



Product Alert 16

Product	All versions of StorNext® Storage Manager
Summary	Possible loss of data on tape
Date	May 2010

Overview

This alert applies only to systems running StorNext Storage Manager (SNSM). It does NOT apply to systems running only the StorNext file system.

Tapes in drives shared between systems may unexpectedly rewind, which may cause data loss. The problem has been observed under the following conditions:

The SCSI device reservation mechanism in SNSM is disabled

AND

- 1 Systems running RedHat Enterprise Linux 4 (RHEL4) updates 1 through 7 have access to shared tape drives. The problem is fixed in RHEL4 update 8 (RedHat bug 468680).

OR

- 2 Systems running SuSE Linux Enterprise Server 10 (SLES10) service pack 2 (sp2) or sp3 have access to shared tape drives.

The problem has not been observed on the initial SLES10 release and SLES10 sp 1. SLES10 sp4 has not been released as of May 2010. It is not known at this time when a fix will be available in a SLES10 service pack.

The problem has not been observed on systems running RedHat Enterprise Linux 5.

Cause

For RHEL4:

A script named `/etc/rc.d/rc.sysinit` runs when RedHat Linux boots. On RHEL4 that script runs an executable called `/sbin/kmodule`. Also, `/usr/sbin/kudzu` may be run during system boot or by a system administrator. Both executables share a library which may reference `/dev/st<x>` (rewind-on-close) tape devices, causing tapes in those drives to rewind.

For SLES10:

A script named `/etc/rc.d/boot.udev` runs when SLES10 boots. That runs `/sbin/udevtrigger`, which references `/etc/udev/rules.d/60-persistent-storage.rules`. The latter contains rules which cause `/sbin/scsi_id` to run. `/sbin/scsi_id` references `/dev/st<x>`, which will causes tapes in the referenced drives to rewind.

Risks Common to All SNSM Servers

SCSI device reservations provide substantial protection from drive access by other systems. Systems running SNSM versions 3.x or 4.x on which SCSI reservations are disabled ("`FS SCSI RESERVE=none;`" is set in `/usr/adic/TSM/config/fs_sysparm_override`) may be at risk of unintended tape rewinds. All SNSM versions prior to 3.0 are at risk because "SCSI device reservations" are not available.

Symptoms

A read label failure is detected when an affected tape is mounted, though read label failures may also occur for other unrelated reasons. Look for the following in the `/usr/adic/TSM/logs/tac` log files:

```
Apr 9 3:57:32 meta02 sntsm fs_drv[7403]:  
E1200(7)<1807964952>:mdt2tap1657: Invalid label data +000000  
C0239A95 2A000000 01000000 09000000 *.#..*.....*
```

...

```
Apr 9 3:57:32 meta02 sntsm fs_drv[7403]:  
E1006(6)<1804952>:mscudrv552: Drive: /dev/sg35.  
Read Label Failure. Medium: ABC123. errno=0  
Reason: Label info was inconsistent with label type.
```

...

```
Apr 9 3:57:33 meta02 sntsm fs_resource[1209]:  
E1200(7)<1804952>:msc9req6811: {1194}: Media label validation  
failed. Changing medium state to unavailable.
```

Solution

Using SCSI device reservations is highly recommended on ALL systems operating as SNSM servers. If SCSI reservations are disabled, it is important to resolve any underlying issues so that SCSI reservations may be enabled.

Systems running SNSM on RHEL4 updates 1 through 7 should upgrade to RHEL4 update 8 or upgrade to RHEL5. Some architectures and OS combinations are not supported by various versions of StorNext. Please examine the StorNext Release Notes or Installation Guides to identify any upgrade requirements that may exist.

Please see the Workarounds section if upgrading is not possible, or if SLES10sp2 or SLES10sp3 are in use. As an extra precaution, it might be desirable to implement the workaround on all systems running any RHEL4 update or any SLES10 service pack.

Workarounds

Removing the offending references to rewind tape devices from the udev rules will prevent the problem for both RHEL4 and SLES10.

Caution: The workaround requires modifying files used during system boot. This procedure is best initially performed/tested on a non-production system because an error may cause failures during boot.

RHEL4 Workaround

Prevent the `/dev/st*` device files from being created, using `/dev/xst*` instead. Perform the following:

1 Save the original files

```
# cp /etc/udev/rules.d/50-udev.rules /etc/udev/rules.d/50-udev.rules.bak
# cp /etc/udev/permissions.d/50-udev.permissions /etc/udev/permissions.d/50-udev.permissions.bak
```

2 Modify (edit) `/etc/udev/rules.d/50-udev.rules` to add the following:

```
KERNEL="st[0-9]*", BUS="scsi", NAME="xst%n"
```

3 Modify `/etc/udev/permissions.d/50-udev.permissions` to change the following:

```
st*:root:disk:0660
```

to look like

```
xst*:root:disk:0660
```

For additional protection from undesired access on the local system, Quantum recommends changing the 0660 above to 0000.

4 Test, if possible. After a reboot, most system configurations will not have any files named `/dev/st[0-9]*`. Those files should be named `/dev/xst[0-9]*` with the permissions specified above.

5 If any applications use the `/dev/st` device files, either the applications must be modified to use the new `/dev/xst` device file names or the required device files must be returned to their original (`/dev/st`) names. `/etc/rc.d/rc.local` might be a good choice for making the latter change, if necessary.

SLES10 Workaround

Save a copy of the original file (`/etc/udev/rules.d/60-persistent-storage.rules`) and comment out all offending lines.

- 1** `#cp /etc/udev/rules.d/60-persistent-storage.rules /etc/udev/rules.d/60-persistent-storage.rules.save`
- 2** Change the following 2 lines (the portion beginning with "IMPORT" is actually on the same line as "KERNEL"):

```
KERNEL=="st[0-9]*", BUS=="scsi", ENV{ID_SERIAL}=="",  
IMPORT{program}="/sbin/scsi_id -g -x -s %p -d $tempnode"  
  
KERNEL=="st[0-9]*", BUS=="scsi", ENV{ID_SERIAL}=="",  
IMPORT{program}="/sbin/scsi_id -g -x -a -s %p -d $tempnode"
```

to (add a "#" in column 1)

```
#KERNEL=="st[0-9]*", BUS=="scsi", ENV{ID_SERIAL}=="",  
IMPORT{program}="/sbin/scsi_id -g -x -s %p -d $tempnode"  
  
#KERNEL=="st[0-9]*", BUS=="scsi", ENV{ID_SERIAL}=="",  
IMPORT{program}="/sbin/scsi_id -g -x -a -s %p -d $tempnode"
```

- 3 Test, if possible, by positioning a tape beyond beginning-of-tape and running `/etc/udevtrigger`. The tape position should remain unchanged. This test must be performed while no other tape drives are in use.

Caution: Files in `/dev/udev` may be overwritten during installation of the "udev" software package and during operating system updates or upgrades, in which case these modifications will need to be applied again.

Contacting Quantum

More information about StorNext is available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

For further assistance, contact the Quantum Technical Assistance Center:

North America	+1 800-284-5101 Option 5
EMEA	00800 9999 3822
Online Service and Support	www.quantum.com/OSR
World Wide Web	www.quantum.com/ServiceandSupport

(Local numbers for specific countries are listed on the Quantum Service and Support Website.)