

# Optimus.2 Max

## 2.5" SAS, SSD Product Line Generic

## Legal Disclaimer

The Western Digital Corporation or its affiliate's general policy does not recommend the use of its products in life support applications wherein a failure or malfunction of the product may directly threaten life or injury. Without limitation to the foregoing, SanDisk shall not be liable for any loss, injury, or damage caused by use of its products in any of the following applications:

- Special applications such as military related equipment, nuclear reactor control, and aerospace.
- Control devices for transportation equipment including automotive vehicles, trains, ships, and traffic equipment.
- Safety systems for disaster prevention and crime prevention.
- Medical-related equipment (including medical measurement devices).

Accordingly, in any use of SanDisk products in life support systems or other applications where failure could cause damage, injury, or loss of life, the products should only be incorporated in systems designed with appropriate redundancy, fault tolerant or back-up features. Per SanDisk Terms and Conditions of Sale, the user of SanDisk products in life support or other such applications assumes all risk of such use and agrees to indemnify, defend, and hold harmless Western Digital Corporation or its affiliates against all damages.

Security safeguards, by their nature, are capable of circumvention. SanDisk cannot, and does not, guarantee that data will not be accessed by unauthorized persons, and SanDisk disclaims any warranties to that effect to the fullest extent permitted by law.

This document and related material are for information use only and are subject to change without prior notice. Western Digital Corporation or its affiliates assumes no responsibility for any errors that may appear in this document or related material, nor for any damages or claims resulting from the furnishing, performance, or use of this document or related material. absent a written agreement signed by Western Digital Corporation or its affiliates or its authorized representative to the contrary, Western Digital Corporation or its affiliates explicitly disclaims any express and implied warranties and indemnities of any kind that may, or could, be associated with this document and related material, and any user of this document or related material agrees to such disclaimer as a precondition to receipt and usage hereof. Each user of this document expressly waives all guaranties and warranties of any kind associated with this document and/or related materials, whether expressed or implied, including without limitation, any implied warranty of merchantability or fitness for a particular purpose or infringement, together with any liability of Western Digital Corporation or its affiliates and its affiliates under any contract, negligence, strict liability of Western Digital Corporation or its affiliates and its affiliates under any contract, profit or other incidental, punitive, indirect, special, or consequential damages, including without limitation physical injury or death, property damage, lost data, or costs of procurement of substitute goods, technology, or services.

This document and its contents, including diagrams, schematics, methodology, work product, and intellectual property rights described in, associated with, or implied by this document, are the sole and exclusive property of Western Digital Corporation or its affiliates and its applicable subsidiaries ("SanDisk"). No intellectual property license, express or implied, is granted by SanDisk associated with the document recipient's receipt, access and/or use of this document; SanDisk retains all rights hereto.

No work for hire, nor any form of joint ownership, is granted or implied by the document recipient's receipt, access and/ or use of this document.

Any work requested (or implied by the document recipient to be requested) to SanDisk associated with this document and/or its contents, shall be the sole and exclusive property of SanDisk, except to the extent, if any, expressly agreed otherwise by SanDisk in writing referencing this document.

This document and SanDisk's communications to the user associated therewith, shall be treated as SanDisk's proprietary and confidential information, protected by the recipient as such, and used by the recipient only for the purpose authorized in writing by SanDisk. This document shall be covered as SanDisk's confidential information under all applicable nondisclosure agreements between the recipient and SanDisk.

No part of this document may be reproduced, transmitted, transcribed, stored in a retrievable manner, or translated into any language or computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual, or otherwise, without the prior written consent of an officer of Western Digital Corporation or its affiliates.

All parts of the SanDisk documentation are protected by copyright law and all rights are reserved. SanDisk and the SanDisk logo are registered trademarks of Western Digital Corporation or its affiliates, registered in the United States and other countries. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s). Copyright 2016 Western Digital Corporation or its affiliates. All rights reserved.

## Revision History

Date	Revision	Description
November 2016	E	AM91
June 2016	D	Legal Disclaimer updated and Western Digital branding
October 2015	C	AM70
July 2015	B	AM40
April 2015	A	AM35 - Initial Release

Part Number - SKU	Product Description	Firmware Code Name
SDLLOCDR-038T-5CA1	Optimus.2 MAX 3.8TB (0.5 DWPD)	optimus_max_4k_generic_AM91_fw_plus_boot.dob



### ESD Caution – Handling

Static electricity may be discharged through this disk subsystem. In extreme cases, this may temporarily interrupt the operation or damage components. To prevent this, make sure you are working in an ESD- safe environment. For example, before handling the disk subsystem, touch a grounded device, such as a computer case.

