



Intel[®] Xeon[®] D-2700 and D-2800 Processor Families

Specification Update

April 2024

Revision 007US



Intel technologies may require enabled hardware, software or service activation.

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.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Performance varies by use, configuration, and other factors. Learn more on the [Performance Index site](#).

Your costs and results may vary.

"Conflict-free" refers to products, suppliers, supply chains, smelters, and refiners that, based on our due diligence, do not contain or source tantalum, tin, tungsten or gold (referred to as "conflict minerals" by the U.S. Securities and Exchange Commission) that directly or indirectly finance or benefit armed groups in the Democratic Republic of the Congo or adjoining countries.

All product plans and roadmaps are subject to change without notice.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

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.

Results have been estimated or simulated.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visiting the [Intel Resource and Document Center](#).

© 2024 Intel Corporation. Intel, the Intel logo, Xeon, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Contents

Preface	5
Affected Documents/Related Documents	5
Nomenclature	6
Identification Information	7
Component Identification Using Programming Interface	7
Component Marking Information	8
Summary Tables of Changes	9
Codes Used in Summary Tables	9
CPU Errata Details	14
LAN/Ethernet Errata Details	35
Specification Changes	36
Specification Clarifications	37
Documentation Changes	38



Revision History

Date	Revision	Description
April 2024	007US	Added erratum ICXD96 .
December 2023	006US	Added erratum ICXD95 .
November 2023	005US	Added erratum ICXD94 . Added section for LAN/Ethernet Errata Details . Added LAN/Ethernet Errata Summary Table . Added errata LAN1. , LAN2. , LAN3 .
October 2023	004US	Added errata ICXD82 . through ICXD93 . Updated errata ICXD22. , ICXD25. , ICXD35. , ICXD44. , ICXD47. , ICXD58. , ICXD65. , ICXD66. , ICXD81 . Updated Table 1 , "Component Identification Via Registers" on page 7. Updated figure name: Table 2 , "Processor Top Side Marking (Example)" on page 8.
July 2023	003US	Updated title and naming. Updated Component Marking Information . Updated Summary Tables of Changes . Updated errata ICXD21. , ICXD30. , ICXD35. , ICXD58 . Added errata ICXD68 . through ICXD81 .
November 2022	002US	Added errata ICXD65. , ICXD66. , ICXD67 . Errata numbers are changed from ICXD46 through ICXD5S . Updated Table 1 , "Component Identification Via Registers" on page 7. Updated Table 3 , "CPU Errata Summary Table" on page 10.
February 2022	001US	Initial Release.

§

Preface

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

Information types defined in [Nomenclature](#) are consolidated into the specification update and are no longer published in other documents.

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

Affected Documents/Related Documents

Document Title	Document Number/Location
Software Documents	
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual</i>	325462 ¹
<i>Intel® 64 and IA-32 Intel Architectures Optimization Reference Manual</i>	248966 ¹
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	252046 ¹
<i>Intel® Virtualization Technology Specification for Directed I/O Architecture Specification</i>	D51397 ¹

Note:

1. This document can be downloaded from <http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>.



Nomenclature

Errata are design defects or errors. These may cause the SoC 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.

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 on as described in the processor identification information table. Read all notes associated with each S-Spec number.

Qualification Detail Form (QDF) Number A several digit code used to distinguish between engineering samples. These processors are used for qualification and early design validation. The functionality of these parts can range from mechanical only to fully functional. The NDA specification update has a processor identification information table that lists these QDF numbers and the corresponding product sample details.

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

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

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

Note:

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

§

Identification Information


Component Identification Using Programming Interface

The CPU Identification (CPUID) instruction returns processor identification and feature information to the EAX, EBX, ECX, and EDX registers, as determined by input entered in EAX (and in some cases, ECX as well).

Details of the instruction can be found in the Instruction Set Reference portion of the *Intel® 64 and IA-32 Architectures Software Developer’s Manual*.

When CPUID executes with EAX set to 01h, the SoC version information is returned in EAX. See [Figure 1, “Version Information Returned by CPUID in EAX”](#).

Figure 1. Version Information Returned by CPUID in EAX



This table shows the ID corresponding to each of the available steppings.

Table 1. Component Identification Via Registers

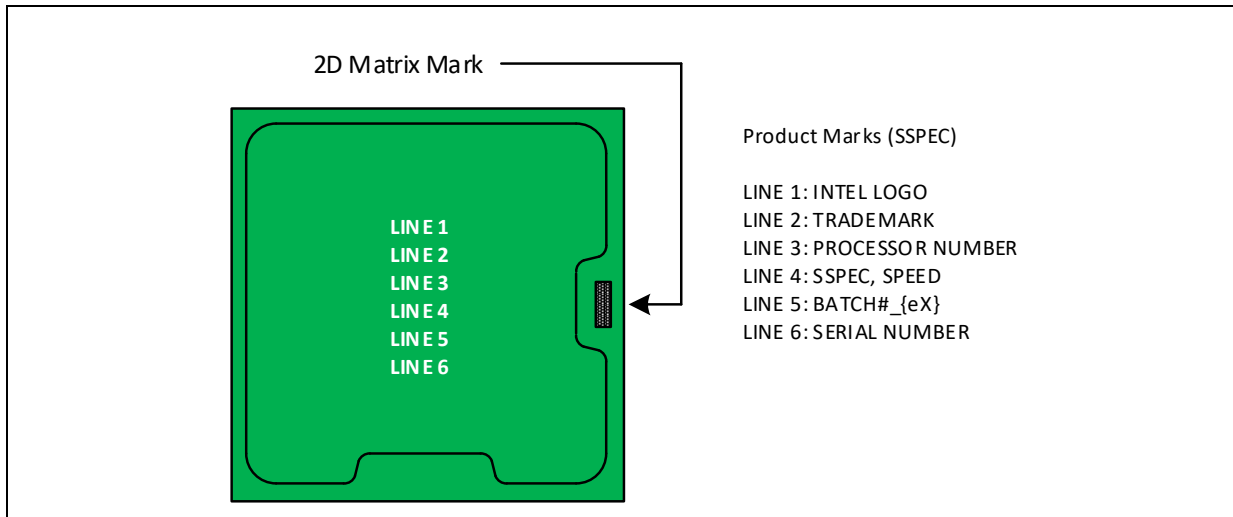
Identification Parameter	Stepping U1, U2
Stepping ID	0x1
CPUID	0x606C1
LPC REVID [B:0, D:31, F:0] + 8h, bits [7:0]	11h
CAPID4 bits [7:6] Physical Chop for Intel Xeon 2700/2800 die	10b

§

Component Marking Information

The processor can be identified by the following component markings.

Figure 2. Processor Top Side Marking (Example)



For Intel® Xeon® D-2700 and D-2800 Processor Families SKUs, see <https://ark.intel.com/content/www/us/en/ark/products/series/87041/intel-xeon-d-processor.html>.

