

Linux® Host Utilities 5.3

Release Notes

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
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: 215-06198_A0
August 2011

Contents

Linux Host Utilities 5.3 Release Notes	7
The 5.3 release of the Linux Host Utilities	7
Information on finding system requirements	7
New features in the 5.3 release of the Linux Host Utilities	8
Updates to the Release Notes	8
Important Cautions	11
Known problems and limitations	13
(FC) Creating partitions on DM-Multipath causes an error	13
(FC) Host freezes after unloading QLogic driver	14
(FC) DM-Multipath failback delayed after giveback	14
(FC) LUN resets observed during interconnect cable pulls	15
(FC) Commands hang when interconnect cable is pulled	15
(FC) sanlun displays error messages with certain Emulex kits	15
(FC) scli qlinstall script fails when running uninstall option	16
(FC) Host creates /dev/sd* if LUN 0 missing	16
(FC) The DM-Multipath driver fails to create device maps for new devices	17
(FC) Using devnode causes blacklist problems on reboots	17
(iSCSI) OCFS2 self-fencing occurs during cluster failover	17
(Red Hat) Known issues when running Red Hat Enterprise Linux	19
(Red Hat, FC) Red Hat 6.1 host fails to add all LUN paths after fabric fault	19
(Red Hat, FC) Support for DM-Multipath root device fails during operating system installation	20
(Red Hat, FC) Support for DM-Multipath root device fails if user_friendly_names is disabled	20
(Red Hat, FC) Automatic mounting of DM-Multipath devices fails	21
(Red Hat, iSCSI) Online GFS grow feature is not supported on iSCSI LUN	21
(Red Hat, iSCSI) service clvmd restart fails	22
(Red Hat, FC) With Red Hat 6 and QLogic, host occasionally crashes at qla2x00_process_completed_request	23
(Red Hat, FC) With Red Hat 6 and QLogic, host becomes unresponsive during I/O with faults	23

(Red Hat, FC) With Red Hat 6 and Emulex, host occasionally panics at lpfc_abort_handler	24
(Red Hat, FC) Host running Red Hat 6 or 5.6 with Emulex 8 Gb hangs during fabric faults	24
(Red Hat) With Red Hat 6, wrong path grouping occurs after storage system fault with I/O	25
(Red Hat) With Red Hat 6, DM-Multipath fails to update paths during I/O with fabric faults	26
(Red Hat, iSCSI) With Red Hat 6, iscsi script fails to show correct status	26
(Red Hat, iSCSI) With Red Hat 5 Update 2, host freezes if multipathd halts	27
(Red Hat, iSCSI) With Red Hat 5 Update 2, host intermittently stops during SAN boot	28
(Red Hat, iSCSI) With Red Hat 4 Update 3, system hangs during shutdown, reboot	29
(Red Hat, iSCSI) With Red Hat 4 Update 2, first LUN must be LUN 0	30
(Red Hat, iSCSI) With Red Hat 4, continuous ISID Rule messages displayed	30
(Red Hat) With Red Hat 4 and a large number of LUNs, DM-Multipath sometimes fails to configure devices	31
(SUSE) Known issues when running SUSE Linux Enterprise Server ..	33
(SUSE, FC) With SUSE 11 SP1, DM-Multipath sometimes fails to update maps after I/O fault	33
(SUSE, FC) With SUSE 10.3 and Emulex 4 Gb, FC ports go offline during controller faults	34
(SUSE, iSCSI) Online resize of file system not supported	34
(SUSE, iSCSI) One or more ext3 file systems not mounted after reboot	35
(SUSE, iSCSI) sanlun does not display all iSCSI LUNs	35
(Veritas) Known issues when running the Host Utilities with Veritas ..	37
(Veritas) With Red Hat 4 Update 5, paths not reenabled after giveback	37
(Veritas) I/O delay seen on systems using iSCSI	37
(Veritas) With Oracle Enterprise Linux 5.4 and Emulex, Veritas volumes go offline	38
(Veritas) With SUSE Linux and iSCSI, volumes might enter disabled state after controller fault	38
Viewing current and fixed product bugs	41
NetApp product documentation	43
How to search for documentation on the NetApp Support Site	43

Links to NetApp product documentation 43

Contact Technical Support 45

 Things to check before you call Technical Support 45

 Information you need to provide to Technical Support 45

 How to contact NetApp Technical Support 45

Copyright information 47

Trademark information 49

How to send your comments 51

Linux Host Utilities 5.3 Release Notes

These Release Notes contain the latest information for the Linux® Host Utilities 5.3, including updates about known problems and limitations, information on tunable parameters, and any important cautions. The Release Notes also include any issues that may have been discovered since the Host Utilities were released.

These Release Notes are updated when new information on using the Linux Host Utilities becomes available. It is a good practice to periodically check the Release Notes to see if new information has been added. See the section *Recent updates to the Release Notes* to determine whether information has been added since you last checked.

The 5.3 release of the Linux Host Utilities

The Linux Host Utilities support several Linux operating systems as well as Veritas™ Storage Foundation. The Host Utilities enable you to connect a Linux host to NetApp storage systems.

You can use the Host Utilities with:

- Red Hat® Enterprise Linux
- SUSE® Linux Enterprise Server
- Oracle® Enterprise Linux
- Oracle VM

Note: For specific information about the operating system versions and updates that the Host Utilities support, see the Interoperability Matrix.

The Host Utilities software is provided as a single .rpm file. You can download the .rpm file and the documentation from the NOW (NetApp on the Web®) site. The *Linux Host Utilities 5.3 Installation and Setup Guide* provides instructions for installing and setting up the Host Utilities to work with your environment and protocol.

Information on finding system requirements

NetApp qualifies additional components between releases of the Host Utilities software.

For the most up-to-date information on the system requirements, see the NetApp Interoperability Matrix.

Related information

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

New features in the 5.3 release of the Linux Host Utilities

The October 2010 release of the Linux Host Utilities 5.3 added several features.

- The Host Utilities now use .rpm as the package file format for the software download file.
- The Host Utilities added support for the following:
 - Kernel-based Virtual Machines (KVMs)
 - Oracle VM
 - Broadcom 10Gb iSCSI offload
 - The iSCSI protocol on Red Hat Enterprise Linux 5 Update 5
 - Veritas Storage Foundation 5.0 MP4 with both FC and iSCSI
 - Veritas Storage Foundation 5.1 RP1 HF2 with both FC and iSCSI
- The Host Utilities no longer provide the diagnostic programs that were previously part of the toolkit. Instead, the Host Utilities support the nSANity Diagnostic and Configuration Data Collector program. This program is not part of the Host Utilities. You can download it from the NOW site if Technical Support asks you to do so.

Some of these changes are transparent to the end-user. The *Linux Host Utilities 5.3 Installation and Setup Guide* contains explanations of these new features where needed.

Note: NetApp continues to work with the Host Utilities and add support for features after the initial release. For the most up-to-date information on which features are supported and which new features have been added, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Updates to the Release Notes

The Release Notes are updated between product releases as new information is uncovered. The following changes have occurred since the Linux Host Utilities 5.3 were released.

