3rd Gen Intel® Xeon® Scalable Processors, Codename Ice Lake

Specification Update

July 2023

Reference Number: 637780-020US

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.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

See backup for workloads and configurations. Results may vary.

For workloads and configurations visit www.intel.com/PerformanceIndex. Results may vary.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

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 Optane, and Xeon are trademarks of Intel Corporation or its subsidiaries.

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

Contents

Revision History

Date	Revision	Description
July 2023	020us	Added errata ICX152. to ICX166. Removed ICX52. duplicate of ICX38. Removed ICX132. duplicate of ICX122. Removed ICX134. duplicate of ICX112.
April 2023	019us	Added errata ICX150.to ICX151.
March 2023	018us	Added erratum ICX149.
February 2023	017us	Added errata ICX147. to ICX148. Updated erratum ICX142.
November 2022	016us	Added errata ICX145. to ICX146.
October 2022	015us	Added erratum ICX144.
September 2022	014us	Added errata ICX142. to ICX143.
August 2022	013us	Added errata ICX138. to ICX141. Updated the Table 3, "List of Microcode Revisions".
July 2022	012us	Added errata ICX131. to ICX137.
June 2022	011us	Added errata ICX122. to ICX130. Removed erratum ICX121. Updated erratum ICX3.
May 2022	010us	Added errata ICX116. through ICX121. Updated Nomenclature
March 2022	009us	Added errata ICX109. through ICX115.
February 2022	008us	Added errata ICX108. Updated ICX101.
January 2022	007us	Added errata ICX102. through ICX107. Updated ICX3., ICX36.
October 2021	006us	Added errata ICX97. through ICX101.
September 2021	005us	Added errata ICX90. through ICX96.
August 2021	004us	Added errata ICX82. through ICX89.
July 2021	003us	Added errata ICX74. through ICX81. Updated ICX39. and ICX58.
June 2021	002us	Added errata ICX47. through ICX73.
May 2021	001us	Initial Release

Preface

This document is an update to the specifications contained in the following Related Documents table. 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.

Related Documents

Document Title	Document Number/ Location
Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture	253665[1]
Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference, A-M	253666[1]
Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference, N-Z	253667[1]
Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide, Part 1	253668[1]
Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3B: System Programming Guide, Part 2	253669[1]
Advanced Configuration Power Interface (ACPI) Specifications	www.acpi.info[2]

1. Documents are available publicly at https://www.intel.com/content/www/us/en/design/resource-design-center.html.
2. Document available at www.uefi.org.

Nomenclature

Errata are design defects or errors. These may cause the 3rd Gen Intel® Xeon® Scalable Processors' 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, all notes associated with each S-Spec number.

Qualification Detail Form (QDF) Number is 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 forth.

Identification Information

Component Identification via Programming Interface

The 3rd Gen Intel® Xeon® Scalable Processors stepping can be identified by the following register contents:

Reserved	Extended Family[1]	Extended Model[2]	Reserved	Processor Type[3]	Family Code[4]	Model Number[5]	Stepping ID[6]
31:28	27:20	19:16	15:13	12:12	11:8	7:4	3:0
	00000000b			0b			Varies per stepping

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® Pro, Pentium® 4, Intel® Core™ processor families, Intel® Core™ ix families, and Intel® Xeon® processor families.
2. The Extended Model, bits [19:16] in conjunction with the Model Number, specified in bits [7:4], is used to identify the model of the processor within the processor's family.
3. The Processor Type, specified in bit [12], indicates whether the processor is an original OEM processor, an Intel OverDrive processor, or a dual processor (capable of being used in a dual processor system).
4. The Family Code corresponds to bits [11:8] of the Extended Data Register (EDX) after RESET, bits [11:8] of the Extended Accumulator Register (EAX) 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.
5. 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.
6. 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.

When EAX is set to a value of "1," the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number, and Stepping ID together referred to as the processor signature value, in the EAX register. Note that after reset, the EDX processor will report the processor signature value in both the EDX and the EAX registers.

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

Table 1. Component Identification via Registers

Physical Chip	Stepping	Segment Wayness	CPUID	CAPID0(Segment)			CAPID0(Wayness)		CAPID4(Chop)	
				B:31, D:30, F:3, O:84h					B:31, D:30, F:3, O:94H	
				5	4	3	1	0	7	6
XCC	D-2	Server, 1S	606A6	1	1	1	0	1	1	1
	D-2	Server, 2S	606A6	1	1	1	1	0	1	1
HCC	M-1	Server, 1S	606A6	1	1	1	0	1	1	0
	M-1	Server, 2S	606A6	1	1	1	1	0	1	0

Component Marking Information

The 3rd Gen Intel® Xeon® Scalable Processors can be identified by the following register markings.

Figure 1. Processor Preliminary Top Side Marking (Example)

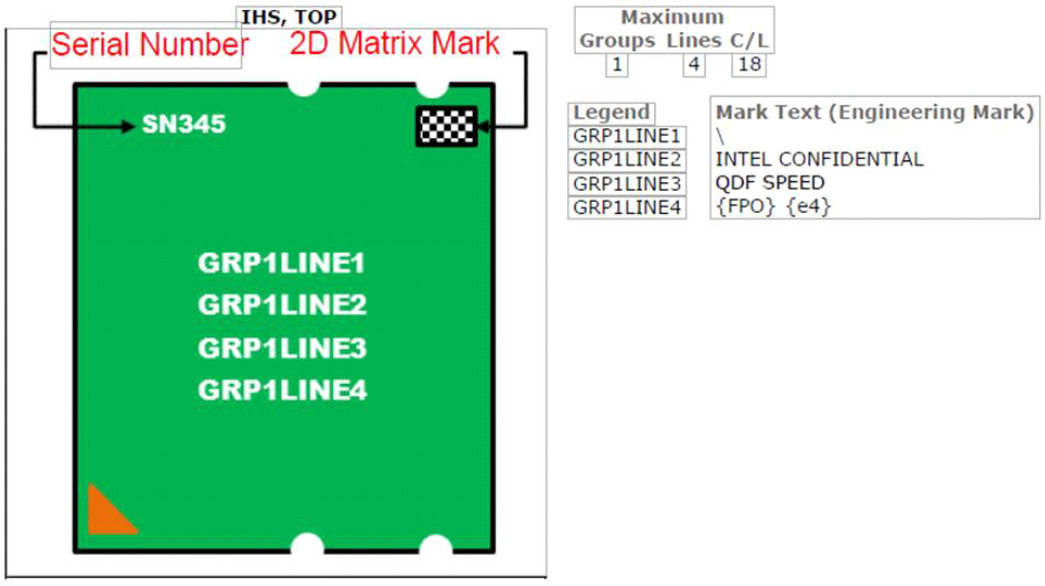

For 3rd Gen Intel® Xeon® Scalable Processors SKUs, see https://ark.intel.com/content/www/us/en/ark/products/series/204098/3rd-generation-intel-xeon-scalable-processors.html.

Summary Tables of Changes

The following tables indicate the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the 3rd Gen Intel® Xeon® Scalable Processors product. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations:

Codes Used in Summary Tables

Stepping	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.
No Fix	There are no plans to fix this erratum.

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX1.	No Fix	No Fix	Memory Errors in a VLS Region on a Certain Device May Not be Properly Corrected
ICX2.	No Fix	No Fix	Wrong Page Access Semantics May be Reported When Intel® SGX ENCLU[EMODPE] Instruction Generates Page Fault (#PF) Exception
ICX3.	No Fix	No Fix	Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May Cause a #GP
ICX4.	No Fix	No Fix	VMREAD/VMWRITE Instructions May Not Fail When Accessing an Unsupported Field in VMCS
ICX5.	No Fix	No Fix	VERR Instruction Inside VM-Entry May Cause DR6 to Contain Incorrect Values
ICX6.	No Fix	No Fix	Vector Masked Store Instructions May Cause Write Back of Cache Line Where Bytes Are Masked
ICX7.	No Fix	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store
ICX8.	No Fix	No Fix	SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
ICX9.	No Fix	No Fix	Single Correctable Error Can be Logged Twice if Patrol Scrub Reads Address When Read Transaction is in Flight to Same Address
ICX10.	No Fix	No Fix	Processor May Hang if Warm Reset Triggers During BIOS Initialization
ICX11.	No Fix	No Fix	Placing Posted-Interrupt Descriptors Within The PRMRR May Result in a Processor Hang
ICX12.	No Fix	No Fix	Placing Page Table Information in The APIC-Access Page May Lead to Unexpected Page Faults While Performing Enclave Accesses
ICX13.	No Fix	No Fix	Performance Monitoring Load Latency Events May be Inaccurate For Gather Instructions
ICX14.	No Fix	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
ICX15.	No Fix	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
ICX16.	No Fix	No Fix	Overlap Between APIC And SMRR2 Memory-Mapped Registers Will Not Signal a #GP
ICX17.	No Fix	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set
ICX18.	No Fix	No Fix	Intel SGX Enclave Accesses to The APIC-Access Page May Cause APIC-Access VM Exits
ICX19.	No Fix	No Fix	Intel® PT TIP or FUP Packets May be Dropped Without OVF Packet
ICX20.	No Fix	No Fix	Intel PT CBR Packet May be Delayed or Dropped