§

Summary Tables of Changes

The following tables indicate the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the Intel® Xeon® D-2700 and D-2800 Processor Families. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations:

Codes Used in Summary Tables

Stepping	Description
(No mark) or (Blank box)	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status	Description
Doc	Document change or update will be implemented
Planned Fix	This erratum may be fixed in a future stepping of the product
Fixed	This erratum has been previously fixed in Intel hardware, firmware or software
No Fix	There are no plan to fix this erratum



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

Erratum ID	D-2700/ D-2800 Steppings	Title
	U1, U2	
ICXD1.	No Fix	Wrong Page Access Semantics May be Reported When Intel® SGX ENCLU[EMODPE] Instruction Generates Page Fault (#PF) Exception
ICXD2.	No Fix	Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May Cause a #GP
ICXD3.	No Fix	VMREAD/VMWRITE Instructions May Not Fail When Accessing an Unsupported Field in VMCS
ICXD4.	No Fix	VERR Instruction Inside VM-Entry May Cause DR6 to Contain Incorrect Values
ICXD5.	No Fix	Vector Masked Store Instructions May Cause Write Back of Cache Line Where Bytes Are Masked
ICXD6.	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store
ICXD7.	No Fix	SMRAM State-Save Area Above The 4 GB Boundary May Cause Unpredictable System Behavior
ICXD8.	No Fix	Single Correctable Error Can be Logged Twice if Patrol Scrub Reads Address When Read Transaction is in Flight to Same Address
ICXD9.	No Fix	Overlap Between APIC And SMRR2 Memory-Mapped Registers Will Not Signal a #GP
ICXD10.	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set
ICXD11.	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
ICXD12.	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
ICXD13.	No Fix	IA32_RTIT_STATUS.FilterEn Bit Might Reflect a Previous Value
ICXD14.	No Fix	IA32_MC1_STATUS MSR May Not Log Errors When IA32_MC1_CTL MSR is Set to Not Signal Errors
ICXD15.	No Fix	False MC1 Error Reported in The Shadow of a Internal Timer Error
ICXD16.	No Fix	Debug Exceptions May Be Lost or Misreported When MOV SS or POP SS Instruction is Not Followed by a Write to SP
ICXD17.	No Fix	CPUID TLB Information is Inaccurate
ICXD18.	No Fix	Configuring The PRMRR as Non-WB Might Lead to Incorrect VM-Exit Interruption Error Code
ICXD19.	No Fix	#GP on Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code
ICXD20.	No Fix	IERR Not Logged Correctly When Ubox Requested to Signal MSMI
ICXD21.	Fixed	CPU Complex PCIe* Surprise Link Down Events May Not be Reported
ICXD22.	No Fix	CPU Complex PCIe RPPPIO May Contain Incorrect Tag Value
ICXD23.	No Fix	Uncore MC Bank Registers Corrected Error Count Field May Not Have a Sticky Most Significant Bit
ICXD24.	No Fix	Poisoned Locked Bus Transactions May Not Allow Warm Reset to Correctly Reset The Processor
ICXD25.	No Fix	Spurious CPU Complex PCIe Link Parity Errors May be Logged
ICXD26.	No Fix	IBIST Receiver Error Overflow Register Field Cannot be Cleared by Software
ICXD27.	No Fix	MBM May Report Incorrect Bandwidth For Certain Access Strides
ICXD28.	No Fix	Correctable Errors May Set The Overflow Bit
ICXD29.	No Fix	Enabled Error May Not be Logged When Other Errors Are Disabled
ICXD30.		Removed

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

Erratum ID	D-2700/ D-2800 Steppings	Title
	U1, U2	
ICXD31.	No Fix	Machine Check Bank Status MSR May Not Set Overflow Bit When Multiple Uncorrectable Errors Occur
ICXD32.	No Fix	FERR Registers Are Not Getting Cleared When CHANERR Register is Being Cleared
ICXD33.	No Fix	IOMMU Translation Requests to Interrupt Range May Fail
ICXD34.	No Fix	Internal Firmware Errors May Not Set Error Enable Bit
ICXD35.	No Fix	CPU Complex PCIe Surprise Link Down Logging May be Unexpectedly Blocked
ICXD36.	No Fix	MBA 2.0 May Cause MMIO Traffic to be Throttled
ICXD37.	No Fix	Cannot Inject Errors Into PCLS Bits
ICXD38.	No Fix	CAP Error And ECC Error During ADC/ADDDC Sparing May Not be Corrected
ICXD39.	No Fix	NSR Field Attribute Does Not Comply With PCIe Base Specification 4.0
ICXD40.	Fixed	Intermittent Correctable Memory Errors May be Observed
ICXD41.	No Fix	Unmapped Intel® Quick Data Technology DMA Reads May Result in Unpredictable System Behavior
ICXD42.	No Fix	Unexpected System Behavior May Occur During INVD Instruction Execution
ICXD43.	No Fix	Configuring ADL May Prevent Package C-States Entry
ICXD44.	No Fix	CPU Complex PCIe Rx Common Mode Impedance May be Too Low During Reset or Power-Down
ICXD45.	No Fix	Uncore Semaphore Capability May Not Generate a Semaphore Error Machine Check Exception
ICXD46.	Fixed	WBINVD Delays May Lead to a Machine Check Exception
ICXD47.	Fixed	CPU Complex PCIe EB Error May Be Escalated to Receiver Errors
ICXD48.	Fixed	Unpredictable System Behavior May Occur Due to Memory Read Roundtrip Latency
ICXD49.	No Fix	A PECCI Request May Receive an Incorrect Response
ICXD50.	Fixed	IFU Internal Parity Error
ICXD51.	No Fix	CHA UCNA Errors May be Incorrectly Controlled by MCI_CTL Enable Bits
ICXD52.	Fixed	Mesh to Memory Timeout May Occur When TME is Enabled
ICXD53.	No Fix	Inaccurate Mesh to Memory Corrected Error Count
ICXD54.	Fixed	CHA Errors May be Reported Incorrectly After a Warm Reset
ICXD55.	Fixed	Processor May Not Successfully Enter ADR
ICXD56.	No Fix	Debug_Has_Occurred Bit May be Asserted
ICXD57.	Fixed	IOSFSB Timeout May Lead to MCE
ICXD58.	Fixed	SRIS-Configured CPU Complex PCIe Link May Fail to Train
ICXD59.	No Fix	Unexpected Rollover in MBM Counters
ICXD60.	No Fix	DPC Trigger Status Bit May Not Cleared
ICXD61.	No Fix	Processor May Not Allow Intel® DCI Tool to Operate at 100 MHz
ICXD62.	No Fix	Four Unsuccessful Global Reset Attempts Needed For S5 Entry
ICXD63.	No Fix	New_Century Bit of the TCO1_STS Register Should Not be Cleared
ICXD64.	No Fix	xHCI Host Controller Reset May Cause a System Hang
ICXD65.	No Fix	HSUART May Stop Functioning When DMA is Activated



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

Erratum ID	D-2700/ D-2800 Steppings	Title
	U1, U2	
ICXD66.	Fixed	PCH PCIe Link Partner May Not Enumerate
ICXD67.	No Fix	VREFCA Tolerance May Violate The JEDEC JESD79-4 Specification
ICXD68.	No Fix	Setting MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT Does Not Prevent The Three-strike Counter From Incrementing
ICXD69.	Fixed	SGX Enclave Attestation And Unseal May Fail
ICXD70.	Fixed	Warm Reset Will Set C1E Enable Without Changing TMRT
ICXD71.	Fixed	HWPM Max Ratio May Not be Capped at P1
ICXD72.	Fixed	DRAM Performance May be Reduced Following Microcode Update
ICXD73.	Fixed	PCIe Bandwidth Reduction When Intel VT-d is Enabled
ICXD74.	Fixed	PC6 Entry is Not Prevented With PECI Pkg C-state Entry Control
ICXD75.	No Fix	A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an Intel AVX Gather or REP MOVSB Instruction
ICXD76.	No Fix	Incorrect MCACOD For L2 Prefetch MCE
ICXD77.	No Fix	Call Instruction Wrapping Around The 32-Bit Address Boundary May Return to Incorrect Address
ICXD78.	Fixed	FEATURE_TUNING_1 May be Reset to Factory Defaults
ICXD79.	Fixed	Accesses to CHA Configuration Space Beyond the CHA Logical Limit May Fail
ICXD80.	No Fix	Monitor May Not be Triggered
ICXD81.	No Fix	Incorrect CPU Complex PCIe RCB Advertisement
ICXD82.	No Fix	Unsupported Request Error Will be Logged in IEH
ICXD83.	Fixed	PCH PCIe Root Ports May Operate at 8.0 GT/s Rather Than 5.0 GT/s
ICXD84.	Fixed	PCH PCIe Devices With LUR May Not Link
ICXD85.	Planned Fix	After a Warm Reset, The SMBus Host Controller May Incorrectly Set The CPE Bit
ICXD86.	No Fix	PCH PCIe Root Port Cluster 0 Does Not Correctly Process MCTP Messages
ICXD87.	Fixed	PCH PCIe 3.0 Link May Observe Link Errors After Speed Change
ICXD88.	Fixed	PECI Transactions to Certain Devices Will Time Out
ICXD89.	Fixed	System May Hang With USB 3.2 Ports Enabled
ICXD90.	Fixed	PCH PCIe Link Speed May Be Limited to 5.0 GT/s
ICXD91.	Fixed	Branch Predictor May Produce Incorrect Instruction Pointer
ICXD92.	No Fix	DRNG May Erroneously Return Poisoned Data
ICXD93.	Fixed	System May Hang During ADR C2F
ICXD94.	Fixed	IERR Maybe Seen During Warm Reset After OS Path Load of Microcode Revision IDs of 0x1000230 or 0x1000260
ICXD95.	No Fix	Reading TCO_RLD Register Sets SMBus Host Status In Use Status Bit
ICXD96.	No Fix	A Write to The TSC_Deadline MSR May Cause an Unexpected Timer Interrupt

