



Product Alert 16

Product	All versions of StorNext® Storage Manager
Summary	Possible loss of data on tape
Date	May 2008

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 if one or more of those systems are running RedHat Enterprise Linux 4 (RHEL4). If the drives are in use by another system (which could be running ANY operating system) when the rewind occurs, data loss may occur. Systems running SNSM version 2.x in a High Availability (HA) configuration in conjunction with RHEL4 are particularly vulnerable to this problem.

See also: Product Alert 6 (“Running KUDZU on a Linux Machine Can Cause Data Loss”)

Cause: A script called `/etc/rc.d/rc.sysinit` runs when RedHat Linux boots. On RHEL4, that script runs an executable named `/sbin/kmodule`, which shares common code with `kudzu`. `/sbin/kmodule` may access `/dev/st<x>` (rewind-on-close) tape devices, causing tapes in those drives to rewind. Which tape drives, if any, are accessed by `/sbin/kmodule` is configuration dependent.

If a system sharing those drives is writing to a tape when the rewind occurs, then data, including the label at beginning of tape, may be overwritten. On most tape drives, data beyond the newly written EOD tape mark can no longer be accessed.

SNSM version 3.0 or later running on all supported versions of Linux utilize the “scsi reserve” drive reservation mechanism, which provides substantial protection from access by another system. SNSM versions 2.x and those running on Solaris in a StorNext HA configuration do not utilize “scsi reserve” and are vulnerable to rewinds caused by any rebooting RHEL4 system which shares the tape drives.

Symptoms: A read label failure is detected when an affected tape is mounted, though that may not be unique to this particular problem. 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 *.#..* .....
```

...

```
+000480 2CE9FFBF 7F000000 5FE9FFBF 7F000000 *,....._.....*
```

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 3:57:32 meta02 sntsm fs_resource[1209]:
E1202(9)<1804952>:msc9utl607: {1194}: Child Process(0) exited with status
255

Apr 9 3:57:33 meta02 sntsm fs_resource[1209]:
E1002(2)<1804952>:msc9req6770: {1194}: Could not obtain label
information for media ID ABC123.

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: Ensure that /sbin/kmodule cannot access any rewind-on-close tape devices. This modification applies only to systems running RHEL4. No modifications are needed for SuSE Linux or for other RedHat versions because they do not have /sbin/kmodule.

This solution is based on bullet 2) at the bottom of the following web page:

<http://forums11.itrc.hp.com/service/forums/questionanswer.do?admit=109447626+1209044618243+28353475&threadId=1048563>

<p>Caution: This solution requires modifying files used during system boot. This procedure is best initially performed/tested on a non-production system because an error may cause a system to not boot.</p>
--

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 (edit) /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 names. `/etc/rc.d/rc.local` might be a good choice for making the latter change, if desired.

Caution: Files in `/etc/udev` may be overwritten during an upgrade of RHEL4 to a newer update, in which case this modification will need to be applied again.

Recommendation (all operating systems): If tape drives used by StorNext Storage Manager are connected via a fibre switch, Quantum recommends that the switch be zoned to allow only the system(s) running SNSM to have access to the tape drives. Otherwise, ALL systems running RHEL4 with tape drive access (not just those running SNSM), will require the above modification. That will help to ensure that a “rogue” system cannot cause a tape to be rewound unexpectedly.