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX21.	No Fix	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
ICX22.	No Fix	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
ICX23.	No Fix	No Fix	In eMCA2 Mode, When The Retirement Watchdog Timeout Occurs CATERR# May be Asserted
ICX24.	No Fix	No Fix	IA32_RTIT_STATUS.FilterEn Bit Might Reflect a Previous Value
ICX25.	No Fix	No Fix	IA32_MC1_STATUS MSR May Not Log Errors When IA32_MC1_CTL MSR is Set to Not Signal Errors
ICX26.	No Fix	No Fix	False MC1 Error Reported in The Shadow of a Internal Timer Error
ICX27.	No Fix	No Fix	Debug Exceptions May Be Lost or Misreported When MOV SS or POP SS Instruction is Not Followed by a Write to SP
ICX28.	No Fix	No Fix	CPUID TLB Information is Inaccurate
ICX29.	No Fix	No Fix	CPUID L2 Cache Information May be Inaccurate
ICX30.	No Fix	No Fix	Configuring The PRMRR as Non-WB Might Lead to Incorrect VM-Exit Interruption Error Code
ICX31.	No Fix	No Fix	An Invalid Algorithm in The CRYPTO_ALG Field of The MKTME_KEY_PROGRAM_STRUCT Will Not Cause an INVALID_ENC_ALG Failure as Expected
ICX32.	No Fix	No Fix	A Spurious APIC Timer Interrupt May Occur After Timed MWAIT
ICX33.	No Fix	No Fix	#GP on Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code
ICX34.	No Fix	No Fix	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
ICX35.	No Fix	No Fix	IERR Not Logged Correctly When Ubox Requested to Signal MSMI
ICX36.	No Fix	No Fix	PCIe* Surprise Link Down Events May Not be Reported
ICX37.	No Fix	No Fix	PCIe RPPIO May Contain Incorrect Tag Value
ICX38.	No Fix	No Fix	Uncore MC Bank Registers Corrected Error Count Field May Not Have a Sticky Most Significant Bit
ICX39.	No Fix	No Fix	UPI Correctable Error May be Observed After Cold Resets And TCRH Events
ICX40.	No Fix	No Fix	Poisoned Locked Bus Transactions May Not Allow Warm Reset to Correctly Reset The Processor
ICX41.	No Fix	No Fix	TOR Timeout During WBINVD May Cause System Hang

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX42.	No Fix	No Fix	Intel® PT VMentry Indication Depends on The Incorrect VMCS Control Field
ICX43.	No Fix	No Fix	Intel PT Trace May Drop Second Byte of CYC Packet
ICX44.	No Fix	No Fix	Intel PT TIP.PGD May Not Have Target IP Payload
ICX45.	No Fix	No Fix	Intel PT PSB+ May be Lost
ICX46.	No Fix	No Fix	Intel Processor Trace PSB+ Packets May Contain Unexpected Packets
ICX47.	No Fix	No Fix	Spurious PCIe Link Parity Errors May be Logged
ICX48.	No Fix	No Fix	IBIST Receiver Error Overflow Register Field Cannot be Cleared by Software
ICX49.	No Fix	No Fix	MBM May Report Incorrect Bandwidth For Certain Access Strides
ICX50.	No Fix	No Fix	Correctable Errors May Set The Overflow Bit
ICX51.	No Fix	No Fix	Enabled Error May Not be Logged When Other Errors Are Disabled
ICX52.	No Fix	No Fix	Removed
ICX53.	No Fix	No Fix	Machine Check Bank Status MSR May Not Set Overflow Bit When Multiple Uncorrectable Errors Occur
ICX54.	No Fix	No Fix	FERR Registers Are Not Getting Cleared When CHANERR Register is Being Cleared
ICX55.	No Fix	No Fix	IOMMU Translation Requests to Interrupt Range May Fail
ICX56.	No Fix	No Fix	Internal Firmware Errors May Not Set Error Enable Bit
ICX57.	No Fix	No Fix	UPI PH_PLS.SRstRcvdP Value May be Incorrect Following a UPI PHY Reset
ICX58.	No Fix	No Fix	Spurious Write Data Parity Errors May be Logged
ICX59.	No Fix	No Fix	PKG_MIN_PWR May be Incorrect
ICX60.	No Fix	No Fix	PN_POWER_OF_SKU May be Incorrect
ICX61.	No Fix	No Fix	PCIe Surprise Link Down Logging May be Unexpectedly Blocked
ICX62.	No Fix	No Fix	MBA May Incorrectly Throttle Threads When Hyperthreading is Enabled

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX63.	No Fix	No Fix	MBA 2.0 May Cause MMIO Traffic to be Throttled
ICX64.	No Fix	No Fix	Cannot Inject Errors Into PCLS Bits
ICX65.	No Fix	No Fix	SNC2 And Hemisphere Mode do Not Work Correctly on Processors With 38 CHAs
ICX66.	No Fix	No Fix	CAP Error And ECC Error During ADC/ADDDC Sparing May Not be Corrected
ICX67.	No Fix	No Fix	Processor IERR Condition Followed by Warm Reset May Encounter Subsequent Fatal Error
ICX68.	No Fix	No Fix	NSR Field Attribute Does Not Comply With PCIe Base Specification 4.0
ICX69.	No Fix	No Fix	Intermittent Correctable Memory Errors May be Observed
ICX70.	No Fix	No Fix	Intel® UPI CRC Errors May be Detected During Power Management Transitions
ICX71.	No Fix	No Fix	Unmapped QDT DMA Reads May Result in Unpredictable System Behavior
ICX72.	No Fix	No Fix	Unexpected System Behavior May Occur During INVD Instruction Execution
ICX73.	No Fix	No Fix	Incorrect Intel® AVX2 And Intel® AVX-512 High Priority Frequencies May Be Reported
ICX74.	No Fix	No Fix	Configuring ADL May Prevent Package C-States Entry
ICX75.	No Fix	No Fix	DDR-T Interrupts During Warm Reset May Lead to a System Hang
ICX76.	No Fix	No Fix	Intermittent Intel UPI Test Failures May Be Observed During BSCAN
ICX77.	No Fix	No Fix	PCIe Rx Common Mode Impedance May be Too Low During Reset or Power-Down
ICX78.	No Fix	No Fix	Uncore Semaphore Capability May Not Generate a Semaphore Error Machine Check Exception
ICX79.	No Fix	No Fix	Uncore FIVR Fault May be Logged Incorrectly
ICX80.	No Fix	No Fix	Intel® SST-CP May Not be Dynamically Configurable
ICX81.	No Fix	No Fix	CHA Data Parity Error or Writeback LLC Miss May be Reported
ICX82.	No Fix	No Fix	Multiple SGX_Doorbell_Errors on Ubox Response Mismatch
ICX83.	No Fix	No Fix	Intel SGX Doorbell MCA Error May Incorrectly be Masked

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX84.	No Fix	No Fix	LMCE May Hang With Intel SGX Disabled
ICX85.	No Fix	No Fix	PECI Writes to 64-Bit Registers May Fail
ICX86.	No Fix	No Fix	Monitor Snoop May be Missed When LLC Prefetching is Enabled
ICX87.	No Fix	No Fix	Mismatch Between UboxErrMisc and MCI_STATUS Registers Error Logs
ICX88.	No Fix	No Fix	System Address Logged for WDB Parity Errors May be Incorrect
ICX89.	No Fix	No Fix	Multiple CHA Errors May be Reported Incorrectly
ICX90.	No Fix	No Fix	WBINVD Delays May Lead to a Machine Check Exception
ICX91.	No Fix	No Fix	Mesh to Memory Transactions Timeout During Memory Stress Test
ICX92.	No Fix	No Fix	Certain Processor Units May Not Reach Intel® AVX-512 P1 All-cores Base Frequency
ICX93.	No Fix	No Fix	PCIe EB Error May Be Escalated to Receiver Errors
ICX94.	No Fix	No Fix	Disabling All Processor Cores on First Physical Row May Lead to an MCE
ICX95.	No Fix	No Fix	Unpredictable System Behavior May Occur Due to Memory Read Roundtrip Latency
ICX96.	No Fix	No Fix	A PECI Request May Receive an Incorrect Response
ICX97.	No Fix	No Fix	IFU Internal Parity Error
ICX98.	No Fix	No Fix	CHA UCNA Errors May be Incorrectly Controlled by MCi_CTL Enable Bits
ICX99.	No Fix	No Fix	Mesh to Memory Timeout May Occur When TME is Enabled
ICX100.	No Fix	No Fix	Inaccurate Mesh to Memory Corrected Error Count
ICX101.	No Fix	No Fix	CHA Errors May be Reported Incorrectly After a Warm Reset
ICX102.	No Fix	No Fix	Processor May Not Successfully Enter ADR
ICX103.	No Fix	No Fix	Debug_Has_Occurred Bit May be Asserted
ICX104.	No Fix	No Fix	PCIe Completion Timeout Error May Occur

