

Intel[®] X58 Express Chipset

Specification Update

| *October 2009*

Notice: The Intel[®] X58 Express Chipset IOH may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

| Doc Number: 320839-010



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® X58 Express Chipset may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

I²C is a two-wire communications bus/protocol developed by Philips. SMBus is a subset of the I²C bus/protocol and was developed by Intel. Implementations of the I²C bus/protocol may require licenses from various entities, including Philips Electronics N.V. and North American Philips Corporation.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain computer system software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

Intel, Intel® Virtualization Technology for Directed I/O (Intel® VT-d), and the Intel logo are trademarks of Intel Corporation in the United States and other countries.


*Other names and brands may be claimed as the property of others.

Copyright © 2008-2009, Intel Corporation. All Rights Reserved.



Contents


Preface	5
Affected Documents	5
Nomenclature	5
Summary Tables of Changes	6
Codes Used in Summary Tables	6
Stepping	6
Status	6
Row	6
Errata	7
Specification Changes	9
Specification Clarifications	9
Documentation Changes	9
Identification Information	10
Component Identification via Programming Interface	10
Component Marking Information	10
Errata	11
Specification Changes	27
Specification Clarifications	28
Documentation Changes	29



Revision History

Revision	Description	Date
-001	Initial Release.	November 2008
-002	Add errata 38-45	December 2008
-003	Added errata 46-49	January 2009
-004	Added B3 stepping and errata 51-60	March 2009
-005	Updated Erratum 31	April 2009
-006	Added Errata 61-65	May 2009
-007	Removed Erratum 1 Added Errata 66 & 67	June 2009
-008	Added Errata 68 & 69	July 2009
-009	Added Erratum 70 Added Specification Clarification 1	August 2009
-010	Added Specification Clarification 2	October 2009

§



Preface

This document is an update to the specifications contained in the [Affected Documents](#) table below. This document is a compilation of device errata and documentation corrections, 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.

Affected Documents

Document Title	Document Number/ Location
Intel® X58 Express Chipset Datasheet	320838-001

Nomenclature

Errata are design defects or errors. These may cause the Intel® X58 Express Chipset IOH's behavior 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 on 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.

Note: Errata remain in the specification update throughout the product's lifecycle, or until a particular stepping is no longer commercially available. Under these circumstances, errata removed from the specification update are archived and available upon request. Specification changes, specification clarifications and documentation changes are removed from the specification update when the appropriate changes are made to the appropriate product specification or user documentation (datasheets, manuals, and so forth).



Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the IOH 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 use the following notations:

Codes Used in Summary Tables

Stepping

- X: An erratum 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.

Status

- Doc: Document change or update will be implemented.
- Plan 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.



Errata (Sheet 1 of 3)

Number	Steppings		Status	Errata
	B2	B3		
1	X	X	No Fix	PCIe* Link bandwidth permanently set.
2	X	X	No Fix	CPURST bit does not get cleared by hardware
3			Fixed	PCIe RX L0s entry/exit issues.
4	X	X	No Fix	V _{TX-RCV-DETECT} pulse too large during receiver detection.
5			Fixed	PCIe Gen2 differential peak-to-peak transmit voltage swing is too low
6	X	X	No Fix	Testlo[20] must be pulled high.
7	X	X	No Fix	Surprise down error status not being flagged
8	X	X	No Fix	Extended Error Detect Mask Registers of all PCIe root ports mask error logging by default.
9	X	X	No Fix	Long XOR chain is broken
10	X	X	No Fix	PCIe Header of a malformed TLP is not logged
11	X	X	No Fix	PCIe PMCSR Power State fields allow writing D1 and D2
12	X		No Fix	DEVCON2[3:0] and CTOCTRL not accurately set the completion timeout value.
13			Fixed	In Intel® QuickPath Interconnect L1 Power management mode, Link layer sends retryable flits after L1.Ack.
14	X	X	No Fix	PCIe Gen2 Tx Return loss fails spec
15	X	X	No Fix	Persistent Jtag error reported at MIERRST register
16	X	X	No Fix	Interop issue of some PCIe Gen1 cards with Gen2 Devices
17			Fixed	PCIe ASPM L1 tests reporting lost credits
18	X	X	No Fix	Intel QuickPath Interconnect fails to retrain to L0 slow mode after 32 in-band resets
19			Fixed	ATS address field in the translation completion returned in reversed byte order
20			Fixed	SLTCON bit 11 returning wrong value when read
21	X		Fixed	Boot-dependent Intel QuickPath Interconnect CRC errors
22	X	X	No Fix	PCIe Inbound Msg that should be ignored is treated as Unsupported Request
23	X	X	Doc	Memory writes to a certain address range are considered advisory non-fatal
24			Fixed	Inbound Write fetches may not make forward progress
25	X	X	No Fix	PCIe lane-to-lane performance variation
26	X	X	No Fix	Transactions to addresses above TOCM are not setting the Master Abort
27			Fixed	SERR_EN ignored during unsupported request error escalation
28	X	X	No Fix	Setting bits 24 and 25 of the MISCCTRLSTS does not result in expected behavior
29	X	X	No Fix	Timeout values much larger than specified
30			Fixed	XPGLBERRPTR not accurately indicating first error in XPGLBERRSTS
31	X	X	No Fix	Failure during operation at PCI Express* L1 power management state
32	X	X	No Fix	Bandwidth very low for write traffic with noSnoop attribute set
33	X	X	No Fix	Intel QuickPath Interconnect Queue/Table overflow or underflow error observed.
34			Fixed	IOH Course Grained Clock Gating requires an ITP
35	X		Non-Si	Some Gen2 endpoints will not complete the training in Gen2 mode.



Errata (Sheet 2 of 3)

