



4th Gen Intel[®] Xeon[®] Processor Scalable Family

Specification Update

Rev. 012US

December 2024



Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

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

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

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

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Your costs and results may vary.

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

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

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

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

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

SPEC[®], SPECint[®], and SPECfp[®] are registered trademarks of the Standard Performance Evaluation Corporation.

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

Copyright © 2024, Intel Corporation. All rights reserved.

Contents

Revision History	6
1.0 Preface	7
1.1 Related Documents.....	7
1.2 Nomenclature.....	7
2.0 Identification Information	9
2.1 Component Identification via Programming Interface.....	9
3.0 Component Marking Information	12
4.0 Summary Tables of Changes	13
4.1 Codes Used in Summary Tables.....	13
4.2 Errata Summary Table.....	13
5.0 Errata Details	20
6.0 Specification Changes	60
7.0 Specification Clarifications	61
8.0 Documentation Changes	62



Figures

1	Processor Preliminary Top Side Marking (Example).....	12
---	---	----



Tables

1	Component Identification via Capability Registers.....	10
2	Errata Summary Table.....	13
3	Specification Changes.....	18
4	Specification Clarifications.....	19
5	Documentation Changes.....	19

Revision History

Date	Revision	Description
December 2024	012US	<ul style="list-style-type: none"> Added errata SPR141 and SPR143.
November 2024	011US	<ul style="list-style-type: none"> Added errata SPR137 to SPR140.
August 2024	010US	<ul style="list-style-type: none"> Added errata SPR134 to SPR136.
June 2024	009US	<ul style="list-style-type: none"> Added errata SPR129 to SPR133.
May 2024	008US	<ul style="list-style-type: none"> Added errata SPR125 to SPR128.
March 2024	007US	<ul style="list-style-type: none"> Added errata SPR123 and SPR124.
October 2023	006	<ul style="list-style-type: none"> Added errata SPR120 to SPR122.
August 2023	005	<ul style="list-style-type: none"> Added errata SPR118 to SPR119. Updated errata SPR91.
July 2023	004	<ul style="list-style-type: none"> Added errata SPR115 to SPR117.
May 2023	003	<ul style="list-style-type: none"> Added errata SPR89 to SPR114. Updated errata status for SPR44, SPR66, SPR80, and SPR81.
April 2023	002	<ul style="list-style-type: none"> Added errata SPR82 to SPR88. Updated the Component Identification via Registers
February 2023	001	<ul style="list-style-type: none"> Initial Release

1.0 Preface

This document is an update to the specifications contained in the following 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](#) on page 7 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.

1.1 Related Documents

Document Title	Document Number/ Location
Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 1: Basic Architecture	671436 ¹
Intel® 64 and IA-32 Architectures Software Developer’s Manual Combined Volumes 2A, 2B, 2C, and 2D: Instruction Set Reference, A-Z	671110 ¹
Intel® 64 and IA-32 Architectures Software Developer’s Manual Combined Volumes 3A, 3B, 3C, and 3D: System Programming Guide	671447 ¹
<i>Eagle Stream Platform BIOS Writer’s Guide</i>	613938 ¹
ACPI Specifications	www.acpi.info ²

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

1.2 Nomenclature

Errata are design defects or errors. These may cause the 4th 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.

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 life cycle, 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).

2.0 Identification Information

2.1 Component Identification via Programming Interface

The 4th Gen Intel® Xeon® Scalable Processors stepping can be identified by the following register contents:

CPUID (Offset:1Ah-19h)	Extended Family ID ¹	Extended Model ²	Reserved	Processor Type ³	Processor Family ⁴	Processor Model ⁵	Processor Stepping ⁶
Bit	27:20	19:16	15:14	13:12	11:8	7:4	3:0
XCC E-5/ E-4, MCC S-3/ S-2, HBM B-3	00000000b	1000b		0b	0110b	1111b	1111b

NOTES

1. The Extended Family, bits [27:20] are used in conjunction with the Processor Family, 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], are 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 Processor Family corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
5. The Processor Model, bits [7:4] 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 Processor Stepping, bits [3:0] indicates the revision number of that model. See [Table 1](#) on page 10 for the processor stepping ID number in the CPUID information.

To find the mapping between a processor's CPUID and its Family/Model number, see the [Intel® 64 and IA-32 Architectures Software Developer Manual Combined Volumes](#).

A complete description of the processor identification and feature determination is located in Chapter 20.

When EAX is set to a value of `1`, the CPUID instruction returns the Processor Family, Extended Model ID, Processor Type, Family, Model, and Stepping together referred 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 TLB descriptor parameters are provided in the EAX, EBX, ECX, and EDX general purpose registers after the CPUID instruction is executed with a 2 in the EAX register. Special uses of general purpose registers include: EAX (Accumulator for operands and results data), EBX (Pointer to data in the DS segment), ECX (Counter for string and loop operations), and EDX (I/O pointer).

The 4th Gen Intel® Xeon® Scalable Processors Stepping can also be identified in *4th Gen Intel® Xeon® Processor Scalable Family, Codename Sapphire Rapids Registers Specification*

Table 1. Component Identification via Capability Registers

Physical Chip	Stepping	SEGMENT, WAYNESS	CPUID	SEGMENT ¹ [Bits 5:3]			WAYNESS ² [Bits 1:0]		PHYSICAL_CHO P ³ [Bits 7:6]		
				Offset [B:31, D:30, F:3] + 84h						Offset [B:31, D:30, F:3] + 94h	
				5	4	3	1	0	7	6	
XCC	E-5/E-4	Server, 1S	806F8/806F7	1	1	1	0	0	1	1	
XCC	E-5/E-4	Server, 2S	806F8/806F7	1	1	1	0	1	1	1	
XCC	E-5/E-4	Server, 4S	806F8/806F7	1	1	1	1	0	1	1	
XCC	E-5/E-4	Server, 8S	806F8/806F7	1	1	1	1	1	1	1	
MCC	S-3/S-2	Server, 1S	806F8/806F7	1	1	1	0	0	0	1	
MCC	S-3/S-2	Server, 2S	806F8/806F7	1	1	1	0	1	0	1	
MCC	S-3/S-2	Server, 4S	806F8/806F7	1	1	1	1	0	0	1	
HBM	B-3	Server, 1S	806F8	1	1	1	0	0	1	1	
HBM	B-3	Server, 2S	806F8	1	1	1	0	1	1	1	
HBM	B-3	Server, 4S	806F8	1	1	1	1	0	1	1	

NOTES

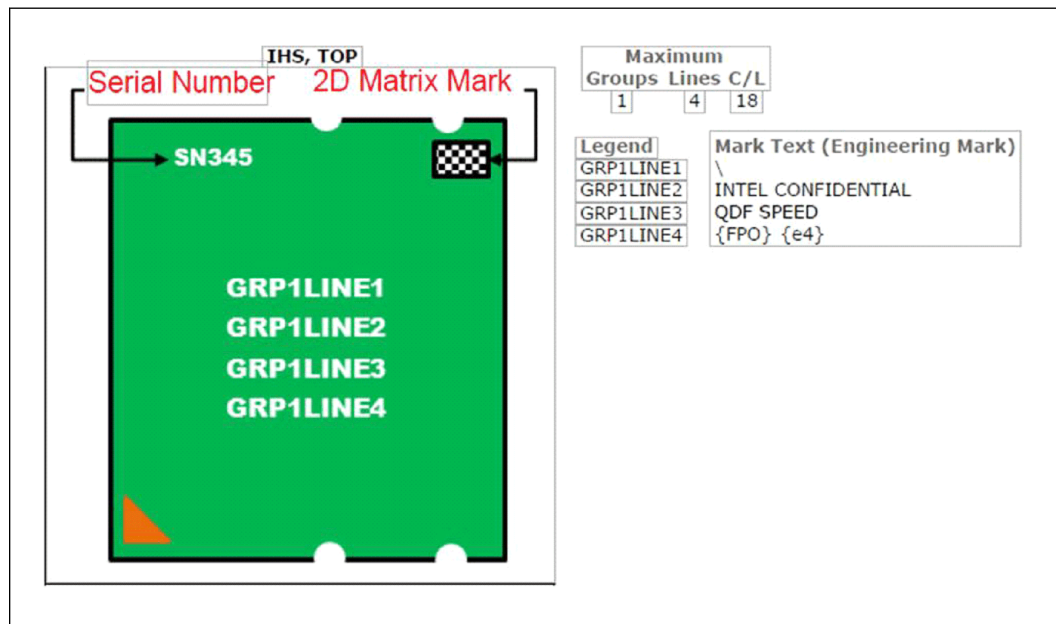
1. The SEGMENT, bits [5:3] corresponds to 111: Server; 011: Server-FPGA; 001: Server-Fabric; 100: HPC; 110: Server-Atom; 010: Workstation; 000: HEDT; Others: Reserved.
2. The WAYNESS, bits [1:0] corresponds to 00=1S, 01=2S, 10 = 4S, 11 = 8S.
3. The PHYSICAL_CHOP, bits [7:6] corresponds to 11:XCC; 01:MCC; 11:HBM.

The 4th Gen Intel® Xeon® Scalable Processors Capability Registers can also be identified in *4th Gen Intel® Xeon® Processor Scalable Family, Codename Sapphire Rapids Registers Specification* .

3.0 Component Marking Information

The 4th Gen Intel® Xeon® Processor Scalable Family can be identified by the following register markings.

Figure 1. Processor Preliminary Top Side Marking (Example)



For 4th Gen Intel® Xeon® Processor Scalable Family, Codename Sapphire Rapids SKUs, see [4th Gen Intel® Xeon® Scalable Processors](#).

4.0 Summary Tables of Changes

The following tables indicate the Specification Changes, Errata, Specification Clarifications, or Documentation Changes which apply to the 4th 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:

4.1 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 plan to fix this erratum.

4.2 Errata Summary Table

Table 2. Errata Summary Table

Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR1	No Fix	No Fix	No Fix	IPSR May Not Function Correctly
SPR2	No Fix	No Fix	No Fix	Poison Data Reported Instead of a CS Limit Violation
SPR3	No Fix	No Fix	No Fix	Monitor Instructions to Legacy VGA Region May Fail
SPR4	No Fix	No Fix	No Fix	TILEDATA State May Be Saved Incorrectly
SPR5	No Fix	No Fix	No Fix	A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an AMX Tile-Load or AVX Gather or REP MOVSB Instruction
SPR6	No Fix	No Fix	No Fix	IFS MSRs Will Ignore a Non-Zero EDX Value And Not Signal a #GP
SPR7	No Fix	No Fix	No Fix	Processor May Signal Spurious #GP Fault
SPR8	No Fix	No Fix	No Fix	A Break Point May be Hit Twice When a VM Exit Without Commit Occurs
SPR9	No Fix	No Fix	No Fix	Faulted XRSTORS Instruction May Result in Unexpected X87 FTW Value
SPR10	No Fix	No Fix	No Fix	Error Conditions Detected During Cold Reset May Not be Cleared by Subsequent Warm Reset

continued...



Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR11	No Fix	No Fix	No Fix	DSA/IAX Does Not Log The E2E Prefix Bit And The Prefix-Type Bits in AERTLPPLOG1
SPR12	No Fix	No Fix	No Fix	The Processor May Drop Noncompliant Posted Peer-to-peer Transactions
SPR13	No Fix	No Fix	No Fix	Certain Bits in IA32_MC5_STATUS Register Will Always Return 0
SPR14	No Fix	No Fix	No Fix	Occupancy Interrupt Handle is Not Checked Against Interrupt Table Size
SPR15	No Fix	No Fix	No Fix	Processor May Incorrectly Set PFD Assisted in Correction Bit in Memory Controller
SPR16	No Fix	No Fix	No Fix	DSA CMDSTATUS Register May Not Reflect Correct Hardware Status
SPR17	No Fix	No Fix	No Fix	Remapping Hardware May Set Access/Dirty Bits in a First-stage Page-table Entry
SPR18	No Fix	No Fix	No Fix	System Software May Not Receive Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Fault SPT.3 For Non-Zero Writes to b[191:HAW +128]
SPR19	No Fix	No Fix	No Fix	APCTL.APNGE Should be RW Instead of RWS
SPR20	No Fix	No Fix	No Fix	CXL Device May Not Receive Viral
SPR21	No Fix	No Fix	No Fix	OOBMSM TSC Will be 320ns Behind The Globally Aligned Counter
SPR22	No Fix	No Fix	No Fix	Performance Monitoring Event Coherent_ops May Undercount
SPR23	No Fix	No Fix	No Fix	PCIe Link Re-Equalization May Not Occur if Link is in L1 State
SPR24	No Fix	No Fix	No Fix	Machine Check Bank 4 UCNA Errors May Not be Signaled
SPR25	No Fix	No Fix	No Fix	DSA/IAA Use of Priv and PASID
SPR26	No Fix	No Fix	No Fix	Reserved(0) Check For a PASID Table Entry May Not Happen For a DMA Request
SPR27	No Fix	No Fix	No Fix	Remapping Hardware May Not Generate a Page Request Group Response Message While Operating in Legacy Mode or Abort DMA Mode
SPR28	No Fix	No Fix	No Fix	Remapping Hardware May Abort ZLR to Second-Stage Write Only Pages
SPR29	No Fix	No Fix	No Fix	Remapping Hardware with Major Version Number 6 Incorrectly Advertises the ESRTPS Support
SPR30	No Fix	No Fix	No Fix	Platform May Hang if System Software Sends a Page Group Response or DevTLB Invalidation to Non-existent Requester ID
SPR31	No Fix	No Fix	No Fix	Remapping Hardware Does Not Perform Reserved (0) Check in Page Response Descriptor
SPR32	No Fix	No Fix	No Fix	Remapping Hardware Implements b[31:16] of the three Event Data Registers (VTDBAR offsets 0x3C, 0xA4, and 0xE4) as Read-Writable
SPR33	No Fix	No Fix	No Fix	IAA Do Not Report Overlap Errors For AECS Size of 2GB or Greater
SPR34	No Fix	No Fix	No Fix	DSA/IAA Invalid TC Not Reported in The SWERROR Register
SPR35	No Fix	No Fix	No Fix	IAA Unaligned Completion Record Address Error is Not Reported in SWERROR Register
SPR36	No Fix	No Fix	No Fix	Intel® UPI Link Not Resetting When L1 Mismatch Occurs Between Local and Remote Sockets
SPR37	No Fix	No Fix	No Fix	DSA/IAA May Fail to Log an MDPE Error For Back-to-Back Parity Errors
continued...				

Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR38	No Fix	No Fix	No Fix	Relaxed Ordering Not Disabled by DEVCTL.ERO bit for DSA/IAA Upstream Transactions
SPR39	No Fix	No Fix	No Fix	System Address Logged For WDB Parity Errors May be Incorrect
SPR40	No Fix	No Fix	No Fix	Incorrect MCACOD For L2 MCE
SPR41	No Fix	No Fix	No Fix	System May Hang Due to Full LLRB
SPR42	No Fix	No Fix	No Fix	IAA May Fail to Properly Decode Data With a Large Header
SPR43	No Fix	No Fix	No Fix	Memory Controller Violates JEDEC RCD tCSALT Timing
SPR44	Fixed	Fixed	Fixed	Wrong CKE Signal Used on 1 DPC 3DS 4H Configs
SPR45	No Fix	No Fix	No Fix	Address May Not be Logged For a UCR Error Detected in The MLC
SPR46	No Fix	No Fix	No Fix	VT-d DMA Remapping Hardware May Hang if it Encounters Page Request Queue Overflow Condition
SPR47	No Fix	No Fix	No Fix	Receiver Common Mode Input Impedance May be Below Specification When Interface is Powered Down
SPR48	No Fix	No Fix	No Fix	Remapping Hardware Will Not Report The PASID Value For RTA.2 Faults in Modes Other Than Scalable Mode
SPR49	No Fix	No Fix	No Fix	Remapping Hardware Does Not Perform a Reserved(0) Check in Interrupt Remap Table Entry
SPR50	No Fix	No Fix	No Fix	Processor PCIe Root Port Link Spurious Data Parity Error May be Reported
SPR51	No Fix	No Fix	No Fix	Mismatch Between UboxErrMisc and MCI_STATUS Registers Error Logs
SPR52	No Fix	No Fix	No Fix	CHA UCNA Errors May be Incorrectly Controlled by MCI_CTL Enable Bits
SPR53	No Fix	No Fix	No Fix	Reading The PPERF MSR May Not Return Correct Values
SPR54	No Fix	No Fix	No Fix	No #GP Will be Signaled When Setting MSR_MISC_PWR_MGMT.ENABLE_SDC if MSR_MISC_PWR_MGMT.LOCK is Set
SPR55	No Fix	No Fix	No Fix	System May Experience an Internal Timeout Error When an Internal Parity Error Occurs While Working With Intel® AMX
SPR56	No Fix	No Fix	No Fix	Last Branch Records May Not Survive Warm Reset
SPR57	No Fix	No Fix	No Fix	Single Step on Branches Might be Missed When VMM Enables Notification On VM Exit
SPR58	No Fix	No Fix	No Fix	Incorrect #CP Error Code on UIRET
SPR59	No Fix	No Fix	No Fix	#GP May be Serviced Before an Instruction Breakpoint
SPR60	No Fix	No Fix	No Fix	Unexpected #PF Exception Might Be Serviced Before a #GP Exception
SPR61	No Fix	No Fix	No Fix	VMX-Preemption Timer May Not Work if Configured With a Value of 1
SPR62	No Fix	No Fix	No Fix	User Interrupt Might be Delayed
SPR63	No Fix	No Fix	No Fix	VM Exit Qualification May Not be Correctly Set on APIC Access While Serving a User Interrupt
SPR64	No Fix	No Fix	No Fix	Software Tuning That Relies on PCLS Values May Experience Inaccurate Event Counts
SPR65	No Fix	No Fix	No Fix	Multiple SGX_Doorbell_Errors on Ubox Response Mismatch
SPR66	Fixed	Fixed	Fixed	ECS Readout Fails on Mixed Mode Systems

continued...



Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR67	No Fix	No Fix	No Fix	Intel® DSA/Intel® IAA Completion Record is Not Written For Non-Completion Record Invalid Traffic Classes
SPR68	No Fix	No Fix	No Fix	Intel® IAA Expand Operation With PRLE Format Input May Return an Error
SPR69	No Fix	No Fix	No Fix	Intel® IAA Compression with Compress Bit Order Set May Produce an Odd Number of Bytes
SPR70	No Fix	No Fix	No Fix	Intel® IAA Source 2 Not Written Properly When Source 2 Size is 32 Bytes
SPR71	No Fix	No Fix	No Fix	Intel® IAA May Not Report Invalid Filter Flags Status Code When Source 2 Bit Order Field is Set
SPR72	No Fix	No Fix	No Fix	Intel® IAA Does Not Allow Source 1 Size to be 0 For Expand Operation
SPR73	No Fix	No Fix	No Fix	Intel® DSA/Intel® IAA And WQ Configuration Registers May be Incorrectly Updated
SPR74	No Fix	No Fix	No Fix	Invalid Flags Field of The Completion Record May Not be Set Correctly For Intel IAA Compression Operation
SPR75	No Fix	No Fix	No Fix	With Intel® SGX Disabled, Software That Relies on ENCLVexiting May Not Function as Expected
SPR76	No Fix	No Fix	No Fix	Headers Logged in AERHDRLOG for an AER Error for Intel DSA/IAA may be Incorrect
SPR77	No Fix	No Fix	No Fix	Intel® DSA/Intel® IAA May Fail to Send an ERR_FATAL Message if a Non-Fatal Error Occurs in The Same Cycle
SPR78	No Fix	No Fix	No Fix	Intel® DSA/Intel® IAA May Fail to Log an Unexpected Completion Error For an Invalid ATS Response
SPR79	No Fix	No Fix	No Fix	Intel® IAA Compression Output Buffer Overflow Error May be Incorrectly Reported
SPR80	No Fix	No Fix	No Fix	Intel® QuickAssist Technology Accelerator May Violate ATS Invalidation Completion Ordering
SPR81	No Fix	No Fix	No Fix	Intel® QuickAssist Technology Accelerator Device May Not Invalidate PASID SupervisorPrivilege Translations Privilege Translations
SPR82	Fixed	Fixed	Fixed	The Time-Stamp Counter May Report an Incorrect Value
SPR83	No Fix	No Fix	No Fix	UPI Machine Check Bank May Not Report The Most Recently Logged Error
SPR84	No Fix	No Fix	No Fix	PECI Wire Host may Continuously Receive a Completion Code of 0x80
SPR85	No Fix	No Fix	No Fix	DDR5 9x4 DIMMs ECS Data May be Reported Incorrectly
SPR86	No Fix	No Fix	No Fix	RETRY_RD_ERR_LOG_MISC.DDR5_9x4_half_device Bit Maybe Incorrect
SPR87	No Fix	No Fix	No Fix	PIROM Reports The Wrong 2 DPC Speed For Processors With Less Than 4800 MT/s 1 DPC Speed
SPR88	No Fix	No Fix	No Fix	An MDF Parity Error May Incorrect Set The Overflow Bit
SPR89	No Fix	No Fix	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
SPR90	No Fix	No Fix	No Fix	CPUID Reports Incorrect Number of Ways For The Load DTLB
SPR91	No Fix	No Fix	No Fix	Intel® PT Trace May Contain Incorrect Data When Configured With Single Range Output Larger Than 4KB
SPR92	No Fix	No Fix	No Fix	On Instructions Longer Than 15 Bytes, #GP Exception is Prioritized And Delivered Over #CP Exception
<i>continued...</i>				

Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR93	No Fix	No Fix	No Fix	Mismatch on DR6 Value When Breakpoint Match is on Bitmap Address
SPR94	No Fix	No Fix	No Fix	RTM Abort Status May be Incorrect For INT1/INT3 Instructions
SPR95	No Fix	No Fix	No Fix	WRMSR to Reserved Bits of IA32_L3_QOS_Mask_15 Will Not signal a #GP
SPR96	No Fix	No Fix	No Fix	x87 FDP Value May be Saved Incorrectly
SPR97	No Fix	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
SPR98	No Fix	No Fix	No Fix	Exit Qualification For EPT Violations on Instruction Fetches May Incorrectly Indicate That The Guest-physical Address Was Writeable
SPR99	No Fix	No Fix	No Fix	Processor May Generate Spurious Page Faults On Shadow Stack Pages
SPR100	No Fix	No Fix	No Fix	Processor May Hang if Warm Reset Triggers During BIOS Initialization
SPR101	No Fix	No Fix	No Fix	IA32_MC1_STATUS MSR May Not Log Errors When IA32_MC1_CTL MSR is Set to Not Signal Errors
SPR102	No Fix	No Fix	No Fix	System May Hang When Bus-Lock Detection Is Enabled And EPT Resides in Uncacheable Memory
SPR103	No Fix	No Fix	No Fix	OFFCORE_REQUESTS_OUTSTANDING Performance Monitoring Events May be Inaccurate
SPR104	No Fix	No Fix	No Fix	Incorrect MCACOD For L2 Prefetch MCE
SPR105	No Fix	No Fix	No Fix	Call Instruction Wrapping Around The 32-bit Address Boundary May Return to Incorrect Address
SPR106	No Fix	No Fix	No Fix	ADDDC Reverse Sparing May Lead to Incorrect Data
SPR107	No Fix	No Fix	No Fix	SGX ENCLU[EACCEPT] Will Not Cause #GP When TCS.PREVSP is Non-zero
SPR108	No Fix	No Fix	No Fix	Setting MISC_FEATURE_CONTROL.DISABLE_THREE_STRIKE_CNT Does Not Prevent The Three-strike Counter From Incrementing
SPR109	No Fix	No Fix	No Fix	Mismatch Between UboxErrMisc and MCI_STATUS Registers Error Logs
SPR110	No Fix	No Fix	No Fix	SST-TF May Fail to Report an Error if Turbo is Disabled
SPR111	No Fix	No Fix	No Fix	Unexpected Rollover in MBM Counters
SPR112	Fixed	Fixed	Fixed	In 96bit ECC Mode, a Correctable Error May be Incorrectly Logged in RETRY_RD_ERR_LOG With UC Bit Set
SPR113	No Fix	No Fix	No Fix	System Crash Observed on Host When TD Private Pages Are Not Zeroed Out if Reused For Non-TD SW
SPR114	No Fix	No Fix	No Fix	IA32_MC2_ADDR And IA32_MC2_MISC MSRs Will be Cleared on Warm Reset
SPR115	No Fix	No Fix	No Fix	Multiple Write CRC Errors May Lead to System Hang
SPR116	No Fix	No Fix	No Fix	Intel DSA/IAA CTO Errors May Inconsistently Update the Prefix Log and Prefix Log Present Flag
SPR117	No Fix	No Fix	No Fix	RTIT_CTL.TRACE_EN May be Disabled at BIOS_DONE Even if it Was Previously Enabled
SPR118	No Fix	No Fix	No Fix	WRMSR to a Few Core MSRs Might be Overwritten
SPR 119	Fixed	Fixed	Fixed	Virtualize IA32_SPEC_CTRL VM-execution Control Does Not Properly Virtualize Bits [63:32] of IA32_SPEC_CTRL

continued...