## Overview

This document describes the fixes and known issues for the Optimus.2 Max firmware.

### AM91 Firmware

The AM91 firmware contains the following fixes:

**Title: Allow SAS negotiation to continue if reserved fields in the IDENTIFY Frame are set.**

**Failure Scenario:** Following the IDENTIFY sequence between host and drive, SAS negotiation failed to progress further (with the drive resetting the link).

**Root Cause:** A previously-reserved field in the IDENTIFY Address Frame was being verified to contain all 0x00 values. With the usage of the POWER CAPABLE bit field in one of those reserved bytes in newly developed HBAs, the reserved-field check was failing and thus the drive would not acknowledge a valid IDENTIFY frame from the host.

**Change Description:** Suppress reserved-field checking in the IDENTIFY Address frame to allow handshaking to continue.

Likelihood: low

Severity: major

**Title: Port Layer Improvements for Multi-Initiator System Setups**

**Failure Scenario:** Running in a multi-initiator queued-workload environment, drive will occasionally time out while sending status for a command.

**Root Cause:** Context switching and routing between multiple initiators' outstanding workloads could occasionally cause SSP Response Frames to be incorrectly routed, or dropped before delivery.

**Change Description:**

Multi initiator improvements:

- (1) Increased the priority of WWN interrupts in the Port layer interrupt service routine.
- (2) Fixed a possible code hang in the "wait for receive path idle" code. Allow for instances where this path never goes idle.
- (3) Made Port layer improvements to address issues seen running multi-initiator firmware update tests.
- (4) Added logic to detect and record instances where the incoming SAS frame WWN is not what the drive was expecting. Logging this event would indicate an error in the port initiator context control logic.
- (5) Added logic to detect pending WWN interrupts when interrupts are disabled. The new logic temporarily re-enables interrupts to service the WWN.

General port layer updates:

- (1) Made improvements to the Port layer handling of BREAK primitives.
- (2) Removed obsolete recovery logic from the data frame handler.
- (3) Made improvements to the command abort path that eliminates a race condition detected during aggressive hard reset testing.

Likelihood: medium

Severity: major

**Title: Significant performance degradation when data retention recycling begins**

**Failure Scenario:** Reduced performance and long latency commands detected when data retention recycling starts during host commands.

**Root Cause:**

- (1) A multiplier was added to accelerate reclamation recycling when necessary. This multiplier was incorrectly applied to data retention recycling as well.
- (2) Unnecessary reclamation recycling operations can start on blocks that have just completed data integrity recycling.
- (3) At power-on initialization, multiple superblocks can be queued for data retention recycling, even though data integrity recycling does not begin until initialization completes.
- (4) Superblocks with host data and empty superblocks can both hit the data retention limit and require recycling. The old but empty superblock can be selected as the destination superblock, and may immediately require retention recycling after filling.

**Change Description:**

- (1) Reduce the data retention minimum ratio to account for the multiplier used for reclamation recycling.
- (2) Prevent initiation of unnecessary reclamation recycling while waiting for completion of the data integrity recycling of the same superblock.
- (3) Do not queue data retention recycling superblocks during initialization.

(4) If available, choose an unused superblock from the data retention recycle list to process before selecting superblocks containing host data.

Likelihood: low

Severity: minor

**Title: Performance degradation when processing UNMAP commands**

**Failure Scenario:** Processing a large UNMAP command degrades host workload performance by an unacceptable amount.

**Root Cause:** Within the FTL, unmap commands are broken up into smaller subtasks and scheduled in a loop that includes other tasks, including host commands. When successive unmap subtasks are being sent to the scheduling loop, host commands do not get enough time to run.

**Change Description:** Change the scheduler to skip the unmap task a fixed number of times when it appears in the scheduler loop, to allow enough time to process host commands. This change only affects the scheduling loop when an unmap task is present.

Likelihood: medium

Severity: minor

**Title: Multi initiator frame misdirection**

**Failure Scenario:** In a multi initiator system the drive could direct an XFER\_RDY to the wrong initiator leading to the command being aborted for an Initiator response time out.

