

GRIDScaler[™]

Spectrum Scale Metadata Scanning Failure Puts File System Data or Metadata at Risk

ALERT! *Installation of the Spectrum Scale 4.2.3.4 PTF is MANDATORY* for customers running GRIDScaler 3.2.4 or 4.2.0. Please contact DDN Support for assistance at your earliest convenience.

ALERT! *DO NOT RUN the commands `mmdeldisk`, `mmrpldisk`, `mmrestripefs`, or `mmaddisk -r`* on GRIDScaler 3.2.4 or 4.2.0 until the Spectrum Scale 4.2.3.4 PTF has been applied. Running these commands before patching the file system will put your data at risk.

Issue Summary

IBM has identified a problem with the GPFS file system metadata scanning function in IBM Spectrum Scale versions 4.2.3.0 through 4.2.3.3. The metadata scanning problem may result in silent file system data corruption or metadata corruption under certain conditions. This issue affects DDN GRIDScaler versions 3.2.4 and 4.2.0.

Affected Systems

The metadata scanning issue affects storage systems meeting **all** of the following conditions:

- **System is running GRIDScaler 3.2.4 or 4.2.0** (with Spectrum Scale 4.2.3.2 or 4.2.3.0, respectively).

▪ **Any of the following commands have been run and failed:**

- **mmdeldisk** (shown as **tsdeldisk** in the **mmfs.log** file)
- **mmrpldisk** (shown as **tsrpldisk** in the **mmfs.log** file)
- **mmrestripefs** (except **-c** option for replica compare or **-z** option for compressing user files)
- **mmadddisk -r** (restripe run automatically after adding disks)

NOTE If **all** executions of these commands have run **successfully** since GRIDScaler 3.2.4 or 4.2.0 were installed, your system has **not** been affected.

ALERT! **DO NOT RUN the commands mmdeldisk, mmrpldisk, mmrestripefs, or mmadddisk -r** on GRIDScaler 3.2.4 or 4.2.0 until the Spectrum Scale 4.2.3.4 PTF has been applied. Running these commands before patching the file system will put your data at risk.

Determining Whether Your System Has Been Affected

Example Error Messages

An unrecoverable metadata scanning error may appear in command output. For example:

```
$ mmrestripefs gpfs1 -m
Scanning file system metadata, phase 1 ...
Error processing inodes.
Not enough memory to allocate internal data structure.
mmrestripefs: Command failed. Examine previous error messages to determine cause.
```

Failed commands are also logged in the **mmfs.log** files. For example:

```
[I] Command: mmrestripefs /dev/gpfs1 -m
[E] Command: err 12: mmrestripefs /dev/gpfs1 -m
[E] Not enough memory to allocate internal data structure.

[I] Command: tsdeldisk /dev/fs0 -F /var/mmfs/tmp/diskfile.mmdeldisk.9961542 -r
[E] Command: err 12: tsdeldisk /dev/fs0 -F /var/mmfs/tmp/diskfile.mmdeldisk.9961542 -r
[E] Not enough memory to allocate internal data structure.
```

After a metadata scanning failure has caused corruption, a highly varied cascade of subsequent errors can follow.

For example, after an **mmrestripefs** command has failed with “**err 12**” (the out-of-pagepool-memory error), a second attempt to run **mmrestripefs** may trigger an **MMFS_FSSTRUCT** error, regardless of success or failure of the command. This will be logged in the syslog (**/var/log/messages**):

```
Error=MMFS_FSSTRUCT, ID=0x94B1F045, Tag=7866532:  
Invalid disk data structure. Error code 1108. Volume fs0
```

In another example, after `mmdeildisk` (shown as `tsdeildisk` in logs) fails with “**Not enough memory to allocate internal data structure**”, an offline `mmfsck` may fail in log recovery and cause the file system to be unable to mount. The resulting errors might look like this:

```
# mmfsck fs0 -xa -nvc -t /log  
GPFS: 6027-700 [E] Log recovery failed.  
GPFS: 6027-699 [E] Inconsistency in file system metadata.  
mmfsck: 6027-1639 Command failed. Examine previous error messages to determine cause.  
  
$ mmmount fs0  
6027-1623 mmmount: Mounting file systems ...  
mount: Stale file handle  
mmmount: 6027-1639 Command failed. Examine previous error messages to determine cause.
```

Checking Logs

To determine whether your system has been affected by the metadata scanning issue, check all `mmfs.log` files and syslog files (`/var/log/messages`) for any failure to execute the `mmdeildisk` (`tsdeildisk`), `mmrpldisk` (`tsrpldisk`), `mmrestripefs`, or `mmaddisk -r` commands. Also look for any `MMFS_FSSTRUCT` error messages in the syslog.

If these logs do not completely cover the period since GRIDScaler 3.2.4 or 4.2.0 were installed, checking the logs will not be sufficient to make a final determination. System administrators in this case should be alert for unusual error reports in logs or from users.

ALERT! If you suspect that a file system has already been affected by the metadata scanning issue, **please contact DDN Support immediately** for assistance. Support will run an **offline mmfsck** to confirm. **Do not run mmfsck without checking with DDN Support first.**

Resolution

This issue is resolved by the Spectrum Scale 4.2.3.4 PTF, which may be applied directly to GPFS without upgrading your entire GRIDScaler installation.

Please be aware that if data has been corrupted prior to installation of the PTF, data recovery must be undertaken separately. Installing the patch will **not** recover corrupted data.

ALERT! **Installation of the Spectrum Scale 4.2.3.4 PTF is MANDATORY** for customers running GRIDScaler 3.2.4 or 4.2.0. Please contact DDN Support for assistance at your earliest convenience.

Contacting DDN Technical Support

Please contact DDN Technical Support at any time if you have questions or require assistance.

Support can be reached by phone, by email, or on the web as listed below.

Web

DDN Community Support Portal
Portal Assistance

<https://community.ddn.com/login>
webportal.support@ddn.com

Telephone

DDN Support Worldwide Directory

<http://www.ddn.com/support/contact-support>

Email

Support Email

support@ddn.com

Bulletins

Support Bulletins

End-of-Life Notices

Bulletin Subscription Requests

<http://www.ddn.com/support/technical-support-bulletins>

<http://www.ddn.com/support/end-of-life-notices>

support-tsb@ddn.com