

# Intel<sup>®</sup> Xeon<sup>®</sup> D-2100 Processor Product Family

**Specification Update**

---

*August 2020*

**Order Number: 338854-012**



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described 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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm>

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

No computer system can be absolutely secure.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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


Copyright © 2020, Intel Corporation. All rights reserved.



## Contents

---

<b>Revision History</b> .....	<b>4</b>
<b>Preface</b> .....	<b>5</b>
Affected Documents/Related Documents .....	5
Nomenclature .....	6
<b>Summary Tables of Changes</b> .....	<b>7</b>
Codes Used in Summary Tables .....	7
Stepping .....	7
Page .....	7
Status .....	7
Row .....	7
<b>Intel® Xeon® D-2100 Processor Product Family CPU Errata</b> .....	<b>8</b>
<b>Intel® Xeon® D-2100 Processor Product Family PCH Errata</b> .....	<b>13</b>
<b>Identification Information</b> .....	<b>16</b>
Component Identification via Programming Interface .....	16
Component Marking Information .....	17
<b>CPU Errata</b> .....	<b>18</b>
<b>PCH Errata</b> .....	<b>41</b>
<b>Specification Changes</b> .....	<b>59</b>
<b>Specification Clarifications</b> .....	<b>60</b>
<b>Documentation Changes</b> .....	<b>61</b>



## Revision History

---

Date	Revision	Description
August 2020	012	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD36S</a>.</li><li>Removed <a href="#">SKXD26S</a>. and <a href="#">SKXD11</a>. due to duplicate entries.</li></ul>
July 2020	011	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD35S</a>.</li></ul>
June 2020	010	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD34S</a>.</li></ul>
May 2020	009	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD32S</a>. and <a href="#">SKXD33S</a>.</li></ul>
April 2020	008	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD31S</a>.</li></ul>
December 2019	007	<ul style="list-style-type: none"><li>Updated CPU errata <a href="#">SKXD30S</a>.</li></ul>
November 2019	006	<ul style="list-style-type: none"><li>Updated PCH errata <a href="#">LBG22</a>.</li><li>Added PCH errata <a href="#">LBG61</a>.</li><li>Added CPU errata <a href="#">SKXD27S</a>., <a href="#">SKXD28S</a>., and <a href="#">SKXD29S</a>.</li></ul>
October 2019	005	<ul style="list-style-type: none"><li>Updated CPU errata <a href="#">SKXD25S</a>.</li></ul>
September 2019	004	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD25S</a>. and <a href="#">SKXD26S</a>.</li><li>Updated PCH errata <a href="#">LBG60</a>.</li></ul>
August 2019	003	<ul style="list-style-type: none"><li>Added CPU errata <a href="#">SKXD24S</a>.</li></ul>
May 2019	002	Updated the following: <ul style="list-style-type: none"><li>Updated <a href="#">SKXD14S</a>.; removed duplicate entry.</li></ul> Added the following Errata: <ul style="list-style-type: none"><li>Added <a href="#">SKXD22S</a>.</li><li>Added <a href="#">SKXD23S</a>.</li></ul>
March 2019	001	Updated the following: <ul style="list-style-type: none"><li>Updated <a href="#">Component Marking Information</a>.</li><li>Updated <a href="#">SKXD55</a>.</li><li>Updated <a href="#">SKXD57</a>.</li></ul>





## 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 sighting, 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/Related Documents

Document Title	Document Number/ Location
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual (all seven volumes)</i>	325462 <sup>1</sup>
<i>Intel® 64 and IA-32 Intel Architectures Optimization Reference Manual</i>	248966 <sup>1</sup>
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	252046 <sup>1</sup>
<i>Intel® Virtualization Technology Specification for Directed I/O Architecture Specification</i>	D51397 <sup>1</sup>
<i>ACPI Specifications</i>	<a href="http://www.acpi.info">www.acpi.info</a>

**Notes:**

1. This document can be downloaded from <https://software.intel.com/en-us/articles/intel-sdm>.



## Nomenclature

**S-Spec Number** is a five-digit code used to identify products. Products are differentiated by their unique characteristics, such as, core speed, L2 cache size, package type, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

**QDF Number** is a four digit code used to distinguish between engineering samples. These samples are used for qualification and early design validation. The functionality of these parts can range from mechanical only to fully functional. This document has a processor identification information table that lists these QDF numbers and the corresponding product details.

**Errata** are design defects or errors. These may cause the Intel® Xeon® D-2100 Processor Product Family'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, etc.).



## Summary Tables of Changes

---

The following tables indicate the sightings, specification changes, specification clarifications, and documentation changes which apply to the Intel® Xeon® D-2100 Processor Product Family. Intel may fix some of the sightings 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:	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.
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.



## Intel® Xeon® D-2100 Processor Product Family CPU Errata

**Table 1. CPU Errata (Sheet 1 of 5)**

Errata Number/ HSD Number	Steppings	Status	Errata
	M-1		
SKXD1.	X	No Fix	A CAP Error While Entering Package C6 Might Cause DRAM to Fail to Enter Self-Refresh
SKXD2.	X	No Fix	PCIe* Lane Error Status Register Might Log False Correctable Error
SKXD3.	X	No Fix	In Memory Mirror Mode, DataErrorChunk Field Might be Incorrect
SKXD4.	X	No Fix	PCIe Port Might Incorrectly Log Malformed_TLP Error
SKXD5.	X	No Fix	CMCI Might Not be Signaled for Corrected Error
SKXD6.	X	No Fix	Intel® CAT/CDP Might Not Restrict Cacheline Allocation Under Certain Conditions
SKXD7.	X	No Fix	Credits Not Returned For PCIe Packets That Fail ECRC Check Problem
SKXD8.	X	No Fix	PCIe Link Might Fail to Train
SKXD9.	X	No Fix	IODC Entry 0 Cannot be Masked
SKXD10.	X	No Fix	With eMCA2 Enabled a 3-Strike Might Cause an Unnecessary CATERR# Instead of Only MSMI
SKXD11.			Removed
SKXD12.	X	No Fix	CSRs SVID and SDID Are Not Implemented For Some DDRIO and PCU Devices
SKXD13.	X	No Fix	Register Broadcast Read From DDRIO May Return a Zero Value
SKXD14.	X	No Fix	Intel® CMT Counters May Not Count Accurately
SKXD15.	X	No Fix	Intel® CAT Might Not Restrict Cacheline Allocation Under Certain Conditions
SKXD16.	X	No Fix	Intel PCIe Corrected Error Threshold Does Not Consider Overflow Count When Incrementing Error Counter
SKXD17.	X	No Fix	IIO RAS VPP Hangs During The Warm Reset Test
SKXD18.	X	No Fix	A Core 3-Strike Event May Be Seen Under Certain Test Conditions





**Table 1. CPU Errata (Sheet 2 of 5)**

Errata Number/ HSD Number	Steppings	Status	Errata
	M-1		
SKXD19.	X	No Fix	DDR4 Memory Bandwidth May be Lower Than Expected at 2133 and 1866 MHz
SKXD20.	X	No Fix	Intel Sparing Per-Rank Error Masking Does Not Mask Correctable Errors
SKXD21.	X	No Fix	PCIe Root Port Electromechanical Interlock Control Register Can Be Written
SKXD22.	X	No Fix	Performance Monitoring M2MEM Counters For Memory Controller Reads/Writes Are Not Counting Read/Write Retries
SKXD23.	X	No Fix	System Hangs May Occur When IPQ and IRQ Requests Happen at The Same Time
SKXD24.	X	No Fix	IIO VPP May Hang During Warm Reset
SKXD25.	X	No Fix	Machine Check Events May Be Logged In Banks 9, 10 And 11 That Do Not Represent Actual Errors
SKXD26.	X	No Fix	Lower Than Expected Performance May Be Seen With Some Intel AVX Workloads Due To Incorrect Uncore Frequency Scaling
SKXD27.	X	No Fix	Spurious Corrected Errors May Be Reported
SKXD28.	X	No Fix	Writing To LT_LOCK_MEMORY And LT_UNLOCK_MEMORY MSRs Simultaneously May Have Inconsistent Results
SKXD29.	X	No Fix	Masked Bytes In A Vector Masked Store Instructions May Cause Write Back Of A Cache Line
SKXD30.	X	No Fix	ERROR_N[2:0] Pins May Not Be Cleared After a Warm Reset
SKXD31.	X	No Fix	CRC Store Operation Corner Case May Result In Hang
SKXD32.	X	No Fix	Intel PCIe Slot Presence Detect and Presence Detect Changed Logic Not PCIe Specification Compliant
SKXD33.	X	No Fix	In Patrol Scrub System Address Mode, Address is Not Loaded From CSRs After Re-enable
SKXD34.	X	No Fix	The Corrected Error Count Overflow Bit In IA32_MC0_STATUS Is Not Updated When The UC Bit Is Set
SKXD35.	X	No Fix	SMRAM State-Save Area Above The 4 GB Boundary May Cause Unpredictable System Behavior
SKXD36.	X	No Fix	POPCNT Instruction May Take Longer To Execute Than Expected
SKXD37.	X	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
SKXD38.	X	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
SKXD39.	X	No Fix	Incorrect Branch Predicted Bit In BTS/BTM Branch Records
SKXD40.	X	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When A MOV/POP SS Is Followed By A Store Or An MMX Instruction
SKXD41.	X	No Fix	Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions




**Table 1. CPU Errata (Sheet 3 of 5)**

Errata Number/ HSD Number	Steppings	Status	Errata
	M-1		
SKXD42.	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 Is Set To 1
SKXD43.	X	No Fix	x87 FPU Exception (#MF) May Be Signaled Earlier Than Expected
SKXD44.	X	No Fix	CPUID TLB Associativity Information Is Inaccurate
SKXD45.	X	No Fix	Vector Masked Store Instructions May Cause Write Back Of Cache Line Where Bytes Are Masked
SKXD46.	X	No Fix	Incorrect FROM_IP Value For An RTM Abort In BTM Or BTS May Be Observed
SKXD47.	X	No Fix	MOVNTDQA From WC Memory May Pass Earlier Locked Instructions
SKXD48.	X	No Fix	#GP On Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code
SKXD49.	X	No Fix	BNDLDX And BNDSTX May Not Signal #GP On Non-Canonical Bound Directory Access
SKXD50.	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests May be Incorrect
SKXD51.	X	No Fix	Branch Instructions May Initialize MPX Bound Registers Incorrectly
SKXD52.	X	No Fix	A Spurious APIC Timer Interrupt May Occur After Timed MWAIT
SKXD53.	X	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
SKXD54.	X	No Fix	Use of VMX TSC Scaling or TSC Offsetting Will Result in Corrupted Intel PT Packets
SKXD55.	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
SKXD56.	X	No Fix	Memory May Continue to Throttle after MEMHOT# De-Assertion
SKXD57.	X	No Fix	Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values
SKXD1S.	X	No Fix	Certain PCIe Cards May Not Function as Expected When Executing a Hot Plug Add/Remove Sequence
SKXD2S.	X	No Fix	PROCHOT Input to Drive Processor Throttling May Not Function As Intended
SKXD3S.	X	No Fix	Reading Some C-state Residency MSRs May Result in Unpredictable System Behavior
SKXD4S.	X	No Fix	Intel MBA Read After MSR Write May Return Incorrect Value
SKXD5S.	X	No Fix	In eMCA2 Mode, When the Retirement Watchdog Timeout Occurs CATERR# May be Asserted
SKXD6S.	X	No Fix	Intel MBA May Incorrectly Throttle All Threads
SKXD7S.	X	No Fix	Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang



**Table 1. CPU Errata (Sheet 4 of 5)**

Errata Number/ HSD Number	Steppings	Status	Errata
	M-1		
SKXD8S.	X	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
SKXD9S.	X	No Fix	SpecPower Workloads on Native without Legacy Mode May See a Higher than Expected Loadline Power Rating
SKXD10S.	X	No Fix	Error Injection Testing with NVDIMM Present in System May Cause a System Reset
SKXD11S.	X	No Fix	MSCOD Error Code =0x2E May Be Seen During Warm Reset Testing
SKXD12S.	X	No Fix	Processor May Hang When Executing Code In an HLE Transaction Region
SKXD13S.	X	No Fix	IDI_MISC Performance Monitoring Events May Be Inaccurate
SKXD14S.	X	No Fix	Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB
SKXD15S.	X	No Fix	Intel® PT VM-entry Indication Depends on the Incorrect VMCS Control Field
SKXD16S.	X	No Fix	System May Hang Due to Lock Prefixes on Instructions That Access IIO's MMCFG
SKXD17S.	X	No Fix	With Intel ME Recovery Mode enabled, The Processor May Incorrectly Throttle
SKXD18S.	X	No Fix	VccSA Voltage May Increase After A Demoted Warm Reset Cycle
SKXD19S.	X	No Fix	An IERR May be Seen When The CPU Attempts Consecutive C6 Entries
SKXD20S.	X	No Fix	ZMM/YMM Registers May Contain Incorrect Values
SKXD21S.	X	No Fix	Intel® PT PSB+ Packets May be Omitted on a C6 Transition
SKXD22S.	X	No Fix	Unexpected Uncorrected Machine Check Errors May Be Reported
SKXD23S.	X	No Fix	Intel® PT Trace May Drop Second Byte of CYC Packet
SKXD24S.	X	No Fix	Intel MBM Counters May Report System Memory Bandwidth Incorrectly
SKXD25S.	X	No Fix	Processor May Behave Unpredictably on Complex Sequence of Conditions Which Involve Branches That Cross 64 Byte Boundaries
SKXD26S.			Removed
SKXD27S.	X	No Fix	DMI and PCIe Interfaces May See Elevated Bit Error Rates
SKXD28S.	X	No Fix	Unexpected Page Faults in Guest Virtualization Environment
SKXD29S.	X	No Fix	STIBP May Not Function as Intended
SKXD30S.	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation



**Table 1. CPU Errata (Sheet 5 of 5)**

Errata Number/ HSD Number	Steppings	Status	Errata
	M-1		
SKXD31S.	X	No Fix	Direct Branches With Partial Address Aliasing May Lead to Unpredictable System Behavior
SKXD32S.	X	No Fix	PCIe Function Level Reset May Generate a NMI# Exception
SKXD33S.	X	No Fix	Performance Monitoring General Counter 2 May Have Invalid Value Written When TSX Is Enabled
SKXD34S.	X	No Fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of the Same Priority Completes
SKXD35S.	X	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May Be Incorrectly Set
SKXD36S.	X	No Fix	A Fixed Interrupt May Be Lost When a Core Exits C6