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX105.	No Fix	No Fix	IOSFSB Timeout May Lead to MCE
ICX106.	No Fix	No Fix	CHA/IDI Parity Error Machine Check Exceptions May Occur
ICX107.	No Fix	No Fix	SRIS-Configured PCIe Link May Fail to Train
ICX108.	No Fix	No Fix	Unexpected Rollover in MBM Counters
ICX109.	No Fix	No Fix	Incorrect MCACOD For L2 MCE
ICX110.	No Fix	No Fix	Poison Data Reported Instead of a CS Limit Violation
ICX111.	No Fix	No Fix	Intel PT Trace May Contain Incorrect Data When Configured With Single Range Output Larger Than 4KB
ICX112.	No Fix	No Fix	ADDDC Reverse Sparing May Lead to Incorrect Data
ICX113.	No Fix	No Fix	PCIE Completion Timeouts May Occur Under Memory Stress
ICX114.	No Fix	No Fix	Unnecessary PkgC6 Exit Events
ICX115.	No Fix	No Fix	IAA May Fail to Properly Decode Data With a Large Header
ICX116.	No Fix	No Fix	System May Experience an Internal Timeout Error When Directing Intel® PT to a Small, Uncacheable, Single-Range Output Buffer
ICX117.	No Fix	No Fix	Setting MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT Does Not Prevent The Three-strike Counter From Incrementing
ICX118.	No Fix	No Fix	SGX Enclave Attestation And Unseal May Fail
ICX119.	No Fix	No Fix	ADR May Fail to Complete
ICX120.	No Fix	No Fix	Processor May Not Report Accurate Thermal Data Via PECI
ICX121.	No Fix	No Fix	Removed
ICX122.	No Fix	No Fix	System May Hang After MCA_DISPATCHER_RUN_BUSY_TIMEOUT MCA Occurs
ICX123.	No Fix	No Fix	Processor May Signal Spurious #GP Fault
ICX124.	No Fix	No Fix	RTM Abort Status May be Incorrect For INT1/INT3 Instructions
ICX125.	No Fix	No Fix	Warm Reset Will Set C1E Enable Without Changing TMRT
ICX126.	No Fix	No Fix	HWPM Max Ratio May Not be Capped at P1
ICX127.	No Fix	No Fix	DRAM Performance May be Reduced Following Microcode Update
ICX128.	No Fix	No Fix	PCIe Bandwidth Reduction When Intel® VT-d is Enabled
ICX129.	No Fix	No Fix	WRMSR to a Few Core MSRs Might be Overwritten
ICX130.	No Fix	No Fix	PMON May Overcount M2Mem Ingress Events
ICX131.	No Fix	No Fix	PC6 Entry Is Not Prevented With PECI Pkg C-state Entry Control

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX132.	No Fix	No Fix	Removed
ICX133.	No Fix	No Fix	MCA_DISPATCHER_RUN_BUSY_TIMEOUT May be Seen During Warm Reset
ICX134.	No Fix	No Fix	Removed
ICX135.	No Fix	No Fix	A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an AVX Gather or REP MOVS Instruction
ICX136.	No Fix	No Fix	Incorrect MCACOD For L2 Prefetch MCE
ICX137.	No Fix	No Fix	Call Instruction Wrapping Around The 32-bit Address Boundary May Return to Incorrect Address
ICX138.	No Fix	No Fix	PUNIT Dispatcher May Hang
ICX139.	No Fix	No Fix	FEATURE_TUNING_1 May be Reset to Factory Defaults
ICX140.	No Fix	No Fix	SST-TF May Fail to Report an Error if Turbo is Disabled
ICX141.	No Fix	No Fix	BSP May Not be The Lowest Numbered APIC ID
ICX142.	No Fix	No Fix	Accesses to CHA Configuration Space Beyond the CHA Logical Limit May Fail
ICX143.	No Fix	No Fix	Monitor May Not be Triggered
ICX144.	No Fix	No Fix	Incorrect PCIe RCB Advertisement
ICX145.	No Fix	No Fix	PCIe Link Width May Degrade After a Hot Reset or Link Disable
ICX146.	No Fix	No Fix	Intel SGX TCB Recovery Boot Flow May Fail
ICX147.	No Fix	No Fix	Mapped Out PMEM Module May Result in Unpredictable System Behavior
ICX148.	No Fix	No Fix	VCCIN Current Draw May Exceed Specification
ICX149.	Np Fix	No Fix	DRNG May Erroneously Return Poisoned Data
ICX150.	Np Fix	No Fix	Branch Predictor May Produce Incorrect Instruction Pointer
ICX151.	Np Fix	No Fix	TOR_TIMEOUT MCE May Occur Under Sustained High-bandwidth PCIe Traffic
ICX152.	No Fix	No Fix	MASKMOV* Instruction To a Physical Memory Location Mapped By Two Linear Addresses of Different Page Sizes May Result In Unpredictable System Behavior
ICX153.	No Fix	No Fix	x87 FDP Value May be Saved Incorrectly
ICX154.	No Fix	No Fix	System May Hang When Bus-Lock Detection Is Enabled And EPT Resides in Uncacheable Memory
ICX155.	No Fix	No Fix	THERMTRIP Will be Asserted When Conditions to Assert MEMTRIP Are Met
ICX156.	No Fix	No Fix	MCA Registers May Report an Incorrect Number of Corrected Errors

Errata Summary Table

Erratum ID	Processor Line / Steppings		Title
	HCC	XCC	
	M-1	D-2	
ICX157.	No Fix	No Fix	Incorrect Memory Transaction Type May be Logged When Data Poisoning is Disabled
ICX158.	No Fix	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
ICX159.	No Fix	No Fix	IA32_MC6_CTL2 CMCI Enable Bit is Not Implemented Correctly
ICX160.	No Fix	No Fix	CPUID L2 Cache Information May Be Inaccurate
ICX161.	No Fix	No Fix	VM Entry That Clears TraceEn May Generate a FUP
ICX162.	No Fix	No Fix	DPC Trigger Status Bit May Not Cleared
ICX163.	No Fix	No Fix	Boot Guard ACM Authentication Failure May Result In Endless Reset Cycles
ICX164.	No Fix	No Fix	Writing To Apic Timer Initial Count Arms The Counter
ICX165.	No Fix	No Fix	VMExit Might Be Missed On INT3/INTO Preceded by MOV/POP SS
ICX166.	No Fix	No Fix	Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP

Specification Changes

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

Specification Clarifications

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

Documentation Changes

Number	Documentation Changes
1	None for this revision of the specification update.

Errata Details

ICX1. Memory Errors in a VLS Region on a Certain Device May Not be Properly Corrected

Problem: Under complex micro-architectural conditions, when Adaptive Data Correction (ADC) or Adaptive Double Device Data Correction (ADDDC) is enabled, and the system has spared out DRAM device 0, 1, 3, 4, 5, 8, 13, 15, or 16, and the system is in Virtual Lockstep (VLS) mode, then if a limited subset of multi-bit errors are detected on primary DRAM device 16 in the VLS region, those errors may not be properly corrected.

Implication: The system may experience unpredictable system behavior. Intel has only observed this behavior under synthetic testing conditions.

Workaround: None identified.

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

ICX2. 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 Summary Tables of Changes.

ICX3. 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: It may be possible for BIOS to contain a workaround for this erratum.

Status: For the Steppings affected, refer to the Summary Tables of Changes.

ICX4. 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 Summary Tables of Changes.

ICX5. 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 Summary Tables of Changes.

ICX6. 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 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: It is recommended not to map MMIO range as WB. If WB is used for MMIO range, OS or VMM should not map such MMIO page adjacent to a regular WB page (adjacent on the linear address space, before or after the I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

Workaround: None identified.

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

ICX7. 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 Summary Tables of Changes.

ICX8. 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) may 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 Summary Tables of Changes.

ICX9. 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 Summary Tables of Changes.

ICX10. Processor May Hang if Warm Reset Triggers During BIOS Initialization

Problem: Under complex micro-architectural conditions, when the processor receives a warm reset during BIOS initialization, the processor may hang with a machine check error reported in IA32_MCi_STATUS, with MCACOD (bits [15:0]) value of 0400H, and MSCOD (bits [31:16]) value of 0080H.

Implication: Due to this erratum, the processor may hang. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

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

ICX11. Placing Posted-Interrupt Descriptors Within The PRMRR May Result in a Processor Hang

Problem: Posted-interrupt processing is a virtualization feature for interrupts which requires configuring addresses in the posted-interrupt descriptor fields in the Virtual Machine Control Structure (VMCS). Configuring posted-interrupt descriptors addresses that are within the Processor Reserved Memory Range Register (PRMRR), defined by MSR 1F4H and MSR 1F5H) may result in a logical processor hang.

Implication: This erratum may result in a processor hang. Intel has not observed this erratum with any commercially available software.

Workaround: Virtual Machine Monitor (VMM) software should not use addresses within the PRMRR for posted-interrupt descriptors.

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

ICX12. Placing Page Table Information in The APIC-Access Page May Lead to Unexpected Page Faults While Performing Enclave Accesses

Problem: Guest-physical access using a guest-physical address that translates to an address on the APIC-access page (as identified by the APIC-access address field in the VMCS) should cause an APIC-access VM exit. This includes page table information accesses done as part of page translation (page walks). Due to this erratum placing page table information in the APIC-access page may result in a page fault instead of VM exit when the page translation is done as part of an enclave access.

Implication: Software that places page table information in the APIC access page may get page faults on executing enclave accesses, instead of exiting to the Virtual-Machine Monitor (VMM). Intel has not observed this erratum with any commercially available software.

Workaround: Software should not place page table information in the APIC access page.

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

ICX13. Performance Monitoring Load Latency Events May be Inaccurate For Gather Instructions

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

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

Workaround: None identified.

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

ICX14. Performance Monitoring Counters May Undercount When Using CPL Filtering

Problem: Performance Monitoring counters configured to count only OS or only USR events (by setting only one of bits 16 or 17 in IA32_PERFEVTSELx) may undercount for a short cycle period of typically less than 100 processor clock cycles after the processor transitions to a new CPL. Events affected may include those counting CPL transitions (by additionally setting the edge-detect bit 18 in IA32_PERFEVTSELx).

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

Workaround: None identified.

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

ICX15. PEBS Eventing IP Field May be Incorrect After Not-Taken Branch

Problem: When a Precise-Event-Based-Sampling (PEBS) record is logged immediately after a not-taken conditional branch (Jcc instruction), the Eventing IP field should contain the address of the first byte of the Jcc instruction. Due to this erratum, it may instead contain the address of the instruction preceding the Jcc instruction.

