



SCS Engineering Release Notice

Phase5 GCA Release Version 05.00.00.00 - SAS3FW_MASTER_DEV (SCGCQ00700566)

(SCGCQ00700566) - Phase5 GCA Release Version 05.00.00.00 - SAS3FW MASTER DEV
(SCGCQ00698224) - Phase5 Beta Release Version 04.250.07.00 - SAS3FW MASTER DEV
(SCGCQ00690399) - Phase5 Beta Release Version 04.250.06.00 - SAS3FW MASTER DEV
(SCGCQ00684452) - Phase5 Beta Release Version 04.250.05.00 - SAS3FW MASTER DEV
(SCGCQ00675507) - Phase5 Alpha Release Version 04.250.04.00 - SAS3FW MASTER DEV
(SCGCQ00665063) - Phase5 Alpha Release Version 04.250.03.00 - SAS3FW MASTER DEV
(SCGCQ00649048) - Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW MASTER DEV

(SCGCQ00638883) - Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW MASTER DEV



SCS Engineering Release Notice

Phase5 GCA Release Version 05.00.00.00 - SAS3FW_MASTER_DEV (SCGCQ00700566)

Defects=0, Enhancements=0 (Version Change Only)



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.07.00 - SAS3FW_MASTER_DEV (SCGCQ00698224)

Change Summary (Defects=1)

SCGCQ00692652 (DFCT) - Fault 6230 While Disabling Expander Phys and Running Stress in Multi-Path Failover Environment



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.07.00 - SAS3FW_MASTER_DEV (SCGCQ00698224)

Total Defects Resolved (1)

(SCGCQ00692652)		Defect 1/1
HEADLINE:	Fault 6230 While Disabling Expander Phys and Running Stress in Multi-Path Failover Environment	
DESC OF CHANGE:	Made a change in Tx transport layer error handling code to call the appropriate SMP engine function.	
TO REPRODUCE:	<p>Topology includes:</p> <ul style="list-style-type: none">2 LSI Gen 3 SAS/SATA controllers - FW 04.250.06.002 host systems running Linux2 LSI Gen 3 expanders - FW 0.4.204.0 <p>Each port of the controllers should make a separate path going into 1 Cobra, followed by 8 enclosures, connected in a round robin fashion, including 203 physical sas drives.</p>	
ISSUE DESC:	<p>Run heavy IOs and periodically glitch (temporarily disable) random phys on the expanders.</p> <p>There's an SMP passthrough request that encounters a no destination count threshold exceeded error, which means the expander that was intended to receive the SMP couldn't be found, probably due to the glitching of expander phys being done in this test. During the cleanup of this error, the non-immediate SMP case was not accounted for, causing the wrong function to be called in the SMP engine code. This resulted in the timer for the SMP not being removed. When the timer bucket expired, controller firmware could not determine where the IO had come from, resulting in a fault.</p>	



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.06.00 - SAS3FW_MASTER_DEV (SCGCQ00690399)

Change Summary (Defects=1 Enhancements=1)

SCGCQ00669775 (DFCT) - Fault 0x6230 While Running Heavy IOs and Cable Breaking

SCGCQ00690282 (CSET) - NVDATA: Adding NVDATA configurations for Active Cable Management



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.06.00 - SAS3FW_MASTER_DEV (SCGCQ00690399)

Total Defects Resolved (1)

(SCGCQ00669775)		Defect 1/1
HEADLINE:	Fault 0x6230 While Running Heavy IOs and Cable Breaking	
DESC OF CHANGE:	Made some changes to the timer code to account for this situation where the SMP cannot be fully cleaned up until later.	
TO REPRODUCE:	LSI Gen 3 SAS/SATA Controller - FW: 04.250.02.00 Modified version of Fedora Linux 1 LSI Gen 3 Cobra R C1 expander (4 SAS drives 4 SATA) -> Quarch cable breaker -> cascaded to 3 enclosures (33 SAS drives, 12 SATA) All drives are running heavy IOs through block move test. While IO's are running, cable breaker is connected to separate Linux host running perl script to break 1 random link every 5 seconds.	
ISSUE DESC:	There was an SMP request that timed out. During hardware cleanup of the SMP request, it was found in a non-current context of Rx Context Manager, which means the cleanup couldn't be completed at that time. This also means the timer for the SMP could not be removed until later. This showed up as an IO in the expired timer bucket that wasn't accounted for, causing the fault.	



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.06.00 - SAS3FW_MASTER_DEV (SCGCQ00690399)

Total Enhancements Implemented (1)

(SCGCQ00690282 - Port of SCGCQ00689826)	Enhancement 1/1
---	-----------------

HEADLINE:	NVDATA: Adding NVDATA configurations for Active Cable Management
NEW FUNCTIONALITY:	<p>Adding new NVDATA files to support active cable management, which includes active copper cables and optical cables. The following boards now support it:</p> <p>sas9300-8e sas9300-16e sas9300-4i4e sas9311-4i4e</p> <p>All files that have active cable management enabled will have the abbreviation "acm" inserted into the file name, i.e. SAS9300-8e_it_acm_p.xml.</p> <p>Active cable management must be actively selected by the customer. There are qualified 12G cables that report 6G functionality under active cable management, and the reported values are used.</p>

Change Summary (Defects=14 NVDATA=1)

- SCGCQ00556766 (DFCT) - IOP: Successive resets may cause Invader/Fury to stop responding on PCIe
- SCGCQ00570820 (DFCT) - When phy power management is turned on, it has been intermittently seen that a HARD_RESET primitive is never sent by the HBA when they send the TARGET RESET function. However, the HBA informs them that the command was sent and the status was good.
- SCGCQ00659829 (DFCT) - Allocation Length for REPORT LUNS SCSI Command is not Handled Properly for SATA Drive
- SCGCQ00663176 (DFCT) - IOP: PCIe DPA Status Configuration Register does not default to zero after a link down
- SCGCQ00663185 (DFCT) - IOP: PCIe Link Down can cause Transmit Margin field to not reset
- SCGCQ00675342 (DFCT) - IOP: MCTP: EIDAllocStatus and EIDAssignStatus in MCTP Control Set EID response swapped
- SCGCQ00675583 (DFCT) - IOP: MCTP: Fault 0x2667 when set PCIe BusMaster, then clear, and send MCTP Request with PCIe VDM
- SCGCQ00679566 (DFCT) - SCSI Command is Not Failed with Expected Response When Reserved Fields Set as Non-Zero Value for SATA Disk Under Format
- SCGCQ00680506 (DFCT) - PL: Firmware is not clearing RxDmaEng host parity interrupt
- SCGCQ00681116 (DFCT) - Support for Power Management States still advertised in PCIe DPA Configuration Register when FW does not support
- SCGCQ00675166 (CSET) - Read log ext command wasn't issued by PL FW to recover from error during the failure of write verify command
- SCGCQ00676251 (CSET) - Change SATL handling for SCSI Format command to issue Write DMA Ext instead of Write FPDMA
- SCGCQ00677682 (CSET) - IOP: I2C Read of non-existent device causes infinite-loop and FW is stalled
- SCGCQ00682737 (CSET) - 12G SAS BC train fails with 6M cable between Fury and Cobra
- SCGCQ00674207 (NVDATA) - Update NVData to add Overtemp Fault Support

Total Defects Resolved (14)**(SCGCQ00556766)** Defect 1/14

HEADLINE: IOP: Successive resets may cause Invader/Fury to stop responding on PCIe

DESC OF CHANGE: Modified a hardware workaround to not disable the reset interrupt during the initialization and bootloader portions of firmware and then periodically check to see if a reset request was made. If such a reset was requested, firmware will then safely reset the controller.

TO REPRODUCE: Boot the controller and then reset it. Shortly after resetting, issue another reset. Subsequently, attempt to read a register on the controller, and the request may not be responded to.

ISSUE DESC: If the controller is successively reset fast enough, the controller may stop responding to requests on the PCI-Express bus.

(SCGCQ00570820) Defect 2/14

HEADLINE: When phy power management is turned on, it has been intermittently seen that a HARD_RESET primitive is never sent by the HBA when they send the TARGET RESET function. However, the HBA informs them that the command was sent and the status was good.

DESC OF CHANGE: There are two interrupts that the hardware can send to the firmware when a power condition change has occurred, Power Management Initiated and Power Management Wake. There can be three types of transition.

Active --> Partial --> Active
Partial --> Slumber --> Active
Slumber --> Active --> Partial

TO REPRODUCE: When FW gets notified of PM Init and PM Wake, it has to determine what the current power condition is. In the first two cases we are transitioning to Active power condition, so we need to send a Hard Link Reset. So now a drive initiated power condition change will result in a Hard Link Reset being sent out.

When WD sends a TARGET RESET task management function across the write we see a HARD_RESET primitive being sent to the drive causing it to reset. When they turn on phy power management, intermittently when they send the TARGET RESET function the HARD_RESET primitive is never sent by the HBA however the HBA informs them that the command was sent and the status was good.

Here's a little more detail on the Target reset issue. The problem is mishandling of TMF TARGET_RESET when the PHY is in low power state.

The attached bus trace shows 2 iterations of the sequence below; one successful and one failure :

1. Enter slumber (HIPM initiated)
2. Send TARGET_RESET TMF
3. Send TUR expecting UA status

In the successful case, the HBA initiates wakeup from slumber with COMINIT and in the process sends HARD_RESET primitive. The drive responds as expected to subsequent TUR command with Unit attention status.

In the failing case, the HBA again initiates wakeup from slumber with COMINIT, but no HARD_RESET primitive is sent. The drive responds to the subsequent TUR command with GOOD status.

ISSUE DESC: When they turn on phy power management, intermittently when they send the TARGET RESET function the HARD_RESET primitive is never sent by the HBA however the HBA informs them that the command was sent and the status was good.

(SCGCQ00659829) Defect 3/14

HEADLINE: Allocation Length for REPORT LUNS SCSI Command is not Handled Properly for SATA Drive

DESC OF CHANGE: Firmware will return enough data for 1 LUN for a Report LUNS command even if allocation length is set to a much larger number.

TO REPRODUCE: Send Report LUNS scsi command to a SATA drive.

ISSUE DESC: When allocation length of Report LUNS command is set to 0xFF, firmware would return 255 byte of data even though we only support 1 LUN.

(SCGCQ00663176) Defect 4/14

HEADLINE: IOP: PCIe DPA Status Configuration Register does not default to zero after a link down

DESC OF CHANGE: Added code to clear the register when a Link Down interrupt is received.

TO REPRODUCE: Write a 1 to Substate Status field of PEPP register 0x19A8
Verify that DPA State of DPA Status config register contains a 1
Perform a link disable/enable
Verify that DPA State of DPA Status config register erroneously contains a 1 instead of the default value of 0

ISSUE DESC: Per the PCI-Express specification Revision 3.0, the default for the Substate Status field of the DPA Status should default to 0x00 after a conventional reset. During a link disable (a.k.a. link down), the registers are not reset and the configuration register does not get reset to its default value.

(SCGCQ00663185) Defect 5/14

HEADLINE: IOP: PCIe Link Down can cause Transmit Margin field to not reset
DESC OF CHANGE: Added code to clear the register value when a link down event occurs.
TO REPRODUCE: Write a 1 to the Transmit Margin Field of the PCIe config Link Status/Control 2 register
Perform a link disable/enable
Verify that the transmit margin field erroneously contains a 1 instead of a default value of 0
ISSUE DESC: Per the PCI-Express specification, the Transmit Margin field must reset to 0 when a link down event occurs. The hardware is not clearing this automatically, so PCI-Express validation tests are failing.

(SCGCQ00675342) Defect 6/14

HEADLINE: IOP: MCTP: EIDAllocStatus and EIDAssignStatus in MCTP Control Set EID response swapped
DESC OF CHANGE: Modified the code that defined the EIDAllocStatus and EIDAssignStatus fields so that they are in the correct location.
TO REPRODUCE: Send an MCTP Control Set EID request.
ISSUE DESC: The EIDAllocStatus and EIDAssignStatus in MCTP Control Set EID response were swapped.

(SCGCQ00675583) Defect 7/14

HEADLINE: IOP: MCTP: Fault 0x2667 when set PCIe BusMaster, then clear, and send MCTP Request with PCIe VDM
DESC OF CHANGE: Several modifications were made. Firmware cannot process any MCTP over PCIe VDM TLPs until the BusMaster bit is first set. If the BusMaster bit is cleared, then firmware cannot send any responses as PCIe messages.
TO REPRODUCE: Set the Bus Master bit in the PCIe Config Space.
Clear the Bus Master bit.
Send a MCTP Request over PCIe VDM, then a 0x2667 fault occurs.
ISSUE DESC: If the Bus Master bit in the PCIe Config Space is enabled, then disabled, and a MCTP Request is sent in over PCIe VDM, then a 0x2667 fault occurs. No PCIe messages can be sent out when the Bus Master bit is disabled.

