



Product Alert 50

Product	StorNext 6.2.0 through 6.3.1
Summary	Potential data missing on tape device
Date	January 2020

Overview

This section provides the following information:

- [Affected Systems](#)
- [Problem Description](#)
- [For ANTF \(Quantum Default\) Formatted Tapes](#)
- [For LTFS Formatted Tapes:](#)
- [Time Line/Example](#)

Affected Systems

Affected system include:

- This issue only applies to systems running StorNext 6.2.0 through StorNext 6.3.1 using Storage Manager to store files to a tape device.
- Releases prior to StorNext 6.2.0 are not affected by this issue.
- Storage Manager may encounter this problem while storing a group of files to a tape device by **policy** or **fsstore**.
- Systems using Storage Manager or FlexTier to store data exclusively in public cloud, private cloud, or Storage Disk (SDisk) are not affected by this issue.

6-00960-157 Rev B, January 2020, Product of USA.

© 2020 Quantum Corporation. All rights reserved. Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law. Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, DXiV-Series, FlexTier, Lattus, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners. Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks

Problem Description

When policy runs, it selects a group of files to be written to tape. The default file count is 3000 files, but it may vary because it is a configurable policy setting parameter.

If the **last file** in the selected group of files is deleted from the disk after store processing starts and prior to writing the files to tape, there is a possibility that some or all files may not be completely written to tape. The files may not be accessible from the tape, depending upon the tape format, but the store processing will report success. This can occur whether the file was deleted manually or by an application.

For ANTF (Quantum Default) Formatted Tapes

There is a possibility that the next to last file (or last few files if small files) may not be completely written to tape.

All of the files are written to the tape, but the last portion of the 2nd to last file may still be in the memory buffer and error handling causes that buffer not to be flushed to the tape device. It is possible that this could be multiple files if the sum of their sizes is less than the buffer size (typically 512K).

As an example, **fspolicy** runs and selects 3000 files to store to a tape device. The last 2 files in the list are `file2999` and `file3000`. While this policy is running, `file3000` is removed. There is a possibility that `file2999` may be written to the tape device, but the last portion of the file may not be written.

For LTFS Formatted Tapes:

Caution: DO NOT MOUNT or RE-MOUNT THESE TAPES - LTFS may overwrite the index.

The data for the files may have been written to the tape, but there is the possibility that the LTFS index is not written to tape. This results in the cluster of files being inaccessible. On the next mount of the LTFS partition, the LTFS recovery algorithm, because the index is not present, will attempt to overwrite that cluster of data.

Note: If the LTFS tape is not mounted after the error situation, the data may still exist on the tape.

Time Line/Example

To help, here is a more detailed time line/example:

T1: Customer or application creates 10 files.

file1, file2, file3, file4, file5, file6, file7, file8, file9, file10

T2: 4 of these files are selected to store:

fsstore file1 file2 file3 file4

(Or this could be an fspolicy store)

T3: fsstore will verify that all of the files exist on disk and meet the criteria. Note that this could be fspolicy as well.

T4: The time window for the issue occurring opens now.

T5: Some time elapses to allocate the resources for the group of files being written to tape.

NOTE: Depending upon the workload on the system, this could be a short period of time OR it could be a longer period of time depending upon the operations ahead of the store.

T6: Once the tape is allocated and the request passed to the process performing the I/O, the process does initial file validation.

T7: The time window for this issue occurring closes now.

T8: The I/O process writes the files to tape.

Example

The 4 files are stored in order:

- file1
- file2
- file3
- file4

Note: That is the order they will be written to tape.

For ANTF formatted tapes:

If between T4 & T7 the last file (file4) is removed, then it's possible that file3 will be written to tape, but if there's any residual data that crosses a block boundary, that residual data will not be written to tape.

For LTF5 formatted tapes:

If between T4 & T7 the last file (file4) is removed, then it is possible that the index for the cluster may not be written to the tape, resulting in the error situation.

Caution: DO NOT MOUNT or RE-MOUNT THESE TAPES - LTFs may overwrite the index. Contact Quantum Support for assistance with how to proceed. See [Contacting Quantum](#) on page 10.

