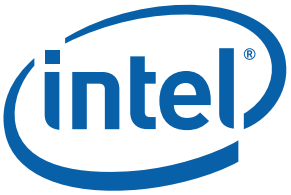

Intel® Xeon® E7-8800/4800 v3 Processor Product Family

Specification Update

February 2020

Reference Number: 332317-022US

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

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

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.

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.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see http://www.intel.com/content/www/us/en/architecture-andtechnology/turbo-boost/turbo-boost-technology.htmlhttp://www.intel.com/technology/turboboost

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 www.intel.com/design/literature.htm.

Intel, the Intel logo, Intel Core, Pentium, Pentium 4, Xeon, Intel QuickPath Interconnect (Intel QPI), Intel Hyper-Threading Technology (Intel HT Technology), Intel Turbo Boost Technology, Intel Trusted Execution Technology (Intel TXT), Intel Transactional Synchronization Extensions (Intel TSX), Intel QuickData Technology, Intel 386, Intel 486, Intel C102, Intel 104, Intel 112, Intel 114, 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.

Contents

Revision History

Date	Revision	Description
February 2020	022	• Added erratum HSX82
April 2019	021	• Added erratum HSX81
March 2019	020	• Added erratum HSX80
June 2018	019	• Added erratum HSX79
April 2018	018	• Updated erratum HSX75
January 2018	017	• Added erratum HSX78
March 2017	016	• Added erratum HSX77
December 2016	015	• Added errata HSX74-HSX76
November 2016	014	• Added errata HSX72-HSX73 • Updated erratum HSX51
October 2016	013	• Added errata HSX70-HSX71
September 2016	012	• Added erratum HSX69
August 2016	011	• Added erratum HSX68
July 2016	010	• Added errata HSX66 - HSX67
May 2016	009	• Added errata HSX64 - HSX65
April 2016	008	• Added errata HSX61 - HSX63
March 2016	007	• Updated erratum HSX49 • Added errata HSX59 - HSX60
February 2016	006	• Revised erratum HSX56 title • Added erratum HSX58
December 2015	005	• Added errata HSX56 - HSX57
November 2015	004	• Added erratum HSX55
September 2015	003	• Added erratum HSX54
August 2015	002	• Added errata HSX48 - HSX53. • Updated Table 1.
May 2015	001	• Initial Release

§

Preface

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

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

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

Affected Documents

Document Title	Document Number/Location
Intel® Xeon® Processor E7-4800/8800 v3 Product Families Datasheet - Volume 1: Electrical, Mechanical and Thermal	332314
Intel® Xeon® Processor E7-8800/4800 v3 Product Families Datasheet Volume 2: Registers	332315

Nomenclature

Errata are design defects or errors. These may cause the Intel® Xeon® E7-8800/4800 v3 Processor Product Family 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.

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, and so forth, as described in the processor identification information table. Read all notes associated with each S-Spec number.

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

§

Identification Information

Component Identification via Programming Interface

The Intel® Xeon® Processor E7 v3 Product Family stepping can be identified by the following register contents:

Reserved	Extended Family	Extended Model	Reserved	Processor Type	Family Code	Model Number	Stepping ID
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0011b		00b	0110b	1111b	0100b

Notes:
1. The Extended Family, Bits [27:20] are used in conjunction with the Family Code, specified in Bits [11:8], to indicate whether the processor belongs to the Intel® 386, Intel® 486, Pentium®, Pentium 4, or Intel® Core™ Processor Family.
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 Extended Data Register (EDX) register after RESET, Bits [11:8] of the Extended Accumulator Register (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 1 for the processor stepping ID number in the CPUID information.
6. Refer to the *Haswell-EN/EP/EP 4S/EX Processor BIOS Writer's Guide (BWG), Combined Volumes: 1-3* for additional information.Refer to the Intel® 64 and IA-32 Architectures Software Developer's Manual documentation for additional 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.

The Cache and the Translation Lookaside Buffer (TLB) descriptor parameters are provided in the EAX, the Extended Base Register (EBX), the Extended Count Register (ECX) and the EDX registers after the CPUID instruction is executed with a 2 in the EAX register.

Component Marking Information

For Intel® Xeon® Processor E7 v3 Product Family SKUs, see https://ark.intel.com/content/www/us/en/ark/products/series/78585/intel-xeon-processor-e7-v3-family.html

Figure 1. Intel® Xeon® Processor E7 v3 Product Families Top-Side Markings (example)

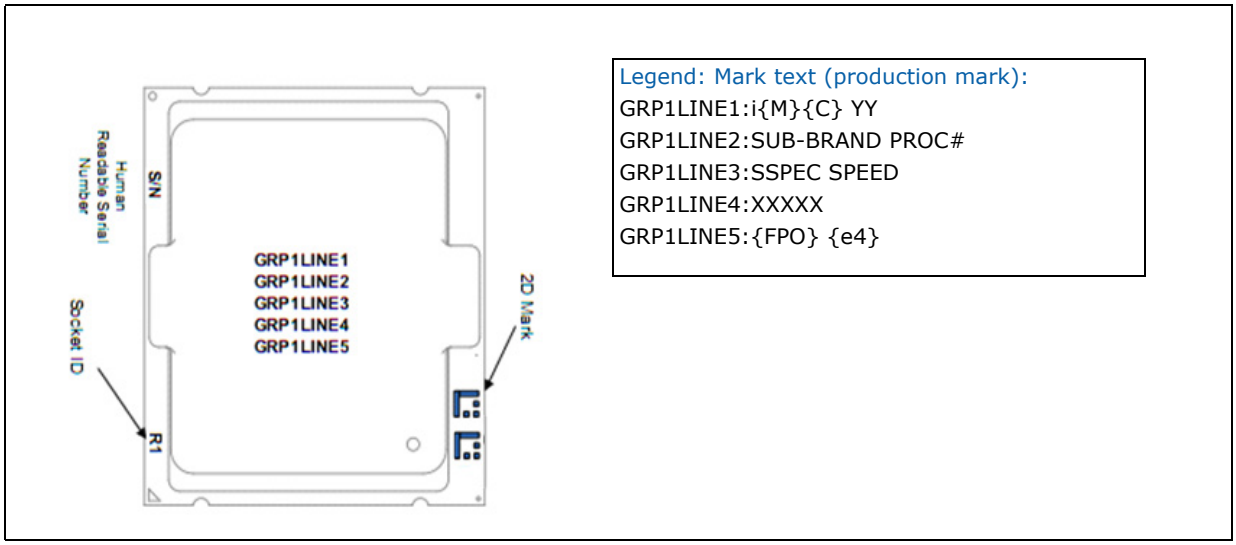

Legend: Mark text (production mark):
GRP1LINE1:i{M}{C} YY
GRP1LINE2:SUB-BRAND PROC#
GRP1LINE3:SSPEC SPEED
GRP1LINE4:XXXXX
GRP1LINE5:{FPO} {e4}

The product family stepping can be identified by the following component markings. Refer to the Dear Customer Letter (DCL) for additional details and conditions of test support.