(SCGCQ00679566) Defect 8/14

HEADLINE: SCSI Command is Not Failed with Expected Response When Reserved Fields Set as Non-Zero Value for SATA Disk Under Format
DESC OF CHANGE: Added an additional check in SATL code to look for this case and fail the command with check condition, illegal request, invalid field in CDB.
TO REPRODUCE: Send a SCSI Inquiry command to a SATA drive currently undergoing a format operation, and in the SCSI CDB, set the reserved and/or control fields to a non-zero value. The command will complete successfully, when it should be failed.
ISSUE DESC: When a SCSI Inquiry command was sent to a SATA drive currently undergoing a format operation, and specifically when the command had non-zero values for the reserved and/or control fields, the command was not being failed as was intended.

(SCGCQ00680506) Defect 9/14

HEADLINE: PL: Firmware is not clearing RxDmaEng host parity interrupt
DESC OF CHANGE: Clear the RxDmaEng host parity interrupt after firmware handles it.
TO REPRODUCE: None. This is found by code inspection.
ISSUE DESC: Firmware does not clear the RxDmaEng host parity interrupt after handling it. This could cause the firmware to keep handling the interrupt over and over.

(SCGCQ00681116) Defect 10/14

HEADLINE: Support for Power Management States still advertised in PCIe DPA Configuration Register when FW does



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.05.00 - SAS3FW_MASTER_DEV (SCGCQ00684452)

not support

DESC OF CHANGE: When Power Management is disabled, the DPA config register now reports only supporting one DPA State.

TO REPRODUCE:

1. Disable Power Management
2. During the PCI-SIG PCIECV compliance test, the test system sees in the DPA config register that the HBA supports three DPA states.

ISSUE DESC: When power management is disabled, the DPA config register is reporting that the HBA supports three DPA States. It should report only supporting the Full Power State.

(SCGCQ00675166 - Port of SCGCQ00623315)

Defect 11/14

HEADLINE: Read log ext command wasn't issued by PL FW to recover from error during the failure of write verify command

DESC OF CHANGE: PL FW code has been modified to issue Write Ext command instead of Write FPDMA (NCQ mode) for the write phase of write verify command.

TO REPRODUCE:

1. Issue write verify command to an invalid volume
2. FW receive abort error

ISSUE DESC: Write Verify command is issuing Write FPDMA (NCQ mode) for the write phase and the Write FPDMA is failed, then PL needs to issued the read log ext to recover from the error, but FW did not send Read Log Ext command to recover from the error

(SCGCQ00676251 - Port of SCGCQ00660847)

Defect 12/14

HEADLINE: Change SATL handling for SCSI Format command to issue Write DMA Ext instead of Write FPDMA

DESC OF CHANGE: Modified SATL code to issue Write DMA Ext instead of Write FPDMA for SCSI Format command to avoid NCQ error handling to recover.

TO REPRODUCE: Issue SCSI Format command to SATA drive.

ISSUE DESC: As part of SCSI format command to SATA drive, firmware issues Write FPDMA (NCQ mode) for the write phase and if the Write FPDMA is failed, then firmware needs to issue Read Log Ext command to recover from the error, but currently in firmware there is no Read Log Ext command issued for the SCSI format command.

(SCGCQ00677682 - Port of SCGCQ00673161)

Defect 13/14

HEADLINE: IOP: I2C Read of non-existent device causes infinite-loop and FW is stalled

DESC OF CHANGE: Fix the I2C Read code that attempted retries with non-existent devices, but reset internal timeouts. The lower level I2C Read code no longer attempts any retries with non-existent devices. This is left up to the higher layers that called the I2C Read code.

TO REPRODUCE: Send an ISTWI Toolbox MPI Request to a controller. The request must be for a non-existent device.

ISSUE DESC: An I2C Read of a non-existent device on an I2C bus causes an infinite-loop and stalls FW.

(SCGCQ00682737 - Port of SCGCQ00678053)

Defect 14/14

HEADLINE: 12G SAS BC train fails with 6M cable between Fury and Cobra

DESC OF CHANGE: Firmware will program the preload value based on nvdata. On link down interrupt, firmware will also load the same preload values used at start-of-day.

TO REPRODUCE: Connect a Fury HBA to Cobra using a 6-meter cable.

ISSUE DESC: Some phys on some HBA may not link up with Cobra at 12G with a 6-meter cable.



SCS Engineering Release Notice

Phase5 Beta Release Version 04.250.05.00 - SAS3FW_MASTER_DEV (SCGCQ00684452)

Total NVDATA Changes (1)

(SCGCQ00674207)

NVDATA 1/1

HEADLINE: Update NVData to add Overtemp Fault Support

NEW FUNCTIONALITY: Update NVData to add Overtemp Fault Support



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.04.00 - SAS3FW_MASTER_DEV (SCGCQ00675507)

Change Summary (Defects=13)

SCGCQ00661386 (DFCT) - SR-IOV: Devices invisible to VF are listed in SAS IO Unit Page 0

SCGCQ00661564 (DFCT) - IOP: Incorrect PCIe SERDES settings

SCGCQ00664517 (DFCT) - IOP SR-IOV: Fault 0x265D seen when issuing Function Level Resets after a firmware fault

SCGCQ00666811 (DFCT) - PL: Low read/write throughput for SATA when interspersed with non-automated commands.

SCGCQ00668564 (DFCT) - IOP SR-IOV: Firmware fails to clear visibility settings on removed devices

SCGCQ00668686 (DFCT) - IOP MCTP: Uninitialized variables need correcting

SCGCQ00668692 (DFCT) - PL: Add fault code to indicate coding error in cable management code

SCGCQ00669712 (DFCT) - PL: Physical port structure is allocated too large

SCGCQ00670932 (DFCT) - IOP SR-IOV: Reply Descriptor Post Queue Array from VFs during heavy load causes 0x0A05 fault

SCGCQ00672407 (DFCT) - PL: SCSI Read Command of Zero Transfer Length to SATA Returns Data

SCGCQ00673173 (DFCT) - IOP: MCTP: Packets with valid MCTP SeqCount values are dropped

SCGCQ00672212 (CSET) - Unexpected IOCStatus in Error Reply when NAK is received by Fury in target mode

SCGCQ00672646 (CSET) - IOP: PowerPC ITLB Miss Errata with Pre-fetch using start of TLB padding



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.04.00 - SAS3FW_MASTER_DEV (SCGCQ00675507)

Total Defects Resolved (13)

(SCGCQ00661386) Defect 1/13

HEADLINE: SR-IOV: Devices invisible to VF are listed in SAS IO Unit Page 0
DESC OF CHANGE: In SR-IOV firmware (only), added code to not populate the AttachedDevHandle field if the attached device is not visible and it is an end device. For non-SR-IOV firmware, no change is made.
TO REPRODUCE: Use SR-IOV firmware on a host with some directly attached drives. Do not give any virtual function permission to access the drives. Perform a configuration page read of SAS IO Unit Page 0 using the virtual function and observe device handles are reported in the AttachedDevHandle field.
ISSUE DESC: When using SR-IOV firmware, Virtual Functions reading SAS IO Unit Page 0 are able to see devices which they do not have permission to see. Similar problems existing in SAS Phy Page 0 and SAS Phy Page 4.

(SCGCQ00661564) Defect 2/13

HEADLINE: IOP: Incorrect PCIe SERDES settings
DESC OF CHANGE: Changed firmware to load the correct PCI-Express SERDES settings.
TO REPRODUCE: Identified by code inspection.
ISSUE DESC: An incorrect value was being loaded in the PCI-Express SERDES registers.

(SCGCQ00664517) Defect 3/13

HEADLINE: IOP SR-IOV: Fault 0x265D seen when issuing Function Level Resets after a firmware fault
DESC OF CHANGE: Modified firmware to ignore function level resets until after a large portion of initialization. This prevents firmware from accessing uninitialized pointers.
TO REPRODUCE: Use SR-IOV firmware and configure several virtual machines to use several virtual functions. Cause firmware to fault. If a virtual function issues its reset after the physical function, firmware will fault with code 0x265D.
ISSUE DESC: When using SR-IOV firmware, if firmware faults, the host drivers will issue resets. If a Function Level Reset is received after the diagnostic reset from the Physical Function, firmware will fault with code 0x265D.

(SCGCQ00666811) Defect 4/13

HEADLINE: PL: Low read/write throughput for SATA when interspersed with non-automated commands.
DESC OF CHANGE: Fixed the Fast Path Engine state machine so that the non-read/write command(s) pending in the firmware are started immediately after zero IO count notification is received from the Fast Path hardware.
TO REPRODUCE: Attach a SATA drive to the controller. Issue and queue up write commands to the drive so that peak throughput is achieved. Next, intersperse the write commands with a non-read/write command such as the SCSI Inquiry at regular intervals (say every 5 seconds). Monitor the write throughput. Observe a steep drop in the throughput (~70%-80% drop) around the time the non-read/write command is issued.
ISSUE DESC: When read/write commands (hardware automated) interspersed with non-read/write commands (firmware handled) are issued to a SATA drive, then the read/write throughput drops drastically from the peak throughput.

(SCGCQ00668564) Defect 5/13

HEADLINE: IOP SR-IOV: Firmware fails to clear visibility settings on removed devices
DESC OF CHANGE: Corrected a bug where a variable which should have been constant was being changed, causing an if() statement to fail when processing more entries in the SAS Topology Change List Event.
TO REPRODUCE: Use SR-IOV firmware and attach several expander attached devices on sequential expander phys. Set visibility to at least one virtual function. Remove the expander from the topology. Wait for the device missing delay timer to expire and the devices get removed. Reconnect the expander to the controller. When the driver configures visibility for the same virtual function, a device add event is not generated to the virtual function.
ISSUE DESC: When running SR-IOV firmware, when devices are removed from the topology, firmware may not clear the visibility flags from the devices. This may cause issues when the devices are re-inserted whereby add events may not be sent to the virtual functions.

(SCGCQ00668686) Defect 6/13

HEADLINE: IOP MCTP: Uninitialized variables need correcting

DESC OF CHANGE: Both variables were either properly set before used or a different variable was checked to generate the proper action.

TO REPRODUCE: One issue would be triggered by sending an ABORT command that failed to find the AppMsgTag to abort. The other would be triggered by the I2C binding in slave-response mode when a response to start transmission with a Flow Control Resume could not be found.

ISSUE DESC: There were two instances of uninitialized variables which were not checked in MCTP code.

(SCGCQ00668692) Defect 7/13

HEADLINE: PL: Add fault code to indicate coding error in cable management code

DESC OF CHANGE: Firmware would fault if it ever runs into this condition.

TO REPRODUCE: Found by code analysis tool.

ISSUE DESC: There's a variable in cable management code that could potentially get out of bound.

(SCGCQ00669712) Defect 8/13

HEADLINE: PL: Physical port structure is allocated too large

DESC OF CHANGE: Removed extra allocated entry. Also modified code which loops expecting this extra entry.

TO REPRODUCE: Found by code analysis.

ISSUE DESC: An internal data structure is allocated with one entry per port, but an extra entry is allocated for no reason.

(SCGCQ00670932) Defect 9/13

HEADLINE: IOP SR-IOV: Reply Descriptor Post Queue Array from VFs during heavy load causes 0x0A05 fault

DESC OF CHANGE: Added code to request a critical frame which should always be available for use for the RDPQ Array. This should eliminate the possibility of not getting memory to copy the RDPQ Array locally during initialization.

TO REPRODUCE: Create a system with multiple virtual functions used and start IO on all but one function. On the remaining virtual function, issue an IOC Init request using the Reply Descriptor Post Queue Array feature. The firmware will fault with code 0x0A05 if there are not enough resources to move the array local to the controller.

ISSUE DESC: In a SR-IOV configuration with many virtual functions in use and sending heavy IO, a virtual function issuing a MPI IOC Init message may cause firmware to fault with code 0x0A05 if the IOC Init message uses the Reply Descriptor Post Queue Array feature of MPI.

(SCGCQ00672407) Defect 10/13

HEADLINE: PL: SCSI Read Command of Zero Transfer Length to SATA Returns Data

DESC OF CHANGE: Made a change in the Rx frame manager interrupt handler so that the IO will be failed with SCSI_status = success and IOC_status = underrun, with the transfer count field set to 0.

TO REPRODUCE: Using an LSI gen 3 controller, send a SCSI read command to a SATA drive, with transfer length set to 0 but data length in the SCSI IO request set to non-zero. The return status will be success with data transfer.