Implication: Performance monitoring software using PEBS may incorrectly attribute PEBS events that occur on a Jcc to the preceding instruction.

Workaround: None identified.

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

ICX16. 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 Summary Tables of Changes.

ICX17. 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 Summary Tables of Changes.

ICX18. Intel SGX Enclave Accesses to The APIC-Access Page May Cause APIC-Access VM Exits

Problem: In VMX non-root operation, Intel® Speed Select Technology - Core Power (Intel® SST-CP) Intel® SGX enclave accesses to the APIC-access page may cause APIC-access VM exits instead of page faults.

Implication: A Virtual-Machine Monitor (VMM) may receive a VM exit due to an access that should have caused a page fault, which would be handled by the guest OS.

Workaround: A VMM avoids this erratum if it does not map any part of the Enclave Page Cache (EPC) to the guest's APIC-access address; an operating system avoids this erratum if it does not attempt indirect enclave accesses to the APIC.

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

ICX19. Intel® PT TIP or FUP Packets May be Dropped Without OVF Packet

Problem: The Intel® Processor Trace (Intel® PT) Overflow (OVF) packet may not be generated when only Target IP Packets (TIPs) and/or Flow Update Packets (FUPs) are lost due to internal buffer overflow.

Implication: A decoder error will result from the missing FUP and/or TIP packets.

Workaround: None identified.

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

ICX20. Intel PT CBR Packet May be Delayed or Dropped

Problem: Due to a complex set of microarchitectural conditions, the Intel® PT Core:Bus Ratio (CBR) packet generated on a frequency change may be dropped, without an Overflow (OVF) packet, or may be inserted into the trace late, after other packets (including possibly another CBR) that were generated after the frequency change completed.

Implication: An Intel® PT decoder may report an incorrect core:bus ratio to a portion of the trace, which may result in an incorrect wall clock time calculation.

Workaround: None identified.

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

ICX21. 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 Summary Tables of Changes.

ICX22. 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 Summary Tables of Changes.

ICX23. 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 steppings affected, refer to the Summary Tables of Changes.

ICX24. 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 may reflect a previous value instead of the current one.

Implication: Due to this erratum, IA32_RTIT_STATUS.FilterEn bit may 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 Summary Tables of Changes.

ICX25. 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 may 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 Summary Tables of Changes.

ICX26. 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 Summary Tables of Changes.

ICX27. 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 Summary Tables of Changes.

ICX28. 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 Summary Tables of Changes.

ICX29. CPUID L2 Cache Information May be Inaccurate

Problem: CPUID extended function 80000006H (EAX=80000006H) inaccurately reports information about the L2 cache in ECX. The function reports that the L2 cache size is 256K divided into 8 ways, while the actual L2 size and structure should be inferred from reading CPUID leaf 04H sub-leaf 02H.

Implication: Software that uses CPUID extended leaf 80000006H L2 cache 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 Summary Tables of Changes.

ICX30. 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 may 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 may 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 Summary Tables of Changes.

ICX31. An Invalid Algorithm in The CRYPTO_ALG Field of The MKTME_KEY_PROGRAM_STRUCT Will Not Cause an INVALID_ENC_ALG Failure as Expected

Problem: The supported encryption algorithms are specified in TME_CAPABILITY MSR (981h). Using PCONFIG instruction to config an invalid algorithm in CRYPTO_ALGS field of the MKTME_KEY_PROGRAM_STRUCT (bits [23:8] in the KEYID_CTRL) will not cause an INVALID_ENC_ALG failure as expected.

Implication: Due to this erratum, the INVALID_CRYPTO_ALG bit (bit 4) reported by PCONFIG instruction may be incorrect and memory encryption may not commence. Intel has not observed this erratum in any commercially available software.

Workaround: Software should only write to supported CRYPTO_ALG bits as enumerated in TME_CAPABILITY MSR (981h).

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

ICX32. A Spurious APIC Timer Interrupt May Occur After Timed MWAIT

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

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

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

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

ICX33. #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 Summary Tables of Changes.

ICX34. x87 FPU Exception (#MF) May be Signaled Earlier Than Expected

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

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

Workaround: None identified.

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

ICX35. 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 Summary Tables of Changes.

ICX36. PCIe* Surprise Link Down Events May Not be Reported

Problem: When the processor's PCI Express* (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: It may be possible for BIOS to contain a workaround for this erratum.

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

ICX37. PCIe RPPIO May Contain Incorrect Tag Value

Problem: The 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 Summary Tables of Changes.

ICX38. 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..28]_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 Summary Tables of Changes.

ICX39. UPI Correctable Error May be Observed After Cold Resets And TCRH Events

Problem: Intel® Ultra Path Interconnect (Intel® UPI) correctable errors (IA32_MCi_STATUS.MSCOD [bits 31:16] 0x22/0x23/0x30, IA32_MCi_STATUS.MCACOD=0x0E0F [bits 15:0]) may be observed after cold resets and Train Cold Run Hot (TCRH) events.

Implication: Due to this erratum, Intel® UPI Correctable Errors may be logged and lead to a link width change.

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

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

ICX40. 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 Summary Tables of Changes.

ICX41. TOR Timeout During WBINVD May Cause System Hang

Problem: Under complex microarchitectural conditions, a TOR Timeout Error Machine Check Exception (Machine Check banks 9, 10, and 11 [MSRs 0x425, 0x429, and 0x42D] with IA32_MCi_STATUS.MSCOD=0x000C [bits 31:16]) may occur during a WBINVD instruction.

Implication: Due to this erratum, a system hang may occur.

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

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

ICX42. Intel® PT VMentry Indication Depends on The Incorrect VMCS Control Field

Problem: An Intel® PT Paging Information Packet (PIP), which includes indication of entry into non-root operation, will be generated on VMentry as long as the "Conceal VMX in Intel PT" field (bit 19) in Secondary Execution Control register (IA32_VMX_PROCBASED_CTLS2, MSR 048BH) is clear. This diverges from expected behavior, since this PIP should instead be generated only with a zero value of the "Conceal VMX entries from Intel PT" field (Bit 17) in the Entry Control register (IA32_VMX_ENTRY_CTLS MSR 0484H)

Implication: An Intel® PT trace may incorrectly expose entry to non-root operation.

Workaround: A VMM (virtual machine monitor) should always set both the "Conceal VMX entries from Intel® PT" field in the Entry Control register and the "Conceal VMX in Intel® PT" in the Secondary Execution Control register to the same value.

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

ICX43. Intel PT Trace May Drop Second Byte of CYC Packet

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

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

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

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

ICX44. Intel PT TIP.PGD May Not Have Target IP Payload

Problem: When Intel® PT is enabled and a direct unconditional branch clears IA32_RTIT_STATUS.FilterEn (MSR 571H, bit 0), due to this erratum, the resulting TIP.PGD (Target IP Packet, Packet Generation Disable) may not have an IP payload with the target IP.

Implication: It may not be possible to tell which instruction in the flow caused the TIP.PGD using only the information in trace packets when this erratum occurs.

Workaround: The Intel® PT trace decoder can compare direct unconditional branch targets in the source with the FilterEn address range(s) to determine which branch cleared FilterEn.

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

ICX45. Intel PT PSB+ May be Lost

Problem: Intel® PT generates a Packet Stream Boundary+ (PSB+) set of packets periodically, based on the number of trace bytes written out. If the threshold for a PSB+ is reached while Intel® PT is being disabled by clearing IA32_RTIT_CTL.TraceEn[0] (MSR 0570H) either during a VM-exit or after generating fewer than 8 bytes of trace since TraceEn was last set, that PSB+ may be lost.

Implication: An Intel® PT decoder that is scanning for a PSB+ at which to begin decode may have to skip over more trace output bytes before finding one.

Workaround: Software processing the trace at runtime can detect that a PSB+ was dropped by checking that IA32_RTIT_STATUS.PacketByteCnt[48:32] (MSR 0571H) has recently crossed the PSB threshold, while scanning the trace to check that the expected PSB+ was not inserted. When a dropped PSB+ is detected, software can force a PSB+ to be inserted the next time Intel® PT is enabled by clearing PacketByteCnt.70H) to 1, as an internal buffer overflow that loses a CYC packet will generate an OVF.

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

ICX46. Intel Processor Trace PSB+ Packets May Contain Unexpected Packets

Problem: Some Intel® PT packets should be issued only between TIP.PGE (Target IP Packet.Packet Generation Enable) and TIP.PGD (Target IP Packet.Packet Generation Disable) packets. Due to this erratum, when a TIP.PGE packet is generated it may be preceded by a PSB+ (Packet Stream Boundary) that incorrectly includes FUP (Flow Update Packet) and MODE.Exec packets.

Implication: Due to this erratum, FUP and MODE.Exec may be generated unexpectedly.

Workaround: Decoders should ignore FUP and MODE.Exec packets that are not between TIP.PGE and TIP.PGD packets.

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

ICX47. Spurious PCIe Link Parity Errors May be Logged

Problem: The processor may log a spurious parity error into (Local Data Parity Mismatch Status registers (Bus: 1,2,3,4; 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, 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 Summary Tables of Changes.

ICX48. 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].RXERRCNTOVRFLOW [Bus: 4-1; 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 Summary Tables of Changes.

ICX49. 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 Summary Tables of Changes.

ICX50. 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 Summary Tables of Changes.

ICX51. 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 Summary Tables of Changes.

