



Product Bulletin 25

Product	StorNext File System
Summary	Possible data access problems with LUNs using 4K byte sectors
Date	July 2008

Problem

StorNext file systems comprised of disks/LUNs with 4K byte sectors may experience data access problems or performance problems under unusual circumstances. Systems using disks/LUNs with 512 byte sectors (which is the most common configuration,) are not affected.

To determine if any disks/LUNs with 4K byte sectors are in use by StorNext, examine the StorNext file system configuration (.cfg) files on the metadata controllers. "SectorSize 4096" in the "DiskType" sections would identify such a configuration.

Symptoms

A short read (where fewer bytes of data than requested are returned) may occur, even if the file contains a sufficient amount of data to satisfy the read. This is most likely to occur with reads following an "lseek" that specifies an unusual offset within a file.

Messages such as the following may also be observed in Linux system logs when the problem has occurred:

```
Jun 12 13:30:41 devrandom kernel: sdad : block=151483
```

```
Jun 12 13:30:41 devrandom kernel: sd: Bad block number requested  
sd_init_command: disk=sdad, block=151995, count=512
```

Not every occurrence of these messages is an indication of this problem. There are many other reasons (unrelated to StorNext) that "Bad block number" messages could be logged.

Some StorNext configurations will attempt to retry the I/O that caused this particular "Bad block number" event to be logged. In that case, the above messages may continue to repeat in the system log, which may cause significant performance issues until the problem is corrected.

Workaround

Mount file systems using an additional mount option: “memalign=4k”.

<p>Note: StorNext does not support 4K byte sectors on all operating systems. No sector sizes other than 512 bytes and 4K bytes are supported on any StorNext configuration.</p>
--