**Root Cause:** A corner case was found in the connection context switching code when sending XFER\_RDYs for write data. If the lookup of the WWN for the initiator was delayed for more than 810uSec the drive would use the wrong value. This caused the drive to send out an XFER\_RDY to the WWN saved at index 0 of the WWN table.

**Change Description:** Eliminated multiple sources of delay that could cause the 810uSec internal timer to expire in error during context switching.

Likelihood: low

Severity: major

**Title: Invalid tag overlap reported**

**Failure Scenario:** In a multi initiator environment where all the active Initiators are using a small number of similar command tag values the drive could reject an incoming command as a tag overlap when no actual overlap existed.

**Root Cause:** The firmware responsible for removing tags from the active tag table was using an index value that was not always valid in a multi initiator system.

**Change Description:** Fixed the removal of tags from the active tag table to ensure the use of the valid index.

Likelihood: low

Severity: major

**Title: Link Counters reported incorrectly**

**Failure Scenario:** SAS Link counter registers become saturated in a multi port drive environment. The counters may be reported by the drive incorrectly. It is also possible that the drive performance becomes degraded.

**Root Cause:** The drive controller receives an interrupt when a Link counter reaches saturation. The drive logic to process the saturation event used a stale pointer to access the PHY control registers. If the stale pointer happened to be set to the wrong port's PHY control registers, then the link counter saturation event would not be handled properly, resulting in incorrect Link event counters. Also, the interrupt for Link Counter saturation would be repeated for each counter increment, which could affect drive performance due to excessive interrupt handling overhead.

**Change Description:** Ensure that the PHY control register pointer is updated properly before being used in Link counter saturation processing.

Likelihood: medium

Severity: major

**Title: GET LBA STATUS does not handle an ALLOCATION LENGTH < 24 correctly**

**Failure Scenario:** The host issues a GET LBA STATUS command with an ALLOCATION LENGTH less than 24.

**Root Cause:** When the ALLOCATION LENGTH field has a value less than 24, the drive erroneously returns a value of 4 in the parameter data's PARAMETER DATA LENGTH field instead of an allowed value.

**Change Description:** The value returned in the PARAMETER DATA LENGTH field has been corrected. The drive ensures that at least one LBA status descriptor is returned.

Likelihood: medium

Severity: minor

**Title: Add Periodic Logging of FTL Information to the Drive Event Log**

**Change Description:** Periodically record various FTL information in the event log in order to facilitate debug of performance issues. Information captured includes the following:

- Recycle counters
- Spare/Erase block counters
- Recycling time information

Information is recorded to the event log every 20 minutes if there has been related activity.

**Title: Enable Periodic Logging of Front End Information to the Drive Event Log**

**Change Description:** Periodically record various Front End information in the event log in order to facilitate debug of performance issues. Information captured includes:

- Queue depth
- Average and maximum command time information
- Command counters
- Prefetch counters

Data is output to the event log every 20 minutes after power up and upon F7 Dump command if there has been related activity.

**Title: SANITIZE with OVERWRITE service action and IMMEDIATE=1**

**Failure Scenario:** SANITIZE with OVERWRITE service action and IMMEDIATE=1 results in inconsistent LBA data.

**Root Cause:** In the specific case of SANITIZE with OVERWRITE and IMMEDIATE=1, the initialization pattern is read from the wrong place, and read cache is not invalidated, causing inconsistent results.

**Change Description:** Use the correct reference for the initialization pattern buffer, and invalidate the read cache.

Likelihood: low

Severity: critical

**Title: SANITIZE command hangs with a Sanitize operation in progress**

**Failure Scenario:** SANITIZE issued with a Sanitize operation in progress results in a drive hang.

**Root Cause:** Status is not being set correctly for a SANITIZE command that is received while a Sanitize operation is in progress.

**Change Description:** When a Sanitize operation is in progress, respond to a new SANITIZE command with LUN NOT READY, SANITIZE IN PROGRESS (02/04/1B) status.

Likelihood: low

Severity: major

**Title: WRITE SAME command timeout**

**Failure Scenario:** A WRITE SAME command did not return status during high queue depth traffic.

**Root Cause:** A WRITE SAME command did not complete during high queue depth traffic because of resource starvation.

**Change Description:** Properly reschedule a WRITE SAME command's execution when resource becomes available.

Likelihood: low

Severity: major

**Title: Power failure may result in format corrupt**

**Failure Scenario:** Format corrupt after power failure.

**Root Cause:** A loop within a specific function can take 200 msec to complete and power failure was not checked during the loop. Depending upon timing, there may not be time to SCRAM data.

