

iSCSI Red Hat® Enterprise Linux® Host Utilities 3.2.1

Release Notes

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

Part number: 215-04332_A0
October 2008

Copyright and trademark information

Copyright information

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

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

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

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

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

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

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

Trademark information

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

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

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

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

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

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

Table of Contents

Changes in this release 1

Contents of the Host Utilities. 2

Known problems and limitations using RHEL 5 series 4

Known problems and limitations using RHEL 4. 6

Known problems and limitations using RHEL 3 12

Where to find additional information. 18

Changes in this release

New Feature in the Release

The iSCSI Red Hat Enterprise Linux Host Utilities 3.2.1 Release Notes update of October 2008:

- ◆ Support for iSCSI Red Hat Enterprise Linux SAN boot feature.

Linux Host Utilities kit

Enhanced `linux_info` script

Documentation changes

Updated previous documentation related issues.

Contents of the Host Utilities

Software

The Host Utilities includes a set of software programs to help you use your Linux host with NetApp storage systems. You can download this software from the NOW Web site after accepting the NetApp Customer Software License.

The software is distributed as a single file (netapp_linux_host_utilities_3_2.tar.gz), which you install on your Linux host. The installation instructions are provided in the *Setup Guide* and on the NOW download page.

The same software distribution is shared by a number of Host Utilities. Therefore, it contains some software that is not used for Linux iSCSI. The following table lists each file in the distribution:

File	Purpose
install	Installs the Host Utilities software on a Linux host.
uninstall	Removes the Host Utilities software from the Linux host.
sanlun_bin.tar	Contains the files below.
sanlun	Displays information about LUNs on connected NetApp storage systems.
filer_info	Collects information about the storage system in a file that can be submitted to NetApp for problem analysis.
linux_info	Collects information about the Linux host in a file that can be submitted to NetApp for problem analysis.
san_version	Displays the version of the software tools installed for this Host Utilities. The version of the software tools for this Host Utilities is 3.2.1
brocade_info	Not used for iSCSI.
cisco_info	Not used for iSCSI.
mcdata_info	Not used for iSCSI.
Telnet.pm	Support library for *info scripts.

File	Purpose
SHsupport.pm	Support library for *info scripts.
libqlsdm.so	Support library for sanlun.

Documentation

The Host Utilities includes the following documentation.

Document	Purpose
Release Notes	Lists the contents of the Host Utilities, and any known problems or last minute updates. Download from the NOW Web site.
Setup Guide	Explains how to install and configure software to enable the Linux host to use LUNs on the storage system. Download from the NOW Web site.
Command Quick Reference	Lists the most common storage system and Linux commands used to manage iSCSI. Download from the NOW Web site.
man pages	Displays information about the Host Utilities software. The man pages for the <code>sanlun</code> , <code>filer_info</code> , and <code>linux_info</code> commands are copied to the <code>/opt/netapp/man/man1</code> directory when you install the Host Utilities software. Included in Host Utilities software distribution.
NOW pages	Host Utilities installation instructions, information updates, links to additional information. Browse on NOW Web site.

Known problems and limitations using RHEL 5 series

GFS growth is not possible on mounted file system on an iscsi LUN

Problem: GFS provides the online grow feature, where GFS file system can be extended without unmounting the file system, using the `gfs_grow` command. This feature cannot be used on a file system created over an iSCSI and multipath device on a NetApp LUN.

Note

When GFS is used on LVM volume, it can be extended without unmounting the file system by adding new LUNS to the volume group, and expanding the logical volume on which GFS is created.

Though NetApp filers provide dynamic resizing of the LUN, the iSCSI layer in Linux is not capable of detecting the LUN resize. The new size of the LUN gets reflected on Linux only after iSCSI service is restarted on Linux.

