaths dmpnodename [82](#)



NA 210-06244_A0, Printed in USA

GC52-1337-07

