Intel® Communications Chipset 8900 Series

Specification Update

June 2016

Document Number: 328000-006US

Contents

Revision History

Date	Revision	Description
June 2016	006	• Added component to Component Identification table. • Added Documentation Changes section.
May 2014	005	• Added Erratum 83.
December 2013	004	• Removed erratum 72. • Added DH89xxCL, A0 stepping. • Changed erratum 79. • Changed Specification Change 1. • Changed Specification Clarifications 1, 2, 3, 4. • Removed documentation changes.
July 2013	003	• Added errata 81-82.
March 2013	002	• Added Documentation Changes 1 - 2. • Added errata 79 - 80. • Added Specification Updates #1. • Added Specification Clarification #4. • Changed DID for B0:D28:fn when subtractive decode is enabled to 0x244E. • Added Note to Errata #72 "Intel® QuickAssist Technology (QAT): Decompression Adler32 Checksum May Be Calculated Incorrectly" to indicate that the recommended "Workaround" is implemented in software release 1.1 and later versions.
November 2012	001	• Initial Release

Preface

This document is an update to the specifications contained in the Affected Documents/ Related Documents table below. This document is a compilation of device and documentation errata, specification clarifications and changes. It is intended for hardware system manufacturers and software developers of applications, operating systems, or tools.

Information types defined in Nomenclature are consolidated into the specification update and are no longer published in other documents.

This document may also contain information that was not previously published.

In this document:

- Intel® Communications Chipset 89xx Series:
  - DH89xxCC includes Chipset 8900 Series SKUs where 8900 ≤ SKU ≤ 8920
  - DH89xxCL includes Chipset 8900 Series SKUs where 8925 ≤ SKU ≤ 8955

Affected Documents/Related Documents

Document Title	Document Number/ Location
Intel® Communications Chipset 89xx Series - Datasheet	327879-00x

Nomenclature

Errata are design defects or errors. Errata may cause the behavior of the PCH to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present in all devices.

Specification Changes are modifications to the current published specifications. These changes will be incorporated in any new release of the specification.

Specification Clarifications describe a specification in greater detail or further highlight a specification's impact to a complex design situation. These clarifications will be incorporated in any new release of the specification.

Documentation Changes include typos, errors, or omissions from the current published specifications. These will be incorporated in any new release of the specification.

Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the Intel® Communications Chipset 89xx Series product. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables uses the following notations.

Codes Used in Summary Tables

Stepping

X:	Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.
(No mark) or (Blank box):	This erratum is Fixed in listed stepping or specification change does not apply to listed stepping.

Page

(Page):	Page location of item in this document.

Status

Doc:	Document change or update will be implemented.
Future Fix:	This erratum may be Fixed in a future stepping of the product.
Fixed.	This erratum has been previously Fixed.
No Fix:	There are no plans to fix this erratum.

Row

Change bar to left of table row indicates this erratum is either new or modified from the previous version of the document.

Table 1. Errata Summary Table (Sheet 1 of 3)

Erratum Number	DH98xxCC	DH89xxCL	Status	ERRATA
	C1	A0		
1	X	X	No Fix	DMI Port: DMI Degradation Mode not Functional if Lane Reversal is implemented
2	X	X	No Fix	PCIe* Root Port (RP): PCIe RP Degradation Mode not Functional if Lane Reversal is implemented
3	X	X	No Fix	EndPoint (EP): PCIe EP does not check the PCIe TLP Digest (TD) bit
4	X		No Fix	GbE: GbE SMGII Protocol Layer Delays affects System Collision Detection
5	X		No Fix	GbE: GbE MAC may drop the first Receive (RX) packet after a Link Speed Change
6	X	X	No Fix	PCIe Root Port (RP): PCIe RP Degradation is not Functional
7	X	X	No Fix	SATA: SATA Low Power Device Detection
8	X	X	No Fix	USB: USB Port Stall with Bulk and Control Traffic
9	X	X	No Fix	SATA: SATA SYNC Escape Issue
10	X	X	No Fix	USB: USB End of Frame When Retrying Packets Issue
11	X	X	No Fix	USB: USB Classic Device Removal Issue
12	X	X	No Fix	USB: USB RMH Descriptor May Report Incorrect Number of USB Ports
13	X	X	No Fix	PCH: High Precision Event Timer (HPET) Writing Timing Issue
14	X	X	No Fix	USB: USB Full-Speed Port Staggering
15	X	X	No Fix	USB: USB Devices May Slow or Hang
16	X	X	No Fix	USB: USB Low-Speed Control Transactions
17	X	X	No Fix	SATA: SATA Controller May Not Detect Unsolicited SATA COMINITs
18	X	X	No Fix	SATA: SATA Hot Unplug May Not be Detected
19	X	X	No Fix	USB: USB Missing ACK
20	X	X	No Fix	SATA: SATA 6 Gb/s Device Detection
21	X	X	No Fix	PCIe Root Port (RP): PCIe RP Link Disable Bit
22	X	X	No Fix	USB: USB Isochronous In Transfer Error Issue
23	X	X	No Fix	USB: USB Full-Speed/Low-Speed Device Removal Issue
24	X	X	No Fix	USB: USB Babble Detected with Software Overscheduling
25	X	X	No Fix	USB: USB Low-Speed/Full-Speed EOP Issue
26	X	X	No Fix	USB: USB PLL Control FSM not Getting Reset on Global Reset
27	X	X	No Fix	USB: USB Asynchronous Retries Prioritized Over Periodic Transfers
28	X	X	No Fix	USB: Incorrect Data for Low-Speed or Full-Speed USB Periodic IN Transaction
29	X	X	No Fix	USB: USB RMH Delayed Periodic Traffic Timeout Issue
30	X		No Fix	GbE: GbE MNG Reset Clears Resource Grant With No Feedback
31	X		No Fix	GbE MDIO: Com_MDIO and Destination Bits of MDICNFG Register Are Not Loaded Consistently from EEPROM
32	X		No Fix	GbE MNG: I2C Data Out Hold Time Violation
33	X		No Fix	GbE EEPROM: EE_CS_N Control Signal Hold Time Violation
34	X		No Fix	GbE SGMII: Counters Incorrectly Increment on Collision
35	X		No Fix	GbE TSYNC: Auxiliary Timestamp from SDP is Unreliable

Table 1. Errata Summary Table (Sheet 2 of 3)

Erratum Number	DH98xxCC	DH89xxCL	Status	ERRATA
	C1	A0		
36	X	X	No Fix	EndPoint (EP): Spurious MSI/MSIX Interrupt Generated
37	X	X	No Fix	EndPoint (EP): PCIe EP May not Detect Unexpected Completion Packets
38	X		No Fix	GbE: GbE Near End Analog Loopback Not Supported
39	X	X	No Fix	PCIe Root Port (RP): PCIe RP May not Automatically Switch into Compliance Mode
40	X	X	No Fix	EndPoint (EP): PCIe EP Incorrect Completion ID
41	X	X	No Fix	EndPoint (EP): PCIe EP Link Status Register 2 (LNKSTS2) Not Implemented for Functions[1-4]
42	X	X	No Fix	EndPoint (EP): PCIe Malformed Packet Checking is not implemented
43	X	X	No Fix	EndPoint (EP): PCIe EP Virtual Functions Wrongly Reporting Advisory Errors as Correctable Errors
44	X	X	No Fix	EndPoint (EP): PCIe EP Sends Two Error Messages to the Host for IO_RD to Unsupported Address Space
45	X	X	No Fix	EndPoint (EP): PCIe EP May Generate Spurious Error Message
46	X	X	No Fix	EndPoint (EP): PCIe EP Incorrect Handling of Multiple Errors
47	X	X	No Fix	EndPoint (EP): PCIe EP Incorrect Handling of Poisoned Memory Writes.
48	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Detect Poisoned Data for IO_WR Transactions
49	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Reliably Update the Header Log and First Error Pointer
50	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Report the Function Number Associated with a Poisoned TLP Error
51	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Send Error Message for MMIO Read transaction to Unsupported Memory Space
52	X	X	No Fix	EndPoint (EP): Active State Power Management (ASPM) Not Supported
53	X	X	No Fix	EndPoint (EP): Incorrect Logging and Reporting of Data Link Protocol Errors (DLPEs)
54	X	X	No Fix	EndPoint (EP): PCIe EP Transition from D3HOT to L1
55	X	X	No Fix	EndPoint (EP): Some PCIe EP Configuration Registers have Wrong Attribute Assignment
56	X	X	No Fix	EndPoint (EP): PCIe EP Transition from D3HOT to D0uninitialized
57	X	X	No Fix	EndPoint (EP): PCIe EP Link Control Register 2 (PLCNTLR2) Sticky Bits Not Retaining value after Hot Reset
58	X	X	No Fix	USB: PLL Configuration Settings for USB-I/O Interface
59	X	X	No Fix	EndPoint (EP): PCIe Transaction Pending Bit Initialization
60	X	X	No Fix	EndPoint (EP): Incorrect I/O Transaction Response in D3HOT State
61	X	X	No Fix	PCH: IEEE Std. 1149.6 EXTEST_PULSE and EXTEST_TRAIN Instructions Not Supported by Non-EP_JTAG Port
62	X	X	No Fix	PCH: Auxiliary (AUX) and Auxiliary2 (AUX2) Thermal Trip Interrupts not Functional
63	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Report Timeout Errors
64	X	X	No Fix	EndPoint (EP): PCIe EP Does Not Register Correctable Errors
65	X	X	No Fix	EndPoint (EP): PCIe EP Incorrect Handling of Read Completions with Unsupported Request (UR) Status
66	X	X	No Fix	EndPoint (EP): PCIe EP does not Enable the Correct number of Virtual Functions (VFs)

Table 1. Errata Summary Table (Sheet 3 of 3)

Erratum Number	DH98xxCC	DH89xxCL	Status	ERRATA
	C1	A0		
67	X		No Fix	GbE: SGMII Interface Receiver does not meet Input Differential Hysteresis (Vhyst) Specification
68	X	X	No Fix	SATA: Incorrect Number of Supported Ports Reported
69	X	X	No Fix	Endpoint (EP): PCIe EP Does not Indicate Poisoned Data from Reads of Corrupted Memory
70	X	X	No Fix	Endpoint (EP): PCIe EP Link May Not Train
71	X		No Fix	GbE: SGMII Interface Transmit Rise (t_{rise}) and Fall (t_{fall}) AC Timing Specification Violation
72			No Fix	Removed: Intel® QuickAssist Technology (QAT): Decompression Adler32 Checksum May be Calculated Incorrectly.
73	X	X	No Fix	Endpoint (EP): PCIe EP Link Disable requires Hot Reset
74	X		No Fix	GbE: GbE MACs PCI Config Space Does not contain Valid VID/DID Default Values
75	X		No Fix	GbE: GbE Interface Transmit Voltage Level Specification Violation
76	X	X	No Fix	USB: USB Full/Low Speed Port Reset or Clear Transaction Translation (TT) Buffer Request
77	X	X	No Fix	USB: USB RMH Think Time Issue
78	X	X	No Fix	USB: USB RMH False Disconnect Issue
79	X	X	No Fix	Endpoint (EP): Internal pull-ups on GbEx_LED/EP_xxx Signals De-asserted Before Sampling
80	X	X	No Fix	SATA: SATA Signal Voltage Level
81	X	X	No Fix	Endpoint (EP): VF Capability Pointers: Invalid Value for Setup
82	X	X	No Fix	PCH: Possible Increased Bit Error Rate (BER) on SATA, PCIe Root Ports, PCIe EndPoint, and LAN Port Media Interfaces
83	X	X	No Fix	DMI Port: GPIO17 Strap Is not Functional for DMI RX Termination

Table 2. Specification Changes

Number	Specification Change
1	GBE3_LED/EP_RESET_SEQ Strap LOW De-featured

Table 3. Specification Clarifications

Number	Specification Clarification
1	GbE: Use of Wake on LAN Together with Manageability
2	GbE SMBus: Illegal STOP Condition
3	GbE SERDES: AN_TIMEOUT Only Works When Link Partner Idle
4	External Strapping Requirements

Table 4. Documentation Changes

Number	Specification Change
1	BATLOW (Battery Low) feature is not supported in the Intel® Communications Chipset 8900 Series.
2	ME Non-Maskable Wake and some causes of host and global resets are not supported in the Intel® Communications Chipset 8900 Series.

Identification Information

Component Marking Information

The Intel® Communications Chipset 8900 Series PCH components are identified in the component markings in Table 5.

Table 5. Component Identification

PCH Stepping	MM#	S-Spec (GRP1LINE1)	Product	Notes
C1	923820	SLJW2	DH8900CC	Intel® Communications Chipset 8900
	923819	SLJVZ	DH8903CC	Intel® Communications Chipset 8903
	923818	SLJVY	DH8910CC	Intel® Communications Chipset 8910
	923817	SLJVX	DH8920CC	Intel® Communications Chipset 8920
A0	930162	SLK96	DH8925CL	Intel® Communications Chipset 8925
	931957	SLKCJ	DH8926CL	Intel® Communications Chipset 8926
	931963	SLKCK	DH8950CL	Intel® Communications Chipset 8950
	932468	SLKD4	DH8955CL	Intel® Communications Chipset 8955

Figure 1. Top Markings

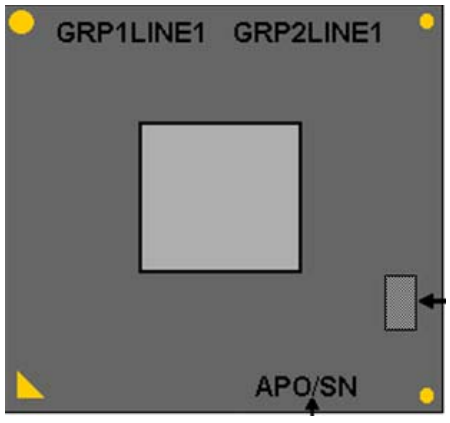

Device and Revision Identification

The Revision ID (RID) is traditionally an 8-bit register located at the offset 08h in the PCI header of every PCI device and function. The assigned value is based on the product stepping.