Workaround: The workaround is to perform the following actions on all nodes:

Step	Action
1	Unmount the file system
2	Resize the LUN on the filer
3	Run the following command: <code>multipath -F</code>
4	Run the following command: <code>service multipath stop</code>
5	Run the following command: <code>service iscsi restart</code>
6	Run the following command: <code>service multipath start</code>
7	Mount the file system

Step	Action
8	Run the following command on one of the nodes: <code>gfs_grow</code>
9	Run the following command on all nodes to verify the new size: <code>running gfs_grow -Tv mountpoint</code>

For more information, see bug 215923 on Bugs Online at

<http://now.netapp.com/NOW/cgi-bin/bol>

Using partitions on Linux with Data ONTAP LUNs may require alignment for best performance

Problem: Using default partition settings on LUNs may result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

Workaround: Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. Refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for specific steps.

iscsi package is not automatically selected

Problem: Installer fails to select iscsi package automatically during SAN boot installation on multipath device.

Workaround: Select iscsi package manually during installation.

Multipathed device is not listed during installation

Problem: During installation multipathed device is not listed if the LUN is discovered through only one path.

Workaround: During mpath installation you must discover the LUN through more than one path.

iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use the `iscsi nodename` to change the default node name of the storage system, ensure the new node name does not exceed 255 characters

Known problems and limitations using RHEL 4

GFS growth is not possible on mounted file system on an iscsi LUN

Problem: GFS provides the online grow feature, where GFS file system can be extended without unmounting the file system, using the `gfs_grow` command. This feature cannot be used on a file system created over an iSCSI and multipath device on a NetApp LUN.

Note_____

When GFS is used on LVM volume, it can be extended without unmounting the file system by adding new LUNS to the volume group, and expanding the logical volume on which GFS is created.

Though NetApp filers provide dynamic resizing of the LUN, the iSCSI layer in Linux is not capable of detecting the LUN resize. The new size of the LUN gets reflected on Linux only after iSCSI service is restarted on Linux.

Workaround: The workaround is to perform the following actions on all nodes:

Step	Action
1	Unmount the file system
2	Resize the LUN on the filer
3	Run the following command: <code>multipath -F</code>
4	Run the following command: <code>service multipath stop</code>
5	Run the following command: <code>service iscsi restart</code>
6	Run the following command: <code>service multipath start</code>
7	Mount the file system

Step	Action
8	Run the following command on one of the nodes: <code>gfs_grow</code>
9	Run the following command on all nodes to verify the new size: <code>running gfs_grow -Tv mountpoint</code>

For more information, see bug 215923 on Bugs Online at

<http://now.netapp.com/NOW/cgi-bin/bol>

Service clvmd restart fails with connect fail on local socket

Problem: If GFS is used with DLM lock, it uses the underlying cluster lvm layer for providing cluster volume manager. In rare cases, volume group does not come up on boot. This might happen if iSCSI LUNs are not available before lvm startup script comes up.

In such cases, you might have to run `vgscan` after bootup to read the volume groups from the iSCSI LUNs. This might also need a `clvmd` restart. If we run `service clvmd restart` to restart `clvmd`, in GFS 4.0, it fails. This is due to timing issue in `clvmd` service start script.

Workaround: The workaround to avoid the `clvmd` restart failure is to stop and start the `clvmd`. To do that, run:

Step	Action
1	Run the following command to stop <code>clvmd</code> : <code>service clvmd stop</code>
2	Run the following command to start <code>clvmd</code> : <code>service clvmd start</code>

OCFS2 self fencing during cluster failover on RHEL 4 Update 3

Problem: During a clustered failover or network disruption, nodes fence themselves by panicking.

Oracle Cluster File System (OCFS2) has a fencing mechanism that relies on each node being able to write a heartbeat value to a specific area on each OCFS2 volume. Every two seconds, each node writes a new heartbeat value, and reads the value of the other node. If a node is unable to write or read the heartbeat value, it retries the operation a specific number of times, and then self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage can become unavailable, and none of the nodes in the OCFS2 cluster will be able to write or read the heartbeat values. As a result, the nodes may fence themselves by panicking.

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

NetApp recommends that, with a dedicated TCP/IP network for iSCSI link between Linux hosts and NetApp filer, a reasonable value in most cases is 61. This corresponds to a total of 120 seconds before a node gives up, and self-fences.

Depending on the load in the iSCSI network used between the Linux host and NetApp filer, the value should be raised. Netapp recommends that a value in the range of 61 to 201 would be ideal for a configuration with OCFS2 on mutipath and iSCSI devices.