Number	Steppings		Status	Errata
	B2	B3		
36			Fixed	Setting bit 10 of PCICMD register having effect on legacy interrupts
37	X	X	No Fix	ACS Violation is not treated as Advisory when severity is set to Non-Fatal
38	X		Fixed	MIERRCNT does not properly count persistent SMBus retry failures and MINNERRST doesn't log the errors (HEDT only)
39	X	X	No Fix	OPI L1 state is taking greater than 15 us from L1 exit to L0 state
40	X		Fixed	EOI to the I/OxAPIC can be blocked
41	X		Fixed	GTIME upper 32 bits can not be read or written.
42	X		Fixed	Intel® QuickPath Interconnect (QPI) L0s and L1 Power Management Link State Fails
43	X		Fixed	PCI Express* 2.0 L0s Link Recovery Fails
44	X		Fixed	IOH May Falsely Assert THERMTRIP_N Signal After a Reset Event
45	X		No Fix	Intel VT-d queue-based invalidation is enabled only when enabled on both channels (WS only)
46	X		No Fix	Intel VT-d does not support the draining of compatibility-format interrupts (WS only)
47	X		No Fix	Hardware applies HPA_LIMIT to upstream memory request when Intel VT-d is disabled (WS only)
48	X		No Fix	Intel VT-d: Memory read request with AT=11 results in malformed TLP
49			Fixed	Intel VT-d: IOTLB Domain-page-selective invalidation not working correctly.
50	X	X	No Fix	Intel VT-d translated write transactions are blocked but not recorded.
51	X		Fixed	The IOH Ignores Snoop Behavior Bit in VT-d Page Table.
52	X	X	No Fix	Intel QuickPath InterconnectLink Training failures in L0 when L0r enabled
53	X		Fixed	Intel VT-d reports an Invalidation Queue Error when the Queue size is 7 (HEDT only)
54	X	X	No Fix	ERRSID not logging ReqID for Inbound ERR_* messages
55	X	X	No Fix	Intel QuickPath Interconnect PhyResets causing Intel QuickPath Interconnect Errors to be flagged (HEDT only)
56	X	X	No Fix	Intel QuickPath Interconnect initialization abort failures logged during power-on resets
57	X	X	No Fix	Intel QuickPath Interconnect errors can occur on inband resets
58	X	X	No Fix	Data Mismatch on Inbound MemWrts after MSI with payload greater than 1 DWORD payload
59	X	X	No Fix	MSI with greater than 1DWord payload is not logged in XPUNCERRSTS bit 8
60	X	X	No Fix	PPF Flag (due to Intel VT-d ISOCH fault) in Fault Status Register not being cleared
61	X	X	No Fix	Lost interrupts when MSI used
62	X	X	No Fix	Intel VT-d: UP Workstation ONLY. Receiving two identical interrupt requests in back to back cycles may corrupt attributes of remapped interrupt, or hang subsequent interrupt-remap-cache invalidation command.
63	X	X	No Fix	Intel QuickPath Interconnect Error Status D3 is observed
64	X	X	No Fix	Unpredictable PCI behavior accessing non-existent memory space
65	X	X	No Fix	Bandwidth changed status errors being escalated to Global RAS
66	X	X	No Fix	TXT writes may not complete as expected



Errata (Sheet 3 of 3)

Number	Steppings		Status	Errata
	B2	B3		
67	X	X	No Fix	IOTLB Invalidation not completing on Intel VT-d Isochronous unit (UP WS ONLY)
68	X	X	No Fix	Intel® VT-d: Address remapping error when DMA/interrupt remapping is active
69	X	X	No Fix	Intel® VT-d: In-flight remap-able interrupts not drained on interrupt invalidation command
70	X	X	No Fix	Source ID for errors internally detected by PCIE root port 3 is not logged as expected

Specification Changes

Number	SPECIFICATION CHANGES
	There are no specification changes in this revision of the specification update

Specification Clarifications

No.	SPECIFICATION CLARIFICATIONS
1	Non-Posted Peer-to-Peer Support between IOH and ICH

Documentation Changes

No.	DOCUMENTATION CHANGES
1	Documentation update to the Intel® X58 Express Chipset Datasheet 320838-001
2	Chapter 17, Configuration Register Space - ERRCAPHDR



Identification Information

Component Identification via Programming Interface

The Intel X58 Express Chipset (IOH) stepping can be identified by the following register contents:


Stepping	Features	Vendor ID	Device ID	Revision Number
B2	36S	8086h	3405h	12h
B3	36S	8086h	3405h	13h

Component Marking Information

The Intel® X58 Express Chipset (IOH) production stepping can be identified by the following component markings:

Stepping	S-Spec	Top Marking	Notes
B2	SLGBT	AC82X58	Production Units
B3	SLGMX	AC82X58 SLGMX 901076	Production Units

§



Errata

2. CPURST bit does not get cleared by hardware.

Problem: SYRE: CPURST bit does not get cleared by hardware.

Background

SYRE: System Reset - This register controls IOH reset behavior.

CPURESET (SYRE bit 10) - When this bit is set to "1", the IOH asserts RESET0_N. The IOH clears this bit when CPURESET timer elapses.

Implication: The IOH will not assert RESET0_N when CPURST is set if it has been set previously.

Workaround: The CPURST bit must be cleared prior to setting it.

Status: For the steppings affected, see the Summary Tables of Changes.

3. PCIe* RX L0s entry/exit issues.

Problem: If tx-L0s is enabled in Gen2 speed for endpoint devices directly connected to the Intel X58 Express Chipset ports, the IOH may fail to properly exit rx-L0s. Transitions to RECOVERY will occur and several types of errors including, but not limited to, receiver errors, bad TLPs and bad DLLPs may be logged.

Implication: May cause unexpected correctable errors on PCIe links.

Workaround: Disable tx-L0s in endpoints directly attached to the Intel X58 Express Chipset ports.

Status: For the steppings affected, see the Summary Tables of Changes.

4. VTX-RCV-DETECT pulse too large during receiver detection.

Problem: The VTX-RCV-DETECT pulse has been measured as high as 700 mV.

Background

VTX-RCV-DETECT - The amount of voltage change allowed during Receiver Detection. The maximum is 600 mV for both 2.5 GT/s and 5 GT/s.

Implication: This may overstress PCIe agents.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

5. PCIe Gen2 differential peak-to-peak transmit voltage swing is too low.

Problem: The PCIe Gen2 transmit buffers are not generating a large enough peak-to-peak transmit voltage swing, thus violating the PCIe Gen2 specification.

Implication: The PCIe Gen2 specification is violated.

Workaround: Fixed in B0. BIOS changes are required as part of the fix as outlined below:

Dev 13:Func 3

1. Set Offset 330h [15:12] to 5, [11:8] to 3.

Dev 14:Func 0

1. Set Offset 330h [15:12] to 5, [11:8] to 3.



Dev 13:Func 0

1. Set Offset 1A0h [23:16] to 5, [7:0] to FFh.
2. Set Offset 11Ch [28:25] to n, (n = 0)
3. For each n setting in step 2, set Offset 1B4h [11:8] to 8, [7:4] to 6, [3:0] to 5.

Dev 13:Func 1

1. Set Offset 1A0h [23:16] to 5, [7:0] to FFh.
2. Set Offset 11Ch [28:25] to n, (n = 0, 1)
3. For each n setting in step 2, set Offset 1B4h [11:8] to 8, [7:4] to 6, [3:0] to 5.