ISSUE DESC: In the case where a SCSI read command is sent to a SATA drive (SATL read), and the transfer length is set to 0, but the data length in the SCSI IO request is not set to 0, the controller returned unexpected data. The return status showed data transfer, when in fact there wasn't data transfer and should not have been any data transfer.

(SCGCQ00673173) Defect 11/13

HEADLINE: IOP: MCTP: Packets with valid MCTP SeqCount values are dropped

DESC OF CHANGE: Removed code requiring when the Start-Of-Message equals 1 to have the SeqCount field be 0.

TO REPRODUCE: Send an MCTP Control command with the SeqCount as a non-zero value.

ISSUE DESC: MCTP packets that contain a non-zero MCTP SeqCount when the Start-Of-Message bit is set are dropped. These are valid packets.

(SCGCQ00672212 - Port of SCGCQ00662771) Defect 12/13

HEADLINE: Unexpected IOCStatus in Error Reply when NAK is received by Fury in target mode



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.04.00 - SAS3FW_MASTER_DEV (SCGCQ00675507)

DESC OF CHANGE: When NAK is received by Fury in target mode, IOC should fail the I/O with IOCStatus:0x71 (NAK_RECEIVED) instead of IOCStatus:0x65 (NO_CONNECTION).

TO REPRODUCE: Topology:
SAS Initiator --- SAS Jammer --- SAS Target(SAS3008)

Procedure:
1) SAS initiator sends READ(10) command (512bytes single data) to SAS3008.
2) SAS3008 receives this command properly.
3) SAS3008 target host driver issues MPI Target Assist request to transfer the data to the initiator.
4) SAS jammer injects CRC error on this data frame.
5) SAS initiator receives the data frame and detects CRC error on that frame.
6) SAS initiator sends NAK to SAS3008 due to CRC error.
7) SAS3008 sends an unexpected error reply to MPI target assist request.

ISSUE DESC: When NAK is received by Fury in target mode, IOC fails I/Os with IOCStatus:0x65 (NO_CONNECTION).

(SCGCQ00672646 - Port of SCGCQ00672307)

Defect 13/13

HEADLINE: IOP: PowerPC ITLB Miss Errata with Pre-fetch using start of TLB padding

DESC OF CHANGE: Added padding to the beginning instead of end of each block of instructions that correspond to an instruction TLB; no instructions may reside in this padding. This prevents the prefetcher from fetching executable instructions from the virtually adjacent instruction TLB.

TO REPRODUCE: This issue has not been reproduced using IT/IR firmware.

ISSUE DESC: One to four extra random instructions might be executed erroneously after an instruction-side translation lookaside buffer (ITLB) miss operation occurs due to SCGCQ00634879, but through the instruction pre-fetch.



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.03.00 - SAS3FW_MASTER_DEV (SCGCQ00665063)

Change Summary (Defects=12 Enhancements=1 NVDATA=1)

SCGCQ00582515 (DFCT) - IOC Fault 0x1609 by Message Unit Reset

SCGCQ00643835 (DFCT) - Fixed Simple Coding Issue Leading to 0x6230 Fault

SCGCQ00649572 (DFCT) - Inquiry SCSI Command is Not Failed with Expected Response When NACA bit in Control Byte is set during During Format

SCGCQ00654721 (DFCT) - Unsupported SCSI Command fails with Incorrect Sense Data When SATA Drive is in Standby Power Condition

SCGCQ00655251 (DFCT) - IOP SR-IOV: Controller reports 16 MSIx vectors in the MSIx Table Size field of the Config Space

SCGCQ00655407 (DFCT) - IOP SR-IOV: Devices in SAS Topology Change List Event are incorrectly ordered

SCGCQ00659190 (DFCT) - PL: Invalid config page setting causes uneven link up in SAS vs SATA

SCGCQ00660285 (DFCT) - IOP: One path to freeing received generic PCIe messages is optimized out

SCGCQ00663100 (DFCT) - IOP: WA for Inbound Message Queue enable bit surviving resets in Trident with reset Inbound Message Queue regs

SCGCQ00663155 (DFCT) - IOP: MCTP: MCTP Control Discovery Notify over PCIe may be transmitted before host enables PCIe messages

SCGCQ00664321 (DFCT) - IOP: SR-IOV Virtual Functions are able to access BIOS configuration pages

SCGCQ00651231 (CSET) - SAS Phy page 1 and 2 error counters don't perform properly

SCGCQ00645681 (ENHREQ) - WW: SGPIO LED Control for NVME devices on the PCIe Switch

SCGCQ00612736 (NVDATA) - Initial NVDATA review for SAS9300-16i

Total Defects Resolved (12)

(SCGCQ00582515)		Defect 1/12
HEADLINE:	IOC Fault 0x1609 by Message Unit Reset	
DESC OF CHANGE:	The algorithm used to reset the queues during MUR tries to free up the MIDs starting from the producer until it reaches the consumer of the queue. The producer pointer is actually pointing to the entry that is not yet been used in the queue (and contains a value of 0xEEEE). When we try to free up this entry, we hit the fault 0x1609.	
TO REPRODUCE:	To fix this issue, the algorithm to free up the entries is modified to start from the consumer pointer and end when the producer pointer is reached.	
	<ul style="list-style-type: none">- Power-on the system.- IOC is in RESET state due to no dedicated Flash, perform the host boot. IOC goes to READY state.- Start-of-day (IOCFacts, IOCInit). IOC goes to OPERATIONAL state.- Post command buffers by using MPI CommandBufferPostBase and CommandBufferPostList. Verify SUCCESS reply is returned from FW.- Read the manufacturing page 17 by using MPI Configuration Request (PAGE_HEADER).- Before the target host driver receives SUCCESS reply from FW, it issues MUR.- IOC goes to FAULT state with 0x1609 error code.	
ISSUE DESC:	Immediately after CommandBufferPost(Target Mode), the host driver issues MUR (Message Unit Reset). As a result, IOC goes to FAULT state with 0x1609 error code.	
(SCGCQ00643835)		Defect 2/12
HEADLINE:	Fixed Simple Coding Issue Leading to 0x6230 Fault	
DESC OF CHANGE:	There's a function that looks to see if any of the IOs that have timed out are SATA init IOs. The change was to remove a break statement in that function that shouldn't have been there. It was causing the function to find at most one IO, when in some cases there was more than one.	
TO REPRODUCE:	Topology as follows: LSI Gen 3 Controller -> cable breaker 1 -> 3 cascaded enclosures -> cable breaker 2 -> LSI Cobra R Expander There are 52 drives total with 4 sata on each expander.	
ISSUE DESC:	Test ran step 1, glitching on cable breaker 1 for 5 ms every 5 seconds.	
	Test failed in step 2, which was breaking every 1 to 5 seconds with 5 sec between breaks.	
	0x6230 fault was added recently, and it means that there's an IO which timed out that we cannot account for (meaning we did not find it among the SATA init IOs, or SMPs, etc).	
(SCGCQ00649572)		Defect 3/12
HEADLINE:	Inquiry SCSI Command is Not Failed with Expected Response When NACA bit in Control Byte is set during During Format	
DESC OF CHANGE:	When NACA bit is set, firmware should fail the command with with CHECK CONDITION status with the sense key set to ILLEGAL REQUEST and the additional sense code set to INVALID FIELD IN CDB.	
TO REPRODUCE:	Send an inquiry command with NACA bit set in control byte during format.	
ISSUE DESC:	Inquiry command during format is not failed with proper response when NACA bit is set.	
(SCGCQ00654721)		Defect 4/12
HEADLINE:	Unsupported SCSI Command fails with Incorrect Sense Data When SATA Drive is in Standby Power Condition	
DESC OF CHANGE:	Changed the returned sense data for this case to follow the spec, which is a sense key of Illegal Request and ASC value of Invalid Command Operation Code.	
TO REPRODUCE:	Connect a SATA drive to an LSI gen 3 SAS/SATA controller. Initialize the SATA drive. Then move the drive to standby state. Now send an illegal SCSI command to the SATA drive (one that the SCSI to ATA translation layer does not support in ATA mode). The sense data returned will be Not Ready, Logical unit not ready, initializing command required.	
ISSUE DESC:	When a SATA drive is in standby mode (not spun up), and an illegal SCSI command (one that SATL does not translate to ATA) is received, the SATL was replying with incorrect sense data.	
(SCGCQ00655251)		Defect 5/12
HEADLINE:	IOP SR-IOV: Controller reports 16 MSIx vectors in the MSIx Table Size field of the Config Space	



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.03.00 - SAS3FW_MASTER_DEV (SCGCQ00665063)

DESC OF CHANGE:	Enabled a hardware work around which ensures this value is set correctly.
TO REPRODUCE:	Load SR-IOV firmware on a card and boot the system. Read the configuration space and observe the value is incorrect.
ISSUE DESC:	SR-IOV firmware reports support for 16 MSIx vectors for all PFs and VFs in the MSIx Table Size field of the PCI configuration space. The controller only supports 8. This only impacts SR-IOV enabled firmware.

(SCGCQ00655407)		Defect 6/12
HEADLINE:	IOP SR-IOV: Devices in SAS Topology Change List Event are incorrectly ordered	
DESC OF CHANGE:	Changed the code which filters the SAS Topology Change List Event to mark phys as having No Change instead of deleting the entry from the list when the attached device is not visible to a VF.	
TO REPRODUCE:	Using SR-IOV firmware, configure visibility for a VF for expander attached drives such that for phys 0-16, devices attached to lower numbered phys are not visible and devices attached to higher numbered phys are visible. Pull a drive from the topology and re-insert it before the Device Missing Delay timer expires. The SAS Topology Change List Event will list the re-inserted drive in the wrong location.	
ISSUE DESC:	In SR-IOV firmware, configuring disks in such a way may cause SAS Topology Change List Events to report devices in a non-phy sequential order as required by MPI. This results in the driver thinking a drive is attached to a different phy than it actually is.	

(SCGCQ00659190)		Defect 7/12
HEADLINE:	PL: Invalid config page setting causes uneven link up in SAS vs SATA	
DESC OF CHANGE:	Added code to error check, and in the case of persistent pages, error correct these settings. Writing incorrect values for this field to the current page will result in an error. If a user has made an incorrect setting in the persistent pages, the values will be rounded up or down to the correct range.	
TO REPRODUCE:	Perform a configuration page write to SAS IO Unit Page 1 and set the maximum/minimum link rate to 1.5Gbps on all controller phys. Write this page to either the persistent configuration or the current configuration. If writing to the persistent pages, reset the controller and observe SAS drives will not link up to the controller and SATA drives will link up at 3Gbps.	
ISSUE DESC:	If the user sets a maximum and minimum link rate of 1.5Gbps, which is not supported, SAS drives will not link up but SATA drives will link up at 3Gbps. The mismatch in behavior is not expected.	

(SCGCQ00660285)		Defect 8/12
HEADLINE:	IOP: One path to freeing received generic PCIe messages is optimized out	
DESC OF CHANGE:	Modified the PCIe credit freeing code to not be optimized out of the firmware.	
TO REPRODUCE:	Send a PCIe message to a controller with MCTP enabled. If this PCIe message is not a VDM message, the PCIe credit is not properly freed.	
ISSUE DESC:	A particular code path that frees a PCIe credit for generic PCIe messages is optimized out. This causes PCIe credits to not be properly freed.	

(SCGCQ00663100)		Defect 9/12
HEADLINE:	IOP: WA for Inbound Message Queue enable bit surviving resets in Trident with reset Inbound Message Queue regs	
DESC OF CHANGE:	Modified FW to clear the sticky enable bit for the Inbound Message Queue when FW receives a PCIe Hot Reset, an Adapter Reset, and the bootloader runs.	
TO REPRODUCE:	Setup FW to receive PCIe messages (e.g. enable MCTP). Reset the controller by either PCIe Hot Reset or an Adapter Reset. Send a PCIe message immediately. A 0x2622 fault will occur in FW and the PCIe Config Space will show a CompleterAbort error.	
ISSUE DESC:	The bit to enable the Inbound Message Queue for FW to receive generic PCIe messages is sticky and survives soft resets. The registers that contain the addresses to save the data do not survive soft resets. This causes an internal error, which sets the CompleterAbort error in PCIe Config Space, and will cause the FW to fault with code 0x2622.	

(SCGCQ00663155)		Defect 10/12
HEADLINE:	IOP: MCTP: MCTP Control Discovery Notify over PCIe may be transmitted before host enables PCIe messages	
DESC OF CHANGE:	FW now only sends MCTP Control Discovery Notify messages over PCIe after the host sets the BusMaster bit.	



