



# Product Alert 35

<b>Product</b>	All versions of StorNext®
<b>Summary</b>	Possible data corruption with large metadata stripe group
<b>Date</b>	April 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.

## 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.cvbt_f_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 will be addressed in a future StorNext release.

In the meantime, 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.

---

## Contacting Quantum

---

More information about StorNext is available on the Quantum Service and Support website at [www.quantum.com/ServiceandSupport](http://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:

<b>North America</b>	+1 800-284-5101 Option 5
<b>EMEA</b>	00800 9999 3822
<b>Online Service and Support</b>	<a href="http://www.quantum.com/OSR">www.quantum.com/OSR</a>
<b>World Wide Web</b>	<a href="http://www.quantum.com/ServiceandSupport">www.quantum.com/ServiceandSupport</a>

(Local numbers for specific countries are listed on the Quantum Service and Support Website.)