ICX52. Removed

ICX53. 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 Summary Tables of Changes.

ICX54. 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 Summary Tables of Changes.

ICX55. 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 Summary Tables of Changes.

ICX56. 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 Summary Tables of Changes.

ICX57. UPI PH_PLS.SRstRcvdP Value May be Incorrect Following a UPI PHY Reset

Problem: The value in the Intel® UPI register S_RST_RCVD_P (KTIREUT_PH_PLS Bus 30; Device 2; Function 2; Offset 160h; Bit 14) is reset and cleared instead of taking the value in S_RST_RCVD (KTI_REUT_PH_CLS Bus 30; Device 2; Function 2; Offset 164h; Bit 14) following an Intel® UPI PHY reset.

Implication: Due to this erratum, software relying on KTIREUT_PH_PLS following an Intel® UPI reset may contain inaccurate logging details.

Workaround: None identified.

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

ICX58. Spurious Write Data Parity Errors May be Logged

Problem: In 2S systems, a spurious Write Data Parity Error (MSCOD = 0x002, MCACOD=0x0405) may be logged in the memory controller machine check bank status register (Bank 13, 14, 17, 18, 21, 22, 25, 26 and MSRs 0x435, 0x439, 0x449, 0x455, 0x459, 0x465 and 0x569). The same errors will be logged in the associated shadow registers at MMIO offset 23440h from the MEM[0-7]_BAR registers (Bus 30, Device 0, Function 1, Offsets D8h, DCh, E0h, E4h, E8h, ECh, F0h, and F4h).

Implication: When poison is enabled (MCG_CONTAIN.POiSON_ENABLE = 1 (MSR 178h, bit 0), a spurious Write Data Parity Error may be logged and will be signaled as a UCNA error, but there will be no poison data in memory. Software that relies on UCNA error reporting may take unnecessary actions. When poison is not enabled, the spurious Write Data Parity Error may result in a fatal machine check exception.

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

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

ICX59. PKG_MIN_PWR May be Incorrect

Problem: Certain processors may report an incorrect value in the PKG_MIN_PWR field (Bits 30:16) of the PACKAGE_POWER_SKU_CFG register (Bus: 31; Device 30; Function 0; Offset 80h).

Implication: Software that utilizes PKG_MIN_PWR to budget processor or platform power may not behave as intended.

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

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

ICX60. PN_POWER_OF_SKU May be Incorrect

Problem: Certain processors may report an incorrect value in the PN_POWER_OF_SKU field (Bits 14:0) of the POWER_LIMIT_MISC_INFO_CFG register (Bus: 31; Device 30; Function 5; Offset E0h).

Implication: Software that utilizes PN_POWER_OF_SKU to budget processor or platform power may not behave as intended.

Workaround: None identified.

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

ICX61. 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 1,2,3,4; 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 Summary Tables of Changes.

ICX62. MBA May Incorrectly Throttle Threads When Hyperthreading is Enabled

Problem: When one logical processor is idle, the Memory Bandwidth Allocation (MBA) feature may select an incorrect MBA throttling value to apply to the core. An idle logical processor may behave as though the CLOS field in its associated IA32_PQR_ASSOC MSR (0xC8F) is set to the CLOS of the last active thread. When this occurs, the MBA throttling value associated with CLOS of the last active thread on the unused logical processor may be incorrectly applied to the physically co-located active logical processor.

Implication: An idle logical processor is interpreted to have the CLOS that the last active thread sets in its IA32_PQR_ASSOC MSR, which affects the calculation for the actual throttling applied to the physical core. When this erratum occurs, the MBA throttling value associated with a given core may be incorrect. This erratum does not affect use cases when both threads on a physical core are assigned the same CLOS.

Workaround: None identified.

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

ICX63. MBA 2.0 May Cause MMIO Traffic to be Throttled

Problem: Memory Bandwidth Allocation (MBA) 2.0 may throttle 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 Summary Tables of Changes.

ICX64. 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 Summary Tables of Changes.

ICX65. SNC2 And Hemisphere Mode do Not Work Correctly on Processors With 38 CHAs

Problem: Sub-NUMA Clustering 2 (SNC2) and Hemisphere Mode (HM) do not work correctly on processors with 38 Caching and Home Agents (CHA).

Implication: On processors with 38 CHAs, if SNC2 or HM is enabled, the processor may exhibit unpredictable system behavior.

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

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

ICX66. 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 Summary Tables of Changes.

ICX67. Processor IERR Condition Followed by Warm Reset May Encounter Subsequent Fatal Error

Problem: Certain fatal error conditions that result in IERR assertion may fail to properly recover following a warm reset.

Implication: When this erratum occurs, the processor may require a global reset to boot the system.

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

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

ICX68. 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 Summary Tables of Changes.

ICX69. 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 Summary Tables of Changes.

ICX70. Intel® UPI CRC Errors May be Detected During Power Management Transitions

Problem: Correctable CRC errors may be detected in machine check banks MC5_STATUS, MC7_STATUS, or MC8_STATUS (MSRs 415h, 41Dh, or 421h) with MSCOD = 0x30 or 0x23 and MCACOD = 0x0E0F when the Intel® UPI) link is entering or exiting L0p or L1 power management states.

Implication: Due to this erratum, correctable CRC errors may be detected during power state transitions.

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

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

ICX71. Unmapped QDT 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 Summary Tables of Changes.

ICX72. 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 Summary Tables of Changes.

ICX73. Incorrect Intel® AVX2 And Intel® AVX-512 High Priority Frequencies May Be Reported

Problem: The processor may incorrectly report High Priority Intel® Advanced Vector Extensions 2 (Intel® AVX2) and Intel® Advanced Vector Extensions 512 (Intel® AVX-512) frequencies higher than the corresponding SSE frequency.

Implication: Due to this erratum, software that reads High Priority Intel® AVX2 and Intel® AVX-512 frequencies may report an incorrect value. The processor's Intel® AVX2 and Intel® AVX-512 frequencies will not exceed the corresponding SSE frequency.

Workaround: None identified.

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

ICX74. 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 Summary Tables of Changes.

ICX75. DDR-T Interrupts During Warm Reset May Lead to a System Hang

Problem: If a DDR-T interrupt occurs during warm reset it may not get serviced, which may lead to a system hang.

Implication: Due to this erratum, a system hang may occur.

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

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

ICX76. Intermittent Intel UPI Test Failures May Be Observed During BSCAN

Problem: Sending a Boundary Scan (BSCAN) external test pattern from an Intel® UPI transmitter may encounter intermittent failures where the pattern is not correctly received by the Intel® UPI receiver on the remote socket.

Implication: A BSCAN pattern from an Intel® UPI transmitter may not be observed correctly at an Intel® UPI receiver.

Workaround:

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

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

ICX77. 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 Summary Tables of Changes.

ICX78. 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 0x184, 0x190, 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 Summary Tables of Changes.

ICX79. Uncore FIVR Fault May be Logged Incorrectly

Problem: If an Uncore FIVR fault occurs prior to the first BIOS Microcode Update (MCU) load, it may be reported in machine check bank IA32_MC4_STATUS (MSR 411h) as an MCA_DISPATCHER_RUN_BUSY_TIMEOUT (MSCOD=0xB0, MCACOD=0x402) instead of an MCA_FIVR_PD_HARDERR (MSCOD=0x56, MCACOD=0x402).

Implication: Due to this erratum, an incorrect FIVR fault error may be logged.

Workaround: None identified.

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

ICX80. Intel® SST-CP May Not be Dynamically Configurable

Problem: The processor supports both the WRITE_PCU_MISC_CONFIG and WRITE_PM_CONFIG mailbox commands. If the WRITE_PCU_MISC_CONFIG mailbox command is used to enable Intel® SST-CP, then the WRITE_PM_CONFIG mailbox command cannot be used to disable it.

Implication: Due to this erratum, Intel® SST-CP may not be dynamically disabled as expected.

Workaround: Software should use only the WRITE_PM_CONFIG mailbox command to dynamically configure Intel® SST-CP.

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

ICX81. CHA Data Parity Error or Writeback LLC Miss May be Reported

Problem: A processor with a reported junction temperature above 75 °C and an uncore frequency > 2 GHz, a fatal Machine Check Exception (MCE) may be observed on Cache Home Agent (CHA) Machine Check Banks (Banks 9, 10, and 11) MCi_STATUS MSRs (425h, 429h, or 42Dh) due to CHA data parity errors (MCACOD = 405h and MSCOD = Ah) or a Writeback LLC miss (MCACOD = 1146h and MSCOD = Bh).

Implication: Due to this erratum, a fatal MCE may be observed.

Workaround:

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

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

ICX82. Multiple SGX_Doorbell_Errors on Ubox Response Mismatch

Problem: In the event of a mismatch between the Intel® UPI LT_Doorbell Response completion and the SGX_Secure_En configuration bit in Ubox, SGX_Doorbell_Errors may overflow the NCEVENTS_CR_UBOX_MCI_STATUS (merged MCA bank 6) register and signal a redundant Machine Check Abort (MCA) with MSCOD 801Ch and MCACOD of 0407h to the cores and PUNIT.

Implication: Due to this erratum, a redundant MCA may be signaled.

Workaround: None identified.

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

ICX83. Intel SGX Doorbell MCA Error May Incorrectly be Masked

Problem: The register NCEVENTS_CR_GL_ERR_CFG1_CFG.MaskSGXDoorbellError (Bus:30; Device:0; Function:0; Offset: 158h) bit [18] masks the assertion of the global fatal status for an Intel® SGX doorbell error. However, the RTL is also incorrectly masking the Machine Check Assertion (MCA) signaling to the cores and punit.