Dev 13:Func 3

1. Set Offset 1A0h [23:16] to 5, [7:0] to FFh.
2. Set Offset 11Ch [28:25] to n, (n = 0, 1, 2, 3)
3. For each n setting in step 2, set Offset 1B4h [11:8] to 8, [7:4] to 6, [3:0] to 5.

Dev 14:Func 0

1. Set Offset 1A0h [23:16] to 5, [7:0] to FFh.
2. Set Offset 11Ch [28:25] to n, (n = 0, 1, 2, 3)
3. For each n setting in step 2, set Offset 1B4h [11:8] to 8, [7:4] to 6, [3:0] to 5.

Note: The above configuration sequence needs to be repeated in the S3 state resume path.

Status: For the steppings affected, see the Summary Tables of Changes.

6. Testlo[20] must be pulled high.

Problem: The Testlo[20] signal is inverted.

Implication: If the workaround is not implemented, the IOH will not correctly function.

Workaround: The Testlo[20] pin must be pulled to the VCCEPW rail via a 10K ohm $\pm 1\%$ resistor.

Status: For the steppings affected, see the Summary Tables of Changes



7. Surprise down error status not being flagged.

Problem: The Surprise down error status is not being flagged during training failure or surprise removal.

Background

UNCSTS: Uncorrectable Error Status - This register identifies uncorrectable errors detected for PCI Express/DMI port.

Surprise Down Error Status (UNCSTS bit 5) - Flags the system when a training failure or surprise removal has occurred.

Implication: The system will not be able to reliably detect if a PCIe card has not trained properly or been inadvertently removed unless the workaround is used.

Workaround: BIOS fix

Status: For the steppings affected, see the Summary Tables of Changes.

8. Extended Error Detect Mask Registers of all PCIe root ports mask error logging by default.

Problem: The mask fields of all PCIe root ports (devices 1 - 10) Extended PCIe Error Detect Mask Registers are set to 1 by default.

Implication: All PCIe Advanced error status logging registers have a corresponding error detect mask register to control if error statuses will be logged. Below are the error status/error detect mask pairs:

Uncorrectable Error Status (offset 104h) and Detect Status Mask (offset 218h).

Correctable Error Status (offset 110h) and Detect Status Mask (offset 21Ch).

Root Port Error Status (offset 130h) and Detect Status Mask (offset 220h).

XP Correctable Error Status (offset 200h) and Detect Mask (offset 228h).

XP Uncorrectable Error Status (offset 208h) and Detect Mask (offset 224h).

With the default values of these error detect mask registers, no PCIe advanced errors will be logged and reported.

Workaround: BIOS must clear registers 218h, 21Ch, 220h, 228h, and 224h to zero in order to log and report corresponding errors.

Status: For the steppings affected, see the Summary Tables of Changes.

9. Long XOR chain is broken.

Problem: The long XOR chain used for high volume manufacturing test of the PCIe interface is not functional.

Implication: Tests dependent on using the long XOR chain for high volume manufacturing test will not function

Workaround: The individual XOR chains for each PCIe port are operable and can be used to cover those areas covered by the long XOR chain.

Status: For the steppings affected, see the Summary Tables of Changes.

**10. PCIe Header of a malformed TLP is not logged.**

Problem: A malformed Transaction Layer Packet (TLP) is logged in the UNCERRSTS register (Device:0-10, Function:0, Offset:0x104h) but the HDRLOG register (Header Log register, Device:0-10, Function:0, Offset 11Ch) do not log the header of the malformed TLP.

Implication: Given that the PCIe specifies that the header of the malformed TLP must be logged, the PCIe specification is violated.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

11. PCIe PMCSR Power State fields allow writing D1 and D2.

Problem: The PMCSR (Devices: 0 to 10, Function: 0, Offset: E4h) bits 1:0 allow states D1 and D2 to be written.

Implication: Given that the IOH does not support the D1 and D2 states, the IOH should not allow these values to be written. The IOH does not change power state from D0 or D3hot when PMCSR bits 1:0 are written to D1 or D2, so there is no functional impact to the IOH when these states are written.

Workaround: Do not write states D1 and D2 to the PMCSR bits 1:0.

Status: For the steppings affected, see the Summary Tables of Changes.

12. DEVCON2[3:0] and CTOCTRL not accurately set the completion timeout value.

Problem: The DEVCON2 (device: 0-10, function: 0, offset: B8h) bits 3:0 in conjunction with CTOCTRL (device: 0-10, function: 0, offset: 1E0h) bits 9:8 does not accurately set the completion timeout value.

Implication: The duration of the PCIe timeout will differ from that programmed.

Workaround: See the documentation change for 3319486.

Status: For the steppings affected, see the Summary Tables of Changes.

13. In Intel® QuickPath Interconnect L1 Power management mode, Link layer sends retryable flits after L1.Ack.

Problem: The Intel X58 Express Chipset Intel QuickPath Interconnect Link layer is sending packets during the window between the Intel X58 Express Chipset sending L1.Ack and the time it receives Inband Reset from Nehalem processor.

Implication: This causes a system hang as these packets after L1.Ack will be dropped by Nehalem processor. As per Intel QuickPath Interconnect Spec, Link layer should not send any packets after L1.Ack.

Workaround: Workaround is implemented in the Intel QuickPath Interconnect Initialization Reference Code (Intel QuickPath Interconnect RC) v0.75 and later. When BIOS enables Intel QuickPath Interconnect L1, it must also set the Intel QuickPath Interconnect RC L1_workaround input parameter to invoke the workaround. Below is an example of adding this input parameter to the csiSetup.c file in the Intel QuickPath Interconnect RC.

Status: For the steppings affected, see the Summary Tables of Changes.



14. PCIe Gen2 Tx Return loss fails spec

Problem: The the differential and common mode transmit return loss fail to meet the PCIe Gen2 specification.

Implication: Some Gen2 PCIe agents may have increased inter-symbol interference, ISI, due to signal reflection at the driver of Intel X58 Express Chipset. Gen1 return loss will meet specifications.

Workaround: Following Intel X58 Express Chipset PDG PCIe Gen2 guidelines will result in no Gen2 performance or functional issues.

Status: For the steppings affected, see the Summary Tables of Changes.

15. Persistent Jtag error reported at MIERRST register

Problem: The MIERRST register (Device: 20, Function: 2, Offset: 380h) bit 2 continuously gets set if the MIERRCTL register (Device: 20, Function: 2, Offset: 384h) bit 2 is set regardless of whether an error truly exists on the JTAG interface or not. Further, the MIERRCNT register (Device: 20, Function: 2, Offset 3C0h) will indicate an overflow via bit 7 as errors are continuously registered.

Implication: The MIERRST and MIERRCNT registers are not dependable sources of information if the IOH is configured to record errors related to the JTAG interface.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.