**Change Description:** Check for pfail within the function loop and terminate when pfail activated.

Likelihood: low

Severity: major

**Title: Write performance could suffer after several HARD RESET events**

**Failure Scenario:** Following several HARD RESET events, the drive's write performance was degraded.

**Root Cause:** HARD RESETs occurring while small write buffers are owned by the SAS port could result in the buffers being lost until the drive is power-cycled. The drive's HARD RESET processing was not properly bookkeeping all cases of small transfer buffers primed into the ports through a HARD RESET.

**Change Description:** Re-sequenced the FIFO resets so that buffers being released by the port were not lost during the HARD RESET processing.

Likelihood: medium

Severity: major

**Title: Drive became unresponsive following firmware download**

**Failure Scenario:** Multi-Initiator code download testing (with high IO volume). The drive stopped responding to SAS commands, resulting in command timeouts.

**Root Cause:** The code download processing was waiting for the SAS Link to be IDLE before proceeding to boot-up with new the code-image. In the failure scenario, the SAS Link was in a suspended state, and the IDLE condition was never met, which prevented the drive from loading the new code-image, and caused pending commands to time-out.

**Change Description:** Updated code download processing to suspend the SAS port's Link state and wait an additional 500us for all active connections to close. This eliminates possible 'race' conditions and properly handles the transition to the new code.

Likelihood: low

Severity: major

**Title: Unaligned 4K Block Write Transfer Failure Fix**

**Failure Scenario:** After a hotswap an unaligned 4K block write failed to transfer data properly.

**Root Cause:** Incorrect LBA bitmask management.

**Change Description:** Corrected the LBA bitmask for unaligned 4K block writes.

Likelihood: medium

Severity: major

**Title: Unmap interrupted by power fail can result in read-only drive**

**Failure Scenario:** Power failure during unmap resulted in read-only drive.

**Root Cause:** The unmap batches up the valid page count (VPC) update to the super block at the end of function. If there is a pFAIL before this happens, then there is a mismatch in VPC from the persistent info (the wrong value) on the next power up as init does not replay the VPC updates of the super blocks.

**Change Description:** Replay the VPC updates if unmap was interrupted by a pFAIL after the L2P was updated but before the VPC updates.

Likelihood: low

Severity: major

**Title: Power failure handling of partition load**

**Failure Scenario:** Format corrupt occurred after power failure.

**Root Cause:** FTL code was searching for properly loaded L2P partitions when power failure was detected which delayed the start of SCRAM.

**Change Description:** Terminate the routine if pfail detected so SCRAM operations can begin.

Likelihood: low

Severity: major

**Title: SAS link reset after firmware download**

**Failure Scenario:** During a firmware download, a link reset occurred.

**Root Cause:** The phy registers were re-programmed during the start-up with the newly downloaded firmware. The hardware register was cleared before writing the desired value. The brief incorrect register values caused some host adapters to bring the link down.

**Change Description:** Changed the post-firmware-download start-up behavior to not re-program the phy registers. The registers are programmed after the next power cycle.

Likelihood: low

Severity: major

**Title: Drive returns incorrect WWN hash address in Read data**

**Failure Scenario:** Drive returns incorrect WWN hash address in Read data in a high-traffic, high queue depth, multi-initiator workload environment.

**Root Cause:** 1. When two incoming SAS frames from two different initiators are received back-to-back, the WWN hash from the second frame's hardware register may overwrite the value from the first frame. 2. The WWN index table was not cleared during LUN Reset, resulting in a stale entry being used.

**Change Description:** 1. Detect the race condition and use the correct WWN value from the index table. 2. When a LUN Reset is received, clear the hash hardware register and clear the WWN index table to force the drive to re-populate the WWN table when an initiator logs in.

Likelihood: low

Severity: major

**Title: Command timeout occurs in workloads with overlapped tags**

**Failure Scenario:** In queued workloads with multiple initiators, where overlapped tags are issued for outstanding commands, the drive may stop responding to I/O commands.

**Root Cause:** In the failure scenario, the SAS Port layer logic identified a tag overlap condition, but the device server logic did not find the associated command to abort. A problem with the abort processing in this case caused commands from the wrong nexus ID to be aborted, resulting in command timeouts to be seen by the host for those commands.

**Change Description:** Addressed the flaw in overlapped tag abort handling to correctly identify the Nexus ID of commands to abort.

