

**ReleaseOrder ID:** SCGCQ01223280  
**Headline:** GCA Release: SAS3FW\_Phase14.0 - 14.00.00.00 IT IR  
**Release Version:** 14.00.00.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_Phase14.0\_Rel  
**Release Type:** GCA  
**State:** Test\_Complete  
**Release Baseline:** SAS3FW\_Phase14.0-2016-11-21-14.00.00.00\_REL\_1479719271@  
 \SAS\_CTRL\_FW  
**Release Date:** 23-NOV-16  
**Date Generated:** Dec 20, 2016

## Release History

- [SCGCQ01220935 - Beta Release: SAS3FW\\_Phase14.0 - 13.250.05.00 Fir](#)
- [SCGCQ01206306 - Beta Release: SAS3FW\\_MASTER\\_DEV - 13.250.04.00 Fi](#)
- [SCGCQ01191205 - Alpha Release: SAS3FW\\_MASTER\\_DEV - 13.250.03.00 F](#)
- [SCGCQ01181673 - Pre-Alpha Release: SAS3FW\\_MASTER\\_DEV - 13.250.02.](#)
- [SCGCQ01169097 - Pre-Alpha Release: SAS3FW\\_MASTER\\_DEV - 13.250.01.](#)

**ReleaseOrder ID:** [SCGCQ01220935](#) Open In CQWeb  
**Headline:** Beta Release: SAS3FW\_Phase14.0 - 13.250.05.00 Fir  
**Release Version:** 13.250.05.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_Phase14.0\_Rel  
**Release Type:** Beta  
**State:** Test\_Complete  
**Release Baseline:** SAS3FW\_Phase14.0-2016-11-16-13.250.05.00\_REL\_1479295461@  
 \SAS\_CTRL\_FW  
**Release Date:** 18-NOV-16  
**Date Generated:** Dec 20, 2016

### Defects Fixed (2):

**ID:** SCGCQ01204905 (Port Of Defect SCGCQ01155407)

**Headline:** PL fault 5854 seen when system boots with all SATA drives in setup spun down

**Description Of Change:** Reverted first command for SATA Initialization that was prior to SMR support, where we sent IDENTIFY as the first command, instead of read Block Device Characteristics log page. Rearranged the SATA Initialization sequence to read the Block Device Characteristics log page after the check power mode to avoid checking any initial command failures during SATA initialization.

**Issue Description:** As part of SMR drive support we had added a new command to read the Block Device Characteristics log page at the start of the SATA Initialization process, replacing the IDENTIFY as the 1st command. We were also not checking for all the errors and command failures of the block device characteristics log page since this was supposed to fail for NON SMR devices. But we also happened to not check for errors related to command not transferred, timeout, drive errors and etc. This lead to further commands sent to drive, though errors were logged in the hardware leading to faults.

**Steps To Reproduce:** 1. Large configof SATA HDDs. OS on VD on the controller  
 2. Mark all the PDs Prepare for Removal except OS VD and make sure they are in spundown mode  
 3. Reboot the server and see if the OS boot fails

**ID:** SCGCQ01205282 (Port Of Defect SCGCQ01172334)

**Headline:** PL SATA Only: outstanding SATA passthrough IO can cause target reset from timer callback check if sata Initialization not done yet

**Description Of Change:** In timer call back, do not initiate TM immediately if IO is SATA Pass Through command. Instead initiate TM if the same Passthrough IO still exists even after 4s. Delay initiating TM to make sure IO is stuck.