16. Interop issue of some PCIe Gen1 cards with Gen2 Devices

Note: Note: this is not a Intel X58 Express Chipset issue, and the workaround documented below is to resolve the interoperability issue certain PCIe Gen1 cards have with Gen2 capable Intel X58 Express Chipset PCIe root ports.

Problem: During PCIe physical layer link initialization, Intel X58 Express Chipset Gen2 capable PCIe root ports will advertise 2.0 Specification defined capabilities such as 5.0 GT/s data rate support and link upconfigure in the Training Sequence (TS) Ordered Sets. Some downstream pre-2.0 PCIe Gen1 cards may have trouble dealing with these Gen1 reserved fields that are not reserved for PCIe Gen2.

Implication: These PCIe Gen1 cards may fail to properly train with Intel X58 Express Chipset.

Workaround: Below is the standard method to force Intel X58 Express Chipset PCIe root ports to Gen1 operation:

1. Set LNKCON2 (Dev 1-10: Func 0: Offset C0h) bits [3:0] to 0001b.
2. Clear LNKCON2 (Dev 1-10: Func 0: Offset C0h) bits [6] to 0b
3. Set LNKCON (Dev 1-10: Func 0: Offset A0h) bit [5] to 1b to retrain the link.

If some PCIe Gen1 cards fail to train with the standard method, BIOS should program the following registers to prevent Intel X58 Express Chipset PCIe root ports from advertising 2.0 Specification defined capabilities. There is one register control field for each Intel X58 Express Chipset PCIe root port.

Dev 13: Funcs 6 - 1: Offset 4B4h bit[23] for root ports 6 - 1.

Dev 14: Funcs 3 - 0: Offset 4B4h bit[23] for root ports 10 - 7.

Clear this bit to disable the corresponding root port from advertising upconfigure capability before doing Step 3 above.

Status: For the steppings affected, see the Summary Tables of Changes.

**17. PCIe ASPM L1 tests reporting lost credits.**

Problem: The internal counter for flow control (FC) credit tracking is 8 bits wide. If the link is in a state where it cannot modify FC credits, but receives greater than 256 posted data credits to be updated, the credits can be lost. The most likely case for this to occur is when the link is in an ASPM L1 state.

Implication: Performance is impacted due to loss of credits.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.

18. Intel QuickPath Interconnect fails to retrain to L0 slow mode after 32 in-band resets.

Problem: After 32 Intel QuickPath Interconnect inband resets from the CPU to an IOH port are executed when in slow mode the next inband reset will fail as the Intel QuickPath Interconnect will not retrain to L0 slow mode.

Implication: The Intel QuickPath Interconnect link will fail when operating in slow mode after 32 Intel QuickPath Interconnect inband resets.

Workaround: BIT 5 of device 13, functions 1-0, offset B4Ch must be set every time before an inband reset from the CPU to an IOH port.

Status: For the steppings affected, see the Summary Tables of Changes.

19. ATS address field in the translation completion returned in reversed byte order.

Problem: A PCIe Translation Request returns a completion with the address, but the address is returned from LSB to MSB, when it should be returned MSB to LSB.

Implication: The PCIe Address Translation Specification, version 1.0 is violated.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

20. SLTCON bit 11 returning wrong value when read

Problem: SLTCON (Device:1-10, Function:0, Offset:A8h) register bit 11 should always return zero when read even if this bit is written with one; however, the bit is not reloaded with 0 when a 1 is written.

Implication: The bit 11 will not correctly reflect a value of 0 when read, but the functionality of the bit is not affected due to this. The EMIL is correctly pulses when bit 11 is written with a 1.

Workaround: No work around required. But to reflect the correct value, write SLTCON register bit 11 with zero after writing 1 to it.

Status: For the steppings affected, see the Summary Tables of Changes.


21. Boot-dependent Intel QuickPath Interconnect CRC errors

Problem: Intel QuickPath Interconnect CRC errors may be observed as a function of boot-dependent (probabilistic) phase alignment of internal clock domains

Implication: Intel QuickPath Interconnect port may fail to train or may train with sustained non-zero CRC error rate

Workaround: If CRC's are not user-tolerable, power cycle system.

Status: For the steppings affected, see the Summary Tables of Changes.



22. PCIe Inbound Msg that should be ignored is treated as Unsupported Request.

Problem: PCIe Inbound Messages with Message Codes between 0x40 and 0x48 (with the exception for x42 and 0x46) cause the root port to log errors as Unsupported Request when it should have ignored and discarded them.

Implication: Unsupported Request Error will be logged under the circumstances outlined above. End Point devices are strongly encouraged by PCIe spec 2.0 not to send those messages.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

23. Memory writes to a certain address range are considered advisory non-fatal.

Problem: Accesses above address range 2^{51} (TOCM) are required to be Master Aborted and error is to be logged in UNCERRSTS (Device:0-10, Function:0, Offset:104h). The accesses are aborted and logged; however, when severity of master-abort (UR) is set to non-fatal, memory write accesses are not to be considered advisory non-fatal, rather, they should be considered normal non-fatal. There is a range of address from 2^{51} to $2^{52}-1$ (0x8_0000_0000_0000 to 0xF_FFFF_FFFF_FFFF) for which memory write accesses are logged as advisory non-fatal. In this case, CORERRSTS (Device:0-10, Function:0, Offset:110h) bit 13 is set. Note that this issue does not arise if UR severity is set to Fatal.

Implication: The status of transactions occurring as described above will not be correctly reflected.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

24. Inbound Write fetches may not make forward progress

Problem: Stress traffic causes Outstanding Request Buffer (ORB) to fill up creating a corner condition where Inbound Write fetches that are supposed to make forward progress do not.

Implication: This creates a deadlock situation since CPU is unable to respond to an IOH initiated StopReq (Lock Request) because it waiting on pending WB to complete which IOH cannot because it's unable to make forward progress

Workaround: Reference the BIOS Update Specification.

Status: For the steppings affected, see the Summary Tables of Changes.


25. PCIe lane-to-lane performance variation

Problem: The lane-to-lane performance is inconsistent on the PCIe interface. Some lanes source/sink more current than others leading to analog variations.

Implication: Current compensation is compromised insomuch as not all of the lanes can be counted on to perform consistently. Optimizing the current compensation for one lane will not lead to optimized analog behavior in a neighboring lane. Ultimately, the overall analog behavior is limited by the lowest performing lane, thus overall channel analog behavior is compromised.

Workaround: Under Investigation.

Status: For the steppings affected, see the Summary Tables of Changes.



26. Transactions to addresses above TOCM are not setting the Master Abort

Problem: Inbound Transactions to addresses above TOCM are master aborted, but TBG does not set the C4 bit in the IOHERRST (Device:20, Function:2, Offset: 300h) register.

