FCP Linux® Host Utilities 4.1 for Native OS and Veritas®
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 215-03803_A0
April 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 Network Appliance, 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 Network Appliance, Inc. 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.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
Known problems and limitations for SUSE Linux Enterprise Server hosts . . . 3
Known problems and limitations for Red Hat Enterprise Linux Hosts 10
Known problems and limitations for Veritas Storage Foundation 18
For information about additional problems. 19
Where to find additional information. 20

Changes in this release

New version

The FCP Linux® Host Utilities 4.1 is the follow-on release to the FCP Linux® Host Utilities 4.0.

Linux Versions

For the latest support information, see the FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Supported packages

Use the version of the packages that are shipped along with the specific operating system release, unless explicitly mentioned.

New features

FCP Linux Host Utilities 4.1 supports the following features:

- Support for ALUA
- Support for bundled QLogic and Emulex drivers with Red Hat Enterprise Linux
- Support for SAN Boot with Red Hat Enterprise Linux

Note

You have to download the respective errata multipathing package from the Red Hat Netwok web site at rhn.redhat.com/ for SAN boot support on Red Hat Enterprise Linux 5 Update 1. For more infromation about the errata package, see the Notes section of the NetApp® Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability.

For information about the correct version and supported configuration, see the the NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Rescan script

The qla2xxx_lun_rescan script is no longer a part of Host Utilities kit. You should use the rescan script available from the respective HBA vendor. You can download the rescan script from the respective HBA vendor Web site.

Delayed dm-multipath failback on SUSE Linux Enterprise Server 10 and later hosts

On a SUSE Linux Enterprise Server 10 and later FCP host, after a storage system giveback, the dm-multipath failback sometimes gets delayed by around 4 to 5 minutes. During this time, there is a complete IO stall on the host before IO resumes back on the primary paths.

In certain other cases after the storage system giveback, IO continues on the secondary paths for a similar delayed period before resuming on the primary paths.

During the IO stall, the `multipath -ll` command displays the primary paths status as `[failed][ready]`.

Use *directio path_checker* instead of *readsector0* or *tur* in the /etc/multipath.conf file.

Delayed dm-multipath failback on Red Hat Enterprise Linux 5 Update 1 and later hosts

On a Red Hat Enterprise Linux 5 Update 1 and later FCP host, after a storage system giveback, the dm-multipath failback sometimes gets delayed by around 4 to 5 minutes. During this time, there is a complete IO stall on the host before IO resumes back on the primary paths.

In certain other cases after the storage system giveback, IO continues on the secondary paths for a similar delayed period before resuming on the primary paths.

During the IO stall, the `multipath -ll` command displays the primary paths status as `[failed][ready]`.

Use *directio path_checker* instead of *readsector0* or *tur* in the /etc/multipath.conf file.

Known problems and limitations for SUSE Linux Enterprise Server hosts

The qlinstall script fails displaying an error message

Problem: On a Linux FCP host that has the bundled QLogic driver loaded and in use, running the `uninstall` option of the SANsurfer scli `qlinstall` script would fail displaying the following error message: “ERROR:qla2xxx does not seem to be installed”.

Workaround: Before installing the recommended QLogic driver on a Linux FCP host using the SANsurfer scli package, the following steps should be performed.

Ensure that the QLogic driver is not in use.

Uninstall the QLogic driver by the following method.

- Run the `uninstall` option of the `qlinstall` for a SANsurfer package installed on the host.
- Run the `modprobe -r qla2300/qla2400` command on a freshly installed Linux host with bundled drivers loaded.

Ensure the QLogic driver is successfully unloaded. Run the `qlinstall` script (with no options) of the SANsurfer scli package to install the requisite driver.

Commands hang on interconnect cable pull

Problem: When NetApp controller interconnect cable is pulled, commands like `multipath -ll`, `sanlun lun show` hangs.

Workaround: Make sure that the interconnect cable is properly connected and working.

Problems creating partition on Linux dm-multipath devices

Problem: In some versions of Linux, creating partitions on dm-multipath devices may cause problems. The partitioning software may report an error after a new partition is created. This error message indicates that the new partition is not recognized. As a result, new dm-multipath device will not be created automatically for the new partition.