August 2011 --The Release Notes were updated in August 2011 as a result of new information on the Linux Host Utilities, including:

- Changes to the *Linux Host Utilities 5.3 Installation and Setup Guide*. For pointers to the updated sections see the *Changes to this document: August 2011* section in the install guide. It lists the new features that affect the documentation.
- An issue has been discovered where a host running Red Hat Enterprise Linux 6.1 and FC sometimes fails to add back all the paths to LUNs after it recovers from a fabric fault. For more information, see the section (*Red Hat, FC*) *Red Hat 6.1 host fails to add all LUN paths after fabric fault*.
- An issue has been discovered where a host running SUSE Linux Enterprise Server 10.3 and an Emulex 4-Gb FC HBA. In this environment, FC ports go off-line during storage controller faults.

For more information, see the section *(SUSE, FC) With SUSE 10.3 and Emulex 4 Gb, FC ports go offline during controller faults*.

April 2011 --The Release Notes were updated in April 2011 as a result of updates to the Linux Host Utilities. This update includes the following new information:

- Changes to the *Linux Host Utilities 5.3 Installation and Setup Guide*. This guide was updated to reflect support for Red Hat Enterprise Linux 6 series and SUSE Linux Enterprise Server 11 with Service Pack 1 with the FCoE, FC, and iSCSI protocols. For pointers to the updated sections see the *Changes to this document: April 2011* section in the install guide.

Note: If you are running Data ONTAP 8.0.1 or higher on your storage systems, Red Hat Enterprise Linux 6 automatically recognizes the BLOCKS LIMITS VPD page of the storage system and aligns itself with the limits specified in the VPD page.

- An issue has been discovered on systems running Red Hat Enterprise Linux 6 and QLogic. The host sometimes crashes at `qla2x00_process_completed_request` during I/O with controller faults. This issue has been seen with the QLogic QLE2562 FC HBA. For more information, see the section *(Red Hat, FC) With Red Hat 6 and QLogic, host occasionally crashes at qla2x00_process_completed_request*.
- An issue has been discovered on systems running Red Hat Enterprise Linux 6, QLogic, and DM-Multipath and also has DM-Multipath configured on some LUNs. In this environment, the host sometimes hangs during I/O with controller faults and fails to recover. For more information, see the section *(Red Hat, FC) With Red Hat 6 and QLogic, host becomes unresponsive during I/O with faults*.
- An issue has been discovered on systems running Red Hat Enterprise Linux 6 and a Emulex where the host sometimes panics at the `lpfc_abort_handler` during I/O with controller faults. For more information, see the section *(Red Hat, FC) With Red Hat 6 and Emulex, host occasionally panics at lpfc_abort_handler*.
- Red Hat Enterprise Linux 6 has the same issue as Red Hat Linux 5 Update 6. The host hangs during fabric faults when it is using an 8 Gb Emulex adapter that uses an inbox LPFC driver and firmware v2.00A3. For more information, see the section *(Red Hat, FC) Host running Red Hat 6 or 5.6 with Emulex 8 Gb hangs during fabric faults*.
- An issue has been discovered with Red Hat Enterprise Linux 6 where DM-Multipath wrongly updates the path groups on a host that is running Red Hat Enterprise Linux 6.0. This problem occurs during I/O with a storage system and fabric faults. For more information, see the section *(Red Hat) With Red Hat 6, wrong path grouping occurs after storage system fault with I/O*.
- An issue has been discovered on systems running Red Hat Enterprise Linux 6 and a version of DM-Multipath is using the device-mapper-multipath errata v0.4.9.31.el6_0.2. In these environments, DM-Multipath sometimes fails to update paths during I/O with fabric faults. For more information, see the section *(Red Hat) With Red Hat 6, DM-Multipath fails to update paths during I/O with fabric faults*.
- With Red Hat Enterprise Linux 6, the `/etc/init.d/iscsi` script does not show the correct status for the iSCSI daemon when you install the root file system on an iSCSI LUN. For more information, see the section *(Red Hat) With Red Hat 6, DM-Multipath fails to update paths during I/O with fabric faults*.

- On occasion when a host is running Red Hat Linux 4 and has a large number of newly discovered LUNs, DM-Multipath fails to create all the maps for the LUNs. For more information, see the section *(Red Hat, iSCSI) With Red Hat 6, iscsi script fails to show correct status*.
- An issue has been discovered on systems running SUSE Linux Enterprise Server 11 SP1 with multiple LUNs where DM-Multipath sometimes fails to update its maps during I/O with controller faults. For more information, see the section *(SUSE, FC) With SUSE 11 SP1, DM-Multipath sometimes fails to update maps after I/O fault*.

January 2011 --The Release Notes were updated in January 2011. This update includes the following new information:

- Updates to the *Linux Host Utilities 5.3 Installation and Setup Guide*. The Linux Host Utilities have added support for Veritas Storage Foundation 5.1 SP1 as well as Red Hat Enterprise Linux 5 Update 5 with the FC and FCoE protocols and Red Hat Enterprise Linux 5 Update 6 with the FC, FCoE, and iSCSI protocols. These environments require changes to some of the procedures. Where necessary, the *Linux Host Utilities 5.3 Installation and Setup Guide* has been updated with new steps that apply to these environments. For pointers to the updated sections see the *Changes to this document: January 2011* section in the install guide.
- An issue has been discovered where a host running Red Hat Linux 5 Update 6 hangs during fabric faults. This occurs when the host is using an 8 Gb Emulex adapter that uses an inbox LPFC driver and firmware v2.00A3. For more information, see the section *(Red Hat, FC) Host running Red Hat 5.6 with Emulex 8 Gb hangs during fabric faults*.

November 2010 --The Release Notes were updated in November 2010. This update includes the following new information:

- An issue has been discovered where, when using Veritas Storage Foundation 5.1 with SUSE Linux Enterprise Server and the iSCSI protocol, volumes might go into a disabled state after a storage controller fault. To work around this issue, you need to set the VxDMP parameter `dmp_lun_retry_timeout` to 300. For more information, see the section *(Veritas) With SUSE Linux and iSCSI, volumes disabled after controller fault*.

Important Cautions

There are no important cautions at this time.

Known problems and limitations

Some unexpected and potentially undesired behaviors, as well as, in some cases, workarounds to avoid these behaviors, have been identified.

To request more information about an individual bug and to explore other bug-related tools, see Bugs Online on the NOW site.

(FC) Creating partitions on DM-Multipath causes an error

In some cases when a system is running the FC protocol, the partitioning software reports an error after you create a partition on DM-Multipath.

About this task

For example, the `fdisk` command might produce the following error messages when it cannot update the partition table on the disk.

```
WARNING: Re-reading the partition table failed with error 22: Invalid argument.
```

```
The kernel still uses the old table.
```

```
The new table will be used at the next reboot.
```

The errors indicate that the new partition is not recognized, which means that the system does not automatically create a new DM-Multipath device for the partition.

The following workaround resolves this problem. This workaround does not disturb other multipath maps and only needs to be performed after a new partition is created.