Implication: Such cases of master aborts will not have status recorded. When Intel VT is enabled, there is no workaround.

Workaround: When Intel VT is disabled, program the HPA_LIMT to the maximum value.

Status: For the steppings affected, see the Summary Tables of Changes.

27. SERR_EN ignored during unsupported request error escalation

Problem: If DEVCON (Device: 0-10, Function: 0, Offset: 98h) Unsupported Request Reporting Enable bit (bit 3) is clear and PCICMD (Device: 0-10, Function: 0, Offset: 04h) SERR Enable bit (bit 8) is set then the unsupported request should be escalated.

Implication: If the DEVCON (Device: 0-10, Function: 0, Offset: 98h) Unsupported Request Reporting Enable bit (bit 3) is clear, then the unsupported request will not be escalated.

Workaround: Set DEVCON (Device: 0-10, Function: 0, Offset: 98h) Unsupported Request Reporting Enable bit (bit 3) and unsupported request will be escalated.

Status: For the steppings affected, see the Summary Tables of Changes.

28. Setting bits 24 and 25 of the MISCCTRLSTS does not result in expected behavior

Problem: Setting the bits 24 and 25 in the MISCCTRLSTS (Device: 0-10; Function: 0, Offset: 188h) register does not result in the peer-to-peer memory read/write transactions being disabled as expected.

Implication: Memory read/write transactions may still go through to the memory and the expected Completer Abort message may not be received.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Change.

29. Timeout values much larger than specified

Problem: When using the bits 51:48 of the CSIPCTRL register (Device: 16, Function: 1, Offset 0x4C) to set the configuration retry timeout values, it has been observed that the actual timeout values may be much longer than what is specified for each of the settings.

Implication: The expected timeout signal for transactions exceeding the timelimit may not be observed.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

30. XPGLBERRPTR not accurately indicating first error in XPGLBERRSTS

Problem: The XPGLBERRPTR register may not accurately point to which of the errors indicated by the by XPGLBERRSTS register occurred first.

Implication: Due to this issue it may not be possible to detect which error occurred first in XPGLBERRSTS register when more than one bit is set in the XPGLBERRSTS register.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes

31. Failure during operation at PCI Express* L1 power management state

Problem: There may be intermittent failures observed when exiting from the L1 power management state back to the L0 power management state.



Implication: Due to this issue, successful transition from the L1 state back to the L0 state may not be achieved.

Workaround: Set the bit 26 in the following registers: Device 14, Function 3-0, offset 390h and Device 13, Function 6-0, Offset 390h to 1. Also set bit 15 in the following registers: Device 14, Function 0, Offset 31Ch and Device 13, Function 3, Offset 31Ch to a 0.

Status: For the steppings affected, see the Summary Tables of Changes.

32. Bandwidth very low for write traffic with noSnoop attribute set.

Problem: 100% upstream posted write traffic (PCIe to memory) with the noSnoop attribute set is achieving lower than expected bandwidth. In comparison, the same test with snooped traffic achieves much higher bandwidths.

Implication: Affected systems will achieve lower than expected performance.

Workaround: The PERFCTRLSTS (Device: 1-10, Function:0, Offset: 180h) register, bits 3 and 2 should be set to 0.

Status: For the steppings affected, see the Summary Tables of Changes.

33. Intel QuickPath Interconnect Queue/Table overflow or underflow error observed.

Problem: Intel QuickPath Interconnect Queue/Table overflow or underflow error status (Dev 20:Func 2:Offset 84h.[27:26]) is observed.

Implication: This spurious error status will cause false error reporting if enabled.

Workaround: BIOS should program the following Intel X58 Express Chipset registers to mask off the spurious sub-state status and allow the Intel QuickPath Interconnect Queue/Table overflow or underflow error status to function. Set Dev 13:Func 0/1: Offset F90h (Dword) bit 14 to 1.

Status: For the steppings affected, see the Summary Tables of Changes.

34. IOH Course Grained Clock Gating requires an ITP.

Problem: Course Grained Clock Gating will only function if an ITP is connected to the IOH.

Implication: Any testing involving Course Grained Clock Gating will require an ITP.

Workaround: An ITP must be installed in the system on both the IOH and CPU with the ITP software running.

Status: For the steppings affected, see the Summary Tables of Changes.

35. Some Gen2 endpoints will not complete the training in Gen2 mode.

Problem: Some Gen2 endpoints will not complete training as expected when operating in Gen2 mode.

Note: Related to 3319126.

Implication: Affected Gen2 endpoints will be unusable when operating in Gen2 mode without the workaround.


Workaround: Please refer to the Intel X58 Express Chipset BIOS Sightings Update

Status: For the steppings affected, see the Summary Tables of Changes.

36. Setting bit 10 of PCICMD register having effect on legacy interrupts.

Problem: PCICMD register Interrupt disable bit (bit 10) controls the generation of root port interrupts for internal events. If this bit is set to one, INTx status (bit 3) in PCISTS register PCISTS is cleared. This is incorrect as the INTx status (bit 3) should not be affected when the interrupt disable bit in the PCICMD register is set.

Implication: When bit 10 of the PCICMD register is set, INTX status in PCISTS becomes zero. This is incorrect behavior.



Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

37. ACS Violation is not treated as Advisory when severity is set to Non-Fatal.

Problem: If the severity of an Access Control Services (ACS) Violation is defined as non-fatal and the completer sends a completion with CA completion status then this case must be handled as an Advisory Non-Fatal Error as described in the PCI Express Spec.

Implication: The ACS violations are treated as Non-Fatal errors instead of Advisory Non-Fatal errors.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

38. MIERRCNT does not properly count persistent SMBus retry failures and MINNERRST doesn't log the errors.

Problem: The first time a SMBus invalid access is performed the error is logged on the MIERRST register and the MIERRCNT register counts the first error as expected. However, subsequent SMBUS retry failures are not captured by the MIERRCNT register and also the MINNERRST register does not log the error as expected.

Implication: Persistent SMBus retry failures are not logged as expected.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.

39. QPI L1 state is taking greater than 15 us from L1 exit to L0 state.

Problem: The IOH Intel QuickPath Interconnect interface is taking longer than 15 us to transition from the L1 state to the L0 state.

Implication: As a result IOH L1 Exit latency needs to be reduced to prevent USB underun or potential audio glitches.

Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

40. EOI to the I/OxAPIC can be blocked.

Problem: The End of Interrupt (EOI) message targeted to the I/OxAPIC will be blocked in the case that an INTx from a PCIe device and an MSI are pending inside the Intel X58 Express Chipset.

