

# Intel<sup>®</sup> Xeon<sup>®</sup> E-2100 and E-2200 Processor Family

**Specification Update**

---

**Revision 012**

***September 2020***



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

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

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

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

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

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

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

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see <http://www.intel.com/technology/turboboost>

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

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

Intel, Pentium, Celeron, Core, Xeon, SpeedStep Technology, and the Intel logo are trademarks of Intel Corporation or its subsidiaries.

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


Copyright © 2020, Intel Corporation. All rights reserved.



# Content

---

<b>Revision History</b> .....	4
<b>Preface</b> .....	5
<b>Identification Information</b> .....	7
<b>Summary Tables of Changes</b> .....	11
<b>Errata</b> .....	16
<b>Specification Changes</b> .....	53
<b>Specification Clarifications</b> .....	54
<b>Documentation Changes</b> .....	55



# Revision History

Date	Revision	Description
September 2020	012	Added "Intel® Xeon® E-2100 and Intel® Xeon® E-2200 Turbo Boost 2.0 Frequency" Added erratum CFW136
August 2020	011	Added erratum CFW135
May 2020	010	Added errata CFW133 and CFW134
April 2020	009	Added errata CFW131 and CFW132
February 2020	008	Added errata CFW126 - CFW130 Revised CFW2
August 2019	007	Revised content and title of CFW32 Revised CFW101 affected Stock Keeping Units (SKUs) in <a href="#">Table 5, "S-Processor Line Errata Summary Table"</a> Added errata CFW121 - CFW125
July 2019	006	Added erratum CFW120
May 2019	005	Added erratum CFW117, CFW118, CFW119
February 2019	004	Not available
November 2018	003	Added erratum CFW116
October 2018	002	Removed erratum CFW22 Updated erratum CFW105 Added errata CFW110 - CFW115
September 2018	001	Initial release





## Preface

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

Information types defined in the Nomenclature section 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
<i>Intel® Xeon® E-2100 Processor Family Datasheet, Volume 1 of 2</i>	338012
<i>Intel® Xeon® E-2100 Processor Family Datasheet, Volume 2 of 2</i>	338013

### Related Documents

Document Title	Document Number/ Location
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2A: Instruction Set Reference, A-L -- Refer to section "AP-485, Intel® Processor Identification and the CPUID Instruction"</i>	<a href="http://www.intel.com/sdm">www.intel.com/sdm</a>
<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, A-L</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2B: Instruction Set Reference, M-U</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2C: Instruction Set Reference, V-Z</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3A: System Programming Guide, Part 1</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3B: System Programming Guide, Part 2</i> <i>Intel® 64 and IA-32 Architectures Optimization Reference Manual</i>	<a href="https://software.intel.com/en-us/articles/intel-sdm">https://software.intel.com/en-us/articles/intel-sdm</a>
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	<a href="https://software.intel.com/en-us/download/intel-64-and-ia-32-architectures-software-developers-manual-documentation-changes">https://software.intel.com/en-us/download/intel-64-and-ia-32-architectures-software-developers-manual-documentation-changes</a>
<i>RS - Intel® Virtualization Technology Specification for Directed I/O Architecture Specification</i>	<a href="https://soco.intel.com/docs/DOC-1945654">https://soco.intel.com/docs/DOC-1945654</a>
<i>Advanced Configuration Power Interface (ACPI) Specifications</i>	<a href="http://www.acpi.info">www.acpi.info</a>



## 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.


**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 Celeron®, Pentium®, or Intel® Core™ processor family.
2. The extended model, bits [19:16] in conjunction with the model number, specified in bits [7:4], are used to identify the model of the processor within the processor’s family.
3. The family code corresponds to bits [11:8] of the Extended Data Register (EDX) register after RESET; bits [11:8] of the Extended Accumulator Register (EAX) register after the CPUID instruction are executed with a 1 in the EAX register and with the generation field of the Device ID register accessible through the 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 are executed with a 1 in the EAX register and with 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 the *Skylake, Kaby Lake and Coffee Lake, Whiskey Lake and Comet Lake Processor Family Core and Uncore BIOS Specification Rev 3.2.0*, document number 550049, for additional information. When the EAX is initialized to a value of 1, the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number, and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

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



## Intel® Xeon® E-2100 and Intel® Xeon® E-2200 Turbo Boost 2.0 Frequency

**Table 2. Intel® Xeon® E-2100 Turbo Boost 2.0 Frequency**

Processor SKU	TDP (W)	Base Frequency (GHz)	Cores	1	Maximum Frequency** in GHz (+ x00 MHz over base frequency) by Active Cores					
					2	3	4	5	6	
<b>Standard Power SKUs</b>										
E-2186G	95	3.8	6	+9	+8	+8	+7	+6	+5	
E-2176G	80	3.7	6	+10	+9	+8	+7	+7	+6	
E-2174G	71	3.8	4	+9	+7	+6	+5			
E-2146G	80	3.5	6	+10	+9	+8	+8	+8	+7	
E-2144G	71	3.6	4	+9	+8	+7	+6			
E-2136	80	3.3	6	+12	+11	+10	+10	+10	+9	
E-2134	71	3.5	4	+10	+9	+8	+7			
E-2126G	80	3.3	6	+12	+11	+10	+9	+9	+8	
E-2124G	71	3.4	4	+11	+10	+8	+7			
E-2124	71	3.3	4	+10	+9	+8	+6			
E-2104G	65	3.2	4	+0	+0	+0	+0			