§

Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the "Intel® Xeon® E7-8800/4800 v3 Processor Product Family". Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables uses the following notations:

Codes Used in Summary Tables

Stepping

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

Page

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

Status

Doc:	Document change or update will be implemented.
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.

Table 1. Errata (Sheet 1 of 3)

Number	Stepping	Status	Errata
	E0		
HSX1	X	No Fix	Intel® QuickPath Interconnect (Intel® QPI) Layer May Report Spurious Correctable Errors
HSX2	X	No Fix	Platform Environment Control Interface (PECI) DDR DIMM Digital Thermal Reading Returns Incorrect Value
HSX3	X	No Fix	IIO Control and Status Register (CSR) Lnkcon2 Field Selectable_De_Emphasis Cannot Be Set For DMI2 Mode
HSX4	X	No Fix	PCIe* Receiver May Not Meet the Specification for AC Common Mode Voltage And Jitter
HSX5	X	No Fix	Receiver Termination Impedance On PCIe 3.0 Does Not Comply With The Specification
HSX6	X	No Fix	A Memory Channel With More Than 4 Ranks May Lead to a System Hang
HSX7	X	No Fix	Writing R3QPI Performance Monitor Registers May Fail
HSX8	X	No Fix	Intel® QPI Link Re-training After a Warm Reset or L1 Exit May be Unsuccessful
HSX9	X	No Fix	VCCIN Voltage Regulator (VR) Phase Shedding is Disabled
HSX10	X	No Fix	PECI Commands During Reset May Result in Persistent Timeout Response
HSX11	X	No Fix	System May Hang When Using the Transaction Layer Packet Processing Hint (TPH) Prefetch Hint
HSX12	X	No Fix	TS1s Do Not Convey The Correct Transmitter Equalization Values During Recovery.RcvrLock
HSX13	X	No Fix	MSR_TEMPERATURE_TARGET Model Specific Register (MSR) May Read as '0'
HSX14	X	No Fix	PECI RdIAMSR() Command May Fail After Core C6 State is Entered
HSX15	X	No Fix	Closed Loop Thermal Throttling (CLTT) May Cause the BIOS To Hang On a Subsequent Warm Reset
HSX16	X	No Fix	PCIe Extended Tag Field May be Improperly Set
HSX17	X	No Fix	A MOV to CR3 When Extended Page Tables (EPT) is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
HSX18	X	No Fix	Memory Controller tsod_present Settings Being Improperly Cleared
HSX19	X	No Fix	DDR4 Power Down Timing Violation
HSX20	X	No Fix	Correctable Memory Error Correcting Code (ECC) Errors May Occur at Boot
HSX21	X	No Fix	Backup Tracker (BT) Timeouts May Cause Spurious Machine Checks
HSX22	X	No Fix	PCIe Type 1 Vendor Defined Message (VDM) May be Silently Dropped
HSX23	X	No Fix	CONFIG_TDP_NOMINAL CSR Implemented at Incorrect Offset
HSX24	X	No Fix	A Machine-Check Exception Due to Instruction Fetch May Be Delivered Before an Instruction Breakpoint
HSX25	X	No Fix	Power Consumed During Package C6 May Exceed Specification
HSX26	X	No Fix	Platform Performance Degradation When C1E is Enabled
HSX27	X	No Fix	PCIe Correctable Error Status Register May Not Log Receiver Error at 8.0 GT/s
HSX28	X	No Fix	PCIe Hot-Plug Slot Status Register May not Indicate Command Completed
HSX29	X	No Fix	Local PCIe Peer-to-Peer Traffic on x4 Ports May Cause a System Hang
HSX30	X	No Fix	ILLC Error Conditions May be Dropped or Incorrectly Signaled
HSX31	X	No Fix	A DDR4 Command/Address (C/A) Parity Error in Lockstep Mode May Result in a Spurious Uncorrectable Error
HSX32	X	No Fix	Some Integrated Memory Controller (IMC) and Intel® QPI Functions Have Incorrect PCI Capability Pointer Register (CAPPTR) Values
HSX33	X	No Fix	PCIe TLP Translation Request Errors Are Not Properly Logged For Invalid Memory Writes
HSX34	X	No Fix	Consecutive PECI RdIAMSR Commands When Core C6 is Enabled May Cause a System Hang

Table 1. Errata (Sheet 2 of 3)

Number	Stepping	Status	Errata
	E0		
HSX35	X	No Fix	Enabling Targeted Row Refresh (TRR) With DDR4 LRDIMMs May Lead to Unpredictable System Behavior
HSX36	X	No Fix	Command Address (C/A) Parity Error Injection May Cause the System to Hang
HSX37	X	No Fix	The System May Shut Down Unexpectedly During a Warm Reset.
HSX38	X	No Fix	Patrol Scrubbing of Mirrored Memory May Log Spurious Memory Errors
HSX39	X	No Fix	MSR_TURBO_ACTIVATION_RATIO MSR Cannot be Locked
HSX40	X	No Fix	The System May Shut Down Unexpectedly During a Warm Reset
HSX41	X	No Fix	Invalid Intel® QuickData Technology XOR Descriptor Source Addressing May Lead to Unpredictable System Behavior
HSX42	X	No Fix	Warm Reset May Cause PCIe Hot-Plug Sequencing Failure
HSX43	X	No Fix	PCIe Uncorrectable Response (UR) And Completer Abort (CA) Responses May be Sent Before Link Enters Live Error Recovery (LER) State
HSX44	X	No Fix	Surprise Down Error Status is Not Set Correctly on DMI Port
HSX45	X	No Fix	Intel SMI2 in Half Width Mode With Dual Device Data Correction (DDDC) Enabled Will Not Report RdECC Errors
HSX46	X	No Fix	PCIe SLTCON CSRs electromechanical_interlock_control Field Read as 1
HSX47	X	No Fix	Intel DDR3 SMI2 Command Address Parity (CAP) Errors Are Ignored Leading to Unpredictable System Behavior
HSX48	X	No Fix	PECI RdPkgConfig Command DRAM Services May Behave Incorrectly
HSX49	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events May Undercount
HSX50	X	No Fix	Performance Monitoring OFFCORE_RESPOSE_{1,2} Events May Miscount L3_MISS_REMOTE_HOP
HSX51	X	No Fix	Some DRAM and L3 Cache Performance Monitoring Events May Count Incorrectly
HSX52	X	No Fix	Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations
HSX53	X	No Fix	An IRET Instruction That Results in a Task Switch Does Not Serialize The Processor
HSX54	X	No Fix	A P-State or C-State Transition May Lead to a System Hang
HSX55	X	No Fix	PECI Frequency Limited to 1 MHz
HSX56	X	No Fix	VGATHERQPS That Loads an Element From The Advance Programmable Interrupt Controller (APIC)-Access Page May Load Other Elements From Incorrect Addresses
HSX57	X	No Fix	A Spurious Patrol Scrub Error May be Logged
HSX58	X	No Fix	DRAM Device Failure With Error Flow Registers Enabled May Result in a Machine Check
HSX59	X	No Fix	Monitor Trap Flag (MTF) VM Exit on XBEGIN Instruction May Save State Incorrectly
HSX60	X	No Fix	PEBS Record May Be Generated After Being Disabled
HSX61	X	No Fix	PCIe Ports Do Not Support Data Link Layer (DLL) Link Active Reporting
HSX62	X	No Fix	PCIe Link Speed Negotiation May Fail After Link is Re-enabled
HSX63	X	No Fix	PROCHOT# Assertion During Warm Reset May Cause Persistent Performance Reduction
HSX64	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May be Lost
HSX65	X	No Fix	A Corrected Internal Parity Error May Result in a System Hang
HSX66	X	No Fix	IA32_MC0_STATUS May be Incorrect After a Machine Check Overflow
HSX67	X	No Fix	IA32_MC1_STATUS.MISCV May be Incorrect on a Machine Check Overflow
HSX68	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions

Table 1. Errata (Sheet 3 of 3)

Number	Stepping	Status	Errata
	E0		
HSX69	X	No Fix	Resume Flag (RF) May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or Branch Trace Store (BTS)
HSX70	X	No Fix	An APIC Timer Interrupt During Core C6 Entry May be Lost
HSX71	X	No Fix	MOVNTDQA From Write Combining (WC) Memory May Pass Earlier Locked Instructions
HSX72	X	No Fix	An x87 Store Instruction Which Pends #PE While EPT is Enabled May Lead to an Unexpected Machine Check and/or Incorrect x87 State Information
HSX73	X	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
HSX74	X	No Fix	Certain PerfMon Events May be Counted Incorrectly When The Processor is Not in C0 State
HSX75	X	No Fix	Writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP May #GP When Intel® Transactional Synchronization Extensions (Intel® TSX) is Not Supported
HSX76	X	No Fix	JTAG Boundary Scan For Intel® QPI and PCIe Lanes May Report Incorrect Stuck at 1 Errors
HSX77	X	No Fix	APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode
HSX78	X	No Fix	Debug Exceptions May Be Lost in The Case Of Machine Check Exception
HSX79	X	No Fix	In eMCA2 Mode, When the Retirement Watchdog Timeout Occurs CATERR# May be Asserted
HSX80	X	No Fix	Using the Intel® TSX Instructions May Lead to Unpredictable System Behavior
HSX81	X	No Fix	Spurious Corrected Errors May be Reported
HSX82	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation

Table 2. Specification Clarifications

No.	Specification Clarifications
1	None

Table 3. Specification Changes

No.	Specification Changes
1	None

Table 4. Documentation Changes

No.	Documentation Changes
1	None

§

Errata

HSX1 Intel® QuickPath Interconnect (Intel® QPI) Layer May Report Spurious Correctable Errors

Problem: Intel® QPI may report an inband reset with no width change (error 0x22) correctable error upon exit from the L1 power state as logged in its IA32_MC{5, 20, 21}_STATUS Model Specific Register (MSRs) (415H,451H,455H).

Implication: An unexpected inband reset with no width change error may be logged.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX2 Platform Environment Control Interface (PECI) DDR DIMM Digital Thermal Reading Returns Incorrect Value

Problem: When using the PECI RdPkgConfig() command to read Package Config Space (PCS) Service 14 "DDR DIMM Digital Thermal Reading", the value returned is incorrect.

Implication: Platform thermal management may not behave as expected.

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

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

HSX3 IIO Control and Status Register (CSR) Lnkcon2 Field Selectable_De_Emphasis Cannot Be Set For DMI2 Mode

Problem: The CSR lnkcon2 (Bus 0; Device 0; Function 0, Offset 0x1C0) field selectable_de_emphasis (bit 6) cannot be set for a link when the DMI port is operating at 5 GT/s.The documentation has the attribute of Read, Write Once (RW-O), but the processor incorrectly operates as read-only. This erratum does not occur when link is operating as a PCIe* port.

Implication: When the link is in DMI2 mode, the de-emphasis cannot be changed for an upstream component.

Workaround: None identified.

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

HSX4 PCIe* Receiver May Not Meet the Specification for AC Common Mode Voltage And Jitter

Problem: Due to this erratum, PCIe receivers may not meet the specification for AC common mode voltage (300 mV) and jitter (78.1 ps) at high temperatures when operating at 5 GT/s.

Implication: Specifications for PCIe receiver AC common mode voltage and jitter may not be met. Intel has not observed this erratum on any commercially available system with any commercially available PCIe devices.

Workaround: None identified.

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

HSX5 Receiver Termination Impedance On PCIe 3.0 Does Not Comply With The Specification

Problem: The PCIe Base Specification revision 3.0 defines ZRX-HIGH-IMP-DC-NEG and ZRX-HIGH-IMP-DC-POS for termination impedance of the receiver. The specified impedance for a negative voltage (-150 mV to 0 V) is expected to be greater than 1 Kohm. Sampled measurements of this impedance as low as 400 ohms have been seen. The

specified impedance for a positive voltage (> 200 mV) is greater than 20 Kohms. Sampled measurements of this impedance as low as 14.6 Kohms have been seen.

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

Workaround: None identified.

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

HSX6 A Memory Channel With More Than 4 Ranks May Lead to a System Hang

Problem: A memory controller channel with more than 4 ranks and with a Targeted Row Refresh (TRR) enabled may fail leading to a system hang. This erratum only impacts memory channels with three dual-rank DDR4 RDIMMs.

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

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX7 Writing R3QPI Performance Monitor Registers May Fail

Problem: Due to this erratum, attempting to write R3QPI performance monitor registers (Bus 0; Device 11; Functions 1,2,5,6; Offset 0xA0-0xF7) may be unsuccessful.

Implication: A failed write to one or more R3QPI performance monitor registers is likely to yield incorrect performance events counts.

Workaround: Consecutively write the identified registers twice with the same value before performance monitoring is globally enabled.

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

HSX8 Intel® QPI Link Re-training After a Warm Reset or L1 Exit May be Unsuccessful

Problem: After a warm reset or an L1 exit, the Intel® QPI links may not train successfully.

Implication: A failed Intel® QPI link can lead to reduced system performance or an inoperable system.

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

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

HSX9 VCCIN Voltage Regulator (VR) Phase Shedding is Disabled

Problem: Due to this erratum, the processor does not direct the VCCIN VR to shed phases during low power states.

Implication: Platform power consumption may exceed expected levels during deep package C-states.

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

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

HSX10 PECI Commands During Reset May Result in Persistent Timeout Response

Problem: Due to this erratum, a PECI command other than GetDIB(), Ping(), or GetTemp() received before RESET_N is de-asserted may result in a timeout (0x81 completion code) for all subsequent such commands.

Implication: Future PECI commands other than GetDIB(), Ping(), and GetTemp() will not be serviced after this erratum occurs.

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

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

HSX11 System May Hang When Using the Transaction Layer Packet Processing Hint (TPH) Prefetch Hint

Problem: When all enabled cores on a socket are simultaneously in core C3, core C6, or package C6 state and a PCIe TPH with the prefetch hint set is received, the system may hang.