When this occurs, there is no error indicated back to the user and the issue will not be discovered until they attempt to retrieve file3.

If the file is not truncated from disk, the data is still good on disk. It's only when the file is truncated that it becomes an issue of "data loss".

Symptom

If the problem has occurred on the system, when retrieving a file from the tape device the retrieve may fail due to missing data.

Looking at the TSM tac logs, something similar to the following will be seen.

For ANTF:

```
...fs_moverd[174990]:  
E1201(8)<4875075>:<1:1>:fsScsi::FsScsiLogCheckCond1455:  
{10056:10055}: Check condition: op=08h key=00h asc=00h ascq=01h  
FILEMARK DETECTED  
  
...fs_moverd[174990]:  
E1200(7)<4875075>:<1:1>:ReaderThread::readSync453: {10056:10055}:  
Not enough data read. read 0 of 524288, requested 378944  
  
...fs_moverd[174990]:  
E1200(7)<4875075>:<1:1>:ReaderThread::reader181: {10056:10055}:  
Read failed after 621056 bytes  
  
...fs_moverd[174990]:  
E1200(7)<4875075>:<1:1>:MoverThread::processMoverReq1490:  
{10055:10055}: A fatal error: E_SRCE_RD_MS(12) has been encountered,  
failing all files in request
```

For LTFS:

```
...fs_moverd[174990]:  
E1200(7)<4875478><1:1>:MoverThread::openLTFSFile8122:  
{10093:10093}: Could not open LTFS file: file_key 643962, version 1, path  
_Handle_Only_  
  
...fs_moverd[174990]:  
E1200(7)<4875478><1:1>:MoverThread::copyCheckTransfer2284:  
{10093:10093}: IO transfer mismatch. Expected 1000000, actual: 0  
  
...fs_moverd[174990]:  
E1200(7)<4875478><1:1>:MoverThread::processMoverReq1490:  
{10093:10093}: A fatal error:  
  
E_IO_INCOMPLETE_MS(42) has been encountered, failing all files in request  
  
...fs_moverd[174990]: E1202(9)<4875478><1:1>:Util::logEndCopy102:  
{10093:10093}: END copy of File 'key:643962 seg:1'  
Move_status:E_IO_INCOMPLETE_MS(42)
```

Solution

This issue is addressed in StorNext release 6.3.1.1. Potentially affected systems should be upgraded immediately to prevent future issues.

Note: StorNext release 6.3.1.1 was released on 12/28/2019.

After upgrading, use the following steps for systems expected to have encountered the issue.

Scripts are available from Quantum Service to detect whether your system may have encountered this issue.

- For systems that use ANTF media, the script is more robust and can better detect potential files that have issues.
- For systems that use LTFS media, as of this writing, the only mechanism available is to examine the existing Storage Manager tac logs for particular entries.
- If the logs have rolled and do not exist, Quantum will be unable to detect the issue.
- Engineers are working to determine if there is another mechanism for detecting all files written to LTFS media.

Note: To obtain the detection scripts: Contact your Quantum Service Representative, refer to Product Alert 50 (6-00960-157), and request the scripts related to Bug #76414.

See [Contacting Quantum](#) on page 10.

Scripts

These scripts will help detect if the system may have the issue.

- After it's determined which script to use, run the script on the MDC node.
- The script will indicate which files could have the issue.
- For each file it provides the path (at the time the file was stored), the file key and which copies may be affected.

If it is determined that the issue exists with files, there are a several steps to mitigate the problem. It will need to be understood if each file is truncated or still resides on disk.

For LTFS Formatted Media

Before continuing, to avoid overwriting the index on media that are affected, the media should be made unavailable by running the **fschmedstate** command or physically changing the write protect tab on the media to write-protected. The **fschmedstate** command is as follows:

```
fschmedstate -s unavail <mediaid>
```

Note: When contacting Quantum Support about the issue, let them know that this was done to each of these media as this will need to be adjusted during future recovery. See [Contacting Quantum](#) on page 10.