Table 6. Intel® Communications Chipset 8900 Series Device and Revision ID Table

Device Function	PCH PCIe Devices	Device ID (DID)	Revision ID (RID)	Comments
B0:D28:fn		0x244E		Used by the PCIe* Root Port function that has subtractive decode compatibility set to '1'.
B0:D31:F0	LPC	0x2310	0x08	LPC Controller
B0:D31:F2	SATA	0x2323	0x08	SATA Controller #1
B0:D31:F3	SMBus	0x2330	0x08	SMBus Host Controller
B0:D31:F5	SATA	0x2326	0x08	SATA Controller #2
B0:D31:F6	Thermal	0x2332	0x08	Thermal Subsystem
B0:D31:F7	WDT	0x2360	0x08	WDT for Core Reset
B0:D29:F0	USB	0x2334	0x08	USB Gen 2
B0:D22:F0	Intel® MEI #1	0x2364	0x08	Intel® Management Engine Interface #1
B0:D22:F1	Intel® MEI #2	0x2365	0x08	Intel® Management Engine Interface #2
B0:D28:F0	PCIe Root Port 1	0x2342	0x08	The PCIe Root Port with a subtractive decode will have a different Device ID (See Row 1 of this table).
B0:D28:F1	PCIe Root Port 2	0x2344	0x08	
B0:D28:F2	PCIe Root Port 3	0x2346	0x08	
B0:D28:F3	PCIe Root Port 4	0x2348	0x08	
BM:D0:F0	PCIe EndPoint	0x0434	0x21	PCIe Endpoint and Intel® QuickAssist Technology
BM:D0:F1	GbE MAC 0	0x0436	0x21	The Device ID for the GbE Controllers can be overwritten by the EEPROM. **Note:** Not applicable to the DH89xxCL
BM:D0:F2	GbE MAC 1	0x0436	0x21	
BM:D0:F3	GbE MAC 2	0x0436	0x21	
BM:D0:F4	GbE MAC 3	0x0436	0x21	

Errata

1. DMI Port: DMI Degradation Mode not Functional if Lane Reversal is implemented

Problem: DMI Degradation mode will not work when DMI lanes are routed in reverse order.

Implication: If the DMI Port is routed with the lanes reversed, it will work in full-mode only. The interface will not work if a link problem occurs and causes the interface to degrade.

Workaround: Do not route DMI lanes in reverse order.

Status: No Fix.

2. PCIe* Root Port (RP): PCIe RP Degradation Mode not Functional if Lane Reversal is implemented

Problem: PCIe RP Degradation mode will not work when the PCIe lanes are routed in reverse order.

Implication: If the PCIe RP is routed with the lanes reversed, it will work in full-mode only. The interface will not work if an error occurs and causes the interface to degrade.

Workaround: Do not route PCIe RP lanes in reverse order.

Status: No Fix.

3. EndPoint (EP): PCIe EP does not check the PCIe TLP Digest (TD) bit

Problem: The EP does not check the TD bit (bit 7 of byte 2) in the PCIe Transaction Layer Packet (TLP) header.

Implication: If the EP receives a TLP with the TD bit set, the following will happen:

- If the TLP contains the 4-byte digest field, the EP will not complete the request and will incorrectly respond with a Malformed Packet completion.
- If the TLP does not contain the 4-byte digest field, the EP will complete the request instead of responding with a Malformed Packet completion.

Workaround: Ensure that the TD bit in the TLP header is not set.

Status: No Fix.

4. GbE: GbE SMGII Protocol Layer Delays affects System Collision Detection

Problem: GbE SGMII protocol delays inhibits the proper detection of system collisions.

Implication: Collision detection errors may cause transmission data corruption.

Workaround: Reduce standard network transmission wire length by 10% to offset delays introduced by the SMGII Layer.

Status: No Fix.

5. GbE: GbE MAC may drop the first Receive (RX) packet after a Link Speed Change

Problem: The GbE MAC internal clock synchronization may not be complete before the SGMII link is established following a link speed change. This may result in the first Receive (RX) packet immediately following the link speed change to be truncated in the MAC layer and subsequently dropped due to bad Frame Check Sequence (FCS).

Implication: The Ethernet protocol FCS protects any bad/incomplete packet that may have appeared in MAC as an effect of this issue. Dropped packet due to bad FCS may be recorded in statistics.

Workaround: None.

Status: No Fix.

6. PCIe Root Port (RP): PCIe RP Degradation is not Functional

Problem: The PCIe RP cannot degrade from a x4 link to a x2 or x1 link.

Implication: If the RP is configured as PCIe x4 port and a x2 or x1 device is connected, the link will not train and the device will not be seen.

Workaround: Using the Flash Image Configuration Tool (FITC), configure Soft Strap 9 to bifurcate the root port as 4x1 or 1x4 as needed on the platform.

Status: No Fix.

7. SATA: SATA Low Power Device Detection

Problem: The SATA Low Power Device Detection (SLPD) may not recognize, or may falsely detect, a SATA hot-plug event during a Partial or Slumber Link Power Management (LPM) state.

Implication: On systems that enable LPM, when a SATA device attached to the PCH is configured as External or hot-plug capable, one of the following symptoms may occur:

- Symptom #1: A hot-plug or External SATA device removal which is not detected results in the OS and Intel® Matrix Storage Manager or Intel® Rapid Storage Technology console falsely reporting the device present, or incorrectly identifying an eSATA device.
- Symptom#2: A false hot-plug removal detection may occur resulting in OS boot hang or ODD media playback hang.

Workaround: A Software Driver workaround is implemented.

Status: No Fix.

8. USB: USB Port Stall with Bulk and Control Traffic

Problem: When a single USB bulk device is active on an EHCI controller, and the device has pending control and bulk traffic, the USB controller may not be able to resolve which traffic type is a priority and the association with the device may stall.

The processor must be in C0 for an extended period of time, such as when Cx states are disabled, or if system traffic prevents the system from leaving C0.

Implication: The USB device may appear unresponsive. If Cx states are enabled, the device may recover a short time later.

Note: Intel has only observed this failure on a limited number of devices. Failure only occurs if software associated with a USB device programs the Nak Count Reload bits defined in the *EHCI Specification for USB* Rev 1.0 to 0.

Workaround: BIOS implementation follows Intel® Communications Chipset 8900 Series *BIOS Specification* and program D29:F0:88h[7] = "1b."

Status: No Fix.

9. SATA: SATA SYNC Escape Issue

Problem: When SYNC Escape by a SATA device occurs on a D2H FIS, the Chipset 8900 Series does not set the PxIS.IFS bit to '1.' This deviates from section 6.1.9 of the Rev 1.3 Serial ATA Advanced Host Controller Interface (AHCI).

Implication: There is no known observable impact. Instead of detecting the IFS bit, software will detect a timeout error caused by the SYNC escape and then respond.

Workaround: None.

Status: No Fix.

10. USB: USB End of Frame When Retrying Packets Issue

Problem: If the USB controller encounters a full-speed or low-speed USB transaction with errors, it may retry the transaction without considering if the transaction can finish before the end of the current frame.

Implication: The implication depends on the particular USB device. The USB controller will attempt to recover per error handling specified in Section 4.5.2 of the *USB Specification* 2.0. The device may hang and require cycle to resume normal functionality.

Note: Intel has only observed this behavior on a limited number of USB devices. The implication only occurs if a USB device does not correctly respond to error handling as specified in Section 4.5.2 of the *USB Specification* 2.0.

Workaround: None.

Status: No Fix.

11. USB: USB Classic Device Removal Issue

Problem: If two or more USB full-speed/low-speed devices are connected to the same USB controller, and if the devices are not suspended, then if one device is removed, one or more of the devices remaining in the system may be affected by the disconnect.

Implication: The implication is device-dependent. A device may experience a delayed transaction, stall, and be recovered via software, or it may stall and require a reset such as a hot-plug to resume normal functionality.

Workaround: None.

Status: No Fix.

12. USB: USB RMH Descriptor May Report Incorrect Number of USB Ports

Problem: The PCH supports six ports on RMH #1 and may incorrectly report eight USB ports in the bNbrPorts field of the RMH hub descriptor.

Implication: If AC power is removed while the system is in hibernate, when the system resumes, new USB devices may not be detected, and all devices on RMH #1 may not function.

Note: AC power removal while the system is in hibernate is not a normal usage model, or common occurrence.

Workaround: None.

Status: No Fix.

13. PCH: High Precision Event Timer (HPET) Writing Timing Issue

Problem: A read transaction that immediately follows a write transaction to the HPET register space may return an incorrect value.

Implication: Implication depends on the usage model as noted below:

- For the HPET TIMn_COMP Timer 0 Comparator Value Register and HPET MAIN_CNT—Main Counter Value Register, the issue could result in the software receiving stale data. This may result in undetermined system behavior.

  Note: Timers [1:7] are not affected by this issue.

- For TIMERn_VAL_SET_CNF bit 6 in the TIMn_CONF—Timer n Configuration, there is no known usage model for reading this bit and there are no known functional implications.
- A write to the High Precision Timer Configuration (HPTC) register followed by a read to HPET register space, may return all 0xFFFF_FFFFh.

Workaround: Software workaround has been identified as described below:

- A write to the HPET TIMn_COMP Timer 0 Comparator Value Register should be followed by two reads that are discarded, and a third read where the data can be used.
- A write to the HPET MAIN_CNT - Main Counter Register should be followed by one read that is discarded, and a second read where the data can be used.

Workaround: TIMERn_VAL_SET_CNF bit 6 in the TIMn_CONF - Timer n - There is no known usage model to read this bit, but a write to the bit should be followed by one read that is discarded and a second read where the data can be used.

Status: No Fix.

14. USB: USB Full-Speed Port Staggering

Problem: When USB full-speed/low-speed port staggering is enabled, the USB controller may not wait for the bus to return to an idle state after an End of Packet (EOP), and may incorrectly acknowledge bus noise as a data packet.

Implication: Some full-speed/low-speed devices may fail to enumerate and function.

Note: This issue has been seen with a minimum number of devices on some motherboard ports with certain cable and trace lengths.

Workaround: BIOS should disable USB FS/LS Port Staggering by clearing RCBA+3564h[12].

Status: No Fix.

15. USB: USB Devices May Slow or Hang

Problem: When the processor is in C0, and a single bulk high-speed USB device is active, the port associated with the active device may hang.

Implication: The implication is device driver-dependent. Intel has observed some USB devices may have decreased performance, or the device may hang.

Workaround: BIOS implementation follows Intel® Communications Chipset 8900 Series *BIOS Specification* (section 16.9) and programs D29:F0:88h[7]= "1b."

Status: No Fix.

16. USB: USB Low-Speed Control Transactions

Problem: If the USB control buffers in the PCH Rate Matching Hub(s) are saturated with pending transactions, the buffers may not be serviced in round robin order.

Implication: Some low-speed endpoints may not receive their pending control transactions.

Note: This issue has only been observed in a synthetic test environment. The implication will be device, driver, and operating system specific.

Workaround: None.

Status: No Fix.

17. SATA: SATA Controller May Not Detect Unsolicited SATA COMINITs

Problem: SATA controller may not detect an unsolicited COMINIT from a SATA device.

Implication: The SATA device may not be properly detected and configured, resulting in the device not functioning as expected.

Workaround: BIOS implementation follows Intel® Communications Chipset 8900 Series *BIOS Specification* and programs D31:F2:Offset 98h[20:19] = "11b."

Status: No Fix.

18. SATA: SATA Hot Unplug May Not be Detected

Problem: SATA controller may not detect the unplug of a SATA 3.0 Gb/s device on a hot-plug enabled SATA port.

Implication: The unplugged SATA device may temporarily appear to be available.

Workaround: BIOS should program D31:F2:Offset 98h[6:5] = "00b."

Status: No Fix.

19. USB: USB Missing ACK

Problem: Following system power cycling or S3-S5 resume, if both high-speed and low-speed/full-speed devices are attached to the same controller, the host controller may not respond to a high-speed device ACK during a Get Descriptor request from the host software to a USB high-speed port.

Implication: USB high-speed devices may not be detected after a power cycling or S3-S5 resume.

- Intel has only observed this failure on a limited number of platforms. On a failing platform, the issue occurs infrequently.
- Full-speed and low-speed USB devices are not impacted by this issue.

Workaround: None.

Status: No Fix.

20. SATA: SATA 6 Gb/s Device Detection

Problem: The SATA controller may not be able to complete SATA Out Of Band (OOB) signaling with SATA 6Gb/s devices and down-shift to SATA 3 Gb/s speed.

Implication: SATA controller may not detect a SATA 6 Gb/s device upon power up or resume from S3, S4 or S5 State, resulting in indeterminate system behavior.

Workaround: None.

Status: No Fix.

21. PCIe Root Port (RP): PCIe RP Link Disable Bit

Problem: The PCIe RP may not exit the disable state when the Link Control Register "Link Disable" bit is set and PCIe Device Electrical Idle Exit is detected.

Implication: Port-specific software-directed Hot Plug or Power Management (PM) support using the "Link Disable" bit may cause PCIe RP to be stuck in the "Link Disable state" until a Host Reset with Power Cycling occurs.

Workaround: For PCIe RP port-specific software-directed Hot Plug or Power Management support, use the PCI Power Management Control register $D3_{HOT}$ bits instead of the Link Disable bit.

Status: No Fix.

22. USB: USB Isochronous In Transfer Error Issue

Problem: If a USB full-speed inbound isochronous transaction with a packet length 190 bytes or greater is started near the end of a micro-frame, the USB controller may see more than 189 bytes in the next micro-frame.

Implication: If the USB controller sees more than 189 bytes for a micro-frame, an error will be sent to software and the isochronous transfer will be lost. If a single data packet is lost, no perceptible impact for the end user is expected.

Note: Intel has only observed the issue in a synthetic test environment where precise control of packet scheduling is available, and has not observed this failure in its compatibility validation testing.

- Isochronous traffic is periodic and cannot be retried, thus it is considered good practice for software to schedule isochronous transactions to start at the beginning of a micro-frame. Known software solutions follow this practice.
- To sensitize the system to the issue, additional traffic such as other isochronous transactions or retries of asynchronous transactions would be required to push the inbound isochronous transaction to the end of the micro-frame.

Workaround: None.

Status: No Fix.

23. USB: USB Full-Speed/Low-Speed Device Removal Issue

Problem: If two or more USB full-speed/low-speed devices are connected to the same USB controller, the devices are not suspended, and one device is removed, one or more of the devices remaining in the system may be affected by the disconnect.

Implication: The implication is device-dependent. A device may experience a delayed transaction, stall, and be recovered via software, or it may stall and require a reset such as a hot-plug to resume normal functionality.

Workaround: None.

Status: No Fix.

24. USB: USB Babble Detected with Software Overscheduling

Problem: If software violates USB periodic scheduling rules for Full-Speed isochronous traffic by overscheduling, the RMH may not handle the error condition properly and return a completion split with more data than the length expected.

Implication: If the RMH returns more data than expected, the endpoint will detect packet babble for that transaction, and the packet will be dropped. Since overscheduling occurred to create the error condition, the packet would be dropped regardless of RMH behavior. If a single isochronous data packet is lost, no perceptible impact to the end user is expected.

Note: USB software overscheduling occurs when the amount of data scheduled for a micro-frame exceeds the maximum budget. This is an error condition that violates the USB periodic scheduling rule.

Note: This failure has only been recreated synthetically with USB software intentionally overscheduling traffic to hit the error condition.

Workaround: None.

Status: No Fix.

25. USB: USB Low-Speed/Full-Speed EOP Issue

Problem: If the EOP of the last packet in a USB Isochronous split transaction (defined as a transaction > 189 bytes) is dropped or delayed 3 ms or longer, the following may occur:

- If there are no other pending low-speed or full-speed transactions, the RMH will not send SOF, or Keep-Alive. Devices connected to the RMH will interpret this condition as idle and will enter suspend.
- If there are other pending low-speed or full-speed transactions, the RMH will drop the isochronous transaction and resume normal operation.