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

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

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

HSX12 TS1s Do Not Convey The Correct Transmitter Equalization Values During Recovery.RcvrLock

Problem: The PCIe 3.1 Base Specification requires that TS1s sent during Recovery.RcvrLock following 8.0 GT/s adaptive equalization (EQ) contain the final transmitter preset number and coefficient values that were requested by an endpoint during phase 2 of the EQ. Due to this erratum, TS1s with incorrect transmitter preset number values may be sent during Recovery.RcvrLock following 8.0 GT/s adaptive equalization.

Implication: Endpoints that check these values may, when unexpected values are found, request equalization restart in subsequent TSs it sends. If the EQ requests from the endpoint are supported in the BIOS or OS, the EQ will be restarted and the link may continue this EQ loop indefinitely.

Workaround: None identified.

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

HSX13 MSR_TEMPERATURE_TARGET Model Specific Register (MSR) May Read as '0'

Problem: Due to this erratum, reading the MSR_TEMPERATURE_TARGET MSR (1A2H) may incorrectly return '0'.

Implication: Software that depends on the contents of the MSR_TEMPERATURE_TARGET MSR may not behave as expected.

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

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

HSX14 PECI RdIAMSR() Command May Fail After Core C6 State is Entered

Problem: Reading core Machine Check Bank registers using the PECI RdIAMSR() command may fail after core C6 state has been entered.

Implication: Invalid data may be returned when using a PECI to read core Machine Check Bank registers.

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

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

HSX15 Closed Loop Thermal Throttling (CLTT) May Cause the BIOS To Hang On a Subsequent Warm Reset

Problem: If the CLTT is enabled when a warm reset is requested, due to this erratum, the processor will resume DIMM temperature polling before the memory sub-system has been re-initialized.

Implication: This erratum may lead to a BIOS hang. The warm reset request will fail, along with subsequence warm reset attempts. The failing condition is cleared by a cold reset.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX16 PCIe Extended Tag Field May be Improperly Set

Problem: The Extended Tag field in the Transaction Layer Packet (TLP) Header will not be zero for TLPs issued by PCIe ports 1a, 1b, 2c, 2d, 3c, and 3d even when the Extended Tag Field Enable bit in the Device Control Register (Offset 08H, bit 8) is 0.

Implication: This does not affect ports 0, 2a, 2b, 3a and 3b. This will not result in any functional issues when using device that properly track and return the full 8-bit Extended Tag value with the affected ports. However, if the Extended Tag field is not returned by a device connected to an affected port then this may result in unexpected completions and completion timeouts.

Workaround: None identified.

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

HSX17 A MOV to CR3 When Extended Page Tables (EPT) is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation

Problem: If the Extended Page Tables (EPT) is enabled, a MOV to CR3 or VMFUNC may be followed by an unexpected page fault or the use of an incorrect page translation.

Implication: Guest software may crash or experience unpredictable behavior as a result of this erratum.

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

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

HSX18 Memory Controller tsod_present Settings Being Improperly Cleared

Problem: On single Home Agent configurations, due to this erratum, the processor interferes with TSOD (thermal sensor on DIMM) usage by incorrectly clearing the tsod_present field (bits[7:0]) of the smbcntl_1 CSR (Bus 0; Device 19; Function 0; Offset 0x198) after the BIOS writes that field.

Implication: The CLTT will not work as expected.

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

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

HSX19 DDR4 Power Down Timing Violation

Problem: When DDR4 is operating at 2133 MHz, the processor's memory control may violate the JEDEC t_{PRPDEN} timing specification.

Implication: Violation of timing specifications can lead to unpredictable system behavior; however, Intel has not observed this erratum to impact the operation of any commercially available system using validated DIMMs by Intel Platform Memory Operations.

Workaround: None identified.

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

HSX20 Correctable Memory Error Correcting Code (ECC) Errors May Occur at Boot

Problem: With memory lockstep enabled, the system may experience correctable memory errors during boot with IA32_MCi_STATUS.MCACOD= 0x009x (where x is 0,1,2, or 3 and indicates the channel number reporting the error)

Implication: The system may experience correctable memory errors.

Workaround: None identified.

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

HSX21 Backup Tracker (BT) Timeouts May Cause Spurious Machine Checks

Problem: The BT timeout logic in the Home Agent can trigger spuriously, causing false machine checks indicated by IA32_MCi_STATUS.MSCOD=0x0200.

Implication: Due to this erratum, timeout machine check may occur.

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

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

HSX22 PCIe Type 1 Vendor Defined Message (VDM) May be Silently Dropped

Problem: Due to this erratum, a PCIe Type 1 VDMs is silently dropped unless the vendor ID is the Management Component Transport Protocol (MCTP) value of 0x1AB4.

Implication: PCIe Type 1 VDMs may be unexpectedly dropped. Intel has not observed this erratum to impact the operation of any commercially available system.

Workaround: None identified.

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

HSX23 CONFIG_TDP_NOMINAL CSR Implemented at Incorrect Offset

Problem: The PCIe Base Specification indicates that Configuration Space Headers have a base address register at offset 0x10. Due to this erratum, the Power Control Unit's CONFIG_TDP_NOMINAL CSR (Bus 1; Device 30; Function 3; Offset 0x10) is located where a base address register is expected.

Implication: Software may treat the CONFIG_TDP_NOMINAL CSR as a base address register leading to a failure to boot.

Workaround: None identified.

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

HSX24 A Machine-Check Exception Due to Instruction Fetch May Be Delivered Before an Instruction Breakpoint

Problem: Debug exceptions due to instruction breakpoints take priority over exceptions resulting from fetching an instruction. Due to this erratum, a machine-check exception resulting from the fetch of an instruction may take priority over an instruction breakpoint if the instruction crosses a 32-byte boundary and the second part of the instruction is in a 32-byte poisoned instruction fetch block.

Implication: Instruction breakpoints may not operate as expected in the presence of a poisoned instruction fetch block.

Workaround: None identified.

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

HSX25 Power Consumed During Package C6 May Exceed Specification

Problem: Due to this erratum, the processor power usage may be higher than specified for the VCCIN and/or IIO domains while in Package C6 state.

Implication: Systems may experience increased power consumption while the processor is in Package C6.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX26 Platform Performance Degradation When C1E is Enabled

Problem: Due to this erratum, when C1E is enabled and after the processor has entered Package C1E state, core clock frequency becomes limited to its minimum value (sometimes referred to as Pn) until the system exits Package C3 state (or deeper) or the system is reset.

Implication: When this erratum occurs, operating frequency will be lower than expected.

Note: After a Package C3 exit, re-entering Package C1E state re-imposes this erratum's frequency limit.

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

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

HSX27 PCIe Correctable Error Status Register May Not Log Receiver Error at 8.0 GT/s

Problem: Due to this erratum, correctable PCIe receiver errors may not be logged in the DPE field (bit 15) of the PCISTS CSR (Bus: 0; Device 1,2,3; Function 0-1, 0-3, 0-3; Offset 6H) when operating at 8.0 GT/s.

Implication: Correctable receiver errors during 8.0 GT/s operation may not be visible to the OS or driver software.

