

Optimus.2 Eco/Ascend/Ultra/Extreme



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.

Disclaimer

The SanDisk, a Western Digital brand 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 automotive vehicles, train, ship 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 SanDisk, a Western Digital brand and 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 is for information use only and is subject to change without prior notice. SanDisk, a Western Digital brand 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 SanDisk, a Western Digital brand or its authorized representative to the contrary, SanDisk, a Western Digital brand 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 SANDISK, A WESTERN DIGITAL BRAND AND ITS AFFILIATES UNDER ANY CONTROL, NEGLIGENCE, STRICT LIABILITY OF SANDISK, A WESTERN DIGITAL BRAND 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.

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 SanDisk, a Western Digital brand.

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 SanDisk, a Western Digital brand, 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 SanDisk, a Western Digital brand. All rights reserved.

SanDisk, a Western Digital brand
951 SanDisk Dr.
Milpitas, CA 95035

www.sandisk.com

An ISO 9001 certified company.



SanDisk

Revision History

Date	Revision	Description
May 2016	D	F910
October 2015	C	F820
June 2015	B	F760 (Initial release for Eco E3.2)
August 2014	A	F6C2
April 2014	1	F670 Initial Release

Part Number - SKU	Product Description	Firmware Code Name
SDLKODDR-400G-5CA1	Optimus.2 Eco 2.5" 400GB (3 DWPD)	optimus2_eco_generic_F910_customer.dob
SDLKOCDR-800G-5CA1	Optimus.2 Eco 2.5" 800GB (3 DWPD)	
SDLLOCDR-016T-5CA1	Optimus.2 Eco 2.5" 1600GB (3 DWPD)	
SDLLOCDR-020T-5CA1	Optimus.2 Eco 2.5" 2000GB (1 DWPD)	
SDLKODDR-480G-5CA1	Optimus.2 Eco 2.5" 480GB (1 DWPD)	
SDLKOCDR-920G-5CA1	Optimus.2 Eco 2.5" 920GB (1 DWPD)	
SDLKOEEM-200G-5CA1	Optimus.2 Ascend 2.5" 200GB (10 DWPD)	optimus2_generic_F910_fw_plus_boot.dob
SDLKODDM-400G-5CA1	Optimus.2 Ascend 2.5" 400GB (10 DWPD)	
SDLKOCDM-800G-5CA1	Optimus.2 Ascend 2.5" 800GB (10 DWPD)	
SDLLOCDM-016T-5CA1	Optimus.2 Ascend 2.5" 1.6TB (10 DWPD)	
SDLKOE9W-150G-5CA1	Optimus.2 Ultra 2.5" 150GB (30 DWPD)	
SDLKODGW-300G-5CA1	Optimus.2 Ultra 2.5" 300GB (30 DWPD)	
SDLKOCGW-600G-5CA1	Optimus.2 Ultra 2.5" 600GB (30 DWPD)	
SDLLOCGW-012T-5CA1	Optimus.2 Ultra 2.5" 1.2TB (30 DWPD)	
SDLKOE9W-100G-5CA1	Optimus.2 Extreme 2.5" 100GB (45 DWPD)	
SDLKOD9W-200G-5CA1	Optimus.2 Extreme 2.5" 200GB (45 DWPD)	
SDLKOC9W-400G-5CA1	Optimus.2 Extreme 2.5" 400GB (45 DWPD)	
SDLLOC9W-800G-5CA1	Optimus.2 Extreme 2.5" 800GB (45 DWPD)	

Overview

This document describes the fixes and known issues for the Optimus.2 Eco/Ascend/Ultra firmware.

F910 Firmware contains the following fixes:

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 incorrect 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: 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

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

F820 Firmware

The F820 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 where the internal value used to calculate percent_life_left was incorrectly using 10 DWPD rather than 1 DWPD value. (Eco only)
 - Note: Although the value will decrease to the correct value, the drive will continue to function as advertised from a Total Bytes Written/DWPD perspective.
- 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. (Eco only)
- Fixed Mode Page E8h.
- Fixed an issue that caused false negatives for bad PI Ref Tags.