Implication: The I/OxAPIC will continue to treat a level triggered interrupt as asserted as the EOI message is never received by the I/OxAPIC to clear the interrupt.

Workaround: Avoid using the I/OxAPIC in the Intel X58 Express Chipset and use only the I/OxAPIC in the ICH component.

Status: For the steppings affected, see the Summary Tables of Changes.

41. GTIME upper 32 bits can not be read or written.

Problem: The GTIME register (Device: 20, Function: 2, Offset 0x1D4) always returns 0 and writes 0 to the upper 32 bits of data. The register does count correctly.

Implication: The GTIME upper 32 bits can not be read or written accurately. The GFFERRTIME (Device:20, Function:2, Offset:1E0h) and GNFERRTIME (Device:20, Function:2, Offset:1F0h) registers do get logged correctly given that the register does count correctly.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.



42. Intel QuickPath Interconnect L0s and L1 Power Management Link State Fails.

Problem: The Intel QuickPath Interconnect on the IOH-side of the Intel® X58 Express Chipset may exhibit some instability entering or exiting the link power management states L0s and L1.

Implication: System hang or a low rate of CRC errors on QPI link.

Workaround: BIOS workaround- Disable Intel QPI power management states L0s and L1. CPU Package C-states will be disabled as well. Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

43. PCI Express 2.0 L0s Link Recovery Fails.

Problem: When the Intel X58 Express Chipset receives a request to enter L0s on the PCI Express* 2.0 interface and a subsequent link recovery is initiated the link training may fail. This issue does not affect the DMI interface on the Intel X58 Express Chipset.

Implication: System hang.

Workaround: BIOS workaround- Disable the transmission of L0s requests on devices connected to the Intel X58 Express Chipset PCI Express interface. Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

44. IOH May Falsely Assert THERMTRIP_N Signal After a Reset Event

Problem: When Intel X58 Express Chipset based platforms undergo a reset event, such as a front panel reset or any warm reset, the THERMTRIP_N pin on the IOH may falsely assert, resulting in unexpected system behavior (that is, - possible platform shutdown, IOH initiated platform throttling, and so forth, depending on implementation).

Implication: Unexpected system behavior.

Note: Platforms that have THERMTRIP_N unconnected or connected to a pull-up resistor are NOT impacted.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.

45. Intel VT-d queue-based invalidation is enabled only when enabled on both channels

Problem: Intel VT-d queue-based invalidation is enabled only when queue-based invalidation is enabled on both isochronous and non-isochronous channels.

Implication: This is a violation of the Intel VT-d specification which allows each engine to be configured in either queue-based or register-based invalidation independently.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.


46. Intel VT-d does not support the draining of compatibility-format interrupts

Problem: The Intel VT-d does not support the draining of compatibility-format interrupts which is a violation of the Intel Vt-d specification.

Implication: The Draining of compatibility-format interrupts is used when software wants to convert an interrupt source from compatibility-format to remappable-format.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.



47. Hardware applies HPA_LIMIT to upstream memory request when Intel VT-d is disabled.

Problem: When Intel VT-d DMA translation is enabled, hardware correctly applies the GPA_LIMIT check to the incoming DMA address and HPA_LIMIT to Intel VT-d page-walk addresses. When Intel VT-d DMA translation is not enabled, hardware applies the HPA_LIMIT check to incoming DMA addresses. Note that the reset default value of HPA_LIMIT in hardware is 39 (that is, address bits 38:0 are valid).

Implication: When Intel VT-d DMA translation is not enabled, hardware should ignore the GPA and HPA LIMITs so that hardware does not apply the HPA_LIMIT check to the incoming DMA addresses.

Workaround: The easiest workaround for this hardware bug is to have BIOS program the appropriate HPA_LIMIT in VTGENCTRL register always, irrespective of the BIOS setup option to expose/disable Intel VT-d. Please see the BIOS Update Specification for more details.

Status: For the steppings affected, see the Summary Tables of Changes.

48. Intel VT-d: Memory read request with AT=11 results in malformed TLP

Problem: When in Intel VT-d mode, memory read requests with Address Type 11b result in a malformed Transaction Layer Packet.

Implication: When in Intel VT-d mode, memory read requests with Address Type 11b should be completed with an Unsupported Request; therefore, this is a violation of the PCIe Address Translation Services Specification, Version 1.0.

Workaround: None.

Status: For the steppings affected, see the Summary Tables of Changes.

49. Intel VT-d: IOTLB Domain-page-selective invalidation not working correctly.

Problem: When Intel VT-d is doing the IOTLB Domain-page-selective invalidation for leaf only(IH=1) while Intel VT-d is still in the middle of pagewalk, Intel VT-d needs to invalidate all entries in the cache that are still waiting for the pagewalk data from memory to complete a walk. Because of a logic bug, those entries did not get invalidated and are still waiting to receive the pagewalk data. This can never happen because the data could be thrown away by the invalidation command.

Implication: The pagewalk never completes and no address will be given back to the IOH and this could hang the system.

Workaround: Workaround is to not set the INVADDRREG (Device: MMIO, BAR: VTBAR), bit 6.

Status: For the steppings affected, see the Summary Tables of Changes.

50. Intel VT-d translated write transactions are blocked but not recorded.

Problem: The Intel VT-d engine blocks Intel VT-d Translated write transactions to the interrupt address range (0xFEExxxxx) but does not record the error. Requests outside of the interrupt address range are not affected.


Implication: For Intel VT-d Translated write transactions to the interrupt address range (0xFEExxxxx), the transaction is blocked, but evidence of the blocked transaction will not be available in the error registers.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

51. The IOH ignores Snoop Behavior Bit in Intel VT-d Page Table.

Problem: When the snoop behavior control bit <11> of the Intel VT-d leaf page table entry is set, all inbound untranslated transactions should be treated as snooped, irrespective of the NS (Non-Snoop) attribute in the request. However, the state of this bit appears to not have any affect on snoop behavior.



Implication: If any DMI agent sends non-snoop traffic and the page-table entry configures the request to force snoop behavior, the traffic will remain non-snoop traffic.

Workaround: The BIOS should force the IOH to disable the "snoop control" capability for the nonisochronous DMI Intel VT-d unit. See the BIOS specification update for further details.

Status: For the steppings affected, see the Summary Tables of Changes.

52. Intel QuickPath Interconnect Link Training failures in L0 when L0r enabled.

Problem: Intel QuickPath Interconnect Link Training failures seen when L0r is enabled. The failures observed include repetitive CRC errors resulting in a Link Layer Retry (LLR) Phy Reinit or LLR Abort.

Implication: Certain systems may fail to boot when L0r is enabled.