Implication: Due to this erratum, no errors are logged in the NCEVENTS_CR_UBOXERRSTS2_CFG (Bus:30; Device:0; Function:0; Offset A0h) register. The associated machine check is not reported, and may lead to a system hang.

Workaround: None identified. Software should not use NCEVENTS_CR_GL_ERR_CFG1_CFG.MaskSGXDoorbellError.

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

ICX84. LMCE May Hang With Intel SGX Disabled

Problem: When Intel® SGX is disabled, if a thread receives a Local Machine Check Exception (LMCE), 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 steppings affected, refer to the Summary Tables of Changes.

ICX85. PECI Writes to 64-Bit Registers May Fail

Problem: PECI register writes to the upper 32 bits of certain 64-bit PECI registers (SCF_MEM [Bus: 30; Device: 12; Function: 0; Offset: 238h], a_PCIw_DMI_CBDMA_PCIE [Bus: 0; Device: 4; Function: 0; Offset: 608h, Port0_DMI_x4PCIe_DMI_CBDMA_PCIE [Bus: 0; Device: 3; Function: 0; Offset: 608h) will not successfully complete, returning a PECI Unsuccessful Write response (90h completion code).

Implication: Due to this erratum, PECI writes to a 64-bit register may fail.

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

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

ICX86. Monitor Snoop May be Missed When LLC Prefetching is Enabled

Problem: Under complex microarchitectural conditions, a monitor snoop may be missed, when Last Level Cache (LLC) prefetching is enabled and lines prefetched to the LLC are remotely homed and monitored at the same time.

Implication: Due to this erratum, the thread may not respond to the monitor snoop.

Workaround: None identified.

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

ICX87. Mismatch Between UboxErrMisc and MCI_STATUS Registers Error Logs

Problem: The logging in UboxErr Misc Registers (UboxErrMisc_CFG(Bus: 30; Device: 0; Function: 0; Offset ECh), UboxErrMisc2_CFG (Bus: 30; Device: 0; Function: 0; Offset E8h) and UboxErrMisc3_CFG (Bus: 30; Device: 0; Function: 0; Offset F4h)) and IA32_MC6_STATUS (Offset 419h) may be related to different events when a poisoned MMIO transaction and a poisoned Interrupt transaction occur concurrently due to differences in priority logic for logging into the MCI_STATUS register and logging into the UboxErrMisc registers.

Implication: Due to this erratum, the UboxErrMisc registers may show information for a different transaction than the one logged in MCI_STATUS.

Workaround: None identified.

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

ICX88. System Address Logged for WDB Parity Errors May be Incorrect

Problem: When target XOR enable bit [21] MCMTR.DDR_HALF_XOR_MODE_ENABLE (MEM0_BAR; Offset: 20EF8h) or channel XOR enable bit [20] MCMTR.DDR_XOR_MODE_ENABLE (MEM0_BAR; Offset: 20EF8h) is set or clock gating disable bit [28] [DDRT_CLK_GATING.DIS_REVADDR_LOG_CLKGATING (MEM0_BAR; Offset 20514h)] is not set, the IMC0_POISON_SOURCE (MEM0_BAR; Offset 20E80h) register may log Write Data Buffer/Byte Enable (WDB/BE) Register File parity errors with an incorrect system address.

Implication: Due to this erratum, the IMC0_POISON_SOURCE register may log the incorrect system address when WDB_PARITY_ERR = 1 in IMC0_POISON_SOURCE.

Workaround: None identified. Software may avoid this erratum by disabling clock gating (DDRT_CLK_GATING.DIS_REVADDR_LOG_CLKGATING = 1) and disabling target XOR (MCMTR.DDR_HALF_XOR_MODE_ENABLE = 0) and disabling channel XOR (MCMTR.DDR_XOR_MODE_ENABLE= 0).

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

ICX89. Multiple CHA Errors May be Reported Incorrectly

Problem: If a software write to one of the Cache Home Agent (CHA) Machine Check Banks (Banks 9, 10, and 11) MCi_STATUS MSRs (425h, 429h, or 42Dh) happens on the same cycle when internal logic sets the Overflow bit, the hardware error will be logged in the CHA, but the internal severity of the error may be that of the software write, possibly leading to incorrect severity being reported in the associated machine check bank.

Implication: Due to this erratum, multiple CHA errors may not be reported correctly.

Workaround: None identified. Intel has only observed this in a synthetic test environment.

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

ICX90. 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 Summary Tables of Changes.

ICX91. Mesh to Memory Transactions Timeout During Memory Stress Test

Problem: When the processor is utilizing Intel® Optane™ Persistent Memory (Intel® Optane™ PMem) in Memory Mode, under complex microarchitectural conditions when running memory stress testing a soft hang may occur.

Implication: Due to this erratum, a system hang may occur with a Timeout Error Machine Check reported in MCi_STATUS of machine check banks 12, 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).

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

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

ICX92. Certain Processor Units May Not Reach Intel® AVX-512 P1 All-cores Base Frequency

Problem: Certain processor units may not reach specified base all-cores active Intel® AVX-512 frequency when running benchmarks with high utilization of Intel® AVX-512 instructions.

Implication: Due to this erratum, impacted processors may run at a lower base all-cores active Intel® AVX-512 P1 core frequency.

Workaround:

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

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

ICX93. 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 *PCIe* 4.0 Specification*, Section 4.2.6 and are escalated to receiver errors in the PCIe* Correctable Error Status Register (ERRCORSTS) (Bus: [4-1]; Device: 2; Function: 0; Offset: 110h).

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

Workaround:

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 Summary Tables of Changes.

ICX94. Disabling All Processor Cores on First Physical Row May Lead to an MCE

Problem: When the cores on the first physical row (cores 0-3 and 19-22 for XCC family processors; cores 0-2 and 13-15 for HCC family processors) are disabled in BIOS, a Machine Check Exception (MCE) failure may be logged during reset in Machine Check Banks (Banks 9, 10, and 11) MCi_STATUS MSR's (425h, 429h, or 42Dh).

Implication: Due to this erratum, a fatal MCE may occur and the system may fail to boot.

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

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

ICX95. 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 Summary Tables of Changes.

ICX96. A PECI Request May Receive an Incorrect Response

Problem: An additional request issued on the PECI 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, PECI 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 Summary Tables of Changes.

ICX97. 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_MC0_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 Summary Tables of Changes.

ICX98. 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 MSR's (425h, 429h, or 42Dh) may be incorrectly controlled by the associated MCi_CTL MSR's (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 =FFFFFFFFFh, 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: None identified.

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX99. 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, 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 BIOS to contain a workaround for this erratum.

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX100. 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, 441h, 451h, or 461h) may be inaccurate.

Workaround: None identified.

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

ICX101. 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: SIt may be possible for BIOS to contain a workaround for this erratum

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

ICX102. 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 BIOS to contain a workaround for this erratum.

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

ICX103. 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 Summary Tables of Changes.

ICX104. PCIe Completion Timeout Error May Occur

Problem: When the processor is utilizing Intel® Optane™ PMem in Memory Mode; under complex microarchitectural conditions when running memory stress testing a PCIe* Completion Timeout Error may occur.

Implication: Due to this erratum, a PCIe* timeout may occur which may result in a system hang.

Workaround: It may be possible to adjust the PCIe* Completion Timeout limit. It may be possible for BIOS to contain a workaround for this erratum.

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

ICX105. IOSFSB Timeout May Lead to MCE

Problem: Under complex microarchitectural conditions, IO 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 BIOS to contain a workaround for this erratum

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

ICX106. CHA/IDI Parity Error Machine Check Exceptions May Occur

Problem: A fatal Machine Check Exception (MCE) may be observed on CHA Machine Check Banks (Banks 9, 10, and 11) MCi_STATUS MSRs (425h, 429h, or 42Dh) due to CHA Parity Errors (MCACOD = 405h and MSCOD = Ah or MCACOD=1146h and MSCOD = Ah) or an IDI Parity Error (MCACOD = 405h and MSCOD = 40h).

Implication: Due to this erratum, the system may experience a fatal CHA Parity or IDI Parity MCE.

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

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

ICX107. SRIS-Configured 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 Summary Tables of Changes.

ICX108. 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 Summary Tables of Changes.

ICX109. Incorrect MCACOD For L2 MCE

Problem: Under complex micro-architectural conditions, an L2 poison MCE that should be reported with MCACOD 189h 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 Summary Tables of Changes.

ICX110. Poison Data Reported Instead of a CS Limit Violation

Problem: Under complex micro-architectural conditions, in case of poisoned data on an address that violates the CS (code segment) limit, a poison MCE may be signaled and logged in IA32_MC0_STATUS MSR (MSR 401H, MCACOD 150h) instead of CS limit violation.

Implication: Due to the erratum, the processor may signal an MCE, rather than a higher-priority CS limit violation.

Workaround: None identified.

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

ICX111. Intel PT Trace May Contain Incorrect Data When Configured With Single Range Output Larger Than 4KB

Problem: Under complex micro-architectural conditions, when using Intel® PT with single range output larger than 4KB, disabling Intel® PT and then enabling Intel® PT using the TraceEn bit in IA32_RTIT_CTL MSR (MSR 570h, bit 0) may cause incorrect output values to be recorded.

Implication: Due to this erratum, an Intel® PT trace may contain incorrect values.

Workaround: None identified. Software should avoid using Intel® PT with single range output larger than 4KB.