Implication:

- If there are no other transactions pending, the RMH is unaware a device has entered suspend and may start sending a transaction without waking the device. The implication is device-dependent, but a device may stall and require a reset to resume functionality.
- If there are other transactions present, only the initial isochronous transaction may be lost. The loss of a single isochronous transaction may not result in end user perceptible impact.

Note: Intel has only observed this failure when using software that does not comply with the USB specification and violates the hardware isochronous scheduling threshold by terminating transactions that are already in progress.

Workaround: None.

Status: No Fix.

26. USB: USB PLL Control FSM not Getting Reset on Global Reset

Problem: The USB PLL may not lock if a Global Reset occurs early during a cold boot sequence.

Implication: The USB Port would not be functional and additional cold boot would be necessary to recover.

Workaround: None.

Status: No Fix.

27. USB: USB Asynchronous Retries Prioritized Over Periodic Transfers

Problem: The integrated USB RMH incorrectly prioritizes low-speed and full-speed asynchronous retries over dispatchable periodic transfers.

Implication: Periodic transfers may be delayed or aborted. If the asynchronous retry latency causes the periodic transfer to be aborted, the impact varies depending on the nature of periodic transfer:

- If a periodic interrupt transfer is aborted, the data may be recovered by the next instance of the interrupt or the data could be dropped.
- If a periodic isochronous transfer is aborted, the data will be dropped. A single dropped periodic transaction should not be noticeable by the end user.

  Note: This issue has only been seen in a synthetic environment. The USB specification does not consider the occasional loss of periodic traffic a violation.

Workaround: None.

Status: No Fix.

28. USB: Incorrect Data for Low-Speed or Full-Speed USB Periodic IN Transaction

Problem: The Periodic Frame list entry in DRAM for a USB low-speed or full-speed Periodic IN transaction may incorrectly get some of its data from a prior Periodic IN transaction which was initiated very late into the preceding micro-frame.

It is considered good practice for software to schedule Periodic Transactions at the start of a micro-frame. However Periodic transactions may occur late into a micro-frame due to the following cases:

- Asynchronous transaction starting near the end of the proceeding micro-frame gets asynchronously retried.

  Note: Transactions getting asynchronous retried would only occur for ill-behaved USB device or USB port with a signal integrity issue.

- Two Periodic transactions are scheduled by software to occur in the same micro-frame and the first one needs to push the second Periodic IN transaction to the end of the micro-frame boundary.

Implication: The implication will be device, driver, or operating system specific.

Note: This issue has only been observed in a synthetic test environment.

Workaround: None.

Status: No Fix.

29. USB: USB RMH Delayed Periodic Traffic Timeout Issue

Problem: If an interrupt transaction is pushed to the x+4 micro-frame boundary due to asynchronous retries, the RMH may not wait for the interrupt transaction to time out before starting the next transaction.

IF RMH Transaction Translator (TT) reaches a discard boundary, a timeout may be ignored.

Implication: If the next transaction is intended for the same device targeted by the interrupt, the successful completion of that transaction is device dependent and cannot be guaranteed. The implication may differ depending on the nature of the transaction: (this only impacts TT – low-speed / full-speed).

- If the transaction is asynchronous and the device does not respond, it will eventually be retried with no impact.
- If the transaction is periodic and the device does not respond, the transfer may be dropped. A single dropped periodic transaction should not be noticeable by the end user.

Note: This issue has only been seen in a synthetic environment.

Workaround: None.

Status: No Fix.

30. GbE: GbE MNG Reset Clears Resource Grant With No Feedback

Problem: When accessing the EEPROM (via EEC register), grants may be lost due to deadlock or firmware reset. Software will not be notified of the lost grant. A driver in the middle of a bit bang may renew the request and receive the grant without knowing that it is actually starting a new transaction.

Implication:

1. EEC bit banging transactions may fail.
2. Long transactions may turn into different transactions than expected.

Workaround:

1. Software should not execute bit bang sequences longer than one word at a time.
2. When software reads the EEC, it should make sure that it still has the request and grant. If not, software should renew it and re-start the transaction. This does not cover all cases but reduces the possibility of a problem.

Status: No Fix.

31. GbE MDIO: Com_MDIO and Destination Bits of MDICNFG Register Are Not Loaded Consistently from EEPROM

Problem: The Com_MDIO (bit 30) and Destination (bit 31) bits of the MDICNFG register (0x0E04) are not loaded consistently from the EEPROM. In some cases, the hardware default value of 0b is used instead.

Implication: Software that relies on the EEPROM-loaded value might not be able to initialize an external PHY.

Workaround: Software should assume that the initial values of these bits are undefined and should program them before attempting to initialize an external PHY. The EEPROM bits can be used to determine the intended settings.

Note: A GbE driver workaround has been identified and deployed for Intel drivers.

Status: No Fix.

32. GbE MNG: I^{2}C Data Out Hold Time Violation

Problem: The GbE MNG I^{2}C interface should provide a data out hold time of 50 ns on the SFPx_I2C_DATA pins. The actual hold time is about 16 ns.

Implication: I^{2}C timing specification violation. There have been no reports of failures resulting from this timing. Note that the data input hold time required is zero, so the provided output hold time should be more than enough as long as the I^{2}C CLK and DATA signals are reasonably matched on the board.

Workaround: None.

Status: No Fix.

33. GbE EEPROM: EE_CS_N Control Signal Hold Time Violation

Problem: The EEPROM datasheet indicates a hold time of 250 ns for EE_CS_N relative to the falling edge of EE_SK (tCSH). At the end of a READ or RDSR operation, EE_CS_N is actually negated about 32 ns after the falling edge of EE_SK.

Implication: Although this is a timing specification violation for many EEPROM devices, no malfunction has been reported.

Workaround: None.

Status: No Fix.

34. GbE SGMII: Counters Incorrectly Increment on Collision

Problem: In SGMII mode/half duplex, the statistics counters listed below incorrectly increment when a collision occurs:

Name	Definition	Location
RLEC	Length error counter	0X4040
CRCERRS	CRC error counter	0x4000
RFC	receive frame counter	0x40A8

Implication: Error counters may not be accurate.

Workaround: None.

Status: No Fix.

35. GbE TSYNC: Auxiliary Timestamp from SDP is Unreliable

Problem: The SDP inputs to the timestamp logic are not properly synchronized. As a result, both the Auxiliary Timestamp Register values and the Auxiliary Timestamp Taken bits in TSAUXC are sometimes loaded incorrectly.

Implication: The auxiliary timestamp feature should be considered unreliable.

Workaround: For applications that use the auxiliary timestamp feature to synchronize to an external clock, it might be acceptable to drop some of the samples. For such applications, software can filter out many of the incorrect timestamp values by comparing them to an approximate expected timestamp and discarding unreasonable values.

In addition, the following method can be used to filter out incorrect values:

- Connect the input signal to two SDP inputs for the same port.
- Using the TSSDP register, assign one of the SDP inputs to AUX0 and the other SDP input to AUX1.
- When reading the TSAUXC register to check for new samples, check that both AUTT0 and AUTT1 are set. Otherwise, discard the sample.
- Read both the AUX0 and AUX1 timestamp values and compare the values. Discard the values if they differ by more than the sampling uncertainty of 8 ns (if the SDP inputs are balanced externally) or slightly higher if the external trace lengths differ significantly.

Using this method, along with a software filter for expected values, almost all errors can be filtered out, with the remaining samples having a very high probability of being correct.

When using Port 0, the following combinations of SDP connections to AUX0 and AUX1 allow the above method to filter out all errors.

AUX0 Connection	**AUX1 Connection**
SDP0	SDP1
SDP0	SDP3
SDP1	SDP0
SDP1	SDP2
SDP2	SDP0
SDP2	SDP1
SDP2	SDP3
SDP3	SDP2

Status: No Fix.

36. EndPoint (EP): Spurious MSI/MSIX Interrupt Generated

Problem: When binding an interrupt from the default interrupt core (core 0) to the associated interrupt core, extra interrupts may be generated when clearing the PCI_MSIX_ENTRY_VECTOR_CTRL during an active interrupt session.

Implication: Extra interrupts may be generated during interrupt core binding.

Workaround: Device driver workaround has been identified and deployed to handle the extra interrupts.

Status: No Fix.

37. EndPoint (EP): PCIe EP May not Detect Unexpected Completion Packets

Problem: The PCIe EP may drop packets with unexpected completions if the unexpected completion packets are received back-to-back.

Implication: Unexpected completion armored packets may not be detected and logged.

Workaround: None.

Status: No Fix.

38. GbE: GbE Near End Analog Loopback Not Supported

Problem: The GbE Near End Analog Loopback is not supported by the PCH GbE Interface.

Implication: GbE Near End Analog Loopback is not supported for Debug purposes.

Workaround: None.

Status: No Fix.

39. PCIe Root Port (RP): PCIe RP May not Automatically Switch into Compliance Mode

Problem: The PCIe RP may not automatically switch into compliance mode when the transmitter is terminated to 50-ohm test load.

Implication: Affects compliance testing.

Workaround: Software configuration may be used to force the link to enter compliance mode in both components of the link and then initiate a hot reset on the link.

Status: No Fix.

40. EndPoint (EP): PCIe EP Incorrect Completion ID

Problem: The Bus Number (B) section within a Completion ID (B:D:F) sent to the Host by the PCIe EP is always set to '0' for any completion response. The Device and Function Numbers (D:F) sections are set correctly.

Implication: Violation of PCIe Specification for Type 0 Configuration Write Requests where Functions are required to capture the Bus Number in the Completion ID of the response.

Workaround: None.

Status: No Fix.

41. EndPoint (EP): PCIe EP Link Status Register 2 (LNKSTS2) Not Implemented for Functions[1-4]

Problem: The PCIe LNKSTS2 Register in EP Functions[1-4] is not implemented. Reading the registers via Functions[1-4] returns indeterminate values.

Implication: The LNKSTS2 Register cannot be read via Functions[1-4].

Workaround: Read the LNKSTS2 Register via Function 0.

Status: No Fix.

42. EndPoint (EP): PCIe Malformed Packet Checking is not implemented

Problem: A malformed PCIe TLP (Transaction Layer Packet) is a TLP that violates specific TLP formation rules as defined in the PCI Express Base Specification Revision 2.0 dated December 20, 2006. The PCIe EndPoint does not implement checks for malformed TLPs, and thus no malformed TLP errors will be reported by the device. The PCIe EndPoint does not transmit malformed TLPs.

Implication: Reception of the following types of malformed TLPs by the PCIe EndPoint may cause silent data corruption:

- The size of the data payload of a received TLP as given by the TLP's length field exceeds the length specified by the value in the Max_Payload_Size field of the PCIe EndPoint (Section 2.2.2 of *PCIe Base Specification*).
- The value in the TLP's length field does not match the actual amount of data included in the TLP (Section 2.2.2, Sec 2.2.9 of *PCIe Base Specification*).
- A TLP where the TLP Digest (TD) bit value does not correspond with the observed size (accounting for the data payload, if present), (Section 2.2.3 of *PCIe Base Specification*).
- All received TLPs which use undefined Type field values (Section 2.3 of *PCIe Base Specification*).

Workaround: As of October 2011 there were no published malformed errata/sightings that did not have a workaround for all Root Ports or Downstream ports listed below that would allow a malformed formed packet to be sent to the EndPoint:

PCIe EndPoint directly connected to any of the PCIe Root Ports in these platforms:

- Sandy Bridge EN/EP + Patsburg PCH
- Sandy Bridge EN/EP + Chipset 8900 Series PCH
- Sandy Bridge Gladden + Chipset 8900 Series PCH
- Sandy Bridge Desktop + Cougar Point PCH
- Ivy Bridge EN/EP + Patsburg PCH
- Ivy Bridge Gladden + Chipset 8900 Series PCH
- Ivy Bridge Desktop + Panther Point PCH
- Jasper Forest + IbexPeak PCH
- Nehalem + Tylersburg IOH + ICH10
- Westmere + Tylersburg IOH + ICH10

PCIe EndPoint directly connected to a downstream port of the following PCIe Switches:

- PLX 87* family
- IDT IDT89HPES32NT24AG2 switch

Status: No Fix.

43. EndPoint (EP): PCIe EP Virtual Functions Wrongly Reporting Advisory Errors as Correctable Errors

Problem: The EP wrongly report Advisory Non-Fatal errors as Correctable Errors (ERR_COR_Messages) instead of as Non-Fatal Errors (ERR_NONFATAL Messages).

Implication: System software should handle these Advisory Errors as Non-Fatal Errors.

Workaround: None.

Status: No Fix.

44. EndPoint (EP): PCIe EP Sends Two Error Messages to the Host for IO_RD to Unsupported Address Space

Problem: When an unsupported request is detected with an IO_RD, the PCIe EP may return two error messages to the Host.

Implication: Induces multiple handling for the same error.

Workaround: Software should avoid IO_RD to unsupported address space.

Status: No Fix.

45. EndPoint (EP): PCIe EP May Generate Spurious Error Message

Problem: When clearing an Error status bit while handling an error, the PCIe Endpoint may spuriously generate another Error Message for the same error.

Implication: System software may be invoked for an error that has already been handled.

Workaround: None.

Status: No Fix.

46. EndPoint (EP): PCIe EP Incorrect Handling of Multiple Errors

Problem: If a function has a pending error that has not yet been handled and cleared, and a second error is detected by another function, the EP does not report the second error.

Implication: If the second error is an uncorrectable error, it goes undetected and may affect system operation.

Workaround: None.

Status: No Fix.

47. EndPoint (EP): PCIe EP Incorrect Handling of Poisoned Memory Writes.

Problem: The EP is supposed to handle a poisoned memory write as an Unsupported Request and abort the request, but the EP completes the write to the final destination with the poisoned data, and reports an error.

Implication: May cause unpredictable behavior if the writes target control structures.

Workaround: None.

Status: No Fix.

48. EndPoint (EP): PCIe EP Does Not Detect Poisoned Data for IO_WR Transactions

Problem: IO_WR transactions with poisoned data completes successfully without detection.

Implication: Target Registers may be corrupted.

Workaround: Software should perform a read after a write to verify the register value.

Status: No Fix.

49. EndPoint (EP): PCIe EP Does Not Reliably Update the Header Log and First Error Pointer

Problem: The Header Log and First Error Pointer are required to be updated for masked errors. If an unmasked error occurs after a masked error was detected, the Header Log and First Error Pointer are not be updated with the masked error information.

Implication: System software is unable to determine the PCIe transaction that caused the error.

Workaround: None.

Status: No Fix.

50. EndPoint (EP): PCIe EP Does Not Report the Function Number Associated with a Poisoned TLP Error

Problem: The PCIe EP reports Poisoned TLP error using Function Number 0 regardless of the Function that received the poisoned transaction.

Implication: Wrong error source indicator for errors not associate with Function 0.

Workaround: System software must scan all the PCIe EP functions to identify the error.

Status: No Fix.