Workaround: None identified.

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

HSX28 PCIe Hot-Plug Slot Status Register May not Indicate Command Completed

Problem: The PCIe Base Specification requires a write to the Slot Control register (Offset A8H) to generate a hot plug command when the downstream port is hot plug capable. Due to this erratum, a hot plug command is generated only when one or more of the Slot Control register bits [11:6] are changed.

Implication: Writes to the Slot Control register that leave bits [11:6] unchanged will not generate a hot plug command and will therefore not generate a command completed event. Software that expects a command completed event may not behave as expected.

Workaround: It is possible for software to implement a one-second timeout in lieu of receiving a command completed event.

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

HSX29 Local PCIe Peer-to-Peer Traffic on x4 Ports May Cause a System Hang

Problem: Under certain conditions, Peer-to-Peer traffic with x4 PCIe ports on the same processor (for example, local) may cause a system hang.

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

Workaround: None identified. Local Peer-to-Peer traffic should not be used to or from x4 PCIe ports.

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

HSX30 ILLC Error Conditions May be Dropped or Incorrectly Signaled

Problem: When two Last Level Cache (LLC) errors happen in close proximity, a Uncorrectable No Action Required (UCNA) machine check may be dropped or a spurious machine check or the Corrected Machine Check Interrupt (CMCI) may be issued. Further, when this erratum occurs, the merged CBo LLC machine check bank IA32_MC[17-19]_STATUS MSRs may be incorrect.

Implication: IA32_MC[17-19]_STATUS MSR may not reflect most current error.

Workaround: It is possible for the BIOS to contain a partial workaround for this erratum. The workaround does not address the potential dropped a UCNA machine check.

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

HSX31 A DDR4 Command/Address (C/A) Parity Error in Lockstep Mode May Result in a Spurious Uncorrectable Error

Problem: If a memory C/A parity error occurs while the memory subsystem is configured in lockstep mode then the channel that observed the error will properly log the error but the associated channel in lockstep will incorrectly log an uncorrectable error in its IA32_MCi_STATUS MSR.

Implication: Due to this erratum, incorrect logging of an uncorrectable memory error in IA32_MCi_STATUS may occur.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX32 Some Integrated Memory Controller (IMC) and Intel® QPI Functions Have Incorrect PCI Capability Pointer Register (CAPPTR) Values

Problem: The PCI CAPPTR is defined to contain the offset to the capabilities list structure when the PCI Status Register (PCI PCISTS) bit 4 (Capabilities_List) is set to 1. Due to this erratum, CAPPTR (offset 0x34) should hold a value of 0x40 but is instead zero for these IMC and Intel® QPI device:

Device 8, functions 3,5,6

Device 9, functions 3,5,6

Device 10, functions 3,5,6

Device 19, functions 0-5

Device 20, functions 0-3

Device 21, functions 0-3

Device 22, functions 0-3

Device 23, functions 0-3

Implication: Software that depends on CAPPTR to access additional capabilities may not behave as expected.

Workaround: Software that needs to access these capabilities must take this erratum into account.

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

HSX33 PCIe TLP Translation Request Errors Are Not Properly Logged For Invalid Memory Writes

Problem: A PCIe Memory Write TLP with an AT field value of 01b (address translation request) does not set the UR (Unsupported Request) bit (UNCERRSTS CSR, Bus 0; Device 0; Function 0; Offset 0x14C; Bit 20) as required by the PCIe Base Specification.

Implication: System or software monitoring error status bits may not be notified of an unsupported request. When this erratum occurs, the processor sets the 'advisory_non_fatal_error_status' bit (CORERRSTS CSR, Bus 0; Device 0; Function 0; Offset 0x158; Bit 13) and drops the failing transaction.

Workaround: None identified.

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

HSX34 Consecutive PECI RdIAMSR Commands When Core C6 is Enabled May Cause a System Hang

Problem: Consecutive PECI RdIAMSR commands to access core Machine Check MSRs can result in a system hang when core C6 state is enabled.

Implication: When this erratum occurs, the PECI commands can lead to a system hang.

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

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

HSX35 Enabling Targeted Row Refresh (TRR) With DDR4 LRDIMMs May Lead to Unpredictable System Behavior

Problem: Due to this erratum, the TRR is not compatible with DDR4 LRDIMMs.

Implication: Unpredictable system behavior may occur.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX36 Command Address (C/A) Parity Error Injection May Cause the System to Hang

Problem: When the C/A parity error injections are occurring too frequently, the home agent may be prevented from completing memory transactions. This may result in an internal timer error indicated by IA32_MCi_STATUS. MSCOD=0x0080 and IA32_MCi_STATUS. MCACOD=0x0400.

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

Workaround: Ensure there is at least 30 µs of delay between injections.

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

HSX37 The System May Shut Down Unexpectedly During a Warm Reset.

Problem: Certain complex internal timing conditions present when a warm reset is requested can prevent the orderly completion of in-flight transactions. It is possible under these conditions that the warm reset will fail and trigger a full system shutdown.

Implication: When this erratum occurs, the system will shut down and all machine check error logs will be lost.

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

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

HSX38 Patrol Scrubbing of Mirrored Memory May Log Spurious Memory Errors

Problem: The Patrol scrubber, when mirroring is enabled, may incorrectly identify certain data patterns as poison data or as memory errors.

Implication: Spurious memory errors and poisoned data may be logged when mirroring is enabled.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX39 MSR_TURBO_ACTIVATION_RATIO MSR Cannot be Locked

Problem: Setting the TURBO_ACTIVATION_RATIO_LOCK field (bit 31) of the MSR_TURBO_ACTIVATION_RATIO MSR (64CH) has no effect; it does not block future writes to the MSR_TURBO_ACTIVATION_RATIO MSR.

Implication: Software cannot rely on locking MSR_TURBO_ACTIVATION_RATIO MSR.

Workaround: None identified.

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

HSX40 The System May Shut Down Unexpectedly During a Warm Reset

Problem: Certain complex internal timing conditions present when a warm reset is requested can prevent the orderly completion of in-flight transactions. It is possible under these conditions that the warm reset will fail and trigger a full system shutdown.

Implication: When this erratum occurs, the system will shut down and all machine check error logs will be lost.

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

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

HSX41 Invalid Intel® QuickData Technology XOR Descriptor Source Addressing May Lead to Unpredictable System Behavior

Problem: Intel® QuickData Technology (for example, Crystal Beach DMA v3.2) does not correctly halt and report aborts on illegal source addresses placed in a CBDMA descriptor regardless of type (Legacy or PQ). This abort condition may cause unpredictable system behavior.

Implication: This erratum may lead to unpredictable system behavior.

Workaround: Ensure XOR DMA descriptor source addresses targets valid DRAM memory locations.

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

HSX42 Warm Reset May Cause PCIe Hot-Plug Sequencing Failure

Problem: The Integrated I/O unit uses the Virtual Pin Port (VPP) to communicate with power controllers, switches, and LEDs associated with PCIe Hot-Plug sequencing. Due to this erratum, a warm reset occurring when a VPP transaction is in progress may result in an extended VPP stall, termination of the in-flight VPP transaction, or a transient power down of slots subject to VPP power control.