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.03.00 - SAS3FW_MASTER_DEV (SCGCQ00665063)

TO REPRODUCE:	Enable MCTP over PCIe VDM on the controller. Boot up the server and do not set the Bus Master bit in the Status/Command register in the PCIe Config Space. A PCIe message from the controller can be seen routed to the root complex. Further results depends upon system BIOS behavior.
ISSUE DESC:	When the controller first boots with MCTP over PCIe VDM enabled, the controller needs to transmit an MCTP Control Discovery Notify message as a PCIe VDM type 1 message. Some servers allow the PCIe attached cards to boot much earlier than the system BIOS is ready to receive any PCIe messages and the controller assumes this is already enabled, i.e. before the Bus Master bit in the Command/Status register in the PCIe Config Space.

(SCGCQ00664321) Defect 11/12

HEADLINE:	IOP: SR-IOV Virtual Functions are able to access BIOS configuration pages
DESC OF CHANGE:	Removed the VF_READ permission from the BIOS configuration pages in the firmware. The pages are still accessible from the physical function.
TO REPRODUCE:	Using LSIUtil on a Virtual Function, attempt to read a BIOS configuration page. The request will be honored.
ISSUE DESC:	From a Virtual Function, the host is able to read BIOS configuration pages. This is not allowed per the requirements.

(SCGCQ00651231 - Port of SCGCQ00588750) Defect 12/12

HEADLINE:	SAS Phy page 1 and 2 error counters don't perform properly
DESC OF CHANGE:	Instead of using the firmware counters, read from the hardware counters directly.
TO REPRODUCE:	Inject phy errors. Read sas phy page 1 or 2.
ISSUE DESC:	The error counters will always wrap around when maximum values are reached.



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.03.00 - SAS3FW_MASTER_DEV (SCGCQ00665063)

Total Enhancements Implemented (1)

(SCGCQ00645681)	Enhancement 1/1
-----------------	-----------------

HEADLINE: WW: SGPIO LED Control for NVME devices on the PCIe Switch

NEW FUNCTIONALITY: With compatible hardware, PL code can now use SGPIO to control LEDs for NVME devices on a PCIe switch.



SCS Engineering Release Notice

Phase5 Alpha Release Version 04.250.03.00 - SAS3FW_MASTER_DEV (SCGCQ00665063)

Total NVDATA Changes (1)

(SCGCQ00612736)

NVDATA 1/1

HEADLINE: Initial NVDATA review for SAS9300-16i

NEW FUNCTIONALITY: Added IT NVDATA for the SAS9300-16i board.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

Change Summary (Defects=13 Enhancements=8)

SCGCQ00628087 (DFCT) - WW: Capella-2 FPGA related fixes and cleanup

SCGCQ00628654 (DFCT) - PL fault 0xD073 while running IOs

SCGCQ00632252 (DFCT) - PL FW faults with fault code 0xec11

SCGCQ00637107 (DFCT) - WW: A few NVME completion Q prints were flooding ring buffer, removed them

SCGCQ00637787 (DFCT) - WW: Fault 0x6001 when running I/O and TaskManagment

SCGCQ00640245 (DFCT) - Whirlwind: Set and clear FPE flags correctly for NVMe drives for TM

housekeeping purposes

SCGCQ00641946 (DFCT) - IOP: Recursive fault may not report fault state to host

SCGCQ00645292 (DFCT) - IOP: MCTP messages: Test size of MPI Event Notification Request in MCTP

SCGCQ00645300 (DFCT) - IOP: MCTP messages: RAID SCSI IO Passthrough MPI requests timed out in

IR

SCGCQ00645484 (DFCT) - WW: (Capella-2) NVMe initialization times out.

SCGCQ00637311 (CSET) - NVMe: handling of admin queue full condition during NVMe async event

handling

SCGCQ00643315 (CSET) - when Phase 18 IR FW boot a RAID volume (such as raid 10), 2308 bootup hangs at OROM "Initializing.." and fault 8803 has seen

SCGCQ00648184 (CSET) - Need a method to disable PL Light print in IO path

SCGCQ00577387 (ENHREQ) - Temperature Monitoring with bi-colored LED and over temp faulting.

SCGCQ00584505 (ENHREQ) - RAID Flags and CDB fields are corrupted after Fast Path Engine rejects

FastPath ATA NATIVE COMMAND

SCGCQ00601742 (ENHREQ) - WW: Add support for DirectIO Q pair

SCGCQ00621996 (ENHREQ) - WW: account for data length in SGL during translation to NVMe PRP

SCGCQ00626865 (ENHREQ) - IOP MCTP: Add support for VDM Buffers

SCGCQ00644379 (ENHREQ) - IOP: Add compile option to start the second CPU

SCGCQ00645401 (ENHREQ) - PL: Poll for parity error in FPE PLB Slave

SCGCQ00638811 (CSET) - Additional diag console commands



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

Total Defects Resolved (13)

(SCGCQ00628087)		Defect 1/13
HEADLINE:	WW: Capella-2 FPGA related fixes and cleanup	
DESC OF CHANGE:	Fixed start-of-day configuration for Capella-2 switch FPGA. Fixed the CSR handling code to correctly reply to Configuration Write requests from the host. Updated the host port setup code to allow Completion TLPs from Invader root complex to route back to the host.	
TO REPRODUCE:	Attach a host to Capella-2 switch FPGA. Power on the Capella-2 FPGA and activate the bitstream. Issue CLI command to perform enumeration and setup CSR. Power on the host. The host will not boot.	
ISSUE DESC:	WW does not allow host attached to Capella-2 FPGA to boot.	
(SCGCQ00628654)		Defect 2/13
HEADLINE:	PL fault 0xD073 while running IOs	
DESC OF CHANGE:	Modified the code to check if the bucket count is zero before decrementing it manually. The manual decrements will be done only when the timer bucket count is non zero.	
TO REPRODUCE:	1. Configure 3 VD's, R0, R1 and R5. And keep 3 JBODs. 2. From the Client, mount all the VD's(mix of VMFS5 and VMFS3), so that it reflects in the /vmfs/volumes/ 3. Now, copy chaos directly in the mounted VD's. And run chaos IOs with 20 thread on each VD and JBOD. Remember to mix IO size with 32k,32k and 32k,4096k in some on the VD's. 4. As IOs start to run, trigger chip restart from StorCLI. 5. Wait for sometime and again trigger a restart. Wait for half an hour or so.	
ISSUE DESC:	PL Fault 0xD073 was seen while running IOs	
(SCGCQ00632252)		Defect 3/13
HEADLINE:	PL FW faults with fault code 0xec11	
DESC OF CHANGE:	The SCSI READ/WRITE breakup IO completion handling was added. Also, when issuing a breakup IO, the LBA, FUA, transfer length values were properly pulled from the IO message frame.	
TO REPRODUCE:	1. Connect Invader RC and NVME emulator to the PLX switch 2. Create a R0 logical volume 3. Run full init operation	
ISSUE DESC:	The FW faults with code ec11	
(SCGCQ00637107)		Defect 4/13
HEADLINE:	WW: A few NVME completion Q prints were flooding ring buffer, removed them	
DESC OF CHANGE:	Removed unneeded prints	
TO REPRODUCE:	N/A	
ISSUE DESC:	WW: A few NVME completion Q prints were flooding ring buffer, removed them	
(SCGCQ00637787)		Defect 5/13
HEADLINE:	WW: Fault 0x6001 when running I/O and TaskManagment	
DESC OF CHANGE:	Internally generated Flush command and Abort command as part of a Abort Task sent down by host to an NVMe drive doesnt impact IoCount any longer. This is done to keep it in sync with the same principle followed by PL when it sends a Task IU in SAS/SSP.	
TO REPRODUCE:	Inject random Task Managements while running IOs to a bunch of NVMe drives attached to the switch	
ISSUE DESC:	IoCount wasnt getting accounted for during a FLUSH command cleanup when Abort Task to an NVMe drive timed out and or got picked up by a Target Reset (higher scope) resulting in IoCount imbalance and eventual 6001.	
(SCGCQ00640245)		Defect 6/13
HEADLINE:	Whirlwind: Set and clear FPE flags correctly for NVMe drives for TM housekeeping purposes	
DESC OF CHANGE:	Set and clear FPE flags correctly for NVMe drives for TM housekeeping purposes	
TO REPRODUCE:	n/a	



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

ISSUE DESC: Set and clear FPE flags correctly for NVMe drives for TM housekeeping purposes.

(SCGCQ00641946) Defect 7/13

HEADLINE: IOP: Recursive fault may not report fault state to host
DESC OF CHANGE: Added code in the recursive fault handler to set the doorbell to indicate a fault was hit.
TO REPRODUCE: Cause firmware to fault recursively before the firmware sets the doorbell to indicate fault. The doorbell will never indicate the fault condition.
ISSUE DESC: In some conditions, a recursive fault may occur such that the doorbell is not updated to indicate the controller firmware entered the fault state. This can cause the host to exhibit indeterminate behavior.

(SCGCQ00645292) Defect 8/13

HEADLINE: IOP: MCTP messages: Test size of MPI Event Notification Request in MCTP
DESC OF CHANGE: Added code to properly sanitize the size of an MPI Event Notification Request in MCTP.
TO REPRODUCE: Send MPI Event Notification Requests over MCTP, both with the proper size and too small of size.
ISSUE DESC: The size of an MPI Event Notification Request in MCTP was not properly sanitized. This lead to correct requests being improperly rejected, and improper requests being rejected with the wrong error code.

(SCGCQ00645300) Defect 9/13

HEADLINE: IOP: MCTP messages: RAID SCSI IO Passthrough MPI requests timed out in IR
DESC OF CHANGE: Modified the code that routes the RAID SCSI IO Passthrough requests to pass the MPI Request to IR.
TO REPRODUCE: Send a RAID SCSI IO Passthrough MPI request over MCTP. Other IOs may be timed out and all of them removed by a TM.
ISSUE DESC: RAID SCSI IO Passthrough MPI requests were being timed out by IR. SCSI IOs including RAID SCSI IO Passthroughs received over MCTP were only handled internally within the IOP, but caused other IOs to timeout in IR and removed by a TM. The IO from MCTP would also be removed by the TM.

Additionally, IR has some checks, which prevent writes to the drives, which were being bypassed. Bypassing IR allows SCSI writes to occur that IR would block. The preference is to allow IR to manage these policies.

(SCGCQ00645484) Defect 10/13

HEADLINE: WW: (Capella-2) NVMe initialization times out.
DESC OF CHANGE: Address traps have been added for Invader's inbound MSI-X region for each port on the switch that has an NVMe device attached.
TO REPRODUCE: Attach Invader running Whirlwind firmware to a Capella-2 switch. Attach one or more NVMe devices to the switch. Attach a host to the switch. Allow initial enumeration to complete and the host to boot and load the OS. Issue a port enable request from the host. On port enable complete, read PCIe device page 0 for each NVMe device. Observe a non-zero status in the AccessStatus field of the PCIe device page 0 for each NVMe device.
ISSUE DESC: NVMe initialization times out when Whirlwind is used with Capella-2 PCIe switch.

(SCGCQ00637311 - Port of SCGCQ00600515) Defect 11/13

HEADLINE: NVMe: handling of admin queue full condition during NVMe async event handling
DESC OF CHANGE: Added handling of the admin queue full condition during NVMe async event handling.
TO REPRODUCE: Code inspection.
ISSUE DESC: If an NVMe device's admin queue is found to be full during NVMe asynchronous event handling, an assert (fault) will occur.

(SCGCQ00643315 - Port of SCGCQ00592317) Defect 12/13

HEADLINE: when Phase 18 IR FW boot a RAID volume (such as raid 10), 2308 bootup hangs at OROM "Initializing.." and fault 8803 has seen
DESC OF CHANGE: Changed the allocation of memory for configuration during device discovery.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

TO REPRODUCE: Concatenate customer's NVDATA with high host credit values with raw FW. Fault 0x8803 fault has occurred upon reboot after creating RAID volume.

ISSUE DESC: Concatenate customer's NVDATA with high host credit values with raw FW. Fault 0x8803 fault has occurred upon reboot after creating RAID volume.

(SCGCQ00648184 - Port of SCGCQ00632446)

Defect 13/13

HEADLINE: Need a method to disable PL Light print in IO path

DESC OF CHANGE: Compiled out PL Light print in IO path.

TO REPRODUCE: Run IOs via PL light and dump the ringbuffer. Even after masking all IO path prints, I see the PL light print in the ringbuffer for every IO completion.

ISSUE DESC: iopiDefs is not listing a debug mask to disable the below PL Light prints. We either need to delete this print in IO path or need a mask definition to disable this print:
FPE IO owned by PL Light!! Mid|DevH [8da000e], MidUsed[T/_F] = 0x0.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

Total Enhancements Implemented (8)

(SCGCQ00577387) Enhancement 1/8

HEADLINE: Temperature Monitoring with bi-colored LED and over temp faulting.

