Product Alert 46

Product	StorNext 5 File Systems prior to release 5.3.0 including the Xsan versions 2.2 and 2.3
Summary	Incomplete restore and/or corrupt metadump
Date	September 2015

Problem

StorNext File Systems created prior to Release 4.0 and later upgraded to a version of StorNext 5 prior to release 5.3.0 may have incomplete restore and/or corrupt metadumps.

Symptoms

This issue applies to versions of StorNext 5 file systems prior to release 5.3.0 that meet all of the following four criteria:

Note: If the pre-StorNext 4 file system has not been upgraded to StorNext 5, refer to the Solution section.

In this document, substitute *fsname* with the name of your file system under consideration or repair.

1 The file system(s) was/were originally created with a release of StorNext prior to 4.0. These file systems were created when file system inode data structures were 512 bytes in length (a "small inode" file system).

- StorNext file systems created by versions prior to 4 have inode structures that are 512 bytes long, also known as "small inode" file systems. By contrast, file systems created by StorNext 4 and 5 have much larger inode structures

- For file systems that have not had their metadata restored from a StorNext 5 metadump, it can be determined if the file system is a small inode file system by using cvfsdb's "**show sb**" command and examining field sb_InodeVersion to look for the following string: "small inode":

  # /usr/cvfs/bin/cvfsdb fsname
  cvfsdb> show sb

- If the file system has had its metadata restored from a StorNext 5 metadump, then someone with knowledge of the file system's history must determine if the file system was originally created prior to StorNext 4.0. Finding the date when the file system was created may be helpful in this regard and that can be done using **cvadmin**:

  # cvadmin -e "select fsname"

2 They have subsequently been upgraded to a version of StorNext 5 prior to release 5.3.0 or later.

3 They have **restoreJournal** set to 'true' in the configuration file. This will be the case for all managed file systems but may or may not be the case for non-managed file systems.

   # vi /usr/cvfs/config/fsname.cfgx

   <snfs:restoreJournal msgId="restoreJournal">true</
   snfs:restoreJournal>

4 At the time of the upgrade to StorNext 5, either sb_QuotaInode or sb_NTSecurityIdxInode or both were set to zero in the superblock:

   # /usr/cvfs/bin/cvfsdb fsname
   cvfsdb> show sb

   . . .

    sb_NTSecurityIdxInode = 0x0000000000000000

    sb_NTSecurityDatInode = 0x0000000000000000

    sb_QuotaInode   = 0x0000000000000000

   . . .

Two Scenarios

Caution: The actual estimated time for the two solutions and/or workaround/prevention to complete is based on the size of the file system(s) that must be restored at any particular site.

Restore time depends on many variables including: performance of the storage for the shared file system where the metadump resides, number of transaction bundles pending in the update database, total number of inodes, and total number of file extents, etc.

The file system could take several hours to restore; so, schedule your downtime accordingly.

This section provides the following:

- First Scenario
- Second Scenario

First Scenario

The first scenario has no visible symptoms, but the metadump cannot be used for disaster recovery or metadata migration. In this scenario, the file system has not been restored using a StorNext 5 metadump. Although the file system is not affected, its corresponding metadump cannot be used.

Second Scenario

In the second scenario, the file system's metadata has been restored from the StorNext 5 metadump, transferring any inode corruption that may have existed in the metadump into the file system.

The second scenario has the following three known symptoms:

1 An attempt to delete a directory (folder) may cause the FSM to panic with the following error:

   PANIC: /usr/cvfs/bin/fsm ASSERT failed "error != ENOENT"

2 An attempt to access a regular file (or any non-directory inode) may cause the FSM to panic with the following error:

   FSM PANIC "In_rele: ip(0x… : 0x…) i_ref_count(-1) has gone negative!"

3 After shutting down the FSM, running '**cvfsck -n**' produces errors indicating that inodes are corrupt and/or there are orphaned inodes (files or directories without a parent directory). For example:

   Error: Clearing directory 0x18000000440e5: error opening directory
   Clearing inode 0x18000000440e5, reason: 'corrupt non-root directory'.
   Directory entry 2015.08.19_16.15.52 (p=0x… i=0x… flags=0x…) is an orphan. Removed.

Solution

If customers suspect that they have one or more file systems that meet all four criteria listed in Symptoms on page 1, immediately contact Quantum Support. Refer to Contacting Quantum on page 4.

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:

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.)