Table 4. LAN/Ethernet Errata Summary Table

Erratum ID	D-2700/ D-2800 Steppings	Title
	U1, U2	
LAN1.	No Fix	Inverted FCS Does Not Cause Increment of CRC Error Count
LAN2.	No Fix	Incorrect Calculation of Inner UDP Checksum on Rx Packet
LAN3.	No Fix	Incorrect Receive Length Errors

Table 5. Specification Changes

Number	Specification Change
1	None for this revision of this specification update.

Table 6. Specification Clarifications

Number	Specification Clarification
1	None for this revision of this specification update.

Table 7. Documentation Changes

Number	Documentation Change
1	None for this revision of this specification update.

§

CPU Errata Details

ICXD1. Wrong Page Access Semantics May be Reported When Intel® SGX ENCLU[EMODPE] Instruction Generates Page Fault (#PF) Exception

Problem: When Intel® Software Guard Extensions (Intel® SGX) extends an Enclave Page Cache (EPC) via the page permissions instruction (ENCLU[EMODPE]) and generates a Page Fault (#PF), even though the page permissions instruction access is a read access to the target page, the Page Fault Error Code (#PF's PFEC) will indicate that the fault occurred on a write (PFEC.W bit will be set) instead.

Implication: This erratum may impact debugging Intel SGX enclaves software. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD2. Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May Cause a #GP

Problem: IA32_THERM_STATUS MSR (19CH) includes Read-Only (RO) fields as well as writable fields. Writing a non-zero value to any of the read-only fields may cause a #GP.

Implication: Due to this erratum, software that reads the IA32_THERM_STATUS MSR, modifies some of the writable fields, and attempts to write the MSR back may cause a #GP.

Workaround: None identified.

Status: For the Steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD3. VMREAD/VMWRITE Instructions May Not Fail When Accessing an Unsupported Field in VMCS

Problem: The execution of VMREAD or VMWRITE instructions should fail if the value of the instruction's register source operand corresponds to an unsupported field in the Virtual Machine Control Structure (VMCS). The correct operation is that the logical processor will set the Zero Flag (ZF), write 0CH into the VM-instruction error field and for VMREAD leave the instruction's destination unmodified. Due to this erratum, the instruction may instead clear the ZF, leave the VM-instruction error field unmodified and for VMREAD modify the contents of its destination.

Implication: Accessing an unsupported field in VMCS may fail to properly report an error. In addition, a VMREAD from an unsupported VMCS field may unexpectedly change its destination. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD4. VERR Instruction Inside VM-Entry May Cause DR6 to Contain Incorrect Values

Problem: Under complex micro-architectural conditions, a VERR instruction that follows a VM-entry with a guest-state area indicating MOV SS blocking (bit 1 in the Interruptibility state) and at least one of B3-B0 bits set (bits [3:0] in the pending debug exception) may lead to incorrect values in DR6

Implication: Due to this erratum, DR6 may contain incorrect values. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD5. 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 MMIO (Memory Mapped IO) or non-coherent agents in the following ways:
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.
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: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD6. VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store

Problem: Execution of the VCVTPS2PH instruction with a memory destination may update the MXCSR exceptions flags (bits [5:0]) if the store to memory causes a fault (for example, #PF) or VM exit. The value written to the MXCSR exceptions flags is what would have been written if there were no fault.

Implication: Software may see exceptions flags set in MXCSR, although the instruction has not successfully completed due to a fault on the memory operation. Intel has not observed this erratum to affect any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD7. 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-GB, 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: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD8. Single Correctable Error Can be Logged Twice if Patrol Scrub Reads Address When Read Transaction is in Flight to Same Address

Problem: When patrol scrubbing reads an address with a correctable ECC error, and at the same time a memory read transaction is in flight to that address, a single correctable error may be logged twice. These errors get logged in the RETRY_RD_ERR_LOG register BDF=(U0, 12, 0) offset 0x22C60 and in the MMIO ECC Correctable Error Counter Registers (22C18h-22C24h/26C18h-26C24h).

Implication: When this erratum occurs, an ECC error from one cache line could result in two correctable errors instead of one. Therefore, incorrectly increasing the overall correctable error count.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD9. Overlap Between APIC And SMRR2 Memory-Mapped Registers Will Not Signal a #GP

Problem: Overlapped APIC and SMRR2 Memory-mapped configurations will not cause a General Protection (#GP) exception when configured

Implication: Due to this erratum, a #GP exception will not be triggered. Intel has not observed this erratum with any commercially available software or platform.

Workaround: None identified. Software should not overlap SMRR2 with APIC registers page.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD10. 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: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD11. 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: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD12. 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: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD13. IA32_RTIT_STATUS.FilterEn Bit Might Reflect a Previous Value

Problem: Under complex micro-architectural conditions, reading the IA32_RTIT_STATUS.FilterEn bit (bit 0 in MSR 571h) after entering or exiting an RTIT region might reflect a previous value instead of the current one.

Implication: Due to this erratum, IA32_RTIT_STATUS.FilterEn bit might reflect a previous value. This erratum has not been seen in any commercially available software.

Workaround: Software should perform an LFENCE instruction prior to reading the IA32_RTIT_STATUS MSR to avoid this issue.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD14. IA32_MC1_STATUS MSR May Not Log Errors When IA32_MC1_CTL MSR is Set to Not Signal Errors

Problem: Under complex micro-architectural conditions, IA32_MC1_STATUS MSR (405H) may not log a poison error when the enable bit (bit 0) in the IA32_MC1_CTL MSR (281H) is cleared.

Implication: Due to this erratum, poison errors might not be logged in the MC1 bank. Intel has not observed this erratum in any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD15. False MC1 Error Reported in The Shadow of a Internal Timer Error

Problem: After a internal timer error has been reported in MC3_STATUS MSR (0x40d) with MCACOD (bits [15:0]) value of 0400H, and MSCOD (bits [31:16]) value of 0080H, under complex micro-architectural conditions, a false error may be reported in MC1_STATUS MSR (0x405) with MCACOD 0x174 or MCACOD 0x124.

Implication: Due to this erratum, a false MCE may be reported in MC1_STATUS MSR. Intel has not observed this erratum in a synthetic test environment.

Workaround: Software should ignore the MC1 error when it appears with an internal timer error.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD16. Debug Exceptions May Be Lost or Misreported When MOV SS or POP SS Instruction is Not Followed by a Write to SP

Problem: If a MOV SS or POP SS instruction generated a debug exception, and is not followed by an explicit write to the Stack Pointer (SP), the processor may fail to deliver the debug exception or, if it does, the DR6 register contents may not correctly reflect the causes of the debug exception.