Known issues with F820

Optimus.2 Eco/Ascend/Ultra/Extreme 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

F760 Firmware

The F760 firmware contains the following fixes:

- Implemented a fix for an issue that was causing a command timeout event.
- Fixed an issue where low OP drives have incorrect database maximums.
- Added a safeguard against a memory allocation failure to prevent drives from hanging.
- Changed the default global SCRAM timeout from 5 milliseconds to 6 milliseconds.
- Added a persistent non-volatile counter to track single-bit ECC corrections.

- Enabled write limiter by default to reduce write fragmentation that improves write performance at high QD and xfer-size.
- Added changes to disable rate matching support and increase queue depth support to 252.
- Fixed Logical Block Reference Tag error being falsely reported on 32 byte read command.
- Modified port behavior so that after the new firmware is written to the SPI and before the new firmware gets activated, the port sends out BUSY to all requests from the initiator. We can now safely come back up with the new firmware and continue as normal.
- Fixed an issue with sense data errors in Send diagnostic command.
- Fixed an issue to check the ParameterListLength field in the Unmap CDB.
- Changed Persistent Reservation Out command to correctly ignore non 0 Scope values for Service Actions, Register, Register and Ignore, and Clear.
- Fixed an issue where a response to a read command issued to flash was not being handled correctly while Servicing a verify operation.
- Changed Persistent Reservation Out command to check Allocation Length and return kPARAMETER_LIST_LENGTH_ERROR if it is !=24 bytes.
- Ensured that overlap_chain is always flagged as invalid on all commands, even immediately after reboot when a slot is being used by a read.
- Avoided faulty ACK/NAK Time Out (ANTO) on certain retries exhausted cases.
- Avoided early ReadQueue transmits from accessing unexpected DDR location causing single and double bit DDR errors.
- Resolved an issue where by recovering from a failed read, we read the data from DDR but don't take into account the allocation unit of the recovery.
- 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.
- Moved process_unmap_region back into fast RAM to restore its performance.
- Added safeguard against memory allocation failure to prevent drives from hanging.
- Fixed an issues where not removing tags after errors caused overlapped commands.
- Fixed and issue for SCRAM margining on software reset and code downloads.
- Fixed an issue for Persistent Reservation Out issues seen in Microsoft "SCSI-3 Persistent Reservation requirements test" and "Validate Storage Spaces Persistent Reservations test". Added code additions to the handling of "Register and Ignore" and "Preempt".
- Fixed an issue that caused HEAD OF QUEUE commands to block other SIMPLE commands.
- Added a check to ensure a correct response when a parameter list length resulted in truncation.
- Implemented a WriteLimiter functionality which prevents Write commands from starting if there is too much write data already in flight.
- Fixed an issue where an Ack/Nak timeout on a read could cause IT Nexus loss instead of reporting Aborted Command.
- Shortened the host data transfer lengths for 4Kb+ block size in social_regression_pi2.py that can cause Windows to blue-screen.
- Repaired an inadvertent change to the TOKENIZED_PRINT that generates the FLASH_EXERCISER_ROW entry in the common diagnostic output.
- Added a compile time assert to check whether MAX_DEFECTS_IN_DB can actually guarantee a certain percent of erase blocks.
- Increased the threshold for doing 100% recycles.
- Increased the CUTOFF_AGE_OF_PENDING_LOGICAL_PAGE to take into account the size of the logical pages.
- Increased the precondition of power cycle tests to prevent event injection failure.
- Added supported builds to Makefile and improved layout of help.
- Enabled support for Rate matching. Fixed ITNL and IRTOs in multi-initiator environment with mixed read and write traffic.
- Fixed an issue found in multi-initiator TMF testing.