Workaround: Complete the following steps.

1. Run fdisk to create the new partition.

2. Run the kpartx tool manually to read the partition table on the multipath device.
3. Create any new multipath devices for the partitions.

This workaround does not disturb other multipath maps. This workaround need is required only after creating a new partition.

For more information, see bug 196448 on Bugs Online at http://now.netapp.com/NOW/cgi-bin/bol.

Blacklisting devices for dm-multipath use does not work on SUSE Linux Enterprise Server 10 and later

Problem: The `devnode_blacklist` used in the /etc/multipath.conf file for blacklisting devices for dm-multipath use does not work on SUSE Linux Enterprise Server 10 and later.

Workaround: The multipath tools version has changed in SLES 10. For blacklisting on SUSE Linux Enterprise Server 10 and later, replace `devnode_blacklist` with `blacklist`.

Configuring multipath devices on Linux host fails

Problem: While configuring dm-mp devices on a Linux host, the dm-multipath driver fails to create appropriate device maps if stale sd entries are present in the sysfs. Configuring multipath devices on such a host would fail. Due to this, no dm-multipath entries would show up in the /dev/mapper/ directory as well in the `multipath -l/-ll` output.

Workaround: Identify the stale sd entries for blacklisting.

To identify the stale sd entries for blacklisting, the use the following methods.

- Blacklist those sd entries in the sysfs (or /sys/block/) which are not displayed in the `sanlun lun show` command.
- Blacklist the sd entries that displays errors when the `scsi_id -gus` command is run on them.

Using the devnode method causes incorrect blacklisting on reboot

Problem: In the multipath.conf for dm-multipath using the devnode method may causes incorrect blacklisting. This is because the device registered as sda in one boot may not remain as sda in the next boot.

Workaround: Use the wwid method rather than the devnode method for blacklisting devices in the multipath.conf file.

LUNs not recognized by dm-multipath

Problem: 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. 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 FCP devices are blacklisted, check the `devnode_blacklist` or the `blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

If the FCP 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 228744 on Bugs Online at http://now.netapp.com/NOW/cgi-bin/bol.

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.

Explanation: 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.

Unloading of the QLogic module fails

Problem: On the Linux host, unloading the QLogic module sometimes fails.

When Linux FCP initiators are included in multiple initiator FC switch zones, the HBA module usage count on the specified initiator host always shows a higher value than expected (that is, if it is more than 1), despite no apparent usage of the HBA module.

For example, for a SUSE Linux Enterprise Server 9 SP3 host with its QLogic initiators connected to a multiple initiator zone on the FC switch, the QLogic module usage count on is sometimes shown as follows:

```
qla2400 207360  0

qla2xxx 3654906 5 qla2400
```

Here, the `qla2xxx` usage count is shown as 5 (instead of 1), despite the QLogic module not being used. This could cause unloading of the QLogic module to fails.

Workaround: While connecting Linux FCP hosts to a switch fabric, always ensure that each FCP host initiator is connected to a separate individual initiator zone alone instead of a multiple initiator zone.

LUN resets observed on interconnect cable pull

Problem: During interconnect cable pulls, LUN resets are visible on Linux hosts with Emulex adapters, connected to Netapp controllers running Data Ontap 7.2. This also leads to IO stall on the host and continues till the interconnect cable is plugged back.

The Emulex initiator on a Linux host sends queries on the proxy paths, which fail during an interconnect cable pull leading to the LUN resets. This is because Data ONTAP 7.2 does not respond to queries on the proxy paths during an interconnect cable pull.

Workaround: Set the following options on both the clustered filers by running the following commands:

```
options lun.partner_unreachable.linux.behavior error

options lun.partner_unreachable.linux.hold_time 5

options lun.partner_unreachable.linux.skey 0x2

options lun.partner_unreachable.linux.asc 0x4

options lun.partner_unreachable.linux.ascq 0x1

options lun.partner_unreachable.linux.scsi_status 0x2
```

These options are applicable for all Data ONTAP versions from 7.2.0 to 7.2.3.