**Issue Description:** A SATA Passthrough command issued to a SATA drive requiring initialization results in the controller resetting the SATA drive and failing the passthrough command

**Steps To Reproduce:** Only issue SATA passthrough IO to SATA drive without any TUR/Report LUN etc.

**ReleaseOrder ID:** [SCGCQ01206306](#) Open In CQWeb  
**Headline:** Beta Release: SAS3FW\_MASTER\_DEV - 13.250.04.00 Fi  
**Release Version:** 13.250.04.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_MASTER\_Invdr\_Rel  
**Release Type:** Beta  
**State:** Superseded  
**Release Baseline:** SAS3FW\_MASTER\_DEV-2016-10-20-13.250.04.00\_REL\_1476955420@  
 \SAS\_CTRL\_FW  
**Release Date:** 21-OCT-16  
**Date Generated:** Dec 20, 2016

### Defects Fixed (3):

**ID:** SCGCQ01185613

**Headline:** PL: Fault 0x6003 hits if blocked device has IOs pended before it gets blocked

**Description Of Change:** During starting of pended IOs, if the IO could not be started send failure response to host.

**Issue Description:** During starting of pended IOs, if device is blocked reply is not sent to host but IO is removed from the pend queue. Due to this abort is received due to no response. During TM, 6003 fault is hit since the IO was not found as it was already removed from pend queue.

**Steps To Reproduce:** 1. Boot to OS with the Invader & connect drives.  
 2. Disable Fast Path on 1/2 the SAS drives and 1/2 the SATA devices in the topology.  
 3. Start IOs using Smash tool.  
 4. Start the script to send inquiry cmd to all connected drives.  
 5. Run a script which will do half phy glitch every 10 secs using Cable Breaker  
 Result: Fault 0x6003 seen.

**ID:** SCGCQ01204880 (Port Of Defect SCGCQ01193764)

**Headline:** PL: Fault 5814 while multipath SAS pull/push of 2 enclosures

**Description Of Change:** Updated firmware to avoid MID validation while handling the Tx Frame Manager send timeout interrupt.

**Issue Description:** The FW might fault with code 0x5814 during while multi-path SAS pull/push 2 enclosures on 3516 chip based HBA. Cable pull/push causes internal PL TM and there seems to be race condition when TM cleans Tx Frame manager block and it send the frame in scope of the TM. The Tx Frame Manager buffer could indicate MID in scope of TM, its expected and FSM timeout handler shouldn't validate the MID.

**Steps To Reproduce:** Daisy chain multipath IO stress with cable pull/push on 3516 based HBA, and observe the PL firmware fault 5814 during the course.

ID: SCGCQ01205297 (Port Of Defect SCGCQ01184674)

**Headline:** Ventura: PL fault 5811

**Description Of Change:** Modified FW to not to validate the MID, the MID may be invalid as the HW is working on different frame.

**Issue Description:** The FW might fault with fault x5811 when rebooting server multiple time with 3516 chip based HBA board connected to SAS/SATA expander enclosures. The HW while sending frames on the wire hits timeout due to unstable connection but eventually sends the frame successfully. While handling the interrupt, the FW checks for MID that caused them timeout but it could be invalid as HW has moved on to process the next MID.

**Steps To Reproduce:** Perform multiple reboots on server with 3516 chip based HBA board connected to SAS/SATA expander enclosures and observe the x5811 fault.

---

**ReleaseOrder ID:** [SCGCQ01191205](#) Open In CQWeb  
**Headline:** **Alpha Release: SAS3FW\_MASTER\_DEV - 13.250.03.00 F**  
**Release Version:** 13.250.03.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_MASTER\_Invdr\_Rel  
**Release Type:** Alpha  
**State:** Superseded  
**Release Baseline:** SAS3FW\_MASTER\_DEV-2016-09-28-13.250.03.00\_REL\_1475053785@  
ISAS\_CTRL\_FW  
**Release Date:** 28-SEP-16  
**Date Generated:** Dec 20, 2016

---

### Defects Fixed (3):

ID: SCGCQ01180930 (Port Of Defect SCGCQ01171835)

**Headline:** Storlib/storlibtest stuck on a bad drive while retrieving phyinfo

**Description Of Change:** IO received will be terminated on receiving the Sense Key/ASC/ASCQ value of 0x00044400 as it is non-recoverable hardware failure from bad drive.

