



Product Alert 39

Product	StorNext® 4.0, 4.1 and 4.1.1
Summary	Potential data integrity issues when defragmenting or migrating stub files
Date	August 2011

Overview

Running snfsdefrag on stub files can cause the files to appear to be online without valid data being present in the area beyond the stub. This can lead to read-side data corruption.

To encounter the issue, the following conditions must exist:

- 1 The metadata controllers must be running StorNext 4.0, 4.1 or 4.1.1.
- 2 Stub files must be enabled for one or more managed file systems.
- 3 An operation must be performed on the stub files that causes data movement using the snfsdefrag command such as Migration, Disk-to-Disk Relocation, or Defragmentation.
- 4 Files are subsequently read.

For more information, contact the Quantum Technical Assistance Center and reference CR 32240.

Symptoms

When an application reads an affected file at a file offset beyond the stub, it receives zeros instead of the expected data.

Potentially corrupted files can be identified by comparing the number of allocated blocks with the length of the file and checking whether the file is marked as being online. This can be done using the following commands:

```
# ls -l /stornext/snfs1/mypol/testfile
-rwxr-xr-x 1 root root 1784053 2010-11-30 12:57 /stornext/snfs1/
mypol/testfile
# du testfile
16      testfile
# fsfileinfo /stornext/snfs1/mypol/testfile
-----
File Information Report                               Tue Nov 30 13:00:18 2010
Filename: /stornext/snfs1/mypol/testfile
-----
--
      Last Modification: 30-nov-2010 12:57:25
Owner:      root                               Location:   DISK AND TAPE
Group:      root                               Existing Copies: 1
Access:     755                               Target Copies: 1
Target Stub: 16 (KB)                          Existing Stub: n/a
File size:  1,784,053                          Store:     MINTIME
Affinity:   n/a                               Reloc:     MINTIME
Class:      mypol                              Trunc:     MINTIME
                                           Clean DB Info: NO

Media:  sdisk1(1)
Checksum:      N
```

```
FS0000 30 2033404121 fsfileinfo completed: Command Successful.
```

Note that:

- 1 The file is marked as being online (the fsfileinfo command reports the file as being on DISK AND TAPE).
- 2 The file size (1,784,503 bytes) is greater than bytes allocated on disk (16,384 bytes).

Therefore, the file is potentially corrupt (assuming the file was not originally written as a sparse file).

The following Perl script (called pa39.pl) can also be used to audit an entire system:

```
#!/usr/bin/perl
use File::stat;
open(F, "find @ARGV \! -type l -a \! -type d|" ) ||
    die "cannot execute find: $!";
while(<F>) {
    chop;
    $file = $_;
    $s = stat($file);
    if (!defined($s)) {
        print "cannot stat \"$file\": $!\n";
        next;
    }
    next if ($s->blocks * 512 >= $s->size);
    $ondisk = `fsfileinfo $file | grep -c "Location.*DISK"`;
    print $file, "\n" if ($ondisk != 0)
}
}
```

To run this command, use an editor to write out the above text to a file named **pa39.pl** and then type the following on the Primary metadata controller (MDC):

```
# chmod 755 ./pa39.pl  
# . /usr/adic/.profile  
# ./pa39.pl /stornext/snfs1 /stornext/snfs2
```

Replace "/stornext/snfs1 /stornext/snfs2" above with the mount points of the exposed managed file systems.

Note: The pa39.pl script runs find and fsfileinfo commands that may degrade MDC performance to a degree.

Files showing corruption can be corrected by running the fsmrmdiskcopy command on them. For example,

```
# fsmrmdiskcopy /stornext/snfs1/mypolicy/myfile
```

Running fsmrmdiskcopy will mark the file as offline, and the correct data will be read from tape the next time that data beyond the stub is accessed or the file is explicitly retrieved.

Cause

The problem is due to a regression that exists in StorNext 4.0, 4.1 and 4.1.1.

Solution

Upgrade to StorNext 4.0.1.1, or upgrade to StorNext 4.1.2 or newer. Also, use the procedures described previously to identify and correct any corrupted files.

Workaround

To avoid the corruption, do not defragment or otherwise migrate any stub files. Disabling stub files prevents stubs from being created in the future, but files having existing stubs will retain them (and thus remain exposed) unless or until the files are retrieved.

Deploying workarounds to prevent corruption (including upgrading StorNext) will NOT address existing corrupted files. Such files can be identified and corrected as described previously in the [Overview](#).

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.)