



Product Bulletin 31

Product	All 2.x and 3.x StorNext versions
Summary	A StorNext file system will not start
Date	November 2008

Problem

The error message (“Server could not find any metadata devices”) is somewhat misleading because the metadata device may be fine. The problem arises because one of the LUNs associated with the file system and configured in its configuration file cannot be included in the list of LUNs being built by FSM. When this happens, the FSM sees this as a FATAL error and will not start the file system.

Below are some of the reasons that can cause the failure.

System out of memory:

- Cannot malloc space for local PHYSIO device!
- Cannot malloc space for PHYSIO path
- Cannot malloc space for PHYSIO disk label

I/O Path not operational:

- Cannot get physical disk

Connectivity issue (sockets):

- Cannot connect to the FSMPPM

Multi-path issues:

- Device raw: <device_path> name <StorNext_label> is a DUPLICATE DEVICE - FAIL!

Symptoms

A StorNext file system will not start and the following error is displayed in the cvlog for that file system:

```
[1029 17:09:28] 0x10000 (**FATAL**) Server could not find any Meta-Data devices!
```

Solutions and Workarounds

Finding the correct solution can be challenging because the error message does not provide a specific cause. Below are suggestions for some solutions and workarounds for the causes previously mentioned.

System Out of Memory

Monitor the system's memory usage and determine processes that are consuming a large amount of memory. Reduce usage by those processes, or increase the amount of memory available on the system.

I/O Path Not Operational

Using the data in the file system configuration file, determine which LUNs are required and then validate that system access works for those paths. On *nix systems an easy way to validate access is to use the dd utility to dd the data from the system device path to /dev/null. Monitor the system messages file for any errors. The paths to the LUNs and their label can be found using the SNFS utility /usr/cvfs/bin/cvlabel with the "-l" flag.

Example: dd if=/dev/sdb of=/dev/null

Connectivity Issue (sockets)

Check the file /usr/cvfs/debug/nssdbg.out for connectivity issues. If you need more information, touching the file /usr/cvfs/debug/verbose will cause more debug information to be logged.

Most of the time connectivity issues are caused by mis-matched host name entries, or because the /usr/cvfs/config/fsports file limits the number of ports too much.

Multi-path Issues

Systems using multiple paths to LUNs for redundant access can sometimes have issues. SNFS can see an issue with a multi-path system if the firmware on one controller in a path is different than the firmware on the other paths. The data SNFS uses to make sure there is not a duplicate SNFS label, but a valid multi-path entry, is the "inquiry string" and the WWN serial number. To see what SNFS is looking for, use the utility /usr/cvfs/bin/cvlabel with the "-l" flag. Below is an example of a multi-path system that has mis-matched firmware:

```
Meta_Data /dev/rdisk/20000001ff0202f2/lun77vol/c3p1 163825664 VTOC #
host 3 lun 77 sectors 163825664 sector_size 512 inquiry [DDN S2A 3000
5.28] serial 60001FF02002040902F26BBE30D40202

Data_LUN /dev/rdisk/20000001ff030635/lun155vol/c3p1 1920901248 VTOC #
host 3 lun 155 sectors 1920901248 sector_size 512 inquiry [DDN S2A 8500
5.29] serial 60001FF02002110106364C1FE53F4C01

Meta_Data /dev/rdisk/20000001ff02039a/lun77vol/c7p1 163825664 VTOC #
host 7 lun 77 sectors 163825664 sector_size 512 inquiry [DDN S2A 3000
5.28] serial 60001FF02002040902F26BBE30D40202

Data_LUN /dev/rdisk/20000001ff030636/lun155vol/c7p1 1920901248 VTOC #
host 7 lun 155 sectors 1920901248 sector_size 512 inquiry [DDN S2A 8500
5.28] serial 60001FF02002110106364C1FE53F4C01
```

In the example above the data in RED is the inquiry string SNFS checks, and the data in BLUE is the WWN serial number.

The inquiry string for the first instance of Meta_Data is [DDN S2A 3000 5.28] and the second instance of Meta_data is [DDN S2A 3000 5.28]. The WWN serial numbers in both instances are the same. This is a valid entry for a multi-path device.

The inquiry string for the first instance of the Data_LUN it is [DDN S2A 8500 5.29] and the second instance of the Data_LUN is [DDN S2A 3000 5.28]. These values are different. The WWN serial numbers in both instances are the same. According to the test in the initialization of StorNext, this is considered an invalid multi-path setup, and it is believed there is a duplicate SNFS label present. This will cause the initialization to fail, and the file system will not be allowed to start, preventing data corruption.

The workaround for this particular issue is to remove one of the paths, making it a single path. This will allow the file system to start.

The solution is to make sure the firmware levels are exactly the same on ALL paths being used to this LUN in SNFS.