51. EndPoint (EP): PCIe EP Does Not Send Error Message for MMIO Read transaction to Unsupported Memory Space

Problem: The PCIe EP does not send an error message for MMIO Read access to unsupported memory space even though it indicates the error in the Error Status Register and completes the Read Request with Unsupported Request Status.

Implication: None.

Workaround: System software should not access Unsupported Memory space.

Status: No Fix.

52. EndPoint (EP): Active State Power Management (ASPM) Not Supported

Problem: The PCIe EP does not support Active State Link Power Management.

Implication: If Active State Power Management is enabled and the link enters L0s, the EP may not re-enter L0. This will cause the system to hang.

Workaround: Disable Active State Power Management.

Status: No Fix.

53. EndPoint (EP): Incorrect Logging and Reporting of Data Link Protocol Errors (DLPEs)

Problem: The PCIe Specification requires that non function-specific DLPEs should be logged by all functions in the device which are configured to log the error, and should be reported by the configured functions. However, at the EP, only the last active function (last function which received a Config transaction) reports the error.

Implication: Violation of PCIe Specification for DLPE Error Reporting.

Workaround: None.

Status: No Fix.

54. EndPoint (EP): PCIe EP Transition from $D3_{HOT}$ to L1

Problem: If a MEM_RD or IO_WR request is received when the EP is in $D3_{HOT}$ state with the Link in L1, the EP appropriately transitions to L0 and responds to the request as Unsupported Request (UR). However, the link fails to transition back to the L1 state after the completion and stays in L0 state.

Implication: Violation of PCIe Power Management specifications.

Workaround: None.

Status: No Fix.

55. EndPoint (EP): Some PCIe EP Configuration Registers have Wrong Attribute Assignment

Problem: Some EP configuration registers have been wrongly assigned with Read/Write-Once-Sticky (RWOS) access attributes instead of Read-Only (RO) attributes.

The incorrect attribute affect the following Physical and Virtual Function (PF/VF) registers:

- PCC - PF Class Code Register
- PSVID - PF Subsystem Vendor ID
- PSID - PF Subsystem ID Register
- PLCAPR[0:3] - PF Link Capabilities Register
- PLSR[0:3] - PF Link Status Register
- PSRIOVFDID - PF SRIOV VF Device ID
- VCC[0:15] - VF Class Code Register
- VSID[0:15] - VF Subsystem ID Register

Implication: May fail PCIe Compliance tests.

Workaround: None.

Status: No Fix.

56. EndPoint (EP): PCIe EP Transition from $D3_{HOT}$ to $D0_{uninitialized}$

Problem: When the EP responds to a PCIe transition from $D0_{uninitialized}$ to $D3_{HOT}$, the link appropriately transitions from L0 to L1. However, when the EP is transitioned back from $D3_{HOT}$ to $D0_{uninitialized}$, the link transitions from L1-to-L0-to-L1, instead of L1-to-L0.

Implication: Violation of PCIe Power Management specifications.

Workaround: Perform the transition through $D0_{active}$ ($D0_{unitialized}$-to-$D0_{active}$-to-$D3_{HOT}$).

Status: No Fix.

57. EndPoint (EP): PCIe EP Link Control Register 2 (PLCNTLR2) Sticky Bits Not Retaining value after Hot Reset

Problem: The EP PCIe Configuration space PLCNTLR2 Register (Configuration: Offset 0xA4) bits with Read-Write-Sticky (RWS) attributes are reset to '0b' after Hot Reset.

Implication: Functions that require the bits to remain sticky are impacted. For example, it impacts system compliance testing.

Workaround: None.

Status: No Fix.

58. USB: PLL Configuration Settings for USB-I/O Interface

Problem: For PCH B0 Stepping, the power-up default clock phase may cause transmit clock timing violations.

Implication: The USB transmitter may not function properly if the appropriate clock phase is not selected.

Workaround: BIOS workaround has been identified and should be implemented as follows:

- BIOS should program USBIR<n>, Bit[31] = '1b.'
- BIOS should program PLLCTL2.U2TXPSEL (RCBA, Offset 0x356C), Bits[31:29] = 0x3.
- BIOS should program USBIR<n>, Bit[31] = '0b.'

Status: No Fix.

59. EndPoint (EP): PCIe Transaction Pending Bit Initialization

Problem: The Transaction Pending bit in the PCIe Status Register (PPDSTAT.TP) indicates that the PCIe Function has pending Non-Posted requests awaiting completions. The PCIe Specification requires the bit be cleared when all pending requests have completed successfully or completed with a timeout. At the EP, the bit is cleared only for transactions that complete successfully; it is not cleared if the pending transaction complete with a timeout.

Implication: The Transaction Pending bit is one of the resources that software uses to quiesce a function prior to issuing a Function Level Reset (FLR).

Workaround: None.

Status: No Fix.

60. EndPoint (EP): Incorrect I/O Transaction Response in $D3_{HOT}$ State

Problem: The EP responds normally to all I/O transactions when in $D3_{HOT}$ state, as if in D0 state, instead of responding with Unsurpported Request (UR).

Implication: Violation of PCIe Specifications.

Workaround: Avoid I/O transactions when EP is in $D3_{HOT}$ state.

Status: No Fix.

61. PCH: IEEE Std. 1149.6 EXTEST_PULSE and EXTEST_TRAIN Instructions Not Supported by Non-EP_JTAG Port

Problem: The PCH contains two separate JTAG (TAP) ports. One is used by the Non-Endpoint section on the PCH (JTAG) and the second is for the Endpoint section on the PCH (EP_JTAG). Both JTAG ports are compatible with the IEEE* Std. 1149.1; but only the EP_JTAG is compatible with IEEE Std. 1149.6 instructions set. The Non-EP JTAG (JTAG), which connects to the DMI/SATA/PCIe (Root)/USB legacy I/O interfaces, is not compatible with the IEEE Std. 1149.6 and does not support the EXTEST_PULSE and EXTEST_TRAIN instructions.

Implication: IEEE Std.1149.6 instruction set not supported by the Non-EP JTAG port.

Workaround: None.

Status: No Fix.

62. PCH: Auxiliary (AUX) and Auxiliary2 (AUX2) Thermal Trip Interrupts not Functional

Problem: The generation of PCI interrupts on AUX and AUX2 thermal trip points in the Thermal Sensors (TS0 and TS1) do not work properly.

Implication: Thermal interrupts are limited to Catastrophic and Hot state only.

Workaround: BIOS should disable the generation of interrupts on AUX and AUX2 thermal trips (TS0PIEN/TS1PIEN; TBARB + 0x82/0xC2) as follows:

- BIOS should program TS0PIEN/TS1PIEN, Bit[0,3,4,7] = '0b.'

Status: No Fix.

63. EndPoint (EP): PCIe EP Does Not Report Timeout Errors

Problem: When the PCIe Endpoint transaction timer expires, an error message is not sent to the Host. However, all the relevant error status bits are updated correctly.

Implication: Results in incomplete transaction which may cause the device to hang.

Workaround: Device driver should implement a timeout mechanism to detect device hang condition. This timeout mechanism can then be used by the software to intervene and reset the device.

Status: No Fix.

64. EndPoint (EP): PCIe EP Does Not Register Correctable Errors

Problem: When the PCIe EP receives and detects a Receiver Error it does not set the Correctable Error Detected (CED) bit in the PF PCIe Device Status Register (PPDSTAT[0]). However, a correctable error message is sent to the Host and the Receiver Error bit (RES) is set in the PF PCIe AER Correctable Error Register (PPAERCS[0]).

Implication: Violation of system error reporting.

Note: Since the PCIe EP supports AER features, the device driver can use the AER Correctable Error Status register to identify the cause of the correctable error message.

Workaround: None.

Status: No Fix.

65. EndPoint (EP): PCIe EP Incorrect Handling of Read Completions with Unsupported Request (UR) Status

Problem: When the PCIe Endpoint receives a Read Completion with UR status it does not indicate the UR condition to the Requester of the read. The Requester may wait indefinitely for the completion of the transaction which can cause I/O requests back-up at the device.

Implication: The transaction incompletion may cause the device and system to hang.

Workaround: Device driver should implement a timeout mechanism to detect device hang condition. This timeout mechanism can then be used by the software to intervene and reset the device.

Status: No Fix.

66. EndPoint (EP): PCIe EP does not Enable the Correct number of Virtual Functions (VFs)

Problem: The NUMVF field of PF SRIOV Number of VFs Register (PSRIOVNUMVF), (BM:D0:F0; Offset 150h - 153h, Bits[15:0]) allows software to define the number of VFs assigned to a Physical Function (PF) as part of the process of creating VFs. Programming the NUMVF field does not enable the correct number of VFs as expected. The actual number of Enabled and Expected VFs as shown as follows:

NUMVF	Actual # of VFs Enabled	Expected # of VFs
0x00	0	0
0x01	2	1
0x02	3	2
0x03	4	3
0x04	5	4
0x05	6	5
0x06	7	6
0x07	8	7
----	----	----
0x0F	16	15
0x10	Undefined	16

Implication: Inhibits the capability to enable only 1 VF. Provides capability to enable 2-16 VFs.

Workaround: None.

Status: No Fix.

67. GbE: SGMII Interface Receiver does not meet Input Differential Hysteresis (Vhyst) Specification

Problem: The SGMII Specification requires the receiver to ignore signals below 25 mV hysteresis voltage. However, the receiver may change states on input voltages close the 25 mV hysteresis specification which may cause unwanted state changes with non-monotonic input signals.

Implication: Although this is a specification violation, it does not affect the functional operation of the interface.

Workaround: None.

Status: No Fix.

68. SATA: Incorrect Number of Supported Ports Reported

Problem: The SATA AHCI Host Capabilities Register is configured to indicate that the SATA Controller supports four ports (B0:D31:F2: ABAR + 00h, Bit[4:0] = "03h"). The number of supported ports in the register (NPs) should have been set to "01h" to indicate support for two ports.

Implication: Incorrect reporting of supported ports. All PCH SKUs support only two ports.

Workaround: Driver should ignore the setting in the CAP Register, but instead to use the BIOS configuration of the "Active Ports" in the Ports Implemented Register (B0:D31:F2: ABAR + 0Ch, Bit[5:4]).

Status: No Fix.

69. Endpoint (EP): PCIe EP Does not Indicate Poisoned Data from Reads of Corrupted Memory

Problem: The EP fails to set the Poisoned Data bit (EP bit) in the completion header for memory reads with corrupted data.

Implication: Minimal impact because IA software is notified via interrupts when internal errors are detected.

Workaround: None.

Status: No Fix.

70. Endpoint (EP): PCIe EP Link May Not Train

Problem: The default mux setting selects the wrong analog detect signal which may cause the internal State Machines to get out of sync.

Implication: If the state machines get out of sync, the link will not train.

Workaround: This issue can be fixed with an EEPROM image upgrade.

- Using the Intel starter images:
  - — The EEPROM Version = 3.0 or higher, includes the workaround.

OR

- If you build your own image with EICT, include the following: "LAN0-CSR_Auto_Config_Power_Up_Ptr_0x27-CSR-Bottom_C.txt" version2 or higher in the EICT Pointer Field "LAN0 > LAN0 CSR Auto Config Power Up Ptr."

Note: Intel recommends that the latest EEPROM Version (version 3.0 or higher) should be used.

Status: No Fix.

71. GbE: SGMII Interface Transmit Rise (t_{rise}) and Fall (t_{fall}) AC Timing Specification Violation

Problem: The SGMII interface transmit Rise (t_{rise}) and Fall (t_{fall}) times do not meet the SGMII AC Timing Specification.

- The expected specification timing for t_{rise} and t_{fall} is 100 psec (min) and 200 psec (max) for 20% - 80% voltage rise or fall.
- The measured PCH SGMII interface timing is 85 psec for 20% - 80% voltage rise or fall.

Implication: Although this is a specification violation, it does not affect the functional operation of the interface.

Workaround: None.

Status: No Fix.

72. Removed

73. Endpoint (EP): PCIe EP Link Disable requires Hot Reset

Problem: The EP does not perform an internal upstream port reset on PCIe Link Disable as required by the PCI Express specification.

Implication: If Link Disable is performed the credit consumed counters are not reset, once the link re-trains the device will hang due to lack of available credits.

Workaround: After clearing the Link Disable bit in the Root or Switch Downstream Port connected to the PCH, a Hot Reset must be performed by setting the Secondary Bus Reset bit of the Bridge Control Register associated with the same Root or Switch Downstream Port.

Status: No Fix.

74. GbE: GbE MACs PCI Config Space Does not contain Valid VID/DID Default Values

Problem: The GbE MACs do not have valid VID/DID default values in the PCI configuration space for the PCH C0 and C1 Steppings.

Implication: If an EEPROM with a valid image is not loaded during power-up or system reset, then the GbE MACs VID and DID may be randomly set and invalid.

Note: It is a mandatory requirement to have a pre-programmed EEPROM image in the system; hence this is only an **ALERT** to provide the EEPROM.

Workaround: Ensure that the platform contains a pre-programmed EEPROM with valid VID/DID before powering up the system.

Status: No Fix.

75. GbE: GbE Interface Transmit Voltage Level Specification Violation

Problem: The GbE 1000BASE-BX transmit buffers have been designed to maximize performance and robustness over a variety of routing scenarios. As a result, the transmitter signaling voltage levels may exceed the maximum motherboard TX connector and RX connector peak-to-peak voltage specification of 1350 mV by about 5%. (See Transmitter Specification at TP-T (Table 17), of PICMG 3.1, R1.0 Specification.)

Implication: Although this is a specification violation, it does not affect the functional operation of the interface.

Workaround: Ensure that the signal routing length from TP-1 to TP-4 is greater than 10". Refer to *Intel® Communications Chipset 8900 Series Platform Design Guide (PDG)* for maximum routing length. (See Figure 2.)

Figure 2. 1000BASE-BX/FC-PI Transmitter Electrical Specifications

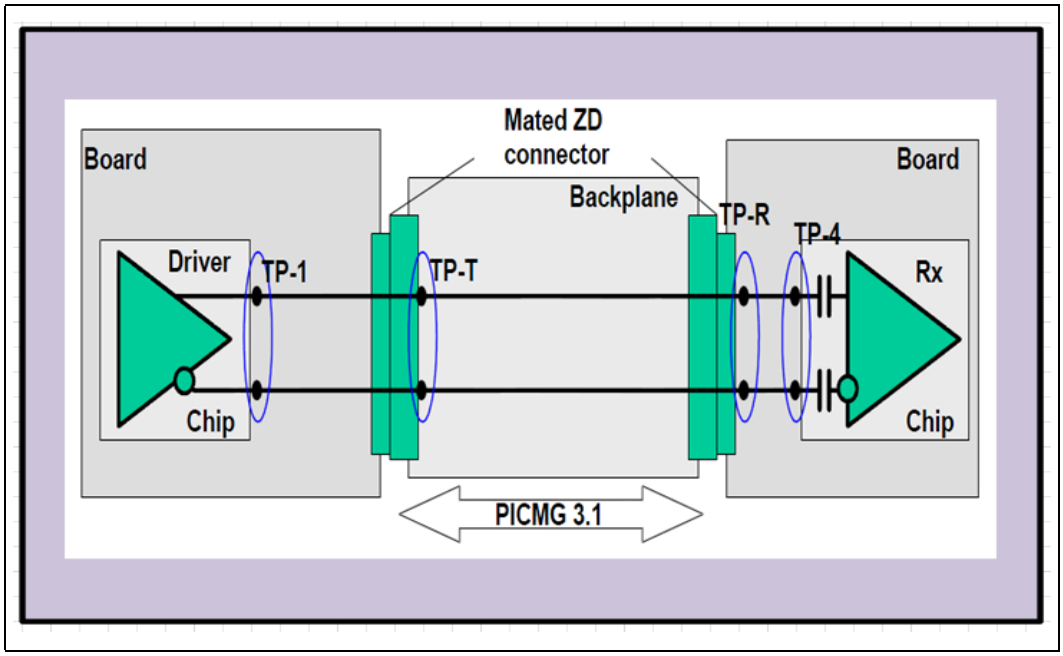

Status: No Fix.

