

DataDirectTM

NETWORKS



SFA OS

MANDATORY FIRMWARE UPGRADE

Enabling DIF During Host I/O Puts Data Integrity at Risk in Versions 1.5.0 to 1.5.3.0

Technical Support Bulletin

Revision A2
10 June 2014

1.0 Overview

1.1 Issue Summary

When the SATAssure DIF function is first enabled, SATAssure performs checksum initialization on the designated pool. It is possible for a medium error to occur during checksum initialization. If such an error occurs, SFA OS will try to correct it. But the medium error correction routine can mistake the initialization state of DIF blocks within a stripe. Incorrect error detection results, and the attempted recovery actions create further medium errors, which are reported to the host.

GPFS may go offline in response to host-level medium errors.

SFA OS versions 1.5.3.1, 1.5.3.2, and 2.0.0.3 resolve the issue.

ALERT SFA OS 1.5.3.1 is a MANDATORY upgrade for users of SFA OS 1.5.3.0.

1.2 Affected Systems

SFA OS versions 1.5.0.0 through 1.5.3.0 can experience the DIF initialization issue. These versions run on S2A6620, SFA10K, SFA10KE, SFA12K, and SFA12KE storage platforms.

2.0 Problem Description and Root Cause

SATAssure is a feature of SFA OS that detects and corrects silent disk errors during read or write operations. Two verification methods are provided:

- **PCOR** (*Parity Check on Read*), which compares existing RAID parity information with parity calculated during each read to detect errors.
- **DIF** (*Data Integrity Field*), which stores a CRC-type checksum for every 65th block in a pool during writes, then recalculates and verifies the checksum during each read to detect errors.

In the absence of specialized hardware acceleration, DIF is faster than PCOR.

When DIF is first enabled, SATAssure initializes all DIF checksums in the designated pool. Host I/O is supported concurrently with DIF initialization. A “fence” location is updated as initialization progresses, so it is always known which blocks have been DIF-initialized and which have not. If a host I/O operation is requested to a location that has already been DIF-initialized, the operation proceeds normally. If host I/O is requested to a location that has not yet been initialized, SATAssure first performs DIF initialization on the affected block, then allows the host I/O operation to proceed.

If a medium error occurs during checksum initialization, the initialization process — and any host I/O that is waiting on DIF initialization — is suspended while SFA OS tries to

correct the medium error. Single-stripe verify is used to perform error correction. However, the stripe verify routine misinterprets the initialization map bit for a DIF-initialized *block* to apply to the entire *stripe* in which it resides, and assumes that all blocks in the stripe have been initialized when many have not. The verify routine adds uninitialized blocks to the Bad Block list when a checksum mismatch is detected, then forces consistent RAID parity. From this point on, any access to these blocks returns a medium error to the host. The medium errors will persist until new data is written to these blocks from the host.

GPFS may go offline in response to these host-level medium errors.

Refer to defect number DE7021 for details.

3.0 Resolution

This issue is resolved in SFA OS versions 1.5.3.1, 1.5.3.2, and 2.0.0.3.

ALERT SFA OS 1.5.3.1 is a MANDATORY upgrade for users of SFA OS 1.5.3.0.

4.0 Workaround

If upgrading the SFA OS firmware must be delayed, use one of the following workarounds.

- **Use PCOR instead of DIF** when running SATAssure under SFA OS 1.5.3.0 or earlier. (There will be a performance impact if this option is chosen.)
- **Disable all host I/O to the storage system before enabling DIF.** Then use the SHOW JOB command to verify that the background initialization job is complete before resuming host I/O.

A running DIF initialization job will look similar to the example below:

```
RAID[0]$ show job

*****
* Background Jobs *
*****

Idx|Type          |Target (Sub) |State   |Fraction|
-----|-----|-----|-----|-----|
 0 |INITIALIZE    |POOL:0 (NA)  |RUNNING| 0%     |50%
-----|-----|-----|-----|-----|
Total Background Jobs: 1
```

5.0 Contacting DDN Technical Support

Please feel free to contact DDN Technical Support for assistance at any time. Support can be reached by telephone, email, or on the web as follows:

Telephone

DDN Support Worldwide Directory <http://www.ddn.com/support/contact-support>

Email

Support Email support@ddn.com

Mailing List Subscriptions support-tsb@ddn.com

Web

Support Portal <https://portal.ddnsupport.com>

Portal Assistance webportal.support@ddn.com

Technical Support Bulletins <http://www.ddn.com/support/technical-service-bulletins>