Steps

1. Create a new partition by executing the following command:

```
fdisk
```

2. Read the partition table on the multipath device by executing the following tool:

```
kpartx -a <partition_name>
```

3. Create new multipath devices for the partitions.

Example of creating a new partition

```
# ls -l /dev/mapper/
total 0
crw----- 1 root root 10, 63 Apr 26 18:08 control
brw-rw---- 1 root disk 253, 0 Apr 27 16:10 mpath0
# kpartx -a /dev/mapper/mpath0
```

```
# ls -l /dev/mapper/
total 0
crw----- 1 root root 10, 63 Apr 26 18:08 control
brw-rw---- 1 root disk 253, 0 Apr 27 16:10 mpath0
brw-rw---- 1 root disk 253, 1 Apr 27 16:11 mpath0p1
```

After you finish

For the latest status on this issue, see bug 196448 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(FC) Host freezes after unloading QLogic driver

The Linux host freezes after you unload or reload the QLogic HBA FC driver even though DM-Multipath devices are mounted on it.

When you use DM-Multipath, only the common qla2xxx module usage count is incremented or decremented. The qla2300/qla2400 usage count is not changed. If you unload or reload the QLogic driver while the qla2xxx module is in use, then the next time you run a command such as `umount /mnt/tmp` or `/etc/init.d/boot.multipath start` on the host, the host freezes until you reboot it.

To prevent this problem, always make sure the qla2xxx module is not in use before you unload or reload the QLogic driver. A usage count of 1 for the qla2xxx module indicates that it is not in use.

(FC) DM-Multipath failback delayed after giveback

In certain environments, DM-Multipath failback is delayed by four to five minutes after a controller giveback occurs. This can happen in environments where the `/etc/multipath.conf` file uses `readsector0` `path_checker` and the host is running the FC protocol and either SUSE Linux Enterprise Server 10, SUSE Linux Enterprise Server 10 SP1, or Red Hat Enterprise Linux 5.

During this time, a complete I/O stall occurs on the host before I/O resumes on the primary paths.

It is also possible that, after a giveback, I/O continues on the secondary paths for the same amount of time as the DM-Multipath failback delay before it resumes on the primary paths.

During this I/O stall, the `multipath -ll` command displays the primary paths status as `[failed]` `[ready]`. In addition, no `multipathd` messages are logged in the syslog, even though `multipathd` has a default polling interval of 5 seconds.

To work around this problem, use one of the following methods:

- **(SUSE Linux, Red Hat Linux 5 Update 1 and later)** In the `/etc/multipath.conf` file, use `path_checker directio` instead of `readsector0` or `tur`. Doing this ensures there are no delays during failback.
- During the delayed failback, update the primary path status by restarting `multipathd`:
`/etc/init.d/multipathd restart`
 After you do that, I/O immediately resumes on the primary paths.

(FC) LUN resets observed during interconnect cable pulls

When your Host Utilities environment uses Emulex adapters, clustered controllers, and a version of Data ONTAP prior to 7.2.4, LUN resets occur each time the interconnect cable is pulled. This causes I/O to stall on the host until the interconnect cable is plugged back in.

The problem occurs because the Emulex initiator on the host sends queries on the proxy paths. These queries fail during an interconnect cable pull. The LUN resets occur because Data ONTAP 7.2 does not respond to queries on proxy paths during an interconnect cable pull.

If you are using Data ONTAP versions 7.2.0 to 7.2.3, you can prevent this problem by running the following command lines on each of the controllers:

```
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
```

(FC) Commands hang when interconnect cable is pulled

When the NetApp controller interconnect cable is pulled, commands such as `multipath -ll` and `sanlun lun show` hang.

To prevent this problem, ensure that the interconnect cable is properly connected and working.

(FC) sanlun displays error messages with certain Emulex kits

When you run the `sanlun fcp show adapter` command on a host that has the Emulex 3.2a5 Application Kit or the 3.2a11 Application Kit installed, you get error messages. This is due to a problem in the hbaapi libraries that the kits provide.

Note: The error messages differ based on the application kit you have installed.

In addition, a warning message is displayed when both these kits are installed on an x86_64 Linux host due to some missing 32-bit libraries. In this environment, the `sanlun fcp show adapter` command fails.

To prevent these problems, use the Emulex 3.2a14 Application Kit or a later version.

If you have an x86_64 host, install the 32-bit `libnl.so` package before you install either the 3.2a14 kit or a later version. Installing the software in this order ensures that the `sanlun fcp show adapter` command functions properly on the x86_64 host.

(FC) `scli qlinstall` script fails when running `uninstall` option

The SANsurfer `scli qlinstall` script fails when running the `uninstall` option on a Linux host that is using FC and a QLogic driver that was part of the Linux operating system distribution.

The script displays an error message similar to the following:

```
ERROR:qla2xxx does not seem to be installed
```

To avoid this problem you must unload the current driver and then install the required driver. Here is the workaround for this issue:

1. Make sure the current QLogic driver is not in use.
 - a. Run the `lsmod | grep qla2` command to display the usage count.
 - b. If the device is in use, stop all the processes using devices that are accessed through the QLogic driver.
2. Uninstall the current QLogic driver by performing the following tasks on the host where the driver is loaded:
 - a. Run the `uninstall` option of the SANsurfer `scli qlinstall` script.
 - b. Run the `modprobe -r {qla2300 | qla2400}` command.
3. Confirm that the QLogic driver has been successfully unloaded.
 - a. Run the `lsmod | grep qla2xxx` command. If the driver does not appear in the output, then it has been successfully unloaded.
4. Install the required driver by running the SANsurfer `scli qlinstall` script without any options.

(FC) Host creates `/dev/sd*` if LUN 0 missing

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.

To resolve the problem, map a LUN 0 to the host and reload the HBA driver.

Note: Adding a LUN 0 without reloading the driver does not resolve the problem.

To prevent this problem, when you map LUNs to the host, make sure the first LUN is always LUN 0. By default, Data ONTAP does this when you do not explicitly specify the LUN identifier.

(FC) The DM-Multipath driver fails to create device maps for new devices

The DM-Multipath driver fails to create device maps after you configure multipath devices on the host. No DM-Multipath entries appear in the `/dev/mapper/` directory or in the `multipath -l/-ll` output.

This problem occurs when there are stale `sd` entries in the `sysfs` file. The DM-Multipath driver cannot create appropriate device maps when there are stale `sd` entries in the `sysfs` file.

To work around this problem you must identify the stale `sd` entries and exclude (blacklist) them. Use one of the following methods to determine which entries to exclude:

- Run the `san lun lun show` command. Then check the `sysfs` (or `/sys/block/`) file. Exclude all the `sd` entries that were not displayed by the `san lun lun show` command.
- Run the `scsi_id -gus` command and exclude the `sd` entries that display errors.

To exclude entries, you can modify the blacklist section of the DM-Multipath configuration file. For more information, see the *Linux Host Utilities Installation and Setup Guide*.

(FC) Using devnode causes blacklist problems on reboots

A problem can occur when you use the `devnode` method to exclude (blacklist) a device in the DM-Multipath configuration file. When you reboot the host after using the `devnode` method, the wrong device is blacklisted.

This problem occurs because a device that is registered as `sda` during one boot process might not be registered as `sda` in the next boot process.

The workaround for this problem is to use the `WWID` method to blacklist the device in the `multipath.conf` file.