76. USB: USB Full/Low Speed Port Reset or Clear Transaction Translation (TT) Buffer Request

Problem: One or more full/low speed USB devices on the same RMH controller may be affected if the devices are not suspended and either (a) software issues a Port Reset OR (b) software issues a Clear TT Buffer request to a port executing a split full/low Speed Asynchronous Out Command.

- The small window of exposure for Full-Speed device is around 1.5 micro-seconds and around 12 micro-seconds for Low-Speed device.

Implication: The affected port may stall or receive stale data for a newly arrived split transfer occurring at the time of the Port Reset or Clear TT Buffer request.

Note: This issue has only been observed in a synthetic test environment.

Workaround: None.

Status: No Fix.

77. USB: USB RMH Think Time Issue

Problem: The USB RMH Think Time may exceed its declared value in the RMH hub descriptor register value of 8 full-speed bit times.

Implication: If the OS USB driver fully subscribes to USB microframe, LS/FS transactions may exceed the microframe boundary.

Note: No functional failures have been observed.

Workaround: None.

Status: No Fix.

78. USB: USB RMH False Disconnect Issue

Problem: The PCH may falsely detect a USB High-Speed (HS) device disconnect if all of the following conditions are met:

1. The HS device is connected through the Rate Matching Hub (RMH) of the PCH's EHCI controller either directly or through a high-speed hub or series of high-speed hubs.
2. The device is resuming from selective suspend or port reset.
3. The resume occurs within a narrow time window during the EOP (End of Packet) portion of the SOF (Start of Frame) Packet on the USB bus.

Implication: Following the false disconnect, the HS device will be automatically re-enumerated. The system implication will depend on the cause of the resume event:

- If the resume event is a port reset, a second port reset will be automatically generated and the device re-enumerated. No end user impact is expected.
- If the resume event is a hardware or software initiated resume from selective suspend, the implication will be device and software specific, which may result in anomalous system behavior.

Note: If the HS device is a hub, then all of the devices behind the hub, independent of the device speed, may also be re-enumerated.

Workaround: None.

Status: No Fix.

79. Endpoint (EP): Internal pull-ups on GbEx_LED/EP_xxx Signals De-asserted Before Sampling

Problem: The internal pull-ups on GbE[3:0]_LED/EP_xxx signals de-assert prior to sampling of internal straps muxed onto these signals.

Implication: Muxed straps may be sampled in a LOW state leading to strap based settings not being appropriately configured for the following straps:

- EP Slave SMBus Address Configuration
- PCIe SR-IOV/ARI Enable/Disable
- PCIe CEM 2.0 Compliant Reset Sequence

Workaround: Implement board level termination on GbE[3:0]_LED_xxx signals. Refer to GBE3_LED/EP_RESET_SEQ Strap LOW De-featured and External Strapping Requirements for GBE[3:0]_LED_xxx updated strap requirements.

Status: No Fix.

80. SATA: SATA Signal Voltage Level

Problem: The 1.5 Gb/s and 3.0 Gb/s SATA transmit buffers have been designed to maximize performance and robustness over a variety of routing scenarios. As a result, the SATA 1.5 Gb/s & 3.0 Gb/s (Gen1i, Gen1m, Gen2i, and Gen2m) transmit signaling voltage levels may exceed the maximum motherboard TX connector and device RX connector voltage specifications (Section 7.2.1 of the *Serial ATA Specification*, Revision 2.5).

Implication: Although this is a specification violation, it does not affect the functional operation of the interface.

Workaround: None.

Status: No Fix.

81. Endpoint (EP): VF Capability Pointers: Invalid Value for Setup

Problem: The values for the capability pointers are incorrect resulting in a broken link list.

Implication: Capabilities will not be visible.

Workaround: None.

Status: No Fix.

82. PCH: Possible Increased Bit Error Rate (BER) on SATA, PCIe Root Ports, PCIe EndPoint, and LAN Port Media Interfaces

Problem: A small subset of parts could be susceptible to an increased occurrence in Bit Error Rate (BER) up to link failure for the SATA, PCIe Root Ports, PCIe EndPoint and the LAN port media interfaces when all supply rails are near maximum voltage (V_{max}).

Implication: If a system has a susceptible part and all of the supply rails are near V_{max}, then increased BER up to link failure may occur.

Workaround: BIOS and EEPROM workarounds have been identified and should be implemented as follows:

- For SATA and PCIe Root Ports update the mPHY setting as documented in the *RS-*Intel® Communications Chipset 8900 Series *BIOS Specification* (Electronic Only Doc# 32417) version 1.3 or later.
- For PCIe EndPoint and LAN port Media interfaces update the EEPROM to version 3.3 or later.
  - The latest EEPROMs may be obtained from the EEPROM Collateral for Intel® Communications Chipset 8900 Series - *Download Instructions* (Doc# 470620).

Status: No Fix.

83. DMI Port: GPIO17 Strap Is not Functional for DMI RX Termination

Problem: Direct Media Interface (DMI) RX (DMI_RX(p/n)[3:0]) termination is referenced to VCCPCPU. The GPIO17 strap does not change the reference to Vss as specified in the Datasheet.

Implication: If using any of the DC coupled processors in Table 7, the signal eye height is reduced and the part is sourcing more current than it was designed to, which has a significant negative impact on the published pFail; therefore, the workaround MUST be applied.

Table 7. Workaround Must Be Applied

Processor	CPU ID[1]
Intel® Xeon® E3-1125C v2	0x000306Ax
Intel® Xeon® E3-1105C v2	0x000306Ax
Intel® Core™ i3 3115	0x000306Ax
Intel® Pentium® B925C	0x000306Ax
Intel® Core™ i7 4702EC	0x000306Cx
Intel® Core™ i7 4700EC	0x000306Cx
Intel® Core™ i5 4402EC	0x000306Cx

[1] "x" is the stepping and must be ignored.

Note: Workaround MUST NOT be applied:

If the DMI interface is AC coupled regardless of processor

Or

If using any of the processors in Table 8.

Table 8. Workaround Must Not be Applied

Processor	CPU ID[1]
Intel® Xeon® E3-1125C	0x000206Ax
Intel® Xeon® E3-1105C	0x000206Ax
Intel® Core™ i3 2115	0x000206Ax
Intel® Pentium® B915C	0x000206Ax
Intel® Celeron® 725C	0x000206Ax

[1] "x" is the stepping and must be ignored.

Workaround: BIOS implementation follows *Intel® Communications Chipset 89xx Series BIOS Specification* (Electronic Only Doc# 33545) version 1.4 or later. See section titled "Early DMI Termination Initialization."

Note: Intel has completed all certifications and validation necessary to ensure the DMI interface will train to allow initial booting to apply the fix if your design is within the Platform Design Guidelines (PDG) guidelines published for the processor used in your design.

If your design is outside of the PDG guidelines, additional analysis is required. See document *Intel® Communications Chipset 89xx Series: Margin for Direct Media Interface (DMI)* (document number 544453) for details.

Status: No Fix.

Specification Changes

1. GBE3_LED/EP_RESET_SEQ Strap LOW De-featured

The GBE3_LED/EP_RESET_SEQ pull-down strap option is de-featured and not supported by Intel® Communications Chipset 89xx Series. The option is re-defined as "RESERVED" in Table 32-8 and is not a supported configurable option.

Update to **Table 32-8. LED, Software Defined, Miscellaneous Signals**

Signal Name	I/O Type	Technology	Ball Count	Internal/ External Resistor Pull-Up/Down	BScan Support	Description
GBE3_LED/ EP_RESET_SEQ[1]	O	LVTTL	1	Internal (Active for 200 µs after assertion of AUX_PWROK) Refer to *Note* under Description	1149.1	Port 3 LED. Programmable LED, mode encoding set by GbE EEPROM. This signal can be use as a strapping function to control reset sequence. 0 ~~Silicon A0 Compatible Reset Sequence (Pull-Down Required)~~ RESERVED 1 PCIe* CEM 2.0 Compliant Reset sequence (~~default~~ Pull-up Required) Sampling occurs during the first 200 ns after assertion of AUX_PWROK. Thereafter the pin functions as a LED output. ***Note:*** External board strap required to guarantee expected configuration setting. Assert strap active until 1ms after assertion of AUX_PWROK.

1. Naming applicable to DH89xxCC/DH89xxCL, respectively

Specification Clarifications

1. GbE: Use of Wake on LAN Together with Manageability

Note: This specification clarification is applicable to the DH89xxCC SKU.

Clarification: The Wakeup Filter Control Register (WUFC) contains the NoTCO bit, which affects the behavior of the wakeup functionality when manageability is in use. Note that if manageability is not enabled, the value of NoTCO has no effect.

When NoTCO contains the hardware default value of 0b, any received packet that matches the wakeup filters will wake the system. This could cause unintended wakeups in certain situations. For example, if Directed Exact Wakeup is used and the manageability shares the host's MAC address, IPMI packets that are intended for the BMC will wake the system, which might not be the intended behavior.

When NoTCO is set to 1b, any packet that passes the manageability filter, even if it also is copied to the host, is excluded from the wakeup logic. This solves the previous problem since IPMI packets will not wake the system. However, with NoTCO = 1b, broadcast packets, including broadcast magic packets, will not wake the system since they pass the manageability filters and are therefore excluded.

The Intel Windows* drivers set NoTCO by default.

Effects of NoTCO Settings WoL	NoTCO	Shared MAC Address	Unicast Packet	Broadcast Packet
Magic Packet	0b	-	OK	OK
Magic Packet	1b	Y	No wake	No wake
Magic Packet	1b	N	OK	No wake
Directed Exact	0b	Y	Wake even if MNG packet. No way to talk to BMC without waking host.	N/A
Directed Exact	0b	N	OK	N/A
Directed Exact	1b	-	OK	N/A

2. GbE SMBus: Illegal STOP Condition

Note: This specification clarification is applicable to the DH89xxCC SKU.

Clarification: It is important to prevent illegal STOP conditions on the SMBus interface, even when resetting the Management Controller (MC).

Specifically, a STOP condition should never be generated by the MC during the high clock phase of an ACK cycle while reading packet data from GbE Controller as part of a Receive TCO LAN packet transaction.

If this situation occurs, the Controller replies with a NACK to all future commands until a power cycle. As a result, the SMBus interface becomes inoperable.

Workaround: Ensure that this illegal sequence does not occur, even during MC reset.

3. GbE SERDES: AN_TIMEOUT Only Works When Link Partner Idle

Note: This specification clarification is applicable to the DH89xxCC SKU.

Clarification: The auto-negotiation timeout mechanism (PCS_LCTL.AN_TIMEOUT_EN) only works if the SerDes partner is sending idle code groups continuously for the duration of the timeout period, which is the usual case.

However, if the partner is transmitting packets, an auto-negotiation timeout will not occur since auto-negotiation is restarted at the beginning of each packet. If the partner has an application that indefinitely transmits data despite the lack of any response, it is possible that a link will not be established.

Workaround: If this is a concern, the auto-negotiation timeout mechanism may be considered unreliable and an additional software mechanism could be used to disable auto-negotiation if sync is maintained without a link being established (PCS_LSTS.SYNC_OK=1b and PCS_LSTS.LINK_OK=0b) for an extended period of time.

4. External Strapping Requirements

Clarification: The internal pull-ups on GBE[3:0]_LED signals de-assert before sampling. External platform straps are required on these signals for expected configuration options. Table 32-8 is updated to define external strap requirements. See Errata #79 “Endpoint (EP): Internal pull-ups on GbEx_LED/EP_xxx Signals De-asserted Before Sampling”.

Update to **Table 32-8. LED, Software Defined, Miscellaneous Signals**

Signal Name	I/O Type	Technology	Ball Count	Internal/ External Resistor Pull-Up/ Down	BScan Support	Description
GBE0_LED/ EP_SMB_ADR2[3]	O	LVTTL	1	Internal (Active for 200 µs after assertion of AUX_PWROK) Refer to *Note* under Description	1149.1	Port 0 LED. Programmable LED, mode encoding set by GbE EEPROM. On power-up this signal becomes an input to Strap SMBus Slave Address bit 2. The EP SMBus slave is accessed with address[7:1] = 1110_XX0. Sampling occurs during the first 200 ns after assertion of AUX_PWROK. Thereafter the pin functions as an LED output. External board strap required to guarantee expected configuration setting. Assert strap active until 1ms after assertion of AUX_PWROK.
GBE1_LED/ EP_SMB_ADR3[3]	O	LVTTL	1	Internal (Active for 200 µs after assertion of AUX_PWROK) Refer to *Note* under Description	1149.1	Port 1 LED. Programmable LED, mode encoding set by GbE EEPROM. On power-up this signal becomes an input to Strap SMBus Slave Address bit 3. The EP SMBus slave is accessed with address[7:1] = 1110_XX0. Sampling occurs during the first 200 ns after assertion of AUX_PWROK. Thereafter the pin functions as an LED output. External board strap required to guarantee expected configuration setting. Assert strap active until 1ms after assertion of AUX_PWROK
GBE2_LED/ EP_VF_ENABLED[3]	O	LVTTL	1	Internal (Active for 200 µs after assertion of AUX_PWROK) Refer to *Note* under Description	1149.1	Port 2 LED. Programmable LED, mode encoding set by GbE EEPROM. This signal can be use as a strapping function to enable or disable PCIe* SRIOV GBE2_LED: 0 Disable SRIOV (Pull-Down Required) 1 Enable SRIOV (~~default~~ Pull-up Required) Sampling occurs during the first 200 ns after assertion of AUX_PWROK. Thereafter the pin functions as an LED output. External board strap required to guarantee expected configuration setting. Assert strap active until 1ms after assertion of AUX_PWROK
GBE3_LED/ EP_RESET_SEQ[3]	O	LVTTL	1	Internal (Active for 200 µs after assertion of AUX_PWROK) Refer to *Note* under Description	1149.1	Port 3 LED. Programmable LED, mode encoding set by GbE EEPROM. This signal can be use as a strapping function to control reset sequence. 0 ~~Silicon A0 Compatible Reset Sequence (Pull-Down Required)~~ RESERVED 1 PCIe* CEM 2.0 Compliant Reset sequence (~~default~~ Pull-up Required) Sampling occurs during the first 200 ns after assertion of AUX_PWROK. Thereafter the pin functions as a LED output. External board strap required to guarantee expected configuration setting. Assert strap active until 1ms after assertion of AUX_PWROK

Notes:
1. For designs where these signals are **NOT** used to drive platform GBE LEDs, the board level terminations may be hardwired to the desired strap configuration
2. Refer to Intel® Xeon® Processor E5-2658 and E5-2448L with Intel® Communications Chipset 8920 Development Kit for implementation reference.
3. Naming applicable to DH89xxCC/DH89xxCL, respectively.

