

Intel[®] Xeon[®] Processor E3-1200 v6 Product Family

Specification Update

August 2019

Revision 019



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

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

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

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

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

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

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

Intel, Pentium, Celeron, Intel Core, Intel Xeon, Enhanced Intel SpeedStep Technology and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.


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

Copyright © 2019, Intel Corporation. All rights reserved.



Content

Revision History	4
Preface	5
Identification Information	7
Microcode Updates	10
Summary Tables of Changes	11
Errata	17
Specification Changes	51
Specification Clarifications	52
Documentation Changes	53



Revision History

Document Number	Revision	Description	Date
335718	001	Initial release	April 2017
335718	002	Added erratum KBW097	May 2017
335718	003	Skipped, no updates	NA
335718	004	Added errata KBW098 - KBW101	August 2017
335718	005	Added erratum KBW102	September 2017
335718	006	Added errata KBW103 - KBW104	October 2017
335718	007	Added erratum KBW105	November 2017
335718	008	Added erratum KBW106	December 2017
335718	009	Updated erratum KBW100 Added errata KBW107 - KBW108	February 2018
335718	010	Added errata KBW109 - KBW114	March 2018
335718	011	Added errata KBW115 - KBW116	June 2018
335718	012	Updated erratum KBW24, KBW116	July 2018
335718	013	Added errata KBW117 - KBW118	August 2018
335718	014	Removed erratum KBW22 Updated erratum KBW115 Added errata KBW119-124	October 2018
335718	015	Added erratum KBW125	November 2018
335718	016	Not available	February 2019
335718	017	Added errata KBW126, KBW127, KBW128	May 2019
335718	018	Added erratum KBW129	July 2019
572033	019	Revised content and title of KBW32 Added errata KBW130, KBW131	August 2019





Preface

This document is an update to the specifications contained in the documents listed in the following Affected Documents/Related Documents table. It is a compilation of device and document errata and specification clarifications and changes, and is intended for hardware system manufacturers and for software developers of applications, operating system, and tools.

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents

Document Title	Document Number/ Location
Intel(R) Xeon(R) Processor E3-1200 v6 Product Family for S Platforms Datasheet, Volume 1 of 2	335695
Intel(R) Xeon(R) Processor E3-1200 v6 Product Family for S Platforms Datasheet, Volume 2 of 2	335696

Related Documents

Document Title	Document Number/ Location
<i>AP-485, Intel® Processor Identification and the CPUID Instruction</i>	http://www.intel.com/design/processor/aplnots/241618.htm
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide</i> <i>Intel® 64 and IA-32 Intel Architecture Optimization Reference Manual</i>	http://www.intel.com/products/processor/manuals/index.htm
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html
<i>Intel® Virtualization Technology Specification for Directed I/O Architecture Specification</i>	D51397-001
<i>ACPI Specifications</i>	www.acpi.info



Nomenclature

Errata are design defects or errors. Errata may cause the processor's behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

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

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

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 the next release of the specifications.

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

Note:

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





Identification Information

Component Identification via Programming Interface

The processor stepping can be identified by the following register contents:

Table 1. S-Processor Line Component Identification

Reserved	Extended Family	Extended Model	Reserved	Processor Type	Family Code	Model Number	Stepping ID
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	0000000b	1001b		00b	0110b	1110b	xxxxb

Notes:


1. The Extended Family, Bits [27:20] are used in conjunction with the Family Code, specified in Bits [11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium 4, or Intel® Core™ processor family.
2. The Extended Model, Bits [19:16] in conjunction with the Model Number, specified in Bits [7:4], are used to identify the model of the processor within the processor's family.
3. The Family Code corresponds to Bits [11:8] of the EDX register after RESET, Bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
4. The Model Number corresponds to Bits [7:4] of the EDX register after RESET, Bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
5. The Stepping ID in Bits [3:0] indicates the revision number of that model.
6. Refer to Kaby Lake Processor BIOS Writers Guide for additional information. When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

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



Component Marking Information

Figure 1. LGA Top-Side Markings



Pin Count: 1151

Package Size: 37.5 mm x 37.5 mm

Production (SSPEC):

GRP1LINE1: Intel logo
GRP1LINE2: BRAND
GRP1LINE3: PROC#
GRP1LINE4: SSPEC SPEED
GRP1LINE5: {FPO} {eX}




Table 2. S-Processor Line

S-Spec #	Processor Number	Stepping	Cache Size	Functional Core	Processor Graphics Cores	Processor Graphics Freq.	Processor Graphics Max Dynamic Freq.	DDR4 Mem. (MT/s)	DDR3L Mem. (MT/s)	Core Freq.	Turbo 1 Core Freq. Rate	Thermal Design Power	Slot / Socket Type
SR325	E3-1280v6	B-0	8 MB	4	0	NA	NA	2400	1600	3.9 GHz	4.2 GHz	72 W	LGA1151
SR326	E3-1270v6	B-0	8 MB	4	0	NA	NA	2400	1600	3.8 GHz	4.2 GHz	72 W	LGA1151
SR327	E3-1240v6	B-0	8 MB	4	0	NA	NA	2400	1600	3.7 GHz	4.1 GHz	72 W	LGA1151
SR328	E3-1230v6	B-0	8 MB	4	0	NA	NA	2400	1600	3.5 GHz	3.9 GHz	72 W	LGA1151
SR329	E3-1220v6	B-0	8 MB	4	0	NA	NA	2400	1600	3.0 GHz	3.5 GHz	72 W	LGA1151
SR32A	E3-1275v6	B-0	8 MB	4	2	0.35 GHz	1.15 GHz	2400	1600	3.8 GHz	4.2 GHz	73 W	LGA1151
SR32B	E3-1245v6	B-0	8 MB	4	2	0.35 GHz	1.15 GHz	2400	1600	3.7 GHz	4.1 GHz	73 W	LGA1151
SR32C	E3-1225v6	B-0	8 MB	4	2	0.35 GHz	1.15 GHz	2400	1600	3.3 GHz	3.7 GHz	73 W	LGA1151

§



Summary Tables of Changes

The following table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed processor stepping. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

Codes Used in Summary Tables

Stepping

X: Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.

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

Page

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

Status

Doc: Document change or update will be implemented.

Plan Fix: This erratum may be fixed in a future stepping of the product.

Fixed: This erratum has been previously fixed.

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

Row

Change bar to left of a table row indicates this erratum is either new or modified from the previous version of the document.



Errata Summary Table

Table 3. S-Processor Line Errata Summary Table (Sheet 1 of 4)

Erratum ID	Stepping	Status	Title
	B-0		
KBW001	X	No Fix	Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures
KBW002	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation
KBW003	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception
KBW004	X	No Fix	The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When The UC Bit is Set
KBW005	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
KBW006	X	No Fix	SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior
KBW007	X	No Fix	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
KBW008	X	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
KBW009	X	No Fix	DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction
KBW010	X	No Fix	Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID
KBW011	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
KBW012	X	No Fix	The SMSW Instruction May Execute Within an Enclave
KBW013	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG Concurrent With an SMX SENTER/SEXIT May Result in a System Hang
KBW014	X	No Fix	Intel® PT TIP.PGD May Not Have Target IP Payload
KBW015	X	No Fix	SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior
KBW016	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
KBW017	X	No Fix	WRMSR May Not Clear The Sticky Count Overflow Bit in The IA32_MCi_STATUS MSR's Corrected Error Count Field
KBW018	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
KBW019	X	No Fix	Debug Exceptions May Be Lost or Misreported Following WRMSR to IA32_BIOS_UPDT_TRIG
KBW020	X	No Fix	Complex Interactions With Internal Graphics May Impact Processor Responsiveness
KBW021	X	No Fix	Intel® Processor Trace PSB+ Packets May Contain Unexpected Packets
KBW023	X	No Fix	VM Entry That Clears TraceEn May Generate a FUP
KBW024	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests May be Incorrect
KBW025	X	No Fix	ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK
KBW026	X	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
KBW027	X	No Fix	ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero
KBW028	X	No Fix	A VMX Transition Attempting to Load a Non-Existent MSR May Result in a Shutdown
KBW029	X	No Fix	Transitions Out of 64-bit Mode May Lead to an Incorrect FDP And FIP
KBW030	X	No Fix	Intel® PT FUP May be Dropped After OVF
KBW031	X	No Fix	ENCLS[ECREATE] Causes #GP if Enclave Base Address is Not Canonical



Table 3. S-Processor Line Errata Summary Table (Sheet 2 of 4)

