



Product Alert 15

Product	All versions of StorNext® prior to release 3.1.2
Summary	Possible problems writing data to tape
Date	April 2008

Overview: This alert applies only to systems running StorNext Storage Manager. It does NOT apply to systems running only the StorNext file system.

In rare instances, problems in the “scsi generic” device driver included with RedHat Enterprise Linux 5 (RHEL5) and SUSE Linux Enterprise Server 10 (SLES10) might prevent data from being written to tape. This includes RHEL5 update 1 and SLES10 service pack 1, which at the time this bulletin was written were the most recent versions.

The “scsi generic” device driver used by StorNext is different than the “scsi tape” device driver commonly used by other utilities (e.g., `/bin/tar`). Devices associated with the former are “`/dev/sg,`” while devices associated with the latter are `/dev/[n]st.`”

Cause: In some instances the scsi generic device driver passes an incorrect length field to the tape drive. Block lengths not a multiple of 512 bytes and larger than the value defined in `/proc/scsi/sg/def_reserved_size` (default 32768 bytes) have been known to encounter this problem.

The firmware in some tape drives will reject commands containing invalid length fields. Device driver differences cause this event to be reported to StorNext in varying ways. In the case of the LSI drivers on SLES10, StorNext is not able to recognize this event as an error. For other drivers on SLES10 and for all known drivers on RHEL5, StorNext recognizes the event as an error, but cannot correct the problem.

Symptoms: In rare instances, due to the device driver problem, the last block of a StorNext file cluster cannot be written to tape. This might have differing implications depending on which tape library, tape drive, i/o interface, device driver and operating system are used.

The problem has been observed on LTO tape drives with IBM firmware. Technically, this firmware does adhere to the scsi tape specification (the error is in the device driver). The problem has not been observed on LTO tape drives with HP firmware.

In some instances, tape libraries acting as fibre bridges have been known to correct the length problem before the command is sent to the tape drive. In such

instances, the problem does not occur. One such library is the Quantum Scalar I2000. Unfortunately, a complete list of these libraries is not available.

Implications: 1) When a tape drive is controlled by a Qlogic, Emulex, or LSI fibre hba/driver on RHEL5 or a Qlogic or Emulex fibre hba/driver on SLES10, the non-recoverable tape i/o error described above may occur. Sometimes, depending on the specific mixture of files and file sizes, the files may never be writable to tape. If this occurs, StorNext will not truncate the on-disk version(s).

Messages such as the following in `/usr/adic/TSM/logs/tac/tac_00*` indicate the problem outlined in 1):

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1200(7)<1735202635>:fsScsi1913:  
Received check condition with no error data. op=0A
```

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1200(7)<1735202635>:mdm3ioc470:  
Flush residue write to destination failed: errno 0
```

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1200(7)<1735202635>:mdm3mov4166:  
Unable flush all of residue buffer to destination.
```

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1200(7)<1735202635>:mdm3mov5550:  
Write error occurred - marking media suspect.
```

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1201(8)<1735202635>:rdb1mediadir1197:  
Medium 000235 was marked as suspect.
```

```
Dec 18 09:32:05 altix2 sntsm fs_fmover[30144]: E1202(9)<1735202635>:mdm3utl1757: END  
copy of File 'num:2 key:23 seg:1' Move_status:15(E_DEST_WR_MS)
```

2) For StorNext versions prior to 3.1 build 20, when a tape drive is controlled by an LSI fibre hba/driver running on SLES10, it is possible that StorNext would not recognize that a command was rejected. In such instances, it is possible the on-disk files contained in the file cluster could be truncated. The problem might not be noticeable until a file is retrieved from tape. In such an instance some data, typically the last block (up to 512K bytes), will be lost.

Messages such as the following in `/usr/adic/TSM/logs/tac/tac_00*` indicate the problem outlined in 2):

```
Dec 15 18:49:30 altix2 sntsm fs_fmover[6592]: E1202(9)<1731970567>:mdm3utl1736: BEGIN  
copy of File '/num:0 key:6238263 seg:1 ver:1 name :_Handle_Only_/key.6238263/class  
index.14/', Copy_type: 2(TAPE2DISK), Copy_mode: 1(DUPLICATE), Srce '/med:402933  
ndx.gen:2987,0 dev:/dev/sg55 cpy:1/', Dest '/DISK/'
```

```
Dec 15 18:49:30 altix2 sntsm fs_fmover[6592]: E1202(9)<1731970567>:mdt1dev3596: Drive /  
dev/sg55 took 0 seconds to open and come ready. Nov 15 18:50:27 cfs sntsm  
fs_fmover[6592]: E1201(8)<1731970567>:fsScsi1371: Check condition: op=08h key=00h  
asc=00h ascq=01h FILEMARK DETECTED
```

```
Dec 15 18:50:27 altix2 sntsm fs_fmover[6592]: E1200(7)<1731970567>:mdm3ioc812:  
mdm3ioc_read_sync: Not enough data read. read 0 of 524288, requested 249008
```

```
Dec 15 18:50:27 altix2 sntsm fs_fmover[6592]: E1200(7)<1731970567>:mdm3ioc1279:  
Read failed after 184320396 bytes
```

```
Dec 15 18:50:28 altix2 sntsm fs_fmover[6592]: E1202(9)<1731970567>:mdm3utl1757: END  
copy of File '/num:0 key:6238263 seg:1' Move_status:10(E_SRCE_RD_MS)
```

```
Dec 15 18:50:30 altix2 sntsm fs_fcopyman[16164]: E1200(7)<1731970567>:mdm1msm504:  
{324658}: Encountered source error during file transfer
```

Solutions: 1) Upgrade to StorNext 3.1.2 (available May, 2008,) or later

OR

2) For versions of StorNext prior to 3.1.2, modify a system startup script to set the following scsi variables as follows to ensure they are set after every reboot:

```
echo 1 > /proc/scsi/sg/allow_dio
```

```
echo 524288 > /proc/scsi/sg/def_reserved_size
```

For SLES10, adding the above commands to `/etc/init.d/boot.local` might be a good choice. For RHEL5, adding the above commands to `/etc/rc.d/rc.local` might be a good choice.

<p>Caution: Some system startup scripts might be replaced by product or operating system upgrades (e.g., installing a new update or service pack,) in which case these modifications could be lost.</p>
--

“allow_dio=1” requests the scsi generic device driver to perform direct i/o rather than buffered i/o, when possible. This avoids executing the problem code. In instances when buffered i/o is required, “def_reserved_size=524288” requests the driver to use 512KB blocks, which is the block size used by StorNext (except for the last block at times). This also avoids executing the problem code. There are no known instances where these settings have been found to impact performance.

The “echo” commands in 2) above may be run as root on the command line to immediately enable the settings, which will remain in place until the system reboots. These commands are run automatically during startup of the TSM component of StorNext Storage Manager versions 3.1.2 and later.