

iSCSI SUSE® Linux® Host Utilities 3.0

Release Notes

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

Part number 215-02058_A0
October 2006

Copyright and trademark information

Copyright information

Copyright © 1994—2006 Network Appliance, Inc. All rights reserved.

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

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

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

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

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

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

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, FilerView, Data ONTAP, MultiStore, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. Network Appliance, and Snapshot are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. FlexClone, FlexVol, NOW, NetApp on the Web, RAID-DP, SecureAdmin, SnapDrive, vFiler, VFM, and Virtual File Manager, are trademarks of Network Appliance, Inc. in the United States and other countries.

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

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

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

Changes in this release 1

iSCSI initiator versions 2

Contents of the Host Utilities. 3

Upgrading from SLES 8 SP3 to SLES 8 SP4 5

Upgrading from SLES 9 to SLES 9 SP1, SP2, or SP3 6

Known problems and limitations 7

Where to find additional information. 14

Changes in this release

Linux Clustering support

iSCSI SUSE® Linux Host Utilities version 3.0 supports Cluster file systems and Clustering suite on Linux hosts.

The usage of cluster suite as framework for the file system, and the usage of iSCSI devices as shared quorum disk(s) by the cluster suites are supported.

The File system supported is:

- ◆ Oracle Cluster File System (OCFS2) 1.2.1 (o2cb the cluster service as part of OCFS2)

The Cluster suite supported is:

- ◆ Oracle Clusterware and Oracle Real Application Cluster (RAC) 10g R2

Documentation changes

There are no changes to the diagnostic tools and scripts from the iSCSI SUSE® Linux Initiator Support Kit version 1.4. This Host Utilities version has only documentation changes.

The documentation changes include:

- ◆ Updates related to Linux Clustering support, and Troubleshooting information in the *Setup Guide*.
- ◆ Updates to Known Problems section of the *Release Notes*.

iSCSI initiator versions

Supported initiator versions

Each version of SUSE Linux includes a different version of the iSCSI initiator. The *Setup Guide* for this Host Utilities version instructs you to verify that you have the correct initiator version installed.

The following table lists each supported version of SUSE Linux and the required initiator version (linux-iscsi userspace RPM). Note that the initiator is included in the SUSE Linux distribution. If you are missing the initiator RPM, be sure to get it from your SUSE Linux distribution.

Note

Do not download an open source initiator and attempt to compile and install it for your SLES host.

Linux Version	Initiator Version	Notes
SLES 8 SP3	3.4.0.3- <i>n</i>	<i>n</i> indicates the platform type. For example, 95 indicates that the system is a dual X86 platform.
SLES 8 SP4	3.6.2-2	
SLES 9	4.0.1-88.13	Reiser FS not supported for iSCSI LUNs.
SLES 9 SP1	4.0.1-88.18	
SLES 9 SP2	4.0.1-88.18	
SLES 9 SP3	4.0.1-88.26	

Supported multipathing RPM versions

If you are using the optional dm-multipath support, you also need the following userspace RPM versions:

- ◆ device-mapper RPM: 1.01.01-1.6
- ◆ multipath RPM: 0.4.5-0.11

Contents of the Host Utilities

Software

The Host Utilities include a set of software programs to help you use your SUSE 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_tools_3_0.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 (formerly, Support Kits). Therefore, it contains some software that is not used for SUSE Linux. 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 software. The version of the software tools for this Host Utilities is 3.0.
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.
SHsupport.pm	Support library for *info scripts.

File	Purpose
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.

Upgrading from SLES 8 SP3 to SLES 8 SP4

Overview of upgrade procedures

If you are running SUSE Linux Enterprise Server 8 (SLES 8) SP3 and are upgrading to SLES 8 SP4, unmount iSCSI LUNs and stop the iSCSI service before you upgrade.

For general information on SLES 8 SP4, see <http://www.novell.com/products/linuxenterpriseserver8/sp4.html>.

Before upgrading to SLES 8 SP4

If you are using iSCSI LUNs with SLES 8 SP3, complete the following steps before upgrading to SLES 8 SP4:

Step	Action
1	Unmount all existing iSCSI LUNs by running the following command on the SLES host: <code>/sbin/iscsi-umountall</code>
2	Stop the iSCSI service on the SLES host: <code>/etc/init.d/iscsi stop</code>
3	Continue with the upgrade to SLES 8 SP4.
4	Create file system labels and add the labels to <code>/etc/fstab.iscsi</code> . See Chapter 4, “Accessing LUNs from the host,” in the <i>Setup Guide</i> for more information.

Upgrading from SLES 9 to SLES 9 SP1, SP2, or SP3

Overview of upgrade procedures

If you are running SUSE Linux Enterprise Server 9 (SLES9) and are upgrading to SLES 9 SP1, SP2, or SP3, unmount iSCSI LUNs and stop the iSCSI service before you upgrade.

For general information on SLES 9 SP1, see <http://www.novell.com/products/server9/sp1.html>

For general information on SLES 9 SP2, see <http://www.novell.com/products/server9/sp1.html>

For general information on SLES 9 SP3, see <http://www.novell.com/products/server9/sp3.html>

Before upgrading to SLES 9 SP1, SP2, or SP3

If you are using iSCSI LUNs with SLES9, complete the following steps before upgrading to SLES9 SP1, SP2, or SP3:

Step	Action
1	Unmount all existing iSCSI LUNs by running the following command on the SLES host: <code>/sbin/iscsi-umountall</code>
2	Stop the iSCSI service on the SLES host: <code>/etc/init.d/iscsi stop</code>
3	Continue with the upgrade to SLES9 SP1, SP2, or SP3.