**Issue Description:** Using storelib/storelibtest tool with a bad drive was taking time to display drive related information and was stuck for almost 15 min.

IR FW after discovering the bad drive, it received SCSI commands like SCSI\_COMMAND\_TEST\_UNIT\_READY (0x00), SCSI\_COMMAND\_READ\_CAPACITY\_16 (0x9E) and READ\_CAPACITY(10) (0x25). IR FW keeps count of every IO which comes in and this count was exceeding the max limit of 45 (SPIN\_UP\_TIMEOUT) and later failing the IO's and completing the same which would restart the counter.

Counter was incremented on IR\_IO\_STATE\_BECOMING\_READY only.

On receiving SCSI\_COMMAND\_TEST\_UNIT\_READY command, Key/ASC/ASCQ read from the bad drive were:

- 0x00020401 (02 - Not Ready, 0x0401 - SCSI\_ASC\_LOGICAL\_UNIT\_IN\_PROCESS\_BECOMING\_READY)
- 0x00020402 (02 - Not Ready, 0x0402 - SCSI\_ASC\_LOGICAL\_UNIT\_NOT\_READY\_INIT\_REQUIRED) and starts the command SCSI\_COMMAND\_START\_STOP\_UNIT (0x1B).
- 0x00020411 (02 - Not Ready, 0x0411 - SCSI\_ASC\_LOGICALUNIT\_NOT\_READY\_SPINUP\_REQUIRED) and starts the command SCSI\_COMMAND\_START\_STOP\_UNIT (0x1B).

In all the above steps, count was getting incremented for step 1 alone in the process of getting information from drive with Sense error code of SCSI\_SENSE\_DATA\_RESPONSE\_CODE\_CURRENT (0x70)

0x00044400 (04 - HARDWARE ERROR, 0x4400 - SCSI\_ASC\_INTERNAL\_TARGET\_FAILURE), it was generating Sense error code of SCSI\_SENSE\_DATA\_RESPONSE\_CODE\_DEFERRED (0x71) and IO state was set to deferred for retrying the same.

On receiving SCSI\_COMMAND\_READ\_CAPACITY\_16 (0x9E) and READ\_CAPACITY(10) (0x25) commands also, we were seeing the above described Key/ASC/ASCQ values.

All these above things was also leading to IO queue full.

**Steps To Reproduce:** 1. At Linux machine Desktop, run ./storelibtest  
2. Option 4  
3. Option 1  
4. Display Intel SSD (Good drive) information (Normal)  
5. Option 0 (meaning get phyinfo from the bad drive)  
6. Display Stuck  
7. After about 10 - 15 min, will display Hitachi SAS HDD information (bad drive)

ID: SCGCQ01183260 (Port Of Defect SCGCQ01164403)

**Headline:** Latest PL firmware(13.0.0.0) causing MR firmware to crash.

**Description Of Change:** While sending SGPIO response data, last parameter i.e flags is properly sent which invokes MR iopi memory move function.

**Issue Description:** MR Firmware crashes after connecting enclosure device.

**Steps To Reproduce:** Connect enclosure to MR card with PL Firmware version 13.0.0.0.

ID: SCGCQ01190643 (Port Of Defect SCGCQ01176300)

**Headline:** Config write request validation adjusted to page header size as the minimum request length

**Description Of Change:** The request length validation calculation has been corrected with the assumption that minimum config write request length could be as less as page header length as opposed to full page size earlier.

**Issue Description:** The config write request was being validated with the assumption that the minimum request length should be as big as the page length. This being not true, this validation is now reduced to the header length which is the true minimum request length possible for config write.