**Table 3. Intel® Xeon® E-2200 Turbo Boost 2.0 Frequency**


Processor SKU	TDP (W)	Base Frequency (GHz)	Cores	1	Maximum Frequency** in GHz (+ x00 MHz over base frequency) by Active Cores						
					2	3	4	5	6	7	8
<b>Standard Power SKUs</b>											
E-2288G	95	3.7	8	+13	+12	+12	+11	+11	+10	+10	+10
E-2286G	95	4	6	+9	+8	+8	+7	+7	+6		
E-2278G	80	3.4	8	+16	+15	+15	+14	+14	+13	+12	+12
E-2276G	80	3.8	6	+11	+10	+10	+9	+9	+8		
E-2274G	83	4	4	+9	+8	+6	+4				
E-2246G	80	3.6	6	+12	+11	+11	+10	+10	+9		
E-2244G	71	3.8	4	+10	+9	+8	+7				
E-2236	80	3.4	6	+14	+13	+13	+12	+12	+11		
E-2234	71	3.6	4	+12	+11	+10	+9				
E-2226G	80	3.4	6	+13	+12	+12	+11	+11	+10		
E-2224G	71	3.5	4	+12	+11	+10	+9				
E-2224	71	3.4	4	+12	+11	+10	+8				





## Component Marking Information

Figure 1. Land Grid Array (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}

**Note:**

**Note:** For the Intel® Xeon® E-2100 and E-2200 Processor Family SKUs, see: <https://ark.intel.com/content/www/us/en/ark/products/series/134861/intel-xeon-e-processor.html>



# 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 5. S-Processor Line Errata Summary Table (Sheet 1 of 4)**

Erratum ID	Stepping		Status	Title
	U-0	R-0		
CFW001	X	X	No Fix	Reported Memory Type May Not Be Used to Access the Virtual-Machine Control Structure (VMCS) and Referenced Data Structures
CFW002	X	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation
CFW003	X	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With an Illegal Value for VEX.vvvv May Produce a #NM Exception
CFW004	X	X	No Fix	The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When the UC Bit is Set
CFW005	X	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
CFW006	X	X	No Fix	SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
CFW007	X	X	No Fix	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
CFW008	X	X	No Fix	Incorrect FROM_IP Value For an Restricted Transactional Memory (RTM) Abort in Branch Trace Message (BTM) or Branch Trace Store (BTS) May be Observed
CFW009	X	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
CFW010	X	X	No Fix	Opcode Bytes F3 0F BC May Execute as TZCNT Even When TZCNT Not Enumerated by CPUID
CFW011	X	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
CFW012	X	X	No Fix	The SMSW Instruction May Execute Within an Enclave
CFW013	X	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG Concurrent With an SMX SENTER/SEXIT May Result in a System Hang
CFW014	X	X	No Fix	Intel® Processor Trace (Intel® PT) TIP.PGD May Not Have Target IP Payload
CFW015	X	X	No Fix	SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
CFW016	X	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
CFW017	X	X	No Fix	WRMSR May Not Clear the Sticky Count Overflow Bit in the IA32_MCI_STATUS MSRs' Corrected Error Count Field
CFW018	X	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
CFW019	X	X	No Fix	Debug Exceptions May Be Lost or Misreported Following WRMSR to IA32_BIOS_UPDT_TRIG
CFW020	X	X	No Fix	Complex Interactions With Internal Graphics May Impact Processor Responsiveness
CFW021	X	X	No Fix	Intel® PT PSB+ Packets May Contain Unexpected Packets
CFW023	X	X	No Fix	VM Entry That Clears TraceEn May Generate a FUP
CFW024	X	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests May be Incorrect
CFW025	X	X	No Fix	ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK
CFW026	X	X	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
CFW027	X	X	No Fix	ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero
CFW028	X	X	No Fix	A VMX Transition Attempting to Load a Non-Existent MSR May Result in a Shutdown
CFW029	X	X	No Fix	Transitions Out of 64-bit Mode May Lead to an Incorrect FDP and FIP
CFW030	X	X	No Fix	Intel® PT FUP May be Dropped After OVF
CFW031	X	X	No Fix	ENCLS[ECREATE] Causes #GP if Enclave Base Address is Not Canonical
CFW032	X	X	No Fix	Graphics Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Hardware May Cache Invalid Entries
CFW033	X	X	No Fix	Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State