Erratum ID	Processor Stepping XCC E-5	Processor Stepping MCC S-3	Processor Stepping HBM B-3	Title
SPR 120	No Fix	No Fix	No Fix	Performance Monitoring Events used by TMA May be Inaccurate
SPR 121	No Fix	No Fix	No Fix	Performance Monitoring Event IDQ.MS_UOPS May Undercount
SPR 122	Fixed	Fixed	Fixed	CHA TOR Timeout May Occur
SPR 123	No Fix	No Fix	No Fix	ECS Reads May Fail to Complete in Certain Memory Configurations and Conditions
SPR 124	Fixed	Fixed	Fixed	A Write to The TSC_Deadline MSR May Cause an Unexpected Timer Interrupt
SPR125	No Fix	No Fix	No Fix	Processor Performance Monitor Discovery Table Incorrectly Reports Data for Intel® UPI Port 3 on MCC SKUs
SPR126	No Fix	No Fix	No Fix	Certain PCIe Devices Are Not Accessible Out-of-Band With Local Bus Number
SPR127	No Fix	No Fix	No Fix	Remapping Hardware May Encounter Incorrect Error Code in Invalidation Queue Error Record Register
SPR128	No Fix	No Fix	No Fix	Disabling The APIC While an Interrupt is Being Delivered May Cause a System Hang
SPR129	No Fix	No Fix	No Fix	RTTO May Occur at Lower Speed And Reduce Link Width
SPR130	No Fix	No Fix	No Fix	Unexpected System Software Behavior on Systems With Intel® HT Enabled
SPR131	No Fix	No Fix	No Fix	Intel® DSA DIF Insert and DIF Strip Operations Do Not Properly Check For Overlapping Buffers
SPR132	No Fix	No Fix	No Fix	PCIe Link Degradation Following a Leaky Bucket Event
SPR133	No Fix	No Fix	No Fix	Intel® DSA or Intel® IAA May Cause Unpredictable System Behavior When Incorrectly Forwarding Data
SPR134	No Fix	No Fix	No Fix	MCA not Logged for WbNoInvd Data ECC Errors
SPR135	No Fix	No Fix	No Fix	Intel® DSA Memory Write with Incorrect Parity May Result in a System Crash
SPR136	Fixed	Fixed	Fixed	Internal Timeout MCE May Occur on Systems with CXL I/O Devices
SPR137	No Fix	No Fix	No Fix	DIMMs May Not be Placed in Self-Refresh Mode Prior to Warm Reset
SPR138	Fixed	Fixed	Fixed	Communication Error With Intel® On Demand Mailbox
SPR139	Fixed	Fixed	Fixed	Intel® DSA Device May Hang
SPR140	No Fix	No Fix	No Fix	A PCIe Device May Not Function After a Warm Reset
SPR141	Fixed	Fixed	Fixed	VM Exit Following MOV to CR8 Instruction May Lead to Unexpected IDT Vectoring-Information
SPR142	Fixed	Fixed	Fixed	System May Hang Waiting For Intel® VT-d Invalidation Wait Descriptor to Complete
SPR143	No Fix	No Fix	No Fix	Power May Be Displayed Incorrectly

Table 3. Specification Changes

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

Table 4. Specification Clarifications

No.	Specification Clarifications
SPRXX	None for this revision of this specification update.

Table 5. Documentation Changes

No.	Documentation Changes
SPRXX	None for this revision of the specification update.

5.0 Errata Details

SPR1. IPSR May Not Function Correctly

Problem: Poison created within the Internal memory controller may not log the correct system address when operating in the following modes of operation, 1-clock gating enabled, 2-channel XOR is enabled, and 3 Intel® Optane™ two-level memory.

Implication: Due to this erratum, software that relies upon the Internal Poison Source Register (IPSR) bit may function incorrectly.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR2. Poison Data Reported Instead of a CS Limit Violation

Problem: Under complex microarchitectural 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, see the [Summary Tables of Changes](#) on page 13.

SPR3. Monitor Instructions to Legacy VGA Region May Fail

Problem: Monitor instructions that target the legacy VGA region (A0000h - BFFFFh) may generate a Machine Check Exception (MCE) in Cache Home Agent (CHA) Machine Check Banks (Banks 9, 10, and 11) MCI_STATUS MSRs (425h, 429h, or 42Dh) with a System Address Decode error (MSCOD = 05h).

Implication: Due to this erratum, a user application or Virtual Machine (VM) guest that is allowed to use MONITOR or UMONITOR instructions to the legacy VGA region may generate a fatal machine check exception.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR4. TILEDATA State May Be Saved Incorrectly

Problem: If execution of XRSTOR or XRSTORS causes a fault or a VM exit, a subsequent execution of XSAVE, XSAVEC, XSAVEOPT, or XSAVES instructions may incorrectly save the TILEDATA state component as all zeroes. This will occur only if the execution of XRSTOR or XRSTORS is attempting to set the TILECFG state component to its initial configuration and to restore the TILEDATA state component from the XSAVE area in memory.

Implication: Due to this erratum, the data saved in the XSAVE area for the TILEDATA state may be incorrect.

Workaround: None identified. Following an execution of XRSTOR or XRSTORS that causes a fault or a VM exit, software should not use the TILEDATA state component saved by a subsequent execution of XSAVE, XSAVEC, XSAVEOPT, or XSAVES that occurs before re-executing the original instruction (after addressing the cause of the fault or VM exit).

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR5. A Poison Data Event May Not be Serviced if a Data Breakpoint Occurs on an Intel® AMX Tile-Load or Intel® AVX Gather or REP MOVBS Instruction

Problem: Under complex microarchitectural conditions, when both data poison and data breakpoint events happen on an Intel® Advanced Matrix Extensions (Intel® AMX) Tile-Load or Intel® Advanced Vector Extensions (Intel® AVX) Gather or REP MOVBS 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: It may be possible for BIOS to contain a mitigation for this erratum. When applying the mitigation a collision between a poison and a data breakpoint will result in skipping the data breakpoint.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR6. IFS MSRs Will Ignore a Non-Zero EDX Value And Not Signal a #GP

Problem: A WRMSR instruction to one of the Intel® In-field scan (IFS) MSRs not in 64-bit mode with a non-zero value in EDX will not trigger a global protection (#GP) fault. This affects COPY_SCAN_HASHES (MSR 2C2h) and AUTHENTICATE_AND_COPY_CHUNK (MSR 2C4h).

Implication: Due to this erratum, IFS software running in a non 64-bit mode and attempting the above WRMSR with non-zero value in EDX will not #GP and instead use a linear address which ignores EDX value.

Workaround: None identified. Software should make sure EDX is clear for the above instructions when not running in 64-bit mode.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR7. 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 128 TB may generate a spurious #GP fault.

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

Workaround: None identified. BIOS should configure PRMRR region to be below 128 TB.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR8. A Break Point May be Hit Twice When a VM Exit Without Commit Occurs

Problem: When a VM Exit happens without a commit in the middle of a guest exception handling, the RF flag (bit 17 of the EFLAGS) will be cleared, if a code breakpoint was configured on the VM entry instruction, the clearing of the RF flag will cause that code breakpoint to be served again when we perform a VM Entry back to the guest.

Implication: Due to this erratum, software may observe a code breakpoint twice on an instruction if a VM Exit without commit occurs.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR9. Faulted XRSTORS Instruction May Result in Unexpected X87 FTW Value

Problem: Under complex microarchitectural conditions, when a #GP fault (General Protection) happens on an XRSTORS instruction that attempts to INIT both x87 and UINTR states, the x87 FPU Tag Word (FTW) may result in an unexpected value.

Implication: Due to this erratum, the value of the FTW state may be incorrect.

Workaround: None identified. Software should rerun the XRSTORS instruction after handling the #GP.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR10. Error Conditions Detected During Cold Reset May Not be Cleared by Subsequent Warm Reset

Problem: Under certain microarchitectural conditions, if an IP_READY_TIMEOUT error (MCACOD = 0402h, MSCOD=0100h) occurs following a cold reset, a subsequent warm reset may not clear out the underlying error conditions, and another error may not be detected.

Implication: Due to this erratum, unpredictable system behavior may occur. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified. BIOS can detect the IERR and force a cold reset to bring the processor back to a known good state.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR11. Intel® DSA / Intel® IAX Does Not Log The E2E Prefix Bit And The Prefix-Type Bits in AERTLPPLOG1

Problem: For internal transactions that have a Process Address Space Identifier (PASID) TLP Prefix and have a TLP error, Data Streaming Accelerator (DSA) or In-Memory Analytics Accelerator (IAX) units do not correctly log the E2E prefix (Bus: 8; Device: 1; Function: 0; Offset: 138h; AERTLPPLOG1, bit[28]) and the prefix-type (AERTLPPLOG1 bits[27:24]).

Implication: Due to this erratum, an incorrect TLP prefix type may be logged in AERTLPPLOG1.

Workaround: None identified. As DSA and IAX only support TLPs with a PASID prefix, software should treat E2E prefix as 1 (AERTLPPLOG1[28]=1) and prefix-type as PASID (AERTLPPLOG1[27:24]=0001b).

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR12. The Processor May Drop Noncompliant Posted Peer-to-peer Transactions

Problem: If the processor receives a noncompliant posted PCIe* peer-to-peer transaction with non-zero upper tag bits [9:8], it may drop the transaction instead of forwarding it to the intended destination.

Implication: Due to this erratum, PCIe* devices that perform peer-to-peer posted transactions may not operate as expected. Intel has not observed this erratum with any commercially available devices.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR13. Certain Bits in IA32_MC5_STATUS Register Will Always Return 0

Problem: IA32_MC5_STATUS register (MSR 415h, bits [36:34]) always return 0 on a read.

Implication: Due to this erratum, software that attempts to write a non-zero value to bits [36:34] in IA32_MC5_STATUS will always read a 0 on subsequent reads.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR14. Occupancy Interrupt Handle is Not Checked Against Interrupt Table Size

Problem: In the In-Memory Analytics Accelerator (IAX) the value of the Occupancy Interrupt Handle (OIH) is not checked against the interrupt table size.

Implication: Due to this erratum, if the OIH is programmed incorrectly by the host driver, an incorrect entry will be used, and an incorrect interrupt will be generated if the entry has the Mask bit cleared.

Workaround: None identified. Software must ensure OIH is programmed correctly.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR15. Processor May Incorrectly Set PFD Assisted in Correction Bit in Memory Controller

Problem: When `RETRY_RD_ERR_LOG_ADDRESS1.FAILED_DEV = 9` (Register `MEM_BAR[0-3]`, Offset: `22C58h`; bits `5:0`) and `RETRY_RD_ERR_LOG_PARITY.PAR_SYN[63:48] = FFFFh` (Register `MEM_BAR[0-3]`, Offset: `22F08h`; bits `63:48`), the `RETRY_RD_ERR_LOG_ADDRESS1.bist_corr` bit (Register `MEM_BAR[0-3]`, Offset: `22C58h`; bit `30`) may be incorrectly set.

Implication: When this erratum occurs, software that relies upon the `bist_corr` bit may behave incorrectly.

Workaround: None identified. If `RETRY_RD_ERR_LOG_ADDRESS1.FAILED_DEV = 9` and `RETRY_RD_ERR_LOG_PARITY.PAR_SYN[63:48] = FFFFh`, software should ignore the value of 1 for

the `bist_corr` bit.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR16. DSA CMDSTATUS Register May Not Reflect Correct Hardware Status

Problem: After submitting a command to the Data Streaming Accelerator (DSA) CMD register (BAR0 offset 0xA0), a subsequent read to CMDSTATUS register (BAR0 offset 0xA8) may incorrectly see a `CMDSTATUS.ACTIVE` (bit 31) value of 0 before it has had a chance to change to 1.

Implication: Due to this erratum, software that relies upon the `CMDSTATUS.ACTIVE` bit may function incorrectly.

Workaround: None identified. Software that relies on the `CMDSTATUS.ACTIVE` bit should read this bit twice, discarding the first read result.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR17. Remapping Hardware May Set Access/Dirty Bits in a First-stage Page-table Entry

Problem: When remapping hardware is configured by system software in scalable mode as Nested (`PGTT=011b`) and with `PWSNP` field Set in the `PASID-table-entry`, it may Set Accessed bit and Dirty bit (and Extended Access bit if enabled) in first-stage page-table entries even when second-stage mappings indicate that corresponding first-stage page-table is Read-Only.