**Steps To Reproduce:** #1 Generate some event, such are restart controller from SLIB  
#2 Register AEN [Option main menu 2 -> 2 -> 9 -> 3 -> event sequence number]  
#3 Try clear event command from Events menu [Option main menu 9 -> 3]

---

**ReleaseOrder ID:** [SCGCQ01181673](#) Open In CQWeb  
**Headline:** **Pre-Alpha Release: SAS3FW\_MASTER\_DEV - 13.250.02.**  
**Release Version:** 13.250.02.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_MASTER\_Invdr\_Rel  
**Release Type:** Pre-Alpha  
**State:** Test Complete  
**Release Baseline:** SAS3FW\_MASTER\_DEV-2016-09-14-13.250.02.00\_REL\_1473839015@  
ISAS\_CTRL\_FW  
**Release Date:** 14-SEP-16  
**Date Generated:** Dec 20, 2016

---

### Defects Fixed (7):

ID: SCGCQ01167407

**Headline:** IOP: MCTP: Outstanding aborted SCSI IO due to too many PE Busy Retries is not being fully aborted within FW

**Description Of Change:** Corrected the PE BUSY RETRY response code for the right PEC.

Added a check when firmware receives an asynchronous response to see if the IO has already been aborted.

**Issue Description:** In MCTP I2C with slave response mode (controller is exclusively an I2C slave), if a SCSI IO is still outstanding to a drive, and it is aborted by a BMC, then it is not actually being fully aborted by the firmware. So the response from the drive can be forwarded back to the BMC.

This was initially caused by firmware messing up the PEC for a specific PE BUSY RETRY (IO was still outstanding) response path, but can still happen generically.

**Steps To Reproduce:** In MCTP I2C slave response mode:  
1. Send a SCSI IO to a drive over MCTP

2. Perform an I2C read and get a bad PEC
3. Abort the SCSI IO
4. Perform I2C reads after the IO was completed, and the BMC will get the response of the aborted IO

---

**ID:** SCGCQ01173592

**Headline:** SATL Only: Write Same SCSI Command Fails With CC for SATA Drive

**Description Of Change:** Corrected the data length check that was added to abort the command if data length is less than SATA sector size.

**Issue Description:** Write Same SCSI Command Fails With Check condition abort for SATA Drive

**Steps To Reproduce:** Issue write same command to SATA drive

---

**ID:** SCGCQ01176386

**Headline:** SATA Only: WRITE AND VERIFY Command Not Failed With CC for LOGICAL BLOCK ADDRESS OUT OF RANGE Condition

**Description Of Change:** The check for successful completion if transfer length zero is moved after the check for LBA out of range.

**Issue Description:** The command passes instead of failing with results of CC (LBA OUT OF RANGE) when LBA beyond MAX LBA for that device is given.

**Steps To Reproduce:** Execute a WRITE AND VERIFY (10) in which LBA address should be beyond the last LBA address and TRANSFER LENGTH set to zero.  
The command passes instead of failing with results of CC (LBA OUT OF RANGE).

---

**ID:** SCGCQ01176401

**Headline:** SATA Only: VERIFY Command Not Failed With CC for LOGICAL BLOCK ADDRESS OUT OF RANGE Condition

**Description Of Change:** The check for successful completion if transfer length zero is moved after the check for LBA out of range.

**Issue Description:** The VERIFY command passes instead of failing with results of CC (LBA OUT OF RANGE) when LBA beyond MAX LBA for that device is given and VERIFICATION LENGTH set to zero.

**Steps To Reproduce:** execute a VERIFY (10) in which LBA address should be beyond the last LBA address and VERIFICATION LENGTH set to zero.  
The command passes instead of failing with results of CC (LBA OUT OF RANGE).

---

**ID:** SCGCQ01125226 (Port Of Defect SCGCQ01125105)

**Headline:** SAS3IRCU not responding after long term of Virtual Disk creation/deletion test

**Description Of Change:** IR firmware was allocating critical frames without de-allocating resource frames. Due to this, message frame allocated by resource queue was not cleared and which exhausted all the message frames and was not able to allocate any more message frames for an IO.  
Due to this reason, IO's received were stalled for ever causing the driver to timeout.