**Table 5. S-Processor Line Errata Summary Table (Sheet 2 of 4)**

Erratum ID	Stepping		Status	Title
	U-0	R-0		
CFW034	X	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
CFW035	X	X	No Fix	ENCLS[EINIT] Instruction May Unexpectedly #GP
CFW036	X	X	No Fix	Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop
CFW037	X	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG May be Counted as Multiple Instructions
CFW038	X	X	No Fix	Branch Instructions May Initialize Intel® Memory Protection Extensions (Intel® MPX) Bound Registers Incorrectly
CFW039	X	X	No Fix	Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled
CFW040	X	X	No Fix	Processor May Run Intel® Advanced Vector Extensions (Intel® AVX) Code Much Slower Than Expected
CFW041	X	X	No Fix	Intel® PT Buffer Overflow May Result in Incorrect Packets
CFW042	X	X	No Fix	Last Level Cache Performance Monitoring Events May be Inaccurate
CFW043	X	X	No Fix	#GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave
CFW044	X	X	No Fix	Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception
CFW045	X	X	No Fix	Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits
CFW046	X	X	No Fix	CR3 Filtering Does Not Compare Bits [11:5] of CR3 and IA32_RTIT_CR3_MATCH in PAE Paging Mode
CFW047	X	X	No Fix	x87 FDP Value May be Saved Incorrectly
CFW048	X	X	No Fix	PECI Frequency Limited to 1 MHz
CFW049	X	X	No Fix	Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults
CFW050	X	X	No Fix	Intel® PT CYCThresh Value of 13 is Not Supported
CFW051	X	X	No Fix	Enabling Virtual Machine Extensions (VMX) Preemption Timer Blocks HDC Operation
CFW052	X	X	No Fix	Integrated Audio Codec May Not be Detected
CFW053	X	X	No Fix	Display Flickering May be Observed with Specific eDP Panels
CFW054	X	X	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
CFW055	X	X	No Fix	MACHINE_CLEAR.SMEMORY_ORDERING Performance Monitoring Event May Undercount
CFW056	X	X	No Fix	CTR_FRZ May Not Freeze Some Counters
CFW057	X	X	No Fix	Instructions and Branches Retired Performance Monitoring Events May Overcount
CFW058	X	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events May Overcount
CFW059	X	X	No Fix	Instructions Fetch #GP After RSM During Intel® PT May Push Incorrect RFLAGS Value on Stack
CFW060	X	X	No Fix	Access to Intel® SGX EPC Page in BLOCKED State is Not Reported as an Intel® SGX Induced Page Fault
CFW061	X	X	No Fix	MTF VM Exit on XBEGIN Instruction May Save State Incorrectly
CFW062	X	X	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
CFW063	X	X	No Fix	Certain Non-Canonical IA32_BNDCFGS Values Will Not Cause VM-Entry Failures
CFW064	X	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions
CFW065	X	X	No Fix	HWP's Guaranteed_Performance Updated Only on Configurable TDP Changes
CFW066	X	X	No Fix	RF May Be Incorrectly Set in the EFLAGS That is Saved on a Fault in PEBS or BTS
CFW067	X	X	No Fix	Intel® PT ToPA Performance Monitoring Interrupt (PMI) Does Not Freeze Performance Monitoring Counters
CFW068	X	X	No Fix	HWP's Maximum_Performance Value is Reset to 0xFF
CFW069	X	X	No Fix	HWP's Guaranteed_Performance and Relevant Status/Interrupt May Be Updated More Than Once Per Second





Table 5. S-Processor Line Errata Summary Table (Sheet 3 of 4)

Erratum ID	Stepping		Status	Title
	U-0	R-0		
CFW070	X	X	No Fix	Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes
CFW071	X	X	No Fix	HWP May Generate Thermal Interrupt While Not Enabled
CFW072	X	X	No Fix	Camera Device Does Not Issue an MSI When INTx is Enabled
CFW073	X	X	No Fix	Attempts to Retrain a PCI Express* (PCIe*) Link May be Ignored
CFW074	X	X	No Fix	PCIe* Port Does Not Support DLL Link Activity Reporting
CFW075	X	X	No Fix	BNDLDX And BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access
CFW076	X	X	No Fix	RING_PERF_LIMIT_REASONS May Be Incorrect
CFW077	X	X	No Fix	Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions
CFW078	X	X	No Fix	Some Bits in MSR_MISC_PWR_MGMT May Be Updated on Writing Illegal Values to This MSR
CFW079	X	X	No Fix	Violations of Intel® SGX Access-Control Requirements Produce #GP Instead of #PF
CFW080	X	X	Fixed	IA32_RTIT_CR3_MATCH MSR Bits[11:5] Are Treated as Reserved
CFW081	X	X	No Fix	The Intel® PT CR3 Filter is Not Re-Evaluated on VM Entry
CFW082	X	X	No Fix	Display Slowness May be Observed Under Certain Display Commands Scenario
CFW083	X	X	No Fix	CPUID TLB Associativity Information is Inaccurate
CFW084	X	X	No Fix	Unpredictable System Behavior May Occur in DDR4 Multi-Rank System
CFW085	X	X	No Fix	Processor May Hang on Complex Sequence of Conditions
CFW086	X	X	No Fix	Potential Partial Trace Data Loss in Intel® Trace Hub (Intel® TH) ODLA When Storing to Memory
CFW087	X	X	No Fix	Using Different Vendors for 2666 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang
CFW088	X	X	No Fix	Spurious Corrected Errors May be Reported
CFW089	X	X	No Fix	Reads From IA32_SGXLEPUBKEYHASH MSRs Return Values in Incorrect Order
CFW090	X	X	No Fix	Masked Bytes in a Vector Masked Store Instructions May Cause Write Back of a Cache Line
CFW091	X	X	No Fix	MOVNTDQA From WC Memory May Pass Earlier MFENCE instructions
CFW092	X	X	No Fix	MOVNTDQA From WC Memory May Pass Earlier Locked Instructions
CFW093	X	X	No Fix	PEBS Record After a WRMSR to IA32_BIOS_UPDT_TRIG May be Incorrect
CFW094	X	X	No Fix	Processor May Incorrectly Assert PROCHOT During PkgC10
CFW095	X	X	No Fix	Writing Non-Zero Values to Read Only Fields in IA32_THERM_STATUS MSR May #GP
CFW096	X	X	No Fix	Precise Performance Monitoring May Generate Redundant PEBS Records
CFW097	X	X	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
CFW098	X	X	No Fix	Intel® SGX ENCLS[EINIT] May Not Signal an Error For an Incorrectly Formatted SIGSTRUCT Input
CFW099	X	X	No Fix	Branch Instruction Address May be Incorrectly Reported on Intel® Transactional Synchronization Extensions (Intel® TSX) Abort When Using Intel® MPX
CFW100	X	X	No Fix	Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP
CFW101	X		No Fix	Hitting a Code Breakpoint Inside an Intel® SGX Debug Enclave May Cause the Processor to Hang
CFW102	X	X	No Fix	Performance Monitoring Anti Side-Channel Interference (ASCI) Status Bit May Be Inaccurate
CFW103	X	X	No Fix	Processor May Hang When Executing Code in an HLE Transaction Region
CFW104	X	X	No Fix	The Processor May Fail to Boot During DDR4 Memory Training
CFW105	X	X	No Fix	Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB
CFW106	X	X	No Fix	Intel® PT VM-entry Indication Depends on the Incorrect VMCS Control Field
CFW107	X	X	No Fix	Certain DDR4 Memory Configurations May Cause Unpredictable System Behavior