Likelihood: low

Severity: major

**Title: Data Abort exception occurs during code download**

**Failure Scenario:** In a multi-initiator environment, while the drive was processing high volume I/O, a code download occurred, and the drive encountered a Data Abort exception and stopped responding to commands, resulting in command timeouts, and code download failure.

**Root Cause:** Due to a command exception, the firmware had performed code download cleanup processing when it was in a state where it had already committed to flash programming. When flash programming processing started, the Data Abort exception occurred.

**Change Description:** Download processing was modified to ensure download cleanup does not occur if flash programming has not yet occurred, and the drive is committed to flash programming.

Likelihood: low

Severity: major

**Title: Firmware Download Error**

**Failure Scenario:** Firmware download may not complete.

**Root Cause:** During firmware download the hardware auto-transfer-ready buffer management detects buffer loss due to a firmware error.

**Change Description:** Firmware has been fixed to properly manage auto-transfer-ready buffers during firmware download.

Likelihood: low

Severity: major

**Title: Drive does not reject SANITIZE command with write protected medium**

**Failure Scenario:** A SANITIZE command issued when the medium is write protected does not fail with data protect (07/xx/xx) status.

**Root Cause:** Drive SANITIZE processing was missing a check for the write protection case.

**Change Description:** Fail a SANITIZE command with the appropriate data protect status (e.g. 07/27/00) if the medium is write protected.

Likelihood: low

Severity: major

**Title: Unmap during initialization causes assert and hang**

**Failure Scenario:** Unmap command issued immediately after power cycle resulted in hang.

**Root Cause:** During initialization, unmap commands that were being processed prior to the L2P partition file completing can cause an assert and drive hang.

**Change Description:** Defer the processing of unmap command till the partition table is completely built.

Likelihood: low

Severity: major

**Title: Drive not ready after interrupted SANITIZE OVERWRITE**

**Failure Scenario:** A SANITIZE OVERWRITE operation that is interrupted by a power fail cannot be restarted.

**Root Cause:** In the case of a SANITIZE OVERWRITE that is interrupted by a power cycle, the state of the medium will be set to Not Ready, which prevents writes from completing, and effectively prevents SANITIZE OVERWRITE from being restarted.

**Change Description:** Set the medium state correctly for interrupted SANITIZE OVERWRITE operation after power on.

Likelihood: low

Severity: major

**Title: Format corrupt after power cycle**

**Failure Scenario:** Power cycle during Data Retention Recycling can lose parity and go format corrupt.

**Root Cause:** The failure is caused because we don't have parity data in order to recover a page in a partially written stripe during init reads of open super blocks. We don't need the parity data if we can recover the page with just doing read retry.

**Change Description:** Even though we don't have parity data for the stripe, attempt a read retry to recover the page.

Likelihood: low

Severity: major

**Title: WWN mismatch when running Multi-Initiator I/O**

**Failure Scenario:** When running multi-initiator I/O at high queue depth, Initiator Response timeouts occurred.

**Root Cause:** Back-to-back frames in the hardware from different initiators caused the hardware to mistakenly change the hash address of the previous frame to that of the incoming frame. This was a hardware bug.

**Change Description:** A hardware feature was invoked to detect the hash mismatch and have the firmware correct the hash address.

Likelihood: medium

Severity: major

**Title: FORMAT command doesn't complete**

**Failure Scenario:** Repeated FORMAT commands can consume resources and result in a FORMAT command not completing.

**Root Cause:** When the format is performed, buffers from host die operation history are not released. So after many formats, write buffers are exhausted and a hang results due to insufficient resources.

**Change Description:** Properly release history buffers when formatting the drive.

Likelihood: low

Severity: major

**Title: Drives not admitted into Storage Spaces in a Microsoft Windows environment**

**Failure Scenario:** Drives were not being admitted into Storage Spaces in a Microsoft Windows environment.

**Root Cause:** Microsoft Windows was not recognizing drives which reported a value of greater than 3 in the LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT field for READ CAPACITY.