For more information, see:

http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html#HEARTBEAT

System hangs when shutting down or rebooting

Problem: The Linux system hangs when shutting down or rebooting. The `/etc/init.d/iscsi stop` command may also report errors. This occurs only in systems running dm-multipath with iSCSI.

Workaround: This is a known problem with the shutdown sequence in RHEL 4 Update 3. If the system is currently hung, disconnect power to force the shutdown to complete. To prevent shutdown problems, complete the following steps before shutting down or rebooting:

1. Unmount all file systems on iSCSI LUNs.
2. Flush the multipath maps using the `multipath -F` command.
3. Stop the iSCSI service using the `/etc/init.d/iscsi stop` command.
4. Enter the shutdown or reboot command.

For more information, see Red Hat Bugzilla 177355 at <https://bugzilla.redhat.com/bugzilla/> or NetApp bug 183045 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol>.

Multiple devices shown for each LUN

Problem: The `iscsi-ls -l` and `sanlun lun show` commands display multiple devices for each LUN that is mapped to the Linux host. This problem occurs only with RHEL 4 Update 2 and RHEL 4 Update 3 without dm-multipath.

Workaround: The iSCSI initiator creates a separate iSCSI session to each iSCSI portal discovered on the storage system. If you are running RHEL 4 Update 3, you can use the dm-multipath support to manage multiple paths.

Otherwise, configure the storage system or network routing so that only one portal (interface) is visible to the Linux host. See the Host Utilities *Setup Guide* for more information.

Always map first LUN as LUN 0

Problem: If the first LUN mapped to a host has an ID other than 0, the `iscsi-ls -l` command shows incorrect information. Adding LUN 0 later does not solve the problem.

Workaround: Always map the first LUN for a host to LUN 0. This is the default behavior of Data ONTAP when you do not explicitly specify the LUN identifier.

For more information, see NetApp bug 164689 at: <http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=164689>.

Using partitions on Linux with Data ONTAP LUNs may require alignment for best performance

Problem: Using default partition settings on LUNs may result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

Workaround: Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. Refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for specific steps.

iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use the `iscsi nodename` to change the default node name of the storage system, ensure the new node name does not exceed 255 characters.

LUNs not recognized by dm-multipath

Problem: LUNs are mapped to the Linux host and are correctly reported by the `iscsi-ls -l` and `sanlun lun show all` commands. However the LUNs are not reported by the `multipath -ll` command. Multipath devices are not created for the LUNs.

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

```
multipath -v -d | grep blacklist
```

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

If the iSCSI devices are not blacklisted, but are still not recognized by the `multipath` command, regenerate the multipath maps by entering the following command:

```
multipath -v3
```

For more information, see bug 201302 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol>

Storage system console displays continuous ISID Rule messages

Problem: The storage system console displays continuous ISID Rule messages, similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1
```

```
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
```

```
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10
```

```
Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

Workaround: Add the following line to the `/etc/iscsi.conf` file. Make sure this line is at the beginning of the file, before any “DiscoveryAddress” entries.

```
Continuous=no
```

Then restart the iSCSI service by entering the following command:

```
/etc/init.d/iscsi restart
```

Known problems and limitations using RHEL 3

Using partitions on Linux with Data ONTAP LUNs may require alignment for best performance

Problem: Using default partition settings on LUNs may result in performance problems. This problem can occur if the `sectors` value for a partition is not a multiple of 8.

Workaround: Manually enter the disk geometry parameters when creating partitions with the `fdisk` command. Refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for specific steps.

Host system performance problems under specific configurations

NetApp testing has found problems with the Linux iSCSI software initiator running under Red Hat Advanced Server (AS) 3.0. The virtual memory (vm) subsystem of the Red Hat AS 3.0 Linux kernels does not interact well with the iSCSI software initiator under certain circumstances. NetApp testing has shown that certain memory and workload configurations may lead to Linux server pauses, and possibly lock up the Linux server completely.