**Table 5. S-Processor Line Errata Summary Table (Sheet 4 of 4)**

Erratum ID	Stepping		Status	Title
	U-0	R-0		
CFW108	X	X	No Fix	VCVTPS2PH To Memory May Update MXCSR in the Case of a Fault on the Store
CFW109	X	X	No Fix	Intel® PT May Drop All Packets After an Internal Buffer Overflow
CFW110	X	X	No Fix	ZMM/YMM Registers May Contain Incorrect Values
CFW111	X	X	No Fix	Data Breakpoint May Not be Detected on a REP MOVSB
CFW112	X	X	No Fix	Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang
CFW113	X	X	No Fix	Performing an XACQUIRE to an Intel® PT ToPA Table May Lead to Processor Hang
CFW114	X	X	No Fix	Intel® PT PSB+ Packets May be Omitted on a C6 Transition
CFW115	X	X	No Fix	Intel® PT PacketEn Change on C-state Wake May Not Generate a TIP Packet
CFW116	X	X	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
CFW117	X	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
CFW118	X	X	No Fix	Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values
CFW119	X	X	No Fix	Intel® PT Trace May Silently Drop Second Byte of CYC Packet
CFW120	X	X	No Fix	Unexpected Uncorrected Machine Check Errors May Be Reported
CFW121	X		No Fix	Processor May Hang at High Temperature With a High-Throughput Graphics Workload
CFW122	X	X	No Fix	Gen9 Graphics Intel® VT-d Hardware May Cache Invalid Entries
CFW123		X	No Fix	Queued Invalidation Is Prevented When Intel® VT-d is Disabled
CFW124		X	No Fix	Processor May Hang During PKG-C8/C9/C10 Exit
CFW125	X	X	No Fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes
CFW126	X	X	No Fix	Executing Some Instructions May Cause Unpredictable Behavior
CFW127	X	X	No Fix	Incorrect Execution of Internal Branch Instructions May Lead to Unpredictable System Behavior
CFW128	X	X	No Fix	Unexpected Page Faults in Guest Virtualization Environment
CFW129	X	X	No Fix	Intel® SGX Key Confidentiality May be Compromised
CFW130	X	X	No Fix	System May Hang Under Complex Conditions
CFW131	X	X	Fixed	PEG PCIe* Link May Fail to Link After Resuming from PKG-C8
CFW132	X	X	No Fix	Incorrect Error Correcting Code (ECC) Reporting Following the Entry to PKG-C7
CFW133	X	X	No Fix	PMU MSR_UNC_PERF_FIXED_CTR is Cleared After Pkg C7 or Deeper
CFW134	X	X	No Fix	Performance Monitoring General Counter 2 May Have Invalid Value Written When Intel® TSX Is Enabled
CFW135	X	X	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set
CFW136	X	X	No Fix	Rare Internal Timing Conditions May Lead to Sporadic Hangs During Graphics VTd Flows



# Errata

---

## **CFW1. Reported Memory Type May Not Be Used to Access the Virtual-Machine Control Structure (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 Virtual Machine Extension (VMX) access to the VMCS or referenced data structures will instead use the memory type that the Memory-Type Range Registers (MTRRs) specify for the physical address of the access.

**Implication:** Bits [53:50] of the IA32\_VMX\_BASIC MSR report that the Write-Back (WB) 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 Tables of Changes"](#).

## **CFW2. Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation**

**Problem:** This erratum may cause a machine-check error (IA32\_MCi\_STATUS.MCACOD=005H with IA32\_MCi\_STATUS.MSCOD=00FH or IA32\_MCi\_STATUS.MCACOD=0150H with IA32\_MCi\_STATUS.MSCOD=00FH) on the fetch of an instruction. It applies only if (1) instruction bytes are fetched from a linear address translated using a 4-Kbyte page and cached in the processor; (2) the paging structures are later modified so that these bytes are translated using a large page (2 Mbyte, 4 Mbyte or 1 GByte) with a different physical address (PA), memory type (PWT, PCD and PAT bits), or User/Supervisor (U/S) bit; and (3) the same instruction is fetched after the paging structure modification but before software invalidates any Translation Lookaside Buffer (TLB) entries for the linear region.