Implication: Due to this erratum, pages mapped as Read-only in second-stage page-tables may be modified by remapping hardware Access/Dirty bit updates.

Workaround: None identified. System software enabling nested translations for a VM should ensure that there are no read-only pages in the corresponding second-stage mappings.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR18. System Software May Not Receive Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Fault SPT.3 For Non-Zero Writes to b[191:HAW+128]

Problem: When the `PASID Granular Translation Type (PGTT)` field (bits 8:6) in the Scalable-Mode `PASID Table` has a value of 010b (Second-level) or 100b (Pass-through), and software writes a non-zero value to `b[191:HAW+128]`, no Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) fault is generated.

Implication: Due to this erratum, system software may not receive Intel® VT-d fault SPT.3 (fault reason 5ah) when software writes a non-zero value to `b[191:HAW+128]`. Intel has not observed any functional implications due to this erratum.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR19. APCTL.APNGE Should be RW Instead of RWS

Problem: Alternate Protocol Negotiation Global Enable (APNGE) field (bit 8) in Alternate Protocol Control register (APCTL) (Bus: 5-0; Device: 1,3,5,7; Function: 0; Offset: B28h) has been implemented as Sticky-Read-Write (RWS) but it should be Read-Write (RW).

Implication: Due to this erratum, the APCTL register is not cleared on warm reset, which violates the PCIe* Base Specification version 5.0. Intel has not observed any functional implications from this erratum.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR20. CXL* Device May Not Receive Viral

Problem: When the processor detects a Data Parity Error in a downstream packet, it may fail to transmit a Viral indication to a Compute Express Link* (CXL*) partner. However, the processor will put the CXL* Link into LinkError state.

Implication: Due to this erratum, the CXL* Link will go into the LinkError state instead of going into Viral.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR21. OOBMSM TSC Will be 320ns Behind The Globally Aligned Counter

Problem: The value in the time stamp counter (TSC) in OOBMSM (Bus: , Device: , Function: , Offset h) will be 320ns behind the globally aligned TSC (BDF).

Implication: Due to this erratum, the TSC value recorded in the OOB Crashlog will differ by 320ns from the globally aligned counter.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR22. Performance Monitoring Event Coherent_ops May Undercount

Problem: The performance monitoring events Coherent_Ops.RFO (Event: 10h, Umask: 08h) or Coherent_Ops.SPECITOM (Event: 10h, Umask:10h) or Coherent_Ops.WBMTTOI (Event: 10h, Umask:40h) or Coherent_Ops.CLFLUSH (Event: 10h, Umask: 80h) may incorrectly undercount when multiple coherent requests occur simultaneously.

Implication: Due to this erratum, certain Coherent_ops events may undercount.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR23. PCIe* Link Re-Equalization May Not Occur if Link is in L1 State

Problem: Link re-equalization may not occur if the PCIe* link is in L1 state and software attempts to initiate a link re-equalization (Bus:5:0; Device: 8:1; Function:0) LINKCTL3.PE(bit 0, offset: A34h) set, and LINKCTL2.TLS(bit 3:0, offset:70h) set to data rate and LINKCTL.RL(bit 5, offset:50h) set.

Implication: Due to this erratum link re-equalization may not occur and the link will retain the previous equalization coefficients.

Workaround: None Identified. Software may disable ASPM L1 prior to initiating re-equalization (see LINKCTL.ASPMCTL), then re-enable ASPM L1 once done performing re-equalization.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR24. Machine Check Bank 4 UCNA Errors May Not be Signaled

Problem: When any UC error is not enabled in machine check bank 4 due its associated bit being 0 in IA32_MC4_CTL (MSR 410h), and the disabled UC error and a UCNA error happen simultaneously, the UC error will be logged with overflow set, but the UCNA error may not be signaled.

Implication: Due to this erratum, when UC errors are disabled in bank 4, UCNA errors may not be signaled.

Workaround: None identified. Software should keep MCAs enabled in IA32_MC4_CTL.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR25. Intel® DSA / Intel® IAA Use of Priv and PASID

Problem: Intel® Data Streaming Accelerator (Intel® DSA) and Intel® Analytics Accelerator (Intel® IAA) do not support concurrent use of user (Priv=0) and supervisor (Priv=1) privileged operations using the same PASID.

Implication: Due to this erratum, if both user-privileged and supervisor-privileged operations are used with the same PASID, the DSA/IAX behavior is undefined. Intel has not observed this erratum to affect any commercially available software.

Workaround: None identified. Use distinct PASIDs for user-privileged operations and supervisor-privileged operations.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR26. Reserved(0) Check For a PASID Table Entry May Not Happen For a DMA Request

Problem: When a DMA Operation encounters any Reserved(0) bits b[95:91] of a PASID table entry as incorrectly Set, the processor may fail to generate Intel® VT-d fault SPT.3, may incorrectly generate Intel® VT-d fault SPT.4, or fail to block the DMA request.

Implication: Due to this erratum, DMA Request may not behave as expected when encounter Reserved(0) of a PASID table. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR27. Remapping Hardware May Not Generate a Page Request Group Response Message While Operating in Legacy Mode or Abort DMA Mode

Problem: Remapping hardware may not generate a Page Request Group Response Message while operating in Legacy mode or Abort DMA mode if a PCIe* device generates a Page Request Message.

Implication: Due to this erratum, when the remapping hardware fails to generate a Page Request Group Response Message may lead to unpredictable device behavior, including a device hang. The remapping hardware will continue to report RTA.3 or RTA.4 faults if it receives these Page Request Group Response Message. Intel has only observed this behavior in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR28. Remapping Hardware May Abort ZLR to Second-Stage Write Only Pages

Problem: Remapping hardware will report non-recoverable Intel® VT-d fault and cause the Zero-Length-Read (ZLR) to be aborted, If a ZLR encounters read-only page in first-stage tables and write-only page in second-stage tables.

Implication: Due to this erratum, device may observe an unexpected abort on a ZLR and an Intel® VT-d fault may be indicated. Intel has not observed this erratum with any commercially available software.

Workaround: None identified. System software should not create write only pages in second-stage page tables.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR29. Remapping Hardware with Major Version Number 6 Incorrectly Advertises the ESRTPS Support

Problem: Remapping hardware Major Version Number 6 (VER_REG.MAJOR_VERSION_NUMBER= 6, VTDBAR offset 0x0, bits 7:4) enables Enhanced Set Root Table Pointer Support (ESRTPS), but CAP_REG.ESRTPS (VER_REG.ESRTPS, VTDBAR offset 0x8, bit 63) is incorrectly reported as 0.

Implication: Due to this erratum, software may incorrectly determine the ESTRPS feature is not supported.

Workaround: None identified. System software can implement ESRTPS feature if VER_REG.MAJOR_VERSION_NUMBER = 6.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR30. Platform May Hang if System Software Sends a Page Group Response or DevTLB Invalidation to Non-existent Requester ID

Problem: When system software submits a Page Group Response or DevTLB Invalidation command to remapping hardware, the remapping hardware forwards commands to Root-Complex so that the Root-Complex may route the command to Requester ID specified by system software. If system software specifies a Requester ID in the command that does not exist on the platform, the command is not correctly aborted and may cause the system to hang.

Implication: Due to this erratum, if system software issues a Page Group Response or DevTLB Invalidation towards Requestor ID that does not exist on the platform, the system may hang. Intel has only observed this behavior in a synthetic test environment.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR31. Remapping Hardware Does Not Perform Reserved (0) Check in Page Response Descriptor

Problem: Remapping hardware will not set Invalidation Queue Error field in the Fault Status Register (VTDBAR offset 0x34) when software writes non-zero value in bits[255:128] and bit[5] of the Page Response descriptor.

Implication: Due to this erratum, system software violating Intel® VT-d architecture requirement by programming non-zero values in bits[255:128] and bit[5] of Page Response descriptor may not fault on current processors but may fault on future processors. Intel has not observed this sighting/erratum with any commercially available system.

Workaround: None identified. .

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR32. Remapping Hardware Implements b[31:16] of the three Event Data Registers (VTDBAR offsets 0x3C, 0xA4, and 0xE4) as Read-Writable

Problem: b[31:16] of the three Event Data registers (VTDBAR offsets 0x3C, 0xA4, and 0xE4) are "Reserved and Zero" (RsvdZ) but are implemented as Read-Writable (RW).

Implication: Due to this erratum, system software may write these bit[31:16] to non-zero values. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR33. Intel® IAA Do Not Report Overlap Errors For AECS Size of 2 GB or Greater

Problem: Intel® In-Memory Analytics Accelerator (Intel® IAA) does not report overlap errors for the AECS (Analytics Engine Configuration and State) within the source or destination data, when a descriptor is submitted with an AECS Size of 2 GB or greater.

Implication: Due to this erratum, software may not be able to rely on the accuracy of the output data if the output buffer overlaps the AECS buffer. Intel has only observed this behavior in a synthetic test environment.

Workaround: None identified. Software should not use an AECS size greater than or equal to 2 GB.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR34. Intel® DSA / Intel® IAA Invalid TC Not Reported in The SWERROR Register

Problem: A Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA) completion record written using an incorrectly configured Traffic Class (TC) will be written using TC0.

Implication: Due to this erratum, an incorrectly configured TC for the completion record is not reported in the SWERROR register.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR35. Intel® IAA Unaligned Completion Record Address Error is Not Reported in SWERROR Register

Problem: When an Intel® In-Memory Analytics Accelerator (Intel® IAA) descriptor is submitted with an unaligned Completion Record Address, the completion record is written to the aligned address (ignoring the lower address bits). The Status byte at the Completion Record Address specified in the descriptor will be written as 0, making it appear to software that the descriptor never completed.

Implication: Due to this erratum, the unaligned Completion Record Address error is not reported in the SWERROR register and unpredictable system behavior may occur.

Workaround: None identified. Software should use properly aligned Completion Record Addresses.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR36. Intel® UPI Link Not Resetting When L1 Mismatch Occurs Between Local and Remote Sockets

Problem: When Intel® Ultra Path Interconnect (Intel® UPI) L1 low power link state is disabled on the local socket and a remote socket requests L1 entry, a link reset should be initiated. However, the local socket does not initiate a link reset and replies with negative acknowledgment while remaining in current link state.

Implication: Due to this erratum, Intel® UPI link reset does not occur. There are no known functional implications due to this erratum. Intel has only observed this behavior in a synthetic test environment.

Workaround: None identified. System Software should configure all links to the same power state.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR37. Intel® DSA / Intel® IAA May Fail to Log an MDPE Error For Back-to-Back Parity Errors

Problem: For a read completion with Error Poisoned set that is preceded back-to-back by a write with an IOSF data parity error, Intel® Data Streaming Accelerator (Intel® DSA) and Intel® In-Memory Analytics Accelerator (Intel® IAA) may fail to set the Master Data Parity Error (MDPE, bit 8) in PCI Status registers (IAA Bus: system design dependent, Device: 2, Function: 0; Offset: 6h, DSA Bus: system design dependent, Device: 1, Function: 0; Offset: 6h).

Implication: Due to this erratum, the PCI Status MDPE bit may not be set. Software that uses this bit may not function as expected.

Workaround: None Identified. Software should use PCIe* Advanced Error Reporting rather than PCI legacy error logging.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR38. Relaxed Ordering Not Disabled by DEVCTL.ERO bit for Intel® DSA / Intel® IAA Upstream Transactions

Problem: The PCIe* configuration register, Device Control Enable Relaxed Ordering (DEVCTL.ERO, Intel® In-Memory Analytics Accelerator (Intel® IAA) Bus: system design dependent, Device: 2, Function: 0, Offset: 48h, Bit 4; Intel® Data Streaming Accelerator (Intel® DSA) Bus: system design dependent, Device: 1, Function: 0, Offset: 48h, Bit 4) bit does not disable relaxed ordering in Intel® DSA / Intel® IAA for upstream writes.

Implication: Due to this erratum, for peer-to-peer traffic, writes from Intel® DSA / Intel® IAA can show up on a PCIe* link with RO=1 even though DEVCTL.ERO is set to '0'. Intel has not observed any functional issues as a result of this erratum.