Documentation Changes

1. BATLOW (Battery Low) feature is not supported in the Intel® Communications Chipset 8900 Series.

The following corrections must be applied to the Intel® Communications Chipset 8900 Series Datasheet.

Table 4-16 Configuration Bits Reset by RTCRST# Assertion (Sheet 1 of 2)

Bit Name	Register	Location	Bit(s)	Default State
Alarm Interrupt Enable (AIE)	Register B (General Configuration) (RTC_REGB)	I/O space (RTC Index + 0Bh)	5	X
Alarm Flag (AF)	Register C (Flag Register) (RTC_REGC)	I/O space (RTC Index + 0Ch)	5	X
SWSMI_RATE_SEL	General PM Configuration 3 Register GEN_PMCON_3	B0:D31:F0:A4h	7:6	0
SLP_S4# Minimum Assertion Width	General PM Configuration 3 Register GEN_PMCON_3	B0:D31:F0:A4h	5:4	0
SLP_S4# Assertion Stretch Enable	General PM Configuration 3 Register GEN_PMCON_3	B0:D31:F0:A4h	3	0
RTC Power Status (RTC_PWR_STS)	General PM Configuration 3 Register GEN_PMCON_3	B0:D31:F0:A4h	2	0
Power Failure (PWR_FLR)	General PM Configuration 3 Register (GEN_PMCON_3)	B0:D31:F0:A4h	1	0
AFTERG3_EN	General PM Configuration 3 Register GEN_PMCON_3	B0:D31:F0:A4h	0	0
Power Button Override Status (PRBTNOR_STS)	Power Management 1 Status Register (PM1_STS)	PMBase + 00h	11	0
RTC Event Enable (RTC_EN)	Power Management 1 Enable Register (PM1_EN)	PMBase + 02h	10	0
Sleep Type (SLP_TYP)	Power Management 1 Control (PM1_CNT)	PMBase + 04h	12:10	0
PME_EN	General Purpose Event 0 Enables Register (GPE0_EN)	PMBase + 2Ch	11	0
~~BATLOW_EN~~	~~General Purpose Event 0 Enables Register (GPE0_EN)~~	~~PMBase + 2Ch~~	~~10~~	~~0~~

Table 4-16 Configuration Bits Reset by RTCRST# Assertion (Sheet 2 of 2)

Bit Name	Register	Location	Bit(s)	Default State
RI_EN	General Purpose Event 0 Enables Register (GPE0_EN)	PMBase + 2Ch	8	0
NEWCENTURY_STS	TCO1 Status Register (TCO1_STS)	TCOBase + 04h	7	0
Intruder Detect (INTRD_DET)	TCO2 Status Register (TCO2_STS)	TCOBase + 06h	0	0
Top Swap (TS)	Backed Up Control Register (BUC)	Chipset Config Registers:Offset 3414h	0	X

Table 4-20 State Transition Rules

Present State	Transition Trigger	Next State
G0/S0/C0	• DMI Msg • SLP_EN bit set • Power Button Override • Mechanical Off/Power Failure	• G0/S0/Cx • G1/Sx or G2/S5 state • G2/S5 • G3
G0/S0/Cx	• DMI Msg • Power Button Override • Mechanical Off/Power Failure	• G0/S0/C0 • S5 • G3
G1/S1, G1/S3, or G1/S4	• Any Enabled Wake Event • Power Button Override • Mechanical Off/Power Failure	• G0/S0/C0 (See Note 2) • G2/S5 • G3
G2/S5	• Any Enabled Wake Event • Mechanical Off/Power Failure	• G0/S0/C0 (See Note 2) • G3
G3	• Power Returns	• Optional to go to S0/C0 (reboot) or G2/S5 (stay off until power button pressed or other wake event). (See Notes 1 and 2)

Notes:
1. Some wake events can be preserved through power failure.
2. ~~Transitions from the S1-S5 or G3 states to the S0 state are deferred until BATLOW# is inactive in mobile configurations.~~

~~4.11.9.4 BATLOW# (Battery Low)~~

~~The BATLOW# input can inhibit waking from S3, S4, and S5 states if power is not sufficient. It also causes an SMI if the system is already in an S0 state.~~

Table 4-42 Event Transitions that Cause Messages

Event	Assertion?	Deassertion?	Comments
INTRUDER# pin	yes	no	Must be in "S1 or hung S0" state
THRM# pin	yes	yes	Must be in "S1 or hung S0" state. The THRM# pin is isolated when the core power is off, thus preventing this event in S3-S5.
Watchdog Timer Expired	yes	no (NA)	"S1 or hung S0" state entered

Table 4-42 Event Transitions that Cause Messages

GPIO[11]/SMBALERT# pin	yes	yes	Must be in "S1 or hung S0" state
~~BATLOW#~~	~~yes~~	~~yes~~	~~Must be in "S1 or hung S0" state~~
CPU_PWR_FLR	yes	no	"S1 or hung S0" state entered

2. ME Non-Maskable Wake and some causes of host and global resets are not supported in the Intel® Communications Chipset 8900 Series.

The following corrections must be applied to the Intel® Communications Chipset 8900 Series Datasheet.

Table 4-24 Causes of Wake Events

Cause	States Can Wake From	How Enabled
RTC Alarm	S1–S5 (Note 1)	Set RTC_EN bit in PM1_EN register
Power Button	S1–S5	Always enabled as Wake event. (Note 2).
GPI[0:15]	S1–S5 (Note 1)	GPE0_EN register Note: GPIs that are in the core well are not capable of waking the system from sleep states when the core well is not powered.
GPIO[27]	S1-S5	Set GP27_EN in GPE0_EN Register
USB*	S1–S4	Set USB1_EN, USB 2_EN, USB3_EN, USB4_EN, USB5_EN, and USB6_EN bits in GPE0_EN register
RI#	S1–S5 (Note 1)	Set RI_EN bit in GPE0_EN register
Secondary PME#	S1–S5	Set PME_EN bit in GPE0_EN register.
PCI_EXP_WAKE#	S1–S5	PCI_EXP_WAKE bit (Note 3)
SATA*	S1	Set PME_EN bit in GPE0_EN register. (Note 4)
PCI_EXP PME Message	S1	Must use the PCI Express* WAKE# pin rather than messages for wake from S3,S4, or S5.
SMBALERT#	S1–S5	Always enabled as Wake event
SMBus Slave Wake Message (01h)	S1–S5	Wake/SMI# command always enabled as a Wake event. Note: SMBus Slave Message can wake the system from S1–S5, as well as from S5 due to Power Button Override. (Note 2).
SMBus Host Notify message received	S1–S5	HOST_NOTIFY_WKEN bit SMBus Slave Command register. Reported in the SMB_WAK_STS bit in the GPEO_STS register.
~~Intel® ME Non-Maskable Wake~~	~~S1–S5~~	~~Always enabled as a wake event. (Note 2).~~

Notes:

1. This is a wake event from S5 only if the sleep state was entered by setting the SLP_EN and SLP_TYP bits via software, or if there is a power failure.
2. If in the S5 state due to a power button override or THRMTRIP#, the possible wake events are due to Power Button, Wake SMBus Slave Message (01h), ~~the ME-Initiated non-maskable Wake,~~ the Integrated WOL Enable Override, Hard Reset Without Cycling (See Command Type 3), Hard Reset System (See Command Type 4).
3. When the WAKE# pin is active and the PCI Express* device is enabled to wake the system, the system will wake the platform.
4. SATA* can only trigger a wake event in S1, but if PME is asserted prior to S3/S4/S5 entry and software does not clear the PME_B0_STS, a wake event would still result.

Table 4-32 Causes of Host and Global Resets

Trigger	Host Reset without Power Cycle	Host Reset with Power Cycle	Global Reset with Power Cycle
Write of 0Eh to CF9h Register when Global Reset Bit=0b	No	Yes	No (Note 1)
Write of 06h to CF9h Register when Global Reset Bit=0b	Yes	No	No (Note 1)
Write of 06h or 0Eh to CF9h Register when Global Reset Bit=1b	No	No	Yes
SYS_RESET# Asserted and CF9h Bit 3 = 0	Yes	No	No (Note 1)
SYS_RESET# Asserted and CF9h Bit 3 = 1	No	Yes	No (Note 1)
SMBus Slave Message received for Reset with Power Cycle	No	Yes	No (Note 1)
SMBus Slave Message received for Reset without Power Cycle	Yes	No	No (Note 1)
TCO Watchdog Timer reaches zero two times	Yes	No	No (Note 1)
Power Failure: PWROK signal or SYS_PWROK signal goes inactive or RSMRST# asserts	No	No	Yes (Note 2)
Special shutdown cycle from CPU causes CF9h-like PLTRST# and CF9h Global Reset Bit = 1	No	No	Yes
Special shutdown cycle from CPU causes CF9h-like PLTRST# and CF9h Bit 3 = 1	No	Yes	No (Note 2)
Special shutdown cycle from CPU causes CF9h-like PLTRST# and CF9h Global Reset Bit = 0	Yes	No	No (Note 1)
~~Intel® Management Engine Triggered Host Reset without power cycle~~	~~Yes~~	~~No~~	~~No (Note 1)~~
~~Intel® Management Engine Triggered Host Reset with power cycle~~	~~No~~	~~Yes~~	~~No (Note 1)~~
~~Intel® Management Engine Triggered Global Reset~~	~~No~~	~~No~~	~~Yes~~
~~Intel® Management Engine Initiated Host Reset with power down~~	~~No~~	~~Yes (Note 3)~~	~~No (Note 1)~~
PLTRST# Entry Timeout	No	No	Yes
~~Intel® Management Engine Watchdog Timer~~	~~No~~	~~Yes (Note 4)~~	~~No (Note 1)~~
Power Management Watchdog Timer	No	Yes (Note 4)	No (Note 1)
CPUPWRGD Stuck Low	No	No	Yes
CPUPWRGD-to-CPURST# Violation	No	No	Yes

Notes:
1. Trigger results in Global Reset with power cycle if the acknowledge message is not received.
2. The system does not send warning message to CPU, reset occurs without delay.
3. ~~The system waits for enabled wake event to complete reset.~~
4. The system allowed to drop this type of reset request if received while the system is in S3/S4/S5.

Table 6-97 Offset PMBASE+00h: PM1_STS—Power Management 1 Status Register (Sheet 1 of 3)

Description:					
View: PCI	**BAR:** PMBASE (IO)		**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 00h **Offset End:** 00h	
Size: 16 bit	**Default:** 0000h			**Power Well:** Bits 0–7: Core, Bits 8–15: Resume, except Bit 11 in RTC	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
15	WAK_STS	Wake Status — This bit is not affected by hard resets caused by a CF9 write, but is reset by RSMRST#. 0 = Software clears this bit by writing a 1 to it. 1 = Set by hardware when the system is in one of the sleep states (via the SLP_EN bit) and an enabled wake event occurs. Upon setting this bit, the PCH will transition the system to the ON state. If the AFTERG3_EN bit is not set and a power failure (such as removed batteries) occurs without the SLP_EN bit set, the system will return to an S0 state when power returns, and the WAK_STS bit will not be set. If the AFTERG3_EN bit is set and a power failure occurs without the SLP_EN bit having been set, the system will go into an S5 state when power returns, and a subsequent wake event will cause the WAK_STS bit to be set. Any subsequent wake event would have to be caused by either a Power Button press, or an enabled wake event that was preserved through the power failure (enable bit in the RTC well).			RWC
14	PCIEXPWAK_STS	PCI Express* Wake Status: 0 = Software clears this bit by writing a 1 to it. If the WAKE# pin is still active during the write or the PME message received indication has not been cleared in the root port, then the bit will remain active (for example, all inputs to this bit are level-sensitive). 1 = This bit is set by hardware to indicate that the system woke due to a PCI Express* wakeup event. This wakeup event can be caused by the PCI Express* WAKE# pin being active or receipt of a PCI Express* PME message at a root port. This bit is set only when one of these events causes the system to transition from a non-S0 system power state to the S0 system power state. This bit is set independent of the state of the PCIEXP_WAKE_DIS bit. Note: This bit does not itself cause a wake event or prevent entry to a sleeping state. Thus if the bit is 1 and the system is put into a sleeping state, the system will not automatically wake.			RWC
13 :12	Reserved	Reserved			
11	PWRBTNOR_STS	Power Button Override Status: 0 = Software clears this bit by writing a 1 to it. 1 = This bit is set any time a Power Button Override occurs (for example, the power button is pressed for at least 4 consecutive seconds), due to the corresponding bit in the SMBus slave message, Intel® ME Initiated Power Button Override, Intel® ME Initiated Host Reset with Power down or due to an internal thermal sensor catastrophic condition. The power button override causes an unconditional transition to the S5 state, as well as sets the AFTERG3_EN bit. The BIOS or SCI handler clears this bit by writing a 1 to it. This bit is not affected by hard resets via CF9h writes, and is not reset by RSMRST#. Thus, this bit is preserved through power failures. If this bit is still asserted when the global SCI_EN is set then an SCI will be generated.			RWC

Table 6-97 Offset PMBASE+00h: PM1_STS—Power Management 1 Status Register (Sheet 2 of 3)