Implication: Debugging software may fail to operate properly if a debug exception is lost or does not report complete information. Intel has not observed this erratum with any commercially available software.

Workaround: Software should explicitly write to the stack pointer immediately after executing MOV SS or POP SS.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).



ICXD17. CPUID TLB Information is Inaccurate

Problem: CPUID leaf 16 (EAX=16H) subleaf 1 (ECX=01H) TLB information inaccurately reports that the instructions' 1st-level TLB is 8-way and supports both 4K and 2M/4M pages, although it is split into 16 sets of 8 ways for 4K pages and 2 sets of 8 ways for 2M/4M pages.

Implication: Software that uses CPUID instructions 1st-level TLB information may operate incorrectly. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD18. Configuring The PRMRR as Non-WB Might Lead to Incorrect VM-Exit Interruption Error Code

Problem: Under complex micro-architectural conditions, while working with PRMRR configured to be non-WB (Write Back), an Asynchronous Enclave Exit (AEX) caused by a page fault (#PF) that is followed by a VM-exit might lead to an incorrect VM-exit interruption error Code.

Implication: Due to this erratum, the error code captured in the VM-exit interruption error code might be incorrect and not indicate a page fault. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD19. #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: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD20. IERR Not Logged Correctly When Ubox Requested to Signal MSMI

Problem: The Ubox can be programmed to signal a Machine Check System Management Interrupt (MSMI) when an IERR is received from the core. In this case, Ubox will signal both IERR and MSMI and log an error into MCERRLOGGINGREG (Bus: 30; Device: 0; Function: 0; Offset: A8h) but not into IERRLOGGINGREG (Bus: 30; Device: 0; Function: 0; Offset: A4h).

Implication: The source of a core 3-strike timeout IERR cannot be identified while decoding the IERRLOGGINGREG registers in each socket.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD21. CPU Complex PCIe* Surprise Link Down Events May Not be Reported

Problem: When the CPU Complex PCIe* root port encounters conditions that should generate a Surprise Link Down (SLD) event, such as LinkUp = 0, the processor may fail to log or report an SLD event in the ERRUNCSTS register (Bus: 1-4; Device: 1; Function: 0; Offset 104h).

Implication: When this erratum occurs, software that relies upon SLD reporting will not behave as intended.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD22. CPU Complex PCIe RPPIO May Contain Incorrect Tag Value

Problem: The CPU Complex PCIe Root Port Programmable Input Output (RPPIO) Header Log 1 (BDFO) tag field may contain a tag value that does not match that transmitted on the PCIe link.

Implication: When this erratum occurs, it may not be possible to associate transactions on the PCIe link with transaction data logged in RPPIO. Intel has not observed any functional implications due to this erratum.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD23. Uncore MC Bank Registers Corrected Error Count Field May Not Have a Sticky Most Significant Bit

Problem: The corrected error count field in IA32_MC[4..19]_STATUS MSR may not contain a sticky most significant bit, and corrected error count may roll over to 0.

Implication: Due to this erratum, there is no indication that the corrected error count has rolled over.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD24. Poisoned Locked Bus Transactions May Not Allow Warm Reset to Correctly Reset The Processor

Problem: On a system with Intel SGX enabled, if the processor receives poisoned data in response to a locked bus transaction while some cores are in or resuming from a Core C6 state, the resulting warm reset may not correctly reset the processor.

Implication: Due to this erratum, the system may not properly reset without a cold reset.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD25. Spurious CPU Complex PCIe Link Parity Errors May be Logged

Problem: The processor may log a spurious parity error into (Local Data Parity Mismatch Status registers (Bus: 1,5; Device: 2; Function: 0; Bits 15:0) G4LDPMSTS (Offset 420H), G4FRDPMSTS (Offset 424H) and G4SRDPMSTS (Offset 428H)) upon exiting Link L1 power states.

Implication: Due to this erratum, CPU Complex PCIe lanes may report spurious data parity mismatches. Intel has not observed any functional implications for this erratum.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD26. IBIST Receiver Error Overflow Register Field Cannot be Cleared by Software

Problem: The receiver error overflow field of the IBIST Error Recovery and Receive Detection Status registers, IBSTERRRCRVSTS[0-3].RXERRCNTTOVRFLOW [Bus: 1,5; Device: 5; Function:0; Offsets: 624h, 62Eh, 634h, 63Ch; Bit 15] cannot be cleared by writing a "1" to clear this bit. A cold reset is required to clear this bit.

Implication: Due to this erratum, software may encounter inaccurate overflow logging.

Workaround: IBIST tests should use a cold reset between tests to clear the error overflow bit.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD27. MBM May Report Incorrect Bandwidth For Certain Access Strides

Problem: Memory Bandwidth Monitoring (MBM) samples the total memory traffic and upscales the results when reporting bandwidth. MBM may report zero to twice the actual memory bandwidth consumed for workloads that primarily access cache lines sequentially with physical address strides that are a multiples of 4 KB.

Implication: Due to this erratum, MBM may report inaccurate bandwidth for workloads that primarily access cache lines sequentially with physical address strides that are a multiples of 4 KB. Actual memory bandwidth is unaffected by this erratum.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD28. Correctable Errors May Set The Overflow Bit

Problem: Machine Check Architecture (MCA) rules were updated to not set the overflow bit (bit 62) of IA32_MCi_STATUS due to Correctable Errors (CE). The overflow bit of IA32_MC[9-11]_STATUS MSRs (425H, 429H, 42D) may be incorrectly set if the first error is a correctable error and the second error is a non-correctable error.

Implication: Due to this erratum, a corrected error may cause the overflow bit to be set.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD29. Enabled Error May Not be Logged When Other Errors Are Disabled

Problem: During the same cycle, a higher priority Uncorrected (UC) error or Software Recoverable Action Optional (SRAO) error which is disabled, may be logged rather than an enabled lower priority UC or SRAO error (for memory controller machine check banks 12-26).

Implication: Due to this erratum, software may not observe the enabled error. Intel has not observed this erratum in any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD30. Removed

ICXD31. Machine Check Bank Status MSR May Not Set Overflow Bit When Multiple Uncorrectable Errors Occur

Problem: The IA32_MC4_STATUS (Offset: 411h) OVER field (bit 62) may not be set if multiple uncorrectable error types occur during the same cycle or if a single uncorrectable error type occurs over multiple cycles.

Implication: Due to this erratum, identification of multiple uncorrectable errors may not be possible.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD32. FERR Registers Are Not Getting Cleared When CHANERR Register is Being Cleared

Problem: The per-channel FERR_(CHANERR/DMACOUNT/INTDMACOUNT/CHANSTS/DESC_CTRL/CADDR/NADDR) (Bus: 0; Device: 1; Function: 7-0 (CBDMA); Offsets: 180h, 184h, 186h, 188h, 190h, 198h, 1A0h) registers (FERR registers) do not get cleared when the corresponding per-channel CHANERR (Bus: 0; Device: 1; Function: 7-0 (CBDMA); Offset: A8h) register gets cleared.

Implication: Due to this erratum, the value in the stale FERR registers may be incorrect if its read when CHANERR=0.

Workaround: Software should not read the per-channel FERR registers unless the corresponding CHANERR register contains non-zero data.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD33. IOMMU Translation Requests to Interrupt Range May Fail

Problem: The Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Architecture Specification specifies that DMA remapping hardware (IOMMU) should return a successful response with U=1 to a Translation Request with address in the interrupt range (0xFEExxxxx).

Implication: Due to this erratum, the DMA remapping hardware provides a response of Completer Abort. Intel has only observed this erratum in a synthetic testing environment.

Workaround: None identified. Devices should always use Untranslated Request when using address from the MSI register of the device.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD34. Internal Firmware Errors May Not Set Error Enable Bit

Problem: The processor does not set the error enable (EN) bit of the IA32_MC6_STATUS MSR (419h; bit 60) when certain internal firmware errors are detected. IA32_MC6_STATUS MSR field MCACOD (bits 15:0) are correctly set to 406h.