Workaround: None identified. Software requiring Strict Ordering can set the Strict Ordering (SO) flag in the descriptors to '1' in order to enforce strict ordering for upstream writes.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR39. System Address Logged For WDB Parity Errors May be Incorrect

Problem: When target XOR enable bit [21] MCMTR.CLUSTER_XOR_ENABLE (MEM0_BAR; Offset: 20EF8h) or channel XOR enable bit [20] MCMTR.CHANNEL_XOR_ENABLE (MEM0_BAR; Offset: 20EF8h) is set or clock gating disable bit [28] [DDRT_CLK_GATING.DIS_REVADDR_LOG_CLKGATING (MEM0_BAR; Offset 21514h)] 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.CLUSTER_XOR_ENABLE = 0) and disabling channel XOR (MCMTR.CHANNEL_XOR_ENABLE = 0).

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR40. Incorrect MCACOD For L2 MCE

Problem: Under complex microarchitectural 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, see the [Summary Tables of Changes](#) on page 13.

SPR41. System May Hang Due to Full LLRB

Problem: The processor may incorrectly fill the Link-Level Retry Buffer (LLRB) with Non-Ack Bearing Flits.

Implication: Due to this erratum, if both the processor and the link partner fill their respective LLRBs with Non-Ack Bearing Flits, the system may hang. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified. This erratum can be mitigated by configuring the LLRB to its maximum size and minimizing link latency.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR42. Intel® IAA May Fail to Properly Decode Data With a Large Header

Problem: If Intel® In-memory Analytics Accelerator (Intel® 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, see the [Summary Tables of Changes](#) on page 13.

SPR43. Memory Controller Violates JEDEC RCD tCSALT Timing

Problem: The processor may violate tCSALT timing (as specified in JEDEC DDR5RCD01 Specification Rev 1.1, Section 8.6.2) by issuing either a Power Down Entry (PDE) or Power Down Exit (PDX) command during the tCSALT window.

Implication: Due to this erratum, Register Clock Drivers that receive PDE/PDX commands during the tCSALT window may not operate as expected.

Workaround: It may be possible to mitigate this issue with a BIOS code change.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR44. Wrong CKE Signal Used on 1 DPC 3DS 4H Configs

Problem: For the specific memory configuration using 1 DIMM per channel (DPC) and 4H 3DS DIMMs, the processor will not correctly assert CKE (Clock Enable) during PPD (Precharge Power Down) mode violating section 4.10.1 of the JEDEC specification revision 1.85.

Implication: Due to this erratum, when the memory subsystem enters PPD mode, the processor may experience uncorrectable memory errors.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR45. Address May Not be Logged For a UCR Error Detected in The MLC

Problem: An Uncorrected No Action Required (UCNA) error logged in machine check bank 3 with MC3_STATUS.MCACOD=0179h (MSR 40Dh, bits 15:0) may not include a valid address in MC3_ADDR (MSR 40Eh) when an ECC Uncorrected Recoverable (UCR) error is detected on an MLC (Mid-level cache) eviction.

Implication: Due to this erratum, the address of the poisoned data produced by the ECC UCR error in the MLC may not be available to software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR46. Intel® VT-d DMA Remapping Hardware May Hang if it Encounters Page Request Queue Overflow Condition

Problem: Intel® VT-d DMA remapping hardware may stop processing new descriptors from the Invalidation Queue and/or stop responding to register reads when it encounters a Page Request Queue overflow condition as indicated by PRS_REG.PRO=1

Implication: Due to this erratum, the system may hang and may not signal a Page Request Queue overflow fault.

Workaround: None identified. Software should size Page Request Queue to avoid overflow condition (PRS_REG.PRO=1). To determine the size of suitable Page Request Queue, software should sum up the value of the Outstanding Page Request Capacity register (Bus: 8-11; Device: 1; Function: 0; Offset: 248h) across all devices where Page Requests are enabled and increment by one. If the result is not a power of 2, then software should round to the nearest higher power of 2. System software that sends Page Response to the device before updating the Page Request Queue Head Register (PQH_REG) will require another doubling of the Page Request Queue Size to avoid overflow condition.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR47. Receiver Common Mode Input Impedance May be Below Specification When Interface is Powered Down

Problem: The processor may fail to meet receiver Common Mode Input Impedance as per PCIe* Specification chapter 8.4.3 when the PCIe* interface is powered down.

Implication: Due to this erratum, link partners may incorrectly detect the processor and initiate link training during platform reset. Intel has not observed any functional issues from this erratum when used with PCIe* compliant link partners.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR48. Remapping Hardware Will Not Report The PASID Value For RTA.2 Faults in Modes Other Than Scalable Mode

Problem: When Remapping Hardware encounters RTA.2 fault condition in modes other than Scalable Mode (RTADDR.TTM==01), the Fault Recording Register (FRCDH_REG_0_0_0_VTD BAR, offset 408h) will incorrectly report a value of 0 in the PASID Present (PP) field (bit 31) and in the PASID Value (PV) field (bits 59:40).

Implication: Due to this erratum, software can not rely on PASID value for RTA.2 faults in modes other than Scalable Mode. Intel has not observed this sighting/erratum to impact the operation of any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR49. Remapping Hardware Does Not Perform a Reserved(0) Check in Interrupt Remap Table Entry

Problem: Remapping hardware does not perform Reserved(0) check on b[127:HAW+64] of the Interrupt Remap Table Entry for a Posted Interrupt.

Implication: Due to this erratum, system software violating Intel® VT-d architecture requirement by programming non-zero reserved values in b[127:HAW+64] of Interrupt Remap Table entry for Posted Interrupt may not fault on current processors but may fault on future processors. Intel has not observed this sighting/erratum with any commercially available system.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR50. Processor PCIe* Root Port Link Spurious Data Parity Error May be Reported

Problem: When the processor's PCIe* root port's link width is down-configured and then subsequently up-configured, the root port may log and report a spurious Local Data Parity Error on the lanes that were disabled and then re-enabled.

Implication: Due to this erratum, the Local data parity error may be observed on PCIe* root port down-configured links in the 16.0 GT/s data parity status registers (Bus 5-0; Device 2; Function 0; Offset 10h/14h/18h). Intel has not observed any functional implications due to this erratum.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR51. 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, see the [Summary Tables of Changes](#) on page 13.

SPR52. 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 = FFFFFFFFh, the UCNA error will be logged and signaled, but will incorrectly set MCI_STATUS.EN. (bit 60). Intel has not observed this erratum to affect any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR53. Reading The PPERF MSR May Not Return Correct Values

Problem: Under complex microarchitectural conditions, a RDMSR instruction to Productive Performance (MSR_PPERF) MSR (Offset 64eh) may not return correct values in the upper 32 bits (EDX register) if Core C6 is enabled.

Implication: Due to this erratum, Software may experience a non-monotonic value when reading the MSR_PPERF multiple times.

Workaround: None identified. Software should not rely on the upper bits of the MSR_PPERF when core C6 is enabled.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR54. No #GP Will be Signaled When Setting MSR_MISC_PWR_MGMT.ENABLE_SDC if MSR_MISC_PWR_MGMT.LOCK is Set

Problem: If the MSR_MISC_PWR_MGMT.LOCK (MSR 1AAh, bit 13) is set, a General Protection Exception (#GP) will not be signaled when MSR_MISC_PWR_MGMT.ENABLE_SDC (MSR 1AAh, bit 10) is cleared while IA32_XSS.HDC (MSR DA0h, bit 13) is set and if IA32_PKERF_HDC_CTL.HDC_PKG_Enable (MSR DB0h, bit 0) was set at least once before.

Implication: Due to this erratum, a #GP will not be signaled even though MSR_MISC_PWR_MGMT.ENABLE_SDC is cleared while the associated LOCK bit is set.

Workaround: None identified. Software should not attempt to clear MSR_MISC_PWR_MGMT.ENABLE_SDC if the above #GP conditions are met.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR55. System May Experience an Internal Timeout Error When an Internal Parity Error Occurs While Working With Intel® AMX

Problem: Under complex microarchitectural conditions, while running Intel® Advanced Matrix Extensions (Intel® AMX), an Internal Parity Error (IA32_MC0_Status (MSR 401n, bits [15:0]) set to 5h) may cause an Internal Timeout Error (IA32_MCI_Status [15:0] set to 400h) in parallel to the reporting of the parity error reporting.

Implication: Due to this erratum, an unexpected Internal Timeout Error may occur.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR56. Last Branch Records May Not Survive Warm Reset

Problem: Last Branch Records (LBRs) are expected to survive warm reset according to Intel® architectures (SDM Vol3 Table 9-2). LBRs may be incorrectly cleared following warm reset if a valid machine check error was logged in one of the IA32_MCI_STATUS MSRs (401h, 405h, 409h, 40Dh).

Implication: Due to this erratum, reading LBRs following warm reset may show zero value even though LBRs were enabled (IA32_LBR_CTL.LBREN[0]=1) before the warm reset.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR57. Single Step on Branches Might be Missed When VMM Enables Notification On VM Exit

Problem: Under complex microarchitectural conditions, "single step on branches" (configured when IA32_DEBUGCTLMR (Offset 1D9h, bit [1]) and TF flag in EFLAGS register are set) while in guest might be missed when VMM enables "notification on VM Exit" (IA32_VMX_PROCBASED_CTL2 MSR, Offset 48Bh, bit [31]) while the dirty access bit is not set for the code page (bit [6] in paging-structure entry).

Implication: Due to this erratum, when "single step on branches" is enabled under the above condition, some single step branches will be missed. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified. When enabling single step on branches for debugging, software should first set the dirty bit of the code page.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR58. Incorrect #CP Error Code on UIRET

Problem: If a #CP exception is triggered during a UIRET instruction execution, the error code on the stack will report NEAR-RET instruction (code 1) instead of FAR-RET instruction (code 2).

Implication: Due to this erratum, an incorrect #CP error code is logged when #CP is triggered during UIRET instruction.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR59. #GP May be Serviced Before an Instruction Breakpoint

Problem: An instruction breakpoint should have the highest priority and needs to be serviced before any other exception. In case an instruction breakpoint is marked on an illegal instruction longer than 15 bytes that starts in bytes 0-16 of a 32B-aligned chunk, and that instruction does not complete within the same 32B-aligned chunk, a General Protection Exception (#GP) on the same instruction will be serviced before the breakpoint exception.

Implication: Due to this erratum, an illegal instruction #GP exception may be serviced before an instruction breakpoint.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR60. Unexpected #PF Exception Might Be Serviced Before a #GP Exception

Problem: Instructions longer than 15 bytes should assert a General Protection Exception (#GP). For instructions longer than 15 bytes, a Page Fault Exception (#PF) from the subsequent page might be issued before the #GP exception in the following cases:

1. The GP instruction starts at byte 1 – 16 of the last 32B-aligned chunk of a page (starting the count at byte 0), and it is not a target of taken jump, and it does not complete within the same 32B-aligned chunk it started in.
2. The GP instruction starts at byte 17 of the last 32B-aligned chunk of a page.

Implication: Due to this erratum, an unexpected #PF exception might be serviced before a #GP exception.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR61. VMX-Preemption Timer May Not Work if Configured With a Value of 1

Problem: Under complex microarchitectural conditions, the VMX-preemption timer may not generate a VM Exit if the VMX-preemption timer value is set to 1.

Implication: Due to this erratum, if the value of the VMX-preemption timer is set to 1, a VM exit may not occur.

Workaround: None identified. Software should avoid programming the VMX-preemption timer with a value of 1.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR62. User Interrupt Might be Delayed

Problem: Under complex microarchitectural conditions, if MOV SS blocking (bit 1 in the guest Interruptibility state) is enabled, when a guest resumes into CPL3 with a user interrupt pending, the awaiting interrupt might be served after the second instruction and not after the first one as expected.

Implication: Due to this erratum, an user interrupt might be delayed. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR63. VM Exit Qualification May Not be Correctly Set on APIC Access While Serving a User Interrupt

Problem: A VM Exit that occurs while the processor is serving a user interrupt in non-root mode should set the “asynchronous to instruction execution” bit in the Exit Qualification field in the Virtual Machine Control Structure (bit 16). However, if a VM Exit occurs during processing a user interrupt due to an APIC access, the bit will not be set.

Implication: Due to this erratum, the “asynchronous to instruction execution” bit will not be set if an APIC Access VM Exit occurs while the processor is serving a user interrupt. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR64. Software Tuning That Relies on PCLS Values May Experience Inaccurate Event Counts

Problem: While monitoring system performance, the processor will incorrectly translate a Prior Cache Line State field value of 0 (no performance detail) in an Intel® UPI data response packet to an one-hop near miss performance event when utilizing an External Node Controller (XNC).