**Implication:** Due to this erratum, an unexpected machine check with error code 0150H with MSCOD 00FH may occur, possibly resulting in a shutdown. This erratum could also lead to unexpected correctable machine check (IA32\_MCi\_STATUS.UC=0) with error code 005H with MSCOD 00FH.

**Workaround:** Software should not write to a paging-structure entry in a way that would change the page size and either the physical address, memory type or user/supervisor bit. It can instead use one of the following algorithms: first clear the P flag in the relevant paging-structure entry (for example, PDE); then invalidate any translations for the affected linear addresses; and then modify the relevant paging-structure entry to set the P flag and establish the new page size. An alternative algorithm: first change the physical page attributes (combination of physical address, memory type and user/supervisor bit) in all 4K pages in the affected linear addresses; then invalidate any translations for the affected linear addresses; and then modify the relevant paging-structure entry to establish the new page size.

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW3. 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 an invalid opcode exception (#UD) 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 device-not-available exception (#NM).

**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 Tables of Changes".

**CFW4. The Corrected Error Count Overflow Bit in IA32\_MC0\_STATUS is Not Updated When the UC Bit is Set**

**Problem:** After an Uncorrected error (UC) 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 Tables of Changes".

**CFW5. 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 Tables of Changes".



**CFW6. SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior**

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

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

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

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

**CFW7. x87 FPU Exception (#MF) May be Signaled Earlier Than Expected**

**Problem:** The 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 event 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 Tables of Changes".

**CFW8. Incorrect FROM\_IP Value For an Restricted Transactional Memory (RTM) Abort in Branch Trace Message (BTM) or Branch Trace Store (BTS) May be Observed**

**Problem:** During RTM operation when branch tracing is enabled using BTM or 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 Tables of Changes".

**CFW9. 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 Tables of Changes".

**CFW10. 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 Bit Scan Forward (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:** For the steppings affected, see the "Summary Tables of Changes".

**CFW11. #GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code**

**Problem:** During a General Protection Exception (#GP), 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 Tables of Changes".

**CFW12. 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 an invalid opcode exception (#UD). 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:** For the steppings affected, see the "Summary Tables of Changes".


**CFW13. 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 Safer Mode Extensions (SMXs) SENTER/SEXIT operation (GETSEC[SENTER] 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 Tables of Changes".



#### **CFW14. Intel® Processor Trace (Intel® PT) TIP.PGD May Not Have Target IP Payload**

**Problem:** When Intel® PT is enabled and a direct unconditional branch clears IA32\_RTIT\_STATUS.FilterEn (MSR 571H, bit 0), due to this erratum, the resulting TIP.PGD (Target IP Packet, Packet Generation Disable) may not have an IP payload with the target IP.

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

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

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

#### **CFW15. 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 invalid-opcode exception (#UD). 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 Tables of Changes".

#### **CFW16. 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 an invalid-opcode exception (#UD). If either the TS or EM flag bits in CR0 are set, a device-not-available exception (#NM) 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 Tables of Changes".

#### **CFW17. 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 Corrected Machine Check Error Interrupt (CMCI) signaling.

**Workaround:** None identified.

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

#### **CFW18. 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 Tables of Changes”.

### **CFW19. 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 Tables of Changes”.

### **CFW20. 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 (Peripheral Component Interconnect [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 Tables of Changes”.


### **CFW21. Intel® PT PSB+ Packets May Contain Unexpected Packets**

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

**CFW23. VM Entry That Clears TraceEn May Generate a FUP**

**Problem:** If VM entry clears Intel® PT IA32\_RTIT\_CTL.TraceEn (MSR 570H, bit 0) while PacketEn is 1, then a FUP will precede the Target IP Packet, Packet Generation Disable (TIP.PGD). VM entry can clear TraceEn if the VM-entry Model Specific Register (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 Tables of Changes".

**CFW24. 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 Tables of Changes".

**CFW25. ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK**

**Problem:** The 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 SGX Enclave Control Structures (SECS's) MISCELECT field.

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

**CFW26. 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 Tables of Changes".

**CFW27. ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero**

**Problem:** The Intel® SGX ENCLU[EREPORT] instruction may cause a general protection exception (#GP) 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 Tables of Changes".

**CFW28. 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 hyper-visor 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 Tables of Changes".

**CFW29. 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. The 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 Tables of Changes".


**CFW30. Intel® PT OVF May be Dropped After OVF**

**Problem:** Some Intel® PT Overflow (OVF) packets may not be followed by a FUP or TIP.PGE.

**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 Tables of Changes".

**CFW31. 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 general protection fault (#GP) 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:** For the steppings affected, see the "Summary Tables of Changes".

**CFW32. Graphics Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Hardware May Cache Invalid Entries**

**Problem:** The processor's graphics Input/Output (I/O) Memory Management Unit (IOMMU) may cache invalid Intel® VT-d context entries. This violates the Intel® VT-d specification for Hardware (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 Intel® VT-d context cache after any update of context table entries.