Erratum ID	Stepping	Status	Title
	B-0		
KBW032	X	No Fix	Graphics VTd Hardware May Cache Invalid Entries
KBW033	X	No Fix	Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State
KBW034	X	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction
KBW035	X	No Fix	ENCLS[EINIT] Instruction May Unexpectedly #GP
KBW036	X	No Fix	Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop
KBW037	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG May be Counted as Multiple Instructions
KBW038	X	No Fix	Branch Instructions May Initialize MPX Bound Registers Incorrectly
KBW039	X	No Fix	Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled
KBW040	X	No Fix	Processor May Run Intel® AVX Code Much Slower Than Expected
KBW041	X	No Fix	Intel® PT Buffer Overflow May Result in Incorrect Packets
KBW042	X	No Fix	Last Level Cache Performance Monitoring Events May be Inaccurate
KBW043	X	No Fix	#GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave
KBW044	X	No Fix	Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception
KBW045	X	No Fix	Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits
KBW046	X	No Fix	CR3 Filtering Does Not Compare Bits [11:5] of CR3 and IA32_RTIT_CR3_MATCH in PAE Paging Mode
KBW047	X	No Fix	x87 FDP Value May be Saved Incorrectly
KBW048	X	No Fix	PECI Frequency Limited to 1 MHz
KBW049	X	No Fix	Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults
KBW050	X	No Fix	Intel® PT CYCThresh Value of 13 is Not Supported
KBW051	X	No Fix	Enabling Virtual Machine Extensions (VMX) Preemption Timer Blocks HDC Operation
KBW052	X	No Fix	Integrated Audio Codec May Not be Detected
KBW053	X	No Fix	Display Flickering May be Observed with Specific eDP Panels
KBW054	X	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
KBW055	X	No Fix	MACHINE_CLEAR.SMEMORY_ORDERING Performance Monitoring Event May Undercount
KBW056	X	No Fix	CTR_FRZ May Not Freeze Some Counters
KBW057	X	No Fix	Instructions And Branches Retired Performance Monitoring Events May Overcount
KBW058	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events May Overcount
KBW059	X	No Fix	Instructions Fetch #GP After RSM During Inter® PT May Push Incorrect RFLAGS Value on Stack
KBW060	X	No Fix	Access to Intel® SGX EPC Page in BLOCKED State is Not Reported as an Intel® SGX Induced Page Fault
KBW061	X	No Fix	MTF VM Exit on XBEGIN Instruction May Save State Incorrectly
KBW062	X	Fixed	Intel® Turbo Boost Technology May be Incorrectly Reported as Supported on Intel® Core™ i3 U/H/S, Select Intel® Mobile Pentium®, Intel® Mobile Celeron®, Intel® Pentium® and Intel® Celeron® Processors
KBW063	X	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
KBW064	X	No Fix	Executing a 256 Bit AVX Instruction May Cause Unpredictable Behavior
KBW065	X	No Fix	System May Hang During Display Power Cycles
KBW066	X	No Fix	Certain Non-Canonical IA32_BNDCFGS Values Will Not Cause VM-Entry Failures



Table 3. S-Processor Line Errata Summary Table (Sheet 3 of 4)

Erratum ID	Stepping	Status	Title
	B-0		
KBW067	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions
KBW068	X	No Fix	HWP's Guaranteed_Performance Updated Only on Configurable TDP Changes
KBW069	X	No Fix	RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
KBW070	X	No Fix	Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters
KBW071	X	No Fix	HWP's Maximum_Performance Value is Reset to 0xFF
KBW072	X	No Fix	HWP's Guaranteed_Performance and Relevant Status/Interrupt May be Updated More Than Once Per Second
KBW073	X	No Fix	Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes
KBW074	X	No Fix	HWP May Generate Thermal Interrupt While Not Enabled
KBW075	X	No Fix	Camera Device Does Not Issue an MSI When INTx is Enabled
KBW076	X	No Fix	An x87 Store Instruction Which Pends #PE May Lead to Unexpected Behavior When EPT A/D is Enabled.
KBW077	X	No Fix	Use of VMASKMOV to Store When Using EPT May Fail
KBW078	X	No Fix	PECI May Not be Functional After Package C10 Resume
KBW079	X	No Fix	Attempts to Retrain a PCIe* Link May be Ignored
KBW080	X	Fixed	PCIe* Expansion ROM Base Address Register May be Incorrect
KBW081	X	No Fix	PCIe* Port Does Not Support DLL Link Activity Reporting
KBW082	X	No Fix	BNDLDX And BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access
KBW083	X	No Fix	RING_PERF_LIMIT_REASONS May be Incorrect
KBW084	X	No Fix	Processor May Exceed VCCCore ICCMAX During Multi-core Turbo
KBW085	X	No Fix	Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions
KBW086	X	No Fix	EDRAM Corrected Error Events May Not be Properly Logged After a Warm Reset
KBW087	X	No Fix	Unpredictable System Behavior May Occur When System Agent Enhanced Intel Speedstep® Technology is Enabled
KBW088	X	No Fix	Processor May Hang Under Complex Scenarios
KBW089	X	No Fix	Some Bits in MSR_MISC_PWR_MGMT May be Updated on Writing Illegal Values to This MSR
KBW090	X	No Fix	Violations of Intel® Software Guard Extensions (Intel® SGX) Access-Control Requirements Produce #GP Instead of #PF
KBW091	X	No Fix	IA32_RTIT_CR3_MATCH MSR Bits[11:5] Are Treated As Reserved
KBW092	X	No Fix	APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode
KBW093	X	No Fix	The Intel® PT CR3 Filter is Not Re-evaluated on VM Entry
KBW094	X	No Fix	Display Slowness May be Observed Under Certain Display Commands Scenario
KBW095	X	No Fix	Short Loops Which Use AH/BH/CH/DH Registers May Cause Unpredictable System Behavior
KBW096	X	No Fix	CPUID TLB Associativity Information is Inaccurate
KBW097	X	No Fix	Processor Graphics May Render Incorrectly or May Hang Following Warm Reset with Package C8 Disabled
KBW098	X	No Fix	Using Different Vendors For 2400 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang
KBW099	X	No Fix	Unpredictable System Behavior May Occur in DDR4 Multi-Rank System



Table 3. S-Processor Line Errata Summary Table (Sheet 4 of 4)

Erratum ID	Stepping	Status	Title
	B-0		
KBW100	X	No Fix	Processor May Hang on Complex Sequence of Conditions
KBW101	X	No Fix	Potential Partial Trace Data Loss in Intel® Trace Hub ODLA When Storing to Memory
KBW102	X	No Fix	Display Artifacts May be Seen With High Bandwidth, Multiple Display Configurations
KBW103	X	No Fix	Spurious Corrected Errors May be Reported
KBW104	X	No Fix	Masked Bytes in a Vector Masked Store Instructions May Cause Write Back of a Cache Line
KBW105	X	No Fix	Processor May Incorrectly Assert PROCHOT During PkgC10
KBW106	X	No Fix	eDP 1.4 Ports With Link Rate 2.16 or 4.32 GHz May Not Resume From Low Power Graphics or System States.
KBW107	X	No Fix	Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May #GP
KBW108	X	No Fix	Precise Performance Monitoring May Generate Redundant PEBS Records
KBW109	X	No Fix	Intel® SGX ENCLS[EINIT] May Not Signal an Error For an Incorrectly Formatted SIGSTRUCT Input
KBW110	X	No Fix	Branch Instruction Address May be Incorrectly Reported on Intel® TSX Abort When Using Intel® MPX
KBW111	X	No Fix	Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP
KBW112	X	No Fix	Hitting a Code Breakpoint Inside an Intel® SGX Debug Enclave May Cause The Processor to Hang
KBW113	X	No Fix	Performance Monitoring ASCII Status Bit May be Inaccurate
KBW114	X	No Fix	Processor May Hang When Executing Code In an HLE Transaction Region
KBW115	X	No Fix	Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB
KBW116	X	No Fix	Intel® PT VM-entry Indication Depends on The Incorrect VMCS Control Field
KBW117	X	No Fix	VCVTSP2PH To Memory May Update MXCSR in The Case of a Fault on The Store
KBW118	X	No Fix	Intel® PT May Drop All Packets After an Internal Buffer Overflow
KBW119	X	No Fix	ZMM/YMM Registers May Contain Incorrect Values
KBW120	X	No Fix	Data Breakpoint May Not be Detected on a REP MOVSB
KBW121	X	No Fix	Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang
KBW122	X	No Fix	Performing an XACQUIRE to an Intel® PT ToPA Table May Lead to Processor Hang
KBW123	X	No Fix	Intel® PT PSB+ Packets May be Omitted on a C6 Transition
KBW124	X	No Fix	Intel® PT PacketEn Change on C-state Wake May Not Generate a TIP Packet
KBW125	X	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
KBW126	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
KBW127	X	No Fix	Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values
KBW128	X	No Fix	Intel® PT Trace May Silently Drop Second Byte of CYC Packet
KBW129	X	No Fix	Unexpected Uncorrected Machine Check Errors May Be Reported
KBW130	X	No Fix	Gen9 Graphics VTd Hardware May Cache Invalid Entries
KBW131	X	No Fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes



Errata

KBW1. Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures

Problem: Bits 53:50 of the IA32_VMX_BASIC MSR report the memory type that the processor uses to access the VMCS and data structures referenced by pointers in the VMCS. Due to this erratum, a VMX access to the VMCS or referenced data structures will instead use the memory type that the MTRRs (memory-type range registers) specify for the physical address of the access.

Implication: Bits 53:50 of the IA32_VMX_BASIC MSR report that the WB (write-back) memory type will be used but the processor may use a different memory type.

Workaround: Software should ensure that the VMCS and referenced data structures are located at physical addresses that are mapped to WB memory type by the MTRRs.

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

KBW2. Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation

Problem: This erratum may cause a machine-check error (IA32_MCi_STATUS.MCACOD=0150H) on the fetch of an instruction that crosses a 4-KByte address boundary. It applies only if (1) the 4-KByte linear region on which the instruction begins is originally translated using a 4-KByte page with the WB memory type; (2) the paging structures are later modified so that linear region is translated using a large page (2-MByte, 4-MByte, or 1-GByte) with the UC memory type; and (3) the instruction fetch occurs after the paging-structure modification but before software invalidates any TLB entries for the linear region.

Implication: Due to this erratum an unexpected machine check with error code 0150H may occur, possibly resulting in a shutdown. Intel has not observed this erratum with any commercially available software.

Workaround: Software should not write to a paging-structure entry in a way that would change, for any linear address, both the page size and the memory type. It can instead use the following algorithm: first clear the P flag in the relevant paging-structure entry (e.g., PDE); then invalidate any translations for the affected linear addresses; and then modify the relevant paging-structure entry to set the P flag and establish the new page size and memory type.