Implication: Due to this erratum, software tuning that relies on correct PCLS values may experience inaccurate Intel® UPI one-hop near miss event counts.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR65. Multiple SGX_Doorbell_Errors on Ubox Response Mismatch

Problem: In the event of a mismatch between Intel® UPI LT_Doorbell Response completion and SGX_Secure_En configuration bit in Ubox, SGX_Doorbell_Errors may overflow the NCEVENTS_CR_UBOX_MCI_STATUS (MCA Bank 6, MSR 419h) register and signal a redundant Machine Check Exception (MCE) with MSCOD 801Ch and MCACOD of 0407h to the cores and PUNIT.

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

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR66. ECS Readout Fails on Mixed Mode Systems

Problem: In platform configurations utilizing both DDR5 and DDRT2 memory technologies on the same channel, accesses to Error Check and Scrub (ECS) data may cause the system to hang.

Implication: Due to this erratum, the system may hang

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR67. Intel® DSA / Intel® IAA Completion Record is Not Written For Non-Completion Record Invalid Traffic Classes

Problem: For Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA), when any Traffic Class (TC) selected by a descriptor is invalid, the completion record is not written and the error is reported in SWERROR.

Implication: Due to this erratum, software that expects a completion record may not function as expected.

Workaround: None Identified. Software should check the status of the SWERROR register if it does not receive a completion record as expected.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR68. Intel® IAA Expand Operation With PRLE Format Input May Return an Error

Problem: Intel® In-Memory Analytics Accelerator (Intel® IAA) Expand operation may parse beyond the required elements of Source 1 and return a Parquet Run Length Encoding (PRLE) Format Error (14h) unexpectedly.

Implication: Due to this erratum, software may receive a spurious error.

Workaround: None Identified. Software should not send non-PRLE encoded stream data in Source 1.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR69. Intel® IAA Compression with Compress Bit Order Set May Produce an Odd Number of Bytes

Problem: When a compress job specifies "Compress Bit Order" flag and not "Stats Mode", the Intel® In-Memory Analytics Accelerator (Intel® IAA) may incorrectly produce an odd number of bytes.

Implication: Due to this erratum, when this occurs, the end of the generated bit-stream is lost.

Workaround: None Identified. Use the Intel® Query Processing Library (Intel® QPL) to handle this case.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR70. Intel® IAA Source 2 Not Written Properly When Source 2 Size is 32 Bytes

Problem: For Intel® In-Memory Analytics Accelerator (Intel® IAA) operations, If the Source 2 size is specified as 32 bytes and if Source 2 is being written, then no Source 2 data will be written.

Implication: Due to this erratum, software that writes 32 bytes of Source 2 data may not function as expected.

Workaround: None Identified. When Source 2 is being written, software should specify Source 2 size to be at least 64 bytes.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR71. Intel® IAA May Not Report Invalid Filter Flags Status Code When Source 2 Bit Order Field is Set

Problem: Intel® In-Memory Analytics Accelerator (Intel® IAA) may not report an error when Source 2 Bit Order field of Filter Flag is erroneously set.

Implication: Due to this erratum, software may not receive an error when expected. Intel has not observed any functional implication as a result of this erratum.

Workaround: None Identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR72. Intel® IAA Does Not Allow Source 1 Size to be 0 For Expand Operation

Problem: Intel® In-Memory Analytics Accelerator (Intel® IAA) Expand operation will incorrectly return an error code when Source 1 size is 0.

Implication: Due to this erratum, software that uses the Expand operation with Source 1 size of 0 may not behave as expected.

Workaround: None Identified. For the Expand operation, software can workaround this issue by supplying Intel® IAA with a dummy Source 1 input that contains at least one byte.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR73. Intel® DSA / Intel® IAA And WQ Configuration Registers May be Incorrectly Updated

Problem: Software may incorrectly update Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA) and Work Queue (WQ) configuration registers when the device state has changed from "Enabled" to "Disable-in-Progress."

Implication: Due to this erratum, unpredictable DSA/IAA device behavior may occur.

Workaround: None Identified. Software should not change Intel® DSA / Intel® IAA device WQ configuration registers until CMDSTATUS Register (offset A8h) bit 31=0.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR74. Invalid Flags Field of The Completion Record May Not be Set Correctly For Intel® IAA Compression Operation

Problem: In the Intel® In-Memory Analytics Accelerator (Intel® IAA) Compress descriptor, if the compression flag Stats Mode = 0 and Read Source 2 flag = 0, then the Intel® IAA Compress operation will return an Invalid Operation Status Code 11h but not set the Invalid Flags field.

Implication: Due to this erratum, software can not tell which flags are invalid based on Invalid Flags field.

Workaround: None Identified. User may examine the descriptor and the documentation to determine which flags are invalid.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR75. With Intel® SGX Disabled, Software That Relies on ENCLVexiting May Not Function as Expected

Problem: On processors with Intel® SGX disabled, the enable ENCLVexiting bit 60 of IA32_VMX_PROCBASED_CTL2 MSR (index 48BH) is incorrectly set as being enabled.

Implication: Due to this erratum, software that relies on the ENCLVexiting bit may not function as expected.

Workaround: None identified. Software should not rely on the ENCLVexiting bit when Intel® SGX is disabled.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR76. Headers Logged in AERHDRLOG for an AER Error for Intel® DSA / Intel® IAA may be Incorrect

Problem: The header log (AERHDRLOG(1-4), [Bus: system design dependent, Device: 1, Function: 0, Offset: 11Ch, 120h, 124h, 128h respectively]) for Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA), will be overwritten as all 1's when a simultaneous and independent Correctable Error occurs.

Implication: Due to this erratum, software relying upon AERHDRLOG (1-4) for an Advanced Error Reporting error may not function as expected.

Workaround: None Identified. Software should ignore AERHDRLOG (1-4) when its value is all 1's.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR77. Intel® DSA / Intel® IAA May Fail to Send an ERR_FATAL Message if a Non-Fatal Error Occurs in The Same Cycle

Problem: When fatal and non-fatal uncorrectable errors occur in the same cycle, an ERR_FATAL message is not sent and only an ERR_NONFATAL message is sent from the Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA).

Implication: Due to this erratum, software may not function as expected due to a fatal error not being reported.

Workaround: None Identified. Software should check for fatal errors during the handling of non-fatal errors.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR78. Intel® DSA / Intel® IAA May Fail to Log an Unexpected Completion Error For an Invalid ATS Response

Problem: Under complex microarchitecture conditions when there are multiple simultaneous errors, Intel® Data Streaming Accelerator (Intel® DSA)/ Intel® In-Memory Analytics Accelerator (Intel® IAA) may fail to log an Unexpected Completion error for an Address Translation Services (ATS) response with an incorrect PASID Privilege Mode Requested value.

Implication: Due to this erratum, when there are multiple simultaneous errors, software may be unaware of an ATS Unexpected Completion error.

Workaround: None Identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR79. Intel® IAA Compression Output Buffer Overflow Error May be Incorrectly Reported

Problem: For Intel® In-Memory Analytics Accelerator (Intel® IAA), when checking for compression output overflow, the upper bits (31: 29) of the Maximum Destination Size descriptor field (Offset 48 - 51) will be discarded.

Implication: Due to this erratum, an output buffer overflow error may be incorrectly reported for buffer sizes greater than or equal to 2000_0000h.

Workaround: None Identified. Software should avoid Maximum Destination Size greater than or equal to 2000_0000h.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR80. Intel® QuickAssist Technology Accelerator May Violate ATS Invalidation Completion Ordering

Problem: Address Translation Service (ATS) invalidations may complete before all in-flight writes are drained from Intel® QuickAssist Technology (Intel® QAT) accelerator.

Implication: Due to this erratum, Intel® QAT accelerator operation with ATS capability enabled may lead to unexpected system behavior.

Workaround: System software (OS/VMM) performing ATS invalidation on Intel® QAT accelerator needs to serially execute a second (duplicate) ATS invalidation request after the first invalidation completes to drain in-flight writes.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR81. Intel® QuickAssist Technology Accelerator Device May Not Invalidate PASID Supervisor-Privilege Translations

Problem: Address Translation Service (ATS) invalidations for Process Address Space ID (PASID) with Supervisor-privilege translations may not correctly invalidate the device TLB on Intel® QuickAssist accelerator (Intel® QAT).

Implication: Due to this erratum, Intel® QAT accelerator operation with ATS capability enabled and Supervisor-privilege PASID may lead to unexpected system behavior.

Workaround: System software (OS/VMM) performing ATS invalidation on Intel® QAT accelerator on behalf of any supervisor-privilege PASID must set the Global Invalidate (G) bit in the ATS invalidation to avoid the erratum.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR82. The Time-Stamp Counter May Report an Incorrect Value

Problem: Under complex microarchitectural conditions, the Time-Stamp Counter (TSC) may incorrectly report the time stamp to be less than the expected time stamp after exiting C6 power saving state.

Implication: Due to this erratum, systems that rely upon a monotonically increasing value reported by the TSC may exhibit unpredictable system behavior.

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](#) on page 13.

SPR83. Intel® UPI Machine Check Bank May Not Report The Most Recently Logged Error

Problem: If multiple Intel® UPI ports log corrected errors and KTI_MCA_CFG LATCH_FIRST_CE is not set for one or more UPI links (Bus: 30; Device: 1-4; Function: 1; Offset 498h; bit: 0=0), the UPI Machine Check Bank (Bank 5; MSRs: 415h - 417h) may not report the most recently logged error.

Implication: Due to this erratum, software can not rely on the most recent error being logged in the Intel® UPI Machine Check Bank.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR84. PECCI Wire Host may Continuously Receive a Completion Code of 0x80

Problem: Regardless of targeted PECCI endpoint, if a PECCI wire host issues a PECCI transaction within one second of a previous PECCI transaction that received a completion code of 0x80 (command response timeout), it may continuously receive a command response timeout.

Implication: Due to this erratum, a PECCI endpoint may be perceived as unresponsive.

Workaround: None identified. A PECCI wire host must wait for at least 1 second if a previous request had timed out before sending command to a different endpoint.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR85. DDR5 9x4 DIMMs ECS Data May be Reported Incorrectly

Problem: For DDR5 9x4 DIMMs, after the memory controller issues a Movable Read Reference (MRR) to device 8, the Error Correctable String (ECS) data will be reported incorrectly in `mr_read_result` (`MEM_BAR [0-3]`, Offsets 22C80h-22C90h or 2AC80h-2AC90h).

Implication: Due to this erratum, the software cannot rely on ECS data for Device 8 with DDR5 9x4 DIMMs. DDR5 10x4 or 5x8 DIMMs are not affected by this erratum.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR86. RETRY_RD_ERR_LOG_MISC.DDR5_9x4_half_device Bit Maybe Incorrect

Problem: On systems using 9x4 DDR5 DIMMs, when Permanent Fault Detection (PFD) is disabled, the `RETRY_RD_ERR_LOG_MISC.DDR5_9x4_half_device` bit (138_MEM_RRD + Offsets 22C54h, 22D80h, 2AC54h, 2AD80h, 22E60h, Bit 7) will always report 0 when an error is detected in device 8.

Implication: Due to this erratum, when an error is detected on device 8, the system software is not able to rely on the value of the `RETRY_RD_ERR_LOG_MISC.DDR5_9x4_half_device` bit.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR87. PIROM Reports The Wrong 2 DPC Speed For Processors With Less Than 4800 MT/s 1 DPC Speed

Problem: PIROM incorrectly reports the 2 DIMM Per Channel (DPC) speed as 400 MT/s less than the top speed on processors with a 1DPC speed of less than 4800 MT/s.

Implication: Due to this erratum, when an error is detected on device 8, the system software is not able to rely on the value of the `RETRY_RD_ERR_LOG_MISC.DDR5_9x4_half_device` bit.

Workaround: None identified. Software should assume that the 2 DPC speed for any processors with 1 DPC speed less than 4800 MT/s is equal to the 1 DPC speed.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR88. An MDF Parity Error May Incorrectly Set The Overflow Bit

Problem: The Overflow bit (bit 62) of IA32_MC[7-8]_STATUS MSRs (41Dh, 421h) may be incorrectly set when a Modular Die Fabric (MDF) parity error occurs (MCACOD = 0405h).

Implication: Due to this erratum, software that relies upon the Machine Check Overflow bit may not operate as expected.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR89. Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

Problem: During RTM (Restricted Transactional Memory) operation when branch tracing is enabled using BTM (Branch Trace Message) or BTS (Branch Trace Store), 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, see the [Summary Tables of Changes](#) on page 13.