Implication: Software that relies on the EN bit may not operate properly. This type of error is always signaled and will result in system shutdown.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD35. CPU Complex PCIe Surprise Link Down Logging May be Unexpectedly Blocked

Problem: In the absence of a power controller, software is still allowed to set the Power Controller Control (PCC) bit to a value of 1 in SLOTCTL (Bus 1,5; Device 2; Function 0; Offset 58h; bit 10). This action blocks logging of Surprise Link Down (SLD) errors regardless of the state of the Power Controller Present (PCP) bit in SLOTCAP (Bus 1,2,3,4; Device 2; Function 0; Offset 54h bit 1). In the event of a PCIe slot losing power, associated SLD errors should only be blocked if the PCP is set.

Implication: Software that relies upon SLD status may not operate as expected. Intel has not observed this erratum in any commercially available software.

Workaround: Software should check that the SLOTCTL.PCP bit is set before writing the PCC bit in SLOTCTL.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD36. MBA 2.0 May Cause MMIO Traffic to be Throttled

Problem: Memory Bandwidth Allocation (MBA) 2.0 may throttle Memory Mapped IO (MMIO) traffic even though this traffic does not consume memory bandwidth.

Implication: Due to this erratum, write throughput to MMIO, particularly using Write Combining memory types, may be adversely impacted.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD37. Cannot Inject Errors Into PCLS Bits

Problem: The DRAM cache line error injection does not inject into Partial Cache Line Sparing (PCLS) bits. After an error injection, the bits covered by PCLS will not be flipped which may change the expected behavior on subsequent reads of the DRAM cache line.

Implication: Due to this erratum, software injecting errors in cache lines may not see the expected errors on reads to those cache lines. The severity of the error may be incorrect or there may be no error.

Workaround: None identified. Software should avoid injecting errors into the PCLS regions.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD38. CAP Error And ECC Error During ADC/ADDDC Sparing May Not be Corrected

Problem: Under complex microarchitectural conditions, during Adaptive Data Correction/Adaptive Double Device Data Correction (ADC/ADDDC) sparing, a correctable Command/Address Parity (CAP) error and a correctable ECC error occurring simultaneously on the last address of the spare copy may not be properly corrected.

Implication: Due to this erratum, correctable CAP errors and correctable ECC errors may not be properly corrected resulting in an uncorrected error or unpredictable system behavior. This erratum has only been observed in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD39. NSR Field Attribute Does Not Comply With PCIe Base Specification 4.0

Problem: The access type of the No Soft Reset bit (bit 3) of the Power Management Control Status Register (PMCSR) (Bus: 1-4; Device: 5; Function: 0; Offset: 84h) is Read/Write/Locked; however, the PCIe Base Specification version 4.0 specifies this bit to be Read Only.

Implication: Due to this erratum, software that relies on the NSR bit may behave unexpectedly.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD40. Intermittent Correctable Memory Errors May be Observed

Problem: The processor may observe intermittent correctable memory errors with non-homogeneous DRAM configurations.

Implication: Due to this erratum, intermittent correctable memory errors may be observed.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD41. Unmapped Intel® Quick Data Technology DMA Reads May Result in Unpredictable System Behavior

Problem: Intel® Quick Data Technology DMA reads to addresses that are not mapped in the IOMMU or reads to peer agents that result in read completion errors. An Unsupported Request/Completer Abort may result in unpredictable system behavior.

Implication: Due to this erratum, the system may exhibit unpredictable system behavior. Intel has not observed this erratum to impact any commercially available software.

Workaround: None identified. Software should not program the Intel Quick Data Technology controller with any source addresses that may receive completion errors.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD42. Unexpected System Behavior May Occur During INVD Instruction Execution

Problem: During the execution of an INVD instruction, if there is a partial cacheline write from the I/O subsystem in progress, the processor may generate an unexpected machine check exception or other unexpected system behavior.

Implication: When this erratum occurs, the processor may experience unexpected system behavior. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD43. Configuring ADL May Prevent Package C-States Entry

Problem: If the Agent Debug Logic (ADL) is configured after Package C-states are enabled, it is possible for future Package C-states (deeper than Package C0) to be blocked.

Implication: Due to this erratum, Package C-states deeper than Package C0 may not be achieved, leading to higher than anticipated platform power consumption.

Workaround: None Identified. It is possible for software to contain code to enable ADL prior to enabling deeper Package C-states.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD44. CPU Complex PCIe Rx Common Mode Impedance May be Too Low During Reset or Power-Down

Problem: A PCIe receiver may exhibit impedance of approximately 2 kΩ - 3 kΩ at reset and 1 kΩ at power-down compared to expected impedance above 20 kΩ (ZRX-HIGH-IMP-DC-POS).

Implication: The processor does not meet the *PCI Express* Base Specification*, Revision 4.0 receiver impedance greater than 20 kΩ. Intel has not observed any functional impact due to this erratum.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD45. Uncore Semaphore Capability May Not Generate a Semaphore Error Machine Check Exception

Problem: This processor's semaphore capability (LOCALAQUSEMP[1:0] and SYSTEMAQUSEMP[1:0] [Bus 30, Device 0, Function 2, Offsets 0x180, 0x184, 0x18C, 0x188]) does not generate a Semaphore Error machine check exception (MCACOD = 0x0407, MSCOD = 0x800E) if the semaphore's tail pointer passes its head pointer.

Implication: Due to this erratum, the system may exhibit unpredictable system behavior instead of a machine check exception.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD46. WBINVD Delays May Lead to a Machine Check Exception

Problem: Executing a WBINVD instruction during IO traffic and lock instructions may result in an Internal Timer Error Machine Check (IA32_MCI_STATUS.MCACOD=80h; bits [31:16], IA32_MCI_STATUS.MCACOD=400h; bits [15:0]).

Implication: An Internal Timer Error Machine Check may be reported during a WBINVD instruction.

Workaround: It may be possible for BIOS to contain code to work around this issue.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD47. CPU Complex PCIe EB Error May Be Escalated to Receiver Errors

Problem: During PCIe configuration and recovery, logging of Elastic Buffer (EB) errors are not masked as per the PCIe 4.0 Spec Sec 4.2.6 and are escalated to receiver errors in the PCIe Correctable Error Status Register (ERRCORSTS) (Bus: [2,5]; Device: 2; Function: 0; Offset: 110h).

Implication: Due to this erratum, an incorrect receiver error may be reported.

Workaround: It may be possible for a BIOS code change to contain a workaround for this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD48. Unpredictable System Behavior May Occur Due to Memory Read Roundtrip Latency

Problem: During DDR4 read operations with data scrambling enabled, setting any of the values in the RT_rank fields in DDRCRINTFROUNDTRIP0_CH1 register (Offset: 13F30h) RT_RANK(0-3) (bits 27:0) or DDRCRINTFROUNDTRIP1_CH1 register (Offset: 13F34h) RT_RANK(4-7) (bits 27:0) to a value greater or equal to 5Ch may lead to unpredictable system behavior.

Implication: When this erratum occurs, the system may exhibit unpredictable behavior.

Workaround: It may be possible for a BIOS code change to contain a workaround for this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD49. A PECCI Request May Receive an Incorrect Response

Problem: An additional request issued on the PECCI interface prior to the previous request completing may lead to the initial response being used as the response for the second request.

Implication: Due to this erratum, PECCI requests may not receive the correct response.

Workaround: None identified. The device needs to wait for the response before issuing another request.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD50. IFU Internal Parity Error

Problem: Under complex micro architectural conditions, it is possible for the processor to generate a spurious IFU parity error machine check exception (IA32_MCO_STATUS register (MSR 401h) MSCOD = Eh and MCACOD = 5h).