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

KBW3. Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception

Problem: The VAESIMC and VAESKEYGENASSIST instructions should produce a #UD (Invalid-Opcode) exception if the value of the vvvv field in the VEX prefix is not 1111b. Due to this erratum, if CR0.TS is "1", the processor may instead produce a #NM (Device-Not-Available) exception.

Implication: Due to this erratum, some undefined instruction encodings may produce a #NM instead of a #UD exception.

Workaround: Software should always set the vvvv field of the VEX prefix to 1111b for instances of the VAESIMC and VAESKEYGENASSIST instructions.

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



KBW4. The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When The UC Bit is Set

Problem: After a UC (uncorrected) error is logged in the IA32_MC0_STATUS MSR (401H), corrected errors will continue to be counted in the lower 14 bits (bits 51:38) of the Corrected Error Count. Due to this erratum, the sticky count overflow bit (bit 52) of the Corrected Error Count will not get updated when the UC bit (bit 61) is set to 1.

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

Workaround: None identified

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

KBW5. VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1

Problem: When "XD Bit Disable" in the IA32_MISC_ENABLE MSR (1A0H) bit 34 is set to 1, it should not be possible to enable the "execute disable" feature by setting IA32_EFER.NXE. Due to this erratum, a VM exit that occurs with the 1-setting of the "load IA32_EFER" VM-exit control may set IA32_EFER.NXE even if IA32_MISC_ENABLE bit 34 is set to 1. This erratum can occur only if IA32_MISC_ENABLE bit 34 was set by guest software in VMX non-root operation.

Implication: Software in VMX root operation may execute with the "execute disable" feature enabled despite the fact that the feature should be disabled by the IA32_MISC_ENABLE MSR. Intel has not observed this erratum with any commercially available software.

Workaround: A virtual-machine monitor should not allow guest software to write to the IA32_MISC_ENABLE MSR

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

KBW6. SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior

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

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

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

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

KBW7. x87 FPU Exception (#MF) May be Signaled Earlier Than Expected

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

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

Workaround: None identified.

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



KBW8. Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

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

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

Workaround: None identified.

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

KBW9. DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction

Problem: If XBEGIN is executed immediately after an execution of MOV to SS or POP SS, a transactional abort occurs and the logical processor restarts execution from the fallback instruction address. If execution of the instruction at that address causes a debug exception, bits [3:0] of the DR6 register may contain an incorrect value.

Implication: When the instruction at the fallback instruction address causes a debug exception, DR6 may report a breakpoint that was not triggered by that instruction, or it may fail to report a breakpoint that was triggered by the instruction.

Workaround: Avoid following a MOV SS or POP SS instruction immediately with an XBEGIN instruction.

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

KBW10. Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID

Problem: If CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 1 then opcode bytes F3 0F BC should be interpreted as TZCNT otherwise they will be interpreted as REP BSF. Due to this erratum, opcode bytes F3 0F BC may execute as TZCNT even if CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 0.

Implication: Software that expects REP prefix before a BSF instruction to be ignored may not operate correctly since there are cases in which BSF and TZCNT differ with regard to the flags that are set and how the destination operand is established.

Workaround: Software should use the opcode bytes F3 0F BC only if CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 1 and only if the functionality of TZCNT (and not BSF) is desired.

Status

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

KBW11. #GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code

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

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

Workaround: None identified.

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



KBW12. The SMSW Instruction May Execute Within an Enclave

Problem: The SMSW instruction is illegal within an Intel® Software Guard Extensions (Intel® SGX) enclave, and an attempt to execute it within an enclave should result in a #UD (invalid-opcode exception). Due to this erratum, the instruction executes normally within an enclave and does not cause a #UD.

Implication: The SMSW instruction provides access to CR0 bits [15:0] and will provide that information inside an enclave. These bits include NE, ET, TS, EM, MP and PE.

Workaround: None identified. If SMSW execution inside an enclave is unacceptable, system software should not enable Intel® SGX.

Status

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

KBW13. WRMSR to IA32_BIOS_UPDT_TRIG Concurrent With an SMX SENTER/SEXIT May Result in a System Hang

Problem: Performing WRMSR to IA32_BIOS_UPDT_TRIG (MSR 79H) on a logical processor while another logical processor is executing an SMX (Safer Mode Extensions) SENTER/SEXIT operation (GETSEC[SENER] or GETSEC[SEXIT] instruction) may cause the processor to hang.

Implication: When this erratum occurs, the system will hang. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

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

KBW14. Intel® PT TIP.PGD May Not Have Target IP Payload

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

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

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

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

KBW15. Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a #UD

Problem: Execution of a 64 bit operand MOVBE instruction with an operand-size override instruction prefix (66H) may incorrectly cause an invalid-opcode exception (#UD).

Implication: A MOVBE instruction with both REX.W=1 and a 66H prefix will unexpectedly cause an #UD (invalid-opcode exception). Intel has not observed this erratum with any commercially available software.

Workaround: Do not use a 66H instruction prefix with a 64-bit operand MOVBE instruction.

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



KBW16. Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception

Problem: Attempt to use FXSAVE or FXRSTOR with a VEX prefix should produce a #UD (Invalid-Opcode) exception. If either the TS or EM flag bits in CR0 are set, a #NM (device-not-available) exception will be raised instead of #UD exception.

Implication: Due to this erratum a #NM exception may be signaled instead of a #UD exception on an FXSAVE or an FXRSTOR with a VEX prefix.

Workaround: Software should not use FXSAVE or FXRSTOR with the VEX prefix.

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

KBW17. WRMSR May Not Clear The Sticky Count Overflow Bit in The IA32_MCI_STATUS MSRs' Corrected Error Count Field

Problem: The sticky count overflow bit is the most significant bit (bit 52) of the Corrected Error Count Field (bits[52:38]) in IA32_MCI_STATUS MSRs. Once set, the sticky count overflow bit may not be cleared by a WRMSR instruction. When this occurs, that bit can only be cleared by power-on reset.

Implication: Software that uses the Corrected Error Count field and expects to be able to clear the sticky count overflow bit may misinterpret the number of corrected errors when the sticky count overflow bit is set. This erratum does not affect threshold-based CMCI (Corrected Machine Check Error Interrupt) signaling.

Workaround: None identified.

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

KBW18. PEBS Eventing IP Field May be Incorrect After Not-Taken Branch

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

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

Workaround: None identified.

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

KBW19. Debug Exceptions May Be Lost or Misreported Following WRMSR to IA32_BIOS_UPDT_TRIG

Problem: If the WRMSR instruction writes to the IA32_BIOS_UPDT_TRIG MSR (79H) immediately after an execution of MOV SS or POP SS that generated a debug exception, the processor may fail to deliver the debug exception or, if it does, the DR6 register contents may not correctly reflect the causes of the debug exception.

Implication: Debugging software may fail to operate properly if a debug exception is lost or does not report complete information.

Workaround: Software should avoid using WRMSR instruction immediately after executing MOV SS or POP SS

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



KBW20. Complex Interactions With Internal Graphics May Impact Processor Responsiveness

Problem: Under complex conditions associated with the use of internal graphics, the processor may exceed the MAX_LAT CSR values (PCI configuration space, offset 03FH, bits [7:0]).

Implication: When this erratum occurs, the processor responsiveness is affected. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

KBW21. Intel® Processor Trace PSB+ Packets May Contain Unexpected Packets

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

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

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

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

KBW23. VM Entry That Clears TraceEn May Generate a FUP

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

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

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

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

KBW24. Performance Monitor Event For Outstanding Offcore Requests May be Incorrect

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

Implication: The performance monitor event OFFCORE_REQUESTS_OUTSTANDING may reflect an incorrect count.

Workaround: None identified.

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

KBW25. ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK

Problem: The Intel® Software Guard Extensions (Intel® SGX) ENCLU[EGETKEY] instruction ignores the MISCMASK field in KEYREQUEST structure when computing a provisioning key, a provisioning seal key, or a seal key.



Implication: ENCLU[EGETKEY] will return the same key in response to two requests that differ only in the value of KEYREQUEST.MISCMASK. Intel has not observed this erratum with any commercially available software.

Workaround: When executing the ENCLU[EGETKEY] instruction, software should ensure the bits set in KEYREQUEST.MISCMASK are a subset of the bits set in the current SECS's MISCSELECT field.

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

KBW26. POPCNT Instruction May Take Longer to Execute Than Expected

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

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

Workaround: None identified

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

KBW27. ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero

Problem: The Intel® Software Guard Extensions (Intel® SGX) ENCLU[EREPORT] instruction may cause a #GP (general protection fault) if any bit is set in TARGETINFO structure's MISCSELECT field.

Implication: This erratum may cause unexpected general-protection exceptions inside enclaves.

Workaround: When executing the ENCLU[EREPORT] instruction, software should ensure the bits set in TARGETINFO.MISCSELECT are a subset of the bits set in the current SECS's MISCSELECT field.

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

KBW28. A VMX Transition Attempting to Load a Non-Existent MSR May Result in a Shutdown

Problem: A VMX transition may result in a shutdown (without generating a machine-check event) if a non-existent MSR is included in the associated MSR-load area. When such a shutdown occurs, a machine check error will be logged with IA32_MCi_STATUS.MCACOD (bits [15:0]) of 406H, but the processor does not issue the special shutdown cycle. A hardware reset must be used to restart the processor.

Implication: Due to this erratum, the hypervisor may experience an unexpected shutdown.

Workaround: Software should not configure VMX transitions to load non-existent MSRs.

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

KBW29. Transitions Out of 64-bit Mode May Lead to an Incorrect FDP And FIP

Problem: A transition from 64-bit mode to compatibility or legacy modes may result in cause a subsequent x87 FPU state save to zeroing bits [63:32] of the FDP (x87 FPU Data Pointer Offset) and the FIP (x87 FPU Instruction Pointer Offset).