SPR90. CPUID Reports Incorrect Number of Ways For The Load DTLB

Problem: CPUID leaf 18H sub-leaf 04H EBX [31:16] reports 4 ways instead of 6 ways for the Load DTLB.

Implication: Due to this erratum, software that relies upon the number of ways in the load DTLB may operate sub optimally.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR91. 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® Processor Trace (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, a 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, see the [Summary Tables of Changes](#) on page 13.

SPR92. On Instructions Longer Than 15 Bytes, #GP Exception is Prioritized And Delivered Over #CP Exception

Problem: A #GP (global protection exception) that results from an instruction being longer than 15 bytes is prioritized and served before a #CP (Controlflow Protection exception) that was created due to a missing ENDBRx instruction at the target of an indirect branch.

Implication: Due to this erratum, during an indirect jump with ENDBRANCH tracking, if the processor lands on an illegal instruction with length longer than 15 bytes or that decodes to a CS limit, the processor will prioritize and deliver a #GP exception over the #CP exception.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR93. Mismatch on DR6 Value When Breakpoint Match is on Bitmap Address

Problem: Under complex microarchitectural conditions, on systems with Control-flow Enforcement Technology (CET) enabled, hitting a predefined data breakpoint may not be reported in B0-B3 (bits 3:0) in the DR6 register if that breakpoint was set on the legacy code page bitmap.

Implication: Due to this erratum, software may not know which breakpoint triggered when setting breakpoints on the legacy code page bitmap.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR94. RTM Abort Status May be Incorrect For INT1/INT3 Instructions

Problem: When Intel® Transactional Synchronization Extensions (Intel® TSX) is enabled, and there is an Restricted Transactional Memory (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, see the [Summary Tables of Changes](#) on page 13.

SPR95. WRMSR to Reserved Bits of IA32_L3_QOS_Mask_15 Will Not signal a #GP

Problem: A General Protection Exception (#GP) will not be signaled when writing non-zero values to the upper 32 bits of IA32_L3_QOS_Mask_15 MSR (Offset C9FH) even though they are defined as reserved bits.

Implication: Due to this erratum, a #GP will not be signaled when the upper bits of IA32_L3_QOS_Mask_15 are written with a non-zero value.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR96. 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, see the [Summary Tables of Changes](#) on page 13.

SPR97. 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: Due to this erratum, 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: None identified. Software should explicitly write to the stack pointer immediately after executing MOV SS or POP SS.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR98. Exit Qualification For EPT Violations on Instruction Fetches May Incorrectly Indicate That The Guest-physical Address Was Writeable

Problem: On EPT violations, bit 4 of the Exit Qualification indicates whether the guest-physical address was writeable. When EPT is configured as supervisory shadow-stack (both bit 60 in EPT paging-structure leaf entry and bit 0 in EPT paging-structure entries are set), non-executable (bit 2 in EPT paging-structure entries is cleared), and non-writeable (bit 1 in EPT paging-structure entries is cleared) a VMExit due to a guest instruction fetch to a supervisory page will incorrectly set bit 4 of the Exit Qualification. Bits 3, 5, and 6 of the Exit Qualification are not impacted by this erratum.

Implication: Due to this erratum, bit 4 of the Exit Qualification may be incorrectly set. Intel has not observed this erratum on any commercially available software.

Workaround: None identified. EPT handlers processing an EPT violation due to an instruction fetch access on a present page should ignore the value of bit 4 of the Exit Qualification.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR99. Processor May Generate Spurious Page Faults On Shadow Stack Pages

Problem: When operating in a virtualized environment, if shadow stack pages are mapped over an APIC page, the processor will generate spurious page faults on that shadow stack page whenever its linear to physical address translation is cached in the Translation Look-aside Buffer.

Implication: When this erratum occurs, the processor will generate a spurious page fault. Intel is not aware of any software that maps shadow stack pages over an APIC page.

Workaround: None identified. Software should avoid mapping shadow stack pages over the APIC page.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR100. 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, see the [Summary Tables of Changes](#) on page 13.

SPR101. IA32_MC1_STATUS MSR May Not Log Errors When IA32_MC1_CTL MSR is Set to Not Signal Errors

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

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

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR102. 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, see the [Summary Tables of Changes](#) on page 13.

SPR103. OFFCORE_REQUESTS_OUTSTANDING Performance Monitoring Events May be Inaccurate

Problem: The OFFCORE_REQUESTS_OUTSTANDING.*DATA_RD performance monitoring events (Event 20h; UMask 08h) counts the number of off-core outstanding data read transactions each cycle. Due to this erratum, an inaccurate count may be observed when Intel® Hyper-Threading Technology (Intel® HT Technology) is enabled and hardware prefetchers are enabled.

Implication: Due to this erratum, OFFCORE_REQUESTS_OUTSTANDING Performance Monitoring Events may be Inaccurate.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR104. 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, see the [Summary Tables of Changes](#) on page 13.

SPR105. Call Instruction Wrapping Around The 32-bit Address Boundary May Return to Incorrect Address

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

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

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR106. ADDDC Reverse Sparing May Lead to Incorrect Data

Problem: When reverse sparing an ADDDC (Adaptive Double Device Data Correction) region to an SDDC (Single Device Data Correction) 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 BIOS to contain a workaround for this Erratum.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR107. SGX ENCLU[EACCEPT] Will Not Cause #GP When TCS.PREVSP is Non-zero

Problem: For processors that support CET (CPUID.(EAX=07H, ECX=0H):ECX[CET_SS), SGX ENCLU[EACCEPT] does not signal a #GP exception when PREVSP field in a Thread Control Structure (TCS) is non-zero.

Implication: Due to this erratum, ENCLU[EACCEPT] will not cause #GP when TCS.PREVSP is non-zero. Intel has not observed any functional implications due to this erratum.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR108. Setting PREFETCH_CTL.DISABLE_THREE_STRIKE_CNT Does Not Prevent The Three-strike Counter From Incrementing

Problem: Setting PREFETCH_CTL.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 PREFETCH_CTL.DISABLE_THREE_STRIKE_CNT is set.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR109. 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, see the [Summary Tables of Changes](#) on page 13.

SPR110. Intel® 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 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, see the [Summary Tables of Changes](#) on page 13.

SPR111. 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 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: Due to this erratum, 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, see the [Summary Tables of Changes](#) on page 13.

SPR112. In 96bit ECC Mode, a Correctable Error May be Incorrectly Logged in RETRY_RD_ERR_LOG With UC Bit Set

Problem: In 96bit ECC mode and when DUE2CE is enabled (LINK_CFG_READ_1.READ_DATA bit (1310_MEM_WP + Offsets 2B458h, 23458h, Bit 26=1)), when a Correctable Error is detected during Adaptive Double Device Data Correction (ADDDC) forward sparing RETRY_RD_ERR_LOG.UC bit (138_MEM_RRD + Offsets 22C60h, 22E54h, Bit 1) may be incorrectly set.

Implication: Due to this erratum, a Correctable Error may incorrectly set the Uncorrected Error (UC) bit in RETRY_RD_ERR_LOG.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR113. System Crash Observed on Host When TD Private Pages Are Not Zeroed Out if Reused For Non-TD SW

Problem: Partial write access to trusted Intel® TDX pages by untrusted software may generate Machine Check Exception with a Data Load or Instruction Fetch SRAR Error.

Implication: Due to this erratum, the system may report SRAR Machine Check Exceptions. This erratum may occur following a Fast Warm Reset event.

Workaround: It may be possible to mitigate this erratum with a combination of system software and a BIOS code change.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR114. IA32_MC2_ADDR And IA32_MC2_MISC MSRs Will be Cleared on Warm Reset

Problem: A non-zero value written to IA32_MC2_ADDR (40Ah) and IA32_MC2_MISC(40Bh) MSRs will be incorrectly cleared following a warm reset.

Implication: Due to this erratum, software that relies on the IA32_MC2_ADDR and IA32_MC2_MISC MSR values may not function correctly after a warm reset. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR115. Multiple Write CRC Errors May Lead to System Hang

Problem: If multiple DRAM Write CRC errors occur, the system may hang with a Mesh to Memory timeout Machine Check Exception reported in IA32_MC12_STATUS (MSR 431h) with MSCOD = 0009h and MCACOD=0400h or a TOR timeout Machine Check Exception reported in IA32_MC9_STATUS, IA32_MC10_STATUS, or IA32_MC11_STATUS (MSRs 425h, 429h, 42Dh) with MSCOD = 000Ch.

Implication: Due to this erratum, the system may hang due to multiple Write CRC errors. Write CRC is a debug only feature and Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR116. Intel® DSA / Intel® IAA CTO Errors May Inconsistently Update the Prefix Log and Prefix Log Present Flag

Problem: For Intel® Data Streaming Accelerator (Intel® DSA) / Intel® In-Memory Analytics Accelerator (Intel® IAA), Completion Timeouts (CTO) will incorrectly set TLP Prefix Log Present (AERCAPCTL.TLPPLP, IAA [Bus: system design dependent, Device: 2, Function: 0, Offset: 118h, Bit 11;] DSA [Bus: system design dependent, Device: 1, Function: 0, Offset: 118h, Bit 11]) bit regardless of whether the offending read had a Process Address Space ID (PASID) TLP prefix. Additionally, the value of the PASID TLP prefix will be stored as all 0's rather than all 1's. The header is correctly stored as all 1's.

Implication: Due to this erratum, software relying upon Advanced Error Reporting (AER) error log AERCAPCTL.TLPPLP for CTO may not function as expected.

Workaround: None identified. Software should treat the AERCAPCTL.TLPPLP as 0 if a CTO error is logged and ignore the associated logged PASID TLP prefix value.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR117. RTIT_CTL.TRACE_EN May be Disabled at BIOS_DONE Even if it Was Previously Enabled

Problem: Under complex microarchitectural conditions, RTIT_CTL.TRACE_EN (bit 0 in MSR 570h) may be disabled when BIOS_UPDT_TRIG.BIOS_DONE is set (MSR 79h bit 0).

Implication: Due to this erratum, Intel® Processor Trace (Intel® PT) may be disabled at BIOS_DONE.

Workaround: None identified. Software that uses PT should ensure it is enabled after BIOS_DONE.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR118. WRMSR to a Few Core MSR's Might be Overwritten

Problem: If any thread is in thread C6 while another thread is updating one of the following MSR's, a subsequent transition from single thread operation to multi-thread operation or vice versa might cause that MSR to revert to its previous value. The affected MSR's 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 MSR's 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, see the [Summary Tables of Changes](#) on page 13.

SPR119. Virtualize IA32_SPEC_CTRL VM-execution Control Does Not Properly Virtualize Bits [63:32] of IA32_SPEC_CTRL

Problem: When the "virtualize IA32_SPEC_CTRL" VM-execution control is enabled, reading IA32_SPEC_CTRL MSR (48h) will return 0 in EDX instead of a value from the IA32_SPEC_CTRL shadow. Similarly, writes to this register will clear bits [63:32] of the IA32_SPEC_CTRL shadow instead of loading them from EDX.

Implication: Due to this erratum, guest software that relies on setting bits [63:32] of IA32_SPEC_CTRL may not function as expected. VMMs that do not enable "virtualize IA32_SPEC_CTRL", or that do not place non-zero values in the upper 32 bits of the IA32_SPEC_CTRL mask or IA32_SPEC_CTRL shadow, are not affected by this erratum. Intel has not observed this erratum with any commercially available software.

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](#) on page 13.

SPR120. Performance Monitoring Events used by TMA May be Inaccurate

Problem: The performance monitoring events TOPDOWN.BACKEND_BOUND_SLOTS (Event A4h UMask 02h) and IDQ_BUBBLES (Event 9Ch) used by the Top-down Microarchitecture Analysis (TMA) method may be inaccurate. Classification of time spent in C0.2 state may be inaccurate when the processor enters the improved power/performance optimized state via TPAUSE or UMWAIT instructions. This erratum also impacts the accuracy of MSR_PERF_METRICS fields Frontend Bound, Backend Bound, and Fetch Latency (MSR 329h, Bits [23:16], [31:24] and [55:48]).

Implication: Due to this erratum, Performance Monitoring events may report inaccurate counts.

Workaround: A workaround is possible using the CPU_CLK_UNHALTED.C02 performance monitoring event. see TMA metrics at github.com/intel/perfmon for details.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR121. Performance Monitoring Event IDQ.MS_UOPS May Undercount

Problem: The performance monitoring event IDQ.MS_UOPS (Event 79H; UMask 20H) and derived events such as IDQ.MS_SWITCHES, IDQ.MS_CYCLES_ANY may undercount MS_UOPS that come from the Decode Stream Buffer (DSB).