For more information and workarounds, see NetApp bug 149391 at: <http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=149391>

File systems on iSCSI LUNs do not automount after reboot

Problem: Red Hat Linux 3 update 4 does not have a standard mechanism for automounting file systems based on iSCSI LUNs after a reboot.

iSCSI LUNs depend on networking services that are typically not available when the file system mounting takes place. Also, the file systems must be mounted by a persistent mechanism; iSCSI devices are assigned device names (`/dev/sda`, `/dev/sdb`...) based on when they are discovered. These device names cannot be guaranteed to map to the same LUN each time.

Workaround: You can work around this limitation by manually issuing the `mount -a` command after networking services become available. You could also create a script that waits for networking services to be available before issuing the `mount -a` command.

For more information, refer to Red Hat Bugzilla 143151 at: https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=143151

Unmount iSCSI LUNs before Linux shutdown

Problem: File systems on iSCSI LUNs are not automatically unmounted when the Linux system is shut down.

The `iscsi-umountall` script used in earlier versions of the iSCSI initiator is not available with Red Hat 3.0 Update 4. There is no default mechanism for automatically unmounting file systems on iSCSI LUNs when the Linux system is shut down. Shutting down Linux with the file systems still mounted can cause a panic of the Linux host.

Workaround: Issue the `umount` command for each file system on an iSCSI LUN before shutting down the Linux host.

For more information, refer to Red Hat Bugzilla 144781 at: https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=144781

Error message: iSCSI: ioctl size 2596 incorrect, expecting 2316

Problem: The error message `iSCSI: ioctl size 2596 incorrect, expecting 2316` is displayed on the Linux host.

Workaround: Remove the SourceForge iSCSI initiator that was installed for an earlier release of Red Hat Linux. The iSCSI initiator is included in the Red Hat Linux version 3 update 4 and version 4 update 2 kernel and distribution. Do not use download an iSCSI initiator from SourceForge.

See the Host Utilities *Setup Guide* for more information.

sanlun does not display all iSCSI LUNs

Problem: The `sanlun` command does not display all iSCSI LUNs mapped to the initiator.

There are two possible explanations for this problem:

- ◆ Number of SCSI devices

The `sanlun` command checks the number of `/dev/sg*` entries on the host. The default number of `/dev/sg*` entries is 32. If you have more than 32 SCSI devices (for example, two physical SCSI hard drives, one SCSI CD-ROM, and 30 LUNs), `sanlun` does not list all the LUNs.

- ◆ Filer is unreachable for a long time

If the time during which the filer is unreachable exceeds the timeouts in the `/etc/iscsi.conf` file, the `sanlun lun show all` command might not provide a complete list of mapped LUNs the first time you run the command.

Workaround: The following table lists the workarounds for this problem.

If...	Then...
You have more than 32 SCSI devices	Create more <code>/dev/sg*</code> entries using the Linux <code>mknod</code> command. For details about the <code>mknod</code> command, see the Linux man page. Example: <code>mknod --mode=660 /dev/sg32 c 21 32</code>
The filer was unreachable	1. Run the <code>sanlun lun show all</code> command a second time. 2. Make sure that the LUNs listed match the devices listed in output of the <code>iscsi-ls -l</code> command.

iscsi-ls does not display all iSCSI LUNs

Problem: The `iscsi-ls -l` command lists detailed information about iSCSI LUNs that are mapped to `/dev/sda` through `/dev/sdz`. LUNs mapped to device names starting with `/dev/sdaa` are listed, but most information (including the device name) is missing.

This is a known bug in the `iscsi-ls` command code.

Workaround: Use the `sanlun` command to view information about LUNs mapped to iSCSI devices. For example: `sanlun lun show all`.

For more information, refer to Red Hat Bugzilla 145141 at: https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=145141

Version numbers of iSCSI kernel module and tools do not match

Problem: The iSCSI initiator kernel module displays as 3.6.1, while the version number of the iSCSI daemon and userspace tools displays as 3.6.2.

The daemon and userspace tools were changed late in the release cycle and the kernel module version was not updated to match. The kernel module version 3.6.1 is identical to kernel module 3.6.2.

This is not a problem and no action is needed.

