



Product Alert 35 (Revised)

Product	All versions of StorNext®
Summary	Possible data corruption with large metadata stripe group
Date	September 2011 (First revision published May 2010)

Overview

Systems with stripe groups larger than 16TB which contain metadata may experience data corruption.

For more information, contact the Quantum Technical Assistance Center and reference CR 31093.

Note: Compared to the first version of this alert published in May 2010, this revision updates the information in the **Solution and Workarounds** section.

Symptoms

The symptoms are varied and can consist of user data corruption, FSM panics and error messages related to directories in cvlog. The size of the pieces of corrupted user data are multiples of the stripe breadth. For example, a 4k file system blocksize and stripe breadth of 2 will show corruption in multiples of 8K.

These are examples of cvlog errors:

```
[0305 22:45:58] 0x4b2c7940 (**Error**) Cannot cvbtree_insert inode  
0x16a498f521e9 into parent 0xa
```

```
[0305 22:45:59] 0x4ac3a940 (**Error**) Invalid inode lookup:  
0x16a86dc96270 markers 0x39ee/0xe25c gen 0x7cbc5619 nextiel  
0x39b5210e31cd32ac
```

This is an example panic message:

```
[0309 05:49:05] 0x4aab7940 (**FATAL**) PANIC: /usr/cvfs/bin/fsm ASSERT  
failed "!fpak.cvbtf_error" file cvbt_insert.c, line 173
```

Cause

A problem in the StorNext user-space I/O library may cause certain I/O requests to be mapped to an incorrect LUN in a metadata stripe group. This problem only affects the user-space applications (FSM, cvfsck, cvupdatefs, cvfsdb) and not the SNFS client I/O routines. Thus, data-only stripe groups are not affected by this issue. In a mixed metadata and data stripe group the LUN mapping error can cause the FSM to write metadata to a location that is allocated for file data. This error will manifest itself as data corruption in files as well as metadata corruption.

The following table shows the sizes at which a stripe group may be at risk, which is dependent on the stripe breadth.

File System Blocksize	StripeBreadth (in FsBlockSize Blocks)		
	1 Block	2 Blocks	4 Blocks
4KB	> 16TB	> 32TB	> 64TB
8KB	> 32TB	> 64TB	> 128TB
16KB	> 64TB	> 128TB	> 256TB
32KB	> 128TB	> 256TB	> 512TB
64KB	> 256TB	> 512TB	> 1024TB

In addition to the stripe group size, to be at risk the number of LUNs in the stripe group must be greater than 1 and not a power of 2. Having 1, 2, 4, 8, 16, 32, 64 or 128 LUNs in a stripe group means data corruption will not occur.

Solution and Workarounds

This issue is corrected in StorNext releases 3.1.4.1, 3.5.2.1 and 4.0.1.

The workaround is to use the table to determine if your configuration is susceptible to the problem. Contact Quantum support for assistance if your configuration is at risk.

If you receive a message referencing Product Alert 35, it means you have the corrected code for the problem described in this Product Alert. There are two reasons for receiving messages referring to this alert:

- 1 It is necessary to determine if a file system was susceptible to corruption from actions occurring prior to upgrading, and then take additional steps to correct any potential problems if the file system was susceptible. If your file system was susceptible, you will receive a message and are asked to contact Quantum Support.
- 2 If the file system was *not* susceptible, it will be marked as such when the file system is started. If you upgrade StorNext and then run a utility prior to starting the file system, the utility may fail.

For example, running cvfsck causes a failure after printing a message referring to Product Alert 35 as follows:

```
*Fatal*: FS_STATUS_PA35_SAFE not set. Aborting
```

You can manually set a file system to contain the information saying it has been checked for susceptibility to this problem by running `cvupdatefs`.

```
cvupdatefs -P on <fsname>
```

This message will appear in the output:

```
Filesystem <fsname> is set for pa35 geometry.
```

Using "-P off" will revert the file system to an unchecked state.

Bulletin and Alert Automatic Notification

To automatically receive future bulletins and/or alerts for Quantum StorNext via email, go to the location below and enter basic information about yourself, your company, and your StorNext product. You will need to provide your StorNext serial number when you sign up.

Go to this location to sign up for automatic notification:

www.quantum.com/stornextbulletins.

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.)