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

**CFW33. Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State**

**Problem:** Voltage glitch of up to 200 mV on the VREF signal lasting for about 1 mS may be observed when entering System S3 state. This violates the JEDEC\* Double Data Rate (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 Tables of Changes".

**CFW34. 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 (for example: following them only with an instruction that writes [E/R]SP).

**Workaround:** None identified.

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

**CFW35. ENCLS[EINIT] Instruction May Unexpectedly #GP**

**Problem:** When using Intel® SGX, the ENCLS[EINIT] instruction will incorrectly cause a General Protection Fault (#GP) 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 State Save Area (SSA). 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 Tables of Changes".

**CFW36. Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop**


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

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

**Workaround:** None identified.

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





### **CFW37. 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 Tables of Changes”.

### **CFW38. Branch Instructions May Initialize Intel® Memory Protection Extensions (Intel® MPX) Bound Registers Incorrectly**

**Problem:** Depending on the current Intel® MPX configuration, execution of certain branch instructions (near CALL, near RET, near JMP, and Jcc instructions) without a BND prefix (F2H) initialize the Intel® MPX bound registers. Due to this erratum, execution of such a branch instruction on a user-mode page may not use the Intel® 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 been executed, a bound-range (#BR) 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 Tables of Changes”.

### **CFW39. Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled**

**Problem:** If 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 Tables of Changes”.

**CFW40. Processor May Run Intel® Advanced Vector Extensions (Intel® AVX) Code Much Slower Than Expected**

**Problem:** After a C6 state exit, the execution rate of 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 Tables of Changes".

**CFW41. Intel® PT Buffer Overflow May Result in Incorrect Packets**

**Problem:** Under complex micro-architectural conditions, an Intel® PT OVF packet may be issued after the first byte of a multi-byte Cycle Count (CYC) 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 and a PSB+. It should then treat the two CYC packets as indicating an overflow.

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

**CFW42. 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 Tables of Changes".


**CFW43. #GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave**

**Problem:** When executing within an Intel® SGX enclave, a #GP may be delivered instead of a Debug exception (#DB) 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 Enclave Page Cache Map (EPCM) 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 Tables of Changes".



#### **CFW44. Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception**

**Problem:** Execution of VAESENCLAST with VEX.L= 1 should signal an invalid opcode exception (#UD); however, due to the erratum, a device-not-available exception (#NM) 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 Tables of Changes".

#### **CFW45. Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits**

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

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

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

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

#### **CFW46. 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 Tables of Changes".

#### **CFW47. 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 FPU Data Pointer (FDP). 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 Tables of Changes".

**CFW48. 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 Tables of Changes".

**CFW49. Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults**

**Problem:** 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 Direct Memory Access (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:** For the steppings affected, see the "Summary Tables of Changes".

**CFW50. Intel® PT CYCThresh Value of 13 is Not Supported**

**Problem:** Intel® PT 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-lead 1H). Due to this erratum, if CYCThresh is set to 13 then the CYC threshold will be 0 cycles instead of 4096 (2<sup>13</sup>-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 Tables of Changes".

**CFW51. 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 Tables of Changes".

**CFW52. Integrated Audio Codec May Not be Detected**

**Problem:** Integrated audio codec may lose power when Low-Power Single Pipe (LPSP) mode is enabled for an embedded DisplayPort\* (eDP\*) 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 Tables of Changes".

**CFW53. 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 Tables of Changes".

**CFW54. Incorrect Branch Predicted Bit in BTS/BTM Branch Records**

**Problem:** BTS and BTM send branch records to the Debug Store (DS) 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 Tables of Changes".

**CFW55. 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 Tables of Changes".

**CFW56. 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 Performance Monitoring Interrupt (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 Tables of Changes".

**CFW57. 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\_ENABLE MSR (1A0H), bit 0 set)
- An MPX #BR exception occurred on BNDLX/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 Tables of Changes".


**CFW58. 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 Tables of Changes".



### **CFW59. Instructions Fetch #GP After RSM During Intel® PT May Push Incorrect RFLAGS Value on Stack**

**Problem:** If Intel® PT is enabled, a general protection exception (#GP) 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 Tables of Changes".

### **CFW60. 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® SGX 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 Tables of Changes".

### **CFW61. MTF VM Exit on XBEGIN Instruction May Save State Incorrectly**

**Problem:** Execution of an XBEGIN instruction, while the Monitor Trap Flag (MTF) 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 Tables of Changes".

### **CFW62. 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:** For the steppings affected, see the "Summary Tables of Changes".

**CFW63. 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 Tables of Changes".

**CFW64. PEBS EventingIP Field May Be Incorrect Under Certain Conditions**

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

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

**Workaround:** None identified.

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

**CFW65. HWP's Guaranteed\_Performance Updated Only on Configurable TDP Changes**

**Problem:** According to Hardware P-states (HWP) 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 Thermal Design Power (TDP), Running Average Power Limit (RAPL), 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 Tables of Changes".

**CFW66. 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 PEBS or 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 Tables of Changes".



**CFW67. Intel® PT ToPA Performance Monitoring Interrupt (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® 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 Tables of Changes".