## Intel® Xeon® D-2100 Processor Product Family PCH Errata

**Table 2. PCH Errata (Sheet 1 of 3)**

Errata Number/ HSD Number	Steppings		Status	Errata
	B2	S1		
LBG1.	X	X	No Fix	The USB3.0 Kernel Debugger And DCI Can Not Be Used At The Same Time
LBG2.	X	X	No Fix	Hot-Plugging And Unplugging Of The DBC Cable Can Lead To An Inoperative Debug Port
LBG3.	X	X	No Fix	Intel® Ethernet Connection X722 PRTDCB_RUP2TC and PRTDCB_TC2PFC Are Not Writable
LBG4.	X	X	No Fix	Manageability Checksum Filtering Of IPv6 Packets In Intel Ethernet Connection X722
LBG5.	X	X	No Fix	INTENA_MSK Setting Might Clear Interrupt In Intel Ethernet Connection X722
LBG6.	X	X	No Fix	Legacy SMBus In The Intel Ethernet Connection X722 Timeout Mechanism Is Not Functional When Not Using ARA Cycle
LBG7.	X	X	No Fix	Intel Ethernet Connection X722 Illegal Byte Error Statistical Counter Inaccuracy
LBG8.	X	X	No Fix	Intel Ethernet Connection X722 Immediate Interrupts May Delay In System
LBG9.	X	X	No Fix	Intel Ethernet Connection X722 L2 Tag Stored In The Wrong RX Descriptor Field
LBG10.	X	X	No Fix	Intel Ethernet Connection X722 Jabber Packets Are Not Counted In Jabber Packets Received Field Of NC-SI Get Controller Packet Statistics Command
LBG11.	X	X	No Fix	HDA Multiple IDMAs Could Cause Audio Corruption
LBG12.	X	X	No Fix	Intel Ethernet Connection X722 TX Performance Degradation For Small Cloud Packets
LBG13.	X	X	No Fix	Intel Ethernet Connection X722 TX Descriptor Might Be Read Twice
LBG14.	X	X	No Fix	Intel Ethernet Connection X722 ECRC Bits Are Not RO When ECRC Is Disable
LBG15.	X	X	No Fix	Intel Ethernet Connection X722 Receive Performance Degradation With Specific Cloud Header
LBG16.	X	X	No Fix	Intel Ethernet Connection X722 Full Switching Table Might Reduce Small Packets Performance
LBG17.	X	X	No Fix	Transaction Pending Bit Is Not Functional In The Intel Ethernet Connection X722
LBG18.	X	X	No Fix	Intel® QuickAssist Technology (Intel® QAT): Signaled System Error (SSE) Bit Not Cleared



**Table 2. PCH Errata (Sheet 2 of 3)**

Errata Number/ HSD Number	Steppings		Status	Errata
	B2	S1		
LBG19.	X	X	No Fix	eSPI's Virtual Wire (VW) Chip Select Counter Is Not Resetting on Slave Link And Channel Recovery (SLCR) De-assertion
LBG20.	X	X	No Fix	Intel QuickAssist Technology Endpoint: (P/V) PAERCTLCAP.TFEP Cannot Be Cleared
LBG21.	X	X	No Fix	WAKE# Assertion Does Not Set PCI Express* Root Port Wake Status (PCIEXP_WAKE_STS)
LBG22.	X	X	No Fix	xHCI Host Controller Reset May Cause A System Hang
LBG23.	X	X	No Fix	Outstanding eSPI SMI/SCI Not Cleared By Warm Reset
LBG24.	X	X	No Fix	PCIe Root Port Interface Not Going Into Loopback Mode
LBG25.	X	X	No Fix	Second eSPI Slave VW And Link Error Cause Registers May Not Update
LBG26.	X	X	No Fix	xHCI Controller Does Not flag Parity Error When A Poison Packet Is Received
LBG27.	X	X	No Fix	OOB PECI (PECI Over PECI Wire) Failing On Boards With High Capacitance Load
LBG28.	X	X	No Fix	eSPI Master May Not Service an ALERT From Slave
LBG29.	X	X	No Fix	The Intel Ethernet Connection X722 Transmitter Does Not Conform To IEEE* 802.3 Clause 72 KR Electrical Specification For Co-Efficient Update
LBG30.	X	X	No Fix	The Intel Ethernet Connection X722 Transmitter Transition Time Does Not Conform To IEEE 802.3 Specification
LBG31.	X	X	No Fix	Sending Data After A RDMA Read In The Intel Ethernet Connection X722 Limited To Less Than 2G
LBG32.	X	X	No Fix	Intel Ethernet Connection X722 Work Request Size
LBG33.	X	X	No Fix	Function-level Reset Fails To Complete In The Intel Ethernet Connection X722
LBG34.	X	X	No Fix	Intel Ethernet Connection X722 LAN Function Disabled By BIOS Might Respond To Management Commands
LBG35.	X	X	No Fix	eSPI 48 MHz Clock Timings
LBG36.	X	X	No Fix	PCI Express Unexpected Completion Status Bit May Get Set On PCIe Root Port
LBG37.	X	X	No Fix	Intel Ethernet Connection X722 MNG Packets Are Dropped While A Function Level Reset To Physical Function 0 (PF 0) Is In Progress
LBG38.	X	X	No Fix	Get Link Status AQ (Admin Queue) Command Might Return Incorrect Status In the Intel Ethernet Connection X722
LBG39.	X	X	No Fix	Intel Ethernet Connection X722 Sticky CFG Space CSRs Are Cleared On Secondary-Bus-Reset When AUX_PWR Is Disabled
LBG40.	X	X	No Fix	Intel Ethernet Connection X722 Legacy Interrupt Config Space Status Bit Not Implemented
LBG41.	X	X	No Fix	I2C Time Between Start and Stop Not Meeting Spec for Intel Ethernet Connection X722 Buses



**Table 2. PCH Errata (Sheet 3 of 3)**

Errata Number/ HSD Number	Steppings		Status	Errata
	B2	S1		
LBG42.	X	X	No Fix	Glitch On GPIO During GLOBR On Intel Ethernet Connection X722 Pins
LBG43.	X	X	No Fix	Intel Ethernet Connection X722 GLQF_PCNT Counters Do Not Wrap Around
LBG44.	X	X	No Fix	PCH Does Not Meet Charged Device Model (CDM) ESD Specification
LBG45.	X	X	No Fix	MSR 0xC80 1A_32_DEBUG_INTERFACE_MSR Enable May Be SET After Clear CMOS
LBG46.	X	X	No Fix	Intel Ethernet Connection X722 RDMA SGE Count Limitations
LBG47.	X	X	No Fix	Program Suspend Instruction and Program Resume Instruction Fields Are Not Used By SPI Controller
LBG48.	X	X	No Fix	Hang On CF9 06 Reset In POST
LBG49.	X	X	No Fix	Intel Ethernet Connection X722 Device Unable To Recover When All Ports Are Disabled
LBG50.	X	X	No Fix	BMC Shared NIC Slow Response In Heavy Network Traffic With Intel Ethernet Connection X722
LBG51.	X	X	No Fix	Memory/IO Reads Targeting BMC On eSPI May Be Completed Incorrectly
LBG52.	X	X	No Fix	System Hangs After BIOS/Intel® Server Platform Services Firmware Flash Update Completes
LBG53.	X	X	No Fix	Legacy GbE Can Cause The PCH To Hang During Boot
LBG54.	X	X	No Fix	Do Not Access Parent Memory Region While Memory Windows Are Active With Intel Ethernet Connection X722 RDMA Applications
LBG55.	X	X	No Fix	Intel Ethernet Connection X722 Activity LED May Blink Regardless If Link Is Up Or Down For A Port
LBG56.	X	X	No Fix	EOI Broadcast From The CPU May Cause Errors To Be Reported On The x16 Uplink
LBG57.	X	X	No Fix	MCTP Broadcast Messages To IE, Intel Ethernet Connection X722 And Intel QuickAssist Technology May Cause Errors To Be Reported
LBG58.	X	X	No Fix	In EFI, Multiple Link Status Change Events In the X722 Ethernet Connection Might Cause An AEN Storm To the BMC
LBG59.	X	X	No Fix	3.3V Deep Sleep Rail Bleed Voltage Onto 1.8V Rail
LBG60.	X	X	No Fix	10GBASE-KR Link Establishment May be Impacted if Link Partner Issues PRESET Request
LBG61.	X	X	No Fix	xHCI Short Packet Event Using Non-Event Data TRB



## Identification Information

---

### Component Identification via Programming Interface

The Intel® Xeon® D-2100 Processor Product Family stepping can be identified by the following register contents:

Reserved	Extended Family <sup>1</sup>	Extended Model <sup>2</sup>	Reserved	Processor Type <sup>3</sup>	Family Code <sup>4</sup>	Model Number <sup>5</sup>	Stepping ID <sup>6</sup>
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0101b		00b	0110b	0101b	L0 = 010b

**Notes:**

1. The Extended Family, bits [27:20] are used in conjunction with the Family Code, specified in bits [11:8], to indicate where the processor belongs to.
2. The Extended Model, bits [19:16] in conjunction with the Model Number, specified in bits [7:4], are used to identify the model of the processor within the processor's family.
3. The Family Code corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
4. The Model Number corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
5. The Stepping ID in bits [3:0] indicates the revision number of that model. See [Table 2](#) for the processor stepping ID number in the CPUID information.

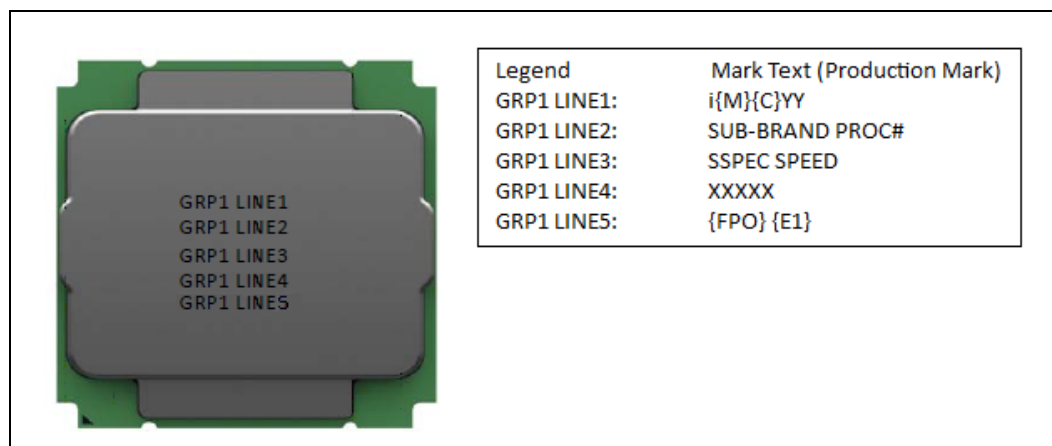
When EAX is initialized to a value of '1', the CPUID instruction returns the *Extended Family, Extended Model, Processor Type, Family Code, Model Number and Stepping ID* value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.





## Component Marking Information

**Figure 1. Intel® Xeon® D-2100 Processor Product Family Top-Side Markings (Example)**



The Intel® Xeon® D-2100 Processor Product Family stepping can be identified by the following component markings.

For Intel® Xeon® D-2100 Processor Product Family SKUs, see

<https://ark.intel.com/products/series/87041/Intel-Xeon-D-Processor>.

**Table 3. Specification Changes**

Number	Specification Change
NA	None to report at this time.

**Table 4. Specification Clarifications**

Number	Specification Clarification
NA	None to report at this time.

**Table 5. Documentation Changes**

Number	Documentation Change
NA	None to report at this time.

§ §



## CPU Errata

---

### **SKXD1. A CAP Error While Entering Package C6 Might Cause DRAM to Fail to Enter Self-Refresh**

**Problem:** A Command/Address Parity (CAP) error that occurs on the command to direct DRAM to enter self-refresh might cause the DRAM to fail to enter self-refresh although the processor enters Package-C6.

**Implication:** Due to this erratum, DRAM might fail to be refreshed, which might result in uncorrected errors being reported from the DRAM.

**Workaround:** None.

**Status:** No Fix.

### **SKXD2. PCIe\* Lane Error Status Register Might Log False Correctable Error**

**Problem:** PCIe LNERSTS (Device 0; Function 0; Offset 258h; bits [3:0]) might log false lane-based correctable errors.

**Implication:** Diagnostics cannot reliably use LNERSTS to report correctable errors.

**Workaround:** None.

**Status:** No Fix.

### **SKXD3. In Memory Mirror Mode, DataErrorChunk Field Might be Incorrect**

**Problem:** In Memory Mirror Mode, DataErrorChunk bits (IA32\_MC7\_MISC register MSR(41FH) bits [61:60]) might not correctly report the chunk containing an error.

**Implication:** Due to this erratum, this field is not accurate when Memory Mirror Mode is enabled.

**Workaround:** None.

**Status:** No Fix.

### **SKXD4. PCIe Port Might Incorrectly Log Malformed\_TLP Error**

**Problem:** If the PCIe port receives a TLP that triggers both a Malformed\_TLP error and an ECRC\_TLP error, the processor should only log an ECRC\_TLP error. However, the processor logs both errors.

**Implication:** Due to this erratum, the processor may incorrectly log Malformed\_TLP errors.

**Workaround:** None.

**Status:** No Fix.



### **SKXD5. CMCI Might Not be Signaled for Corrected Error**

**Problem:** Machine check banks 9, 10, and 11 might not signal Corrected Machine Check Interrupt (CMCI) after the first corrected error is reported in the bank even if the MCI\_STATUS register has been cleared.

**Implication:** After the first corrected error is reported in one of the affected machine check banks, subsequent errors will be logged but may not result in a CMCI.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD6. Intel® CAT/CDP Might Not Restrict Cacheline Allocation Under Certain Conditions**

**Problem:** Under certain microarchitectural conditions involving heavy memory traffic, cache lines might fill outside the allocated L3 capacity bitmask (CBM) associated with the current Class of Service (CLOS).

**Implication:** Cache Allocation Technology/Code and Data Prioritization (CAT/CDP) might see performance side effects and a reduction in the effectiveness of the CAT feature for certain classes of applications, including cache-sensitive workloads than seen on previous platforms.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD7. Credits Not Returned For PCIe Packets That Fail ECRC Check Problem**

**Problem:** The processor's IIO does not return credits back to the PCIe link in case of end-to-end CRC (ECRC) errors.

**Implication:** Due to this erratum, the link might experience degraded performance or might eventually fail due to a loss of credits.

**Workaround:** For processors that support Live Error Recovery (LER) the link would be reset and credits would be restored. Processors that do not support LER should configure ECRC errors to be fatal.

**Status:** No Fix.

### **SKXD8. PCIe Link Might Fail to Train**

**Problem:** When a pin on a PCIe lane is not connected to the link partner, the PCIe port's LTSSM might hang in the detect state.

**Implication:** When this erratum occurs, the PCIe link fails to train and the corresponding link partner(s) are not enumerated.

**Workaround:** None.

**Status:** No Fix.

### **SKXD9. IODC Entry 0 Cannot be Masked**

**Problem:** The individual I/O Directory Cache (IODC) Entry 0 cannot be masked using HA\_COH\_CFG\_1, (Bus 1; Devices 11-8; Functions 7-0, Offset 0x11C, bit 0); therefore, Entry 0 is always allocated.

**Implication:** No functional impact is observed.

**Workaround:** None.

**Status:** No Fix.



**SKXD10. With eMCA2 Enabled a 3-Strike Might Cause an Unnecessary CATERR# Instead of Only MSMI**

**Problem:** When eMCA2 is enabled to cause an MSMI due to a 3-strike event, a pulsed CATERR# and MSMI# event might both be observed on the pins.