NEW FUNCTIONALITY:

- 1) Created a New Function Code in Man Page 6 for Overtemp LED 0x1A.
- 2) Created a New Fault Code to identify the Overtemp Threshold has exceeded in iopiFault.h
#define IFAULT_IOP_OVER_TEMP_THRESHOLD_EXCEEDED (0x2810)
- 3) Updated overtemp threshold monitoring to make Overtemp LED act as the Heartbeat LED when any of the 4 thresholds for any of the discovered sensors has been exceeded.
When temperature is back to normal operating temperature i.e., below all threshold values, then the normal Heartbeat LED will resume its role. This happens only when Overtemp LED is defined in ManPage6.
- 4) When any of the sensors' Threshold 3 is exceeded, the FW will fault the controller with the fault code IFAULT_IOP_OVER_TEMP_THRESHOLD_EXCEEDED (0x2810).

Bi-Colored LED usage is as follows. Primary color is used as Heartbeat LED and the secondary color becomes the Heartbeat LED when an Overtemp condition occurs.
When the temperature is below all thresholds the normal Heartbeat LED will resume its function and the primary color would be used.
SYS_HB_LED signal can be wired up to the primary color and any of the GPIO pins can be wired up to the secondary color.

For additional information see document "SAS Gen-3 Temperature Monitoring".

(SCGCQ00584505) Enhancement 2/8

HEADLINE: RAID Flags and CDB fields are corrupted after Fast Path Engine rejects FastPath ATA NATIVE COMMAND

NEW FUNCTIONALITY: The handler for MPI SCSI IO request ATA Pass-through 12/16 CDBs has been enhanced to store and restore the original IO request.

(SCGCQ00601742) Enhancement 3/8

HEADLINE: WW: Add support for DirectIO Q pair

NEW FUNCTIONALITY: Implemented a new PL interface function pliWWPCleDirectIoCreateQ(), which can be used to create a new IO Queue (Completion/Submission) for an NVMe device. This queue can be later used to send IOs to the NVMe device directly from the host.

ISTATUS pliWWPCleDirectIoCreateQ(PLI_DEV_HANDLE DeviceHandle, PTR_NVME_IO_Q PtrNvmeloQ);

ARGUMENTS:
DeviceHandle - Device the Q create request is for
PtrNvmeloQ - Pointer to the queue creation information

```
typedef struct _NVME_IO_Q
{
    U32 QueueID;
    U32 QueueType; /* Type of IO Q to create Completion/Submission */
    U32 QueueSize;
    U32 PCIQueueBaseAddressHigh;
    U32 PCIQueueBaseAddressLow;

    /* Below fields are applicable only for IO Completion Q */
    U32 PCIInterruptBaseAddressHigh;
    U32 PCIInterruptBaseAddressLow;
    U32 MSIXVector;
    U32 MSIXData;
} NVME_IO_Q, *PTR_NVME_IO_Q;
```

(SCGCQ00621996) Enhancement 4/8

HEADLINE: WW: account for data length in SGL during translation to NVMe PRP

NEW FUNCTIONALITY: Add Datalength to EOL check for SGL termination during IEEE-64 SGL translation to NVMe PRP.

(SCGCQ00626865) Enhancement 5/8

HEADLINE: IOP MCTP: Add support for VDM Buffers



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.02.00 - SAS3FW_MASTER_DEV (SCGCQ00649048)

NEW FUNCTIONALITY: Added support for a Vendor Defined Message Buffer Size. This buffer size eliminates the need for the custom VDM header to be attached to every packet in an MCTP VDM message to a controller. This helps to optimize two use cases, which may be seen together, by sacrificing, but not fully eliminating, the advanced error detection and flow control features:

1. The layer that actually handles the data in the VDM packets and performs the packetization and cannot handle IT firmware's custom VDM header, so the upper layer has to perform this packetization process.
2. The layer that actually transmits each packet experiences a context switch to actually send each packet when this layer is not performing the packetization itself.

(SCGCQ00644379) Enhancement 6/8

HEADLINE: IOP: Add compile option to start the second CPU
NEW FUNCTIONALITY: Add a compile option to create a build which starts the second CPU. This change does not affect standard firmware builds.

(SCGCQ00645401) Enhancement 7/8

HEADLINE: PL: Poll for parity error in FPE PLB Slave
NEW FUNCTIONALITY: Add code to poll for a parity error in the Fast Path Engine's PLB Slave. If a parity error is detected, the firmware will fault.

(SCGCQ00638811 - Port of SCGCQ00398216) Enhancement 8/8

HEADLINE: Additional diag console commands
NEW FUNCTIONALITY: 'pl pci nvmeq' command for dumping NVMe device queues. Also added code in plDiagCmd.c to dump PCI device and device details.

Change Summary (Defects=45 Enhancements=40 NVDATA=1)

- SCGCQ00551008 (DFCT) - Unexpected IOCStatus with invalid config page type
- SCGCQ00559724 (DFCT) - Whirlwind: Fix Inbound MSIx Interrupt firing and NVMe init flow
- SCGCQ00562318 (DFCT) - Whirlwind: Use PliRegister for device registration instead of hardcoded value
- SCGCQ00575816 (DFCT) - WW: Put in WHIRLWIND defines in the version files
- SCGCQ00593463 (DFCT) - Whirlwind: Allow MSIx enable for a device to correctly look at the PIM0 size instead of assuming a 16K PIM0
- SCGCQ00594674 (DFCT) - Whirlwind: Read Inbound MSIx table offset directly from PCI BARs
- SCGCQ00600832 (DFCT) - Fault 265D when running large block IO with higher Qdepth
- SCGCQ00601393 (DFCT) - WW: Invader BAR was only 1MB aligned, needed increased to allow for large PIM0 windows
- SCGCQ00601480 (DFCT) - WW: IOs not unpended if put on pend list due to out of resources
- SCGCQ00604446 (DFCT) - Hitting PL_LOGININFO_CODE_NVME_SGL_INVALID_INT_ADDRESS for READ command
- SCGCQ00605087 (DFCT) - After Controller Reset, Sense Data of REQUEST SENSE SCSI Command data is incorrect when SATA Drive is in Stand By Mode
- SCGCQ00606389 (DFCT) - Whirlwind build was broken after pulling over Invader Phase 4 tip code
- SCGCQ00609576 (DFCT) - Fixed Issue in Transport Layer Retries Code Causing 0x265D and 0x0501
- Faults
- SCGCQ00612665 (DFCT) - WW: Root Complex not enabled by default in NVDATA.
- SCGCQ00613906 (DFCT) - WW code was in different spot in file vs Ventura
- SCGCQ00614457 (DFCT) - NVME devices did not work in all slots of Capella 1 RDK
- SCGCQ00614550 (DFCT) - NVData: MfgPg2 is configured for 16 Virtual Functions with non-SRIOV
- firmware
- SCGCQ00615164 (DFCT) - Whirlwind code was breaking Invader builds
- SCGCQ00622522 (DFCT) - PL: PortWidth field in SAS Port Page 0 is incorrect
- SCGCQ00622844 (DFCT) - Whirlwind incorrectly reporting MPI version as 2.5
- SCGCQ00622857 (DFCT) - Invader build broke due to Whirlwind code added without correct #define
- SCGCQ00626766 (DFCT) - WW: Only 4 NVMe drives reported to host when 5 NVMe drives are attached to the PLX switch.
- SCGCQ00626966 (DFCT) - Fixed multiple order of operation bugs
- SCGCQ00626976 (DFCT) - Allocation Length of SCSI INQUIRY Command for "Power Condition VPD page" is not Handled Properly
- SCGCQ00629210 (DFCT) - Bad TLB entry causes 0x26B5 fault in IT/IR firmware
- SCGCQ00630689 (DFCT) - Whirlwind: Fault 0x6001 While running TM Target Reset to an NVMe drive with device READY_WAIT active
- SCGCQ00631763 (DFCT) - Whirlwind: Fault 6001 during NVMe Target reset while running CM block move.
- SCGCQ00632642 (DFCT) - Remove IT files form WW release orders
- SCGCQ00633121 (DFCT) - WW: IOCFacts was not showing NVME support in ProtocolFlags
- SCGCQ00633161 (DFCT) - WW- Incorporate review comments related to coding standard for the ER
- SCGCQ00601719
- SCGCQ00634797 (DFCT) - Whirlwind: LBA not getting populated for a corner case of breakup IOs
- SCGCQ00638611 (DFCT) - IR firmware fault 0x265D on RAID Volume creation or activation
- SCGCQ00599998 (CSET) - Raid Accelerator IOs can hang if the stream of IOs suddenly stops
- SCGCQ00613946 (CSET) - NVMe: Fix several mode page related bugs
- SCGCQ00614498 (CSET) - 0x6828 fault when using large CDB
- SCGCQ00615115 (CSET) - Few minor backend PCIe related changes
- SCGCQ00620269 (CSET) - MCTP GET_INVENTORY returning wrong deviceType
- SCGCQ00626868 (CSET) - NVMe: IEEE to PRP translator rollback issues

SCGCQ00628040 (CSET) - NVMe: fix function headers
SCGCQ00628045 (CSET) - NVMe: LUN not validated before use in non read write translations
SCGCQ00628047 (CSET) - NVMe: Fault EC0F during IO
SCGCQ00628089 (CSET) - Internal Device Reset storm after controller reset using Long CDB.
SCGCQ00629153 (CSET) - Target Mode: Task IU is sent to the IOP before it knows about the Initiator
SCGCQ00630674 (CSET) - PL: Task management aborted IO search may return faulty information
SCGCQ00635733 (CSET) - IOP: PowerPC ITLB Miss Errata Workaround
SCGCQ00557357 (ENHREQ) - WhirlWind : PL : Remove PCIE core (Trident) registers access from PL -

Part 1

SCGCQ00557883 (ENHREQ) - Whirlwind: Implement Configuration Space Request Redirection (CSR).
SCGCQ00559130 (ENHREQ) - Whirlwind: (Master) One Massive ER
SCGCQ00564420 (ENHREQ) - Convert direct PCIE POM Reads to using PCI mem read instead
SCGCQ00566403 (ENHREQ) - Added WW specific NVDATA and build
SCGCQ00567780 (ENHREQ) - Add support for new IO Unit Control Request (in MPI 2.6)
SCGCQ00567802 (ENHREQ) - Use Man Page 9 configured values for PCIe devices
SCGCQ00567813 (ENHREQ) - Pull updates in scsi.h from SAS3.5 project
SCGCQ00568791 (ENHREQ) - WhirlWind : PL : Completely remove pci msix register usage from PL .
SCGCQ00569101 (ENHREQ) - Add MULTI script to flash WW build
SCGCQ00573598 (ENHREQ) - Implement Read/Write NVM Commands:
SCGCQ00580164 (ENHREQ) - Whirlwind : Remove the PCI Pepp reg (gplHwPtrPciePeppRegs) and FMU reg (gplHwPtrPcieFMURegs) usage from PL.
SCGCQ00580338 (ENHREQ) - WW: Configure GEP BAR0 as 64-bit (Capella-2 FPGA only)
SCGCQ00580340 (ENHREQ) - WW: (Capella-2) CSR buffer allocation.
SCGCQ00583839 (ENHREQ) - Whirlwind: Further IO path tweaks.
SCGCQ00592594 (ENHREQ) - WW: (Capella-2) Establish communication between host and Invader root complex.
SCGCQ00593490 (ENHREQ) - Whirlwind: Add support to dump all PL root complex related debug info in one CLI command
SCGCQ00595519 (ENHREQ) - Selective Extended Image Loading
SCGCQ00601719 (ENHREQ) - WW: Switch EEPROM programming
SCGCQ00605758 (ENHREQ) - Whirlwind: NVMe Task Management Target Reset Support
SCGCQ00606386 (ENHREQ) - WW: (Capella-2) Establish communication to/from NVMe devices.
SCGCQ00610550 (ENHREQ) - Add NVMe encapsulated IO support
SCGCQ00615216 (ENHREQ) - Whirlwind: NVMe target reset (Part II)
SCGCQ00618535 (ENHREQ) - Whirlwind: NVMe Task Abort Implementation
SCGCQ00619353 (ENHREQ) - Add support for Thunderbolt based NVMe emulator
SCGCQ00619545 (ENHREQ) - IOP MCTP: Add basic MCTP Control Commands
SCGCQ00620222 (ENHREQ) - IOP MCTPMsg: Add MPI RAID Action and RAID SCSI IO Passthrough
SCGCQ00620227 (ENHREQ) - IOP MCTP: MPI Event registration, Extended Events command, and transmitting event datagrams
SCGCQ00623300 (ENHREQ) - Whirlwind: Root Complex MSIx setup
SCGCQ00630611 (ENHREQ) - Whirlwind: Improve NVME Initialization Error Handling
SCGCQ00631600 (ENHREQ) - IOP MCTP: Add support for resetting EID in MCTP Control Set EID command
SCGCQ00634740 (ENHREQ) - WW: added BDF 0:5:0 to requestor ID table
SCGCQ00575849 (CSET) - Add in support for CM Build app
SCGCQ00606270 (CSET) - NVMe: Extended Inquiry Data VPD page
SCGCQ00615113 (CSET) - NVMe: Support for Mode Sense (10) LLBAA flag
SCGCQ00615114 (CSET) - NVMe: a couple of mode page tweaks



