



Product Alert 31

Product	StorNext [®] File System
Summary	The new VTOC disk label format in StorNext 3.1.3 for Solaris 10 has implications for IRIX systems and backwards compatibility.
Date	June, 2009

Overview

The default format for VTOC labels changed in StorNext release 3.1.3 from the format used in release 3.1.2.

As of release 3.1.3, StorNext recognizes three VTOC label formats:

- 1 The original format
- 2 A Solaris 10 + patch 118833-17-compatible format
- 3 A Solaris 10-compatible format

The original VTOC format, #1, is not compatible with Solaris 10 systems, except those running a patch level prior to 118833-17.

The default VTOC format used in StorNext release 3.1.2GA is #2. This format is compatible with Solaris 10 patch 118833-17 and beyond, but on many IRIX configurations this format could be incompatible or cause poor performance, and it could cause data corruption when IRIX clients write data under specific failover conditions.

In StorNext 3.1.3, the `cvlabel` command's `-I` option causes the command to write VTOC label format #3, which is compatible with all Solaris 10 systems. However, this is not the default format for `cvlabel` in all versions of StorNext up to and including 3.1.3 and 3.5.0, so you must explicitly invoke its use with the `-I` option. This new format is compatible with IRIX, but IRIX systems will end up with a CTQ depth of 1 in StorNext releases prior to 3.5. This will cause severe performance degradation, so Quantum does not recommend its use. StorNext 3.5 is able to adjust the CTQ depth at mount time, resulting in the expected performance. Label format #3 is not backwards compatible with StorNext versions prior to 3.1.2.

Symptoms

If you use the StorNext GUI to select an unlabeled LUN less than 2TB for use with a new file system or to expand an existing file system, running `cvmkfs/cvupdatefs` will generate an I/O error and the operation will fail.

Following is an example output of the failure:

```
File System Name: snfs1
Mount Point: /stornext/snfs1
Data Migration: Yes
Block Size: 16384
Stripe Group: StripeGroup1
  Stripe Breadth: 64 Kilobytes
Label Type: VTOC
MetaData: Yes
Journal: Yes
User Data: Yes
Disk: /dev/rdisk/c5t203600A0B82665C2d0s2
Re-initializing file system 'snfs1'.
Shared Meta Data File System.
Meta Data Root is on "StripeGroup1".
Resetting Arbitration Block.
Cannot create local device "/dev/rdisk/c5t203600A0B82665C2d0s2" - I/O error.
PiIO: pio_status/5
*Fatal*: Cannot write to ARB block! - I/O error
The creation of file system 'snfs1' failed - I/O error.
Configuration file snfs1.cfg moved to /tmp.
```

Also, running `cvlabel -la` will show disks with errors when previously the disk was working correctly, as shown here:

```
/dev/rdisk/c5t203700A0B82665C2d0s2 [SUN CSM200_R 0619] SNFS "disk001"
Sectors: 209704960. SectorSize: 512. Maximum sectors: 209704960. [Unusable:
Unable to read first data block]
```

Recommendations

Quantum recommends the following based on your operating system and needs:

- **For LUNs less than 2TB and Solaris systems but no IRIX systems and no clients running StorNext releases prior to 3.1.2:**
Run `cvlabel -I` to create Solaris-compatible VTOC labels
- **For LUNs greater than 2TB and Solaris systems but no IRIX systems:**
Generate EFI labels with `cvlabel`
- **For LUNs less than 2TB and IRIX systems but no Solaris systems:**
Run `cvlabel -i` to generate IRIX compatible VTOC labels
- **For IRIX and Solaris systems:**
Upgrade to 3.5.x and generate EFI labels

Although this issue appears in StorNext release 3.5, it has been addressed in release 3.5.1. In StorNext 3.5.x, the cvlabel -I format is compatible with IRIX and has the expected performance of using a larger CTQ.

Workaround

Before running the StorNext GUI, label any new LUNs 2TB (or smaller) through the command line using cvlabel -I.

To create a Solaris-compatible label via the command line, follow these steps:

- 1 Run cvlabel -c to redirect the output to a file. Name the file something appropriate such as "cvlabels".
- 2 Remove from the cvlabels file all of the device entries except for those that are to be configured.
- 3 Modify the line(s) in the cvlabels file containing the devices you want to label: Change the first column in the file to contain a unique StorNext label name.
- 4 Verify that only devices which should be relabeled are present in the cvlabels file, and also verify that the label names are correct and unique.
- 5 Run the command cvlabel -rl <cvlabels file>

The cvlabel command prompts for verification before it writes the label. If the proposed action is correct, answer Y when prompted.

- 6 Reboot the system with the command reboot -- - r to ensure that StorNext recognizes the reconfiguration.

If you are configuring through the StorNext GUI, the newly labeled devices will be present in the configuration wizard.

The man page for cvlabel contains additional background and potentially useful information.