**Issue Description:** Test script (build RAID.sh) will repetitively create, delete and displays the Virtual Disks behind and SAS9311-8i board which directly connects 4 Hard disks.

In each round:

- first use SAS3IRCU create command to create a RAID0 Volume using two physical hard disks.
- then use SAS3IRCU deletevolume command to delete the RAID0 Volume.
- then use SAS3IRCU create command to create a RAID1 Volume using the other two physical hard disks.
- then use SAS3IRCU deletevolume command to delete the RAID1 Volume.
- then use SAS3IRCU display command to show info.

The SAS3IRCU will go not responding after about dozens of rounds of test.  
The SAS3IRCU display command will end up with 'no controller found at index 0  
The SAS3IRCU list command will report 'SAS3IRCU: MPTLIB2 ERROR 52.

**Steps To Reproduce:** Run test script (build RAID.sh) manually which creates, deleted and displays the volume which fails after 39 creates/deletes with SAS3IRCU command timing out and driver starting the task management.

---

**ID:** SCGCQ01163789 (Port Of Defect SCGCQ01084723)

**Headline:** OS failed to boot after hot-plug on the RAID1 member.

**Description Of Change:** Changed the offline condition check for loading current state value of the physical drive from Optimal to Degraded state.

**Issue Description:** RAID1 was created using two drives (A and B) and OS was installed on it using Bios application. Drive B was removed during powered off state and Volume became degraded and OS booted successfully. Power down OS and insert drive B back.  
When powered on, Volume was in optimal state which was wrong because the drive removed has come back.

**Steps To Reproduce:**

1. Create RAID1 using 2 drives(A and B) and install OS(Redhat 7.0 with XFS file-system).
2. Pull out Disk B during power off and when powered on, the RAID is degraded and OS can boot successfully.
3. Power off the OS, and insert the Disk B back.
4. When powered on there is no rebuilding status and RAID1 volume is Optimal and OS failed to boot because of file-system error.

---

**ID:** SCGCQ01163793 (Port Of Defect SCGCQ01096884)

**Headline:** Fault 0x265D encountered on inserting a faulty drive in place of R1 member

**Description Of Change:** When IO was deferred due to bad drive added in place of R1 member, added a Null pointer check since the pointer was already cleared during journal entries and returning from the function without processing.

**Issue Description:** Test scenario is replacing a good drive with a faulty drive during powered off state and hence the rebuild entries were cleared from the drive and this is the IR FW behaviour.  
IR FW does resync of IO periodically to inform the Host and for some reason if Write cache was set to false on the drive.  
IR FW clears the journal which resets the write cache flag. Now when the IO was deferred due to bad drive, volume pointer was cleared and hence it was dangling pointer.

**Steps To Reproduce:** Using 2 SAS drive creating a RAID1. Then power off, insert back a fault drive to replace one of the RAID member.  
And after power on, the fault 0x265D would be encountered after about 15 minutes.

---

## Enhancements Implemented (2):

**ID:** SCGCQ00999243

**Headline:** Updated the NVDATA by aligning the fields of SAS IO UNIT PAGE 4 with the MPI

**Description Of Change:** Updated the NVDATA by aligning the fields of SAS IO UNIT PAGE 4 with respect to the MPI, where some fields were earlier marked as reserved in NVDATA compared to the valid field names in MPI.

---

**ID:** SCGCQ01085063

**Headline:** PL: Support for all the Internal connector types to end devices defined in the SES3 specification revision 13d and the proposed new connectors for SAS4

**Description Of Change:** Added discovery support for all the Internal connector types to end devices defined in the SES3 specification revision 13d, from 0x20 through 0x2F and the proposed new connectors from 0x40 through 0x40F for SAS4.

