



Product Alert 28

Product	StorNext® File System
Summary	Data corruption may occur after bringing UP a stripe group that is marked DOWN.
Date	June 2009

Overview

Corruption may occur after bringing UP an existing stripe group previously marked DOWN by the use of the `cvadmin` command or by setting "Status DOWN" in the file system `.cfg` file. This issue applies to StorNext versions 2.5 through 3.1.3, and 3.5.0.

For more information, contact the Quantum Technical Assistance Center and reference CR 27318.

Symptoms

This issue has the following symptoms:

- This problem may occur only if the file system is mounted (on any client) at the time the stripe group is brought UP
- Only files with data on the formerly DOWN stripe group are subject to corruption
- Files created after the stripe group was brought UP are most likely to encounter corruption, especially if the files were created on clients that did not unmount and remount the file system
- It is possible that corruption will not be evident until the file system is unmounted and remounted
- The system log file may contain messages such as the following:

```
Mar 26 13:36:11 system10 kernel: cvfs 'snfs1': I/O attempt on DOWN/  
OFFLINE stripe group 1 cookie 0xa
```

```
Mar 26 13:36:11 system10 kernel: I/O error on cookie 0xa cvfs error 'IO error'  
(0x3) os status code 0x3 blkno 0x1d0000 offset 0x0 count 0x10000 resid 0x0
```

Solution

Install Stornext 3.1.4 or later, or StorNext 3.5.1 or later. (Although these releases are unavailable as of the the date this alert was written, both releases should be available shortly.)

Workaround

Unmount file system containing the affected stripe group from ALL clients BEFORE bringing the stripe group UP, then remount the file system. It is preferred and recommended that you unmount the file system from all clients before the file system is transitioned to the DOWN state.

Known Issues

A related issue may be encountered when a stripe group is initially transitioning to the DOWN state. StorNext client systems may have data in the buffer cache that has not been written to that stripe group. To prevent I/O errors or lost writes, all clients should cease access (ideally by unmounting the file system) before the stripe group is transitioned DOWN. Access may continue (the file system may be remounted) once the transition is complete.

Be certain the above Solution or Workaround is implemented before bringing the stripe group UP again.