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

SCGCQ00625579 (CSET) - PL: New commands useful for debugging
SCGCQ00625580 (CSET) - SATL: Indicate support for the ATA Device Server Password security protocol
SCGCQ00625583 (CSET) - Print SES page data length and PL SES Diag page buffer size to ring buffer
SCGCQ00625622 (CSET) - Correct PHY ordering in Channel NVData SGPIO configuration
SCGCQ00624122 (CSET) - Gen3 Channel NVDATA: Tighten Outstanding IO range



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

Total Defects Resolved (45)

(SCGCQ00551008) Defect 1/45

HEADLINE: Unexpected IOCStatus with invalid config page type
DESC OF CHANGE: An 'if' condition has been added in the code to differentiate error flag due to invalid type or invalid page.
TO REPRODUCE: Send any IOP config request with an invalid page type. In IOC config reply you will notice invalid page instead of invalid type.
ISSUE DESC: When IOP config request is sent with invalid type it is observed that FW reports error with invalid page while invalid type is expected

(SCGCQ00559724) Defect 2/45

HEADLINE: Whirlwind: Fix Inbound MSix Interrupt firing and NVMe init flow
DESC OF CHANGE: Force firmware to handle admin commands without trying to pass it to Stang first, since Whirlwind doesn't have Stang HW
TO REPRODUCE: Connect an NVMe drive to the PLX switch and try to do NVMe init.
ISSUE DESC: Fix Inbound MSix Interrupt firing and NVMe init flow

(SCGCQ00562318) Defect 3/45

HEADLINE: Whirlwind: Use PliRegister for device registration instead of hardcoded value
DESC OF CHANGE: Use PliRegister for device registration instead of hardcoded value
TO REPRODUCE: n/a
ISSUE DESC: Use PliRegister for device registration instead of hardcoded value

(SCGCQ00575816) Defect 4/45

HEADLINE: WW: Put in WHIRLWIND defines in the version files
DESC OF CHANGE: Modified the version file so that the makefile can dynamically pick the appropriate version based on build type.
TO REPRODUCE: n/a
ISSUE DESC: Whirlwind and non-whirlwind firmware builds need to have different firmware versions.

(SCGCQ00593463) Defect 5/45

HEADLINE: Whirlwind: Allow MSix enable for a device to correctly look at the PIM0 size instead of assuming a 16K PIM0
DESC OF CHANGE: Allow MSix enable for a device to correctly look at the PIM0 size instead of assuming a 16K PIM0
TO REPRODUCE: n/a
ISSUE DESC: Previously, MSix enable for a device for hardcoded for 16K PIM0 size.

(SCGCQ00594674) Defect 6/45

HEADLINE: Whirlwind: Read Inbound MSix table offset directly from PCI BARs
DESC OF CHANGE: Read Inbound MSix table offset directly from PCI BARs
TO REPRODUCE: n/a
ISSUE DESC: Read Inbound MSix table offset directly from PCI BARs

(SCGCQ00600832) Defect 7/45

HEADLINE: Fault 265D when running large block IO with higher Qdepth
DESC OF CHANGE: Ensure that in IO breakup, if we try to allocate memory for SGL, we actually get the memory. If not, just pend the IO. Note there looks to be follow on issue here with IOs being pended forever and timing out, however that will not be resolved in this defect.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

TO REPRODUCE: Run CM bmove with 1-1024 block size and exceeding about ~20 IOs outstanding to 2 NVME devices (emulators)
ISSUE DESC: When running 1-1024 block size and exceeding about ~20 IOs outstanding to 2 NVME devices (emulators), fault 265D is seen

(SCGCQ00601393) Defect 8/45

HEADLINE: WW: Invader BAR was only 1MB aligned, needed increased to allow for large PIM0 windows
DESC OF CHANGE: Change Invader BAR to be 32MB aligned, which should allow for any feasible sized PIM0 window (we only have 6MB internal memory).
TO REPRODUCE: N/A
ISSUE DESC: Invader BAR was only 1MB aligned, needed increased to allow for large PIM0 windows

(SCGCQ00601480) Defect 9/45

HEADLINE: WW: IOs not unpended if put on pend list due to out of resources
DESC OF CHANGE: Correctly unpending IO if it was pended prior due to out of resource.

NOTE: there is still follow on issue of check conditions occurring when we exceed the ~20 IO threshold, that will be handled separate
TO REPRODUCE: Run >20 Qdepth to 2 NVME devices with 1-1024 IO size
ISSUE DESC: IOs not unpended if put on pend list due to out of resources

(SCGCQ00604446) Defect 10/45

HEADLINE: Hitting PL_LOGININFO_CODE_NVME_SGL_INVALID_INT_ADDRESS for READ command
DESC OF CHANGE: Added support to store the complete 64 bit PLB address into the internal PIM POM table, so that it could be used for the PLB -> PCI translation, when translating the PLB address in the SGEs to PCI address for the PRPs.
TO REPRODUCE: 1. Connect the Invader RootComplex and NVMe emulator to the PLX switch
2. Allow the FW to discover the NVMe device
3. During the device initialization, IOP is sending a READ command which failed with PL_LOGININFO_CODE_NVME_SGL_INVALID_INT_ADDRESS
ISSUE DESC: Hitting PL_LOGININFO_CODE_NVME_SGL_INVALID_INT_ADDRESS for READ command.

(SCGCQ00605087) Defect 11/45

HEADLINE: After Controller Reset, Sense Data of REQUEST SENSE SCSI Command data is incorrect when SATA Drive is in Stand By Mode
DESC OF CHANGE: During Sata Initialization we record the current Standby Mode of the Sata Drive. This is used later on when a Request Sense is received.
TO REPRODUCE: Connect a SATA Drive to SAS3 controller and execute a SSU command to move the drive from active to stand by mode. Execute a Request Sense command to verify that drive is in stand by power condition. Now perform controller reset using lsiutils. Now execute a request sense command to SATA drive. Observe that the drive state move from stand by to stopped which seems to be incorrect as per section 8.10 of SAT3R04

Setup:
HBA(Invader) ----> SATA Drives
ISSUE DESC: Observe that the drive state move from stand by to stopped which seems to be incorrect as per section 8.10 of SAT3R04

(SCGCQ00606389) Defect 12/45

HEADLINE: Whirlwind build was broken after pulling over Invader Phase 4 tip code
DESC OF CHANGE: Rename field so field name is correct per Invader Phase 4 and latest Ventura as well
TO REPRODUCE: N/A
ISSUE DESC: Whirlwind build was broken after pulling over Invader Phase 4 tip code, a structure field was not merged properly



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00609576) Defect 13/45

HEADLINE: Fixed Issue in Transport Layer Retries Code Causing 0x265D and 0x0501 Faults
DESC OF CHANGE: Before freeing a node on the doubly linked list which holds TLR info, update (nextNode)->PrevNode.
TO REPRODUCE: Use a topology with cascaded LSI gen 3 expanders connected to an LSI gen 3 Invader HBA, with a combination of SAS and SATA drives off each expander. Run read only stress to all drives, with an IO size of 512k bytes. Perform link breaks on random links between the Invader and first expander, 500ms-5sec duration, every 10-20 seconds.
ISSUE DESC: There is a doubly linked list that holds info about mids undergoing TLR (transport layer retries). Before a node is freed (nextNode)->PrevNode was not being updated. This caused the linked list to get messed up, leading to the two different faults.

(SCGCQ00612665) Defect 14/45

HEADLINE: WW: Root Complex not enabled by default in NVDATA.
DESC OF CHANGE: In NVDATA of WW, Root Complex is not enabled by default. Byte 0x10 in Man page 2 has a value of 0xDF when it should be 0xFF.
TO REPRODUCE: n/a
ISSUE DESC: In NVDATA of WW, Root Complex is not enabled by default. Byte 0x10 in Man page 2 has a value of 0xDF when it should be 0xFF.

(SCGCQ00613906) Defect 15/45

HEADLINE: WW code was in different spot in file vs Ventura
DESC OF CHANGE: Moved to correct spot
TO REPRODUCE: N/A
ISSUE DESC: WW code was in different spot in file plcontextglobals vs Ventura, making comparison difficult

(SCGCQ00614457) Defect 16/45

HEADLINE: NVME devices did not work in all slots of Capella 1 RDK
DESC OF CHANGE: Added upper slot Bus Device Function values to the Capella 1 Requestor ID table to allow them to send data to host for IO.
TO REPRODUCE: Put devices in upper slots of RDK, it won't be able to do IO to host
ISSUE DESC: NVME devices did not work in all slots of Capella 1 RDK

(SCGCQ00614550) Defect 17/45

HEADLINE: NVData: MfgPg2 is configured for 16 Virtual Functions with non-SRIOV firmware
DESC OF CHANGE: Modified the NVDATA to test if SRIOV FW is built. If so, use 0x10 VFs, otherwise use 0x00.
TO REPRODUCE: Read Manufacturing Page 2 and analyze SBR byte 0x47.
ISSUE DESC: In Manufacturing Page 2 the SBR byte 0x47 indicates 0x10, which configures the controller with 16 Virtual Functions, even with non-SRIOV firmware loaded. We should be indicating 0 Virtual Functions for non-SRIOV configurations.

(SCGCQ00615164) Defect 18/45

HEADLINE: Whirlwind code was breaking Invader builds
DESC OF CHANGE: Put WW code in PL_PCI_DEVICE_SUPPORT define and removed code in Invader build that depended on MPI 2.6 defines
TO REPRODUCE: N/A
ISSUE DESC: Whirlwind code was breaking Invader builds

(SCGCQ00622522) Defect 19/45

HEADLINE: PL: PortWidth field in SAS Port Page 0 is incorrect
DESC OF CHANGE: Made a change to have firmware track the number of linked up phys in each port. This new value will be reported to the host instead of the total number of phys, both up and down, in the port.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

TO REPRODUCE: Configure the controller to use Auto Port Configuration. Connect a single phy to phy 0. Observe SAS Port Page 0 reports a portwidth greater than 1.

ISSUE DESC: The PortWidth and PhysicalPortWidth fields in SAS Port Page 0 are reporting more phys than are linked up to the controller. It is expected it will report the number of phys which are linked up and in the port.

(SCGCQ00622844) Defect 20/45

HEADLINE: Whirlwind incorrectly reporting MPI version as 2.5

DESC OF CHANGE: Update code to use MPI 2.6 version define on Whirlwind build.

TO REPRODUCE: view IOC Fact reply, see version is MPI 2.5

ISSUE DESC: Whirlwind incorrectly reporting MPI version as 2.5

(SCGCQ00622857) Defect 21/45

HEADLINE: Invader build broke due to Whirlwind code added without correct #define

DESC OF CHANGE: Added correct #define around the Whirlwind only code

TO REPRODUCE: N/A

ISSUE DESC: Invader build broke due to Whirlwind code added without correct #define

(SCGCQ00626766) Defect 22/45

HEADLINE: WW: Only 4 NVMe drives reported to host when 5 NVMe drives are attached to the PLX switch.

DESC OF CHANGE: Modified the default values of Manufacturing Page 9 parameter for PCIe device allocation such that memory resources are allocated to accommodate at least 8 NVMe drives.

TO REPRODUCE: Attach more than 4 NVMe drives to the PLX switch in the Whirlwind setup. Issue Port Enable request to the controller. Observe only 4 NVMe drives are reported/seen at the host.

ISSUE DESC: Controller running Whirlwind firmware presents only 4 NVMe drives to the host even though more than 4 NVMe drives are attached to the PLX switch.

(SCGCQ00626966) Defect 23/45

HEADLINE: Fixed multiple order of operation bugs

DESC OF CHANGE: Added parenthesis to ensure the original intent of the code was being performed.

TO REPRODUCE: Not Applicable

ISSUE DESC: Code inspection revealed a number of order of operation issues throughout the firmware.