Implication: When this erratum occurs, the processor will report an uncorrectable error. Intel has not observed this erratum to occur in commercially available software.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD51. CHA UCNA Errors May be Incorrectly Controlled by MCI_CTL Enable Bits

Problem: UCNA (Uncorrectable No Action) errors reported in Cache Home Agent (CHA) Machine Check Banks (Banks 9, 10, and 11) MCI_STATUS MSRs (425h, 429h, or 42Dh) may be incorrectly controlled by the associated MCI_CTL MSRs (424h, 428h, or 42Ch).

Implication: Due to this erratum, when MCI_CTL = 0, the UCNA error will be logged but not signaled. When MCI_CTL = FFFFFFFFh, the UCNA error will be logged and signaled, but will incorrectly set MCI_STATUS.EN. (bit 60). Intel has not observed this erratum to affect any commercially available software.

Workaround: It may be possible for a BIOS code change to contain a workaround for this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD52. Mesh to Memory Timeout May Occur When TME is Enabled

Problem: Under complex micro architectural conditions with Total Memory Encryption (TME) enabled, the processor may hang and signal a Internal Timeout Error Machine Check in MCI_STATUS of machine check banks 12 or 16, 20, or 24 (MSRs 431h, 441h, 451h, 461h with MSCOD (bits[31:16]) value of 0009h and MCACOD (bits[15:0]) value of 0400h).

Implication: When this erratum occurs, the system may hang with a Machine Check Exception.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD53. Inaccurate Mesh to Memory Corrected Error Count

Problem: When the Ubox Correctable System Management Interrupt (CSMI) thresholding is disabled in the register (EXRAS_CONFIG.CFGMCACMCIONCORRCOUNTTHR (Bus:30; Device:12; Function:0; Offset 2B4h; bit: 4) for Mesh to Mem Machine Check Banks (Banks 12, 16, 20, or 24) and MCI_CTL2.CMCI_CTL (MSR 28Ch, 290h, 294h, 298h; bit 32) is enabled, the processor will generate two CSMI events for each correctable memory error.

Implication: Due to this erratum, the Corrected Error Count reported in MCI_STATUS MSRs (bits[52:38] in MSR 431h or 441h, 451h, or 461h) may be inaccurate.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD54. CHA Errors May be Reported Incorrectly After a Warm Reset

Problem: If a warm reset occurs after a Cache Home Agent (CHA) error is logged, then the error may be incorrectly reported or not be reported in the CHA Machine Check Banks (Banks 9, 10, and 11) MCI_STATUS MSR's (425h, 429h, or 42Dh) after the warm reset. Further errors that occur post warm reset may also not be reported in the CHA Machine Check Banks.

Implication: Due to this erratum, CHA errors may not be reported or be reported incorrectly after warm reset.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).



ICXD55. Processor May Not Successfully Enter ADR

Problem: For systems configured to use PCode Assisted Asynchronous DRAM Refresh (ADR), if the platform signals an Imminent Power Loss, the processor may not successfully complete the ADR flow.

Implication: When this erratum occurs, ADR does not complete and the processor will increment the LDSC counter. Upon restart, software can determine that the previous shutdown did not complete successfully.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD56. Debug_Has_Occurred Bit May be Asserted

Problem: The processor may report Debug_Has_Occurred in DEBUG_INTERFACE (MSR C80h, bit 31 is set) irrespective of whether debug is enabled or any debug has occurred.

Implication: Due to this erratum, Debug_Has_Occurred bit may be asserted.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD57. IOSFSB Timeout May Lead to MCE

Problem: Under complex microarchitectural conditions, I/O Scalable Fabric Side Band (IOSFSB) timeouts may occur resulting in a fatal Machine Check Exception (MCACOD 402h, MSCOD 5800h/3e00h/5d00h)

Implication: Due to this erratum, a system reset may occur. Intel has not observed this in any commercially available software.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD58. SRIS-Configured CPU Complex PCIe Link May Fail to Train

Problem: The PCIe link may fail to train if the PCIe controller is configured to use SRIS (Separate Reference Clock (Refclk) with Independent Spread Spectrum Clocking) mode.

Implication: Due to this erratum, the PCIe link may fail to train.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD59. Unexpected Rollover in MBM Counters

Problem: When using Intel[®] Resource Director Technology (Intel[®] RDT), unexpected rollover can occur when Memory Bandwidth Monitoring (MBM) counter values are close to the maximum allowed counter value. A rollover is when a MBM counter value read in the n+1th iteration is lower than nth iteration.

Implication: Bandwidth computed from successive MBM readings representing a rollover may not be accurate.

Workaround: None identified. Software should discard the memory bandwidth computed over a rollover interval.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD60. DPC Trigger Status Bit May Not Cleared

Problem: The Downstream Port Containment Trigger Status bit (bit 0) of the DPC Status register (Bus: 0; Device: 9-12, 16-23; Function: 0; Offset: 198h) cannot be cleared while any bits in the error status registers in the Advanced Error Reporting (AER) structure are set.

Implication: Due to this erratum, the DPC Trigger status bit cannot be cleared unless all the AER error status registers are cleared first.

Workaround: Software must clear all the AER error status registers before clearing the DPC Trigger Status bit in the DPC Status register.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD61. Processor May Not Allow Intel® DCI Tool to Operate at 100 MHz

Problem: The processor may incorrectly detect Low Frequency Periodic Signaling (LFPS) when Intel® Direct Connect Interface (Intel® DCI) external debug host [Intel® DCI Out of Band (OOB) adapter] is configured to run at 100 MHz.

Implication: Due to this erratum, the DCI tool may not work properly when configured to run at 100 MHz.

Workaround: Configure Intel DCI tool to operate at a valid frequency other than 100 MHz.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD62. Four Unsuccessful Global Reset Attempts Needed For S5 Entry

Problem: Four instead of three unsuccessful global reset attempts are necessary to set Bit 4 of the Power Management Controller (PMC) Global Reset Causes 1 (GBLRST_CAUSE1) register (BDF 0:31:2, offset PWRMBASE + 1928h, bit [4]).

Implication: Due to this Erratum the system may hang and not enter S5 after three unsuccessful global reset attempts.

Workaround: None.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD63. New_Century Bit of the TCO1_STS Register Should Not be Cleared

Problem: The NEWCENTURY_STS bit in the Power Management Controller (PMC) TCO1_STS register (TCOBASE + 4h, bit 7) cannot be cleared by software writing a 1 back to the bit position. If the bit is written again, the system may hang.

Implication: When this erratum occurs, the system may hang. The failure only occurs when the Year byte rolls over from 99 to 00.

Workaround: None.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD64. xHCI Host Controller Reset May Cause a System Hang

Problem: Within 1 ms of setting the Host Controller Reset bit (HCRST bit1) of the USB Command Register (MBAR + 80h) the xHCI host controller may fail to respond to register accesses.

Implication: The system may hang.

Workaround: 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 (xHCIBAR + 80h) and must add a 120 ms delay in between consecutive xHCI host controller resets.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD65. HUART May Stop Functioning When DMA is Activated

Problem: The High-Speed Universal Asynchronous Receiver/Transmitter (HUART) may stop functioning when the HUART DMA is active ([MEMBA] + Offset 84h, C4h bits = 0) and the HUART is receiving/transmitting information.

Implication: When this erratum occurs, the HUART will stop receiving/transmitting information and may cause a Table Of Request (TOR) Timeout Error Machine Check Exception (Machine Check banks 9, 10, and 11 [MSRs 425h, 429h, and 42Dh] with IA32_MCI_STATUS.MSCOD=0x000C [bits 31:16]).

Workaround: None identified. Software should not activate the HUART DMA.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD66. PCH PCIe Link Partner May Not Enumerate

Problem: A PCIe link partner may fail to detect the PCH PCIe root ports after a warm reset.

Implication: Due to this erratum, a PCIe link partner may not enumerate after a warm reset.