Implication: Leaving 64-bit mode may result in incorrect FDP and FIP values when x87 FPU state is saved.

Workaround: None identified. 64-bit software should save x87 FPU state before leaving 64-bit mode if it needs to access the FDP and/or FIP values.

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



KBW30. Intel® PT FUP May be Dropped After OVF

Problem: Some Intel® PT (Intel® Processor Trace) OVF (Overflow) packets may not be followed by a FUP (Flow Update Packet) or TIP.PGE (Target IP Packet, Packet Generation Enable).

Implication: When this erratum occurs, an unexpected packet sequence is generated.

Workaround: When it encounters an OVF without a following FUP or TIP.PGE, the Intel® PT trace decoder should scan for the next TIP, TIP.PGE, or PSB+ to resume operation.

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

KBW31. ENCLS[ECREATE] Causes #GP if Enclave Base Address is Not Canonical

Problem: pointer in the PAGEINFO structure, which is referenced by the RBX register. Due to this erratum, the instruction causes a #GP (general-protection fault) if the SECS attributes indicate that the enclave should operate in 64-bit mode and the enclave base linear address in the SECS is not canonical.

Implication: System software will incur a general-protection fault if it mistakenly programs the SECS with a non-canonical address. Intel has not observed this erratum with any commercially available software.

Workaround: System software should always specify a canonical address as the base address of the 64-bit mode enclave.

Status

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

KBW32. Graphics VTd Hardware May Cache Invalid Entries

Problem: The processor's graphics IOMMU (I/O Memory Management Unit) may cache invalid VTd context entries. This violates the VTd specification for HW Caching Mode where hardware implementations of this architecture must not cache invalid entries.

Implication: Due to this erratum, unpredictable system behavior and/or a system hang may occur.

Workaround: Software should flush the Gfx VTd context cache after any update of context table entries.

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

KBW33. Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State

Problem: Voltage glitch of up to 200mV on the VREF signal lasting for about 1mS may be observed when entering System S3 state. This violates the JEDEC DDR specifications.

Implication: Intel has not observed this erratum to impact the operation of any commercially available system.

Workaround: None identified.

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

KBW34. DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction

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



data breakpoints matched during the MOV/POP SS when the following instruction is either an MMX instruction that uses a memory addressing mode with an index or a store instruction.

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

Workaround: None identified.

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

KBW35. ENCLS[EINIT] Instruction May Unexpectedly #GP

Problem: When using Intel® Software Guard Extensions (Intel® SGX), the ENCLS[EINIT] instruction will incorrectly cause a #GP (general protection fault) if the MISCSELECT field of the SIGSTRUCT structure is not zero.

Implication: This erratum may cause an unexpected #GP, but only if software has set bits in the MISCSELECT field in SIGSTRUCT structure that do not correspond to extended features that can be written to the MISC region of the SSA (State Save Area). Intel has not observed this erratum with any commercially available software.

Workaround: When executing the ENCLS[EINIT] instruction, software should only set bits in the MISCSELECT field in the SIGSTRUCT structure that are enumerated as 1 by CPUID.(EAX=12H,ECX=0):EBX (the bit vector of extended features that can be written to the MISC region of the SSA).

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

KBW36. Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop

Problem: If an Intel® Processor Trace (Intel® PT) internal buffer overflow occurs immediately before software executes a taken branch or event that enters an Intel® PT TraceStop region, the OVF (Overflow) packet may be lost.

Implication: The trace decoder will not see the OVF packet, nor any subsequent packets (e.g., TraceStop) that were lost due to overflow.

Workaround: None identified.

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

KBW37. WRMSR to IA32_BIOS_UPDT_TRIG May be Counted as Multiple Instructions

Problem: When software loads a microcode update by writing to MSR IA32_BIOS_UPDT_TRIG (79H) on multiple logical processors in parallel, a logical processor may, due to this erratum, count the WRMSR instruction as multiple instruction-retired events.

Implication: Performance monitoring with the instruction-retired event may over count by up to four extra events per instance of WRMSR which targets the IA32_BIOS_UPDT_TRIG register.

Workaround: None identified.

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

KBW38. Branch Instructions May Initialize MPX Bound Registers Incorrectly

Problem: Depending on the current Intel® Memory Protection Extensions (Intel® MPX) configuration, execution of certain branch instructions (near CALL, near RET, near JMP,



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

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

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

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

KBW39. Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled

Problem: If Intel® Processor Trace (Intel® PT) is enabled, WRMSR will not cause a general-protection exception (#GP) on an attempt to write a non-canonical value to any of the following MSRs:

- MSR_LASTBRANCH_{0 - 31}_FROM_IP (680H – 69FH)
- MSR_LASTBRANCH_{0 - 31}_TO_IP (6C0H – 6DFH)
- MSR_LASTBRANCH_FROM_IP (1DBH)
- MSR_LASTBRANCH_TO_IP (1DCH)
- MSR_LASTINT_FROM_IP (1DDH)
- MSR_LASTINT_TO_IP (1DEH) Instead the same behavior will occur as if a canonical value had been written. Specifically, the WRMSR will be dropped and the MSR value will not be changed.

Implication: Due to this erratum, an expected #GP may not be signaled.

Workaround: None identified.

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

KBW40. Processor May Run Intel® AVX Code Much Slower Than Expected

Problem: After a C6 state exit, the execution rate of Intel® Advanced Vector Extensions (Intel® AVX) instructions may be reduced.

Implication: Applications using Intel® AVX instructions may run slower than expected.

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

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

KBW41. Intel® PT Buffer Overflow May Result in Incorrect Packets

Problem: Under complex micro-architectural conditions, an Intel® Processor Trace (Intel® PT) OVF (Overflow) packet may be issued after the first byte of a multi-byte CYC (Cycle Count) packet, instead of any remaining bytes of the CYC.

Implication: When this erratum occurs, the splicing of the CYC and OVF packets may prevent the Intel® PT decoder from recognizing the overflow. The Intel® PT decoder may then encounter subsequent packets that are not consistent with expected behavior.



Workaround: None Identified. The decoder may be able to recognize that this erratum has occurred when a two-byte CYC packet is followed by a single byte CYC, where the latter 2 bytes are 0xf302, and where the CYC packets are followed by a FUP (Flow Update Packet) and a PSB+ (Packet Stream Boundary+). It should then treat the two CYC packets as indicating an overflow.

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

KBW42. Last Level Cache Performance Monitoring Events May be Inaccurate

Problem: The performance monitoring events LONGEST_LAT_CACHE.REFERENCE (Event 2EH; Umask 4FH) and LONGEST_LAT_CACHE.MISS (Event 2EH; Umask 41H) count requests that reference or miss in the last level cache. However, due to this erratum, the count may be incorrect.

Implication: LONGEST_LAT_CACHE events may be incorrect.

Workaround: None identified. Software may use the following OFFCORE_REQUESTS model-specific sub events that provide related performance monitoring data:

DEMAND_DATA_RD, DEMAND_CODE_RD, DEMAND_RFO, ALL_DATA_RD,
L3_MISS_DEMAND_DATA_RD, ALL_REQUESTS.

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

KBW43. #GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave

Problem: When executing within an Intel® Software Guard Extensions (Intel® SGX) enclave, a #GP (general-protection exception) may be delivered instead of a #DB (debug exception) when an instruction breakpoint is detected. This occurs when the instruction to be executed spans two pages, the second of which has an entry in the EPCM (enclave page cache map) that is not valid.

Implication: Debugging software may not be invoked when an instruction breakpoint is detected.

Workaround: Software should ensure that all pages containing enclave instructions have valid EPCM entries.

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

KBW44. Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception

Problem: Execution of VAESENCLAST with VEX.L= 1 should signal a #UD (Invalid Opcode) exception, however, due to the erratum, a #NM (Device Not Available) exception may be signaled.

Implication: As a result of this erratum, an operating system may restore Intel® AVX and other state unnecessarily.

Workaround: None identified.

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

KBW45. Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits

Problem: In VMX non-root operation, Intel® Software Guard Extensions (Intel® SGX) enclave accesses to the APIC-access page may cause APIC-access VM exits instead of page faults.

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



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

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

KBW46. CR3 Filtering Does Not Compare Bits [11:5] of CR3 and IA32_RTIT_CR3_MATCH in PAE Paging Mode

Problem: In PAE paging mode, the CR3[11:5] are used to locate the page-directory-pointer table. Due to this erratum, those bits of CR3 are not compared to IA32_RTIT_CR3_MATCH (MSR 572H) when IA32_RTIT_CTL.CR3Filter (MSR 570H, bit 7) is set.

Implication: If multiple page-directory-pointer tables are co-located within a 4 KB region, CR3 filtering will not be able to distinguish between them so additional processes may be traced.

Workaround: None identified.

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

KBW47. 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.

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 Table of Changes*.

KBW48. PECI Frequency Limited to 1 MHz

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

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

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

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

KBW49. Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults

Problem: Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) specification specifies setting the Fault Processing Disable (FPD) field in the context (or extended-context) entry of IOMMU to mask recording of qualified DMA remapping faults for DMA requests processed through that context entry. Due to this erratum, the IOMMU unit for Processor Graphics device may record DMA remapping faults from Processor Graphics device (Bus: 0; Device: 2; Function: 0) even when the FPD field is set to 1.

Implication: Software may continue to observe DMA remapping faults recorded in the IOMMU Fault Recording Register even after setting the FPD field.

Workaround: None identified. Software may mask the fault reporting event by setting the Interrupt Mask (IM) field in the IOMMU Fault Event Control register (Offset 038H in GFXVTBAR).