**CFW68. HWP's Maximum\_Performance Value is Reset to 0xFF**

**Problem:** According to 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 Tables of Changes".

**CFW69. HWP's Guaranteed\_Performance and Relevant Status/Interrupt May Be Updated More Than Once Per Second**

**Problem:** According to HWP 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 Tables of Changes".

**CFW70. 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 Tables of Changes"](#).

**CFW71. Hardware P-states (HWP) May Generate Thermal Interrupt While Not Enabled**

**Problem:** Due to this erratum, the conditions for 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 (for example: 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 Tables of Changes"](#).


**CFW72. 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 Tables of Changes"](#).

**CFW73. Attempts to Retrain a PCI Express\* (PCIe\*) Link May be Ignored**

**Problem:** A 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 Tables of Changes".

**CFW74. 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, the 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 Tables of Changes".

**CFW75. 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 a general protection exception (#GP) when the address is not canonical (for example: 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 Tables of Changes".

**CFW76. 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 Tables of Changes".

**CFW77. 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 Tables of Changes".

**CFW78. 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 a general protection exception (#GP) 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:** For the steppings affected, see the "Summary Tables of Changes".

**CFW79. Violations of Intel® SGX Access-Control Requirements Produce #GP Instead of #PF**

**Problem:** 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 Tables of Changes".


**CFW80. 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® 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 Tables of Changes".

**CFW81. 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® 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 Tables of Changes".

**CFW82. Display Slowness May be Observed Under Certain Display Commands Scenario**

**Problem:** Back to back access to the Video Graphics Array (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 Tables of Changes".

**CFW83. CPUID TLB Associativity Information is Inaccurate**

**Problem:** CPUID leaf 2 (EAX=02H) TLB information inaccurately reports that the shared second 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 Tables of Changes".

**CFW84. Unpredictable System Behavior May Occur in DDR4 Multi-Rank System**

**Problem:** Due to incorrect configuration of DDR4 On-Die Termination (ODT) by BIOS, it is possible for a multi-rank system to violate Section 4.27 of the *DDR4 JEDEC Specification*, 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 Tables of Changes".

**CFW85. 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 Tables of Changes".

**CFW86. Potential Partial Trace Data Loss in Intel® Trace Hub (Intel® TH) ODLA When Storing to Memory**

**Problem:** When the 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 Tables of Changes".

**CFW87. Using Different Vendors for 2666 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang**

**Problem:** When using 2666 MHz DDR4 UDIMMs from different vendors or mixing single rank and dual rank DIMMs, 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 and do not mix single rank and dual rank for 2666 MHz UDIMMs.

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


**CFW88. 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 Corrected Machine Check Interrupt (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 Tables of Changes".



### **CFW89. Reads From IA32\_SGXLEPUBKEYHASH MSRs Return Values in Incorrect Order**

**Problem:** The IA32\_SGXLEPUBKEYHASH [0,1,2,3] (8CH, 8DH, 8EH, 8FH) MSRs allow software to select an alternative Intel® SGX Launch Enclave provider signing key. Each of four MSRs is used to specify a 64-bit part of the 256 bit Intel® SGX Launch Enclave signing key hash, where the lower 32 bits of each 64-bit component are to be provided in the EAX register, and the higher 32 bits of each 64-bit component are to be provided in the EDX register. Due to this erratum, reads from IA32\_SGXLEPUBKEYHASH MSRs will return the lower 32 bits of the 64-bit component in the EDX register, and the higher 32 bits of the 64-bit component in the EAX register.

**Implication:** Software may incorrectly identify the currently active Intel® SGX Launch enclave provider on the platform.

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

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

### **CFW90. 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 Machine Check Exception (MCE) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.
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 I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

### **CFW91. MOVNTDQA From WC Memory May Pass Earlier MFENCE instructions**

**Problem:** An execution of (V)MOVNTDQA (streaming load instruction) that loads from a Write Combining (WC) memory may appear to pass an earlier MFENCE instruction.

**Implication:** Serialization of the (V)MOVNTDQA with earlier MFENCE may not be enforced.

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

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW92. MOVNTDQA From WC Memory May Pass Earlier Locked Instructions**

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

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

**Workaround:** Software should not rely on a locked instruction to fence subsequent executions of MOVNTDQA. Software should insert an MFENCE instruction if it needs to preserve order between streaming loads and other memory operations.

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW93. PEBS Record After a WRMSR to IA32\_BIOS\_UPDT\_TRIG May be Incorrect**

**Problem:** A PEBS record generated by a WRMSR to IA32\_BIOS\_UPDT\_TRIG MSR (79H) may have an incorrect value in the Eventing EIP field if an instruction prefix was used on the WRMSR.

**Implication:** The Eventing EIP field of the generated PEBS record may be incorrect. Intel has not observed this erratum with any commercially available software.

**Workaround:** Instruction prefixes have no architecturally-defined function for the WRMSR instruction; instruction prefixes should not be used with the WRMSR instruction.

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW94. 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 Tables of Changes"](#).

**CFW95. 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 Tables of Changes"](#).



**CFW96. Precise Performance Monitoring May Generate Redundant PEBS Records**

**Problem:** 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 Tables of Changes".

**CFW97. Load Latency Performance Monitoring Facility May Stop Counting**

**Problem:** The performance monitoring events MEM\_TRANS\_RETIRED.LOAD\_LATENCY\_\* (Event CDH; UMask 01H; any latency) count load instructions whose latency exceed a predefined threshold, where the loads are randomly selected using the Load Latency facility (PEBS extension). However, due to this erratum, load latency facility may stop counting load instructions when Intel® HT Technology is enabled.