(iSCSI) OCFS2 self-fencing occurs during cluster failover

During a clustered failover or network disruption on a system that is using iSCSI, Oracle Cluster File System (OCFS2) nodes might fence themselves by panicking.

OCFS2 uses 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 cannot write or read the heartbeat value, it retries the

operation a specific number of times before it self-fences by panicking the Linux kernel as it fails to obtain quorum.

During a clustered failover or network disruption, all paths to the storage system can become unavailable, which means that none of the nodes in the OCFS2 cluster is able to write or read a heartbeat value. This can cause the nodes to fence themselves by panicking.

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 do not self-fence.

With a dedicated TCP/IP network for the iSCSI link between Linux hosts and the storage system, a reasonable value is usually 181. This value corresponds to a total of 360 seconds before a node self-fences.

Depending on the load in the iSCSI network, you should raise the value. It is best to use a value in the range of 61 to 201, which is ideal for a configuration with OCFS2 on multipath and iSCSI devices.

For more information, see the OCFS FAQs.

Related information

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

(Red Hat) Known issues when running Red Hat Enterprise Linux

The following issues might occur when you are running Red Hat Enterprise Linux. In many cases, workarounds to avoid these issues have been identified.

(Red Hat, FC) Red Hat 6.1 host fails to add all LUN paths after fabric fault

A host that is running FC and Red Hat Enterprise Linux 6.1 sometimes fails to add back all the paths to LUNs after it recovers from a fabric fault. This issue occurs on hosts that are supporting I/O and have only a few LUNs.

In this situation, the host recovers from the fabric fault, but does not detect all the SCSI devices.

If you execute the `multipath -ll` command for all the DM-Multipath devices that initially had four paths to each LUN, it returns output similar to the following:

```
360a98000572d58656d4a6441566b4156 dm-2 NETAPP,LUN
size=10G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|  - 5:0:1:6 sdn 8:208 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
|  - 5:0:0:6 sdy 65:128 active ready running
|  - 4:0:0:6 sdaa 65:160 active ready running
```

At this time, there is no workaround for this issue. If you encounter this issue, you should rescan the SCSI bus by executing the `rescan-scsi-bus.sh` script. This allows you to manually rediscover and add the lost SCSI devices.

For example, you might enter the command:

`rescan-scsi-bus.sh -L max_luns`

`max_luns` is the maximum number of LUNs mapped to the host.

Red Hat is tracking this issue in Red Hat Bugzilla 713682.

For NetApp's latest status on this issue, see bug 511638 at Bugs Online.

Related information

Red Hat Bugzilla 713682 - https://bugzilla.redhat.com/show_bug.cgi?id=713682

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, FC) Support for DM-Multipath root device fails during operating system installation

If you have only one physical path to the boot LUN, support for the DM-Multipath root device fails when you install Red Hat Enterprise Linux 5 Update 1. This failure happens even if you specify `mpath` on the kernel boot line during the operating system installation.

Red Hat Enterprise Linux 5 Update 1 boots from the SCSI device, not the overlying DM-Multipath device. This is true even after you enable multiple physical paths to the boot LUN and then re-create the `initrd` image. The boot LUN shows up as a SCSI device instead of a DM-Multipath device.

To work around this problem, set up the boot LUN so that multiple physical paths are available for use. You can check the HBA BIOS to determine the number of paths to the LUN. Once you have multiple physical paths available, install Red Hat Enterprise Linux 5 Update 1.

For instructions on how to set up a SAN boot LUN, see the *Linux Host Utilities 5.3 Installation and Setup Guide*.

(Red Hat, FC) Support for DM-Multipath root device fails if `user_friendly_names` is disabled

Support for the DM-Multipath root device fails on Red Hat Enterprise Linux 5 Update 1 if `user_friendly_names` is disabled.

To enable support for the DM-Multipath root device, the `user_friendly_names` option is set to `yes` by default when you install Red Hat Enterprise Linux 5 Update 1 using Anaconda.

If you change this value to `no`, the operating system fails to boot and displays an error message similar to the following:

```
Checking filesystems fsck.ext3: No such file or directory while trying to
open /dev/mapper/mpath0p1
```

The boot operation fails because of the `/dev/mapper/mpathp1` entry for `/boot` in `fstab`.

If you need to disable `user_friendly_names` on the Red Hat Enterprise Linux 5 Update 1 root when you have a multipath host, complete the following steps:

1. In the `/etc/multipath.conf` file, either set `user_friendly_names` to `no` or remove it from the file.
2. Change the `/boot` entry in `/etc/fstab` from the `/dev/mapper/mpath` name to the corresponding `/dev/mapper/<DevId>` name.
3. Re-create the `initrd` by running the following command:

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

(Red Hat, FC) Automatic mounting of DM-Multipath devices fails

Automatic mounting of DM-Multipath devices by using `/dev/mpath/` entries in `fstab` fails on Red Hat Enterprise Linux 5 hosts.

If the `user_friendly_names` option is set to `yes` in the DM-Multipath configuration file (`multipath.conf`) on a Red Hat Enterprise Linux 5 host, it causes problems during the boot process. No `/dev/mpath` entries are available during the boot process, which causes the automatic mounting to fail.

Use `/dev/mapper/` entries to set up automatic mounting of DM-Multipath devices on a Red Hat Enterprise Linux host instead of `/dev/mpath/` entries.

(Red Hat, iSCSI) Online GFS grow feature is not supported on iSCSI LUN

If you have a file system that is created over an iSCSI device or a multipath device on a LUN, you cannot use the online GFS grow feature in the Global File System (GFS) .

About this task

When you use GFS on the LVM volume, you can extend the volume without unmounting the file system. You do this by adding new LUNs to the volume group and expanding the logical volume on which the GFS is created.

NetApp storage systems also allow you to resize a LUN dynamically; however, the iSCSI layer in Linux is not capable of detecting the change in the LUN size. To display the new size of the LUN you must restart the iSCSI service. But a restart of the iSCSI service fails if any iSCSI devices are mounted at that time. Because of this, even if the LUN on which the GFS file system is located is resized, you must unmount the file system and then restart iSCSI and the multipathing service before Linux can detect the new size.

The workaround for growing a GFS consists of the following steps.

Steps

1. Unmount the file system.
2. Resize the LUN on the storage system.
3. To flush the multipath, run the following command:

```
multipath -F
```

4. To stop the multipath service, run the following command:
service multipath stop
5. To start the iSCSI session, run the following command:
service iscsi restart
6. To restart the multipath service, run the following command:
service multipath start
7. Mount the file system.
8. On one of the nodes, run the following command:
gfs_grow
9. To verify the new size, run the following command on all the nodes:
gfs_grow -Tv <mountpoint>

After you finish

For the latest status on this issue, see bug 215923 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(Red Hat, iSCSI) service clvmd restart fails

A timing issue in the `clvmd` script causes it to fail if you run the `service clvmd restart` command in a GFS 4.0 environment.