Status

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

KBW50. Intel® PT CYCThresh Value of 13 is Not Supported

Problem: Intel® Processor Trace (Intel® PT) Cycle Count (CYC) threshold is configured through CYCThresh field in bits [22:19] of IA32_RTIT_CTL MSR (570H). A value of 13 is advertised as supported by CPUID (leaf 14H, sub-leaf 1H). Due to this erratum, if CYCThresh is set to 13 then the CYC threshold will be 0 cycles instead of 4096 (2¹³-1) cycles.

Implication: CYC packets may be issued in higher rate than expected if threshold value of 13 is used.

Workaround: None identified. Software should not use value of 13 for CYC threshold.

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

KBW51. Enabling Virtual Machine Extensions (VMX) Preemption Timer Blocks HDC Operation

Problem: Hardware Duty Cycling (HDC) will not put the physical package into the forced idle state while any logical processor is in VMX non-root operation and the "activate VMX-preemption timer" VM-execution control is 1.

Implication: HDC will not provide the desired power reduction when the VMX-preemption timer is active in VMX non-root operation.

Workaround: None identified.

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

KBW52. Integrated Audio Codec May Not be Detected

Problem: Integrated Audio Codec may lose power when LPSP (Low-Power Single Pipe) mode is enabled for an eDP* (embedded DisplayPort) or DP/HDMI ports. Platforms with Intel® Smart Sound Technology (Intel® SST) enabled are not affected.

Implication: The Audio Bus driver may attempt to do enumeration of Codecs when eDP or DP/HDMI port enters LPSP mode, due to this erratum, the Integrated Audio Codec will not be detected and audio maybe be lost.

Workaround: Intel® Graphics Driver 15.40.11.4312 or later will prevent the Integrated Audio Codec from losing power when LPSP mode is enabled.

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

KBW53. Display Flickering May be Observed with Specific eDP Panels

Problem: The processor may incorrectly configure transmitter buffer characteristics if the associated eDP panel requests VESA equalization preset 3, 5, 6, or 8.


Implication: Display flickering or display loss maybe observed.

Workaround: Intel® Graphics Driver version 15.40.12.4326 or later contains a workaround for this erratum.

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

KBW54. Incorrect Branch Predicted Bit in BTS/BTM Branch Records

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



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

Workaround: None identified.

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

KBW55. MACHINE_CLEARS.MEMORY_ORDERING Performance Monitoring Event May Undercount

Problem: The performance monitoring event MACHINE_CLEARS.MEMORY_ORDERING (Event C3H; Umask 02H) counts the number of machine clears caused by memory ordering conflicts. However due to this erratum, this event may undercount for VGATHER*/VPGATHER* instructions of four or more elements.

Implication: MACHINE_CLEARS.MEMORY_ORDERING performance monitoring event may undercount.

Workaround: None identified.

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

KBW56. CTR_FRZ May Not Freeze Some Counters

Problem: IA32_PERF_GLOBAL_STATUS.CTR_FRZ (MSR 38EH, bit 59) is set when either (1) IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI (MSR 1D9H, bit 12) is set and a PMI is triggered, or (2) software sets bit 59 of IA32_PERF_GLOBAL_STATUS_SET (MSR 391H). When set, CTR_FRZ should stop all core performance monitoring counters from counting. However, due to this erratum, IA32_PMC4-7 (MSR C5-C8H) may not stop counting. IA32_PMC4-7 are only available when a processor core is not shared by two logical processors.

Implication: General performance monitoring counters 4-7 may not freeze when IA32_PERF_GLOBAL_STATUS.CTR_FRZ is set.

Workaround: None identified.

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

KBW57. Instructions And Branches Retired Performance Monitoring Events May Overcount

Problem: The performance monitoring events INST_RETIRE (Event C0H; any Umask value) and BR_INST_RETIRE (Event C4H; any Umask value) count instructions retired and branches retired, respectively. However, due to this erratum, these events may overcount in certain conditions when:

- Executing VMASKMOV* instructions with at least one masked vector element
- Executing REP MOVS or REP STOS with Fast Strings enabled (IA32_MISC_ENABLES MSR (1A0H), bit 0 set)
- An MPX #BR exception occurred on BNDLDX/BNDSTX instructions and the BR_INST_RETIRE (Event C4H; Umask is 00H or 04H) is used.


Implication: INST_RETIRE and BR_INST_RETIRE performance monitoring events may overcount.

Workaround: None identified.

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

KBW58. Some OFFCORE_RESPONSE Performance Monitoring Events May Overcount

Problem: The performance monitoring events OFFCORE_RESPONSE (Events B7H and BBH) should count off-core responses matching the request-response configuration specified in MSR_OFFCORE_RSP_0 and MSR_OFFCORE_RSP_1 (1A6H and 1A7H, respectively)



for core-originated requests. However, due to this erratum, DMND_RFO (bit 1), DMND_IFETCH (bit 2) and OTHER (bit 15) request types may overcount.

Implication: Some OFFCORE_RESPONSE events may overcount.

Workaround: None identified. Software may use the following model-specific events that provide related performance monitoring data: OFFCORE_REQUESTS (all sub-events), L2_TRANS.L2_WB and L2_RQSTS.PF_MISS.

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

KBW59. Instructions Fetch #GP After RSM During Inter® PT May Push Incorrect RFLAGS Value on Stack

Problem: If Intel® Processor Trace (Intel® PT) is enabled, a #GP (General Protection Fault) caused by the instruction fetch immediately following execution of an RSM instruction may push an incorrect value for RFLAGS onto the stack.

Implication: Software that relies on RFLAGS value pushed on the stack under the conditions described may not work properly.

Implication: None identified.

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

KBW60. Access to Intel® SGX EPC Page in BLOCKED State is Not Reported as an Intel® SGX Induced Page Fault

Problem: If a page fault results from attempting to access a page in the Intel® Software Guard Extensions (Intel® SGX) Enclave Page Cache (EPC) that is in the BLOCKED state, the processor does not set bit 15 of the error code and thus fails to indicate that the page fault was Intel® SGX induced.

Implication: Due to this erratum, software may not recognize these page faults as being Intel® SGX induced.

Workaround: Before using the EBLOCK instruction to marking a page as BLOCKED, software should use paging to mark the page not present.

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



KBW61. MTF VM Exit on XBEGIN Instruction May Save State Incorrectly

Problem: Execution of an XBEGIN instruction while the monitor trap flag VM-execution control is 1 will be immediately followed by an MTF VM exit. If advanced debugging of RTM transactional regions has been enabled, the VM exit will erroneously save as instruction pointer the address of the XBEGIN instruction instead of the fallback instruction address specified by the XBEGIN instruction. In addition, it will erroneously set bit 16 of the pending-debug-exceptions field in the VMCS indicating that a debug exception or a breakpoint exception occurred. Using the monitor trap flag to debug or trace transactional regions may not operate properly. Intel has not observed this erratum with any commercially available software.

Implication: None identified

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

KBW62. Intel® Turbo Boost Technology May be Incorrectly Reported as Supported on Intel® Core™ i3 U/H/S, Select Intel® Mobile Pentium®, Intel® Mobile Celeron®, Intel® Pentium® and Intel® Celeron® Processors

Problem: These processors may incorrectly report support for Intel® Turbo Boost Technology via CPUID.06H.EAX bit 1.

Implication: The CPUID instruction may report Turbo Boost Technology as supported even though the processor does not permit operation above the Base Frequency.

Workaround: None identified.

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

KBW63. Performance Monitoring Counters May Undercount When Using CPL Filtering

Problem: Performance Monitoring counters configured to count only OS or only USR events by setting exactly one of bits 16 or 17 in IA32_PERFEVTSELx MSRs (186H-18DH) may not count for a brief period during the transition to a new CPL.

Implication: A measurement of ring transitions (using the edge-detect bit 18 in IA32_PERFEVTSELx) may undercount, such as CPL_CYCLES.RING0_TRANS (Event 5CH, Umask 01H). Additionally, the sum of an OS-only event and a USR-only event may not exactly equal an event counting both OS and USR. Intel has not observed any other software-visible impact

Workaround: None identified.

Status

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

KBW64. Executing a 256 Bit AVX Instruction May Cause Unpredictable Behavior

Problem: Under complex micro-architectural conditions, executing a 256 AVX bit instruction may result in unpredictable system behavior.

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

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

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



KBW65. System May Hang During Display Power Cycles

Problem: When the display is turned on after being shutoff to save power or when the display is exiting PSR (Panel Self Refresh) mode, the system may hang.

Implication: When this erratum occurs the system may hang.

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

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

KBW66. Certain Non-Canonical IA32_BNDCFGS Values Will Not Cause VM-Entry Failures

Problem: If the VM-entry controls Load IA32_BNDCFGS field (bit 16) is 1, VM-entry should fail when the value of the guest IA32_BNDCFGS field in the VMCS is not canonical (that is, when bits [63:47] are not identical). Due to this erratum, VM-entry does not fail if bits 63:48 are identical but differ from bit 47. In this case, VM-entry loads the IA32_BNDCFGS MSR with a value in which bits [63:48] are identical to the value of bit 47 in the VMCS field.

Implication: If the value of the guest IA32_BNDCFGS field in the VMCS is not canonical, VM-entry may load the IA32_BNDCFGS MSR with a value different from that of the VMCS field.

Workaround: None identified.

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

KBW67. PEBS EventingIP Field May Be Incorrect Under Certain Conditions

Problem: The EventingIP field in the Processor Event-Based Sampling (PEBS) record reports the address of the instruction that triggered the PEBS event. Under certain complex microarchitectural conditions, the EventingIP field may be incorrect.

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

Workaround: None identified.