**Implication:** When this erratum occurs, an unnecessary CATERR# pulse might be observed.

**Workaround:** None.

**Status:** No Fix.

**SKXD11. Removed**

**SKXD12. CSRs SVID and SDID Are Not Implemented For Some DDRIO and PCU Devices**

**Problem:** The DDRIO (Bus: 3; Device {19,22}; Function {6,7} and Bus: 0; Device: {20,23}; Function: {4,5,6,7};) and PCU (Bus: 3; Device 31; Functions {0,2}) do not implement the SVID (Offset 0x2C) and SDID (Offset 0x2E) CSRs. Read accesses to these register locations return all zeros.

**Implication:** Software relying on DDRIO and PCU SVID and SDID CSR support might not function correctly.

**Workaround:** Do not use SVID and SDID for these devices and functions.

**Status:** No Fix.

**SKXD13. Register Broadcast Read From DDRIO May Return a Zero Value**

**Problem:** When performing a BIOS broadcast register read to DDRIO a value of zero is always returned.

**Implication:** When this erratum occurs, BIOS might not be able to proceed due to always reading a value of zero.

**Workaround:** Use unicast register read for each instance instead of broadcast register read for all instances at once.

**Status:** No Fix.

**SKXD14. Intel® CMT Counters May Not Count Accurately**

**Problem:** Under complex micro-architectural conditions, the Cache Monitoring Technology (CMT) counters might over-count.

**Implication:** Software relying on CMT registers to enable resource allocation might not operate correctly. This can lead to reporting of more cachelines used than the cache supports or the counter wrapping and returning a too small value. WBINVD might not result in the CMT counters being zeroed. Intel has not observed this erratum in commercially available software.

**Workaround:** None.

**Status:** No Fix.

**SKXD15. Intel® CAT Might Not Restrict Cacheline Allocation Under Certain Conditions**

**Problem:** Under certain micro-architectural conditions involving heavy memory traffic, cachelines might fill outside the allocated L3 Capacity Bit-Mask (CBM) associated with the current Class of Service (CLOS).

**Implication:** CAT might appear less effective at protecting certain classes of applications, including cache-sensitive workloads than on previous platforms.

**Workaround:** None Identified.

**Status:** No Fix.



### **SKXD16. Intel PCIe Corrected Error Threshold Does Not Consider Overflow Count When Incrementing Error Counter**

**Problem:** The PCIe corrected error counter feature does not take the overflow bit in the count (bit 15 of XPCORERRCOUNTER (Bus; RootBus Device; 0 Function; 0 Offset; 4D0h)) into account when comparing the count to the threshold in XPCORERRTHRESHOLD.ERROR\_THRESHOLD. Therefore, users end up with another interrupt once the counter has rolled over and hit the threshold + 0x8000.

**Implication:** Due to this erratum, the PCIe corrected error signaling might occur even after the error count has exceeded the corrected error count threshold, not just a single time when reaching the threshold. Intel has not observed this erratum with any commercially available system.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD17. IIO RAS VPP Hangs During The Warm Reset Test**

**Problem:** When VPPCL bit 0 of VPP\_reset\_Mode (Bus 1; Device 30; Function 5; Offset 0xF0) bit is set to 0, and the CPU is undergoing reset flow while PCIe hot-plug operation is in process, the Virtual Pin Port (VPP) hot-plug commands might stop responding.

**Implication:** Due to this erratum, during CPU reset hot-plug commands might not complete.

**Workaround:** Do not set VPP reset mode to zero.

**Status:** No Fix.

### **SKXD18. A Core 3-Strike Event May Be Seen Under Certain Test Conditions**

**Problem:** When running some stress tests and/or related applications, a core 3-strike event may be seen. This similar 3-strike event may also occur when system is at idle.

**Implication:** A core 3-strike event may be seen resulting in a system hang and/or a shutdown.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD19. DDR4 Memory Bandwidth May be Lower Than Expected at 2133 and 1866 MHz**

**Problem:** A DDR4 transaction credit imbalance between memory controllers may result in a lower than expected available memory bandwidth.

**Implication:** Due to this erratum, DDR4 Memory Bandwidth may be lower than expected at 2133 and 1866 MHz.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD20. Intel Sparing Per-Rank Error Masking Does Not Mask Correctable Errors**

**Problem:** The Integrated Memory Controller (IMC) Sparing Per-Rank Error Masking (PREM) capability does not mask off correctable error logging and signaling as expected.

**Implication:** Due to this erratum, errors will continue to be logged and signaled despite per-rank error masking. Per-rank error counters are still masked.

**Workaround:** None Identified

**Status:** No Fix.



### **SKXD21. PCIe Root Port Electromechanical Interlock Control Register Can Be Written**

**Problem:** Electromechanical Interlock Control (bit 11) in the Slot Control register (B: Root Port; D: 0-3; F: 0 bits offset 0x18) in the PCIe Capability table should be read-only and always return 0. Due to this erratum, this register can be written.

**Implication:** Writes to this bit can cause later reads to return the written value. However, this has no other effect on functionality.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD22. Performance Monitoring M2MEM Counters For Memory Controller Reads/Writes Are Not Counting Read/Write Retries**

**Problem:** PMON M2MEM counters for read and write events do not account for scrub reads and scrub writes during the error flow.

**Implication:** Due to this erratum, a mismatch in the counters for Read/Write retries in M2MEM and iMC (integrated memory controller) may be observed.

**Workaround:** When doing error injection testing, counting reads and writes in the presence of ECC errors will only be precise using the iMC counter, not the M2MEM counter

**Status:** No Fix.

### **SKXD23. System Hangs May Occur When IPQ and IRQ Requests Happen at The Same Time**

**Problem:** When IPQ and IRQ requests happen at the same time, and the IPQ request is starved due to PAMatch/NotAllowSnoop on a Table of Request ID (TORID) then the IRQ request that is waiting for the TORID's SF/LLC may become invalid.

**Implication:** Due to this erratum, if IPQ and IRQ requests do not need to snoop any cores, then IPQ requests may block IRQ requests resulting in a system hang. Intel has only observed this erratum in a synthetic test environment.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD24. IIO VPP May Hang During Warm Reset**

**Problem:** When VPP\_Reset\_Mode bit 0 of VPPCTL (Bus 1; Device 30; Function 5; Offset 0xF0) is set to 0, and there is a PCIe hot-plug event in progress, if the processor performs a warm reset, the Virtual Pin Port hot-plug flow may hang.

**Implication:** Due to this erratum, the Virtual Pin Port may hang.

**Workaround:** Do not set VPP\_Reset\_Mode to 0.

**Status:** No Fix.

### **SKXD25. Machine Check Events May Be Logged In Banks 9, 10 And 11 That Do Not Represent Actual Errors**

**Problem:** In some previous CPU Microcode + BIOS code combinations MCEs in banks 9, 10 and 11 may be seen. These do not represent actual errors and normally are processed out by early BIOS execution.

**Implication:** MCEs may be seen on banks 9, 10 and 11 that represent incorrect error data. These MCEs have the potential to be forwarded to the OS and may be end-user visible while not representing actual errors.

**Workaround:** None Identified.

**Status:** No Fix.



### **SKXD26. Lower Than Expected Performance May Be Seen With Some Intel AVX Workloads Due To Incorrect Uncore Frequency Scaling**

**Problem:** Due to a problem with Uncore Frequency Scaling (UFS), lower than expected performance may be seen with some Intel AVX workloads. The CPU may not ramp uncore frequency when running some Intel AVX workloads depending on the number of active cores.

**Implication:** Lower than expected performance may be seen with some Intel AVX workloads due to incorrect uncore frequency scaling.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD27. Spurious Corrected Errors May Be Reported**

**Problem:** Due to this erratum, spurious corrected errors may be logged in the IA32\_MC0\_STATUS MSR (401H) register with the valid field (bit 63) set, the uncorrected error field bit (bit 61) not set, a Model Specific Error Code (bits [31:16]) of 0x0001, and an MCA Error Code (bits [15:0]) of 0x0005. If CMCI is enabled, these spurious corrected errors also signal interrupts.

**Implication:** When this erratum occurs, software may see an unusually high rate of reported corrected errors. As it is not possible to distinguish between spurious and non-spurious errors, this erratum may interfere with reporting non-spurious corrected errors.

**Workaround:** The BIOS code integrated with Microcode 0x30 contains a workaround for this erratum.

**Status:** No Fix.

### **SKXD28. Writing To LT\_LOCK\_MEMORY And LT\_UNLOCK\_MEMORY MSRs Simultaneously May Have Inconsistent Results**

**Problem:** Writing to LT\_LOCK\_MEMORY MSR (2e7H) and to LT\_UNLOCK\_MEMORY MSR (2e6H) simultaneously from different physical cores may have inconsistent results. Some of the memory ranges may get locked as requested by the write to LT\_LOCK\_MEMORY MSR while some may get unlocked as requested by the write to LT\_UNLOCK\_MEMORY MSR.

**Implication:** Writing to LT\_LOCK\_MEMORY MSR and to LT\_UNLOCK\_MEMORY MSRs may not operate as expected if they are done on different cores simultaneously. Intel has not observed this erratum in any commercially available system.

**Workaround:** None Identified. Software (BIOS) should write to these MSRs only on the BSP (boot strap processor).

**Status:** No Fix.

### **SKXD29. Masked Bytes In A Vector Masked Store Instructions May Cause Write Back Of A Cache Line**

**Problem:** Vector masked store instructions to WB (write-back) memory-type that cross cache lines may lead to CPU writing back cached data even for cache lines where all of the bytes are masked.

**Implication:** The processor may generate writes of un-modified data. This can affect Memory Mapped I/O (MMIO) or non-coherent agents in the following ways:

1. For MMIO range that is mapped as WB memory type, this erratum may lead to Machine Check Exception (MCE) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.
2. If the CPU cached data is stale, for example in the case of memory written directly by a non-coherent agent (agent that uses non-coherent writes), this erratum may lead to writing back stale cached data even if these bytes are masked.



Workaround: Platforms should not map MMIO memory space or non-coherent device memory space as WB memory. If WB is used for MMIO range, software or VMM should not map such MMIO page adjacent to a regular WB page (adjacent on the linear address space, before or after the I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

Status: No Fix.

### **SKXD30. ERROR\_N[2:0] Pins May Not Be Cleared After a Warm Reset**

Problem: The processor's ERROR\_N[2:0] pins may not be cleared after a warm reset.

Implication: Due to this erratum, the ERROR\_N[2:0] pins may incorrectly indicate a pending error after a warm reset.

Workaround: The BIOS can contain code changes to work around this erratum.

Status: No Fix.

### **SKXD31. CRC Store Operation Corner Case May Result In Hang**

Problem: Intel® QuickData Technology Local and Remote CRC Store operations may result in a DMA channel hang when the CRC Store transfer size is less than 32 bytes and the destination offset is not DWORD-aligned.

Implication: Due to this erratum, the processor may hang.

Workaround: Software must configure Intel QuickData Technology Local and Remote CRC Store operations to have descriptor destination offset addresses DWORD-aligned.

Status: No Fix.

### **SKXD32. Intel PCIe Slot Presence Detect and Presence Detect Changed Logic Not PCIe Specification Compliant**

Problem: When Hot-Plug Surprise is set in the Slot Capabilities register (Bus: RootBus, Dev: 1-3, Function: 0, Offset: A4h, Bit: 5), the Presence Detect State and Presence Detect Change in the Slot Status register (Bus: RootBus, Dev: 1-3, Function: 0, Offset: A2h), incorrectly ignores the out-of-band presence detect mechanism and only reflects the Physical Layer in-band presence detect mechanism.

Implication: Due to this erratum, if the Hot-Plug Surprise bit is set in the Slot Capabilities register, software will not be able to detect the presence of an adapter inserted while a slot is powered down. Therefore, Hot-Plug Surprise must only be set in configurations where the slot power is always enabled.

Workaround: None Identified.

Status: No Fix.

### **SKXD33. In Patrol Scrub System Address Mode, Address is Not Loaded From CSRs After Re-enable**

Problem: The patrol scrub starting address registers [scrubaddressshi (Bus 2; Devices 12, 10; Function 0; Offset 910) and scrubaddresslo Bus 2; Devices 12, 10; Function 0; Offset 90c] should indicate when the first memory address from which patrol logic should start scrubs [when scrubctl.startscrub (Bus 2; Devices 12, 10; Function 0; Offset 914; Bit 24) is set]. Due to this erratum, after patrol is disabled, if the patrol scrub engine is re-enabled in System Address Mode with scrubctl.startscrub set, the patrol scrubbing engine may ignore the starting address registers. Re-enabling patrol after S3 exit or other warm reset event is not impacted by this.

Implication: Due to this erratum, when configured in system address mode, Patrol scrubs will not start from the address specified in the starting address registers. This may cause certain memory lines to be scrubbed more or less frequently than expected. Intel has not seen this erratum to affect the operation of any commercially available software.





Workaround: None identified. Contact your Intel representative for details of possible mitigations.

Status: [No Fix.](#)

#### **SKXD34. The Corrected Error Count Overflow Bit In IA32\_MC0\_STATUS Is Not Updated When The UC Bit Is Set**

Problem: After a UC (uncorrected) error is logged in the IA32\_MC0\_STATUS MSR (401H), corrected errors will continue to be counted in the lower 14 bits (bits 51:38) of the Corrected Error Count. Due to this erratum, the sticky count overflow bit (bit 52) of the Corrected Error Count will not get updated when the UC bit (bit 61) is set to 1.

Implication: The Corrected Error Count Overflow indication will be lost if the overflow occurs after an uncorrectable error has been logged.

Workaround: None Identified.

Status: [No Fix.](#)

#### **SKXD35. SMRAM State-Save Area Above The 4 GB Boundary May Cause Unpredictable System Behavior**

Problem: If the BIOS uses the RSM instruction to load the SMBASE register with a value that would cause any part of the SMRAM state-save area to have an address above 4-GBytes, subsequent transitions into and out of System-Management Mode (SMM) might save and restore processor state from incorrect addresses.

Implication: This erratum may cause unpredictable system behavior. Intel has not observed this erratum with any commercially available system.

Workaround: Ensure that the SMRAM state-save area is located entirely below the 4 GB address boundary.

Status: [No Fix.](#)

#### **SKXD36. POPCNT Instruction May Take Longer To Execute Than Expected**

Problem: POPCNT instruction execution with a 32 or 64 bit operand may be delayed until previous non-dependent instructions have executed.

Implication: Software using the POPCNT instruction may experience lower performance than expected.

Workaround: None Identified.

Status: [No Fix.](#)

#### **SKXD37. Load Latency Performance Monitoring Facility May Stop Counting**

Problem: The performance monitoring events MEM\_TRANS\_RETIRED.LOAD\_LATENCY\_\* (Event CDH; UMask 01H; any latency) count load instructions whose latency exceed a predefined threshold, where the loads are randomly selected using the Load Latency facility (PEBS extension). However due to this erratum, load latency facility may stop counting load instructions when Intel® Hyper-Threading Technology (Intel® HT Technology) is enabled.

Implication: Counters programmed with the affected events stop incrementing and do not generate PEBS records.

Workaround: None Identified.