(SCGCQ00626976) Defect 24/45

HEADLINE: Allocation Length of SCSI INQUIRY Command for "Power Condition VPD page" is not Handled Properly

DESC OF CHANGE: We changed the check to look at the DataLength instead of the Transfer Length. When determining a data underrun. This data underrun condition is interpreted by the host driver as success or fail based on the allocation length and data transferred.

TO REPRODUCE: Connect a SATA Drive to SAS3 Controller where drive supports VPD page 0x8a (Power Condition VPD page). Execute an Inquiry with ALLOCATION LENGTH=0x00 to Power Condition VPD page which pass with data whereas command should pass with no data.
Setup:-
HBA(Invader) ----> Cobra ----> 2 SATA Drives

ISSUE DESC: As per Section 4.2.5.6 of spc4r36, the command should be passed but not data shall be transferred.

(SCGCQ00629210) Defect 25/45

HEADLINE: Bad TLB entry causes 0x26B5 fault in IT/IR firmware

DESC OF CHANGE: Compiled out a Whirlwind TLB which does not belong in IT/IR firmware.

TO REPRODUCE: Boot the controller with Phase 5 IT/IR firmware. The controller will fault with code 0x26B5 during its initialization.

ISSUE DESC: Firmware will fault with code 0x26B5 when the IT/IR firmware boots.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00630689) Defect 26/45

HEADLINE: Whirlwind: Fault 0x6001 While running TM Target Reset to an NVMe drive with device READY_WAIT active
DESC OF CHANGE: Pause NVMe init until TM has finished if the device becomes ready while a TM is in progress.
TO REPRODUCE: Run repeated target reset with short gap between successive TMs to a single NVMe drive.
ISSUE DESC: A 6001 fault occurs if an NVMe device becomes ready while a target reset is in progress on that device.

(SCGCQ00631763) Defect 27/45

HEADLINE: Whirlwind: Fault 6001 during NVMe Target reset while running CM block move.
DESC OF CHANGE: When a breakup IO has to be aborted by a TM, decrement IoCount before completing the original IO with error to retain IO count balance
TO REPRODUCE: Run CM block move test + CM Target reset test to a single NVMe drive.
ISSUE DESC: A break up IO getting added to the Abort List by a TM. When the TM completes, the orinal IO for the breakup IO is correctly getting completed to the host. However, the ioCount is not getting correctly drecemented, resulting in IoCount disbalance and Fault 6001.

(SCGCQ00632642) Defect 28/45

HEADLINE: Remove IT files form WW release orders
DESC OF CHANGE: Removed IT files form WW release orders
TO REPRODUCE: N/A
ISSUE DESC: Remove IT files form WW release orders

(SCGCQ00633121) Defect 29/45

HEADLINE: WW: IOCFacts was not showing NVME support in ProtocolFlags
DESC OF CHANGE: Add NVME support to IOCFacts ProtocolFlags
TO REPRODUCE: Inspect IOCFacts reply
ISSUE DESC: WW: IOCFacts was not showing NVME support in ProtocolFlags

(SCGCQ00633161) Defect 30/45

HEADLINE: WW- Incorporate review comments related to coding standard for the ER SCGCQ00601719
DESC OF CHANGE: Incorporated all the comments
TO REPRODUCE: NA
ISSUE DESC: WW- Incorporate review comments related to coding standard for the ER SCGCQ00601719

(SCGCQ00634797) Defect 31/45

HEADLINE: Whirlwind: LBA not getting populated for a corner case of breakup IOs
DESC OF CHANGE: Get LBA info for the resource MID before forming the NVM PRP.
TO REPRODUCE: n/a
ISSUE DESC: LBA not getting populated for a corner case of breakup IOs

(SCGCQ00638611) Defect 32/45

HEADLINE: IR firmware fault 0x265D on RAID Volume creation or activation
DESC OF CHANGE: Moved a check for enabling/disabling MCTP before a potentially uninitialized variable is dereferenced.
TO REPRODUCE: Flash IR firmware on a controller along with the LSI BIOS CU. Attach empty drives to the system. Enter the BIOS CU and configure a RAID Volume. Firmware will fault before the volume is created.
ISSUE DESC: Firmware faults with code 0x265D when a RAID Volume is created or activated. If the LSI BIOS is configured, this may stall the boot process.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00599998 - Port of SCGCQ00503412)

Defect 33/45

HEADLINE: Raid Accelerator IOs can hang if the stream of IOs suddenly stops

DESC OF CHANGE: If more than 5 IOs complete on one completion interrupt and no further IOs are started, the code will now complete any IOs that remain after completing the first 5 IOs.

TO REPRODUCE: Run heavy Raid Accelerator IO, suddenly stop sending new IOs and monitor for all IOs to complete. If the timing window is hit, there will be a burst of completions followed by nothing where one or more outstanding IOs will not complete (as long as new IOs are not started).

ISSUE DESC: When a stream of Raid Accelerator IOs suddenly stops, there is a chance that a few of the IOs will never be completed back to the host. This only happens if more than 5 IOs all complete at the same time with no further IOs being started or completed.

(SCGCQ00613946 - Port of SCGCQ00570058)

Defect 34/45

HEADLINE: NVMe: Fix several mode page related bugs

DESC OF CHANGE: Fixed several bugs related to NVMe mode page emulation:

- * DPICZ flag should be set in the Control mode page.
- * The short LBA mode parameter block descriptor may be used with either Mode Select (6) or Mode Select (10), not only Mode Select (6).
- * Allow device write cache state to be changed with the WCE bit in the Caching mode page only if the device supports write cache.
- * Don't send a command to the device to change write cache state if the new state is identical to the current state.
- * Validate block descriptor length and count contained within Mode Select parameter data.
- * When validating Mode Select parameter data, use the NVMe version of certain mode pages if the command was sent to an NVMe device.

TO REPRODUCE: Code inspection.

ISSUE DESC: Code inspection revealed several bugs related to NVMe Mode Sense and Mode Select command and parameter data translation.

(SCGCQ00614498 - Port of SCGCQ00580326)

Defect 35/45

HEADLINE: 0x6828 fault when using large CDB

DESC OF CHANGE: Add code to retrieve MID from SSP_PAYLOAD_ONLY dma entry.

TO REPRODUCE: Create a scsi io request with large CDB. Cause long period of open reject no dest for over 50ms on the expander side by using jammer or other means.

ISSUE DESC: In the failing case, there was a port error and firmware called a function to get the MID from a particular tx dma entry but there was no code to handle SSP_PAYLOAD_ONLY tx dma type.

(SCGCQ00615115 - Port of SCGCQ00613904)

Defect 36/45

HEADLINE: Few minor backend PCIe related changes

DESC OF CHANGE: Made several minor modifications to the backend PCIe code.

TO REPRODUCE: Code inspection.

ISSUE DESC: Make several minor fixes that were noticed while working on other things or were suggested improvements.

(SCGCQ00620269 - Port of SCGCQ00615192)

Defect 37/45

HEADLINE: MCTP GET_INVENTORY returning wrong deviceType

DESC OF CHANGE: Fixed a logic error that generates the device type field in an GET_INVENTORY response.

TO REPRODUCE: Send a GET_INVENTORY command after an INIT command to an MCTP enabled controller. Port enable must have run to see attached drives. Attach drives. The device type results for attached drives have too many bits set, indicating the wrong device type.

ISSUE DESC: The SCS MCTP command GET_INVENTORY was returning the wrong device type.

(SCGCQ00626868 - Port of SCGCQ00626787)

Defect 38/45

HEADLINE: NVMe: IEEE to PRP translator rollback issues

DESC OF CHANGE: Properly detect the rollback-across-SGE case and use the correct SgeOffset when it occurs. If rollback would be required across multiple SGEs, throw an error.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

TO REPRODUCE: Code inspection.
ISSUE DESC: The IEEE to NVMe PRP list translation code won't detect a case where it should back up the SGL pointer during SGL translation across an SGE boundary. If it detects another case, an incorrect SgeOffset will be used.

(SCGCQ00628040 - Port of SCGCQ00562246) Defect 39/45

HEADLINE: NVMe: fix function headers
DESC OF CHANGE: Modified function header comments throughout the NVMe code. There is no expected change in functionality from these modifications.
TO REPRODUCE: Code inspection.
ISSUE DESC: Some of the function header comments in the NVMe code are blank or incomplete.

(SCGCQ00628045 - Port of SCGCQ00562314) Defect 40/45

HEADLINE: NVMe: LUN not validated before use in non read write translations
DESC OF CHANGE: Unless the request contains an Inquiry, Report Luns, or Request Sense command, validate the LUN upon receipt of a SCSI IO request to an NVMe device and fail the request if it is invalid.
TO REPRODUCE: Send a non read write SCSI IO request containing an invalid LUN to a backend NVMe device.
ISSUE DESC: If an invalid LUN is passed in with a non read write SCSI IO request to an NVMe device, for most commands the request should be failed. Currently this is not done.

(SCGCQ00628047 - Port of SCGCQ00556768) Defect 41/45

HEADLINE: NVMe: Fault EC0F during IO
DESC OF CHANGE: Clear additional context flags when starting a PL translated read/write IO.
TO REPRODUCE: Send a Write Buffer, Security Protocol In, or Security Protocol Out command to an NVMe device. Then run a blocksize 1-128 read/write IO test against it.
ISSUE DESC: If a non read write IO with a translation that directly translates the SGL is sent to an NVMe device, and PL translated read/write IOs are then sent to it, an EC0F fault will occur.

(SCGCQ00628089 - Port of SCGCQ00626891) Defect 42/45

HEADLINE: Internal Device Reset storm after controller reset using Long CDB.
DESC OF CHANGE: New code was added to check for the special tx dma type when long cdb is used. Firmware is now able to clean up this type of entry from the transport layer.
TO REPRODUCE: Invader->Cobra->Invader target mode
ISSUE DESC: While controller is sending long CDB mpi scsi io requests to a target device, reset the expander. Long CDB makes use of special dma type. There's no code in gen2 or gen3 firmware to clean out this type of dma entry during hw clean up. Therefore this type of entry would get stuck in the transport layer and keep requesting port layer to open a connection. Eventually our open reject timer would expire and send out target reset, but we still couldn't clean out those entries. The same thing will keep repeating itself.

(SCGCQ00629153 - Port of SCGCQ00606248) Defect 43/45

HEADLINE: Target Mode: Task IU is sent to the IOP before it knows about the Initiator
DESC OF CHANGE: - If the normal get frame function fails when trying to send the initiator added device status change event to IOP and upper layers, try to get a critical frame.
- When a target task IU is received, move it to a command buffer, but don't tell the IOP and upper layers about it until they know about the initiator.
TO REPRODUCE: Run heavy I/O to a large SAS topology (100 drives) while simulating cable pulls in the topology by disabling and enabling randomly-selected SAS ports on the attached storage enclosures periodically.
ISSUE DESC: PL was notifying the IOP and upper layers about a task IU before they'd been notified about the initiator who sent it. This was causing an assertion to fail in IOP.

(SCGCQ00630674 - Port of SCGCQ00524975) Defect 44/45



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

HEADLINE: PL: Task management aborted IO search may return faulty information
DESC OF CHANGE: Always search the aborted IO list for each task management associated with the device even if there is a target reset outstanding to it.
TO REPRODUCE: Unit testing of code that uses the offending function and code inspection.
ISSUE DESC: PL task management code may mistakenly report that an IO is being aborted by a TM if there is a target reset outstanding to that device. This may cause other parts of the PL to conclude that the IO is already being handled by task management, which may cause various faults or other problems.

(SCGCQ00635733 - Port of SCGCQ00634879)

Defect 45/45

HEADLINE: IOP: PowerPC ITLB Miss Errata Workaround
DESC OF CHANGE: Changed the firmware and the linker to ensure no function crosses a TLB page boundary. This workaround also prevents firmware from running out of Flash, and firmware will fault if the Enable XMem Copy flag is not set in Manufacturing Page 11. Firmware also cannot run out of DDR, if it is available, though firmware will continue to utilize DDR for dynamic memory allocation.
TO REPRODUCE: This issue has not been reproduced using IT/IR firmware.
ISSUE DESC: One to four extra random instructions might be executed erroneously after an instruction-side translation lookaside buffer (ITLB) miss operation occurs.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

Total Enhancements Implemented (40)

(SCGCQ00557357) Enhancement 1/40

HEADLINE: WhirlWind : PL : Remove PCIE core (Trident) registers access from PL - Part 1

NEW FUNCTIONALITY: The following changes have been done with this ER
- Local configuration space read/write moved to IOP .
- PIM window configuration for PL moved to IOP

(SCGCQ00557883) Enhancement 2/40

HEADLINE: Whirlwind: Implement Configuration Space Request Redirection (CSR).