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

ICX112. ADDDC Reverse Sparing May Lead to Incorrect Data

Problem: When reverse sparing an Adaptive Double Device Data Correction (ADDDC) region to a Single Device Data Correction (SDDC) region of memory, incorrect data values may be copied to the SDDC region.

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

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

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

ICX113. PCIE Completion Timeouts May Occur Under Memory Stress

Problem: Under complex microarchitectural conditions that stress the memory subsystem, a PCIe* Completion Timeout Error may occur.

Implication: Due to this erratum, a PCIe* timeout may occur which may result in a system hang.

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

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

ICX114. Unnecessary PkgC6 Exit Events

Problem: The memory controller may cause unnecessary PkgC6 exit events, thus reducing PkgC6 residency.

Implication: Due to this erratum, lower than expected PkgC6 residency may be observed.

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

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

ICX115. IAA May Fail to Properly Decode Data With a Large Header

Problem: If In-memory Analytics Accelerator (IAA) receives a header that is greater than 256B in size, it may flag a decompression error in the completion record or may incorrectly decompress the data, which will cause a mismatch between the original data CRC and the CRC in the completion record.

Implication: Due to this erratum, software may receive an unexpected data decompression failure. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

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

ICX116. System May Experience an Internal Timeout Error When Directing Intel® PT to a Small, Uncacheable, Single-Range Output Buffer

Problem: A processor hang may result if Intel® PT is enabled with Mini Time Counter (MTC) packets and single range output mode (TraceEn[0]=1, MTCEn[9]=1 and ToPA[8]=0 in IA32_RTIT_CTL MSR (0570h)), while the output buffer is less than 1 KB in size (IA32_RTIT_OUTPUT_MASK_PTRS[31:0] MSR (0561h) < 0400h) and it is mapped as UnCacheable (UC) or Write Protect (WP) memory type in the Memory Type Range Registers (MTRRs).

Implication: Due to this erratum, the system may experience an Internal Timer Error Machine Check (IA32_MCi_STATUS.MCACOD=400H; bits 15:0). Intel has only observed this erratum in a synthetic test environment.

Workaround: Avoid directing Intel® PT output to an uncacheable buffer less than 1KB in size.

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

ICX117. 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 Summary Tables of Changes.

ICX118. 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 BIOS to contain a workaround for this erratum.

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

ICX119. ADR May Fail to Complete

Problem: Under complex microarchitectural conditions, Asynchronous DRAM Refresh (ADR) flow may not successfully complete.

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

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

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

ICX120. Processor May Not Report Accurate Thermal Data Via PECI

Problem: The processor may not report accurate thermal data via the PECI interface using RdPkgConfig index 2 or 9. This issue does not affect the processor's internal thermal management capabilities.

Implication: System components that rely upon accurate processor thermal information over PECI may not behave as expected.

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

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

ICX121. Removed

ICX122. System May Hang After MCA_DISPATCHER_RUN_BUSY_TIMEOUT MCA Occurs

Problem: If the processor incurs a Machine Check Exception with a MCA_DISPATCHER_RUN_BUSY_TIMEOUT (MCOD = b000h, MCACOD = 0402h), a subsequent warm reset may not be able to properly restore system operation.

Implication: Due to this erratum, after a MCA_DISPATCHER_RUN_BUSY_TIMEOUT MCA occurs, a cold reset may be required to restore system operation.

Workaround: None identified.

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

ICX123. Processor May Signal Spurious #GP Fault

Problem: A processor that supports greater than 48-bit physical addressing (CPUID.80000008:EAX[7:0]) operating in Long Mode with 48-bit addressing and maps the PRMRR region above 128TB may generate a spurious #GP fault.

Implication: A #GP fault may be signaled when software accesses physical addresses greater than 128TB. Intel has not observed this erratum in any commercially available software.

Workaround: BIOS should configure PRMRR region to be below 128TB.

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

ICX124. RTM Abort Status May be Incorrect For INT1/INT3 Instructions

Problem: When Intel® Transactional Synchronization Extensions (Intel® TSX) is enabled, and there is a an RTM abort due to an INT1 or INT3 instruction, bit 5 of the RTM abort status (nested transaction execution) will not be set even if the RTM was nested.

Implication: Due to this erratum, software that manages RTM aborts cannot determine whether an abort is nested.

Workaround: None identified.

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

ICX125. 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 TMRT.

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

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

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

ICX126. 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 Hardware Power Management (HWPM) OOB mode.

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

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

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

ICX127. DRAM Performance May be Reduced Following Microcode Update

Problem: When the DRAM RAPL feature is enabled as indicated by MSR_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 BIOS to contain a workaround for this erratum.

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

ICX128. PCIe Bandwidth Reduction When Intel® VT-d is Enabled

Problem: Under certain PCIe* traffic conditions when Intel® Virtualization Technology (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 BIOS to contain a workaround for this erratum.

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

ICX129. WRMSR to a Few Core MSRs Might be Overwritten

Problem: If any thread is in thread C6 while another thread is updating one of the following MSRs, a subsequent transition from single thread operation to multi-thread operation or vice versa may cause that MSR to revert to its previous value. The affected MSRs are: MEMORY_CONTROL (MSR 33h bit 28), QUIESCE_CTL1 (MSR 50h) and QUIESCE_CTL2 (MSR 51h).

Implication: Due to this erratum, the values of the above MSRs may be incorrect. Intel has not observed any functional impact due to this erratum.

Workaround: None identified. Software must ensure that the other thread is not in TC6 when writing this MSR.

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

ICX130. PMON May Overcount M2Mem Ingress Events

Problem: Performance Monitoring (PMON) may incorrectly overcount Mesh-to-Memory (M2Mem) ingress events (AD Ingress and BL Ingress events (Events: 01h-08h)).

Implication: Due to this erratum, software using the PMON counts may not function as expected.

Workaround: None identified.

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

ICX131. PC6 Entry Is Not Prevented With PECI Pkg C-state Entry Control

Problem: Setting Pkg C-state Entry Control to 1 via the PECI 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 BIOS to contain a workaround for this erratum.

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

ICX132. Removed

ICX133. MCA_DISPATCHER_RUN_BUSY_TIMEOUT May be Seen During Warm Reset

Problem: Under complex microarchitectural conditions, when the processor initiates a warm reset event, it may generate a MCA_DISPATCHER_RUN_BUSY_TIMEOUT fatal machine check exception (MCACOD 402h, MSCOD b000h).

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

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

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

ICX134. Removed

ICX135. A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an AVX Gather or REP MOVS Instruction

Problem: Under complex micro-architectural conditions, when both data poison and data breakpoint events happen on an Intel® AVX Gather or REP MOVS 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 Summary Tables of Changes.

ICX136. 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 Summary Tables of Changes.

ICX137. 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: 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 Summary Tables of Changes.

ICX138. PUNIT Dispatcher May Hang

Problem: Under complex micro-architectural conditions, when the CPU enters and exits S-states the Punit dispatcher may hang and report an Machine Check Exception with MSCOD = b000h and MCACOD = 402h.

Implication: Due to this erratum, a system hang may occur.

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

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX139. 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 BIOS to contain a workaround for this erratum.

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

ICX140. SST-TF May Fail to Report an Error if Turbo is Disabled

Problem: If Turbo is disabled in the processor via the TURBO_MODE_DISABLE field in IA32_MISC_ENABLES (MSR 1A0h, bit 38), attempting to enable SST_FE (Intel® Speed Select Technology - Turbo Frequency [Intel® SST-TF]) via the OS mailbox message, the mailbox will incorrectly report that turbo is enabled.

Implication: Due to this erratum, software may incorrectly report that Turbo is enabled when it is disabled.

Workaround: None identified.

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

ICX141. BSP May Not be The Lowest Numbered APIC ID

Problem: The Advanced Programmable Interrupt Controller (APIC) ID numbering may not assign the Boot Strap Processor (BSP) to the lowest numbered APIC ID.

Implication: Due to this erratum, system software that relies on BSP to be the lowest numbered APIC ID may not function as expected.

Workaround: None identified. Software that expects the BSP to be the lowest numbered APIC ID must save and restore the BSP APIC ID when entering and exiting S3.

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

ICX142. 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. CHA configuration registers are located in PCIe* configuration space associated with the CHA bus, device, and function, with the first CHA being located at Bus U1, Device 0, Function 0 and also Bus U1, Device 10,

Function 0. There are two functions in PCI CFG space for each CHA. Accesses to 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 U1, 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 Summary Tables of Changes.

ICX143. 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 Summary Tables of Changes.

ICX144. Incorrect 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: 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 Summary Tables of Changes.

ICX145. PCIe Link Width May Degrade After a Hot Reset or Link Disable

Problem: For devices with PCIe Configuration.LinkWidth.Start timeout values near the specification max, the PCIe Link Width may reduce after a hot reset or link disable.

Implication: Due to this erratum, bandwidth available to certain PCIe endpoints may be reduced after a hot reset or link disable.

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

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

ICX146. Intel SGX TCB Recovery Boot Flow May Fail

Problem: When Post Launch Release 3 (PLR3=0d000375) Microcode Update (MCU) is loaded using Firmware Interface Table (FIT), and Trusted Compute Base (TCB) recovery is attempted for Intel® Software Guard Extensions (Intel® SGX) data generated using previous MCUs, Intel SGX will not be enabled and TCB recovery will not be possible.

Implication: Due to this erratum, TCB recovery may not be possible.

Workaround: It may be Possible for the BIOS to contain a workaround for this erratum.