Known problems and limitations

OCFS2 self fencing during cluster failover on SLES 9 SP3

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

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

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

See also, bug 210340 on Bugs Online at:

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

Automounting and autounmounting of file system on multipath device do not work on SLES 9 SP3

Problem: Automated mounting of multipath-based file systems at boot time using entries in `/etc/fstab.conf` or `/etc/fstab.iscsi` does not work. This is because, the placement of multipath bootup and shutdown scripts in the rc script list in SLES9 SP3 is out of order with respect to the mount and unmount scripts for iSCSI and non-iSCSI devices.

multipathd and iscsid startup scripts are at the same serial number for a run level in `/etc/init.d/rc3.d`. Since alphabetically *i* comes before *m*, iscsid starts before multipathd.

As a result, the multipath bootup and shutdown sequence is not synchronized with iSCSI bootup and shutdown sequence. Multipath devices will not be available when the mount command is run, and the mount command fails to mount these devices.

Similarly in the shutdown path, by the time the unmount script runs, the multipath utility would have already been shutdown.

Workaround: Manual mounting and unmounting is required for multipath-based file systems on SLES9 SP3. After multipath devices are available on bootup, the file systems can be mounted.

For unmounting, make sure that any multipath-based file systems are unmounted before the multipath utility is shutdown.

The following sequence of actions should be performed to ensure proper shutdown of multipath-based iSCSI file systems:

Step	Action
1	Unmount all file systems on iSCSI LUNs
2	Flush the multipath maps by entering the following command: <code>multipath -F</code>
3	Stop the multipath daemon by entering the following command: <code>/etc/init.d/multipathd stop</code>
4	Stop the iSCSI service by entering the following command: <code>/etc/init.d/iscsi stop</code>

For more information, see bug 196997 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol>

LUNs not recognized by dm-multipath

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

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

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

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

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

```
multipath -v3
```

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

SLES 8 does not generate a unique initiator node name

Problem: On i386 platforms, the SLES 8 software initiator does not generate a unique iSCSI initiator node name.

Workaround: If you have more than one SLES 8 host in your network, you must change the initiator node name of each system to a unique value. See the *Setup Guide* for the procedure for changing the node name.

File systems on partitions may have performance problems

Problem: The default `fdisk` command parameters used to create a partition on a LUN may result in poor performance for that LUN.

The default `fdisk` parameters can create a misalignment between the LUN block boundaries and the underlying WAFL® boundaries. This causes the storage system to double its I/O workload for that LUN.

Workaround: You can explicitly specify the `fdisk` parameters when creating a partition to use values that guarantee the correct block alignment. Refer to bug number 156121 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for specific steps.

iscsi-ls command displays error message

Problem: For some SUSE Linux versions, the `iscsi-ls -l` command displays the following error for each LUN:
`Cannot open /etc/sysconfig/devlabel: No such file or directory`

This message is displayed because of a bug in the command. The message does not indicate an actual error and can be safely ignored.

One or more ext3 file systems not mounted after reboot

Problem: After a reboot of the SUSE host, one or more ext3 file systems are not mounted, and the following error message is issued:
`fs type ext3 not supported by kernel`

The module that supports ext3 file systems was not loaded before the system tried to mount the file systems.

Workaround: Issue the `iscsi-mountall` command to mount the file systems.

I/O and file system errors

Problem: When stopping the host's iSCSI service, the host has I/O or file system errors.

Failure to quiesce file systems or stop applications from accessing iSCSI LUNs before stopping the iSCSI service may lead to I/O and file system errors.

Workaround: See bug number 161587 on Bugs Online at <http://now.netapp.com/NOW/cgi-bin/bol> for the latest 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.
- ◆ Storage system is unreachable for a long time
If the time during which the storage system 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: <pre>mknod --mode=660 /dev/sg32 c 21 32</pre>
The storage system 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 the <code>/dev/iscsi</code> directory or in the output of the <code>iscsi-ls</code> command.

File system errors and system shutdown hangs

Problem: If the initiator is accessing LUNs on a storage system and I/O takes a long time (for example the target is unreachable, a CFO is in progress, or the target is heavily loaded), 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 storage system I/O takes a long time.

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 storage system 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 Reiser, 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 storage system'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.

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 where to find additional information about using iSCSI on your storage system.

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.shtml ◆ <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>
Information about configuring multipathing on SLES	The Novell knowledgebase article <i>How to setup / use multipathing on SLES</i> at: http://support.novell.com/techcenter/sdb/en/2005/04/sles_multipathing.html
Information about configuring multipathing on SLES	The Novell knowledgebase article <i>How to setup / use multipathing on SLES</i> at: http://support.novell.com/techcenter/sdb/en/2005/04/sles_multipathing.html
Information on OCFS2	◆ Oracle Cluster File System (OCFS2) User's Guide http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_users_guide.pdf ◆ OCFS2 - Frequently Asked Questions http://oss.oracle.com/projects/ocfs2/dist/documentation/ocfs2_faq.html (The FAQs on OCFS2)

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://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190 http://www.netapp.com/library/tr/3369.pdf http://www.netapp.com/library/tr/3423.pdf