Status: [No Fix.](#)



### **SKXD38. Performance Monitoring Counters May Undercount When Using CPL Filtering**

**Problem:** Performance Monitoring counters configured to count only OS or only USR events (by setting only one of bits 16 or 17 in IA32\_PERFEVTSELx) may undercount for a short cycle period of typically less than 100 processor clock cycles after the processor transitions to a new CPL. Events affected may include those counting CPL transitions (by additionally setting the edge-detect bit 18 in IA32\_PERFEVTSELx).

**Implication:** Due to this erratum, Performance Monitoring counters may report counts lower than expected.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD39. Incorrect Branch Predicted Bit In BTS/BTM Branch Records**

**Problem:** Branch Trace Store (BTS) and Branch Trace Message (BTM) send branch records to the Debug Store management area and system bus, respectively. The Branch Predicted bit (bit 4 of eighth byte in BTS/BTM records) should report whether the most recent branch was predicted correctly. Due to this erratum, the Branch Predicted bit may be incorrect.

**Implication:** BTS and BTM cannot be used to determine the accuracy of branch prediction.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD40. DR6.B0-B3 May Not Report All Breakpoints Matched When A MOV/POP SS Is Followed By A Store Or An MMX Instruction**

**Problem:** Normally, data breakpoints matches that occur on a MOV SS, r/m or POP SS will not cause a debug exception immediately after MOV/POP SS but will be delayed until the instruction boundary following the next instruction is reached. After the debug exception occurs, DR6.B0-B3 bits will contain information about data breakpoints matched during the MOV/POP SS as well as breakpoints detected by the following instruction. Due to this erratum, DR6.B0-B3 bits may not contain information about data breakpoints matched during the MOV/POP SS when the following instruction is either an MMX instruction that uses a memory addressing mode with an index or a store instruction.

**Implication:** When this erratum occurs, DR6 may not contain information about all breakpoints matched. This erratum will not be observed under the recommended usage of the MOV SS,r/m or POP SS instructions (i.e., following them only with an instruction that writes (E/R)SP).

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD41. Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions**

**Problem:** The performance monitoring events MEM\_TRANS\_RETIRED.LOAD\_LATENCY\_\* (Event CDH; UMask 01H; any latency) count load instructions whose latency exceed a predefined threshold, where the loads are randomly selected using the load latency facility (an extension of PEBS). However, due to this erratum, these events may count incorrectly for VGATHER\*/VPGATHER\* instructions.

**Implication:** The Load Latency Performance Monitoring events may be Inaccurate for Gather instructions.

**Workaround:** None Identified.

**Status:** No Fix.



#### **SKXD42. VM Exit May Set IA32\_EFER.NXE When IA32\_MISC\_ENABLE Bit 34 Is Set To 1**

**Problem:** When “XD Bit Disable” in the IA32\_MISC\_ENABLE MSR (1A0H) bit 34 is set to 1, it should not be possible to enable the “execute disable” feature by setting IA32\_EFER.NXE. Due to this erratum, a VM exit that occurs with the 1-setting of the “load IA32\_EFER” VM-exit control may set IA32\_EFER.NXE even if IA32\_MISC\_ENABLE bit 34 is set to 1. This erratum can occur only if IA32\_MISC\_ENABLE bit 34 was set by guest software in VMX non-root operation.

**Implication:** Software in VMX root operation may execute with the “execute disable” feature enabled despite the fact that the feature should be disabled by the IA32\_MISC\_ENABLE MSR. Intel has not observed this erratum with any commercially available software.

**Workaround:** A virtual-machine monitor should not allow guest software to write to the IA32\_MISC\_ENABLE MSR.

**Status:** No Fix.

#### **SKXD43. x87 FPU Exception (#MF) May Be Signaled Earlier Than Expected**

**Problem:** x87 instructions that trigger #MF normally service interrupts before the #MF. Due to this erratum, if an instruction that triggers #MF is executing when an Enhanced Intel SpeedStep® Technology transitions, an Intel® Turbo Boost Technology transitions, or a Thermal Monitor events occurs, the #MF may be taken before pending interrupts are serviced.

**Implication:** Software may observe #MF being signaled before pending interrupts are serviced.

**Workaround:** None Identified.

**Status:** No Fix.

#### **SKXD44. CPUID TLB Associativity Information Is Inaccurate**

**Problem:** CPUID leaf 2 (EAX=02H) TLB information inaccurately reports that the shared second-Level TLB is 6-way set associative (value C3H), although it is 12-way set associative. Other information reported by CPUID leaf 2 is accurate.

**Implication:** Software that uses CPUID shared second-level TLB associativity information for value C3H may operate incorrectly. Intel has not observed this erratum to impact the operation of any commercially available software.

**Workaround:** None Identified. Software should ignore the shared second-Level TLB associativity information reported by CPUID for the affected processors.

**Status:** No Fix.

#### **SKXD45. Vector Masked Store Instructions May Cause Write Back Of Cache Line Where Bytes Are Masked**

**Problem:** Vector masked store instructions to write-back (WB) memory-type that cross cache lines may lead to CPU writing back cached data even for cache lines where all of the bytes are masked.

This can affect Memory Mapped I/O (MMIO) or non-coherent agents in the following ways:

1. For MMIO range that is mapped as WB memory type, this erratum may lead to Machine Check Exception (MCE) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.
2. If the CPU cached data is stale, for example in the case of memory written directly by a non-coherent agent (agent that uses non-coherent writes), this erratum may lead to writing back stale cached data even if these bytes are masked.

**Implication:** CPU may generate writes into MMIO space which lead to MCE, or may write stale data into memory also written by non-coherent agents.



Workaround: It is recommended not to map MMIO range as WB. If WB is used for MMIO range, OS or VMM should not map such MMIO page adjacent to a regular WB page (adjacent on the linear address space, before or after the I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

Status: [No Fix.](#)

#### **SKXD46. Incorrect FROM\_IP Value For An RTM Abort In BTM Or BTS May Be Observed**

Problem: During Restricted Transactional Memory (RTM) operation when branch tracing is enabled using Branch Trace Message (BTM) or Branch Trace Store (BTS), the incorrect EIP value (From\_IP pointer) may be observed for an RTM abort.

Implication: Due to this erratum, the From\_IP pointer may be the same as that of the immediately preceding taken branch.

Workaround: None Identified.

Status: [No Fix.](#)

#### **SKXD47. MOVNTDQA From WC Memory May Pass Earlier Locked Instructions**

Problem: An execution of (V)MOVNTDQA (streaming load instruction) that loads from Write Combining (WC) memory may appear to pass an earlier locked instruction to a different cache line.

Implication: Software that expects a lock to fence subsequent (V)MOVNTDQA instructions may not operate properly.

Workaround: Software should not rely on a locked instruction to fence subsequent executions of MOVNTDQA. Software should insert an MFENCE instruction if it needs to preserve order between streaming loads and other memory operations.

Status: [No Fix.](#)

#### **SKXD48. #GP On Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code**

Problem: During a #GP (General Protection Exception), the processor pushes an error code on to the exception handler's stack. If the segment selector descriptor straddles the canonical boundary, the error code pushed onto the stack may be incorrect.

Implication: An incorrect error code may be pushed onto the stack. Intel has not observed this erratum with any commercially available software.

Workaround: None Identified.

Status: [No Fix.](#)



### **SKXD49. BNDLDX And BNDSTX May Not Signal #GP On Non-Canonical Bound Directory Access**

**Problem:** BNDLDX and BNDSTX instructions access the bound's directory and table to load or store bounds. These accesses should signal #GP (general protection exception) when the address is not canonical (i.e., bits 48 to 63 are not the sign extension of bit 47). Due to this erratum, #GP may not be generated by the processor when a non-canonical address is used by BNDLDX or BNDSTX for their bound directory memory access.

**Implication:** Intel has not observed this erratum with any commercially available software.

**Workaround:** Software should use canonical addresses for bound directory accesses.

**Status:** No Fix.

### **SKXD50. Performance Monitor Event For Outstanding Offcore Requests May be Incorrect**

**Problem:** The performance monitor event OFFCORE\_REQUESTS\_OUTSTANDING (Event 60H, any Umask Value) should count the number of offcore outstanding transactions each cycle. Due to this erratum, the counts may be higher or lower than expected.

**Implication:** The performance monitor event OFFCORE\_REQUESTS\_OUTSTANDING may reflect an incorrect count.

**Workaround:** None Identified.

**Status:** No Fix.

### **SKXD51. Branch Instructions May Initialize MPX Bound Registers Incorrectly**

**Problem:** Depending on the current Intel® Memory Protection Extensions (Intel® MPX) configuration, execution of certain branch instructions (near CALL, near RET, near JMP, and Jcc instructions) without a BND prefix (F2H) initialize the MPX bound registers. Due to this erratum, execution of such a branch instruction on a user-mode page may not use the MPX configuration register appropriate to the current privilege level (BNDCFGU for CPL 3 or BNDCFGS otherwise) for determining whether to initialize the bound registers; it may thus initialize the bound registers when it should not, or fail to initialize them when it should.

**Implication:** After a branch instruction on a user-mode page has executed, a #BR (bound-range) exception may occur when it should not have or a #BR may not occur when one should have.

**Workaround:** If supervisor software is not expected to execute instructions on user-mode pages, software can avoid this erratum by setting CR4.SMEP[bit 20] to enable supervisor-mode execution prevention (SMEP). If SMEP is not available or if supervisor software is expected to execute instructions on user-mode pages, no workaround is identified.

**Status:** No Fix.

### **SKXD52. A Spurious APIC Timer Interrupt May Occur After Timed MWAIT**

**Problem:** Due to this erratum, a Timed MWAIT that completes for a reason other than the Timestamp Counter reaching the target value may be followed by a spurious APIC timer interrupt. This erratum can occur only if the APIC timer is in TSC-deadline mode and only if the mask bit is clear in the LVT Timer Register.

**Implication:** Spurious APIC timer interrupts may occur when the APIC timer is in TSC-deadline mode.

**Workaround:** TSC-deadline timer interrupt service routines should detect and deal with spurious interrupts.

**Status:** No Fix.



### **SKXD53. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions**

**Problem:** An access to a Guest-Physical Address (GPA) may cause an EPT-violation VM exit. When the "EPT-violation #VE" VM-execution control is 1, an EPT violation may cause a #VE (virtualization exception) instead of a VM exit. Due to this erratum, an EPT violation may erroneously cause a #VE when the "suppress #VE" bit is set in the EPT paging-structure entry used to map the GPA being accessed. This erratum does not apply when the "EPT-violation #VE" VM-execution control is 0 or when delivering an event through the IDT. This erratum applies only when the GPA in CR3 is used to access the root of the guest paging-structure hierarchy (or, with PAE paging, when the GPA in a PDPTe is used to access a page directory).

**Implication:** When using PAE paging mode, an EPT violation that should cause a VMexit in the VMM may instead cause a VE# in the guest. In other paging modes, in addition to delivery of the erroneous #VE, the #VE may itself cause an EPT violation, but this EPT violation will be correctly delivered to the VMM.

**Workaround:** A VMM may support an interface that guest software can invoke with the VMCALL instruction when it detects an erroneous #VE.

**Status:** No Fix.

### **SKXD54. Use of VMX TSC Scaling or TSC Offsetting Will Result in Corrupted Intel PT Packets**

**Problem:** When Intel® Processor Trace (Intel® PT) is enabled within a Virtual Machine Extensions (VMX) guest, and Time Stamp Counter (TSC) offsetting or TSC scaling is enabled for that guest, by setting primary processor-based execution control bit 3 or secondary processor-based execution control bit 25, respectively, in the Virtual Machine Control Structure (VMCS) for that guest, any TMA (TSC/MTC Alignment) packet generated will have corrupted values in the Core Timer Copy (CTC) and FastCounter fields. Additionally, the corrupted TMA packet will be followed by a bogus data byte.

**Implication:** An Intel® PT decoder will be confused when using the TMA packet to align cycle time with wall-clock time. The byte that follows the TMA will likely cause a decoder error for an unexpected or unrecognized packet.

**Workaround:** None identified. If a TMA packet with any reserved payload bits set is encountered by an Intel PT decoder it should be ignored, along with the byte that immediately follows it. Alternatively, Intel PT users may opt to disable MTC and TMA packets by clearing IA32\_RTIT\_CTL.MTCEn[bit 9].

**Status:** No Fix.

### **SKXD55. Using Intel® TSX Instructions May Lead to Unpredictable System Behavior**

**Problem:** Under complex microarchitectural conditions, software using Intel® Transactional Synchronization Extensions (Intel® TSX) may result in unpredictable system behavior. Intel has only seen this under synthetic testing conditions. Intel is not aware of any commercially available software exhibiting this behavior.

**Implication:** Due to this erratum, unpredictable system behavior may occur.

**Workaround:** It is possible for BIOS to contain a workaround for this erratum.

**Status:** No Fix.



### **SKXD56. Memory May Continue to Throttle after MEMHOT# De-Assertion**

**Problem:** When MEMHOT# is asserted by an external agent, the CPU may continue to throttle memory after MEMHOT# de-assertion.

**Implication:** When this erratum occurs, memory throttling occurs even after de-assertion of MEMHOT#.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD57. Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values**

**Problem:** When Restricted Transactional Memory (RTM) is supported (CPUID.07H.EBX.RTM [bit 11] = 1) and when TSX\_FORCE\_ABORT=0, Performance Monitor Unit (PMU) general purpose counter 3 (IA32\_PMC3, MSR C4H and IA32\_A\_PMC3, MSR 4C4H) may contain unexpected values. Further, IA32\_PREFEVTSEL3 (MSR 189H) and IA32\_PERF\_GLOBAL\_CTRL[3] (MSR 38FH) may contain unexpected configuration values; which may also affect IA32\_PERF\_GLOBAL\_INUSE[3] (MSR 392H).

**Implication:** Due to this erratum, software that uses PMU general purposes counter 3 may read an unexpected count and configuration.

**Workaround:** Software can avoid this erratum by writing 1 to bit 0 of TSX\_FORCE\_ABORT (MSR 10FH) which will cause all Restricted Transactional Memory (RTM) transactions to abort with EAX code 0. TSX\_FORCE\_ABORT MSR is available when CPUID.07H.EDX[bit 13]=1.

**Status:** No Fix.

### **SKXD1S. Certain PCIe Cards May Not Function as Expected When Executing a Hot Plug Add/Remove Sequence**

**Problem:** On a hot add/remove, certain PCIe cards have a slower than expected impedance ramp down time. In some cases, the ramp down time may exceed detection limits and a completion timeout may occur.

**Implication:** A completion timeout may be seen when executing a Hot Plug sequence on some PCIe cards.

**Workaround:** It is possible for a BIOS code change to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD2S. PROCHOT Input to Drive Processor Throttling May Not Function As Intended**

**Problem:** When a platform is powered-on with PROCHOT asserted, the PCU is already out of reset before distributed Power Management Agents (PMAs) are also out of reset. This may lead to a scenario where the PMAs do not respond to the, "throttle to processor cores," event sent by PCU based on a PROCHOT input.

**Implication:** PROCHOT input to drive processor throttling may not function as intended.

**Workaround:** None Identified.

**Status:** No Fix.



### **SKXD3S. Reading Some C-state Residency MSRs May Result in Unpredictable System Behavior**

**Problem:** Under complex microarchitectural conditions, an MSR read of MSR\_CORE\_C3\_RESIDENCY MSR (3FCh), MSR\_CORE\_C6\_RESIDENCY MSR (3FDh), or MSR\_CORE\_C7\_RESIDENCY MSR (3FEh) may result in unpredictable system behavior.