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

KBW68. HWP's Guaranteed_Performance Updated Only on Configurable TDP Changes

Problem: According to HWP (Hardware P-states) specification, the Guaranteed_Performance field (bits [15:8]) in the IA32_HWP_CAPABILITIES MSR (771H) should be updated as a result of changes in the configuration of TDP, RAPL (Running Average Power Limit), and other platform tuning options that may have dynamic effects on the actual guaranteed performance support level. Due to this erratum, the processor will update the Guaranteed_Performance field only as a result of configurable TDP dynamic changes.

Implication: Software may read a stale value of the Guaranteed_Performance field.

Workaround: None identified.

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

KBW69. RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS

Problem: After a fault due to a failed Processor Event Based Sampling (PEBS) or Branch Trace Store (BTS) address translation, the Resume Flag (RF) may be incorrectly set in the EFLAGS image that is saved.



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

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

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

KBW70. Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters

Problem: Due to this erratum, if IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI (MSR 1D9H, bit 12) is set to 1 when Intel® Processor Trace (Intel® PT) triggers a Table of Physical Addresses (ToPA) PerfMon Interrupt (PMI), performance monitoring counters are not frozen as expected.

Implication: Performance monitoring counters will continue to count for events that occur during PMI handler execution.

Workaround: PMI handler software can programmatically stop performance monitoring counters upon entry.

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

KBW71. HWP's Maximum_Performance Value is Reset to 0xFF

Problem: According to Hardware P-states (HWP) specification, the reset value of the Maximum_Performance field (bits [15:8]) in IA32_HWP_REQUEST MSR (774h) should be set to the value of IA32_HWP_CAPABILITIES MSR (771H) Highest_Performance field (bits[7:0]) after reset. Due to this erratum, the reset value of Maximum_Performance is always set to 0xFF.

Implication: Software may see an unexpected value in Maximum Performance field. Hardware clipping will prevent invalid performance states.

Workaround: None identified.

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

KBW72. HWP's Guaranteed_Performance and Relevant Status/Interrupt May be Updated More Than Once Per Second

Problem: According to HWP (Hardware P-states) specification, the Guaranteed_Performance field (bits[15:8]) in the IA32_HWP_CAPABILITIES MSR (771H) and the Guaranteed_Performance_Change (bit 0) bit in IA32_HWP_STATUS MSR (777H) should not be changed more than once per second nor should the thermal interrupt associated with the change to these fields be signaled more than once per second. Due to this erratum, the processor may change these fields and generate the associated interrupt more than once per second

Implication: HWP interrupt rate due to Guaranteed_Performance field change can be higher than specified

Workaround: Clearing the Guaranteed_Performance_Change status bit no more than once per second will ensure that interrupts are not generated at too fast a rate

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



KBW73. Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes

Problem: The memory at-retirement performance monitoring events (next listed) may produce incorrect results when a performance counter is configured in OS-only or USR-only modes (bits 17 or 16 in IA32_PERFEVTSELx MSR). Counters with both OS and USR bits set are not affected by this erratum.

The list of affected memory at-retirement events is as follows:

MEM_INST_RETIRE.D.STLB_MISS_LOADS event D0H, umask 11H
MEM_INST_RETIRE.D.STLB_MISS_STORES event D0H, umask 12H
MEM_INST_RETIRE.D.LOCK_LOADS event D0H, umask 21H
MEM_INST_RETIRE.D.SPLIT_LOADS event D0H, umask 41H
MEM_INST_RETIRE.D.SPLIT_STORES event D0H, umask 42H
MEM_LOAD_RETIRE.D.L2_HIT event D1H, umask 02H
MEM_LOAD_RETIRE.D.L3_HIT event D1H, umask 04H
MEM_LOAD_RETIRE.D.L4_HIT event D1H, umask 80H
MEM_LOAD_RETIRE.D.L1_MISS event D1H, umask 08H
MEM_LOAD_RETIRE.D.L2_MISS event D1H, umask 10H
MEM_LOAD_RETIRE.D.L3_MISS event D1H, umask 20H
MEM_LOAD_RETIRE.D.FB_HIT event D1H, umask 40H
MEM_LOAD_L3_HIT_RETIRE.D.XSNP_MISS event D2H, umask 01H
MEM_LOAD_L3_HIT_RETIRE.D.XSNP_HIT event D2H, umask 02H
MEM_LOAD_L3_HIT_RETIRE.D.XSNP_HITM event D2H, umask 04H
MEM_LOAD_L3_HIT_RETIRE.D.XSNP_NONE event D2H, umask 08H

Implication: The listed performance monitoring events may produce incorrect results including PEBS records generated at an incorrect point

Workaround: None identified

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

KBW74. HWP May Generate Thermal Interrupt While Not Enabled

Problem: Due to this erratum, the conditions for Hardware P-states (HWP) to generate a thermal interrupt on a logical processor may generate thermal interrupts on both logical processors of that core.

Implication: If two logical processors of a core have different configurations of HWP (e.g., only enabled on one), an unexpected thermal interrupt may occur on one logical processor due to the HWP settings of the other logical processor.

Workaround: Software should configure HWP consistently on all logical processors of a core.

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

KBW75. Camera Device Does Not Issue an MSI When INTx is Enabled

Problem: When both Message Signaled Interrupts (MSI) and legacy INTx are enabled by the camera device, INTx is asserted rather than issuing the MSI, in violation of the PCI Local Bus Specification.

Implication: Due to this erratum, camera device interrupts can be lost leading to device failure.

Workaround: The camera device must disable legacy INTx by setting bit 10 of PCICMD (Bus 0; Device 5; Function 0; Offset 04H) before MSI is enabled

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



KBW76. An x87 Store Instruction Which Pends #PE May Lead to Unexpected Behavior When EPT A/D is Enabled.

Problem: An x87 store instruction which causes a #PE (Precision Exception) to be pended and updates an Extended Page Tables (EPT) A/D bit may lead to unexpected behavior.

Implication: The VMM may experience unexpected x87 fault or a machine check exception with the value of 0x150 in IA32_MC0_STATUS.MCACOD (bits [15:0] in MSR 401H)

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

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

KBW77. Use of VMASKMOV to Store When Using EPT May Fail

Problem: Use of VMASKMOV instructions to store data that splits over two pages, when the instruction resides on the first page may cause a hang if Extended Page Tables (EPT) is in use, and the store to the second page requires setting the A/D bits in the EPT entry.

Implication: Due to this erratum, the CPU may hang on the execution of VMASKMOV.

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

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

KBW78. PECI May Not be Functional After Package C10 Resume

Problem: When resuming from Package C10, PECI may fail to function properly.

Implication: When this erratum occurs, the PECI does not respond to any command.

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

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

KBW79. Attempts to Retrain a PCIe* Link May be Ignored

Problem: A PCI Express (PCIe*) link should retrain when Retrain Link (bit 5) in the Link Control register (Bus 0; Device 1; Functions 0,1,2; Offset 0xB0) is set. Due to this erratum, if the link is in the L1 state, it may ignore the retrain request

Implication: The PCIe* link may not behave as expected.

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

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

KBW80. PCIe* Expansion ROM Base Address Register May be Incorrect

Problem: After PCIe* 8.0 GT/s Link Equalization on a root port (Bus 0; Device 1; Function 0, 1, 2) has completed, the Expansion ROM Base Address Register (Offset 38H) may be incorrect.

Implication: Software that uses this BAR may behave unexpectedly. Intel has not observed this erratum with any commercially available software.

Workaround: None identified

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

KBW81. PCIe* Port Does Not Support DLL Link Activity Reporting

Problem: The PCIe* Base specification requires Data Link Layer (DLL) Link Activity Reporting when 8 GT/s link speed is supported. Due to this erratum, link activity reporting is not supported



Implication: Due to this erratum, PCIe* port does not support DLL Link Activity Reporting when 8 GT/s is supported.

Workaround: None identified

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

KBW82. BNDLDX And BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access

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

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

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

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

KBW83. RING_PERF_LIMIT_REASONS May be Incorrect

Problem: Under certain conditions, RING_PERF_LIMIT_REASONS (MSR 6B1H) may incorrectly assert the OTHER status bit (bit 8) as well as the OTHER log bit (bit 24).

Implication: When this erratum occurs, software using this register will incorrectly report clipping because of the OTHER reason.

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

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

KBW84. Processor May Exceed VCCCore ICCMAX During Multi-core Turbo

Problem: Due to this erratum, the maximum ring frequency limit is incorrectly configured to be 100 MHz higher than intended.

Implication: VCCCore ICCMAX may be temporarily exceeded when all the cores are executing at a Turbo frequency.

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

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


KBW85. Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions

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

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

Workaround: None identified

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



KBW86. EDRAM Corrected Error Events May Not be Properly Logged After a Warm Reset

Problem: After a warm reset, an EDRAM corrected error may not be logged correctly until the associated machine check register is initialized. This erratum may affect IA32_MC8_STATUS or IA32_MC10_STATUS.

Implication: The EDRAM corrected error information may be lost when this erratum occurs.

Workaround: Data from the affected machine check registers should be read and the registers initialized as soon as practical after a warm reset.

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

KBW87. Unpredictable System Behavior May Occur When System Agent Enhanced Intel Speedstep® Technology is Enabled

Problem: Under complex system conditions, SA-GV (System Agent Enhanced Intel Speedstep® Technology) may result in unpredictable system behavior.

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

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

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

KBW88. Processor May Hang Under Complex Scenarios

Problem: Under a complex micro-architectural conditions, the processor may hang with an internal timeout error (MCACOD 0400H) logged into IA32_MCi_STATUS.

Implication: This erratum may result in a processor hang

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

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

KBW89. Some Bits in MSR_MISC_PWR_MGMT May be Updated on Writing Illegal Values to This MSR

Problem: Attempts to write illegal values to MSR_MISC_PWR_MGMT (MSR 0x1AA) result in #GP (General Protection Fault) and should not change the MSR value. Due to this erratum, some bits in the MSR may be updated on writing an illegal value.