Linux host freeze after unloading QLogic HBA FCP driver

Problem: Linux host freeze after unloading QLogic HBA FCP driver despite having dm-multipath devices mounted on it.

Only the common qla2xxx module usage count is properly incremented/decremented (and not the qla2300/qla2400 usage count) when one uses the dm-multipath utility on a Linux host. This can lead to severe repercussions.

When commands like `umount /mnt/tmp`, `/etc/init.d/boot.multipath start`, and so on are run on the host, the machine would freeze which will make hard boot a necessary.

Workaround: Always ensure that the qla2xxx module is not in use before unloading/reloading the QLogic driver. That is, verify whether the usage count of the qla2xxx module is one before proceeding to unload/reload the QLogic driver on the Linux host.

Error message from sanlun fcp show adapter

Problem: When `sanlun fcp show adapter` is run on Linux host with Emulex adapters, using the 3.2a5 Application Kit and the 3.2a11 Application Kit, you will get error messages. The error messages that you will get on the two application kit are different. This is due to the problem in hbaapi libraries provided by both the kits.

In addition, a warning message is displayed when both these kits are installed on a x86_64 Linux host alone.

This also causes `sanlun fcp show` adapter to fail.

Workaround: Instead of the 3.2a5 or 3.2a11 kits, use the 3.2a14 kit on the Linux host. For a x86_64 host, install the 32bit libnl.so package first before installing the 3.2a14 kit to avoid the warning message. This would also ensure `sanlun fcp show adapter` functions properly on the x86_64 host.

Linux FCP hosts with QLogic initiators fail to detect LUNs

Problem: Linux FCP hosts with QLogic initiators fail to detect LUNs mapped to it when configured to auto negotiate port topology and speed. In hosts using QLogic initiators, the host initiator fails to establish a consistent topology and speed, despite setting both to auto negotiate.

This may cause the host to miss a few LUNs that are mapped to it, during the HBA scan. This behavior persists even after rebooting the host.

Workaround: When using QLogic initiators on a Linux FCP host, ensure the following settings are done for port topology and speed using the SANsurfer scli package as follows:

1. Hard set the QLogic HBA port topology to Point to Point Only while connecting to a NetApp filer operating in SSI, Standby or Partner CF modes, or to Loop Only if the filer is operating in Dual Fabric or Mixed CF modes.

   To do this:

   - Run the `/opt/QLogic_Corporation/SANsurferCLI/scli`.
   - Select the Configure HBA Settings from the Main Menu.
   - For each of the WWPNs listed here, select the ConnectionOptions and set it to Point to Point Only or Loop Only as the case may be.

2. Hard set the QLogic port speed ideally to the highest speed possible, depending on its ceiling and that of the switch/target port it is connected to.

   To do this:

   - Run the `/opt/QLogic_Corporation/SANsurferCLI/scli`.
   - Select the Configure HBA Settings from the Main Menu.

3. For each of the WWPNs listed here, select the Data Rate option and set it to the specified speed.

Always map first LUN as LUN 0

Problem: If the first LUN mapped to a host has an ID other than 0, the host creates a /dev/sd* for the missing LUN 0. 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.

Map a LUN 0 to the host and reload the HBA driver.

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.

Zoning with cfmode standby and mixed

If the active/active configuration storage system is using the standby or mixed cfmode, your switches must be zoned by WWPN.

To prevent failover problems, create zones that include the host HBAs and the primary port WWPNs on storage system. Do not include the standby port WWPNs on the storage system in the zones with the host HBA ports.

The host should never access the standby port WWPNs. During failover, the standby ports of the takeover system assume the WWPNs of the primary ports of the system that is down. The primary ports are already included in existing zones, so the host can access them.

Known problems and limitations for Red Hat Enterprise Linux Hosts

Root device dm-multipath support fails during operating system installation, if only single path to boot LUN is available

Problem: Root device dm-multipath support fails during Red Hat Enterprise Linux 5 Update1 operating system installation, if there is only a single physical path to the boot LUN. This happens despite specifying *mpath* on the kernel boot line during the operating system installation.