**Implication:** Unexpected exceptions or other unpredictable system behavior may occur.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** [No Fix.](#)

### **SKXD4S. Intel MBA Read After MSR Write May Return Incorrect Value**

**Problem:** The Memory Bandwidth Allocation (MBA) feature defines a series of MSRs (0xD50-0xD57) to specify MBA Delay Values per Class of Service (CLOS), in the IA32\_L2\_QoS\_Ext\_BW\_Thrtl\_n MSR range. Certain values when written then read back may return an incorrect value in the MSR. Specifically, values greater than or equal to 10 (decimal) and less than 39 (decimal) written to the MBA Delay Value (Bits [15:0]) may be read back as 10%.

**Implication:** The values written to the registers will be applied; however, software should be aware that an incorrect value may be returned.

**Workaround:** None Identified.

**Status:** [No Fix.](#)

### **SKXD5S. In eMCA2 Mode, When the Retirement Watchdog Timeout Occurs CATERR# May be Asserted**

**Problem:** A Retirement Watchdog Timeout (MCACOD = 0x0400) in Enhanced MCA2 (eMCA2) mode will cause the CATERR# pin to be pulsed in addition to an MSMI# pin assertion. In addition, a Machine Check Abort (#MC) will be pended in the cores along with the MSMI.

**Implication:** Due to this erratum, systems that expect to only see MSMI# will also see CATERR# pulse when a Retirement Watchdog Timeout occurs. The CATERR# pulse can be safely ignored.

**Workaround:** None Identified.

**Status:** [No Fix.](#)

### **SKXD6S. Intel MBA May Incorrectly Throttle All Threads**

**Problem:** When one logical processor is disabled, the Memory Bandwidth Allocation (MBA) feature may select an incorrect MBA throttling value to apply to the core. A disabled logical processor may behave as though the Class of Service (CLOS) field in its associated IA32\_PQR\_ASSOC MSR (0xC8F) is set to zero (appearing to be set to CLOS[0]). When this occurs, the MBA throttling value associated with CLOS[0] may be incorrectly applied to both threads on the core.

**Implication:** When Intel® Hyper-Threading Technology (Intel® HT Technology) is disabled or one logical thread on the core is disabled, the disabled thread is interpreted to have CLOS=0 set in its IA32\_PQR\_ASSOC MSR by hardware, which affects the calculation for the actual throttling value applied to the core. When this erratum occurs, the MBA throttling value associated with a given core may be incorrect.

**Workaround:** To work around this erratum, CLOS[0] should not be used if any logical cores are disabled. Alternately, software may leave all threads enabled.

**Status:** [No Fix.](#)





### **SKXD7S. Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang**

- Problem:** If an Intel® Processor Trace (Intel® PT) Table of Physical Addresses (ToPA) table is placed in UC (Uncacheable) or Uncacheable Speculative Write Combining (USWC) memory, and a ToPA output region is filled during an Intel® Transactional Synchronization Extensions (Intel® TSX) transaction, the resulting ToPA table read may cause a processor hang.
- Implication:** Placing Intel PT ToPA tables in non-cacheable memory when Intel TSX is in use may lead to a processor hang.
- Workaround:** None identified. Intel PT ToPA tables should be located in WB memory if Intel TSX is in use.
- Status:** [No Fix.](#)

### **SKXD8S. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions**

- Problem:** An access to a Guest-Physical Address (GPA) may cause an EPT-violation VM exit. When the "EPT-violation #VE" VM-execution control is 1, an EPT violation may cause a #VE (virtualization exception) instead of a VM exit. Due to this erratum, an EPT violation may erroneously cause a #VE when the "suppress #VE" bit is set in the EPT paging-structure entry used to map the GPA being accessed. This erratum does not apply when the "EPT-violation #VE" VM-execution control is 0 or when delivering an event through the IDT. This erratum applies only when the GPA in CR3 is used to access the root of the guest paging-structure hierarchy (or, with PAE paging, when the GPA in a PDPTE is used to access a page directory).
- Implication:** When using PAE paging mode, an EPT violation that should cause a VMexit in the VMM may instead cause a VE# in the guest. In other paging modes, in addition to delivery of the erroneous #VE, the #VE may itself cause an EPT violation, but this EPT violation will be correctly delivered to the VMM.
- Workaround:** A VMM may support an interface that guest software can invoke with the VMCALL instruction when it detects an erroneous #VE.
- Status:** [No Fix.](#)

### **SKXD9S. SpecPower Workloads on Native without Legacy Mode May See a Higher than Expected Loadline Power Rating**

- Problem:** When running SpecPower workload on native without legacy mode, a higher than expected loadline power may be seen compared to OOB mode. This issue was previously fixed in Skylake CPU Microcode prior to launch. The change was found to be broken in Skylake MB750654\_0200003A. It has been corrected in Skylake MB750654\_02000049.
- Implication:** SpecPower workloads on native without legacy mode may see a higher than expected loadline power rating.
- Workaround:** Intel Xeon Processor Scalable Family CPU Microcode 0x49 contains a fix for this issue.
- Status:** [No Fix.](#)



### **SKXD10S. Error Injection Testing with NVDIMM Present in System May Cause a System Reset**

**Problem:** It is possible to encounter a scenario where an error injection attempt with a NVDIMM present in the system may lead to a system reset. The NVDIMM unexpectedly may not be in self-refresh which can lead to a MSMI hold of the adjacent RDIMM.

**Implication:** Due to this erratum, a system reset may be seen.

**Workaround:** The BIOS integrated with Microcode 0x4D contains a workaround for this erratum.

**Status:** [No Fix.](#)

### **SKXD11S. MSCOD Error Code =0x2E May Be Seen During Warm Reset Testing**

**Problem:** Under certain test conditions, a system may assert an IERR when running warm reset testing. The error may be logged as a Machine Check Bank 4 MSCOD error code =0x2E.

**Implication:** A machine check may be seen during warm reset testing.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** [No Fix.](#)

### **SKXD12S. Processor May Hang When Executing Code In an HLE Transaction Region**

**Problem:** Under certain conditions, if the processor acquires an Hardware Lock Elision (HLE) lock via the XACQUIRE instruction in the Host Physical Address range between 40000000H and 403FFFFFFH, it may hang with an internal timeout error (MCACOD 0400H) logged into IA32\_MCi\_STATUS.

**Implication:** Due to this erratum, the processor may hang after acquiring a lock via XACQUIRE.

**Workaround:** The BIOS can reserve the host physical address ranges of 40000000H and 403FFFFFFH (e.g., map it as UC/MMIO). Alternatively, the Virtual Machine Monitor (VMM) can reserve that address range so no guest can use it. In non-virtualized systems, the OS can reserve that memory space.

**Status:** [No Fix.](#)

### **SKXD13S. IDI\_MISC Performance Monitoring Events May Be Inaccurate**

**Problem:** The IDI\_MISC.WB\_UPGRADE and IDI\_MISC.WB\_DOWNGRADE performance monitoring events (Event FEH; UMask 02H and 04H) counts cache lines evicted from the L2 cache. Due to this erratum, the per logical processor count may be incorrect when both logical processors on the same physical core are active. The aggregate count of both logical processors is not affected by this erratum.

**Implication:** IDI\_MISC performance monitoring events may be inaccurate.

**Workaround:** None identified.

**Status:** [No Fix.](#)



### **SKXD14S. Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB**

- Problem:** Due to a rare microarchitectural condition, generation of an Intel® Processor Trace (Intel® PT) Packet Stream Boundary (PSB) packet can cause a single CYC (Cycle Count) packet, possibly along with an associated Mini Time Counter (MTC) packet, to be dropped.
- Implication:** An Intel PT decoder that is using CYCs to track time or frequency will get an improper value due to the lost CYC packet.
- Workaround:** If an Intel PT decoder is using CYCs and MTCs to track frequency, and either the first MTC following a PSB shows that an MTC was dropped, or the CYC value appears to be 4095 cycles short of what is expected, the CYC value associated with that MTC should not be used. The decoder should wait for the next MTC before measuring frequency again.
- Status:** [No Fix.](#)

### **SKXD15S. Intel® PT VM-entry Indication Depends on the Incorrect VMCS Control Field**

- Problem:** An Intel® Processor Trace Paging Information Packet (PIP), which includes indication of entry into non-root operation, will be generated on VM-entry as long as the "Conceal VMX in Intel PT" field (bit 19) in Secondary Execution Control register (IA32\_VMX\_PROCBASED\_CTLSS2, MSR 048BH) is clear. This diverges from expected behavior, since this PIP should instead be generated only with a zero value of the "Conceal VMX entries from Intel PT" field (Bit 17) in the Entry Control register (IA32\_VMX\_ENTRY\_CTLSS MSR 0484H).
- Implication:** An Intel PT trace may incorrectly expose entry to non-root operation.
- Workaround:** A Virtual Machine Monitor (VMM) should always set both the "Conceal VMX entries from Intel PT" field in the Entry Control register and the "Conceal VMX in Intel PT" in the Secondary Execution Control register to the same value.
- Status:** [No Fix.](#)

### **SKXD16S. System May Hang Due to Lock Prefixes on Instructions That Access IIO's MMCFG**

- Problem:** If a core uses a lock prefix on an access to an IIO's MMCFG space, then it might lead to a hang if that same IIO has a pending level-triggered interrupt.
- Implication:** The system may hang and cause a log a machine check timeout if issuing lock prefixes on MMCFG accesses.
- Workaround:** Do not use lock prefixes on accesses to MMCFG lines.
- Status:** [No Fix.](#)

### **SKXD17S. With Intel ME Recovery Mode enabled, The Processor May Incorrectly Throttle**

- Problem:** When the Intel® Management Engine (Intel® ME) is in Recovery mode, PCH temperature is not accessible. As a result, the processor will assume PHL (PCH Hot Level) has been exceeded and throttle.
- Implication:** Due to this erratum, the processor may incorrectly throttle.
- Workaround:** It is possible for the BIOS to contain a workaround for this erratum.
- Status:** [No Fix.](#)



### **SKXD18S. VccSA Voltage May Increase After A Demoted Warm Reset Cycle**

**Problem:** VccSA may be increased after a Demoted Warm Reset.  
**Implication:** Due to this erratum, VccSA may be higher than expected prior to a cold reset.  
**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.  
**Status:** No Fix.

### **SKXD19S. An IERR May be Seen When The CPU Attempts Consecutive C6 Entries**

**Problem:** When the CPU attempts consecutive C6 entries, it may result in an Internal Timer Error Machine Check Event (MCACOD = 0x400).  
**Implication:** Due to this erratum, the system may hang with an Internal Timer Error.  
**Workaround:** A BIOS code change may be used as a workaround for this erratum.  
**Status:** No Fix.

### **SKXD20S. ZMM/YMM Registers May Contain Incorrect Values**

**Problem:** Under complex microarchitectural conditions values stored in ZMM and YMM registers may be incorrect.  
**Implication:** Due to this erratum, YMM and ZMM registers may contain an incorrect value. Intel® has not observed this erratum with any commercially available software.  
**Workaround:** It is possible for BIOS to contain a workaround for this erratum.  
**Status:** No Fix.

### **SKXD21S. Intel® PT PSB+ Packets May be Omitted on a C6 Transition**

**Problem:** An Intel® Processor Trace (Intel® PT) Packet Stream Boundary+ (PSB+) set of packets may not be generated as expected when IA32\_RTIT\_STATUS.PacketByteCnt[48:32] (MSR 0x571) reaches the PSB threshold and a logical processor C6 entry occurs within the following one KByte of trace output.  
**Implication:** After a logical processor enters C6, Intel PT output may be missing PSB+ sets of packets.  
**Workaround:** None identified.  
**Status:** No Fix.

### **SKXD22S. Unexpected Uncorrected Machine Check Errors May Be Reported**

**Problem:** In rare micro-architectural conditions, the processor may report unexpected machine check errors. When this erratum occurs, IA32\_MCO\_STATUS (MSR 401H) will have the valid bit set (bit 63), the uncorrected error bit set (bit 61), a model specific error code of 03H (bits [31:16]) and an MCA error code of 05H (bits [15:0]).  
**Implication:** Due to this erratum, software may observe unexpected machine check exceptions.  
**Workaround:** It is possible for BIOS to contain a workaround for this erratum  
**Status:** No Fix.



### **SKXD23S. Intel® PT Trace May Drop Second Byte of CYC Packet**

**Problem:** Due to a rare microarchitectural condition, the second byte of a 2-byte CYC (Cycle Count) packet may be dropped without an OVF (Overflow) packet.

**Implication:** A trace decoder may signal a decode error due to the lost trace byte.

**Workaround:** None identified. A mitigation is available for this erratum. If a decoder encounters a multi-byte CYC packet where the second byte has bit 0 (Ext) set to 1, it should assume that 4095 cycles have passed since the prior CYC packet, and it should ignore the first byte of the CYC and treat the second byte as the start of a new packet.

**Status:** No Fix.

### **SKXD24S. Intel MBM Counters May Report System Memory Bandwidth Incorrectly**

**Problem:** Intel Memory Bandwidth Monitoring (MBM) counters track metrics according to the assigned Resource Monitor ID (RMID) for that logical core. The IA32\_QM\_CTR register (MSR 0xC8E), used to report these metrics, may report incorrect system bandwidth for certain RMID values.

**Implication:** Due to this erratum, system memory bandwidth may not match what is reported.

**Workaround:** It is possible for software to contain code changes to work around this erratum. See the white paper titled *Intel® Resource Director Technology (Intel® RDT) Reference Manual* at <https://software.intel.com/en-us/intel-resource-director-technology-rdt-reference-manual> for more information.

**Status:** No Fix.

### **SKXD25S. Processor May Behave Unpredictably on Complex Sequence of Conditions Which Involve Branches That Cross 64 Byte Boundaries**

**Problem:** Under complex micro-architectural conditions involving branch instructions bytes that span multiple 64 byte boundaries (cross cache line), unpredictable system behavior may occur.

**Implication:** When this erratum occurs, the system may behave unpredictably.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.

### **SKXD26S. Removed**

### **SKXD27S. DMI and PCIe Interfaces May See Elevated Bit Error Rates**

**Problem:** The Direct Media Interface (DMI) or Peripheral Component Interconnect Express\* (PCIe\*) interfaces may be subject to a high bit error rate.

**Implication:** Due to this erratum, an elevated rate of packet CRC errors may be observed on these interfaces which may lead to a machine check error and/or may hang the system.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum.

**Status:** No Fix.



### SKXD28S. Unexpected Page Faults in Guest Virtualization Environment

- Problem:** Under complex microarchitectural conditions, a virtualized guest could observe unpredictable system behavior.
- Implication:** When this erratum occurs, systems operating in a virtualization environment may exhibit unexpected page faults (double faults) leading to guest OS shutdown.
- Workaround:** It is possible for the BIOS to contain a workaround for this erratum.
- Status:** No Fix.