Implication: Certain fields may be updated with allowed values when writing illegal values to MSR_MISC_PWR_MGMT. Such writes will always result in #GP as expected.

Workaround: None identified. Software should not attempt to write illegal values to this MSR.
Status

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

KBW90. Violations of Intel® Software Guard Extensions (Intel® SGX) Access-Control Requirements Produce #GP Instead of #PF

Problem: Intel® Software Guard Extensions (Intel® SGX) define new access-control requirements on memory accesses. A violation of any of these requirements causes a page fault (#PF) that sets bit 15 (Intel® SGX) in the page-fault error code. Due to this erratum, these violations instead cause general-protection exceptions (#GP).

Implication: Software resuming from system sleep states S3 or S4 and relying on receiving a page fault from the above enclave accesses may not operate properly.

Workaround: Software can monitor #GP faults to detect that an enclave has been destroyed and needs to be rebuilt after resuming from S3 or S4

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



KBW91. IA32_RTIT_CR3_MATCH MSR Bits[11:5] Are Treated As Reserved

Problem: Due to this erratum, bits [11:5] in IA32_RTIT_CR3_MATCH (MSR 572H) are reserved; an MSR write that attempts to set that field to a non-zero value will result in a #GP fault.

Implication: The inability to write the identified bit field does not affect the functioning of Intel® Processor Trace (Intel® PT) operation because, as described in erratum SKL061, the bit field that is the subject of this erratum is not used during Intel® PT CR3 filtering.

Workaround: Ensure that bits [11:5] of the value written to IA32_RTIT_CR3_MATCH are zero, including cases where the selected page-directory-pointer-table base address has non-zero bits in this range.

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

KBW92. APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode

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

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

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

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

KBW93. The Intel® PT CR3 Filter is Not Re-evaluated on VM Entry

Problem: On a VMRESUME or VMLAUNCH with both TraceEn[0] and CR3Filter[7] in IA32_RTIT_CTL (MSR 0570H) set to 1 both before the VM Entry and after, the new value of CR3 is not compared with IA32_RTIT_CR3_MATCH (MSR 0572H).

Implication: The Intel® Processor Trace (Intel® PT) CR3 filtering mechanism may continue to generate packets despite a mismatching CR3 value, or may fail to generate packets despite a matching CR3, as a result of an incorrect value of IA32_RTIT_STATUS.ContextEn[1] (MSR 0571H) that results from the failure to re-evaluate the CR3 match on VM entry.

Workaround: None identified.

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

KBW94. Display Slowness May be Observed Under Certain Display Commands Scenario

Problem: Back to back access to the VGA register ports (I/O addresses 0x3C2, 0x3CE, 0x3CF) will experience higher than expected latency.

Implication: Due to this erratum, the processor may redraw the slowly when in VGA mode.

Workaround: None identified.

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



KBW95. Short Loops Which Use AH/BH/CH/DH Registers May Cause Unpredictable System Behavior

Problem: Under complex micro-architectural conditions, short loops of less than 64 instructions that use AH, BH, CH or DH registers as well as their corresponding wider register (e.g., RAX, EAX or AX for AH) may cause unpredictable system behavior. This can only happen when both logical processors on the same physical processor are active.

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

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

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

KBW96. CPUID TLB Associativity Information is Inaccurate

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

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

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

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

KBW97. Processor Graphics May Render Incorrectly or May Hang Following Warm Reset with Package C8 Disabled

Problem: Processor Graphics may not properly restore internal configuration after warm reset when package C8 is disabled.

Implication: Due to this erratum Processor Graphics may render incorrectly or hang on warm reset.

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

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

KBW98. Using Different Vendors For 2400 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang

Problem: When using 2400 MHz DDR4 UDIMMs from different vendors within the same channel, a higher rate of correctable errors may occur or the system may hang.

Implication: Due to this erratum, reported correctable error counts may increase or the system may hang.

Workaround: None identified. Use a single vendor for 2400 MHz UDIMMs.

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

KBW99. Unpredictable System Behavior May Occur in DDR4 Multi-Rank System

Problem: Due to incorrect configuration of DDR4 ODT by BIOS, it is possible for a multi-rank system to violate section 4.27 of the DDR4 JEDEC spec revision JESED79-4A.

Implication: Due to this erratum, complex microarchitectural conditions may result in unpredictable system behavior.

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

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



KBW100. Processor May Hang on Complex Sequence of Conditions

Problem: A complex set of architectural and micro-architectural conditions may lead to a processor hang with an internal timeout error (MCACOD 0400H) logged into IA32_MC3_STATUS (MSR 040DH, bits [15:0]). When both logical processors in a core are active, this erratum will not occur in one logical processor unless there is no interrupt for more than 10 seconds to the other logical processor.

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

Workaround: None Identified.

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

KBW101. Potential Partial Trace Data Loss in Intel® Trace Hub ODLA When Storing to Memory

Problem: When the Intel® Trace Hub (Intel® TH) On-Die Logic Analyzer (ODLA) is configured to trace to memory, under complex microarchitectural conditions, the trace may lose a timestamp.

Implication: Some ODLA trace data may be lost. This erratum does not affect other trace data sources. Typically, lost trace data will be displayed as "OVERFLOW." Subsequent timestamps will allow the trace decoder to resume tracing. Intel has not observed this erratum in commercially available software.

Workaround: None identified. For a particular workload, changing the memory buffer size or disabling deep compression may eliminate the microarchitectural condition that causes the erratum.

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

KBW102. Display Artifacts May be Seen With High Bandwidth, Multiple Display Configurations

Problem: With high bandwidth, multiple display configurations, display engine underruns may occur.

Implication: Due to this erratum, the display engine may generate display artifacts.

Workaround: This erratum can be worked around by Intel® Graphics Driver revisions of 15.46.4.64.4749 or later.

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

KBW103. Spurious Corrected Errors May be Reported

Problem: Due to this erratum, spurious corrected errors may be logged in the IA32_MC0_STATUS MSR (401H) register with the valid field (bit 63) set, the uncorrected error field bit (bit 61) not set, a Model Specific Error Code (bits [31:16]) of 0x0001, and an MCA Error Code (bits [15:0]) of 0x0005. If CMCI is enabled, these spurious corrected errors also signal interrupts.

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

Workaround: None Identified.

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



KBW104. Masked Bytes in a Vector Masked Store Instructions May Cause Write Back of a Cache Line

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

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

1. For MMIO range that is mapped as WB memory type, this erratum may lead to MCE (Machine Check Exception) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.

2. If the CPU cached data is stale, for example in the case of memory written directly by a non-coherent agent (agent that uses non-coherent writes), this erratum may lead to writing back stale cached data even if these bytes are masked.

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

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

KBW105. Processor May Incorrectly Assert PROCHOT During PkgC10

Problem: If the PROCHOT# pin is configured as an output-only signal, PROCHOT# may incorrectly be asserted during PkgC10.

Implication: When this erratum occurs, PROCHOT# may be incorrectly asserted. This can lead to the system fan unnecessarily turning on during PkgC10 or other unexpected platform behaviors.

Workaround: None identified.

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

KBW106. eDP 1.4 Ports With Link Rate 2.16 or 4.32 GHz May Not Resume From Low Power Graphics or System States.

Problem: When the Embedded Display Port is operating with link rates 2.16 GHz or 4.32 GHz, the port may not resume from DC5, DC6 display low power states or S3, S4, or S5 system states. This erratum only affects systems with eDP 1.4-compliant display panels.

Implication: Due to this erratum, the system may hang when resuming from system idle or S3/4/5 states.

Workaround: The graphics device driver can contain a workaround for this erratum; Intel® Graphics Driver revisions 15.49 PR1 or later contains this workaround.

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

KBW107. Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May #GP

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

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



Workaround: Software should clear all read-only fields before writing to this MSR

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

KBW108. Precise Performance Monitoring May Generate Redundant PEBS Records

Problem: Processor Event Based Sampling (PEBS) may generate redundant records for a counter overflow when used to profile cycles. This may occur when a precise performance monitoring event is configured on a general counter while setting the Invert and Counter Mask fields in IA32_PERFEVTSELx MSRs (186H - 18DH), and the counter is reloaded with a value smaller than 1000 (through the PEBS-counter-reset field of the DS Buffer Management Area).

Implication: PEBS may generate multiple redundant records, when used to profile cycles in certain conditions.

Workaround: It is recommended for software to forbid the use of the Invert bit in IA32_PERFEVTSELx MSRs or restrict PEBS-counter-reset value to a value of at least 1000.

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

KBW109. Intel® SGX ENCLS[EINIT] May Not Signal an Error For an Incorrectly Formatted SIGSTRUCT Input

Problem: The ENCLS[EINIT] instruction leaf may not signal an error on a specific combination of SIGSTRUCT values even though the signature does not fully comply with RSA signature specifications.

Implication: When this erratum occurs, ENCLS[EINIT] instruction leaf may pass the checks although the SIGSTRUCT signature does not fully comply with RSA signature specifications. This erratum does not compromise the security of Intel® SGX and does not impact normal usage of Intel® SGX.

Workaround: None identified. Software is not expected to be impacted by this erratum.