**Change Description:** Limit LOGICAL BLOCKS PER PHYSICAL BLOCK EXPONENT reporting to 3 for READ CAPACITY.

Likelihood: medium

Severity: major

**Title: Logical Block Application Tag mismatches are not detected**

**Failure Scenario:** Logical Block Application Tag mismatches are not detected by 32-byte SCSI commands.

**Root Cause:** The firmware inadvertently fails to propagate its calculated Application Tag check flag to the hardware.

**Change Description:** The firmware programs its calculated Application Tag check flag into the hardware.

Likelihood: low

Severity: major

**Title: Incorrect sense data in check condition for Write with zero byte payload**

**Failure Scenario:** Write with injected zero-length data payload. Drive was incorrectly reporting 0B/4B/07/01 (BREAK ENCOUNTERED).

**Root Cause:** Incorrect sense data was being returned by drive

**Change Description:** Changed sense data to be 0B/0E/00/00 (INVALID INFORMATION UNIT).

Likelihood: low

Severity: minor

**Title: Fix and enable SCRAM monitoring during code load**

**Failure Scenario:** SCRAM monitoring for code load was previously removed due to an issue causing drive hangs.

**Root Cause:** Attempts to access shared memory by both CPUs resulted in an AXI bus hang.

**Change Description:** Coordinate the timing of CPU restarts such that each processor can receive data from shared memory when requested and restart when safe. Enabled SCRAM monitoring for code load with resolved CPU coordination.

Likelihood: low

Severity: minor

**Title: Program and erase failures may cause critical event**

**Failure Scenario:** Injected random write and erase failures (400 in a row) cause critical event to be logged.

**Root Cause:** Log writes and host writes save sequence numbers when programmed and a program failure can alter what is recorded.

**Change Description:** Make sequence numbers be consistent between memory versions and cached versions for failure cases.

Likelihood: low

Severity: minor

**Title: Poor performance seen in sequential workloads**

**Failure Scenario:** Performance issues seen in 10/90 R/W sequential workload.

**Root Cause:** Recycle logic chose poor candidates to recycle when sequential workloads were invalidating 90% of the superblocks. This caused higher write amplification and low performance.

**Change Description:** Changed the sparse pool delineations to better identify low valid page count superblocks, and allowed logic to choose alternative subdrives and superdevices when much better choices existed.

Likelihood: medium

Severity: minor

**Title: Average number of block erases reported incorrectly**

**Failure Scenario:** Average number of block erases reported incorrectly in log page 31h.

**Root Cause:** Calculation did not account for the number of blocks per die.

**Change Description:** Fixed the average block erase and average page program calculations.

Likelihood: high

Severity: minor

**Title: Reduced performance after quick repeated power cycles**

**Failure Scenario:** Hot swap during write workload can produce drive logs with many critical events logged, which could lead to reduced performance.

**Root Cause:** When drive is power cycled frequently before background initialization completes, incoming host writes could consume free blocks and cause drive to enter critical reclamation mode.

**Change Description:** Change the recycle-during-background-initialization trigger threshold and allow recycle writes during background initialization.

Likelihood: low

Severity: minor

**Title: Excessive data retention recycling reduces performance**

**Failure Scenario:** After an extended idle period, lower performance is observed when the workload resumes.

**Root Cause:** Data retention recycling was performed on superblocks that did not require immediate recycling.

**Change Description:** Select superblocks that will not require data retention recycling immediately.

Likelihood: low

Severity: minor

**Title: Reduced performance after repeatedly reading the same block**

**Failure Scenario:** Repeatedly reading the same block may lead to reduced performance.

**Root Cause:** Repeatedly reading the same block forces read disturb recycling, and will eventually move the superblock to a separate recycle list. The search for a recycle candidate can miss this entry.

**Change Description:** Consider alternate super devices when recycling; select a recycle candidate from the read disturb recycle list when available.

Likelihood: low

Severity: minor

**Title: Command timeout may occur during heavy recycling**

**Failure Scenario:** Write command timeout detected after idle period.

**Root Cause:** Performing different types of recycling concurrently can consume available resources and prevent host writes from completing.

**Change Description:** Reduce number of concurrent recycle operations.

Likelihood: low

Severity: minor

**Title: Persistent reservation key lost**

**Failure Scenario:** - Initiator 0 attempts and successfully makes a persistent reservation. - Initiator 0 successfully reads back the persistent reservation key. - Initiator 1 attempts to make a persistent reservation and is rejected. - Initiator 0 unsuccessfully reads back the persistent reservation key a second time. The persistent reservation key is reported as lost.

