

GRIDScaler[™]

Data Replication Error in AFM Feature of GPFS and Spectrum Scale

ALERT! GRIDScaler customers who implement the active file manager (AFM) feature can experience silent data replication errors that put data at risk.

Issue Summary

DDN has been notified of an issue with the AFM (active file management) feature in GPFS and Spectrum Scale. Under certain conditions, AFM may not completely transfer writes from the AFM cache to the home location of a file. This may cause a mismatch between data in the AFM cache and the data in the home location, thereby putting data at risk.

AFM minimizes the number of write operations queued at a gateway node by either: (1) merging write operations when they apply to the same offset range in the same file; or (2) by creating an optimized list of write operations for the file if the writes cannot be merged. If the write operations being merged or optimized occur in a particular sequence that includes truncation of the file, and if all operations in that sequence take place within the window specified by the `afmAsynchDelay` configuration option (15 seconds by default), data in the home location can become corrupted.

For example:

- Write 1 MB of data at offset 1 MB.
- Write 256 KB of data at offset 0. ← Offset for second write is smaller than first, with no data overlap.
- Truncate file to 512 KB. ← Truncation occurs before offset for first write.

When operations with the characteristics described in red are performed in the order above, and all operations occur within the **afmAsynchDelay** window, the two writes shown are dropped from the optimized list of operations and only the truncation is performed on the file in the home location. The data in AFM cache is correct, but the home file is corrupted. No error is reported.

Subsequent reads will typically be made from the AFM cache and will retrieve the correct data. However, if an outage causes failover from cache to home, reads from the home file will be subject to silent errors. Similarly, if AFM reads data from the home file to cache, cache will become corrupted but no error will be reported.

All AFM replication modes that write to home are affected by this issue, including single-writer (SW), independent-writer (IW), and disaster recovery (DR).

Affected Systems

This issue affects the AFM feature in GPFS versions 3.5.0.0–3.5.0.34 and 4.1.0.0–4.1.0.8, as well as in Spectrum Scale versions 4.1.1.0–4.1.1.16 and 4.2.0.0–4.2.3.4. These versions of GPFS or Spectrum Scale are incorporated into GRIDScaler versions 1.6.0–3.2.4 and 4.0.0–4.2.0.

NOTE If you have **not** implemented the AFM feature, your system is **not** affected.

Resolution

Resolution requires both a fix to the code and recovery of any affected data.

Step 1. Update the software as shown in the table below.

ALERT! *GRIDScaler versions 1.6.0 – 2.3.0 require a full GRIDScaler ISO image upgrade* to GRIDScaler 3.2.0 or higher to resolve this issue. Upgrading just the GPFS bits to version 4.1.1.17 or 4.2.3.5 will create incompatibilities between Spectrum Scale and the remaining, older components of GRIDScaler.

GRIDSCALER VERSION	GPFS / SPECTRUM SCALE VERSION	UPGRADE REQUIRED
GRIDScaler 4.2.0	Spectrum Scale 4.2.3.0	Apply Spectrum Scale PTF 4.2.3.5
GRIDScaler 3.2.1 – 4.1.0	Spectrum Scale 4.2.0.1 – 4.2.1.1	Upgrade Spectrum Scale to version 4.2.3.5
GRIDScaler 3.0.0 – 3.2.0	GPFS 4.1.0.4 – 4.1.1.2	Upgrade GPFS to version 4.1.1.18*
GRIDScaler 1.6.0 – 2.3.0	GPFS 3.5.0.6 – 3.5.0.21	<ul style="list-style-type: none"> Upgrade to GRIDScaler 3.2.0, then upgrade GPFS to version 4.1.1.18* = OR = Upgrade to GRIDScaler 3.2.1 or higher, then upgrade Spectrum Scale to 4.2.3.5

NOTE Spectrum Scale 4.1.1.17 resolves the AFM data corruption issue. However, that version *is subject to a separate bug that can cause file corruption* in the event of an RPC network reconnect. Consequently, DDN recommends upgrading to Spectrum Scale 4.1.1.18 instead.

Step 2. Verify all file checksums between AFM cache and the corresponding home files.

Step 3. For mismatched files, copy the data from cache to home.

ALERT! *Please contact DDN Support for upgrade and data validation 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