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

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD67. VREFCA Tolerance May Violate The JEDEC JESD79-4 Specification

Problem: For systems with unregistered memory modules, the processor may violate the VREFCA tolerance for the DDR[0,1,2]_CAVREF pins.

Implication: Due to this erratum, the VREFCA tolerance may violate the JEDEC JESD79-4 specification. Intel has not observed any functional implications due to this erratum.

Workaround: None Identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD68. Setting MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT Does Not Prevent The Three-strike Counter From Incrementing

Problem: Setting MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT (bit 11 in MSR 1A4h) does not prevent the three-strike counter from incrementing as documented; instead, it only prevents the signaling of the three-strike event once the counter has expired.

Implication: Due to this erratum, software may be able to see the three-strike logged in the MC3_STATUS (MSR 40Dh, MCACOD = 400h [bits 15:0]) even when MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT is set.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD69. SGX Enclave Attestation And Unseal May Fail

Problem: If a warm reset occurs while a core is in a C6 or deeper power state, Intel SGX enclave attestation and unseal operations will fail until the next cold reset.

Implication: Due to this erratum, an Intel SGX enclave may not be able to unseal data that was sealed by it prior to the warm reset. In addition, the platform may not be able to perform any Intel SGX remote attestation or provisioning until the next cold reset.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD70. Warm Reset Will Set C1E Enable Without Changing TMRT

Problem: When C1E is disabled (C1E_ENABLE field (bit 1) of POWER_CTL1 MSR (01FCh)), the processor adjusts the Thermal Monitor Reference Temperature (TMRT) field (bits [23:16]) of TEMPERATURE_TARGET MSR (01A2h) lower to compensate for higher frequency and power consumption during C1 state. On a warm reset, C1E is re-enabled, but the processor does not re-adjust the TMRT.

Implication: Due to this erratum, the TMRT may be set too low after a warm reset and unexpected throttling may occur.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD71. HWPM Max Ratio May Not be Capped at P1

Problem: The platform may be granted a ratio higher than the guaranteed ratio (P1) when the Energy Efficient Turbo Disable bit (19) in the POWER_CTL1 MSR is set to 1h if a ratio higher than P1 is requested in the Hardware Power Management (HWPM) OOB mode.

Implication: Due to this erratum, Turbo mode disable may not be enforced for the HWPM. Intel has not observed any functional failures due to this erratum.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD72. DRAM Performance May be Reduced Following Microcode Update

Problem: When the DRAM RAPL feature is enabled as indicated by DRAM_POWER_LIMIT.PWR_LIM_CTRL_EN=1 (MSR 618h, bit 15), performing a microcode update will result in maximum memory bandwidth throttling.

Implication: Due to this erratum, the DRAM will consistently run in a reduced performance mode with decreased memory bandwidth and increased memory latency.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD73. PCIe Bandwidth Reduction When Intel VT-d is Enabled

Problem: Under certain PCIe traffic conditions when Intel Virtualization Technology (Intel VT) for Directed I/O (Intel VT-d) is enabled and configured for 4 KB page sizes, a reduction in PCIe bandwidth may occur.

Implication: When this erratum occurs, PCIe bandwidth reduction may occur.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD74. PC6 Entry is Not Prevented With PECCI Pkg C-state Entry Control

Problem: Setting Pkg C-state Entry Control to 1 via the PECCI interface (WrPkgConfig Index 54) does not correctly restrict entry into Package-C6 state (PC6).

Implication: When this erratum occurs, the processor may continue to briefly enter PC6 and continue to increment the PC6 residency counter (MSR 3F9h).

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD75. A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an Intel AVX Gather or REP MOVSB Instruction

Problem: Under complex micro-architectural conditions, when both data poison and data breakpoint events happen on an Intel[®] Advanced Vector Extensions (Intel[®] AVX) Gather or REP MOVSB instruction, one of the events may not be signaled.



Implication: Due to this erratum, either a data breakpoint or a poison data event may not be signaled.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD76. Incorrect MCACOD For L2 Prefetch MCE

Problem: Under complex micro-architectural conditions, an L2 prefetch MCE that should be reported with MCACOD 165h in IA32_MC3_STATUS MSR (MSR 40dh, bits [15:0]) may be reported with an MCACOD of 101h.

Implication: Due to this erratum, the reported MCACOD for this MCE may be incorrect.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD77. Call Instruction Wrapping Around The 32-Bit Address Boundary May Return to Incorrect Address

Problem: In 32-bit mode, a call instruction wrapping around the 32-bit address should save a return address near the bottom of the address space (low address) around address zero. Under complex micro-architectural conditions, a return instruction following such a call may return to the next sequential address instead (high address).

Implication: Due to this erratum, in 32-bit mode a return following a call instruction that wraps around the 32-bit address boundary may return to the next sequential IP without wrapping around the address, possibly resulting in a #PF. Intel has not observed this behavior on any commercially available software.

Workaround: None identified. Software should not place call instructions in addresses that wrap around the 32-bit address space in 32-bit mode.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD78. FEATURE_TUNING_1 May be Reset to Factory Defaults

Problem: When a microcode update is performed, the processor may inadvertently reset the SpecI2MEn field in FEATURE_TUNING_1 register (MSR 06dh, bit [30]).

Implication: Due to this erratum, platform-specific performance tuning settings in FEATURE_TUNING_1 may be reset to factory defaults.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD79. Accesses to CHA Configuration Space Beyond the CHA Logical Limit May Fail

Problem: The processor may have more Caching and Home Agent (CHA) physically implemented than are logically available in the processor. The CHA configuration registers are located in the PCIe configuration space associated with the CHA bus, device, and function, with the first CHA being located at Bus 31, Device 0, Function 0 and also Bus 31, Device 10, Function 0. There are two functions in the PCI CFG space for each CHA. Accesses to the CHA configuration space may not return valid results for BDFs beyond the number of logical CHAs supported in the processor as enumerated in CAPID6_CFG (Bus 31, Device 30, Function 3, Offset 9Ch, bits [31:0]) and CAPID7_CFG (Bus 31, Device 30, Function 3, Offset A0h, bits [31:0]).

Implication: Due to the erratum, accesses to CHA configuration spaces, including Device ID, for CHAs beyond the CHA logical limit may not return valid results.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD80. Monitor May Not be Triggered

Problem: Under complex microarchitectural conditions, a monitor that is armed with the MWAIT instruction may not be triggered, leading to a processor hang.

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

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD81. Incorrect CPU Complex PCIe RCB Advertisement

Problem: The root port, which is visible to the endpoint, advertises Read Completion Boundary (RCB) of 64, but the root complex behaves according to the rules that apply when RCB=128.

Implication: When performing transactions that may be treated as malformed, such as Address Translation Services (ATS) operations, some PCIe endpoints may flag certain TLPs as malformed, due to violating PCIe specification rules regarding RCB. This causes fatal errors on the platform.

Workaround: The BIOS should clear bit 33 (RCB128) of ITCCTRL23 register (Bus:0-4; Device:0; Function:0; Offset:548h) to configure the root complex to use an RCB of 64.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD82. Unsupported Request Error Will be Logged in IEH

Problem: When the PCH Satellite Integrated Error Handler (IEH) error reporting capability is enabled and a configuration read is performed to a non-existing or disabled device, a non-fatal Unsupported Request (UR) error will be logged.

Implication: If a configuration read occurs to a non-existing or disabled device, an error will be logged.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).



ICXD83. PCH PCIe Root Ports May Operate at 8.0 GT/s Rather Than 5.0 GT/s

Problem: Under certain PCIe topologies, the PCIe controllers may not honor the target link speed specified in the Link Control 2 register (Bus: 0; Device: 9-12, 20-23; Function: 0; Offset: 70h; bits [3:0]).

Implication: Due to this erratum, certain PCIe ports may operate at 8.0 GT/s, rather than 5.0 GT/s.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD84. PCH PCIe Devices With LUR May Not Link

Problem: Following a warm reset, PCH PCIe root ports may not link with devices that support Link-Up Retry (LUR).