Workaround: The fix is integrated into the Intel QuickPath Interconnect RC 0.90RC1 and above.

Status: For the steppings affected, see the Summary Tables of Changes.

53. Intel VT-d reports an Invalidation Queue Error when the Queue size is 7.

Problem: When the Queue Size is set to 7 in the Invalidation Queue Address register (Addr: MMIO, BAR: VTBAR, Offset: 90h, 1090h), Intel VT-d incorrectly reports Invalidation Queue Error in the Fault Status register (Addr: MMIO, BAR: VTBAR, Offset: 34h, 1034h, bit 4) at the time of fetching a descriptor.

Implication: A fault event will be generated thus leading to an interrupt request if the Fault Event Control Register (Addr: MMIO, BAR: VTBAR, Offset: 38h, 1038h) is programmed to respond to this error.

Workaround: Do not set the Queue Size to 7.

Status: For the steppings affected, see the Summary Tables of Changes.

54. ERRSID not logging ReqID for Inbound ERR_* messages

Problem: Error Source Identification Register (Device: 0-10, Function: 0, Offset: 134h) does not log Requester ID for Inbound ERR_* messages.

Implication: While internally generated error messages in the IOH will have their Requester ID logged correctly in this register, incoming ERR_* messages' Requester ID will not be.

Workaround: Software needs to read downstream devices' error logs to identify the source of the error.

Status: For the steppings affected, see the Summary Tables of Changes.

55. Intel QuickPath Interconnect PhyResets causing Intel QPI Errors to be flagged

Problem: When Intel QuickPath Interconnect PhyResets are executed, unexpected Intel QuickPath Interconnect errors may be logged. The link remains operational regardless of the logged errors. The following errors have been observed: B6 and DG. These errors will be flagged in the QPI[1:0]ERRST (Device: 20, Function: 2, Offset: 280h, 200h) register.

Implication: This issue may manifest itself at all Intel QuickPath Interconnect speeds upon assertion of PhyReset, which should not occur during normal operation. Regardless of the issue, the link remains operational and system does not fail.

Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.



56. Intel QuickPath Interconnect initialization abort failures logged during power-on resets.

Problem: During power-on reset, the IOH may log Intel QuickPath Interconnect initialization abort (D2) failures via the QPI[1:0]ERRST (Device:20, Function:2, Offset:280h, 200h) register, bit 4.

Implication: This has been seen on a small number of parts. If the power-on reset is executed and D2 is logged, then this indicates that at least one Intel QuickPath Interconnect lane hasn't trained correctly.

Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

57. Intel QuickPath Interconnect errors can occur on inband resets.

Problem: On some inband resets (includes L1 and physical layer resets), the IOH may take 1 or 2 extra cycles to determine that the Intel QuickPath Interconnect forwarded clock has stopped.

Implication: The IOH may see CRC errors, system hangs, or link level control errors on a small percentage of L1 Entries.

Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

58. Data Mismatch on Inbound MemWr's after MSI with payload greater than 1 DWORD payload

Problem: If for some reason an MSI is incorrectly sent with a payload greater than 1 DWORD then data mismatches may result on subsequent inbound memory writes.

Implication: This issue may lead to data corruption.

Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

59. MSI with greater than 1DWord payload is not logged in XPUNCERRSTS bit 8.

Problem: When ports are used in a x8 or x16 configuration, bit 8 of the XPUNCERRSTS register (Device:0-10, Function:0, Offset: 0x208) on the port that received the MSI, may not log that an MSI write with a payload greater than 1 DWORD has been received.

Implication: The expected indication is not available on the port that received on the MSI.


Workaround: Please contact your Intel field representative for the latest BIOS Specification Update for details.

Status: For the steppings affected, see the Summary Tables of Changes.

60. PFP Flag (due to Intel VT-d ISOCH fault) in Fault Status Register not being cleared.

Problem: When a Intel VT-d Fault is detected in the ISOCH engine, subsequently setting bit 127 [F fld] of the Fault Record Register (Register:FLTREC[7:0], Addr: MMIO, BAR: VTBAR, Offset:[170h:100h], [1170:1100h]) fails to result in the clearing of the PFP (Primary Fault Pending) status bit in the Fault Status Register (Register: FLTSTS, Addr: MMIO, BAR: VTBAR, Offset:34h, 1034h). In the current design, once the PFP bit is set, it remains set and is not cleared by hardware even after the fault is cleared by software in the Fault Record Register.

Implication: Does not impact normal functionality. Only the first fault can be recorded and reported to software. The additional faults will be ignored.



Workaround: None.

61. **Lost interrupts when MSI used.**

Problem: In a system where Message Signaled Interrupt (MSI) is used to forward PCIe interrupts to the processor, and Uncorrectable (UC) Errors are classified as a mix of either fatal or non-fatal errors, if there are multiple UC Errors occurring in a certain window, a non-related local interrupt waiting to get serviced may get lost.

Implication: This may result in an interrupt getting lost before it is serviced. Please note this issue does not affect any interrupts that are delivered via the ERR# pins or the Global Error Reporting Mechanism.

Workaround: Under Investigation

62. **Intel VT-d: UP Workstation ONLY. Receiving two identical interrupt requests in back to back cycles may corrupt attributes of remapped interrupt, or hang subsequent interrupt-remap-cache invalidation command.**

Problem: If Intel VT-d interrupt-remapping hardware receives two identical back to back interrupt requests the attributes of remapped interrupt returned may be corrupt. This interrupt sequence may hang the system if the software executes a subsequent interrupt-remap-cache invalidation command.

Implication: This scenario may lead to unpredictable external interrupt behavior or a system hang.

Workaround: BIOS to set bit 25 in (Dev/Func: 14/4, Offset: 168h). De-feature interrupt remapping capability by clearing Interrupt-Remap (IR) RWO field (bit 3) in Extended Capability Register and clearing INTR_REMAP bit (bit 0) in Flags field of ACPI DMAR tables used to report Intel VT-d hardware capability to software.

Status: For the steppings affected, see the Summary Tables of Changes.

63. **Intel QuickPath Interconnect Error Status D3 is observed**

Problem: Intel QuickPath Interconnect D3 error status (Dev 20:Func 2:Offset 280h/200h.[11]) is observed due to false sub-state illegal link reset detection

Implication: This spurious error status may cause false error reporting if enabled in system debug scenarios

Workaround: BIOS should program the following IOH registers to mask off the spurious sub-state status and allow the Intel QuickPath Interconnect D3 error status to function normally: Set Dev 13:Func 0/1: Offset F8Ch (Dword) bit[14] to 1. **Note: this workaround will be implemented in future release of Intel QuickPath Interconnect RC.**

Status: For the steppings affected, see the Summary Tables of Changes.