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

Workaround: None identified. Performance monitoring event UOPS_RETIRED.MS may be used instead.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR122. CHA TOR Timeout May Occur

Problem: Under complex microarchitectural conditions, the processor may experience a Caching Home Agent (CHA) Table Of Requests (TOR) timeout machine check exception (IA32_MC9_STATUS, IA32_MC10_STATUS, or IA32_MC11_STATUS (MSRs 425h, 429h, 42Dh) with MSCOD = 000Ch).

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

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR123. ECS Reads May Fail to Complete in Certain Memory Configurations and Conditions

Problem: Error Check and Scrub (ECS) transactions may not complete when 2X refresh mode is active in a 2 DIMM Per Channel (2 DPC), 8 ranks and All Bank Fine Mode memory configuration.

Implication: Due to this erratum, the software which relies on the error count values in Error Check and Scrub Mode Register Reads (MRRs) may not function as expected.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR124. A Write to The TSC_Deadline MSR May Cause an Unexpected Timer Interrupt

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

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

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR125. Processor Performance Monitor Discovery Table Incorrectly Reports Data for Intel® UPI Port 3 on MCC SKUs

Problem: On Medium Core Count (MCC) Processors without Intel® UPI Port 3 (Device 4) of Phy registers in Bus: 30, Device 1,2,3,4, Function: 2, the Performance Monitor discovery mechanism will duplicate Intel® Port 2 (Device 3) details to the Intel® UPI Port 3 entry in the Discovery Table.

Implication: Due to this erratum, software that relies upon Intel® UPI port details may not operate as expected.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR126. Certain PCIe Devices Are Not Accessible Out-of-Band With Local Bus Number

Problem: Work descriptors with an Invalid Opcode and unaligned completion address may cause Intel® DSA and Intel® IAA devices to incorrectly write to memory and hang.

When using Out-of-Band commands RdEndpointConfig() - Local PCI Cfg or WrEndpointConfig() - Local PCI Cfg , the following devices are not accessible:

On systems using PCIe Port 2 on XCC or HBM SKUs:

Port A' PCIeGen4 Registers [Bus: 3; Device: 8; Function: 0]

Non-Transparent Bridge (NTB) Registers [Bus: 3; Device: 9; Function: 0]

On systems using PCIe Port 0 on MCC SKUs:

Port A' PCIeGen4 Registers [Bus: 0; Device: 8; Function: 0]

NTB Registers [Bus: 0; Device: 9; Function: 0]

Implication: Due to this erratum, an Unknown/Invalid/Illegal Request completion code of 90h will be returned.

Workaround: None identified. Software should use RdEndpointConfig() - PCI Cfg and WrEndpointConfig() - PCI Cfg commands to access the following PCI configuration space registers, where Bus (x) denotes the ROOTBUSSx field programmed in the CPUBUSNOx registers following BIOS enumeration.

On systems using PCIe Port 2 on XCC or HBM SKUs:

Port A' PCIeGen4 Registers [Bus: 3; Device: 8; Function: 0] NTB Registers [Bus: 3; Device: 9; Function: 0]

On systems using PCIe Port 0 MCC SKUs:

Port A' PCIeGen4 Registers [Bus: 0; Device: 8; Function: 0] NTB Registers [Bus: 0; Device: 9; Function: 0]

Software should use RdEndpointConfig() - MMIO and WrEndpointConfig() - MMIO commands to access MMIO registers for these devices, where Bus (x) denotes the ROOTBUSSx field programmed in the CPUBUSNOx registers following BIOS enumeration.

On systems using PCIe Port 2 on XCC or HBM SKUs:

Port A' PCIeGen4 [Segment 255; Bus: 30; Device: 1; Function: 7]

NTB Registers [Segment 255; Bus: 30; Device: 1; Function: 6]

On systems using PCIe Port 0 on MCC SKUs:

Port A' PCIeGen4 [Segment 255; Bus: 30; Device: 0; Function: 7]

NTB Registers [Segment 255; Bus: 30; Device: 0; Function: 6]

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR127. Remapping Hardware May Encounter Incorrect Error Code in Invalidation Queue Error Record Register

Problem: When fetching a new descriptor from the Invalidation Queue, if Intel® DMA remapping hardware observes an unsupported value in the Translation Table Mode (TTM) field, it may report an invalid descriptor width programmed in the Invalidation Queue (code 5) instead of invalid value in the TTM field of the Root Table Address (code 7) in the Invalidation Queue Error Info (IQEI) register of the IQERCD_REG (VTDBAR offset 0xB0).

Implication: Due to this erratum, Software that distinguishes between error code 5 and error code 7 may not function as expected.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR128. Disabling The APIC While an Interrupt is Being Delivered May Cause a System Hang

Problem: If software disables the APIC by clearing APIC global enable flag (bit 11) in IA32_APIC_BASE (MSR 1Bh) while an interrupt is being delivered, the system may hang with a machine check exception 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 system may hang. Intel has not observed this erratum in any commercial available software.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR129. RTTO May Occur at Lower Speed And Reduce Link Width

Problem: When a 32 GT/s, x16 PCIe port is configured to operate in lower speed and reduced link width (such as 2.5 GT/s, x1 mode), Data Link Layer Packets (DLLPs), including transaction ACK packets, may incur large latencies.

Implication: Due to this erratum, large latencies may lead to Replay Timer Timeout (RTTO) failures from the link partner.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR130. Unexpected System Software Behavior on Systems With (Intel® TH) Enabled

Problem: On systems with Intel® Hyper-Threading Technology (Intel® HT) enabled, instruction execution may become delayed on one logical thread due to long latency instructions and certain microarchitectural conditions generated by the peer logical thread on the same core.

Implication: Due to this erratum, system software may not behave as expected.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR131. Intel® DSA DIF Insert and DIF Strip Operations Do Not Properly Check For Overlapping Buffers

Problem: For Intel® Data Streaming Accelerator (Intel® DSA) Data Integrity Field (DIF) Insert and DIF Strip operations, when performing the check for buffer overlap, Transfer Size (as described in Appendix B of the Intel® Data Streaming Accelerator Architecture Specification) in the descriptor is used instead of the actual destination buffer size.

Implication: Due to this erratum, in situations where the destination buffer immediately precedes the source buffer, DIF Insert may overwrite a portion of the source buffer, instead of reporting an overlapping buffers error (16h) in the Status field of the Completion Record and DIF Strip may incorrectly report an overlapping buffers error prior to the operation.

Workaround: None Identified. For DIF Insert operations, if the destination buffer immediately precedes the source buffer, software must ensure that $N \times 8$ bytes (where N is the number of blocks in the source data) are added to the source buffer when calculating the Destination Address in the descriptor to avoid overwriting the source buffer. For DIF Strip operations, if the destination buffer immediately precedes the source buffer, software must ensure there are at least $N \times 8$ bytes between the end of the destination buffer and the start of the source buffer to avoid generating an overlapping buffers error.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR132. PCIe Link Degradation Following a Leaky Bucket Event

Problem: Ring 0 software does not have the ability to clear the Leaky Bucket Event Disable LBEGxDIS (bits 31, 30, 29, 28) in the Leaky Bucket Action registers for all PCIe ports (Bus: 5-0; Device: 8-1; Function: 0; Offset: DC4h) after a link downgrade due to a Leaky bucket event.

Implication: Due to this erratum, Ring 0 software may not recover the link to full speed after link degradation following a Leaky Bucket event.

Workaround: None identified. Software can avoid this erratum by using PECI commands or SMM to recover the link speed.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR133. Intel® DSA And Intel® IAA May Cause Unpredictable System Behavior When Incorrectly Forwarding Data

Problem: Hardware logic may incorrectly forward data to either an Intel® Data Streaming Accelerator (Intel® DSA) or Intel® Analytics Accelerator (Intel® IAA) device which may cause unpredictable system behavior.

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

Workaround: None identified. A mitigation for this erratum may be available by limiting access to Intel® DSA and Intel® IAA devices to only software with the highest privilege level.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR134. MCA not Logged for WbNoInvd Data ECC Errors

Problem: When WbNoInvd instruction results in an ECC Data Error being detected by the CHA, the error will neither be logged nor signaled.

Implication: Due to this erratum, software that monitors corrected errors may observe fewer corrections than have actually occurred. For uncorrected errors, on systems with Poison enabled, software will not see the uncorrected error reported, but the data will be Poisoned, and subsequent accesses may report an error. On systems with Poison disabled, uncorrected errors may not be reported leading to unpredictable system behavior.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR135. Intel® DSA Memory Write with Incorrect Parity May Result in a System Crash

Problem: When executing a copy operation, if an Intel® Data Streaming Accelerator (Intel® DSA) device receives a poisoned data response to a memory read request (Bus 8:11; Device: 1; Function: 0; Offset 104h, ERRUNCSTS.PTLP, bit[12]), an associated destination memory write with incorrect parity may be generated.

Implication: Due to this erratum, the memory write with incorrect parity may result in a machine check error leading to a system crash.

Workaround: None Identified. It may be possible for the BIOS to contain a mitigation for this erratum.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR136. Internal Timeout MCE May Occur on Platforms with CXL I/O Devices

Problem: Platforms that include CXL I/O devices may hang with an Internal Timeout Machine Check Exception (IA32_MCi_STATUS.MCACOD=400h and IA32_MCi_STATUS.MSCOD=09h) when entering Package C6 state.

Implication: Due to this erratum, the platform may unexpectedly hang.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR137. DIMMs May Not be Placed in Self-Refresh Mode Prior to Warm Reset

Problem: When multiple ECC errors and/or Command and Parity (CAP) errors are being processed by the Integrated Memory Controller (IMC) as a warm reset occurs on the system, the IMC may not place DIMMs in self-refresh mode.

Implication: Due to this erratum, the system may hang on the next boot cycle with a general protection fault (MEM0_BAR + 23430h, bits 31:16 = 00a0h). A cold reset may be required to recover the system.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR138. Communication Error With Intel® On Demand Mailbox

Problem: Intel On Demand provisioning internal states may be lost following an OS runtime microcode update.

Implication: Due to this erratum, an Intel® Out-of-Band Management Services Module (Intel® OOBM Services Module) communication error with the Intel® On Demand mailbox will occur.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR139. Intel® DSA Device May Hang

Problem: Under complex microarchitecture conditions, submitting a Create Delta (07h) descriptor to an Intel® DSA device may cause the device to hang.

Implication: Due to this erratum, an Intel® DSA device may not complete the Create Delta operation and hang.

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR140. A PCIe Device May Not Function After a Warm Reset

Problem: If bit 0 of the PCIe Downstream Port Containment Trigger Status (DPCTS) register (Bus: 5-0; Device: 8-0; Function: 0; Offset 1A8h) is set before a warm reset, PCIe devices connected to the associated port may not be able to enumerate following the warm reset.

Implication: Due to this erratum, PCIe devices connected to the associated port may not function as expected.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR141. VM Exit Following MOV to CR8 Instruction May Lead to Unexpected IDT Vectoring-Information

Problem: Under certain conditions, a VM exit following execution of the MOV to CR8 instruction may unexpectedly result in setting the Valid bit (bit 31) of the IDT-Vectoring Information Field in the Virtual Machine Control Structure (VMCS).

Implication: Depending on the operation of the virtual-machine monitor (VMM), this may result in unexpected VM behavior.

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](#) on page 13.

SPR142. System May Hang Waiting For Intel® VT-d Invalidation Wait Descriptor to Complete

Problem: If the virtual machine monitor (VMM) uses posted interrupts, Intel® VT-d hardware may not complete an Invalidation Wait descriptor.

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

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

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

SPR143. Power May Be Displayed Incorrectly

Problem: During Operating System Patch Load, if bit 0 in the IA32_MISC_PACKAGE_CTLs MSR (BCh) is set to 1, the Energy Consumed values in bits [31:0] of the PKG_ENERGY_STATUS MSR (611h) and the TOTAL_ENERGY_CONSUMED values in bits [31:0] of the PACKAGE_ENERGY_TIME_STATUS MSR (612h) may not be updated for 21.5 seconds after OSPL.

Implication: Due to this erratum, software should not rely on the power values in the Energy Consumed and TOTAL_ENERGY_CONSUMED MSRs for 21.5 seconds after OSPL.

Workaround: None identified.

Status: For the steppings affected, see the [Summary Tables of Changes](#) on page 13.

6.0 Specification Changes

There are no specification changes in this specification update revision.

7.0 Specification Clarifications

There are no specification changes in this specification update revision.

8.0 Documentation Changes

There are no documentation changes in this specification update revision.