### SKXD29S. STIBP May Not Function as Intended

- Problem:** When the Single Thread Indirect Branch Predictors bit (IA32\_SPEC\_CTL[STIBP] (MSR 48H, bit 1)) is set on one logical processor, then under specific microarchitectural conditions, one logical processor may be able to control the predicted targets of indirect branches on the other logical processors.
- Implication:** Software relying on STIBP to mitigate against the cross-thread speculative branch target injection may allow an attacker running on one logical processor to induce another logical processor on the same core to speculatively execute a disclosure gadget that could allow protected data to be inferred through a side-channel method called Branch Target Injection. This erratum does not affect processors with Intel® Hyper-Threading Technology (Intel® HT Technology) disabled or enabling the cross-thread protections of Indirect Branch Restricted Speculation bit (IA32\_SPEC\_CTL[IBRS] (MSR 48H, bit 0)).
- Workaround:** It is possible for BIOS to contain a workaround for this erratum.
- Status:** No Fix.

### SKXD30S. Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation

- Problem:** This erratum may cause a machine-check error (IA32\_MCi\_STATUS.MCACOD=005H with IA32\_MCi\_STATUS.MSCOD=00FH or IA32\_MCi\_STATUS.MCACOD=0150H with IA32\_MCi\_STATUS.MSCOD=00FH) on the fetch of an instruction. It applies only if (1) instruction bytes are fetched from a linear address translated using a 4-Kbyte page and cached in the processor; (2) the paging structures are later modified so that these bytes are translated using a large page (2-Mbyte, 4-Mbyte or 1-GByte) with a different Physical Address (PA), memory type (PWT, PCD and PAT bits), or User/Supervisor (U/S) bit; and (3) the same instruction is fetched after the paging structure modification but before software invalidates any TLB entries for the linear region.
- Implication:** Due to this erratum an unexpected machine check with error code 0150H with MSCOD 00FH may occur, possibly resulting in a shutdown. This erratum could also lead to unexpected correctable machine check (IA32\_MCi\_STATUS.UC=0) with error code 005H with MSCOD 00FH.
- Workaround:** Software should not write to a paging-structure entry in a way that would change the page size and either the physical address, memory type, or User/Supervisor bit. It can instead use one of the following algorithms: first clear the P flag in the relevant paging-structure entry (e.g., PDE), then invalidate any translations for the affected linear addresses, and then modify the relevant paging-structure entry to set the P flag and establish the new page size. An alternative algorithm: first change the physical page attributes (combination of physical address, memory type and User/Supervisor bit) in all 4K pages in the affected linear addresses, then invalidate any translations for the affected linear addresses, and then modify the relevant paging-structure entry to establish the new page size.
- Status:** No Fix.



### **SKXD31S. Direct Branches With Partial Address Aliasing May Lead to Unpredictable System Behavior**

**Problem:** Under complex micro-architectural conditions involving direct branch instructions with partial address aliasing, unpredictable system behavior may occur. Intel has only seen this under synthetic testing conditions. Intel has not observed this under any commercially available software.

**Implication:** When this erratum occurs, unpredictable system behavior may occur.

**Workaround:** None identified.

**Status:** No Fix.

### **SKXD32S. PCIe Function Level Reset May Generate a NMI# Exception**

**Problem:** Under either of the following conditions, the processor will log a parity error in the OTC\_IRP\_DAT\_PAR register bit (RootBus, Device 5, Function 2, Offset 0X288, bits 11):

1. If CRS Software Visibility Enable bit (RootBus, Device [0-3], Function 0, Offset ACh, bit[4]) is not set,
2. Or if the first transaction sent to the endpoint is not a CfgRd (Configuration Read) transaction following a PCIe Function Level Reset or Secondary Bus Reset event affecting the root port.

**Implication:** When this erratum occurs, the processor will generate an unexpected #NMI exception, which may lead to a system hang or shutdown.

**Workaround:** Software should set the CRS Software Visibility Enable bit to 1. Alternatively, software must ensure that the initial request targeting the end point is a CfgRd request.

**Status:** No Fix.

### **SKXD33S. Performance Monitoring General Counter 2 May Have Invalid Value Written When TSX Is Enabled**

**Problem:** When Transactional Synchronization Extensions (TSX) is enabled, and there are aborts (HLE or RTM) overlapping with access or manipulation of the IA32\_PMC2 general-purpose performance counter (Offset: C3h), it may return invalid value.

**Implication:** The software may read invalid value from IA32\_PMC2.

**Workaround:** None identified

**Status:** No Fix.

### **SKXD34S. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of the Same Priority Completes**

**Problem:** Resuming from the C6 Sleep-State with fixed interrupts of the same priority queued (in the corresponding bits of the IRR and ISR APIC registers), the processor may dispatch the second interrupt (from the IRR bit) before the first interrupt has completed and written to the EOI register, causing the first interrupt to never complete.

**Implication:** Due to this erratum, software may behave unexpectedly when an earlier call to an interrupt handler routine is overridden with another call (to the same interrupt handler) instead of completing its execution.

**Workaround:** None identified.

**Status:** No Fix.



### **SKXD35S. Overflow Flag in IA32\_MC0\_STATUS MSR May Be Incorrectly Set**

**Problem:** Under complex micro-architectural conditions, a single internal parity error seen in IA32\_MC0\_STATUS MSR (401h) with MCACOD [bits 15:0] value of 5h and MSCOD [bits 31:16] value of 7h, may set the overflow flag [bit 62] in the same MSR.

**Implication:** Due to this erratum, the IA32\_MC0\_STATUS overflow flag may be set after a single parity error. Intel has not observed this erratum with any commercially available software.

**Workaround:** None identified.

**Status:** [No Fix.](#)

### **SKXD36S. A Fixed Interrupt May Be Lost When a Core Exits C6**

**Problem:** Under complex micro-architectural conditions, when performance throttling happens during a core C6 exit, a fixed interrupt may be lost.

**Implication:** Due to this erratum, a fixed interrupt may be lost when internal throttling happens during a core C6 exit. Intel has only observed this erratum in synthetic test conditions.

**Workaround:** None identified.

**Status:** [No Fix.](#)





## PCH Errata

---

### **LBG1. The USB3.0 Kernel Debugger And DCI Can Not Be Used At The Same Time**

**Problem:** When both DCI and kernel debugger are enabled, DCI will take precedence over the kernel debugger. This keeps kernel debugger from being used.

**Implication:** The BIOS may not be able to use the USB 3.0 kernel debugger if DCI is enabled.

**Workaround:** Disabling DCI via BIOS to use the USB3.0 kernel debugger.

**Status:** [No Fix.](#)

### **LBG2. Hot-Plugging And Unplugging Of The DBC Cable Can Lead To An Inoperative Debug Port**

**Problem:** If the DBC cable is plugged and unplugged while in the S0 state, in approximately one out of 100 cases it is possible to get into a condition where the debug port gets to a state where it locks up and only a power cycle will restore functionality.

**Implication:** Plugging and unplugging the DBC cable while the system is in S0 state can lead to an inoperative debug port.

**Workaround:** Connect or disconnect the cable when the system is off.

**Status:** [No Fix.](#)



**LBG3. Intel® Ethernet Connection X722 PRTDCB\_RUP2TC and PRTDCB\_TC2PFC Are Not Writable**

**Problem:** The PRTDCB\_RUP2TC (0x1C09A0) and PRTDCB\_TC2PFC (0x1C0980) CSRs cannot be written correctly by software when CSR Protection is enabled.

**Implication:** Programming this CSR is required if the software is configuring DCB on the device.

**Workaround:** For P2TDCB\_RUP2TC: Write to PRTDCB\_RUP2TC as usual, then use a direct admin command with the following values to complete the write transaction. For PRTDCB\_TC2PFC: Write to PRTDCB\_TC2PFC as usual, then use a Direct Admin command with the following values to complete the write transaction.

Field	Byte	Value PRTDCB_RUPTC	value PRTDCB_TC2PFC
Flags	0,1	0x0	0x0
Opcode	2,3	0xFF04	0xFF04
Data Length	4,5	0x0	0x0
Return Value/VFID	6,7	0x0	0x0
Cookie	8-15	Arbitrary value defined by software	Arbitrary value defined by software
Param 0	16-19	0x0	0x0
Param 1	20-23	(0x000AC440 + 0X4 * PRT)	(0x000AC200 + 0X4 * PRT)
Data Address High	24-27	0x0	0x0
Data Address Low	28-31	<CSR write DATA.	<CSR write DATA.

**Status:** No Fix.

**LBG4. Manageability Checksum Filtering Of IPv6 Packets In Intel Ethernet Connection X722**

**Problem:** The IPv6 checksum calculation could be incorrect for received packets that contain either a Routing (type 2) Extension Header or a Destination Options Extension header that includes a home address option.

**Implication:** If the manageability filtering is configured to drop packets with checksum errors, IPv6 manageability packets with the extension headers described above could be incorrectly dropped.

**Workaround:** Do not enable checksum filtering for manageability if the IPv6 Extension Headers described above are used on manageability traffic.

For SMBus: The Enable Xsum Filtering to MNG bit should be 0b in the Update Management Receive filter parameters command, and in the Set Common Filters Receive Control Bytes command if these commands are used.

For NC-SI: Do not use the Enable Checksum Offloading command (Intel OEM command 0x23)

**Status:** No Fix.



**LBG5. INTENA\_MSK Setting Might Clear Interrupt In Intel Ethernet Connection X722**

Problem: A write access to an xxINT\_DYN\_CTLx CSR with INTENA\_MSK bit (bit31) == 1'b0 clears the corresponding interrupt bit in the PBA array.

Implication: There is a possibility of missing an interrupt. However, current Intel software implementation has this bit set to 1 except when enabling or disabling interrupts.

Workaround: INTENA\_MSK should be set in all CSR write accesses other than INTENA bit change.

Status: [No Fix.](#)

**LBG6. Legacy SMBus In The Intel Ethernet Connection X722 Timeout Mechanism Is Not Functional When Not Using ARA Cycle**

Problem: In Legacy SMBus mode, the BMC may get an indication of outstanding events through the ALERT line. The BMC should then do an ARA cycle to get the indicating function. It can instead read the status of all functions. If the BMC fails to do read the status of all ports, and only reads a single function status, then the ALERT line will never de-assert, even if the timeout expires.

Implication: Alert is not de-asserted.

Workaround: Poll all functions' status.

Status: [No Fix.](#)

**LBG7. Intel Ethernet Connection X722 Illegal Byte Error Statistical Counter Inaccuracy**

Problem: Short packets with bad symbols that arrive back-to-back might not be counted by GLPRT\_ILLERRC.

Implication: GLPRT\_ILLERRC is inaccurate.

Workaround: None.

Status: [No Fix.](#)

**LBG8. Intel Ethernet Connection X722 Immediate Interrupts May Delay In System**

Problem: In a case where they are ten or more active queues in the system, and some of the queues are assigned with immediate interrupts, the interrupt delay may exceed the value expected.

Implication: Low performance impact.

Workaround: None.

Status: [No Fix.](#)



### **LBG9. Intel Ethernet Connection X722 L2 Tag Stored In The Wrong RX Descriptor Field**

**Problem:** If two L2 tags (for example VLAN and S-TAG) are programmed to be extracted to the receive descriptor, and the receive descriptor includes only a single L2 tag, the extracted L2 tag is always posted in the L2TAG1 field if L2TSEL is set to 1b, or to L2TAG2 if L2TSEL is set to 0b.

**Implication:** In the following cases, there are no implications:

- If the receive data flow always includes two L2 tags.
- If the receive data might include packets with a single L2 tag, but are always the same tag type (first or second).

If the receive data flow that might include packets with only one L2 tag (which can be either the first or second tag), software cannot identify which of the two enabled L2 tags was extracted to the receive descriptor.

**Workaround:** If the receive data flow includes packets with only one L2 tag, and software is not able to identify if it is the first or second tag, it should not enable more than a single L2 tag to be extracted to the receive descriptor.

**Status:** No Fix.

### **LBG10. Intel Ethernet Connection X722 Jabber Packets Are Not Counted In Jabber Packets Received Field Of NC-SI Get Controller Packet Statistics Command**

**Problem:** Upon issuing NC-SI Get Controller Packet Statistics command, return value counter 12 - Jabber Packets Received should reflect the number of packets received which are larger than the maximum frame size. This counter does not work and the return value is always 0.

**Implication:** Cannot get the number of jabber packets received using the NC-SI Get Controller Packet Statistics command.

**Workaround:** None.

**Status:** No Fix.

### **LBG11. HDA Multiple IDMA's Could Cause Audio Corruption**

**Problem:** When there are multiple IDMA operations in progress, and if the total bandwidth occupied by the multiple IDMA's are fully subscribing the max SDI bandwidth of 424 bytes, there may be data corruption observed in the IDMA where its data is transferred near the end of the Intel® High Definition Audio (Intel® HD Audio) 48 kHz frame boundary.

**Implication:** Corrupted audio input stream

**Workaround:** Write 0x1C to the INPAY register at 0006h.

**Status:** No Fix.



**LBG12. Intel Ethernet Connection X722 TX Performance Degradation For Small Cloud Packets**

Problem: This is for GRE+IPv6+TCP packets without payload. Aggregate TX performance decreases to 33 Gb/s instead of 34 Gb/s.

Implication: It is possible to get less than expected Tx performance. Since this is not a typical packet format it is not expected to be observed in most use cases.

Workaround: None.

Status: [No Fix.](#)

**LBG13. Intel Ethernet Connection X722 TX Descriptor Might Be Read Twice**

Problem: A TX Descriptor might be read more than once in corner case conditions.

Implication: None.

Workaround: None.

Status: [No Fix.](#)

**LBG14. Intel Ethernet Connection X722 ECRC Bits Are Not RO When ECRC Is Disable**

Problem: End-to-end CRC (ECRC) bits in the PCIe\* AER registers are writable even when ECRC is disabled.

Implication: This violated specification. No functional impact is observed.

Workaround: None.

Status: [No Fix.](#)

**LBG15. Intel Ethernet Connection X722 Receive Performance Degradation With Specific Cloud Header**

Problem: A small performance degradation is expected when receiving back-to-back GRE+IPv6+TCP cloud frames with 128-byte header and almost no payload.

Implication: Aggregate performance reduces from 34 Gb/s to 33 Gb/s.

Workaround: None.

Status: [No Fix.](#)

**LBG16. Intel Ethernet Connection X722 Full Switching Table Might Reduce Small Packets Performance**

Problem: If the switching table is relatively full, it might reduce performance with a continuous stream of packets smaller than 160 bytes. A data stream that includes a mix of small and big packets should not experience any degradation.

Implication: Small packets performance impact.

Workaround: Avoid filling up the switching table.

Status: [No Fix.](#)



### **LBG17. Transaction Pending Bit Is Not Functional In The Intel Ethernet Connection X722**

**Problem:** The transaction pending indications and their reflection in the VMPEND registers do not reflect the right value of transaction pending for the PF (Physical Function)/VF (Virtual function) access.

**Implication:** Wrong indication of the transaction pending status.

**Workaround:** Use a timeout mechanism to ensure no transactions are pending.

**Status:** No Fix.

### **LBG18. Intel® QuickAssist Technology (Intel® QAT): Signaled System Error (SSE) Bit Not Cleared**

**Problem:** Intel QuickAssist Technology SSE device status registers PPCISTS and VPCISTS cannot be cleared with a single write.