Even after enabling multiple physical paths to the boot LUN and then recreating the initrd image, the Red Hat Enterprise Linux 5 Update1 operating system boots from the SCSI device instead of the overlying dm-multipath device.

The boot LUN shows up as a SCSI device instead of a dm-multipath device.

Workaround: Ensure that multiple physical paths are available to the boot LUN during the Red Hat Enterprise Linux 5 Update1 operating system installation.

Root device dm-multipath support fails if *user_friendly_names* is disabled

Problem: Root device dm-multipath support fails on Red Hat Enterprise Linux 5 Update1 if *user_friendly_names* is disabled. For root device dm-multipath support, the *user_friendly_names* is set to *yes* by default during the Red Hat Enterprise Linux 5 Update1 operating system installation using anaconda.

Due to this the operating system fails to boot with the following error:
`Checking filesystems fsck.ext3: No such file or directory while trying to open /dev/mapper/mpath0p1`

The boot fails due to the `/dev/mapper/mpathp1` entry for /boot in the fstab.

Workaround: To disable *user_friendly_names* on the Red Hat Enterprise Linux 5 Update1 root on multipath host, perform the following steps:

1. Set *user_friendly_names* to *no* or remove it from the /etc/multipath.conf.
2. Change the /boot entry in the /etc/fstab from the /dev/mapper/mpath name to the corresponding /dev/mapper/<DevId> name.
3. Recreate the initrd by running the `mkinitrd` command.

Problems creating partition on Linux dm-multipath devices

Problem: In some versions of Linux, creating partitions on dm-multipath devices may cause problems. The partitioning software may report an error after a new partition is created. This error message indicates that the new partition is not recognized. As a result, new dm-multipath device will not be created automatically for the new partition.

Workaround: Complete the following steps.

1. Run fdisk to create the new partition
2. Run the kpartx tool manually to read the partition table on the multipath device
3. Create any new multipath devices for the partitions.

This workaround does not disturb other multipath maps. This workaround need is required only after creating a new partition.

For more information, see bug 196448 on Bugs Online at http://now.netapp.com/NOW/cgi-bin/bol.

Linux host freeze after unloading QLogic HBA FCP driver

Problem: Linux host freeze after unloading QLogic HBA FCP driver despite having dm-multipath devices mounted on it.

Only the common qla2xxx module usage count is properly incremented/decremented (and not the qla2300/qla2400 usage count) when one uses the dm-multipath utility on a Linux host. This can lead to severe repercussions.

When commands like `umount /mnt/tmp`, `/etc/init.d/boot.multipath start`, and so on, are run on the host, the machine would freeze which will make hard boot a necessary.

Workaround: Always ensure that the qla2xxx module is not in use before unloading/reloading the QLogic driver. That is, verify whether the usage count of the qla2xxx module is one before proceeding to unload/reload the QLogic driver on the Linux host.

Delayed dm-multipath failback on Red Hat Enterprise Linux 5.0 hosts

Problem: On a Red Hat Enterprise Linux 5 FCP host, after a storage system giveback, the dm-multipath failback sometimes gets delayed by around 4 to 5 minutes. During this time, there is a complete IO stall on the host before IO resumes back on the primary paths.

In certain other cases after the storage system giveback, IO continues on the secondary paths for a similar delayed period before resuming on the primary paths.

During the IO stall, the `multipath -ll` command displays the primary paths status as `[failed][ready]`.

Workaround: During the delay, the primary paths status can be updated by restarting the `multipathd` daemon by running the following command:

```
/etc/init.d/multipathd restart
```

The IO will immediately resume back on the primary paths.

LUN resets observed on interconnect cable pull

Problem: During interconnect cable pulls, LUN resets are visible on Linux hosts with Emulex adapters, connected to Netapp controllers running Data Ontap 7.2. This also leads to IO stall on the host and continues till the interconnect cable is plugged back.

The Emulex initiator on a Linux host sends queries on the proxy paths, which fail during an interconnect cable pull leading to the LUN resets. This is because Data ONTAP 7.2 does not respond to queries on the proxy paths during an interconnect cable pull.