**Root Cause:** A code path was incorrectly decrementing a reservation count when a persistent reservation was rejected.

**Change Description:** Correctly decrement the reservation count when an initiator unregisters it.

Likelihood: low

Severity: minor

**Title: Sequential Read Prefetch Improvements**

**Failure Scenario:** Queued multiple overlapping sequential read streams (e.g. multiple host threads all issuing a sequential read stream from the same start LBA). Performance across the range of queue depths may be lower than expected.

**Root Cause:** In queued multiple overlapping read stream workloads, the earlier-issued streams can get ahead of the later-issued streams in LBA space. In such a workload, the prefetch algorithm would sometimes re-trigger at a lower LBA to service one of the trailing streams instead of maintaining prefetch ahead of the forward-most stream. This resulted in the same data being prefetched multiple times and a degradation in performance.

**Change Description:** A change to the prefetch algorithm was made to not trigger a new prefetch if there has been a recent hit on the current prefetch data stream. In the high queue depth, multiple overlapping read stream case, prefetch is now maintained ahead of the forward-most stream.

Likelihood: medium

Severity: minor

**Title: Incorrect error code may be returned for uncorrectable read**

**Failure Scenario:** After WRITE LONG is used to create an unrecoverable block, a subsequent READ may result in a DATA BCRC ERROR DETECTED (0B/47/82) error instead of a MARKED BAD BY APPLICATION CLIENT (03/11/14) error if the bad block is not at the beginning of the requested read range.

**Root Cause:** The unrecoverable code path was different depending upon whether the host data transfer had started or not. In the case of when the transfer has started, there was debug code that incorrectly caused the wrong error code to be returned.

**Change Description:** Removed the debug code.

Likelihood: medium

Severity: minor

**Title: Long initialization time**

**Failure Scenario:** Sparse writes followed by power cycling can result in long initialization time.

**Root Cause:** Sparsely written data generates metadata that requires a lot of re-arrangement and copying at init time.

**Change Description:** Improve meta-data processing during data block recycling by utilizing hardware-assisted DMA copy to speed-up the copying process time.

Likelihood: low

Severity: minor

**Title: Hang after power cycle during I/O**

**Failure Scenario:** Hang could occur after power cycle while running mixed I/O workloads with UNMAP commands.

**Root Cause:** The log recycle metadata buffer is incorrectly addressed and could result in lost buffers when aborting and restarting a metadata read operation.

**Change Description:** When restarting log recycling, operate on the correct logging superblock reference.

Likelihood: low

Severity: minor

**Title: GET LBA STATUS can report the wrong provisioning status soon after startup**

**Failure Scenario:** GET LBA STATUS commands issued immediately after firmware startup can report a provisioning status of Mapped for LBAs that should be reported as Anchored.

**Root Cause:** The wrong criteria is used to decide whether a logical page's context has been loaded during flash layer initialization, such that pages with incomplete context are reported as Mapped.

**Change Description:** Use the proper criteria to decide whether the provisioning status for a logical page can be reported.

Likelihood: low

Severity: minor

**Title: HLBA error on READ command**

**Failure Scenario:** HLBA mismatch error on read error recovery during mixed workloads.

**Root Cause:** During completion of host reads using media queues, both CPUs can simultaneously access a shared memory structure. The Flash layer CPU doesn't post a completion message to the Front End CPU in case of media queue reads, assuming that the Front End CPU is notified of the completion by the port layer. This leads to a window in which both CPUs can access the structure simultaneously.

**Change Description:** The Flash layer will make a copy of the required information in shared memory when the command is initially received from the Front End CPU, so that the access collision will not occur at completion time.

Likelihood: low

Severity: minor

**Title: Incorrect drive handling of NAK error on XFER\_RDY**

**Failure Scenario:** When a NAK was received in response to an XFR RDY from the drive, the drive responded with 0B/47/01/01 (Data Phase CRC Error Detected - NAK Sent). This is incorrect. It should be 0B/4B/04 (NAK Received).

**Root Cause:** The drive was waiting for the host to send all the data and then responded with Data Phase CRC Error Detected - NAK Sent.