**Implication:** The interrupt would not be cleared and the interrupt handler would continue to be called to service the interrupt.

**Note:** This issue only impacts Intel C620 series chipset PCH SKUs that support Intel QuickAssist Technology and only when the Operating System (OS) has Advanced Error Reporting (AER) disabled.

**Workaround:** OS or bus drivers have to consecutively write twice to clear SSE.

**Status:** No Fix.

### **LBG19. eSPI's Virtual Wire (VW) Chip Select Counter Is Not Resetting on Slave Link And Channel Recovery (SLCR) De-assertion**

**Problem:** When 2 eSPI slaves are enabled, and there is a fatal link on the VW packet to port 0 caused by the slave on Port 0, after slave recovery the next downstream VS message is only sent to the 2nd slave.

**Implication:** When this situation occurs a fatal error will be logged and escalated to a global reset.

**Workaround:** None.

**Status:** No Fix.

### **LBG20. Intel QuickAssist Technology Endpoint: (P/V) PAERCTLCAP.TFEP Cannot Be Cleared**

**Problem:** The First Error Pointer (TFEP) field in the PF and VF PCI Express\* AER Control and Capability Register (PAERCLT.CAP) Control Status Register (CSR) does not get cleared when the corresponding error is cleared.

**Implication:** The TFEP field we never be cleared after the first error shows up. This behavior does not follow the guidance of the *PCIe Specification*.

**Note:** This issue only impacts SoC that support the Intel QuickAssist Technology.

**Workaround:** A reset event will clear the field to its default value.

**Status:** No Fix.



**LBG21. WAKE# Assertion Does Not Set PCI Express\* Root Port Wake Status (PCIEXP\_WAKE\_STS)**

**Problem:** When the WAKE# pin is asserted to indicate a PCI Express root port wakeup event, the PCI Express Wake Status Bit (B0:D31:F2: Offset 0x0, bit 14) is not set as expected.

**Implication:** The wake even will still occur, but the status will not be logged in the PCI Express Wake Status Bit. BIOS should not rely on PCI Express Wake Status bit to be set after WAKE# assertion.

**Workaround:** None.

**Status:** No Fix.

**LBG22. xHCI Host Controller Reset May Cause A System Hang**

**Problem:** Within 1 ms of setting the Host Controller Reset bit (HCRST bit 1) of the USB Command Register (xHCI BAR + 80h), the xHCI host controller may fail to respond to register accesses.

**Implication:** The system may hang.

**Workaround:** None identified.

**Note:** Software must not make any accesses to the xHCI Host Controller registers for 1 ms after setting the HCRST bit 1 of the USB Command Register (xHCI BAR + 80h) and must add a 120 ms delay in between consecutive xHCI host controller resets.

**Status:** No Fix.

**LBG23. Outstanding eSPI SMI#/SCI# Not Cleared By Warm Reset**

**Problem:** Before a host warm reset entry, if the BMC/EC sends an SMI#/SCI# assertion eSPI Virtual Wire (VW) message without following with a deassertion SMI#/SCI# VW, later when the system exits the warm reset and if the BMC/EC sends another SMI#/SCI# assertion VW, the new assertion SMI#/SCI# will be lost.

**Implication:** In the above corner case, the first SMI#/SCI# VW from the BMC/EC will be lost after warm reset exits.

**Workaround:** 1) The BMC/EC always sends a deassertion SMI#/SCI# VW before it acknowledges the warm reset's HOST\_RST\_WARN handshake, 2) On warm reset exit, SW triggers a dummy SMI#/SCI# assertion/deassertion at the beginning of the boot.

**Status:** No Fix.

**LBG24. PCIe Root Port Interface Not Going Into Loopback Mode**

**Problem:** When the PCIe root ports are set up for loopback mode, and the port is configured as either a x4 or x2, the PCIe controller will not get into the "Loopback Active" state.

**Implication:** The PCIe port will not go into loopback mode, which is a requirement of PCIe Rx compliance testing.

**Workaround:** Reconfigure the port to be 4x1.

**Status:** No Fix.



### **LBG25. Second eSPI Slave VW And Link Error Cause Registers May Not Update**

**Problem:** When eSPI detects either NON-FATAL or Type 2 FATA or Type 1 Link FATA error on the first slave virtual wire channel followed by the same error on the second eSPI slave before the first eSPI slave error status is cleared, the cause of the second slave error will not be updated in the cause register.

**Implication:** Software may not be able to determine that an error on the second eSPI slave also occurred.

**Workaround:** None.

**Status:** No Fix.

### **LBG26. xHCI Controller Does Not flag Parity Error When A Poison Packet Is Received**

**Problem:** The Detected Parity Error bit (D20:F0:06 bit 15) in the xHCI controller will not be set when a parity error is received and the Parity Error Response bit (D20:F0:04, bit 6) is set to 1'b0. The Detected Parity Error bit should be set regardless of the state of the Parity Error Response bit.

**Implication:** A parity error will not be logged properly in the device status register if the Parity Error Response bit is not set.

**Workaround:** Always set the Parity Error Response bit.

**Status:** No Fix.

### **LBG27. OOB PECCI (PECCI Over PECCI Wire) Failing On Boards With High Capacitance Load**

**Problem:** On boards with higher capacitance loads, for example 300+ pF as found in 4S and 8S systems, the delay in the falling edge of the PECCI signal can cause internal logic to incorrectly read the wrong value on the wire, causing a check failure.

**Implication:** The PECCI controller will see a lost arbitration error and halt all further PECCI transactions until reset. No problems are seen with dual socket boards. 4 socket boards could be marginal. 8 socket boards will mostly fail.

**Workaround:** Decrease the value of the pull-down resistor on 4 socket boards to 220 ohms. On 8 socket boards, in-band PECCI (PECCI over DMI) will need to be used.

**Status:** No Fix.

### **LBG28. eSPI Master May Not Service an ALERT From Slave**

**Problem:** If two slave devices are used on the eSPI interface, an ALERT from a slave may be incorrectly ignored, preventing the eSPI controller from issuing a GET\_STATUS command.

**Implication:** The eSPI controller may stall if slave ALERT is ignored.

**Workaround:** A workaround has been implemented to send dummy virtual wire commands (Index 0x43) periodically to avoid a eSPI channel stall. The dummy virtual wire commands are sent when the eSPI interface is enabled, PLTRST# is deasserted, and periodically at least every 10 mS.

**Note:** This issue only impacts eSPI with two slave devices. This workaround can be enabled or disabled with PCH soft strap #121, bit 9. Set to 1 when eSPI used in Dual Slave Attached Flash mode. Set to 0 when eSPI interface used in Master Attached Flash mode or Slave Attached Flash mode with only a single SPI Flash. The change in PCH soft strap





#121 will be reflected in an upcoming release of the SPI programming guide and FITc release.

Status: [No Fix.](#)

**LBG29. The Intel Ethernet Connection X722 Transmitter Does Not Conform To IEEE\* 802.3 Clause 72 KR Electrical Specification For Co-Efficient Update**

Problem: The X722 KR transmitter does conform to IEEE 802.3 KR electrical specification as mentioned in section 72.7.1 with exception on following requirements:

- According to section 72.7.1.11 - For any coefficient update, the magnitudes of the changes in v1, v2, and v3 shall be within 5 mV of each other. In X722 magnitude of changes in v1, v2, and v3 is within 11 mV of each other.
- According to Table 72-8 - Pre cursor equalization ratio (Rpre) for preset coefficient settings [c(1) - disabled, c(0) - maximum, c(-1) - disabled] to be in range 0.95 - 1.05. In X722 Rpre ratio at preset coefficient settings range is 0.95 to 1.08.

Implication: Compliance Specification not met and possible impact on choice of transmitter coefficient setting and Rx equalization during training. Intel has not observed any function or performance impact due to this erratum.

Workaround: None.

Status: [No Fix.](#)

**LBG30. The Intel Ethernet Connection X722 Transmitter Transition Time Does Not Conform To IEEE 802.3 Specification**

Problem: The X722 transmitter does not conform to IEEE 802.3 clause 70 1000BASE-X specification for transition time compliance. The specification states that the transition time should be between 60 ps and 320 ps. The worst case has been found to be faster than specification at approximately 36 ps.

Implication: Compliance testing may report specification violations. Intel has not observed any functional impact due to this erratum.

Workaround: None.

Status: [No Fix.](#)

**LBG31. Sending Data After A RDMA Read In The Intel Ethernet Connection X722 Limited To Less Than 2G**

Problem: While an RDMA Read is outstanding, sending a large amount of data (2 GB or more) can cause the QP (Queue Pair) to hang.

Implication: When this occurs the traffic for that QP may hang.

Workaround: When using an RDMA Read, limit the amount of subsequent data being transmitted to less than 2 GB until the Read completes. This problem is less likely to occur on low latency networks.

Status: [No Fix.](#)



### **LBG32. Intel Ethernet Connection X722 Work Request Size**

**Problem:** Send and Receive Work Requests are limited to 3 fragments. Send Work Requests can support inline data size of 48 bytes.

**Implication:** Three fragments may be limiting for some applications.

**Workaround:** When more than three fragments are required, break the message into multiple work requests. Since vendors have different size of inline data, applications should already be able to adjust for different inline data sizes.

**Status:** No Fix.

### **LBG33. Function-level Reset Fails To Complete In The Intel Ethernet Connection X722**

**Problem:** In rare cases, a function-level reset - Physical Function Reset (PFR), Virtual Function Reset (VFR), or Virtual Machine Reset (VMR) might fail to complete.

**Implication:** PFR: Software times out while waiting for the PFR to complete. The firmware gets stuck and the firmware watchdog timer expires, triggering an Embedded Management Processor Reset (EMPR).

VFR/VMR: Software times out while waiting for the reset to complete.

**Workaround:** PFR: Re-initialize the device after the EMPR.

VFR/VMR: After a timeout waiting for the reset to complete, clear and then set the reset trigger bit (GLGEN\_VFRTRIG.VFSWR for VFR or VSIGEN\_RTRIG.VMSWR for VMR) to retry the reset. Restart the polling for reset completion. After 3 attempts, abort with an error.

**Status:** No Fix.

### **LBG34. Intel Ethernet Connection X722 LAN Function Disabled By BIOS Might Respond To Management Commands**

**Problem:** The X722 Ethernet Controller LAN functions can be disabled by the BIOS. However, when the LAN function is disabled, management commands sent to the disabled function might receive a response instead of being ignored.

**Implication:** The X722 might respond to a management command from the BMC for a port that has been disabled by the BIOS.

**Workaround:** Do not disable a LAN function in BIOS if used for management communication.

**Status:** No Fix.

### **LBG35. eSPI 48 MHz Clock Timings**

**Problem:** The eSPI specification lists the clock period as being 40/60. When running at 48 MHz, the PCH will not meet this spec. The clock period can be 34 high/66 low. Other eSPI clock periods are not affected. This issue is only on 48 MHz eSPI. The PCH will spec the clock times for 48 MHz as being 34/66.

**Implication:** eSPI devices that can not support a clock period of 34/66 may not function.

**Workaround:** Reduce eSPI bus speed to 30 MHz.

**Status:** No Fix.



**LBG36. PCI Express Unexpected Completion Status Bit May Get Set On PCIe Root Port**

**Problem:** A PCI Express device replaying a completion TLP may incorrectly cause an Unexpected Completion error. Note: This has only been observed when a PCIe device causes frequent link corruptions and recovery events to occur.

**Implication:** Bit 16 Unexpected Completion Status (UC) may get set in the Uncorrectable Error Status (UES) Register. (D28:F0/F1/F2/F3/F4/F5/F6/F7; D29:F0/F1/F2/F3/F4/F5/F6/F7; D27:F0/F1/F2/F3 Offset 104h).

**Workaround:** System software may set the Unexpected Completion Mask (UC) (bit 16) in the Uncorrectable Error Mask (UEM) Registers (D28:F0/F1/F2/F3/F4/F5/F6/F7; D29:F0/F1/F2/F3/F4/F5/F6/F7; D27:F0/F1/F2/F3 Offset 108h).

**Status:** [No Fix.](#)

**LBG37. Intel Ethernet Connection X722 MNG Packets Are Dropped While A Function Level Reset To Physical Function 0 (PF 0) Is In Progress**

**Problem:** When a Function Level Reset (FLR) is applied to PF 0, it also resets the LAN-to-BMC pass-through flow.

**Implication:** LAN-to-BMC pass-through traffic stops while FLR is applied to PF 0.

**Workaround:** None.

**Status:** [No Fix.](#)

**LBG38. Get Link Status AQ (Admin Queue) Command Might Return Incorrect Status In the Intel Ethernet Connection X722**

**Problem:** If there is an I<sup>2</sup>C access error when executing the Get Link Status AQ command, the X722 might falsely provide a link down response.

**Implication:** A transient error in accessing the external module via I<sup>2</sup>C causes the software device driver to report a link flap to the system.

**Workaround:** If a Get Link Status AQ response shows a link de-assertion, the Get Link Status command should be repeated.

**Status:** [No Fix.](#)

**LBG39. Intel Ethernet Connection X722 Sticky CFG Space CSRs Are Cleared On Secondary-Bus-Reset When AUX\_PWR Is Disabled**

**Problem:** The hot reset and PERST are merged before it enters PCIe cluster, this violates the PCIe base specification which states that when AUX\_PWR is disabled these bits should not get reset by hot reset (secondary\_bus\_reset).

**Implication:** This violates the sticky bit conditions during AUX\_PWR is disabled when asserting hot reset.

**Workaround:** None.

**Status:** [No Fix.](#)



#### **LBG40. Intel Ethernet Connection X722 Legacy Interrupt Config Space Status Bit Not Implemented**

**Problem:** Legacy interrupt status bit in PCIe config primary status register, offset 0x6, is not implemented and will not be set as described in the *PCIe Base Specification*.

**Implication:** System will not be able to determine which legacy device has issued an interrupt.

**Workaround:** None.

**Status:** No Fix.

#### **LBG41. I<sup>2</sup>C Time Between Start and Stop Not Meeting Spec for Intel Ethernet Connection X722 Buses**

**Problem:** I<sup>2</sup>C Tbuf parameter (time between Start and Stop) is less than the 20 us defined in the *Sff-8431 Specification*.

**Implication:** The *SFF-8431 Specification* requires that the minimum time between STOP and START on an I<sup>2</sup>C bus (Tbuf) should be at least 20 us. The time measured in the 722 Series is less than required by the specification. No functional failure is observed.

**Workaround:** None.

**Status:** No Fix.

#### **LBG42. Glitch On GPIO During GLOBR On Intel Ethernet Connection X722 Pins**

**Problem:** GPIO pins that are defined as SDP outputs (PIN\_FUNC is 000b and PIN\_DIR is 1b in GLGEN\_GPIO\_CTL) can have a high-to-low glitch during GLOBR if OUT\_CTL is 0b. The same applies when the port specified in GLGEN\_GPIO\_CTL.PRT\_NUM is enabled/disabled.

**Implication:** The implication depends on the use of the SDP. For example, an SDP used as a QSFP+ reset signal might cause the module to malfunction due to a short reset assertion.