Workaround: Set the following options on both the clustered filers by running the below commands -

```
options lun.partner_unreachable.linux.behavior error
options lun.partner_unreachable.linux.hold_time 5
options lun.partner_unreachable.linux.skey 0x2
options lun.partner_unreachable.linux.asc 0x4
options lun.partner_unreachable.linux.ascq 0x1
options lun.partner_unreachable.linux.scsi_status 0x2
```

These options are applicable for all Data ONTAP versions from 7.2.0 to 7.2.3.

Commands hang on interconnect cable pull

Problem: When NetApp controller interconnect cable is pulled, commands like `multipath -ll`, `sanlun lun show`, `vxdctl enable` hangs.

Workaround: Make sure that the interconnect cable is properly connected and working.

Error message from sanlun fcp show adapter

Problem: When `sanlun fcp show adapter` is run on Linux host with Emulex adapters, using the 3.2a5 Application Kit and the 3.2a11 Application Kit, you will get error messages. The error messages that you will get on the two application kit are different. This is due to the problem in hbaapi libraries provided by both the kits.

In addition, a warning message is displayed when both these kits are installed on a x86_64 Linux host alone.

This also causes `sanlun fcp show adapter` to fail.

Workaround: Instead of the 3.2a5 or 3.2a11 kits, use the 3.2a14 kit on the Linux host. For a x86_64 host, install the 32bit libnl.so package first before installing the 3.2a14 kit to avoid the warning message. This would also ensure `sanlun fcp show adapter` functions properly on the x86_64 host.

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.

Explanation: 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.

Unloading of the QLogic module fails

Problem: On the Linux host, unloading the QLogic module sometimes fails.

When Linux FCP initiators are included in multiple initiator FC switch zones, the HBA module usage count on the specified initiator host always shows a higher value than expected (that is, if it is more than 1), despite no apparent usage of the HBA module.

For example, for a Linux host with its QLogic initiators connected to a multiple initiator zone on the FC switch, the QLogic module usage count on is sometimes shown as follows:

```
qla2400 207360  0

qla2xxx 3654906 5 qla2400
```

Here, the `qla2xxx` usage count is shown as 5 (instead of 1), despite the QLogic module not being used. This could cause unloading of the QLogic module to fails.

Workaround: While connecting Linux FCP hosts to a switch fabric, always ensure that each FCP host initiator is connected to a separate individual initiator zone alone instead of a multiple initiator zone.

Automounting of dm-multipath devices fails

Problem: Automounting dm-multipath devices using /dev/mpath/ entries in the fstab fails on Red Hat Enterprise Linux 5 hosts.

If *user_friendly_names* is set to *yes* in the multipath.conf file on a Red Hat Enterprise Linux 5 host, it causes problems at boot up. No /dev/mpath entries are available at boot up which causes the automounting to fail.

Workaround: Use the /dev/mapper/ entries for automounting dm-multipath devices on a red Hat Enterprise Linux host instead of /dev/mpath/ entries.

The qlinstall script fails displaying an error message

Problem: On a Linux FCP host that has the bundled QLogic driver loaded and in use, running the `uninstall` option of the SANsurfer scli `qlinstall` script would fail displaying the following error message: “ERROR:qla2xxx does not seem to be installed”.

Workaround: Before installing the recommended QLogic driver on a Linux FCP host using the SANsurfer scli package, the following steps should be performed.

Ensure that the QLogic driver is not in use.

Uninstall the QLogic driver by the following method.

- Run the `uninstall` option of the `qlinstall` for a SANsurfer package installed on the host.
- Run the `modprobe -r qla2300/qla2400` command on a freshly installed Linux host with bundled drivers loaded.

Ensure the QLogic driver is successfully unloaded. Run the `qlinstall` script (with no options) of the SAN surfer scli package to install the requisite driver.

Blacklisting devices for dm-multipath use does not work on Red Hat Enterprise Linux 5

Problem: The `devnode_blacklist` used in the /etc/multipath.conf file for blacklisting devices for dm-multipath use does not work on Red Hat Enterprise Linux 5.

Workaround: The multipath tools version has changed in Red Hat Enterprise Linux 5. For blacklisting on Red Hat Enterprise Linux 5, replace `devnode_blacklist` with `blacklist`.