Implication: During or shortly after a warm reset, when this erratum occurs, PCIe Hot-Plug sequencing may experience transient or persistent failures or slots may experience unexpected transient power down events. In certain instances, a cold reset may be needed to fully restore operation.

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

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

HSX43 PCIe Uncorrectable Response (UR) And Completer Abort (CA) Responses May be Sent Before Link Enters Live Error Recovery (LER) State

Problem: Completions with the UR and the CA status should trigger the LER. Further, these packets should be dropped upon entering LER. Due to this erratum, these completions may not be dropped when LER is triggered.

Implication: Since these packets contain no data, there is no loss of error containment. These packets will trigger LER mode; the link will be disabled.

Workaround: None identified.

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

HSX44 Surprise Down Error Status is Not Set Correctly on DMI Port

Problem: Due to this erratum, the Surprise_down_error_status (UNCERRSTS Device0; Function); Offset 0x14C; bit5) is not set to 1 when DMI port detects a surprise down error.

Implication: Surprise down errors will not be logged for the DMI port. This violates the PCIe Base Specification. Software that relies on this status bit may not behave as expected.

Workaround: None identified.

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

HSX45 Intel SMI2 in Half Width Mode With Dual Device Data Correction (DDDC) Enabled Will Not Report RdECC Errors

Problem: When a RdECC error occurs on an Intel SMI2 channel operating at a 1:1 ratio half-width mode and the DDDC enabled, the error logging in the Intel® C102/104/112/114 Scalable Memory Buffer and the processor is not properly coordinated.

Implication: Although the error flow is correct, error isolation may be affected because the processor may log a RdECC error while the Intel C102/104/112/114 Scalable Memory Buffer does not log an error.

Workaround: None identified.

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

HSX46 PCIe SLTCON CSRs electromechanical_interlock_control Field Read as 1

Problem: The PCI Express Base Specification rev 3.1 requires that the SLTCON (Bus 0;Device 3-0;Function 3-0;Offset 0xA8) CSRs' electromechanical_interlock_control (bit 11) "always returns aa 0 when read". Due to this erratum, a read of this bit returns the last value written.

Implication: Software expecting a value of 0 may not function as expected. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: Software should ignore read values returned from this register field.

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

HSX47 Intel DDR3 SMI2 Command Address Parity (CAP) Errors Are Ignored Leading to Unpredictable System Behavior

Problem: An Intel SMI2 DDR3 the CAP error, rather than initiating a mirroring event as expected, is ignored leading to unpredictable system behavior.

Implication: When this erratum occurs, it may lead to unpredictable system behavior. Note that this behavior does not affect DDR4 memory subsystems.

Workaround: It is possible for the BIOS to contain processor configuration data and code changes as a workaround for this erratum.

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

HSX48 PECI RdPkgConfig Command DRAM Services May Behave Incorrectly

Problem: The PECI RdPkgConfig command may return incorrect results when accessing the DRAM Thermal Interface (indices 14 and 22).

Implication: Thermal monitoring and control using a PECI may not behave as expected.

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

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

HSX49 Some OFFCORE_RESPONSE Performance Monitoring Events May Undercount

Problem: The performance monitoring events OFFCORE_RESPONSE (Events B7H and BBH) should count uncore responses matching the request-response configuration specified in MSR_OFFCORE_RSPs (1A6H and 1A7H, respectively) for core-originated requests. However due to this erratum, COREWB (bit 3), PF_L3_DATA_RD (bit 7), PF_L3_RFO (bit 8), PR_L3_CODE_RD (bit 9), SPLIT_LOCK_UC_LOCK (bit 10), and STREAMING_STORES (bit 11) request types may undercount.

Implication: These performance monitoring events may not produce reliable results for the listed request types.

Workaround: None identified.

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

HSX50 Performance Monitoring OFFCORE_RESPOSE_{1,2} Events May Miscount L3_MISS_REMOTE_HOP

Problem: When a Performance Monitoring counter is configured to count OFF_CORE_RESPONSE_{1,2} (Events B7H and B8H), data obtained for remote DRAM may be attributed to L3_MISS_REMOTE_HOP0 (as programmed by MSR_OFFCORE_RSP_{1,2} (MSRs 1A6H, 1A7H) bit 27) instead of L3_MISS_REMOTE_HOP1 (bit 28) or L3_MISS_REMOTE_HOP2P (bit 29). Data provided from remote caching agent associated with remote DRAM is unaffected.

Implication: L3_MISS_REMOTE_HOP0 may over count, while L3_MISS_REMOTE_HOP1 and L3_MISS_REMOTE_HOP2P may undercount.

Workaround: None identified. Set all three configuration bits (L3_MISS_REMOTE_HOP0, L3_MISS_REMOTE_HOP1, L3_MISS_REMOTE_HOP2P) to obtain the total count of data supplied by remote agents.

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

HSX51 Some DRAM and L3 Cache Performance Monitoring Events May Count Incorrectly

Problem: Due to this erratum, the supplier information may become stale, and the following events may count incorrectly:
MEM_LOAD_UOPS_RETIRED.L3_HIT (Event D1H Umask 04H)
MEM_LOAD_UOPS_RETIRED.L3_MISS (Event D1H Umask 20H)
MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_MISS (Event D2H Umask 01H)
MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HIT (Event D2H Umask 02H)
MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HITM (Event D2H Umask 04H)
MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_NONE (Event D2H Umask 08H)
MEM_LOAD_UOPS_L3_MISS_RETIRED.LOCAL_DRAM (Event D3H Umask 01H)
MEM_TRANS_RETIRED.LOAD_LATENCY (Event CDH Umask 01H)
PAGE_WALKER_LOADS.DTLB_L3 (Event BCH Umask 14H)
PAGE_WALKER_LOADS.ITLB_L3 (Event BCH Umask 24H)
PAGE_WALKER_LOADS.DTLB_Memory (Event BCH Umask 18H)
PAGE_WALKER_LOADS.ITLB_Memory (Event BCH Umask 28H)

Implication: The affected events may count incorrectly, resulting in inaccurate memory profiles. For the affected events that are precise, PEBS records may be generated at incorrect points. Intel has observed incorrect counts by as much as 40%.

Workaround: None identified.

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

HSX52 Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations

Problem: If the VM exit occurs from a guest with primary processor-based VM-execution control "activate secondary controls" set to 0 and the secondary processor-based VM-execution control "enable VPID" set to 1, then after a later VM entry with VPID fully enabled ("activate secondary controls" and "enable VPID" set to 1), the processor may use stale linear address translations.

Implication: The processor may incorrectly translate linear addresses. Intel has not observed this erratum with any commercially available software.

Workaround: Software should not enter a guest with "enable VPID" set to 1 when "activate secondary controls" is set to 0.

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

HSX53 An IRET Instruction That Results in a Task Switch Does Not Serialize The Processor

Problem: An IRET instruction that results in a task switch by returning from a nested task does not serialize the processor (contrary to the Software Developer's Manual Vol. 3 section titled "Serializing Instructions").