About this task

In rare situations where a volume group does not come up after the boot process, you might need to restart `clvmd`. For example, if the iSCSI LUNs are not available before the LVM start-up script runs, then the volume group will not come up after the boot process. To read the volume groups from the iSCSI LUNs, you need to run `vgscan`. In addition, you might need to restart `clvmd`. However, if you are using GFS 4.0 and you run `service clvmd restart`, it fails.

Note: GFS and the underlying cluster LVM layer provide cluster management.

To avoid a `clvmd` restart failure with GFS 4.0, you must first stop `clvmd` and then use the `service clvmd start` command to start it. Here is the workaround for this issue.

Steps

1. Stop `clvmd` by running the command:

```
service clvmd stop
```

2. Start clvmd by running the command:

```
service clvmd start
```

(Red Hat, FC) With Red Hat 6 and QLogic, host occasionally crashes at qla2x00_process_completed_request

When a host is running Red Hat Enterprise Linux 6 and QLogic, the host sometimes crashes at qla2x00_process_completed_request during I/O with controller faults. This issue has been seen with the QLogic QLE2562 FC HBA.

This issue has been fixed in the 6.0.z errata kernel v2.6.32-71.18.1.el6 (qla2xxx driver version: 8.03.01.05.06.0-k8).

Red Hat is tracking this issue in Red Hat Bugzilla 639254.

For NetApp's latest status on this issue, see bug 452179 at Bugs Online.

Related information

Red Hat Bugzilla 639254 - https://bugzilla.redhat.com/show_bug.cgi?id=639254

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, FC) With Red Hat 6 and QLogic, host becomes unresponsive during I/O with faults

On a host that is running Red Hat Enterprise Linux 6, QLogic, and DM-Multipath and also has DM-Multipath configured on some LUNs, the host sometimes hangs during I/O with controller faults and fails to recover.

This issue has been fixed in the 6.0.z errata kernel v2.6.32-71.18.1.el6 (qla2xxx driver version: 8.03.01.05.06.0-k8).

Red Hat is tracking this issue in Red Hat Bugzilla 634948.

For NetApp's latest status on this issue, see bug 448149 at Bugs Online.

Related information

Red Hat Bugzilla 634948 - https://bugzilla.redhat.com/show_bug.cgi?id=634948

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, FC) With Red Hat 6 and Emulex, host occasionally panics at lpfc_abort_handler

On systems running Red Hat Enterprise Linux 6 and Emulex, the host sometimes panics at the `lpfc_abort_handler` during I/O with controller faults.

This issue has been fixed in the native LPFC driver v8.3.5.17, which is available in 6.0.z errata kernel-2.6.32-71.18.1.el6.

Red Hat is tracking this issue in Red Hat Bugzilla 633231.

For NetApp's latest status on this issue, see bug 447023 at Bugs Online.

Related information

Red Hat Bugzilla 633231 - https://bugzilla.redhat.com/show_bug.cgi?id=633231

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, FC) Host running Red Hat 6 or 5.6 with Emulex 8 Gb hangs during fabric faults

Linux hosts running Red Hat Enterprise Linux 6 or Red Hat Enterprise Linux 5 Update 6 with an 8-Gb Emulex adapter using a LightPulse Fibre Channel (LPFC) driver shipped in box and firmware v2.00A3 hang when a fabric fault occurs.

There are two workarounds you can use to avoid this problem. Perform one of the following workarounds:

- Workaround 1: Downgrade the 8-Gb adapter firmware to firmware v1.11A5.
- Workaround 2: Disable the LPFC driver's N-Port-ID Virtualization (NPIV) by setting the `lpfc_enable_npiv` option to 0.

There two methods for setting this value to 0. Use one of the following methods:

- Method 1: Disable the setting and ensure that it remains disabled during system boots by adding the following line to the `/etc/modprobe.conf` file:

```
options lpfc lpfc_enable_npiv=0
```

Now rebuild the initrd image.

- Method 2: During load time, call the `lpfc_enable_npiv=0` option in either the `insmod` command line or the `modprobe` command line.

(Red Hat Enterprise Linux 5 Update 6) Red Hat is tracking this issue in Red Hat Bugzilla 660278.

NetApp provides information on this issue in the following bugs at Bugs Online:

- Red Hat Enterprise Linux 6: bug 485099
- Red Hat Enterprise Linux 5 Update 6: bug 468229

Related information

Red Hat Bugzilla 660278 - https://bugzilla.redhat.com/show_bug.cgi?id=660278

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat) With Red Hat 6, wrong path grouping occurs after storage system fault with I/O

DM-Multipath sometimes wrongly updates the path groups on a host that is running Red Hat Enterprise Linux 6.0. This problem occurs during I/O with a storage system and fabric faults.

When this situation occurs, the `multipath -ll` output may look similar to the following:

```
# multipath -ll 360a98000486e2f65686f5a4b544e3379
360a98000486e2f65686f5a4b544e3379 dm-34 NETAPP,LUN
size=6.0G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=4 status=enabled
|  - 1:0:1:27 sdgg 131:192 active ready running
|+- policy='round-robin 0' prio=1 status=enabled
|  - 1:0:2:27 sdab 65:176 active ready running
|  - 1:0:4:27 sdde 70:192 active ready running
|  - 0:0:3:27 sdhu 134:64 active ready running
|  - 0:0:1:27 sdkw 67:320 active ready running
|+- policy='round-robin 0' prio=4 status=active
|  - 0:0:2:27 sdbp 68:48 active ready running
|  - 0:0:4:27 sder 129:48 active ready running
|  - 1:0:3:27 sdji 8:448 active ready running
```

To prevent this problem use device-mapper-multipath errata v0.4.9-31.el6_0.3. This errata contains a fix for this problem.

Red Hat is tracking this issue in Red Hat Bugzilla 658937.

For NetApp's latest status on this issue, see bug 485101 at Bugs Online.

Related information

Red Hat Bugzilla 660278 - https://bugzilla.redhat.com/show_bug.cgi?id=658937

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat) With Red Hat 6, DM-Multipath fails to update paths during I/O with fabric faults

If your system is running Red Hat Enterprise Linux 6 and DM-Multipath is using the device-mapper-multipath errata v0.4.9.31.el6_0.2, DM-Multipath sometimes fails to update paths during I/O with fabric faults.

When this situation occurs, the `multipath -ll` output may look similar to the following:

```
# multipath -ll /dev/sdce
360a98000486e2f65686f6246516e6859 dm-22 NETAPP,LUN
size=5.0G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|+- policy='round-robin 0' prio=50 status=active
|  |- 3:0:1:19 sdce 69:32 active ready running
|  `-- 2:0:1:19 sdca 68:224 active ready running
|+- policy='round-robin 0' prio=10 status=enabled
|  `-- 2:0:0:19 sdaj 66:48 active ready running
|+- policy='round-robin 0' prio=10 status=enabled
|  `-- 3:0:0:19 sdao 66:128 failed ready running
#
```

To prevent this problem use device-mapper-multipath errata v0.4.9-31.el6_0.3 . This errata contains a fix for this problem.

Red Hat is tracking this issue in Red Hat Bugzilla 681144.

For NetApp's latest status on this issue, see bug 488284 at Bugs Online.