Description:					
View: PCI	**BAR:** PMBASE (IO)		**Bus:Device:Function:** B0:D31:F0	**Offset Start:** 00h **Offset End:** 00h	
Size: 16 bit	**Default:** 0000h			**Power Well:**	Bits 0–7: Core, Bits 8–15: Resume, except Bit 11 in RTC
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
10	RTC_STS	RTC Status — This bit is not affected by hard resets caused by a CF9 write, but is reset by RSMRST#: 0 = Software clears this bit by writing a 1 to it. 1 = Set by hardware when the RTC generates an alarm (assertion of the IRQ8# signal). Additionally if the RTC_EN bit (PMBASE + 02h, bit 10) is set, the setting of the RTC_STS bit will generate a wake event.			RWC
09	~~ME_STS~~ Reserved	Reserved ~~Management Engine Status — This bit is set when the Intel® Management Engine generates a Non-Maskable wake event, and is not affected by any other enable bit. When this bit is set, the Host Power Management logic wakes to S0.~~ ~~This bit is only set by hardware and can only be reset by writing a one to this bit position. This bit is not affected by hard resets caused by a CF9 write, but is reset by RSMRST#.~~			~~RWC~~
08	PWRBTN_STS	Power Button Status — This bit is not affected by hard resets caused by a CF9 write: 0 = If the PWRBTN# signal is held low for more than 4 seconds, the hardware clears the PWRBTN_STS bit, sets the PWRBTNOR_STS bit, and the system transitions to the S5 state with only PWRBTN# enabled as a wake event. 1 = This bit can be cleared by software by writing a one to the bit position. This bit is set by hardware when the PWRBTN# signal is asserted Low, independent of any other enable bit. In the S0 state, while PWRBTN_EN and PWRBTN_STS are both set, an SCI (or SMI# if SCI_EN is not set) will be generated. In any sleeping state S1–S5, while PWRBTN_EN (PMBASE + 02h, bit 8) and PWRBTN_STS are both set, a wake event is generated. If the PWRBTN_STS bit is cleared by software while the PWRBTN# signal is sell asserted, this will not cause the PWRBN_STS bit to be set. The PWRBTN# signal must go inactive and active again to set the PWRBTN_STS bit.			RWC
07 :06	Reserved	Reserved			
05	GBL _STS	Global Status: 0 = The SCI handler should then clear this bit by writing a 1 to the bit location. 1 = Set when an SCI is generated due to BIOS wanting the attention of the SCI handler. BIOS has a corresponding bit, BIOS_RLS, which will cause an SCI and set this bit.			RWC

Table 6-97 Offset PMBASE+00h: PM1_STS—Power Management 1 Status Register (Sheet 3 of 3)

Description:					
View: PCI	**BAR:** PMBASE (IO)	**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 00h **Offset End:** 00h		
Size: 16 bit	**Default:** 0000h		**Power Well:** Bits 0–7: Core, Bits 8–15: Resume, except Bit 11 in RTC		
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
04	BM_STS	Bus Master Status — This bit will not cause a wake event, SCI or SMI#. 0 = Software clears this bit by writing a 1 to it. Set by the PCH when a PCH-visible bus master requests access to memory or the BM_BUSY# signal is active.			RWC
03 :01	Reserved	Reserved			
00	TO_STS	Timer Overflow Status: 0 = The SCI or SMI# handler clears this bit by writing a 1 to the bit location. 1 = This bit gets set any time bit 22 of the 24-bit timer goes high (bits are numbered from 0 to 23). This will occur every 2.3435 seconds. When the TMROF_EN bit (PMBASE + 02h, bit 0) is set, then the setting of the TMROF_STS bit will additionally generate an SCI or SMI# (depending on the SCI_EN).			RWC

Table 6-101 Offset PMBASE + 20h: GPE0_STS—General Purpose Event 0 Status Register (Sheet 1 of 4)

Description:					
View: PCI	**BAR:** PMBASE (IO)	**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 20h **Offset End:** 20h		
Size: 64 bit	**Default:** 0000000000000000h		**Power Well:** Resume		
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
63 :36	Reserved	Reserved			
35	GPIO27_STS	0 = Disable. 1 = Set by hardware and can be reset by writing a one to this bit position or a resume well reset. This bit is set at the level specified in GP27IO_POL. GPIO27 is always monitored as an input for the purpose of setting this bit, regardless of the actual GPIO configuration.			RWC
34 :32	Reserved	Reserved			

Table 6-101 Offset PMBASE + 20h: GPE0_STS—General Purpose Event 0 Status Register (Sheet 2 of 4)

Description:					
View: PCI	**BAR:** PMBASE (IO)		**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 20h **Offset End:** 20h	
Size: 64 bit	**Default:** 0000000000000000h			**Power Well:** Resume	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
31 :16	GPIOn_STS	0 = Software clears this bit by writing a 1 to it. 1 = These bits are set any time the corresponding GPIO is set up as an input and the corresponding GPIO signal is high (or low if the corresponding GP_INV bit is set). If the corresponding enable bit is set in the GPE0_EN register, then when the GPIO[n]_STS bit is set: - If the system is in an S1–S5 state, the event will also wake the system. - If the system is in an S0 state (or upon waking back to an S0 state), a SCI will be caused depending on the GPIO_ROUT bits (B0:D31:F0:B8h, bits 31:30) for the corresponding GPI. Mapping is as follows: bit 31 corresponds to GPIO[15]... and bit 16 corresponds to GPIO[0].			RWC
15 :14	Reserved	Reserved			
13	PME_B0_STS	This bit will be set to 1 by the PCH when any internal device with PCI Power Management capabilities on bus 0 asserts the equivalent of the PME# signal. Additionally, if the PME_B0_EN bit is set, and the system is in an S0 state, then the setting of the PME_B0_STS bit will generate an SCI (or SMI# if SCI_EN is not set). If the PME_B0_STS bit is set, and the system is in an S1–S4 state (or S5 state due to SLP_TYP and SLP_EN), then the setting of the PME_B0_STS bit will generate a wake event, and an SCI (or SMI# if SCI_EN is not set) will be generated. If the system is in an S5 state due to power button override, then the PME_B0_STS bit will not cause a wake event or SCI. The default for this bit is 0. Writing a 1 to this bit position clears this bit. Note: HD audio wake events are reported in this bit. ~~Intel® Management Engine "maskable" wake events are also reported in this bit.~~			RWC
12	Reserved	Reserved			
11	PME_STS	0 = Software clears this bit by writing a 1 to it. 1 = Set by hardware when the PME# signal goes active. Additionally, if the PME_EN bit is set, and the system is in an S0 state, then the setting of the PME_STS bit will generate an SCI or SMI# (if SCI_EN is not set). If the PME_EN bit is set, and the system is in an S1–S4 state (or S5 state due to setting SLP_TYP and SLP_EN), then the setting of the PME_STS bit will generate a wake event, and an SCI will be generated. If the system is in an S5 state due to power button override or a power failure, then PME_STS will not cause a wake event or SCI.			RWC
10	Reserved	Reserved			

Table 6-101 Offset PMBASE + 20h: GPE0_STS—General Purpose Event 0 Status Register (Sheet 3 of 4)

Description:					
View: PCI	**BAR:** PMBASE (IO)		**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 20h **Offset End:** 20h	
Size: 64 bit	**Default:** 0000000000000000h			**Power Well:** Resume	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
09	PCI_EXP_STS	0 = Software clears this bit by writing a 1 to it. 1 = Set by hardware to indicate that: - The PME event message was received on one or more of the PCI Express* ports - An Assert PMEGPE message received from the Processor via DMI Notes: 1 - The PCI WAKE# pin has no impact on this bit. 2 - If the PCI_EXP_STS bit went active due to an Assert PMEGPE message, then a Deassert PMEGPE message must be received prior to the software write in order for the bit to be cleared. 3 - If the bit is not cleared and the corresponding PCI_EXP_EN bit is set, the level-triggered SCI will remain active. 4 - A race condition exists where the PCI Express* device sends another PME message because the PCI Express* device was not serviced within the time when it must resend the message. This may result in a spurious interrupt, and this is comprehended and approved by the *PCI Express* Specification, Revision 1.0a*. The window for this race condition is approximately 95-105 milliseconds.			RWC
08	RI_STS	0 = Software clears this bit by writing a 1 to it. 1 = Set by hardware when the RI# input signal goes active.			RWC
07	SMB_WAK_STS	SMBus Wake Status — The SMBus controller can independently cause an SMI# or SCI, so this bit does not need to do so (unlike the other bits in this register). Software clears this bit by writing a 1 to it. 0 = Wake event Not caused by the PCH's SMBus logic. 1 = Set by hardware to indicate that the wake event was caused by the PCH's SMBus logic.This bit will be set by the WAKE/SMI# command type, even if the system is already awake. The SMI handler should then clear this bit. Notes: 1. The SMBus controller will independently cause an SMI# so this bit does not need to do so (unlike the other bits in this register). 2. This bit is set by the SMBus slave command 01h (Wake/SMI#) even when the system is in the S0 state. Therefore, to avoid an instant wake on subsequent transitions to sleep states, software must clear this bit after each reception of the Wake/SMI# command or just prior to entering the sleep state. 3. The SMBALERT_STS bit (D31:F3:I/O Offset 00h:Bit 5) should be cleared by software before the SMB_WAK_STS bit is cleared.			RWC
06	TCOSCI_STS	Software clears this bit by writing a 1 to it. 0 = TOC logic or thermal sensor logic did Not cause SCI. 1 = Set by hardware when the TCO logic or thermal sensor logic causes an SCI.			RWC
05 :03	Reserved	Reserved			
02	SWGPE_STS	The SWGPE_CTRL bit (bit 1 of GPE_CTRL reg) acts as a level input to this bit.			RWC

Table 6-101 Offset PMBASE + 20h: GPE0_STS—General Purpose Event 0 Status Register (Sheet 4 of 4)

Description:					
View: PCI	**BAR:** PMBASE (IO)		**Bus:Device:Function:** B0:D31 :F0	**Offset Start:** 20h **Offset End:** 20h	
Size: 64 bit	**Default:** 0000000000000000h			**Power Well:** Resume	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
01	HOT_PLUG_STS	0 = This bit is cleared by writing a 1 to this bit position. 1 = When a PCI Express* Hot-Plug event occurs. This will cause an SCI if the HOT_PLUG_EN bit is set in the GEP0_EN register.			RWC
00	Reserved	Reserved			

Table 7-40 Offset 3310h: PRSTS—Power and Reset Status

Description:					
View: PCI	**BAR:** RCBA		**Bus:Device:Function** B0:D31 **:** :F0	**Offset Start:** 3310h **Offset End:** 3313h	
Size: 32 bit	**Default:** 02020000h			**Power Well:**	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
31 :16	Reserved	Reserved			
15		Power Management Watchdog Timer — This bit is set when the Power Management watchdog timer causes a global reset.			RWC
14 :08	Reserved	Reserved			
07		VE Watchdog Timer Status — This bit is set when the VE watchdog timer causes a global reset.			RWC
06		Intel® Management Engine Watchdog Timer Status — This bit is set when the Intel® Management Engine watchdog timer causes a global reset.			RWC
05 :	Reserved	Reserved			
02	ME_HRST_WARM_STS	ME Host Reset Warm Status — This bit is set when the Intel® Management Engine generates a Host reset without power cycling. Software clears this bit by writing a 1 to this bit position.			RWC
01	ME_HRST_COLD_STS	ME Host Reset Cold Status — This bit is set when the Intel® Management Engine generates a Host reset with power cycling. Software clears this bit by writing a 1 to this bit position.			RWC
00	~~ME_WAKE_STS~~ Reserved	Reserved ~~ME WAKE STATUS — This bit is set when the Intel® Management Engine generates a Non-Maskable wake event, and is not affected by any other enable bit. When this bit is set, the Host Power Management logic wakes to S0.~~			~~RWC~~

3. PCIe Non-Common Clock Mode is not supported in the Intel® Communications Chipset 8900 Series.

The following corrections must be applied to the Intel® Communications Chipset 8900 Series Datasheet.

Table 12-30 Offset 50h: Link Control Register (PCI Express*—B0:D28:F0/F1/F2/F3)

Description:					
View: PCI	**BAR:** Configuration	**Bus:Device:Function:** B0:D28:F0/F1/F2/F3		**Offset Start:** 50h **Offset End:** 51h	
Size: 16 bit	**Default:** 0000h			**Power Well:**	
Bit Range	**Bit Acronym**	**Bit Description**	**Sticky**	**Bit Reset Value**	**Bit Access**
15 :10	Reserved	Reserved			
09		Hardware Autonomous Width Disable – Hardware never attempts to change the link width except when attempting to correct unreliable Link operation.			RO
08	Reserved	Reserved			
07	SE	Extended Synch: 0 = Extended synch disabled. 1 = Forces extended transmission of FTS ordered sets in FTS and extra TS2 at exit from L1 prior to entering L0.			RW
06	CCC	Common Clock Configuration: 0 = ~~The PCH and device are not using a common reference clock.~~ Reserved. 1 = The PCH and device are operating with a distributed common reference clock. Note: This bit field must be set.			RW
05	RL	Retrain Link: 0 = This bit always returns 0 when read. 1 = The root port will train its downstream link. Software uses LSTS.LT (B0:D28:F0/F1/F2/F3:52, bit 11) to check the status of training. It is permitted to write 1b to this bit while simultaneously writing modified values to other fields in this register. If the LTSSM is not already in Recovery or Configuration, the resulting Link training must use the modified values. If the LTSSM is already in Recovery or Configuration, the modified values are not required to affect the Link training that is already in progress.			RW
04	LD	Link Disable: 0 = Link enabled. 1 = The root port will disable the link.			RW
03	RCBC	Read Completion Boundary Control) — Indicates the read completion boundary is 64 bytes.			RO
02	Reserved	Reserved			
01 :00	APMC	Active State Link PM Control — Indicates whether the root port should enter L0s or L1 or both. Bits Definition 00 Disabled 01 L0s Entry Enabled 10 L1 Entry Enabled 11 L0s and L1 Entry Enabled			RW

4. GPIO [5:2] outputs are Open Drain.

The following corrections must be applied to the Intel® Communications Chipset 8900 Series Datasheet.

Table 32-22 General Purpose I/O Interface Signals (Sheet 1 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
BMBUSY#/ GPIO0	I/O	BMBUSY#	BT45	1	I	CORE		**Bus Master Busy.** • This signal is used to support the C3 state. It indicates that a bus master device is busy. When this signal is asserted, the BM_STS bit will be set. If this signal goes active in a C3 state, it is treated as a break event. ***Note:*** This signal is internally synchronized using the PCICLK and a two-stage synchronizer. It does not need to meet any particular setup or hold time. This signal can also be used as GPIO Port 0.
GPIO1	I/O	GPI	BN47	1	I	CORE		General Purpose I/O Port 1.
GPIO2 [a]	I/O	GPI	BK43	1	I	CORE		General Purpose I/O Port 2.
GPIO3 [a]	I/O	GPI	BR42	1	I	CORE		General Purpose I/O Port 3.
GPIO4 [a]	I/O	GPI	BJ43	1	I	CORE		General Purpose I/O Port 4.
GPIO5 [a]	I/O	GPI	BM45	1	I	CORE		General Purpose I/O Port 5.
GPIO6	I/O	GPI	BP53	1	I	CORE		General Purpose I/O Port 6.
GPIO7	I/O	GPI	BT47	1	I	CORE		General Purpose I/O Port 7.
GPIO8	I/O	GPO	BU16	1	O (High)	SUS	Weak Internal pull-up for strap.	General Purpose I/O Port 8. • This signal has a weak internal pull-up and must not be pulled low during boot up.
GPIO9[b]	I/O	Native	BN21	1	I	SUS		General Purpose I/O Port 9[c]
GPIO10	I/O	Native	BG22	1	I	SUS		General Purpose I/O Port 10.
MST_SMBALERT#/ GPIO11	I/O	MST_SMBALERT#	BG15	1	I	SUS	External pull-up required	**Host SMBus Alert.** • This signal is used to wake the system or generate SMI#. • External pull-up resistor to VCCSUS3P3 is required. • Resistor value should be calculated based on the bus load, (See the platform design collateral for Resistor value) This signal can also be configured to GPIO Port 11.