**Change Description:** Fixed code to respond with proper sense data when a NAK is received for an XFER\_RDY.

Likelihood: low

Severity: minor

**Title: Implement GET LBA STATUS command**

Failure Scenario: N/A

Root Cause: N/A

**Change Description:** Implemented GET LBA STATUS command.

Likelihood: N/A

Severity: N/A

## AM70 Firmware

The AM70 firmware contains the following fixes:

- Implemented Read/Write (transitional) performance improvement.
- Fixed an issue where a TRIM at PFAIL along with Host Writes to the same LBA's can cause Read Only drives. Fixed an issue where the drive ignores initiator Open Request and becomes non-responsive during link up/down testing. Enabled normal link operations by setting the 'kReset\_IDX\_Register' bit in the Link Control Register after an Abandon Class Primitive is received.
- Fixed an issue where Read Capacity 10, byte 1 and byte 8, bit 0x0 obsolete field must return a Good Status. Fixed an error in MC\_Release\_Rx\_Resources that mis-handled the abort of a non-block data command with an RX AES context slot assigned to it.
- Fixed an issue found during internal stress testing where a Critical Event was noticed during Notify Power Loss Expected Primitive insertion between Write Data.
- Fixed an issue when a synchronize cache(16) returned GOOD status when the LBA specified exceeded the Maximum LBA of the drive.
- Fixed an issue that caused a Request Sense command to increment the Non-Medium Error Log Page counter (06h).
- Added OP (Over Provision percentage), DWPD (Drive Write Per Day) and EOL PE (End-of-Life Program Erase) count to Inquiry page E9h.
- Fixed an issue found during internal stress testing where Good Status was being returned even if the OPEN for Status frame was being denied longer than the IT Nexus Loss timeout.
- Fixed an issue where percent of life used was not being reported correctly in Log Page 11h immediately following a power cycle.

## Known issues with AM70

Optimus.2 Max drives have the following known issues:

- SanDisk FWQA testing indicates a potential of seeing a port lockup condition in multi-initiator environments with high stress levels. These issues are still under investigation at SanDisk and will be driven to root cause and a solution. SanDisk currently believes these lockups would not be seen within reasonable operating conditions and are a result of the extreme conditions presented by the FWQA regression tests:

Test Environment: Multiple Initiators, through an expander to the drives while subjecting the drives under test to:

- Varied I/O workloads submitted to the drive while injecting the following:
  - Varied queue depths 1, 8, 16, 32, 64, 126 (depth on each port). Task Set Full conditions at QD 126.
  - Extra traffic – random non-block commands (non-block commands including illegal requests).
  - Hard Resets injected at fixed intervals / Cable Pulls via a hardware signal controller.
  - Abort task set injected.
  - Firmware update testing.

## AM40 Firmware

The AM40 firmware contains the following fixes:

- Corrected the Drive Life-Used calculation to correctly report endurance for Optimus MAX.
  - Prior to this fix, life-used would report less than actual.
- Implemented a fix related to the detection of pFail that could result in a delayed response resulting in the drive going format corrupt.
- Added additional support for S.M.A.R.T. trip reporting via a unit attention.
- Corrected a port layer logic error that left Notify Enable Spin-up primitive detection disable after a hard reset, resulting in the drive reporting a 02/04/11 sense code.
- Corrected an issue where IT\_Nexus\_Loss for a link down could abort commands on the other port.
- Enabled support for Rate Matching.
- Fixed Rate Matching logic so that it will restore the Link state for normal operations when a RETRY EXHAUSTED (or CONNECTION DENIED with no extended status) event happens.
  - This will keep the drive from going unresponsive.

## Known issues from firmware AM35 fixed with firmware AM40

- Fix IO Hang situation seen during internal FWQA regression testing.
- Fixed Multi-Port Stress failure seen during internal FWQA regression testing
- Implemented a fix for a data validation failure seen during FWQA power cycle regression testing.
- Fixed an issue of the drive going unresponsive during FWQA regression Link Up/Down Single/Multi Initiator testing.

## AM35 Firmware

The AM35 firmware contains the following fixes:

- Fixed an issue that was causing us not to report the correct CPU PC address in the scram margining area.
- Implemented a fix to correct an issue where we were getting a superdevice state for an unwritten page. This led to an assert for a request for a superblock for an invalid physical address.
- Implemented a fix that will now disable TART until fast initialization is fully complete.

## Known issues with AM35

Optimus.2 Max drives have the following known issues:

- Failure with IOHung in I/O script, HARD Reset & TMF with extra traffic in the Oakgate tester. These are not typical I/O workloads.
- Multi-Port Stress Failure seen on one port running I/O while second port was sent a combination of additional I/O commands, non-block commands, and illegal command requests.
- Data validation failure on a drive running srt powercycle Seq 32K-50% from start.
- Drive became unresponsive during Link Up/Down Single/Multi Initiator.