Related information

Red Hat Bugzilla 681144 - https://bugzilla.redhat.com/show_bug.cgi?id=681144

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, iSCSI) With Red Hat 6, iscsi script fails to show correct status

When you are using Red Hat Enterprise Linux 6 and you install the root file system on an iSCSI LUN, the `/etc/init.d/iscsi` script does not show the correct status for the iSCSI daemon. This issue does not affect the iSCSI functionality.

For example, the iSCSI daemon status is displayed as "Stopped" even though all the iSCSI sessions are logged in.

This issue exists in the following Red Hat Enterprise Linux packages:

- Red Hat Enterprise Linux 6.0 `scsi-initiator-utils` package
- Red Hat Enterprise Linux 6.0z `scsi-initiator-utils` package

The fix for this issue occurs in the Red Hat Enterprise Linux 6.1 iSCSI package: iscsi-initiator-utils-6.2.0.872-14.el6 and later versions.

Red Hat is tracking this issue in Red Hat Bugzilla 658428.

For NetApp's latest status on this issue, see bug 466322 at Bugs Online.

Related information

Red Hat Bugzilla 658428 - https://bugzilla.redhat.com/show_bug.cgi?id=658428

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, iSCSI) With Red Hat 5 Update 2, host freezes if multipathd halts

When you are using Red Hat Enterprise Linux 5 Update 2 and the root volume is on an iSCSI multipath device, the host freezes if the multipathd daemon stops running. You must perform a hard reboot to get it started again.

Note: This problem only occurs with Red Hat Enterprise Linux 5 Update 2. It is fixed in Red Hat Enterprise Linux 5 Update 3 and later.

If you are using a root volume on an iSCSI multipath device, at least one patch to the root volume must remain up in order to keep the system running. If the system loses all paths to the root volume and multipathd is not running, the root LUN becomes disconnected, which leads to a system hang.

This problem can occur during controller faults. When it happens, all the paths go down for a period of time. If multipathd is not running, the root volume becomes unavailable, which leads to a system freeze.

The problem has also been seen when the system is booting up. If the iscsid service starts before the multipathd service, there is a small window of time where the host might lose its connection with the iSCSI root volume, which leads to a system freeze.

If the multipathd daemon continues to run, then the `queue_if_no_path` feature will hold I/O on the root volume until some paths are reinstated. At that point, the root volume becomes accessible again.

To prevent this problem, modify the `multipathd` `init` script to ensure that it continues to run. To do this, install the patch that is available from Red Hat.

For the latest status on this issue, see bug 295559 at Bugs Online.

Related information

Patch - <https://bugzilla.redhat.com/attachment.cgi?id=314924>

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, iSCSI) With Red Hat 5 Update 2, host intermittently stops during SAN boot

There is a problem with the `iscsid` and `multipathd` sequence during the boot process on Red Hat Enterprise Linux 5 Update 2. Because the `iscsid` service is started before the `multipathd` service, there is a small window when the host sometimes loses its connection with the iSCSI root volume, which leads to a system freeze.

About this task

The workaround for this problem is to configure the system after the installation, as shown in the following steps.

Steps

1. During the first boot, switch to single user mode by modifying the grub parameter
2. Add `single` to the kernel parameters:

```
kernel /vmlinuz-2.6.18-92.el5 ro root=/dev/VolGroup00/LogVol00 rhgb  
quiet single:
```
3. Create the `/etc/multipath.conf` file as described in the *Linux Host Utilities 5.3 Installation and Setup Guide*.
4. Modify the `/etc/iscsi/iscsid.conf` file as described in the *Linux Host Utilities 5.3 Installation and Setup Guide*.
5. Make sure that the `multipathd` service is started before the `iscsid` service.

To do this you must modify the `/etc/init.d/multipathd` script and change the start priority from `SXX` to `S06`. This assumes that the default iSCSI installation sets the `iscsid` start priority to `S07`. If not, you need to ensure that the start priority of `multipathd` is less than the start priority of `iscsid`.

You can use the `chkconfig` command to do this.

6. Run the following commands in the order shown:

```
chkconfig multipathd off  
chkconfig multipathd on
```

7. Edit the `/etc/fstab` file and add the `_netdev` mount option for the root device.

Example

Specify the line as shown in the following example:

```
/dev/mapper/mpath0p1 / ext3 _netdev,defaults 1 1
```

8. Reboot the system.

After you finish

For the latest status on this issue, see bug 295559 at Bugs Online.

Related information

Patch - <https://bugzilla.redhat.com/attachment.cgi?id=314924>

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, iSCSI) With Red Hat 4 Update 3, system hangs during shutdown, reboot

The Linux system stops responding when shutting down or rebooting. The `/etc/init.d/iscsi stop` command might also report errors. This problem occurs only in systems running DM-Multipath with iSCSI and Red Hat Enterprise Linux 4 Update 3.

About this task

This is a known problem with the shutdown sequence in Red Hat Enterprise Linux 4 Update 3. If the system is not responding, disconnect power to force the shutdown to complete.

To prevent shutdown problems, complete the following steps before shutting down or rebooting.

Steps

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

After you finish

For more information, see Red Hat Bugzilla 177355.

For NetApp's latest status on this issue, see bug 183045 at Bugs Online.

Related information

Red Hat Bugzilla 177355 - <https://bugzilla.redhat.com/bugzilla/>

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

(Red Hat, iSCSI) With Red Hat 4 Update 2, first LUN must be LUN 0

If the first LUN mapped to a host has an ID other than 0, the `iscsi-ls -l` command shows incorrect information. Adding a LUN 0 later does not resolve the problem. This problem occurs when you're using iSCSI and Red Hat Enterprise Linux 4 Update 2.

To resolve the problem, map a LUN 0 to the host and reload the HBA driver.

Note: Adding a LUN 0 without reloading the driver does not resolve the problem.

To prevent this problem, when you map LUNs to the host, make sure the first LUN is always LUN 0. By default, Data ONTAP does this when you do not explicitly specify the LUN identifier.

For the latest status on this issue, see bug 164689 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(Red Hat, iSCSI) With Red Hat 4, continuous ISID Rule messages displayed

In some cases when you are running Red Hat Enterprise Linux 4, the storage system console displays continuous ISID Rule messages.

The messages are similar to the following:

```
Wed Jun 11 13:42:47 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 1
Wed Jun 11 13:42:48 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator iqn.1987-05.com.cisco:rachman at IP addr
10.10.10.10
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
session 0
Wed Jun 11 13:42:50 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
New session from initiator iqn.1987-05.com.cisco:rachman at IP addr
10.10.10.10
Wed Jun 11 13:42:51 GMT [ss1: iscsi.notice:notice]: ISCSI: iswta,
ISID Rule: new connection from same initiator, shutting down old
```

To prevent this problem, add the following line to the `/etc/iscsi.conf` file. This line must be at the *beginning* of the file, before any "DiscoveryAddress" entries.

Continuous=no

Now restart the iSCSI service by running the command:

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

(Red Hat) With Red Hat 4 and a large number of LUNs, DM-Multipath sometimes fails to configure devices

On occasion, DM-Multipath fails to create all the maps for newly discovered LUNs. This happens when a large number of LUNs, for example, 512, are mapped to a host running Red Hat Enterprise Linux 4 series.