File system errors and system shutdown hangs

Problem: If the initiator is accessing LUNs on a filer that is unreachable (for example, during a cluster takeover or giveback), file system errors or system shutdown hangs might occur.

File system errors and system shutdown hangs depend on disk command timeout and connection failure timeout settings in the `/etc/iscsi.conf` file.

By default, the disk command and connection failure timeouts are not listed in the `/etc/iscsi.conf` file. As a result, these timeout values are infinite and the host infinitely retries these I/O commands. The infinite timeouts prevent file system errors, but stopping the iSCSI service or a shutdown hangs if the filer becomes unreachable.

If you did not create Linux file systems on your LUNs, and are using raw devices instead, you can set the timeouts to a value of 180 seconds to prevent system shutdown hangs. However, any application using these LUNs might receive SCSI I/O errors and file system errors if a filer is unreachable for a long time.

To set the I/O command timeout values, complete the following steps:

Caution _____
Do not change the timeout values in the following section if you are using any Linux file systems on your storage system LUNs. This includes ext2 and ext3 file systems.

Step	Action
1	Ensure that you are logged in to the host with root privileges.
2	With a text editor, open the host's <code>/etc/iscsi.conf</code> file.
3	At the beginning the file, add the following lines. Make sure that you add these lines before the filer's <code>DiscoveryAddress</code> settings: <code>MaxDiskCommandTimeout=180</code> <code>ConnFailTimeout=180</code>

iSCSI node names

The iSCSI Linux initiator does not support node names with more than 255 characters.

If you use the `iscsi nodename` to change the default node name of the storage system, ensure the new node name does not exceed 255 characters.

Storage system console displays continuous ISID Rule messages

Problem: The storage system console displays continuous ISID Rule messages, similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0

Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1

Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10

Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0

Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator ign.1987-05.com.cisco:rachman at IP addr
10.10.10.10

Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

Workaround: Add the following line to the `/etc/iscsi.conf` file. Make sure this line is at the beginning of the file, before any “DiscoveryAddress” entries.

```
Continuous=no
```

Then restart the iSCSI service by entering the following command:

```
/etc/init.d/iscsi restart
```

For information about additional problems

Occasionally, additional problems are reported with the Host Utilities, or known bugs are fixed. The problems are listed in an online system called Bugs Online.

For a list of known and fixed bugs, access Bugs Online by completing the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com .
2	Click Service & Support.
3	Click Bugs Online & Release Tools.

Step	Action
4	In the Bug Search Tools section, choose iSCSI-Linux from the Bug Type(s) menu.
5	Click Go. Result: The Bugs Online - Advanced Search Results lists all bugs for Linux iSCSI initiators.

Where to find additional information

Where to go for more information

The following table describes documents with the most current information about host and storage system requirements and additional documentation.

If you want...	Go to...
The most current system requirements for the Linux host	<i>iSCSI Host Compatibility Matrix</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml
The supported storage system models for Data ONTAP licensed with iSCSI	◆ <i>FCP/iSCSI Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.cgi ◆ <i>System Configuration Guide</i> for Data ONTAP at http://now.netapp.com
Information about how to configure and manage the iSCSI service on the storage system	The following documents for your Data ONTAP release: ◆ <i>Data ONTAP Block Access Management Guide for iSCSI</i> ◆ <i>Data ONTAP Release Notes</i>
Best practices or configuration issues	http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190
Information about dm-multipath	The Multipath example files included with the device-mapper-multipath RPM are in /usr/share/doc/device-mapper-multipath-<version> where <version> is the latest version number shipped with the RHEL release.
Information on Red Hat Cluster and GFS	http://www.redhat.com/docs/manuals/csgfs/
Information on OCFS2	◆ Oracle Cluster File System (OCFS2) User's Guide: http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf ◆ OCFS2 - Frequently Asked Questions: http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html (The FAQs on OCFS2)

If you want...	Go to...
Information on Oracle RAC	http://www.oracle.com/technology/pub/articles/smiley_rac10g_install.html
NetApp Knowledge Base for Oracle RAC	http://www.netapp.com/library/tr/3339.pdf