Always map first LUN as LUN 0

Problem: If the first LUN mapped to a host has an ID other than 0, the host creates a /dev/sd* for the missing LUN 0. 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.

Map a LUN 0 to the host and reload the HBA driver.

Configuring multipath devices on Linux host fails

Problem: While configuring dm-mp devices on a Linux host, the dm-mp driver fails to create appropriate device maps if stale sd entries are present in the sysfs. Configuring multipath devices on such a host would fail. Due to this, no dm-mp entries would show up in the /dev/mapper/ directory as well in the `multipath -l/-ll` output.

Workaround: Identify the stale sd entries for blacklisting.

To identify the stale sd entries for blacklisting, the use the following methods.

- Blacklist those sd entries in the sysfs (or /sys/block/) which are not displayed in the `sanlun lun show` command.
- Blacklist the sd entries that displays errors when the `scsi_id -gus` command is run on them.

Using the devnode method causes incorrect blacklisting on reboot

Problem: In the multipath.conf for dm-multipath using the devnode method may causes incorrect blacklisting. This is because the device registered as sda in one boot may not remain as sda in the next boot.

Workaround: Use the wwid method rather than the devnode method for blacklisting devices in the multipath.conf file.

Multiple devices shown for each LUN

Problem: The `sanlun lun show` command displays multiple devices for each LUN that is mapped to the Linux host. This problem occurs only with Red Hat Enterprise Linux 4 Update 2 and later without dm-multipath.

Workaround: Use the dm-multipath support to manage multiple paths. Otherwise, configure the storage system or zoning so that there is only one path to each LUN.

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.

Zoning with cfmode standby and mixed

If the active/active configuration storage system is using the standby or mixed cfmode, your switches must be zoned by WWPN.

To prevent failover problems, create zones that include the host HBAs and the primary port WWPNs on storage system. Do not include the standby port WWPNs on the storage system in the zones with the host HBA ports.

The host should never access the standby port WWPNs. During failover, the standby ports of the takeover system assume the WWPNs of the primary ports of the system that is down. The primary ports are already included in existing zones, so the host can access them.

LUNs not recognized by dm-multipath

Problem: 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. 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 FCP devices are blacklisted, check the `devnode_blacklist or blacklist` section of the `/etc/multipath.conf` file. Ensure that the entries are correctly specified.

If the FCP 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 228744 on Bugs Online at http://now.netapp.com/NOW/cgi-bin/bol.

Problem: Secondary paths identified incorrectly as primary paths while configuring DMP.

Workaround: Verify whether Array Support Library (ASL) for NetApp storage is installed correctly. That is, the ASL is successfully installed without any error messages.

Problem: VxVM's enclosure based naming feature is enabled, but the DMP devices do not contain a NetApp storage model string.

Workaround: Verify whether the Array Support Library (ASL) for NetApp storage is installed correctly. That is, the ASL is successfully installed without any error messages.

Problem: SAN boot installation to NetApp LUN fails.

Workaround: Verify whether the host can only see one path to the NetApp LUN during installation.

Problem: Encapsulation of root disk fails after SAN boot installation.

Workaround: Make sure to follow the partitioning scheme given in the *Veritas™ Volume Manager Administrator's Guide* for Linux.

Problem: When NetApp controller interconnect cable is pulled, commands like `multipath -ll`, `sanlun lun show`, `vxdctl enable` hangs.

Workaround: Make sure that the interconnect cable is properly connected and working.

For information about additional problems

Occasionally, additional problems are reported with the Support Kit 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.
4	In the Bug Search Tools section, choose FCP-Linux from the Bug Type(s) menu.
5	Click Go. **Result:** The Bugs Online - Advanced Search Results lists all bugs for Linux FCP hosts.

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	*FCP Host Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml
The supported storage system models for Data ONTAP licensed with FCP	◆ *FCP Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ ◆ *System Configuration Guide* for Data ONTAP at http://now.netapp.com
Information about how to configure and manage the FCP service on the storage system	The following documents for your Data ONTAP release: ◆ *Data ONTAP Block Access Management Guide* ◆ *Data ONTAP Release Notes*