- Added a feature which will allow events in the event log to consume less space, which will allow for more events to be stored in the log.
- Added support for temperature throttling on PCB-based NAND sensors.
- Set the subpage write caching feature to enable by default.
- Resolved an issue related to stranded bytes in the hardware FIFO that could result in an 03/0C/00/06 sense key.
- Disabled the scram margining feature which caused the erasing of the area after each code load.
- Reduced code load time.
- Fixed an issue with the SMART min/max data being unnecessarily saved after every power cycle. It should be saved only when one of the min or max values change.
- Changed the SDF ID of the saved write amplification min/max values so that any previous (and false) values will be forgotten.
- Fixed an issue where Read Buffer mode 0 was returning an incorrect status.
- Resolved an issue with the SPI flash where a sector was stuck in a state that appeared to be erased, but writing to that sector caused program errors. This would happen if the drive lost power while that flash section was being erased.
- Fixed an issue where the Read and Verify counters were getting modified due to false unrecovered errors introduced by Write Long command.
- Made modifications that provide steadier performance at low queue depths.
- Made changes so that the drive checks for data alignment for the Write Buffer command.
- Added support to notify CPU0 of recycle failure.
- Added support for a SMART trip in the event of a SPI flash write failure.
- Fixed an issue to protect the SCSI format against power failures.
- Fixed an issue of CPU0 data aborts in the event log handler when 15 parameters are used.
- Fixed an issue to give INITIATOR RESPONSE TIMEOUT priority over IT NEXUS LUN TIMEOUT when they are set the same.
- Added updates for high queue depth and PI Stripping.
- Avoided false PI Stripping detection and could lead to false ITNL and bad link behavior.

Known issues with F760

- Drive failed during SRT Powercycle.
- Drive hit a mismatched Valid Page Count for Closed SB failure during FlashManagement-FI_WriteInterrupt Test.
- Drive failed with ERROR_DATA_OVER_UNDER_FLOW during TMF multi initiator tests using Oak Gate.
- Drive failed with a IO Hung Error during IO Multi-Initiator tests using Oak Gate.
- Drive became inactive or offline during Cable Pull signal only.
- When formatting with Type 2 PI, under rare conditions the drives may cause a command timeout on reads.

F6C2 Firmware

The F6C2 firmware contains the following fixes:

- Fixed an issue with Send Diag Selftest that caused various data errors on verify, read, and write verify commands while running command loop tests.
- Improved performance boost when reading and verifying unmapped LBAs.
- Fixed an issue in the format path that was causing LBA counts to be truncated.
- Fixed an issue in the detection of reserve byte/bit violations in performance commands.
- Fixed an issue so that the drive does not return a Reservation conflict if the Standard Reservation holder issues a second Reserve command.
- Fixed an issue where the drive would return an unexpected 0B/44/00/03 DELIVERY_TIME_LIMIT_EXCEEDED after having run large IOs on commands that may take longer than 100ms.
- Fixed an issue that caused the drive to return invalid 05/24/00 errors under mixed workloads using 10 and 16 byte commands.

- Removed un-needed 5 sec delay from Send Diag Short selftest.
- Fixed an issue with the handling of skipwrites and overlap detection.
- Added 48-bit LBA support to overwrite.
- Fixed an issue with Write Buffer command for a transferlength of 0.
- Improved write bandwidth at 16K to 32K transfer sizes by increasing AXR allocation from 8K to 16K and by disabling AXR on xfers larger than 16K.
- Added support for an RSA signature on firmware images.
- Fixed an issue that was causing Write Same to fail.
- Added support for 4096 + PI2 sector size.
- Made performance enhancements to WriteSame in uninitialized/unmapped pages.
- Modified the code to reject a Write Buffer command if the allocation length is larger than the Write Buffer Max Download Size.
- Changed percent drive life used SMART trip threshold from 0% to 2%.
- Added support for parameter 5 in the Verify Error Counter log page (05h).
- Fixed an issue when a drive temperature goes over the threshold it will report the warning or trip immediately instead of delaying it by up to 15 minutes.
- Fixed an issue where Linux mounted Optimus ECO resulted in an aborted command during read/write.
- Added support for the Exit Failure Mode and Overwrite service actions in the Sanitize command.
- Added an event log for changing the SMART sense code.
- Added support for the PBLOCK bit in the Read Long and Write Long commands.
- Disabled SCSI Protection Information so that formatting the drive with PI enabled is illegal. When the drive is formatted with PI disabled the customer will not see the problem that can occur when PI is being stripped during a Read command. If the drive remains formatted with PI enabled after loading this code, its behavior will not be changed except to make formatting with PI enabled illegal.
- Fixed an issue where the Interrupted SANITIZE operations will start again after a power cycle.