Implication: Due to this erratum, PCH PCIe devices that support LUR may not link after a warm reset.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD85. After a Warm Reset, The SMBus Host Controller May Incorrectly Set The CPE Bit

Problem: After warm reset, the SMBus host controller may incorrectly set the CSR Parity Error Status (CPE) field (bit 0) of the ERRSTS register (SMTBAR offset 18h).

Implication: Due to this erratum, software may receive unexpected interrupts from the SMBus host controller.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD86. PCH PCIe Root Port Cluster 0 Does Not Correctly Process MCTP Messages

Problem: The PCH PCIe Cluster 0 (Root Ports [RPs] 0, 1, 2, and 3) does not process Management Component Transport Protocol (MCTP) messages correctly.

Implication: Due to this erratum, uncorrectable errors may be reported if the MCTP messages are sent to or received from Cluster 0.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD87. PCH PCIe 3.0 Link May Observe Link Errors After Speed Change

Problem: If software initiates a PCH PCIe root port speed change from 8 GT/s to 2 GT/s or 4 GT/s and then back to 8 GT/s, the PCIe link may become unstable.

Implication: Due to this erratum, PCIe link correctable and uncorrectable link errors may occur.

Workaround: It may be possible for a BIOS code change to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD88. PEFI Transactions to Certain Devices Will Time Out

Problem: PEFI transactions to the following devices will return a Command Response Timeout error (code 0x80):

Serial Peripheral Interface (SPI) Controller (Bus: 0; Device: 31; Function: 5)

Enhanced SPI Master Controller (eSPI-MC) (Bus: 0; Device: 31; Function: 0)

Low Pin Count (LPC) Controller (Bus: 0; Device: 31; Function: 5)

Intel Management Engine (Intel ME) (Bus: 0; Device: 24; Function: 0)

Implication: Due to this erratum, PEFI transactions to the above PCH devices will time out.

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD89. System May Hang With USB 3.2 Ports Enabled

Problem: If the processor PCH is configured with USB 3.2 enabled on at least one of the Ports 1-3 without USB 3.2 Port 0 enabled and without any PCIe Ports enable, the processor may hang.

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

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD90. PCH PCIe Link Speed May Be Limited to 5.0 GT/s

Problem: If the PCH PCIe port fails to link to 8.0 GT/s link speed, subsequent attempts to train the link may be limited to 5.0 GT/s link speed.

Implication: PCH PCIe links may operate at a lower link speed than expected.

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD91. Branch Predictor May Produce Incorrect Instruction Pointer

Problem: Under complex microarchitectural conditions, the branch predictor may produce an incorrect instruction pointer leading to unpredictable system behavior.

Implication: Due to this erratum, the system may exhibit unpredictable behavior.

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD92. DRNG May Erroneously Return Poisoned Data

Problem: Under complex microarchitectural conditions, when executing workloads using the RDRAND instruction, the Digital Random Number Generator (DRNG) may erroneously return poisoned data.

Implication: Due to this erratum, a fatal error may occur if poisoned data had previously been observed in the system. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).



ICXD93. System May Hang During ADR C2F

Problem: Copy to Flash (C2F) operation may not complete during an Asynchronous DRAM Refresh (ADR) event.

Implication: Due to this erratum, the system may hang during the ADR C2F operation.

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD94. IERR Maybe Seen During Warm Reset After OS Path Load of Microcode Revision IDs of 0x1000230 or 0x1000260

Problem: On a system with microcode revision ID 0x01000211 or earlier loaded, performing an OS Patch Load of microcode revision IDs 0x01000230 or 0x01000260 will result in an IERR upon the next warm reset.

Implication: Due to this erratum, the processor may report an IERR during a warm RESET.

Workaround: It may be possible for a BIOS to workaround this erratum.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD95. Reading TCO_RLD Register Sets SMBus Host Status In Use Status Bit

Problem: When software reads the TCO_RLD (TCOBASE Offset 0), hardware sets the In Use Status (Host Status; SMBMAR Offset 0 or SBA Offset 0; bit 6) in the SMBus controller D31:F4.

Implication: Due to this erratum, the SMBus controller may appear to be in use, causing software to not be able to use the SMBus controller.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

ICXD96. A Write to The TSC_Deadline MSR May Cause an Unexpected Timer Interrupt

Problem: Under complex micro-architectural conditions, writing a non-zero value to the Time-Stamp Counter (TSC) Deadline counter, IA32-TSC_DEADLINE MSR (6E0h), may cause timer interrupt following the write.

Implication: Due to this erratum, a unexpected timer interrupt may be signaled.

Workaround: None identified.

Status: For the steppings affected, refer to the [CPU Errata Summary Table](#).

§

LAN/Ethernet Errata Details

LAN1. Inverted FCS Does Not Cause Increment of CRC Error Count

Problem: An incorrect value of the Ethernet FCS on a received Ethernet packet is detected by the Network Interface and Scheduler (NIS). Normally, the packet is dropped and the Port CRC Error Count (GLPRT_CRCERRS) counter is incremented. In the specific case where the FCS is the bitwise inverse of the correct FCS value, the packet is dropped but the counter is not incremented.

Implication: Receiving the bitwise inverse of the FCS is almost always due to an intentional modification of the FCS; there is an extremely low probability of random bit errors causing an inverted FCS field. Therefore, the lack of counting these FCS errors affects testing environments, but should not affect detection of bit errors during operation.

Workaround: When testing the CRC error reporting, do not use the inverted FCS value.

Status: For the steppings affected, see the [LAN/Ethernet Errata Summary Table](#).

LAN2. Incorrect Calculation of Inner UDP Checksum on Rx Packet

Problem: If a tunneled UDP packet is received in which there is L2 padding in the inner packet, the offloaded inner UDP checksum calculation is incorrect. A checksum error is reported on a packet that has the correct checksum.

Implication: When receiving an indication of a checksum error, the Intel OS drivers report to the OS that the checksum was not offloaded, and the OS stack performs the checksum validation. Therefore, there is a no effect on the packet reception other than a slight increase in CPU usage. The driver reports the error count to the OS, possibly resulting in a high error rate which is reported by the OS.

Workaround: The driver should not report the checksum error count to the OS. This is implemented in the NIC 1.2 LEK release.

Status: For the steppings affected, see the [LAN/Ethernet Errata Summary Table](#).

LAN3. Incorrect Receive Length Errors

Problem: If an Ethernet packet is received with a packet length larger than 64 bytes and the Length/Type field contains a value that is smaller than the number of bytes remaining in the packet, an Rx length error is detected. This type of packet could be generated if a short packet is padded to 64 bytes and then a VLAN tag is inserted without removing any padding bytes. When an Rx length error is detected, the packet is considered a bad packet and is discarded unless storing bad packets is enabled. Also, the GLPRT_RLEC counter is incremented.

Implication: Packets with the structure defined above are dropped. These packets are often local diagnostic packets that might be discarded by the software anyway. The driver reports the error count to the OS, possibly resulting in a high error rate that is reported by the OS.

Workaround: The driver should not report the Rx length error count to the OS. This is implemented in the NIC 1.2 LEK release.

Status: For the steppings affected, see the [LAN/Ethernet Errata Summary Table](#).

§



Specification Changes

There are no specification changes in this specification update revision.

§

Specification Clarifications

There are no specification changes in this specification update revision.

§



Documentation Changes

There are no specification changes in this specification update revision.

§

Preface	5
Affected Documents/Related Documents	5
Nomenclature	6
Identification Information	7
Component Identification Using Programming Interface.....	7
Component Marking Information	8
Summary Tables of Changes	9
Codes Used in Summary Tables.....	9
CPU Errata Details	14
LAN/Ethernet Errata Details	35
Specification Changes	36
Specification Clarifications	37
Documentation Changes	38

Working page only. Do not distribute.