Implication: Software which depends on the serialization property of IRET during task switching may not behave as expected. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified. Software can execute an MFENCE instruction immediately prior to the IRET instruction if serialization is needed.

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

HSX54 A P-State or C-State Transition May Lead to a System Hang

Problem: For a small subset of parts under elevated die temperature conditions, a P-state or C-state transition may result in a system timeout or system shutdown.

Implication: When this erratum occurs, the system may shutdown or report a timeout error; Intel has observed transaction completion timeouts and other internal timeouts.

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

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

HSX55 PECI Frequency Limited to 1 MHz

Problem: The PECI 3.1 specification's operating frequency range is 0.2 MHz to 2 MHz. Due to this erratum, the PECI may be unreliable when operated above 1 MHz.

Implication: Platforms attempting to run the PECI above 1 MHz may not behave as expected.

Workaround: None identified. Platforms should limit the PECI operating frequency to 1 MHz.

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

HSX56 VGATHERQPS That Loads an Element From The Advance Programmable Interrupt Controller (APIC)-Access Page May Load Other Elements From Incorrect Addresses

Problem: If the "virtualize APIC accesses" VM-execution control is 1, a 256-bit VGATHERQPS with an element that maps to the APIC-access page may use incorrect addresses to load other elements.

Implication: Loading from an incorrect address can result in unexpected behavior with respect to data, faults or VM exits. This erratum will occur only if a guest operating system attempts to access the APIC using the VGATHERQPS instruction. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

HSX57 A Spurious Patrol Scrub Error May be Logged

Problem: When a memory ECC error occurs, a spurious patrol scrub error may also be logged on another memory channel.

Implication: A patrol scrub correctable error may be incorrectly logged.

Workaround: None Identified.

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

HSX58 DRAM Device Failure With Error Flow Registers Enabled May Result in a Machine Check

Problem: During a high rate of correctable DRAM error events, a DDR4 CAP error may result in an unrecoverable machine check with an IA32_MCi_STATUS.MCACOD (bits [15:0]) value of 0000 0000 100x xxxx (where x can be 0 or 1) and an IA32_MCi_STATUS.MSCOD (bits [31:16]) value of 0x0200 or with an IA32_MCi_STATUS.MCACOD value of 0000 0001 xxxx xxxx and an IA32_MCi_STATUS.MSCOD value of 0x000C.

Implication: A high rate of correctable errors, typically the result of a DRAM device failure, may result in an unrecoverable machine check.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum.

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

HSX59 Monitor Trap Flag (MTF) VM Exit on XBEGIN Instruction May Save State Incorrectly

Problem: Execution of an XBEGIN instruction while the "monitor trap flag" VM-execution control is 1 will be immediately followed by an MTF VM exit. If advanced debugging of RTM transactional regions has been enabled, the VM exit will erroneously save the address of the XBEGIN instruction as the instruction pointer (instead of the fallback instruction address specified by the XBEGIN instruction). In addition, it will erroneously set bit 16 of the pending-debug-exceptions field in the VMCS indicating that a debug exception or a breakpoint exception occurred.

Implication: Software using the monitor trap flag to debug or trace transactional regions may not operate properly. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

HSX60 PEBS Record May Be Generated After Being Disabled

Problem: A performance monitoring counter may generate a Precise Event Based Sampling (PEBS) record after disabling PEBS or the performance monitoring counter by clearing the corresponding enable bit in IA32_PEBS_ENABLE MSR (3F1H) or IA32_PERF_GLOBAL_CTRL MSR (38FH).

Implication: A PEBS record generated after a VMX transition will store into memory according to the post-transition the Debug Store (DS) configuration. These stores may be unexpected if PEBS is not enabled following the transition.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. A software workaround is possible through disallowing PEBS during VMX non-root operation and disabling PEBS prior to VM entry.

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

HSX61 PCIe Ports Do Not Support Data Link Layer (DLL) Link Active Reporting

Problem: The PCIe Base Specification requires every "Downstream Port that supports Link speeds greater than 5.0 GT/s" to support the DLL Link Active Reporting, However, the PCIe ports do not support DLL Link Active Reporting.

Implication: Due to this erratum, the PCIe ports do not support DLL Link Active Reporting. This may be reported by a PCIe compliance test.

Workaround: None identified.

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

HSX62 PCIe Link Speed Negotiation May Fail After Link is Re-enabled

Problem: If a PCIe link is established then disabled and the link partner's advertised speeds are changed while the link is disabled, the link may fail to correctly negotiate link speed when it is re-enabled.

Implication: Due to this erratum, the PCIe link speed negotiation may fail after re-enabling a disabled port.

Workaround: None identified.

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

HSX63 PROCHOT# Assertion During Warm Reset May Cause Persistent Performance Reduction

Problem: Assertion of PROCHOT# after RESET# de-assertion but before the BIOS has completed reset initialization (indicated by CPL3) may result in persistent processor throttling. Asserting PROCHOT# during and after RESET# assertion for the Fault Resilient Boot (FRB) tri-stating of the processor is not affected by this erratum.

Implication: When this erratum occurs, the resultant persistent throttling substantially reduces the processor's performance.

Workaround: None identified.

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

HSX64 Data Breakpoint Coincident With a Machine Check Exception May be Lost

Problem: If a data breakpoint occurs coincident with a machine check exception, then the data breakpoint may be lost.

Implication: Due to this erratum, a valid data breakpoint may be lost.

Workaround: None identified.

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

HSX65 A Corrected Internal Parity Error May Result in a System Hang

Problem: A corrected Internal Parity Error (IA32_MC0_STATUS.MCACOD=0005H and IA32_MC0_STATUS.MSCOD=000FH, MSR 401H bits [15:0] and bits [31:16] respectively) may cause a system hang.

Implication: Due to this erratum, a corrected internal parity error may cause a system hang. Reset, SMI, or INIT will end the system hang.

Workaround: None identified.

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

HSX66 IA32_MC0_STATUS May be Incorrect After a Machine Check Overflow

Problem: When a Software Recoverable Action Required (SRAR) error is logged in IA32_MC0_STATUS (MSR 401H), a subsequent fatal error will correctly set the processor context corrupted (PCC) (bit 57) and OVER (bit 62) flags but may fail to update other fields in IA32_MC0_STATUS.

Implication: When this erratum occurs, the error logged in IA32_MC0_STATUS is an invalid mixture of an SRAR error with the PCC flag set. Software that reads IA32_MC0_STATUS may not behave as expected.

Workaround: None identified.

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

HSX67 IA32_MC1_STATUS.MISCV May be Incorrect on a Machine Check Overflow

Problem: If software writes machine check bank 1 to contain a valid UC error with IA32_MC1_STATUS.MISCV cleared, another UC machine check may set IA32_MC1_STATUS.MISCV in violation of the overwrite rules.

Implication: When there is a machine check overflow in MC1, the IA32_MC1_STATUS.MISCV cannot be relied upon to indicate that IA32_MC1_MISCV is valid. All UC errors in MC1 will set MISCV.