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

ICX147. Mapped Out PMEM Module May Result in Unpredictable System Behavior

Problem: During a warm reset event, if an ADR event incurs a WDB flush failure, then, subsequent to the warm reset event, Intel® Optane™ Persistent Memory (Intel® Optane™ PMem) Modules may fail to be detected.

Implication: The system may experience unpredictable system behavior as the PMem module may be mapped out as non-functional. A cold reset is required to restore the system to operating with the PMem.

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

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

ICX148. VCCIN Current Draw May Exceed Specification

Problem: The processor's VCCIN current draw may exceed specification after a Surprise Warm Reset but prior to the subsequent first code fetch.

Implication: Due to this erratum, the platform may unexpectedly shut down due to an Over Current Protection event in the platform voltage regulator.

Workaround: It may be Possible for the BIOS to contain a workaround for this erratum.

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX149. 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 Summary Tables of Changes.

ICX150. 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 the BIOS to contain a workaround for this erratum.

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

ICX151. TOR_TIMEOUT MCE May Occur Under Sustained High-bandwidth PCIe Traffic

Problem: Under conditions of sustained high-bandwidth PCIe* traffic, the processor may signal a TOR_TIMEOUT machine check exception (MCE).

Implication: Due to this erratum, a TOR_TIMEOUT MCE (MCACOD 1136h, MSCOD Ch) may occur which may result in a system hang.

Workaround: It may be Possible for the BIOS to contain a workaround for this erratum.

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

ICX152. MASKMOV* Instruction To a Physical Memory Location Mapped By Two Linear Addresses of Different Page Sizes May Result In Unpredictable System Behavior

Problem: Under complex micro-architectural conditions, executing a MASKMOVQ or MASKMOVDQU instruction to a physical memory location mapped by two linear addresses of different page sizes pages may result in unpredictable system behavior if either accessed flag (A flag) or the dirty flag (D flag) of one of those pages are cleared or the transaction is to a uncacheable memory.

Implication: When this erratum occurs, the system may behave unpredictably. Intel has not observed this erratum with any commercially available software.

Workaround: Software that uses MASKMOVQ or MASKMOVDQU instructions should invalidate the TLB entries (using an INVLPG instruction) containing an address that could be accessed as part of two different page sizes after each paging-structure change that affects those pages.

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

ICX153. x87 FDP Value May be Saved Incorrectly

Problem: Execution of the FSAVE, FNSAVE, FSTENV, or FNSTENV instructions in real-address mode or virtual-8086 mode may save an incorrect value for the x87 FDP (FPU data pointer). This erratum does not apply if the last non-control x87 instruction had an unmasked exception.

Implication: Software operating in real-address mode or virtual-8086 mode that depends on the FDP value for non-control x87 instructions without unmasked exceptions may not operate properly. Intel has not observed this erratum in any commercially available software.

Workaround: None identified. Software should use the FDP value saved by the listed instructions only when the most recent non-control x87 instruction incurred an unmasked exception.

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

ICX154. System May Hang When Bus-Lock Detection Is Enabled And EPT Resides in Uncacheable Memory

Problem: On processors that support bus-lock detection (CPUID.(EAX=7, ECX=0).ECX[24]) and have it enabled (bit 2 in the IA32_DEBUGCTL MSR (1D9h)), and employ an Extended Page Table (EPT) that is mapped to an uncacheable area (UC), and the EPT_AD is enabled (bit 6 of the EPT Pointer is set), if the VMM performs an EPT modification on a predefined valid page while a virtual machine is running, the processor may hang.

Implication: Due to this erratum, the system may hang when bus-lock detection is enabled. Intel has not observed this erratum in any commercially available software.

Workaround: None identified. VMM should not map EPT tables to Uncacheable memory while using EPT_AD.

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

ICX155. THERMTRIP Will be Asserted When Conditions to Assert MEMTRIP Are Met

Problem: Regardless of the value in the OFFPKG_MEMTRIP_TO_THERMTRIP_EN field (bit 7) or the INPKG_MEMTRIP_TO_THERMTRIP_EN field (bit 6) in THERMTRIP_CONFIG_CFG register (Bus: 31; Device: 30; Function: 2; Offset: F8h), the processor will assert THERMTRIP when conditions to assert MEMTRIP are met.

Implication: When conditions to assert MEMTRIP are met, the processor will assert THERMTRIP.

Workaround: None identified.

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

ICX156. MCA Registers May Report an Incorrect Number of Corrected Errors

Problem: Machine check banks 9 through 27 may underreport corrected errors in the IA32_MCi_STATUS registers Corrected Error Count field (bits [52:38]) when two or more errors happen contemporaneously.

Implication: MCA register banks may report an incorrect number of corrected errors if the errors occur close together.

Workaround: None identified.

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

ICX157. Incorrect Memory Transaction Type May be Logged When Data Poisoning is Disabled

Problem: When MCG_CONTAIN.Poison = 0 (MSR 178h, Bit 0) and a memory location with Poison data is accessed, the processor may report the incorrect Memory Transaction Type in the MMM field (bits 6:4) of IA32_MCi_STATUS for IMC (Integrated Memory Controller) banks 13, 14, 17, 18, 21, 22, 25, and 26 (MSRs 435h, 439h, 445h, 449h, 455h, 459h, 465h, and 469h). A write Memory Transaction Type (MMM = 2h) may be reported as a read Memory Transaction Type (MMM = 1h), or a read Memory Transaction Type may be reported as a write Memory Transaction Type.

Implication: Due to this erratum, the Memory Transaction Type may not match the expected type of the memory access. Intel has not observed the errata to affect any commercially available software.

Workaround: None identified.

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

ICX158. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions

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

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

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

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

ICX159. IA32_MC6_CTL2 CMCI Enable Bit is Not Implemented Correctly

Problem: The CMCI_EN field (bit 30) in UBox MCA Control register (IA32_MC6_CTL2, MSR 286h) should return zero when written to one, but it returns one.

Implication: Due to this erratum, software may believe that CMCI is possible from machine check bank 6 when it is not.

Workaround: None identified.

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

ICX160. CPUID L2 Cache Information May Be Inaccurate

Problem: CPUID extended function 80000006H (EAX=80000006H) inaccurately reports information about the L2 cache in ECX. The function reports that the L2 cache is 8-way associative and the cache size is 256K, although the L2 cache is 20-way associative and the cache size is 1.25M.

Implication: Software that uses CPUID extended leaf 80000006H L2 cache information may operate incorrectly. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified. Software should ignore the L2 cache size and associativity information reported by CPUID extended leaf 80000006H for the affected processors.

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

ICX161. VM Entry That Clears TraceEn May Generate a FUP

Problem: If VM entry clears Intel® Processor Trace (Intel® PT) IA32_RTIT_CTL.TraceEn (MSR 570H, bit 0) while PacketEn is 1 then a Flow Update Packet (FUP) will precede the Target IP Packet, Packet Generation Disable (TIP.PGD). VM entry can clear TraceEn if the VM-entry MSR-load area includes an entry for the IA32_RTIT_CTL MSR.

Implication: When this erratum occurs, an unexpected FUP may be generated that creates the appearance of an asynchronous event taking place immediately before or during the VM entry.

Workaround: The Intel® PT trace decoder may opt to ignore any FUP whose IP matches that of a VM entry instruction.

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

ICX162. DPC Trigger Status Bit May Not Cleared

Problem: The Downstream Port Containment Trigger (DCP) Status bit (bit 0) of the DPC Status register (Bus: 0; Device: 9-12, 20-23; 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 Summary Tables of Changes

ICX163. Boot Guard ACM Authentication Failure May Result In Endless Reset Cycles

Problem: Boot guard Authenticated Code Module (ACM) authentication failure will result in a reset cycle instead of the expected shutdown event.

Implication: Due to this erratum, the system may enter an endless loop of reset cycles due to repeating authentication failures. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX164. Writing To Apic Timer Initial Count Arms The Counter

Problem: Using the APIC Timer requires writing to LVT TIMER MSR (0x832) and to Initial Count MSR (0x838). While a write to the LVT TIMER configures the timer, the write to the initial count is the actual trigger for the counter to start.

Implication: Due to this erratum, APIC timer configurations in which the initial count MSR is not the last to be written needs to be reconsidered.

Workaround: Software should write to the initial count MSR as the last action in each APIC Timer setup.

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX165. VMExit Might Be Missed On INT3/INTO Preceded by MOV/POP SS

Problem: Under complex micro-architectural conditions, a MOV SS or POP SS instruction will create a trap, that trap should be proceed on the boundary of the following instruction, if that instruction happens to be an INT0 or an INT3 instruction the trap will not cause a VMEXIT as expected.

Implication: Due to this erratum, a VMEXIT might not be generated, and once a VMEXIT will occur (due to a different reason), the VMEXIT exit_reason will also include the trap which did not initiate the VMEXIT. Intel has not observed this erratum in any commercially available software.

Workaround: None identified.

Status: For the steppings affected, refer to the Summary Tables of Changes

ICX166. Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP

Problem: Bit 63 of IA32_PERF_GLOBAL_STATUS_SET MSR (391H) is reserved. Due to this erratum, setting the bit will not result in General Protection Fault (#GP).

Implication: Software that attempts to set bit 63 of IA32_PERF_GLOBAL_STATUS_SET MSR do not generate #GP. There are no other system implications to this behavior.

Workaround: None identified.

Status: For the steppings affected, refer to the Summary Tables of Changes

Specification Changes

There are no specification clarifications in this specification update revision.

Specification Clarifications

There are no specification clarifications in this specification update revision.

Documentation Changes

There are no documentation changes in this specification update revision.