Table 32-22 General Purpose I/O Interface Signals (Sheet 2 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
GPIO12	I/O	Native	BF15	1	O (Low)	SUS		General Purpose I/O Port 12.
GPIO13[d]	I/O	GPI	BJ25	1	I	SUS	Weak internal pull-down	General Purpose I/O Port 13.
GPIO14	I/O	Native	BP23	1	I	SUS		General Purpose I/O Port 14.
GPIO15	I/O	GPO	BJ10	1	O (Low)	SUS		General Purpose I/O Port 15.
SATA4_GP/ GPIO16	I/O	GPI	BC47	1	I	CORE		**Serial ATA 4 General Purpose.** • This is an input pin which can be configured as an interlock switch corresponding to SATA Port 4. When used as an interlock switch status indication, this signal should be drive to '0' to indicate that the switch is closed and to '1' to indicate that the switch is open. If interlock switches are not required, this pin can be configured as GPIO Port 16.
GPIO17	I/O	GPI	BF38	1	I	CORE	Platform Dependent	**DMI Voltage Strap.** • This signal strapping sets the DMI termination voltage. • See PDG for additional information. This signal can be used as a GPIO Port 17
GPIO18[e]	I/O	Native	BG56	1	I	CORE		General Purpose I/O Port 18.
GPIO19	I/O	GPI	BF43	1	I	CORE		General Purpose I/O Port 19.
GPIO20	I/O	Native	BC50	1	I	CORE		General Purpose I/O Port 20.
GPIO21	I/O	GPI	BF49	1	I	CORE		General Purpose I/O Port 21.
SCLOCK/ GPIO22	I/O	GPI	BR49	1	I	CORE		**SGPIO Reference Clock.** • The SATA controller uses rising edges of this clock to transmit serial data, and the target uses the falling edge of this clock to latch data. If SCLOCK interface is not used, this signal can be used as a GPIO Port 22.

Table 32-22 General Purpose I/O Interface Signals (Sheet 3 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
LDRQ1#/ GPIO23	I/O	LDRQ1#	BG33	1	I	CORE		**LPC Serial DMA/Master Request Input Bit 1**. • Used by LPC devices, such as Super I/O chips, to request DMA or bus master access. This signal is typically connected to external Super I/O device. An internal pull-up resistor is provided on these signals. If LDRQ1# interface is not used, this signal can be used as a GPIO Port 23.
GPIO24	I/O	GPO	BP7	1	O (Low)	SUS		General Purpose I/O Port 24.
GPIO25	I/O	Native	BK25	1	I	SUS		General Purpose I/O Port 25.
GPIO26	I/O	Native	BM26	1	I	SUS		General Purpose I/O Port 26.
GPIO27	I/O	GPO	BR6	1	O (Low)	SUS		General Purpose I/O Port 27.
GPIO28	I/O	GPI	BG17	1	I	SUS		General Purpose I/O Port 28.
GPIO30	I/O	GPI	BM10	1	I	SUS	Internal pull-down	General Purpose I/O Port 30.
GPIO31	I/O	GPI	BK15	1	I	SUS		General Purpose I/O Port 31.
GPIO32	I/O	Native	BG38	1	O (High)	CORE		General Purpose I/O Port 32.
GPIO33	I/O	GPO	BN41	1	O (High)	CORE	Weak Internal pull-up	**Flash Descriptor Security Overwrite**. • This signal is used to set the security override strap on the PCH. If sampled low, the Flash Descriptor Security will be overridden. If high, the security measures defined in the Flash Descriptor will be in effect. This strap should only be enabled (pulled low) in manufacturing environments using an external pull-down resistor. **GPIO33** 0 = Enable (Pull-Down Required) 1 = Disable (Default) • When the Security Overwrite is enable, it allows permission to every master to read and write to the entire Flash Components including areas outside the defined regions.
GPIO34	I/O	GPI	BN49	1	I	CORE		General Purpose I/O Port 34.
GPIO35	I/O	GPO	BF46	1	O (Low)	CORE		General Purpose I/O Port 35.
GPIO36	I/O	GPI	BN54	1	I	CORE		General Purpose I/O Port 36.

Table 32-22 General Purpose I/O Interface Signals (Sheet 4 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
ADR/ GPIO37	I/O	GPI	BG53	1	I	CORE		General Purpose I/O Port 37. • Can be used for ADR (Asynchronous DRAM Refresh) trigger on platform. Only supported if processor supports ADR. • HARDWARE activation mechanism that triggers memory controller of CPU to put SDRAM into self-refresh mode. Activation of ADR flushes contents of some write data buffers to the DIMM before self refresh entry.
SLOAD/ GPIO38	I/O	GPI	BU49	1	I	CORE		**SATA Serial GPIO Load**. • The controller drives a '1' at the rising edge of SCLOCK to indicate either the start or end of a bit stream. A 4-bit vendor specific pattern will be transmitted right after the signal assertion. If SLOAD interface is not used, this signal can be used as a GPIO Port 38.
SDATAOUT 0/GPIO39	I/O	GPI	BR52	1	I	CORE		**SATA Serial GPIO Data Out 0.** • Driven by the controller to indicate the drive status in the following sequence: drive 4, 5...4,5... If SDATAOUT0 interface is not used, the signals can be used as GPIO Port 39.
OC1#/ GPIO40	I/O	OC1#	BN4	1	I	SUS		**Overcurrent Indicators.** • These signals set corresponding bits in the USB* controllers to indicate that an overcurrent condition has occurred. • OC[3:0]# may optionally be used as GPIO Ports [42,41,40,59]. ***Notes:*** 1. OC# pins are 3.3V and NOT 5 V tolerant. 2. OC# pins must be shared between ports 3. OC#[3:0] can only be used for EHCI controller #1

Table 32-22 General Purpose I/O Interface Signals (Sheet 5 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
OC2#/ GPIO41	I/O	OC2#	BK10	1	I	SUS		**Overcurrent Indicators.** • These signals set corresponding bits in the USB* controllers to indicate that an overcurrent condition has occurred. • OC[3:0]# may optionally be used as GPIO Ports [42,41,40,59]. ***Notes:*** 1. OC# pins are 3.3V and NOT 5 V tolerant. 2. OC# pins must be shared between ports 3. OC#[3:0] can only be used for EHCI controller #1
OC3#/ GPIO42	I/O	OC3#	BM18	1	I	SUS		**Overcurrent Indicators.** • These signals set corresponding bits in the USB* controllers to indicate that an overcurrent condition has occurred. • OC[3:0]# may optionally be used as GPIO Ports [42,41,40,59]. ***Notes:*** 1. OC# pins are 3.3V and NOT 5 V tolerant. 2. OC# pins must be shared between ports 3. OC#[3:0] can only be used for EHCI controller #1
GPIO43	I/O	Native	BK22	1	I	SUS		General Purpose I/O Port 43.
GPIO44	I/O	Native	BT23	1	I	SUS		General Purpose I/O Port 44.
GPIO45	I/O	Native	BT21	1	I	SUS		General Purpose I/O Port 45.
GPIO46	I/O	Native	BM20	1	I	SUS		General Purpose I/O Port 46.
GPIO47	I/O	Native	BU9	1	I	SUS		General Purpose I/O Port 47.
SDATAOUT 1/ GPIO48	I/O	GPI	BT51	1	I	CORE		**SATA Serial GPIO Data Out 1**. • Driven by the controller to indicate the drive status in the following sequence: drive 4, 5...4,5... If SDATAOUT1 interface is not used, the signals can be used as GPIO Port 48.

Table 32-22 General Purpose I/O Interface Signals (Sheet 6 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
SATA5_GP/ TEMP_ALERT#/ GPIO49	I/O	GPI	BC46	1	I	CORE		**Serial ATA 5 General Purpose.** • This is an input pin which can be configured as an interlock switch corresponding to SATA Port 5. When used as an interlock switch status indication, this signal should be drive to '0' to indicate that the switch is closed and to '1' to indicate that the switch is open. **Temperature Alert.** • Used as an alert (active low) to indicate to the external controller (such as, EC or SIO) that temperatures are out of range for the PCH or Memory Controller or the processor core. If interlock switches or Temp Alert are not required, this pin can be configured as GPIO Port 49.
GPIO50	I/O	Native	BG43	1	I	CORE		General Purpose I/O Port 50.
BBS1/ GPIO51	I/O	BBS1	BM38	1	O (High)	CORE	Weak Internal pull-up	**BIOS Boot Strap 1**. BBS1 BBS0 0 0 NOT VALID 0 1 NOT VALID 1 0 NOT VALID 1 1 SPI BOOT If BBS1 interface is not used, the signals can be used as GPIO Port 51.
GPIO52	I/O	Native	BF36	1	I	CORE		General Purpose I/O Port 52. Not Multiplexed.
GPIO53	I/O	Native	BM42	1	O (High)	CORE	Platform Dependent. Weak internal pull-up	**DMI Coupling Strap.** 0 = AC Coupling (Pull-Down Required) 1 = DC Coupling (Default) • See PDG for additional information.
GPIO54	I/O	Native	BR46	1	I	CORE		General Purpose I/O Port 54.

Table 32-22 General Purpose I/O Interface Signals (Sheet 7 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
GPIO55	I/O	Native	BP45	1	O (High)	CORE	Weak Internal pull-up	**BIOS Boot-Block Swap** • This mode allows the PCH to swap the Top-Block in the SPI (the boot block) with another location. **GPIO55** 0 = Enable Top-Block Swap (Pull-down Required) 1 = Disable Top-Block Swap (Default) ***Note:*** The internal pull-up is disabled after PLTRST# de-asserts. If the signal is sampled low, this indicates that the system is strapped to the "Top-Block Swap" mode (the PCH inverts A16 for all cycles targeting BIOS space). • The status of this strap is readable via the Top Swap bit (Chipset Config Registers:Offset 3414h:bit 0). Software will not be able to clear the Top-Swap bit until the system is rebooted without GPIO55 being pulled down.
GPIO56	I/O	Native	BR9	1	I	SUS		General Purpose I/O Port 56.
GPIO57	I/O	GPI	BN11	1	I	SUS		General Purpose I/O Port 57.
SML1CLK/ GPIO58	I/O	SML1CLK	BJ20	1	I	SUS	When used as SMBUS: External pull-up required[f]	**System Management Link 1 Clock:** • SMBus link to external BMC. • External pull-up resistor to VCCSUS3P3 is required. • Resistor value should be calculated based on the bus load, see the platform design collateral. If SML1CLK interface is not used, the signals can be used as GPIO Port 58.

Table 32-22 General Purpose I/O Interface Signals (Sheet 8 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
OC0#/ GPIO59	I/O	OC0#	BK12	1	I	SUS		**Overcurrent Indicators.** • These signals set corresponding bits in the USB* controllers to indicate that an overcurrent condition has occurred. • OC[3:0]# may optionally be used as GPIO Ports [42,41,40,59]. ***Notes:*** 1. OC# pins are 3.3V and NOT 5 V tolerant. 2. OC# pins must be shared between ports 3. OC#[3:0] can only be used for EHCI controller #1
GPIO60	I/O	Native	BR12	1	I	SUS		General Purpose I/O Port 60.
SUS_STAT #/ GPIO61	I/O	SUS_STAT#	BM13	1	O (High)	SUS		**Suspend Status**: • This signal is asserted by the PCH to indicate that the system will be entering a low power state soon. This can be monitored by devices with memory that need to switch from normal refresh to suspend refresh mode. It can also be used by other peripherals as an indication that they should isolate their outputs that may be going to powered-off planes. If SUS_STAT# interface is not used, this signal can be used as a GPIO Port 61.
SUS_CLK/ GPIO62	I/O	SUS_CLK	BP13	1	O (Low)	SUS		**Suspend Clock.** • This clock is an output of the RTC generator circuit. It is used by other chips for refresh clock. If SUS_CLK interface is not used, the signals can be used as GPIO Port 62.
SLP_S5#/ GPIO63	I/O	SLP_S5#	BF20	1	O (High)	SUS		**S5 Sleep Control.** • SLP_S5# is for power plane control. • This signal is used to shut power off to all non-critical systems when in the S5 (Soft Off) states. If SLP_S5# interface is not used, the signals can be used as GPIO Port 63.
GPIO72	I/O	Native	BT13	1	I	SUS		General Purpose I/O Port 72.

Table 32-22 General Purpose I/O Interface Signals (Sheet 9 of 9)

Intel® Communications Chipset 89xx Series								
DH89xxCC: (8900 ≤ *SKU* ≤ 8920) DH89xxCL: (8925 ≤ *SKU* ≤ 8955)								
Signal Name **DH89xxCC & DH89xxCL**	**I/O Type**	**Default Mode**	**Ball Map**	**Ball Count**	**Default Direction and Logic State**	**Power Well**	**Internal/ External Resistor Pull-Up/ Down**	**Description**
GPIO73	I/O	Native	BR22	1	I	SUS		General Purpose I/O Port 73.
SML1ALERT #/ GPIO74	I/O	SML1ALERT#	BU12	1	I	SUS	When used as SMBUS: External pull-up required.	**System Management Link 1 Alert.** • This signal can be connected to an external BMC. • External pull-up resistor to VCCSUS3P3 is required. • Resistor value should be calculated based on the bus load, see the platform design collateral. If SML1ALERT# interface is not used, the signals can be used as GPIO Port 74.
SML1DAT/ GPIO75	I/O	SML1DAT	BK20	1	I	SUS	When used as SMBUS: External pull-up required.	**System Management Link 1 Data.** • SMBus link to external BMC. • External pull-up required to VCCSUS3P3 is required. • Resistor value should be calculated based on the bus load, see the platform design collateral. If SML1DAT interface is not used, the signals can be used as GPIO Port 75.
TOTAL				67				

a. When this signal is configured as GPO, the output stage is an open drain.
b. When the multiplexed GPIO is used as GPIO functionality, care should be taken to ensure the signal is stable in its inactive state of the native functionality, immediately after reset until it is initialized to GPIO functionality. Multiplexed signals is visible or Intel reserved.
c. For GPIOs where Native Mode is configured using SPI Soft Strap, the corresponding GPIO_USE_SEL bits for these GPIOs have no effect. The GPIO_USE_SEL bits for these GPIOs may change to reflect the Soft-Strap configuration even though GPIO Lockdown Enable (GLE) bit is set.
d. The functionality that is multiplexed with the GPIO may not be used in desktop configuration.
e. GPIO18 is configured as an input in default mode. GPIO18 will toggle at a frequency of approximately 1 Hz when the signal is programmed as an output (via GP_IO_SEL Register) by BIOS or system configuration.
f. See the platform design collateral for resistor values.

GPIO[29], GPIO[71:64], and GPIO[95:76] do not exit or are reserved and cannot be used.

§ §