**Workaround:** One of the following:

- If the SDP is supposed to be high during GLOBR, set OUT\_CTL to 1b.
- For a general-purpose 2-state SDP output (PHY\_PIN\_NAME is 0x3F), set PIN\_FUNC to 001b (LED) and use the LED\_MODE field (0000b or 1111b) to control the output value.

**Status:** No Fix.

#### **LBG43. Intel Ethernet Connection X722 GLQF\_PCNT Counters Do Not Wrap Around**

**Problem:** GLQF\_PCNT counters do not wrap around.

**Implication:** The implication depends on the use of the SDP. For example, an SDP used as a QSFP+ reset signal might cause the module to malfunction due to a short reset assertion. The GLQF\_PCNT counters saturate at the maximum value (0xFFFFFFFF). This means that SW has to periodically clear the counters. However, the counters are cleared by a write, so there can be packet counts missed between the last read of the counter and the write that clears it.

**Workaround:** Software should periodically clear these counters by writing any value.

**Status:** No Fix.



**LBG44. PCH Does Not Meet Charged Device Model (CDM) ESD Specification**

**Problem:** PCH fails when tested to an ESD level of 250V for Charged Device Model (CDM) testing on the high speed differential signals (USB3, SATA, SSATA, DMI, PCIe). The maximum ESD passing level is 200V.

**Implication:** The PCH does not meet JEDEC CDM specification target level of 250V.

**Workaround:** None.

**Status:** No Fix.

**LBG45. MSR 0xC80 IA\_32\_DEBUG\_INTERFACE\_MSR Enable May Be SET After Clear CMOS**

**Problem:** After clear CMOS, the C620 Chipset may send an incorrect message to the CPU, which could cause the MSR 0xC80 IA\_32\_DEBUG\_INTERFACE\_MSR enable bit to be set.

**Implication:** After clear CMOS, debug may be inadvertently enabled. Intel® Trusted Execution Technology (Intel® TXT) launch fail.

**Workaround:** Workaround in Intel® Server Platform Services release 4.0.3.202.0 or later.

**Status:** No Fix.

**LBG46. Intel Ethernet Connection X722 RDMA SGE Count Limitations**

**Problem:** The maximum number of Scatter Gather Elements (SGEs) for Send and Receive Work Queue Elements (WQEs) is 3. This also limits amount of inline data to 48 bytes.

**Implication:** The amount of inline data and the number of SGEs supported for Send and Receive WQEs varies across devices and vendors. It is normal for applications to query the device for its characteristics and the industry standard APIs have queries for this purpose. Applications are already expected to be written flexibly to use a variable number of SGEs.

**Workaround:** If an application needs more than 3 SGEs, it will need to break the message into multiple messages.

**Status:** No Fix.

**LBG47. Program Suspend Instruction and Program Resume Instruction Fields Are Not Used By SPI Controller**

**Problem:** The 13<sup>th</sup> double word (DW) of the Serial Flash Descriptor Parameter (SFDP) contains the opcodes used for Suspend Instruction [31:24] (write or erase type operation), Resume instruction [23:16] (write or erase type operation), Program Suspend Instruction [15:8] (program operation), and Program Resume Instruction [7:0] (program operation). The Intel C620 Series Chipset PCH SPI controller only reads bits 31:16, and ignores bits [15:0].

**Implication:** If bits [31:16] are different than bits 15:0, then the suspend/resume feature cannot be used.

**Note:** A survey of the major flash vendors has shown they program bits [31:16] the same as bits [15:0].

**Workaround:** Disable the suspend/resume feature using the appropriate soft strap.

**Status:** No Fix.



#### **LBG48. Hang On CF9 06 Reset In POST**

**Problem:** It is possible to power gate DFX logic via the Host SW PG Control Register 1 (HSWPGCR1) in the PMC memory mapped registers.

**Implication:** If power gating is enabled, hangs during warm resets have been seen.

**Workaround:** Keep bit 0 of HSWPGCR (the DFX SW PG Req control bit) cleared to 0.

**Status:** No Fix.

#### **LBG49. Intel Ethernet Connection X722 Device Unable To Recover When All Ports Are Disabled**

**Problem:** When all four ports of the LAN device are disabled from the BIOS, the LAN device goes into a hung state and cannot be recovered. Ports can not be re-enabled.

**Implication:** LAN device does not recover, possible boot impairments.

**Workaround:** At least have one port enabled at all times if BIOS is being used to enabled/disable ports. If all 4 ports need to be disabled, use the descriptor method.

**Status:** No Fix.

#### **LBG50. BMC Shared NIC Slow Response In Heavy Network Traffic With Intel Ethernet Connection X722**

**Problem:** In PXE mode, when the Intel Ethernet Connection X722 LAN's Rx pipe is set to no-drop mode, packets are held in pipe until processed by the host. The BMC traffic in the MAC receive data buffer shares the same pipe with host traffic. While the MAC is receiving heavy traffic such as ARP packet storm, since the UEFI driver processes packets very slowly, the BMC traffic is delayed/dropped in the Rx pipe.

**Implication:** Because the BMC traffic is delayed/dropped in the Rx pipe, the ICMP packet would get Destination Unreachable or Time Exceeded, and BMC would fail in DHCP.

**Workaround:** There is a partial fix in X722 NVM version 3.49 or later that applies a dynamic drop mode mechanism. It detects congestion in Rx buffer in such case then it switches Rx-path into drop mode periodically. Dynamic drop mode improves the BMC shared NIC slow response so that the issue cannot be reproduced in managed switch or receiving moderate APP traffic. In a critical test case like ARP packet storm, the BMC can still get DHCP failures or response ICMP packet timeout.

**Status:** No Fix.

#### **LBG51. Memory/IO Reads Targeting BMC On eSPI May Be Completed Incorrectly**

**Problem:** If the following sequence of steps occurs:

1. Any TPM read.
2. PCI configuration write to D31:F5 (SPI controller) and there is no PCI configuration read or memory read to D31:F5 before step 1.
3. I/O read/write cycle to BMC on eSPI.

Then memory I/O cycles targeting the BMC may be completed incorrectly. Once in this state, the error condition can only be cleared by a global reset. Host resets will not clear out the SPI/eSPI logic.

If LPC is used instead of eSPI, there is no problem.



Implication: Incorrect data will be returned from the BMC with unpredictable results when eSPI interface is used.

Workaround: After step 1 or step 2 do a configuration or memory read to D31:F5 (SPI controller).

Status: [No Fix](#).

### **LBG52. System Hangs After BIOS/Intel® Server Platform Services Firmware Flash Update Completes**

Problem: If a global reset occurs during a SPI erase cycle, the SPI flash component will not respond back to a SPI SFDP read properly because it is busy doing a SPI erase.

Implication: The read of the descriptor will fail, so no soft straps or Intel ME FW will be loaded, resulting in various failure signatures based upon the system design. If power gating is enabled, hangs during warm resets have been seen.

In a system with no X722 LAN, it is possible to not have the 25 MHz crystal populated and this will result in a hang condition.

If the LAN crystal is populated:

1. Soft straps will fail to load that could cause invalid system configurations and global resets.
2. Intel ME may not load PMC patch/FW resulting in an invalid system configuration.

Workaround: If the global reset is the result of FW initiating it, FW will need to wait until the SPI device has finished the erase cycle and is ready to take further actions.

If the global reset is the result of a HW initiated reset, there is no workaround.

Status: [No Fix](#).

### **LBG53. Legacy GbE Can Cause The PCH To Hang During Boot**

Problem: If the legacy GbE controller is enabled via fuse/soft strap, then the PMC can enter a hang condition causing the PCH to hang during boot. In order to avoid this situation, the default soft straps for enabling the legacy GbE was redefined to be reserved = 0, keeping the legacy GbE disabled and a new soft strap was defined for the PMC to read in order to enable the legacy GbE later on.

Implication: A new softstrap location was defined for enabling/disabling GbE. This was done and implemented in FITC and the SPI Programming Guide at the time of distribution of C620 A0 parts.

BIOS enable/disable of the legacy GbE works in conjunction with the original soft strap. Since that must be kept as 0, disabling legacy GbE, the BIOS enabling and disabling of the legacy GbE will not function anymore.

Workaround: If the legacy GbE needs to be disabled, it must be done via Soft Strap.

Status: [No Fix](#).



**LBG54. Do Not Access Parent Memory Region While Memory Windows Are Active With Intel Ethernet Connection X722 RDMA Applications**

**Problem:** Using memory windows while at the same time referencing the parent memory region from RNIC hardware is not common but it is legal. RDMA applications need to avoid this scenario.

**Implication:** When using memory windows, referencing the parent memory region via SQ, RQ or remote operations may cause system hangs.

**Workaround:** The parent memory region should not be accessed from a remote application. To avoid remote reference to a local parent memory region, an RDMA application should not advertise the parent memory region's STag to the remote application.

**Status:** No Fix.

**LBG55. Intel Ethernet Connection X722 Activity LED May Blink Regardless If Link Is Up Or Down For A Port**

**Problem:** X722 Activity LEDs toggle as a result of incoming or out-going packets, regardless of the link state of a port. The packets can be received either from the host or the Base Management controller (BMC).

**Implication:** When the LAN cable is disconnected, packets from the BMC will toggle the Activity LED. It is expected of the BMC to send traffic to a LAN port only if the link is up.

**Workaround:** Have the BMC check link status of port prior to sending traffic. For link indication, X722 supports both Link Status Change AEN and Get Link Status command per NC-SI specification.

**Status:** No Fix.

**LBG56. EOI Broadcast From The CPU May Cause Errors To Be Reported On The x16 Uplink**

**Problem:** When the CPU broadcasts the EOI message down the PCIe x16 uplink to the PCH, and the 10/1 GbE LAN and all three Intel® QuickAssist Technology (Intel® QAT) endpoints aren't enabled, then the PSF error handlers will report an error. There no functional errors that occur because of this.

**Implication:** The PCIe x16 uplink Advanced Error Reporting register will set bit 20 (Unsupported Request) in the Uncorrectable Error Status register (offset 104h) and bit 3 (Unsupported Request Detected) in the Device Status register (offset 4Ah).

**Workaround:** Have the BIOS block the forwarding of the EOI on the CPU root ports attached to the PCH PCIe uplink (MISCCTRLSTS\_0 register, offset 0x188h, bit 26 set to 1).

**Status:** No Fix.

**LBG57. MCTP Broadcast Messages To IE, Intel Ethernet Connection X722 And Intel QuickAssist Technology May Cause Errors To Be Reported**

**Problem:** When MCTP broadcast messages are issued by Intel® Server Platform Services and are sent to the following functions (IE, X722 Ethernet Connection, QuickAssist Endpoint 0, QuickAssist Endpoint1, QuickAssist Endpoint 2) and any one of those functions are disabled or not visible, then the PSF error handlers will report and error. There are no functional errors that occur because of this.

**Implication:** This varies due to how error handling is enabled. Possible implications are the following

- 1) For cases where the broadcast MCTP message does not make it to the X722 LAN or any one of the three Intel QuickAssist endpoints, then the PCIe uplink Advanced Error Reporting register will set bit 20 (unsupported request) in the uncorrectable error





status register (offset 104h) and bit 3 (Unsupported Request Detected) in the Device Status register (offset 4Ah). This can occur in both the x16 uplink and optional x8 uplink.

2) In the case of IE being disabled, if PCH Server Error Mode is enabled (Bit 8 of the General Interrupt Control Register (offset 31FC)) then an ERR\_NONFATAL message is sent up DMI to the CPU. .

Workaround: 1) Set the mask bit for these errors so no error reporting is done. 2) Contact your Intel Representative for the version of Intel Server Platform Services that corrects this problem.

For IE, disable PSF1 multicast controller from forwarding multicast messages to PSF6/IE.

Status: [No Fix.](#)

### **LBG58. In EFI, Multiple Link Status Change Events In the X722 Ethernet Connection Might Cause An AEN Storm To the BMC**

Problem: While in EFI, repeatedly plugging and unplugging an Ethernet cable may cause multiple link status change events to be sent to the UEFI driver and the BMC. Note that the same AQ command is sent to both the UEFI driver and the BMC, and the UEFI driver has limited space in a buffer designated for AQ commands that is not emptied on a scheduled basis. When link status changes 32 times or more, the X722 firmware attempts to send a link status change to both the UEFI driver and the BMC. However, the driver buffer is already full, which results in firmware continually repeating the notification to both BMC and UEFI driver. This causes an AEN storm to the BMC.

Implication: The BMC receives multiple AEN requests that might impact the performance or in some cases might reset the BMC.

Workaround: While in UEFI, limit the number of times the Ethernet cables is plugged in and unplugged. If the driver buffer is full, unload and then reload the UEFI driver to clear the buffer.

Status: [No Fix.](#)

### **LBG59. 3.3V Deep Sleep Rail Bleed Voltage Onto 1.8V Rail**

Problem: When the 3.3V Deep Sleep Power rail powers, a small percentage of material observes that the on chip LDOs that provide 1.8V power internally will turn on too soon.

Implication: The 3.3V deep sleep rail may start pulling the 1.8V rail as high as 2.6V, which could cause the system to not boot if the 1.8V Vreg sense this as an overvoltage condition. When the BIOS initializes the chip, the 3.3V deep sleep rail will stop pulling the 1.8V rail.

Workaround: To keep the 1.8V rail from being pulled higher, the 1.8V Vreg design must be able to sink 5 mA of current. This can be done via the Vreg itself or by having a load resistor on the 1.8V rail to ground.

Status: [No Fix.](#)



**LBG60. 10GBASE-KR Link Establishment May be Impacted if Link Partner Issues PRESET Request**

**Problem:** At the beginning of the IEEE 802.3 Clause 72 10GBASE-KR training sequence, the link partner may issue TxFFE coefficient requests that are acted upon correctly. Approximately 45 ms after link training has commenced, the X722 transmitted signal reverts back to its INITIALIZE setting when no INITIALIZE request was issued by the link partner. Further link training requests are responded to correctly. The link partner may also complete its training and achieve the Receiver Ready state prior to the reversion to INITIALIZE.

**Implication:** This behavior is more prominent to a link partner who issues a PRESET request at the beginning of 10GBASE-KR link training. As a result the 10GBASE-KR link establishment may take a long time to link, may fail, or a link may be established with non-optimal coefficients.

**Workaround:** The workaround is to configure the 10GBASE-KR INITIALIZE coefficients to the same value as the PRESET coefficients. This workaround should only be used with link partners which issue a PRESET request at the beginning of 10GBASE-KR link training.

Contact your Intel representative to obtain workaround NVM images.

**Status:** No Fix.

**LBG61. xHCI Short Packet Event Using Non-Event Data TRB**

**Problem:** The xHCI may generate an unexpected short packet event for the last transfer's Transfer Request Block (TRB) when using Non-Event Data TRB with multiple TRBs.

**Implication:** Transfer may fail due to the packet size error.

**Note:** This issue has only been observed in a synthetic environment. No known implication has been identified with commercial software.

**Workaround:** None identified. Intel recommends software to use Data Event TRBs for short packet completion.

**Status:** No Fix.



## **Specification Changes**

---

There are no specification clarifications at this time.

**§ §**



## **Specification Clarifications**

---

There are no specification clarifications at this time.

**§ §**



## **Documentation Changes**

---

There are no documentation changes at this time.

**§ §**