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

KBW110. Branch Instruction Address May be Incorrectly Reported on Intel® TSX Abort When Using Intel® MPX

Problem: When using Intel® Memory Protection Extensions (Intel® MPX), an Intel® Transactional Synchronization Extensions (Intel® TSX) transaction abort will occur in case of legacy branch (that causes bounds registers INIT) when at least one Intel® MPX bounds register was in a NON-INIT state. On such an abort, the branch Instruction address should be reported in the FROM_IP field in the Last Branch Records (LBR), Branch Trace Store (BTS) and Branch Trace Message (BTM) as well as in the Flow Update Packets (FUP) source IP address for Processor Trace (PT). Due to this erratum, the FROM_IP field in LBR/BTS/BTM, as well as the Flow Update Packets (FUP) source IP address that correspond to the Intel® TSX abort, may point to the preceding instruction.

Implication: Software that relies on the accuracy of the FROM_IP field/FUP source IP address and uses Intel® TSX may operate incorrectly when Intel® MPX is used.

Workaround: None identified.

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

KBW111. Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP

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



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

Workaround: None identified.

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

KBW112. Hitting a Code Breakpoint Inside an Intel® SGX Debug Enclave May Cause The Processor to Hang

Problem: Under complex microarchitecture conditions, the processor may hang when hitting code breakpoint inside a Intel® SGX debug enclave. This may happen only after opt-out entry into a Intel® SGX debug enclave and when the execution would set the accessed bit (A-bit) in any level of the paging or Extended Page Table (EPT) structures used to map the code page, and when both logical processors on the same physical core are active.

Implication: Due to this erratum, the processor may hang while debugging an Intel® SGX debug enclave.

Workaround: None identified.

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

KBW113. Performance Monitoring ASCI Status Bit May be Inaccurate

Problem: The Anti Side-Channel Interference (ASCI) field in IA32_PERF_GLOBAL_STATUS (MSR 38EH, bit 60) should be set when the count in any of the configured performance counters (i.e., IA32_PMCx or IA32_FIXED_CTRx) was altered due to direct or indirect operation of Intel® SGX. Due to this erratum, the ASCI bit may not be set properly when IA32_FIXED_CTR0 is used.

Implication: Software that relies on the value of the ASCI bit in IA32_PERF_GLOBAL_STATUS for its operation may not operate correctly when IA32_FIXED_CTR0 is used.

Workaround: None identified.

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

KBW114. Processor May Hang When Executing Code In an HLE Transaction Region

Problem: Under certain conditions, if the processor acquires an Hardware Lock Elision (HLE) lock via the XACQUIRE instruction in the Host Physical Address range between 40000000H and 403FFFFFFH, it may hang with an internal timeout error (MCACOD 0400H) logged into IA32_MCi_STATUS.

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

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

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

KBW115. Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB

Problem: Due to a rare microarchitectural condition, generation of an Intel® Processor Trace (Intel® PT) PSB (Packet Stream Boundary) packet can cause a single CYC (Cycle Count) packet, possibly along with an associated MTC (Mini Time Counter) packet, to be dropped.



Implication: An Intel® PT decoder that is using CYCs to track time or frequency will get an improper value due to the lost CYC packet.

Workaround: If an Intel® PT decoder is using CYCs and MTCs to track frequency, and either the first MTC following a PSB shows that an MTC was dropped, or the CYC value appears to be 4095 cycles short of what is expected, the CYC value associated with that MTC should not be used. The decoder should wait for the next MTC before measuring frequency again.

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

KBW116. Intel® PT VM-entry Indication Depends on The Incorrect VMCS Control Field

Problem: An Intel® Processor Trace (Intel® PT) Paging Information Packet (PIP), which includes indication of entry into non-root operation, will be generated on VM-entry as long as the "Conceal VMX in Intel® PT" field (bit 19) in Secondary Execution Control register (IA32_VMX_PROCBASED_CTLDS2, MSR 048BH) is clear. This diverges from expected behavior, since this PIP should instead be generated only with a zero value of the "Conceal VMX entries from Intel® PT" field (Bit 17) in the Entry Control register (IA32_VMX_ENTRY_CTLDS MSR 0484H).

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

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

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

KBW117. VCVTSP2PH To Memory May Update MXCSR in The Case of a Fault on The Store

Problem: Execution of the VCVTSP2PH instruction with a memory destination may update the MXCSR exceptions flags (bits [5:0]) if the store to memory causes a fault (e.g., #PF) or VM exit. The value written to the MXCSR exceptions flags is what would have been written if there were no fault.

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

Workaround: None identified.

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

KBW118. Intel® PT May Drop All Packets After an Internal Buffer Overflow

Problem: Due to a rare microarchitectural condition, an Intel® Processor Trace (Intel® PT) Table of Physical Addresses (ToPA) entry transition can cause an internal buffer overflow that may result in all trace packets, including the Overflow (OVF) packet, being dropped.

Implication: When this erratum occurs, all trace data will be lost until either PT is disabled and re-enabled via IA32_RTIT_CTL.TraceEn [bit 0] (MSR 0570H) or the processor enters and exits a C6 or deeper C state.

Workaround: None identified.

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

KBW119. ZMM/YMM Registers May Contain Incorrect Values

Problem: Under complex microarchitectural conditions values stored in ZMM and YMM registers may be incorrect.



Implication: Due to this erratum, YMM and ZMM registers may contain an incorrect value. Intel® has not observed this erratum with any commercially available software.

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

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

KBW120. Data Breakpoint May Not be Detected on a REP MOVSB

Problem: A REP MOVSB instruction that causes an exception or a VM exit may not detect a data breakpoint that occurred on an earlier memory access of that REP MOVSB instruction.

Implication: A debugger may miss a data read/write access if it is done by a REP MOVSB instruction.

Workaround: Software that relies on data breakpoint for correct execution should disable fast-strings (bit 0 in IA32_MISC_ENABLE MSR).

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

KBW121. Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang

Problem: If an Intel® PT (Processor Trace) ToPA (Table of Physical Addresses) table is placed in UC (Uncacheable) or USWC (Uncacheable Speculative Write Combining) memory, and a ToPA output region is filled during an Intel® Transaction Synchronization Extensions (Intel® TSX) transaction, the resulting ToPA table read may cause a processor hang.

Implication: Placing Intel® PT ToPA tables in non-cacheable memory when Intel® TSX is in use may lead to a processor hang.

Workaround: None identified. Intel® PT ToPA tables should be located in WB memory if Intel® TSX is in use.

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

KBW122. Performing an XACQUIRE to an Intel® PT ToPA Table May Lead to Processor Hang

Problem: If an XACQUIRE lock is performed to the address of an Intel® PT (Processor Trace) ToPA (Table of Physical Addresses) table, and that table is later read by the CPU during the HLE (Hardware Lock Elision) transaction, the processor may hang.

Implication: Accessing ToPA tables with XACQUIRE may result in a processor hang.

Workaround: None identified. Software should not access ToPA tables using XACQUIRE. An OS or hypervisor may wish to ensure all application or guest writes to ToPA tables to take page faults or EPT violations.

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

KBW123. Intel® PT PSB+ Packets May be Omitted on a C6 Transition

Problem: An Intel® PT (Processor Trace) PSB+ (Packet Stream Boundary+) set of packets may not be generated as expected when IA32_RTIT_STATUS.PacketByteCnt[48:32] (MSR 0x571) reaches the PSB threshold and a logical processor C6 entry occurs within the following one KByte of trace output.

Implication: After a logical processor enters C6, Intel® PT output may be missing PSB+ sets of packets.

Workaround: None identified.

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



KBW124. Intel® PT PacketEn Change on C-state Wake May Not Generate a TIP Packet

Problem: A TIP.PGE (Target IP, Packet Generation Enabled) or TIP.PGD (Target IP, Packet Generation Disabled) packet may not be generated if Intel® PT (Processor Trace) PacketEn changes after IA32_RTIT_STATUS.FilterEn (MSR 571H, bit 0) is re-evaluated on wakeup from C6 or deeper sleep state.

Implication: When code enters or exits an IP filter region without a taken branch, tracing may begin or cease without proper indication in the trace output. This may affect trace decoder behavior.

Workaround: None identified. A trace decoder will need to skip ahead to the next TIP or FUP packet to determine the current IP.

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

KBW125. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions

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

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

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

Status: For the steppings affected, see the *Summary Table of Changes*

KBW126. Using Intel® TSX Instructions May Lead to Unpredictable System Behavior

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

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

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

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

KBW127. Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values

Problem: When Restricted Transactional Memory (RTM) is supported (CPUID.07H.EBX.RTM [bit 11] = 1) and when TSX_FORCE_ABORT=0, Performance Monitor Unit (PMU) general purpose counter 3 (IA32_PMC3, MSR C4H and IA32_A_PMC3, MSR 4C4H) may contain unexpected values. Further, IA32_PREFEVTSEL3 (MSR 189H) may also contain unexpected configuration values.



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

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

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

KBW128. Intel® PT Trace May Silently Drop Second Byte of CYC Packet

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

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

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

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

KBW129. Unexpected Uncorrected Machine Check Errors May Be Reported

Problem: In rare micro-architectural conditions, the processor may report unexpected machine check errors. When this erratum occurs, IA32_MC0_STATUS (MSR 401H) will have the valid bit set (bit 63), the uncorrected error bit set (bit 61), a model specific error code of 03H (bits [31:16]) and an MCA error code of 05H (bits [15:0]).

Implication: Due to this erratum, software may observe unexpected machine check exceptions.

Workaround: None Identified.

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

KBW130. Gen9 Graphics VTd Hardware May Cache Invalid Entries

Problem: The Gen9 graphics subsystem may cache invalid VTd context entries.

Implication: Due to this erratum, unpredictable system behavior and/or a system hang may occur.

Workaround: Software should flush the Gfx VTd context cache after any update of context table entries.

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

KBW131. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

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

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

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

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


S



Specification Changes

There are no specification changes in this Specification Update revision.


§



Specification Clarifications

There are no specification clarifications in this Specification Update revision.

§



Documentation Changes

There are no documentation changes in this Specification Update revision.

§