NEW FUNCTIONALITY: Implemented Configuration Space Request Redirection (CSR).

(SCGCQ00559130) Enhancement 3/40

HEADLINE: Whirlwind: (Master) One Massive ER

NEW FUNCTIONALITY: One Massive ER to pull in everything from Whirlwind Pre-alpha to Whirlwind Master project.

(SCGCQ00564420) Enhancement 4/40

HEADLINE: Convert direct PCIE POM Reads to using PCI mem read instead

NEW FUNCTIONALITY: POM reads are not feasible long term due to lack of error handling in the HW for this path. Converting to use a PCIE memory access instead, via an accessor function.

(SCGCQ00566403) Enhancement 5/40

HEADLINE: Added WW specific NVDATA and build

NEW FUNCTIONALITY: Added WW specific NVDATA and build

(SCGCQ00567780) Enhancement 6/40

HEADLINE: Add support for new IO Unit Control Request (in MPI 2.6)

NEW FUNCTIONALITY: With this ER , the obsolete 'SAS IO Cotrol unit request' (MPI2.5 specific) is replaced with IO Unit Control request (MPI2.6) .

(SCGCQ00567802) Enhancement 7/40

HEADLINE: Use Man Page 9 configured values for PCIe devices

NEW FUNCTIONALITY: Added the following to ManPage9:
1. A new resource - RESOURCE_BACKEND_PCIE_DEVICES
2. Bit 4 in Flag - BACKEND_NVME_DEVICE_SUPPORT

Used the ManPage9 configured values during SOD runtime config for PCIe devices.

(SCGCQ00567813) Enhancement 8/40

HEADLINE: Pull updates in scsi.h from SAS3.5 project

NEW FUNCTIONALITY: Ported the scsi.h file changes from SAS3.5 project.

(SCGCQ00568791) Enhancement 9/40

HEADLINE: WhirlWind : PL : Completely remove pci msix register usage from PL .



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

NEW FUNCTIONALITY: Pci Msix registers used in PL during the MSIX vector table initialization .With this ER teh vector table initialization moved to IOP , which help us to remove the MSIX register usage completely from PL.

(SCGCQ00569101) Enhancement 10/40

HEADLINE: Add MULTl script to flash WW build

NEW FUNCTIONALITY: Added MULTl script to flash WW build

(SCGCQ00573598) Enhancement 11/40

HEADLINE: Implement Read/Write NVM Commands:

NEW FUNCTIONALITY: Start IO side

-
1. SCSI Read/Write IO to NVM equivalent translation
 2. IEEE-64 SGL to NVM PRP List conversion
 3. NVM Command creation
 4. Posting to NVMe drive's submission queue

Completion IO path

-
1. Receive MSIx int for IO completion
 2. Read the completion entry from the Completion queue of the NVMe device
 3. Parse the entry for failures and convert thm to appropriate SCSI check conditions
 4. Complete the associated SCSI IO back to the host.

Limitations of this ER:

- =====
1. Break up los not supported
 2. Only supports upto 2 PRP entries.

(SCGCQ00580164) Enhancement 12/40

HEADLINE: Whirlwind : Remove the PCI Pepp reg (gplHwPtrPciePeppRegs) and FMU reg (gplHwPtrPcieFMURegs) usage from PL.

NEW FUNCTIONALITY: Removed the Pepp regs and FMU reg as part of this ER. Following functions have been removed to accomplish this.

- plPciRclsLinkUp()
- plPciRcSetPCleSpeed()
- plPciRcSetPcieWidth()
- plPciRcReadPdbAHB()
- plPciRcWritePdbAHB()

(SCGCQ00580338) Enhancement 13/40

HEADLINE: WW: Configure GEP BAR0 as 64-bit (Capella-2 FPGA only)

NEW FUNCTIONALITY: Implemented code to configure management end point's (GEP) BAR0 as a 64-bit BAR.

(SCGCQ00580340) Enhancement 14/40

HEADLINE: WW: (Capella-2) CSR buffer allocation.

NEW FUNCTIONALITY: A buffer aligned at 4KB boundary is allocated at start-of-day by IOP for use by Capella-2's CSR.

(SCGCQ00583839) Enhancement 15/40

HEADLINE: Whirlwind: Further IO path tweaks.

NEW FUNCTIONALITY: Further IO path tweaks.Extended support for larger IOs and break up los. Also, modified completion path to scan the IO Completion queue for all devices for every MSIx interrupt,



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00592594) Enhancement 16/40

HEADLINE: WW: (Capella-2) Establish communication between host and Invader root complex.

NEW FUNCTIONALITY: Implemented PCIe TLP routing between host and Invader root complex.

(SCGCQ00593490) Enhancement 17/40

HEADLINE: Whirlwind: Add support to dump all PL root complex related debug info in one CLI command

NEW FUNCTIONALITY: (1) Add support to dump all PL Root Complex related debug info in one CLI command
(2) Add support to dump relevant PL Root Complex related debug info as part of "pl dbg" CLI command

(SCGCQ00595519) Enhancement 18/40

HEADLINE: Selective Extended Image Loading

NEW FUNCTIONALITY: This firmware changes how the bootloader loads extended images into on chip memory. Instead of blindly copying the entire firmware image with all of its attached extended images, this code will now only copy the extended images which are needed by firmware to run. This should reduce the memory usage by the firmware.

(SCGCQ00601719) Enhancement 19/40

HEADLINE: WW: Switch EEPROM programming

NEW FUNCTIONALITY: Following major functionalities are implemented with this ER
- PLX Switch SEEPROM detection and validation
- PLX SEEPROM SEEPROM Read and Write
- Programming SEEPROM image
- Reading SEEPROM image

(SCGCQ00605758) Enhancement 20/40

HEADLINE: Whirlwind: NVMe Task Management Target Reset Support

NEW FUNCTIONALITY: NVMe Task Management Target Reset Support

(SCGCQ00606386) Enhancement 21/40

HEADLINE: WW: (Capella-2) Establish communication to/from NVMe devices.

NEW FUNCTIONALITY: Implemented PCIe TLP routing between NVMe and Invader.
Implemented PCIe TLP routing between NVMe and Host.

(SCGCQ00610550) Enhancement 22/40

HEADLINE: Add NVMe encapsulated IO support

NEW FUNCTIONALITY: Added the MPI NVMe Encapsulated message support. Now, the Admin/IO NVMe commands can be issued to the NVMe device directly (as passthrough).

(SCGCQ00615216) Enhancement 23/40

HEADLINE: Whirlwind: NVMe target reset (Part II)

NEW FUNCTIONALITY: NVMe target reset (Part II) - IOCount management during TM Target Reset

(SCGCQ00618535) Enhancement 24/40

HEADLINE: Whirlwind: NVMe Task Abort Implementation

NEW FUNCTIONALITY: NVMe Task Abort Implementation



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00619353) Enhancement 25/40

HEADLINE: Add support for Thunderbolt based NVMe emulator
NEW FUNCTIONALITY: Added support to discover and configure Thunderbolt based controller/board as NVMe emulator.

(SCGCQ00619545) Enhancement 26/40

HEADLINE: IOP MCTP: Add basic MCTP Control Commands
NEW FUNCTIONALITY: Adds Basic MCTP Control Commands such as Set and Get EID

The following MCTP Control Commands are included, but unsupported. Support will be available in future releases.
- Endpoint Discovery
- Discovery Notify

(SCGCQ00620222) Enhancement 27/40

HEADLINE: IOP MCTPMsg: Add MPI RAID Action and RAID SCSI IO Passthrough
NEW FUNCTIONALITY: Added support for the following IR specific MPI commands:
- MPI RAID Action
- MPI RAID SCSI IO Passthrough

(SCGCQ00620227) Enhancement 28/40

HEADLINE: IOP MCTP: MPI Event registration, Extended Events command, and transmitting event datagrams
NEW FUNCTIONALITY: Added support to allow the SP to register for MPI Events. This instructs the FW that it should send Event datagrams to the SP when registered events occur.

Added support for the Extended Events command in the SCS MCTP Spec. This contains the RAID events, which are intended to be obtained by the SP after sending an asynchronous Event datagram to the SP.

Added new event types, including some RAID events and Log Entry, were added to the standard event list.

(SCGCQ00623300) Enhancement 29/40

HEADLINE: Whirlwind: Root Complex MSIx setup
NEW FUNCTIONALITY: Added code to enable MSIx in root complex mode. In a non-root complex mode the MSIx enable bit in the MSI Control register is setup by the host software. For whirlwind, this needs to be enabled by the IOP code after the MSIx vector table is programmed.

(SCGCQ00630611) Enhancement 30/40

HEADLINE: Whirlwind: Improve NVME Initialization Error Handling
NEW FUNCTIONALITY: In NVME initialization, changed bad reg read fault cases to fail the init instead. That will lead to a device reset being done.

(SCGCQ00631600) Enhancement 31/40

HEADLINE: IOP MCTP: Add support for resetting EID in MCTP Control Set EID command
NEW FUNCTIONALITY: Added support to MCTP Control's Set EID command for resetting the EID to the originally configured static (preconfigured) EID.

(SCGCQ00634740) Enhancement 32/40

HEADLINE: WW: added BDF 0:5:0 to requestor ID table
NEW FUNCTIONALITY: added BDF 0:5:0 to requestor ID table



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

(SCGCQ00575849 - Port of SCGCQ00526831)

Enhancement 33/40

HEADLINE: Add in support for CM Build app
NEW FUNCTIONALITY: added command file to make IT and PL for Whirlwind

(SCGCQ00606270 - Port of SCGCQ00563906)

Enhancement 34/40

HEADLINE: NVMe: Extended Inquiry Data VPD page
NEW FUNCTIONALITY: Added Extended Inquiry Data VPD page (page code 0x86) to the Inquiry translation. Translation is based on that found in the NVMe-SCSI Translation Reference 1.1.

(SCGCQ00615113 - Port of SCGCQ00571441)

Enhancement 35/40

HEADLINE: NVMe: Support for Mode Sense (10) LLBAA flag
NEW FUNCTIONALITY: Generate and return a long LBA mode parameter block descriptor if the LLBAA bit is set in a Mode Sense (10) CDB.

(SCGCQ00615114 - Port of SCGCQ00572583)

Enhancement 36/40

HEADLINE: NVMe: a couple of mode page tweaks
NEW FUNCTIONALITY: Save off an NVMe device's initial write cache state and use it for the value of the WCE bit when the Caching mode page is read with default mode parameters.

Update the Rigid Disk Geometry page with a value reflecting an NVMe device's capacity while building the page.

(SCGCQ00625579 - Port of SCGCQ00228558)

Enhancement 37/40

HEADLINE: PL: New commands useful for debugging
NEW FUNCTIONALITY: New UART commands have been added to the controller UART interface for debugging problems in the field.

(SCGCQ00625580 - Port of SCGCQ00465302)

Enhancement 38/40

HEADLINE: SATL: Indicate support for the ATA Device Server Password security protocol
NEW FUNCTIONALITY: When the Supported Security Protocol List is requested from a SATA drive using a SECURITY PROTOCOL IN command with Security Protocol value of 0 and Security Protocol Specific value of 0:

If the device supports the ATA Trusted Computing Group feature set and the ATA Security feature set, append security protocol code 0xEF to the list returned by the device.

If the device does not support the TCG feature set but does support the Security feature set, fabricate a list containing security protocol codes 0 and 0xEF.

If the device supports neither the TCG feature set nor the Security feature set, fabricate a list containing security protocol code 0 only.

(SCGCQ00625583 - Port of SCGCQ00557625)

Enhancement 39/40

HEADLINE: Print SES page data length and PL SES Diag page buffer size to ring buffer
NEW FUNCTIONALITY: This CSET adds ring buffer debug print to dump the SES req length, current SES diag buffer size, OpCode/PageCode and SkipCount while building a SES request.

(SCGCQ00625622 - Port of SCGCQ00616246)

Enhancement 40/40

HEADLINE: Correct PHY ordering in Channel NVData SGPIO configuration
NEW FUNCTIONALITY: Corrected the SGPIO PHY ordering for the internally connected channel boards.



SCS Engineering Release Notice

Phase5 Pre-Alpha Release Version 04.250.01.00 - SAS3FW_MASTER_DEV (SCGCQ00638883)

Total NVDATA Changes (1)

(SCGCQ00624122 - Port of SCGCQ00621446)	
HEADLINE:	Gen3 Channel NVDATA: Tighten Outstanding IO range
NEW FUNCTIONALITY:	Tightened the range for the number of Outstanding IOs in Manufacturing Page 9 to a range of between 8,000 and 10,000.

NVDATA 1/1