64. **Unpredictable PCI behavior accessing non-existent memory space**

Problem: Locked instructions whose memory reference is split across cache line boundaries and are aborted on PCI behind ICH may cause subsequent PCI writes to be noticeably unpredictable.

Implication: Aborted split lock accesses to non-existent PCI memory space behind ICH may cause PCI devices to become inoperable until a platform reset. Intel has not observed this issue with commercially available software and has only observed this in a synthetic test environment.

Workaround: None

Status: For the steppings affected, see the Summary Tables of Changes.

65. **Bandwidth changed status errors being escalated to Global RAS**

Problem: It has been observed that setting the bit 0 in the XPCORERRMSK register (Device: 0-10, Function: 0, Offset: 204h) only prevents bandwidth changed status errors from being escalated to Global error registers for Ports 0, 3 and 7.

Implication: Setting the bit 0 in the XPCORERRMSK register (Device: 0-10, Function: 0, Offset: 204h) does not prevent the bandwidth changed status errors from being escalated to Global error registers for Ports 1,2,4,5,6,8,9 and 10.



Workaround: The bandwidth status changed error is a correctable error and under normal operation is permitted to be escalated to the global error register.

Status: For the steppings affected, see the Summary Tables of Changes.

66. TXT writes may not complete as expected

Problem: When TXT is enabled, in a system with back to back Configuration Retry transactions and certain TXT write requests, the TXT write request may not complete as expected.

Implication: This may result in incorrect system behavior. *Please note that TXT is not supported on S-Spec SLGMX.*

Workaround: None Identified. BIOS must not enable TXT in processors via IA32_FEATURE_CONTROL MSR for S-Spec SLGMX See the Chapter 6 of Intel® 64 and IA-32 Architecture Software Developer's Manual Volume 2B for details regarding this MSR.

Status: For the steppings affected, see the Summary Tables of Changes.

67. IOTLB Invalidations not completing on Intel VT-d Isochronous unit (UP WS Only)

Problem: IOTLB invalidation on the Intel VT-d Isochronous unit may not complete and cause IVT field (bit 63 of IOTLBINV register) to not be cleared as expected. As a result, software continues to poll this bit and not detect successful invalidation completion.

Implication: When the IOH maps Integrated High Definition Audio traffic on DMI to the Intel VT-d Isochronous unit, software requests for IOTLB invalidation while audio traffic is active may result in system hang. Intel has not encountered the problem with currently available commercial Intel VT-d enabled software.

Workaround: None Identified.

Status: For the steppings affected, see the Summary Tables of Changes.

68. Intel® VT-d: Address remapping error when DMA/interrupt remapping is active

Problem: With Intel-VT-d enabled, when software updates the root table pointer or interrupt remapping table pointer while DMA/interrupt remapping is active, it is possible that the address used to access the page-table structure for DMA requests or interrupt remapping could be corrupted and cause an address remapping error.

Implication: Software cannot update the root table pointer or interrupt remapping table pointer while DMA/interrupt remapping is active

Workaround: None Identified.

Status: For the steppings affected, see the Summary Tables of Changes.


69. Intel® VT-d: In-flight remap-able interrupts not drained on interrupt invalidation command

Problem: The Intel® VT-d hardware is not draining in-flight remapped interrupts when an interrupt invalidation command occurs. The software expectation is that when it performs invalidation, all in-flight interrupts are received and acknowledged by the processor.

Implication: As a result of this issue, software may lose interrupts or receive spurious interrupts.

Workaround: BIOS to set bit 25 in (Dev/Func: 14/4, Offset: 168h). Disable interrupt remapping capability by clearing Interrupt-Remap (IR) RWO field (bit 3) in Extended Capability Register and clearing INTR_REMAP bit (bit 0) in Flags field of ACPI DMAR tables used to report VT-d hardware capability to software.

Status: For the steppings affected, see the Summary Tables of Changes.



70. Source ID for errors internally detected by PCIe root port 3 is not logged as expected

Problem: The Source ID for errors internally detected by PCIe root port 3 is not logged in the ERRSID register (Device: 0-10, Function: 0, Offset: 134h) as expected. However the error message details are logged correctly as specified by the PCIe Specification in the root port 3 RPERRSTS register (Device: 0-10, Function: 0, Offset: 130h) as expected.

Implication: The root port 3 ERRSID register does not log the internally detected errors as expected.

Workaround: Other methods to determine if root port 3 detected an error: Reading the UNCERRSTS (Device: 0-10, Function: 0, Offset: 104h) and CORERRSTS (Device: 0-10, Function: 0, Offset: 110h) registers will allow software to determine if root port 3 internally detected an error as well as the type of error.

Status: For the steppings affected, see the Summary Tables of Changes.



Specification Changes

No Specification Changes in this revision of the specification update.



Specification Clarifications

1. **Non-Posted Peer-to-Peer Support between IOH and ICH.**

Specification Clarification for the document: *Intel® X58 Express Chipset Datasheet #320838-003*.

- Page 17) Feature Section -

Currently published: Full peer-to-peer support between PCI Express interfaces.

Clarification: IOH to ICH Non-Posted peer-to-peer transactions are not supported.



Documentation Changes

1. Documentation update to the *Intel® X58 Express Chipset Datasheet #320838*. The next revision of the Datasheet will reflect the following specification changes to the documentation.

- Remove line- *“Support for root entry, context entry, and default context.”*
- Remove line- *“Support for Intel VT-d read prefetching/snarfing, that is, translations within a cacheline are stored in an internal buffer for reuse for subsequent transactions.”*
- Remove line- *“Support for PCISIG endpoint caching (ATS).”*
- Remove line- *“No support for Intel VT-d translation bypass address range (such usage models need to be resolved with VMM help in setting up the page tables correctly).”*
- Remove line- *“Support for queue-based invalidation interface.”*
- Change *“No support for 1 or 2 level page walks for 1, 2, or 3 level walks for non-isoch remap engine”* to *“No support for 1 or 2 level page walks for isoch remap engine, and 1, 2, or 3 level walks for non-isoch remap engine”*

2. Chapter 17, Configuration Register Space

- ERRCAPHDR - PCI Express Enhanced Capability Header Register
ERRCAPHDR[31:20] = 150h

Register: ERRCAPHDR Device:0-10 Function:0 Offset:100h			
Bit	Attr	Default	Description
31:20	RO	00h 150h	Next Capability Offset This field points to the next Capability in extended configuration space.
19:16	RO	1h	Capability Version Set to 1h for this version of the PCI Express logic
15:0	RO	0001h	PCI Express Extended CAP_ID Assigned for advanced error reporting

§