---

---

**ReleaseOrder ID:** [SCGCQ01169097](#) [Open In CQWeb](#)  
**Headline:** *Pre-Alpha Release: SAS3FW\_MASTER\_DEV - 13.250.01.*  
**Release Version:** 13.250.01.00  
**UCM Project:** SAS3FW\_MASTER\_DEV  
**Sub UCM Project:** SAS3FW\_Phase14.0  
**UCM Stream:** SAS3FW\_MASTER\_Invdr\_Rel  
**Release Type:** Pre-Alpha  
**State:** Superseded  
**Release Baseline:** SAS3FW\_MASTER\_DEV-2016-08-26-13.250.01.00\_REL\_1472218771@  
SAS\_CTRL\_FW  
**Release Date:** 29-AUG-16  
**Date Generated:** Dec 20, 2016

---

## Defects Fixed (4):

ID: SCGCQ01123275 (Port Of Defect SCGCQ01118894)

**Headline:** PL: NetApp Wembley expansion enclosures are seeing large number of drive timeout

**Description Of Change:** Implemented a hardware workaround to stop and reset the Tx frame manager so that things are un-stalled.

**Issue Description:** This issue presents as a lock up in the SAS core, where the IOC responds to all open requests with open reject, retry. It's caused by a hardware issue in the Tx frame manager, where it sometimes gets stalled after sending out a frame. Sometimes it recovers and sends the next frame and there's no issue; sometimes it doesn't and everything gets stalled.

**Steps To Reproduce:** Not totally clear. Only certain systems have reproduced it. The customer setup involves two IOCs doing IO to each other (initiator and target) and IO to expander attached drives. There's some suspicion that this may be enabled by drive behavior (if the drives send a lot of opens all at once, at just the right time).

---

ID: SCGCQ01138058 (Port Of Defect SCGCQ01134948)

**Headline:** PL: MidUsed not cleared when handling FPE Requestor Error

**Description Of Change:** Added code to clear the bit in this handling case to ensure the internal state is consistent.

**Issue Description:** An internal status bit could be set but not cleared for a MID during specific error handling relating to FPE timers. This causes the internal status to get out of sync with reality and could cause issues during future operations, including handling FPE timeouts for that MID.

**Steps To Reproduce:** Issue was found during code inspection, no reproduction is known. The code path requires FPE timers to be used.

---

ID: SCGCQ01145612 (Port Of Defect SCGCQ01143329)

**Headline:** PL: 0x5813 and 0x5814 Faults after implementing a hardware workaround

**Description Of Change:** Made a couple changes in SAS core cleanup code and the Tx frame manager interrupt handler to prevent these faults from occurring.

**Issue Description:** After implementing a hardware workaround for an issue with the Tx frame manager, 5813 and 5814 faults were later seen in customer testing.

**Steps To Reproduce:** Perform IO stress testing on customer system.

---

ID: SCGCQ01162833 (Port Of Defect SCGCQ01152796)

**Headline:** Added new PL fault code and also set default 60 secs for SATA hinting timer in PL

**Description Of Change:** - Added new PL fault code for SATA hinting timer creation failure, instead of generic SATA Hinting Failed fault code.  
- Set default 60 secs value in PL for SATA hinting timer if the NVDATA does not have set any value.

**Issue Description:** - To differentiate the 2 faults within the SATA Hinting timer logic, there was need to create a new Fault code.  
- If certain IOPs fail to set NVDATA value for SATA Hinting timer, then discovery of SATA devices could misbehave as the value for timer will become 0. As a precaution, PL would need to set default 60 secs, which will be overridden by the value set in the NVDATA. IOPs' need to set a non zero value in the NVDATA.

**Steps To Reproduce:** NA

---

## Enhancements Implemented (1):

ID: SCGCQ01163873

**Headline:** PL:SCT Write Same support on native 4K drives

**Description Of Change:** For native 4K drive, the SCT data transfer log should transfer 4K data(1 sector size).  
As part of SCT data transfer log (log address 0xE1), 8 data frames of 512 byte(4K) is transferred to native 4K drive.

---