Normally, when a new LUN is mapped and discovered on a host, DM-Multipath automatically creates a map for it. However, if a large number of LUNs are mapped to a host, DM-Multipath sometimes fails to create maps for all the LUNs.

If you encounter this problem, you must manually create the missing maps. After the host discovers the LUNs, run the `multipath -ll` command.

For NetApp's latest status on this issue, see bug 205600 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(SUSE) Known issues when running SUSE Linux Enterprise Server

The following issues might occur when you are running SUSE Linux Enterprise Server. In many cases, workarounds to avoid these issues have been identified.

(SUSE, FC) With SUSE 11 SP1, DM-Multipath sometimes fails to update maps after I/O fault

In FC environments where a SUSE Linux Enterprise Server 11 SP1 host is configured with multiple LUNs and each LUN has multipath paths, DM-Multipath sometimes fails to update its maps during I/O with controller faults. In addition, some paths in the multipath maps remain in inconsistent state.

For example, this problem has been seen in situations where multiple faults occurred with a configuration of 10 LUNs, each of which had 8 paths for a total of 80 SCSI devices.

This issue has also been seen in configurations using a SAN boot LUN. In some cases, the SAN boot LUN becomes unresponsive after the controller faults.

The problem occurs because of issues with the ALUA handler and the multipath-tools package.

To prevent this problem, take the following two actions:

- Install the maintenance update that was released Jan. 13, 2011, which fixes these problems. The following packages have been confirmed to have the fixes that resolve this problem:
 - kernel-default-2.6.32.27-0.2.2
 - device-mapper-1.02.27-8.17.20
 - multipath-tools-0.4.8-40.25.1
 - kpartx-0.4.8-40.25.1
- In addition, use the `tur` option as the value of the `path_checker` parameter in the `multipath.conf` file. The `tur` option replaces the `directio` option as the default value for this parameter.

Novell is tracking this issue in Bugzilla 600043.

For NetApp's latest status on this issue, see bug 414050 at Bugs Online.

Related information

*Novell Bugzilla 600043 - https://bugzilla.novell.com/show_bug.cgi?id=600043
Bugs Online - now.netapp.com/NOW/cgi-bin/bol/*

(SUSE, FC) With SUSE 10.3 and Emulex 4 Gb, FC ports go offline during controller faults

In environments using SUSE Linux Enterprise Server 10.3 with an Emulex 4-Gb, FC HBA, FC ports go offline during storage controller faults. This issue occurs because of a problem in driver version 8.2.0.48.2p with the msi interrupt handler.

There are two workarounds for this issue:

- Upgrade to the SLES10SP4 kernel and use inbox Emulex driver version 8.2.0.92.1p or later.
- If you are using driver version is 8.2.0.48.2p, change the value of the driver parameter `lpfc_use_msi` to 0.

The default value of this parameter is 2.

To set the driver parameter `lpfc_use_msi` to 0, complete the following steps:

1. In the `/etc/modprobe.conf.local` file, append the following line:

```
options lpfc lpfc_use_msi=0
```

2. Re-create the `initrd` by entering the following command:

```
mkinitrd
```

Note: It is a good practice to make a backup copy of the current `initrd` before running the `mkinitrd` command.

3. Reboot the host.
4. Verify that the value of the driver parameter `lpfc_use_msi` by entering the following command:

```
cat /sys/class/scsi_host/host<adapter_number>/lpfc_use_msi
```

For the latest status on this issue, see bug 516749 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(SUSE, iSCSI) Online resize of file system not supported

When you are using SUSE Linux Enterprise Server and iSCSI, resizing a file system while it is online is not supported.

To resize a file system, you must unmount the file system and then use the `resize` utility to extend or shrink the size of the file system.

However, if you resize a LUN on a NetApp storage system, you can display information about its new size on the host by rescanning the iSCSI sessions.

For more information, see Novell bugzilla number 386325.

(SUSE, iSCSI) One or more ext3 file systems not mounted after reboot

After a reboot of the host, the ext3 file systems are not mounted. This issue is seen on hosts running SUSE Linux Enterprise Server and iSCSI.

An error message similar to the following is displayed:

```
fs type ext3 not supported by kernel
```

This problem occurs when the module that supports ext3 file systems was not loaded before the system tried to mount the file systems.

To work around this problem, run the `iscsi-mountall` command to mount the file systems.

(SUSE, iSCSI) sanlun does not display all iSCSI LUNs

On hosts running SUSE Linux Enterprise Server and iSCSI, the `sanlun` command does not display all iSCSI LUNs mapped to the initiator.

There are two possible explanations for this problem. The workaround you use depends on which situation caused the problem.

- More than 32 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), the `sanlun` command does not list all the LUNs.

The workaround in this situation is to create more `/dev/sg*` entries. You can use the Linux `mknod` command to do this. For more information on this command, see the man page.

The following is an example of using the `mknod` command:

```
mknod --mode=660 /dev/sg32 c 21 32
```

- Storage system is unreachable for a long time

If storage system has been unreachable for a time period that 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 it.

To recover from this situation, run the `sanlun lun show all` command again.

Then make sure that the LUNs listed match the devices listed in the `/dev/iscsi` directory or in the output of the `iscsi -ls` command.

(Veritas) Known issues when running the Host Utilities with Veritas

The following issues might occur when you are running the Host Utilities with Veritas Storage Foundation. In many cases, workarounds to avoid these issues have been identified.

(Veritas) With Red Hat 4 Update 5, paths not reenabled after giveback

When you are running Red Hat Enterprise Linux 4 Update 5 and Veritas Storage Foundation Cluster File System for Oracle RAC (SFCFSRAC) 5.0 MP3, some paths are not reenabled after the giveback process completes. This problem happens intermittently when you have 1024 or more paths under Veritas DMP.

This issue has not been seen when there are fewer than 1024 paths under Veritas DMP control.

The workaround for this problem is to run the `vxdctl enable` command after the giveback completes. Doing this re-enables all the paths.

For the latest status on this issue, see bug 322127 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(Veritas) I/O delay seen on systems using iSCSI

On hosts running Veritas Storage Foundation 5.1 with the iSCSI protocol, the resumption of I/O might be delayed after an iSCSI session failure.

The iSCSI parameter `node.session.timeo.replacement_timeout` is set to 120 seconds, which is the default. When a cable failure or iSCSI port block occurs, I/O is held at the iSCSI layer for 120 seconds. Testing has shown that, in some cases, there is a delay of more than 120 seconds before I/O resumes on other active VxDMP paths.

This delay is not seen in takeover/giveback scenarios, where I/O resumes as soon as the takeover completes.

For the latest status on this issue, see bug 399840 at Bugs Online.

Related information

[Bugs Online - now.netapp.com/NOW/cgi-bin/bol/](http://now.netapp.com/NOW/cgi-bin/bol/)

(Veritas) With Oracle Enterprise Linux 5.4 and Emulex, Veritas volumes go offline

When you are using Oracle Enterprise Linux 5.4 with Emulex 4-Gb FC HBA, FC ports go offline intermittently due to a bug in driver version 8.2.0.48.2p with msi interrupt handler. As a result, Veritas volumes go offline if there is a coincident storage controller fault.