Quantum engineering has written scripts that perform the majority of the repair steps listed below. Contact Quantum Support for these scripts.

Determine if the File is Truncated or Not

Run the **fsfileinfo** command and examine the location entry:

```
fsfileinfo <filename> | grep Location
```

For example, you will see one of the following responses:

```
Owner: root   Location:  DISK AND ARCHIVE
Owner: root   Location:  DISK
Owner: root   Location:  ARCHIVE
```

- If the location is only ARCHIVE, then the file is truncated. Continue to the section [Truncated File](#).
- If the location is DISK AND ARCHIVE, the file is stored to tape and is not truncated. Continue to the section [File is Not Truncated](#).
- If the location is DISK, the file is newer and has not been stored yet and nothing further needs to be done as this new version of the file will be stored by policy.

File is Not Truncated

If the file is not truncated, remove all copies of the file from tape using the **fsrmcopy** command.

```
fsrmcopy -a <filename>
```

Note: The **-a** option is key here. If not specified, the file will be truncated from disk.

Allow the policy to re-write the file to tape OR manually re-write the file to tape using the **fsstore** command.

Truncated File

If There Are Other Good Copies

If there are multiple stored copies of the file and the script did not show that all copies are potentially affected: Select a copy # that is not specified by the script and attempt to retrieve that specific copy.

```
fsretrieve -c <good file copy> <filename>
```

When successfully retrieved, remove each suspected copy from tape using the **fsrmcopy** command.

```
fsrmcopy -c <copy # from script> <filename>
```

Note: The **-c** option is key here. If not specified, the file will be truncated from disk.

Allow the policy to re-write the copies of the file to the tape device OR manually re-write the file to tape using the **fsstore** command.

Another Option for ANTF formatted media if time permits:

- Instead of removing the copy, go ahead and try to retrieve the copy specified by the script and see if it retrieves successfully.
- If it retrieves successfully then the file copy is not affected.
- This retrieve can be accomplished using the following command:

```
fsretrieve -c <copy # from script> <filename>
```

If There Are No Other Good Copies

Note: Refer to the section below for the appropriate media format type.

For ANTF Formatted Media

These steps should only be utilized if the media is an ANTF formatted media. Attempt to retrieve the file using the **fsretrieve** command without a copy # specified. This will try all copies until success.

```
fsretrieve <filename>
```

If successful, then remove all copies of the file using the **fsrcopy** command.

```
fsrcopy -a <filename>
```

Note: The **-a** option is key here. If not, specified, the file will be truncated from disk.

Note: It is possible that running multiple retrieves will result in the drives taken offline due to the errors encountered attempting to retrieve the files.

It is recommended that the drive state be monitored and brought back online if this does occur.

Allow the policy to re-write the file to the tape device OR manually re-write the file to the tape device using the **fsstore** command.

- If the retrieve fails and the file can be re-constructed through other mechanisms, please do so.
- If the file cannot be re-constructed, the initial portion of the file may be able to be obtained using the partial retrieve command.

- Please contact support for guidance.
- See [Contacting Quantum](#) on page 10.
- If there are older stored versions of the file, one of these versions could be retrieved and used.

For LTFS Formatted Media

Caution: DO NOT MOUNT or RE-MOUNT THESE TAPES - LTFS may overwrite the index. QUARANTINE THE LTFS TAPES so they are not damaged more.

It is possible that trying to retrieve a file, may result in the cluster being overwritten on tape if it has not been already.

Note: Contact Quantum Support and they can assist in determining if there is data that exists on tape that can be recovered. See [Contacting Quantum](#) on page 10.

Workaround

If the system cannot be upgraded to 6.3.1.1 but is running 6.2.0 or later, avoid deleting files from disk until the files are stored to tape.

Prior to removing a file, to examine if the file has been stored to tape, use the **fsfileinfo** command on the MDC to examine if the file is present, or on a client use the **snfileinfo** command.

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 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:

United States	800-284-5101 Option 5 (toll free)
EMEA	00800 7826 8888 (toll free) 49 6131 3241 1164
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.)