Workaround: Software should not write UC errors into IA32_MC1_STATUS with IA32_MC1_STATUS.MISCV=0. Software may write IA32_MC1_MISC=0 to indicate no valid data in IA32_MC1_MISC.

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

HSX68 PEBS EventingIP Field May Be Incorrect Under Certain Conditions

Problem: The EventingIP field in the PEBS record reports the address of the instruction that triggered the PEBS event. Under certain complex microarchitectural conditions, the EventingIP field may be incorrect.

Implication: When this erratum occurs, performance monitoring software may not attribute the PEBS events to the correct instruction.

Workaround: None identified.

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

HSX69 Resume Flag (RF) May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or Branch Trace Store (BTS)

Problem: After a fault due to a failed PEBS or a BTS address translation, the RF may be incorrectly set in the EFLAGS image that is saved.

Implication: When this erratum occurs, a code breakpoint on the instruction following the return from handling the fault will not be detected. This erratum only happens when the user does not prevent faults on PEBS or BTS.

Workaround: Software should always prevent faults on PEBS or BTS.

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

HSX70 An APIC Timer Interrupt During Core C6 Entry May be Lost

Problem: Due to this erratum, an APIC timer interrupt coincident with the core entering C6 state may be lost rather than held for servicing later.

Implication: A lost APIC timer interrupt may lead to missed deadlines or a system hang.

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

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

HSX71 MOVNTDQA From Write Combining (WC) Memory May Pass Earlier Locked Instructions

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

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

Workaround: None identified. Software that relies on a locked instruction to fence subsequent executions of (V)MOVNTDQA should insert an MFENCE instruction between the locked instruction and subsequent (V)MOVNTDQA instruction.

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

HSX72 An x87 Store Instruction Which Pends #PE While EPT is Enabled May Lead to an Unexpected Machine Check and/or Incorrect x87 State Information

Problem: The execution of an x87 store instruction which causes a Precision Exception (#PE) to be pended and also causes a VM-exit due to an EPT violation or misconfiguration may lead the VMM logging a machine check exception with a cache hierarchy error (IA32_MCi_STATUS.MCACOD = 0150H and IA32_MCi_STATUS.MSCOD = 000FH). Additionally, FSW.PE and FSW.ES (bits 5 and 7 of the FPU Status Word) may be incorrectly set to 1, and the x87 Last Instruction Opcode (FOP) may be incorrect.

Implication: When this erratum occurs, the VMM may receive an expected machine check exception and software attempting to handle the #PE may not behave as expected.

Workaround: None identified.

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

HSX73 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: For the affected steppings, see the Summary Tables of Changes.

HSX74 Certain PerfMon Events May be Counted Incorrectly When The Processor is Not in C0 State

Problem: Due to this erratum, the PerfMon events listed below may be counted when the logical processor is not in C0 State.
IDQ.EMPTY (Event 79H, Umask 02H)
IDQ_UOPS_NOT_DELIVERED.CORE (Event 9CH, Umask 01H)
RESOURCE_STALLS.ANY (Event A2H, Umask 01H)

Adding the following 3 events to existing erratum:
CYCLE_ACTIVITY.CYCLES_LDM_PENDING (Event A3H, Umask 02H, Cmask 02H)
CYCLE_ACTIVITY.CYCLES_NO_EXECUTE (Event A3H, Umask 04H, Cmask 04H)
CYCLE_ACTIVITY.STALLS_LDM_PENDING (Event A3H, Umask 06H, Cmask 06H)

Implication: The count will be higher than expected.

Workaround: None identified.

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

HSX75 Writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP May #GP When Intel® Transactional Synchronization Extensions (Intel® TSX) is Not Supported

Problem: Due to this erratum, on processors that do not support Intel® TSX (CPUID.07H.EBX bits 4 and 11 are both zero), writes to MSR_LASTBRANCH_x_FROM_IP (MSR 680H to 68FH) and MSR_LER_FROM_LIP (MSR 1DDH) may #GP unless bits[62:61] are equal to bit[47].

Implication: The value read from MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP is unaffected by this erratum; bits [62:61] contain IN_TSX and TSX_ABORT information respectively. Software restoring these MSRs from saved values are subject to this erratum.

Workaround: Before writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP, ensure the value being written has bit[47] replicated in bits[62:61]. This is most easily accomplished by sign extending from bit[47] to bits[62:48].

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

HSX76 JTAG Boundary Scan For Intel® QPI and PCIe Lanes May Report Incorrect Stuck at 1 Errors

Problem: Boundary Scan testing of the Intel® QPI and PCIe interfaces may incorrectly report a recurring stuck at 1 failure on Intel® QPI and PCIe receiver lanes. This erratum only affects Boundary Scan testing and does not affect functional operation of the Intel® QPI and PCIe interfaces.

Implication: This erratum may result in Boundary Scan test failures reported on one or more of the Intel® QPI and PCIe lanes.

Workaround: None identified.

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

HSX77 APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode

Problem: After writing to the IA32_TSC_ADJUST MSR (3BH), any subsequent write to the IA32_TSC_DEADLINE MSR (6E0H) may incorrectly process the desired deadline. When this erratum occurs, the resulting timer interrupt may be generated at the incorrect time.

Implication: When the local APIC timer is configured for TSC-Deadline mode, a timer interrupt may be generated much earlier than expected or much later than expected. Intel has not observed this erratum with most commercially available software.

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

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

HSX78 Debug Exceptions May Be Lost in The Case Of Machine Check Exception

Problem: If both a machine check exception and a debug exception are pending on the same instruction boundary, then the machine check exception gets priority and the debug exception may be lost, even if the PCC field is cleared in all of the machine check banks (bit 57=0 in all IA32_MCi_STATUS MSR). This can happen in the case that an instruction triggered a data breakpoint while an unrelated machine check event was received.

Implication: Debugging software may fail to operate

Workaround: None identified.

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

HSX79 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: For the affected steppings, see the Summary Tables of Changes.

HSX80 Using the Intel® TSX Instructions May Lead to Unpredictable System Behavior

Problem: Under complex microarchitectural conditions, software using the 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 the BIOS to contain a workaround for this erratum.

Workaround: Performance Monitoring Impact of the Intel® TSX Memory Ordering Issue

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

HSX81 Spurious Corrected Errors May be Reported

Problem: Due to this erratum, spurious corrected errors may be logged in the IS32_MC0_Status register with the valid field (bit 63) set, the uncorrected error field (bit 61) not set, a Model Specific Error code (bits [31:16]) of 0x000F, and an MCA Error Code (bits [15:0]) of 0x0005. If the CMCI is enable, these spurious corrected errors also may signal interrupts.

Implication: When this erratum occurs, software may see corrected errors that are benign. these corrected errors may be safely ignored.

Workaround: None Identified.

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

HSX82 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 (for example, 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: For the affected steppings, see the Summary Tables of Changes.

§

Specification Changes

There are no specification changes in this specification update revision.

§

Specification Clarifications

There are no specification clarifications in this specification update revision.

§

Documentation Changes

There are no documentation changes.

§