About this task

To work around this issue, do one of the following:

- Update the Oracle Enterprise Linux 5.4 with the latest kernel so that it is using Emulex driver version 8.2.0.48.3p or later.
- If the driver version is 8.2.0.48.2p, change the value of the driver parameter `lpfc_use_msi` to 0 (the default value of this parameter is 2).

The following steps tell you how to set the Emulex driver parameter `lpfc_use_msi` value to 0.

Steps

1. In the `/etc/modprobe.conf` file, append the following line:

```
options lpfc lpfc_use_msi=0
```

2. Re-create the `initrd` by running the following command:

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

Note: It is a good practice to make a backup copy of the current `initrd` before running the `mkinitrd` command.

3. Reboot the host.
4. Verify the value of `lpfc_use_msi` by using the following command:

```
cat /sys/class/scsi_host/host<adapter_number>/lpfc_use_msi
```

(Veritas) With SUSE Linux and iSCSI, volumes might enter disabled state after controller fault

When you are using SUSE Linux Enterprise Server with the iSCSI protocol, Veritas Storage Foundation 5.1 volumes might enter a disabled state after a storage controller fault occurs.

An issue has been found with SUSE Linux Enterprise Server with the latest kernels where the SCSI layer does not honor the timeout value when handling I/O errors during iSCSI session failures. The I/O error is immediately propagated up to the multipathing layer (VxDMP), which disables the path.

If the I/O error is seen on all the VxDMP paths that belong to same LUN, VxDMP sends the error to the volume manager. The volume manager then disables all the volumes.

To work around this issue, you must set the VxDMP parameter `dmp_lun_retry_timeout` to 300.

NetApp engineering is working with Novell to resolve the issue. For more information, see Bugzilla ID# 645616.

For the latest status on this issue, see bug 442853 at Bugs Online.

Related information

Bugzilla ID# 645616 - https://bugzilla.novell.com/show_bug.cgi?id=645616/

Bugs Online - now.netapp.com/NOW/cgi-bin/bol/

Viewing current and fixed product bugs

NetApp provides an online tool that lets you search for the most current information about a known bug. You can also get a list of current or fixed bugs for a particular product.

You can use the Bugs Online tool to determine if new information has been uncovered about a bug or if the bug has been fixed.

You can search Bugs Online in the following ways:

- Enter the bug number.
- Enter keywords related to the bug.
- Select a software product from the list in the Bug Type(s) field or the Product Type(s) field. Some products are listed in one field and some in the other field. Be sure to check both to find the product you want.

Keyword combinations and Bug Type(s) that you might want to use include:

- FCP - Linux
- ISCSI - Linux

You can also set Bugs Online to automatically display all bugs of interest to you. Use the New Bug Alerts Profiler at the bottom of the Bugs Online page to set bug searching criteria that matches your requirements.

Related information

[Bugs Online & Release Tools - now.netapp.com/NOW/cgi-bin/bol](http://now.netapp.com/NOW/cgi-bin/bol)

NetApp product documentation

You can access documentation for all NetApp products on the NOW site at support.netapp.com.

In addition to product documentation, NetApp provides other product information, such as technical reports and white papers.

How to search for documentation on the NetApp Support Site

NetApp provides a variety of documentation on the support.netapp.com site that can be helpful in working with the Host Utilities.

You can use the NOW Natural Language PowerSearch tool, located on the Search and Support page on the support.netapp.com site, to search through the following information sources:

- Product documents
- Knowledgebase solutions
- Bugs Online
- NOW forums
- Tools and other documentation
- Partner knowledgebases

Related information

[Search and Support -- now.netapp.com/eservice/SupportHome.jsp](http://now.netapp.com/eservice/SupportHome.jsp)

Links to NetApp product documentation

The Access and Order Product Documentation page contains pointers to the most current NetApp documentation for Host Utilities and storage systems.

Links to specific Host Utilities documentation are in the SAN/IPSAN Information Library and available from the NetApp NOW Description Pages for each Host Utilities product.

The following table describes some of the documents on the NetApp NOW Web site that contain current information about host and storage system requirements, supported configurations, and best practices, as well as procedures for installing and managing your storage environment.

If you need more information about...	Go to...
The latest supported configurations	• NetApp Interoperability Matrix • <i>System Configuration Guide</i>
Configuring and managing the storage system	• SAN/IP SAN Best Practices Index • The following documents for your Data ONTAP release in the Data ONTAP Information Library • The <i>Data ONTAP Software Setup Guide</i> • The <i>Data ONTAP Block Access Management Guide for iSCSI and FC</i> • <i>Data ONTAP Release Notes</i> • The <i>Data ONTAP Commands: Manual Page Reference, Volume 1</i> • <i>Data ONTAP Commands: Manual Page Reference, Volume 2</i> • The <i>Data ONTAP Upgrade Guide</i>
Supported Fibre Channel SAN topologies	The <i>FC and iSCSI Configuration Guide</i>
Verifying storage systems and environments	<i>Site Requirements Guide</i>
Configuring your host for the Host Utilities	The Host Utilities documents are in the SAN/IPSAN Information Library

Related information

Access and Order Product Documentation page - <http://now.netapp.com/NOW/knowledge/docs/docs.cgi>

SAN/IPSAN Library Information Library - <http://now.netapp.com/NOW/knowledge/docs/san/>

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

System Configuration Guide - <http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/syscfg/>

SAN/IP SAN Best Practices Index - http://now.netapp.com/NOW/knowledge/docs/bpg/san_ipsan/

Data ONTAP Information Library - http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

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

Site Requirements Guide - <http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/site/index.htm>

Contact Technical Support

If you have a question that has not been resolved by these Release Notes, contact NetApp Technical Support.

Things to check before you call Technical Support

Before you call Technical Support, there are several things you can try in effort to solve the problem yourself.

They include:

- Go through these *Release Notes* to see if they contain information on your problem.
- Check all cables to make sure that they are connected properly.
- If you are using switches, check the power switches to make sure that the system is turned on. You should also make sure the system components were turned on in the correct order.
- Review the troubleshooting information in the documentation for this product.

Information you need to provide to Technical Support

Before you contact Technical Support, you need to gather information about your system and your problem.

You should have the following information available when you contact Technical Support:

- Your contact information.
- A list of all the NetApp products you are using.
- All the error messages from the system.
- Information on your system setup.
- Your licensing information.

How to contact NetApp Technical Support

You can contact NetApp Technical Support by telephone or from the NetApp NOW site.

To contact Technical Support by telephone, use the following numbers:

- 888 4-NETAPP (US and Canada)
- 00.800.44.NETAPP (EMEA/Europe)
- +800.800.80.800 (Asia/Pacific)

To contact Technical Support from the Web, use the following address:

- <http://now.netapp.com>

For a detailed listing of Technical Support phone numbers, go to the NetApp Support page at <http://www.netapp.com/us/support/ngs-contacts.html>

Related information

Technical Support - <http://www.netapp.com/us/support/ngs-contacts.html>

Copyright information

Copyright © 1994–2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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

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

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

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

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

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, 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, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, 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.

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

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

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

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

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

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

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support Telephone: +1 (888) 4-NETAPP