**Implication:** Counters programmed with the affected events stop incrementing and do not generate PEBS records.

**Workaround:** None identified.

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

**CFW98. 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 Tables of Changes".

**CFW99. Branch Instruction Address May be Incorrectly Reported on Intel® Transactional Synchronization Extensions (Intel® TSX) Abort When Using Intel® MPX**

**Problem:** When using Intel® MPX, an 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 LBR, BTS and BTM as well as in the FUP source IP address for Intel® PT. Due to this erratum, the FROM\_IP field in LBR/BTS/BTM, as well as the 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 Tables of Changes"](#).

**CFW100. 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 Tables of Changes"](#).

**CFW101. Hitting a Code Breakpoint Inside an Intel® SGX Debug Enclave May Cause the Processor to Hang**

**Problem:** Under complex microarchitectural conditions, the processor may hang when hitting code breakpoint inside an Intel® SGX debug enclave. This may happen only after opt-out entry into an 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 Tables of Changes"](#).


**CFW102. Performance Monitoring Anti Side-Channel Interference (ASCI) Status Bit May Be Inaccurate**

**Problem:** The ASCI field in IA32\_PERF\_GLOBAL\_STATUS (MSR 38EH, bit 60) should be set when the count in any of the configured performance counters (for example: 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 Tables of Changes"](#).

**CFW103. 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 (for example: 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 Tables of Changes".

**CFW104. The Processor May Fail to Boot During DDR4 Memory Training**

**Problem:** The BIOS may fail to properly configure the required DQ/DQS timing parameters for certain DDR4 DIMMs with large length deltas between byte lanes (long fly-by topology).

**Implication:** An incorrect timing parameter value may cause DDR to mis-sample incoming data during write operations, leading to memory training failures and subsequent system boot hangs. Board designs with DDR4 DIMMs short delta lengths between byte lanes (short fly-by topology) are not impacted.

**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 Tables of Changes".

**CFW105. Intel® PT CYC Packet Can be Dropped When Immediately Preceding PSB**

**Problem:** Due to a rare microarchitectural condition, generation of an Intel® PT PSB packet can cause a single CYC packet, possibly along with an associated Mini Time Counter (MTC) 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 Tables of Changes".

**CFW106. Intel® PT VM-entry Indication Depends on the Incorrect VMCS Control Field**

**Problem:** An Intel® PT Paging Information Packet (PIP), which includes indication of entry into non-root operation, will be generated on VM-entry as long as the "Conceal VMX in Intel® PT" field (bit 19) in Secondary Execution Control register (IA32\_VMX\_PROCBASED\_CTL2, 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\_CTL2 MSR 0484H).

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

**Workaround:** A 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 Tables of Changes".

**CFW107. Certain DDR4 Memory Configurations May Cause Unpredictable System Behavior**

**Problem:** When using SODIMM or UDIMM in DDR4 2N CMD timing mode, the processor may incorrectly de-emphasize the first CMD bit transmitted.

**Implication:** When this erratum occurs, memory commands may not complete, potentially leading to system hang or 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 Tables of Changes".

**CFW108. VCVTPS2PH To Memory May Update MXCSR in the Case of a Fault on the Store**

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

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

**Workaround:** None identified.

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


**CFW109. Intel® PT May Drop All Packets After an Internal Buffer Overflow**

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

**Implication:** When this erratum occurs, all trace data will be lost until either Intel® 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 Tables of Changes".

**CFW110. 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 Tables of Changes".

**CFW111. Data Breakpoint May Not be Detected on a REP MOVS**

**Problem:** A REP MOVS 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 MOVS instruction.

**Implication:** A debugger may miss a data read/write access if it is done by a REP MOVS 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 Tables of Changes".

**CFW112. Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang**

**Problem:** If an Intel® PT ToPA table is placed in Uncacheable (UC) or Uncacheable Speculative Write Combining (USWC) memory, and a ToPA output region is filled during an 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 Tables of Changes".

**CFW113. 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 ToPA and that table is later read by the CPU during the HLE 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 Tables of Changes".

**CFW114. Intel® PT PSB+ Packets May be Omitted on a C6 Transition**

**Problem:** An Intel® PT PSB+ 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 Tables of Changes".

**CFW115. 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 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 Tables of Changes".


**CFW116. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions**

**Problem:** An access to a Guest-Physical Address (GPA) may cause an EPT-violation VM exit. When the "EPT-violation #VE" VM-execution control is 1, an EPT violation may cause a #VE (virtualization exception) instead of a VM exit. Due to this erratum, an EPT violation may erroneously cause a #VE when the "suppress #VE" bit is set in the EPT paging-structure entry used to map the GPA being accessed. This erratum does not apply when the "EPT-violation #VE" VM-execution control is 0 or when delivering an event through the Interrupt Descriptor Table (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 Tables of Changes".

**CFW117. Using Intel® TSX Instructions May Lead to Unpredictable System Behavior**

**Problem:** Under complex microarchitectural conditions, software using 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 Tables of Changes".

**CFW118. Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values**

**Problem:** When 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 Tables of Changes".

**CFW119. 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 packet may be dropped without an OVF packet.

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

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

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

**CFW120. Unexpected Uncorrected Machine Check Errors May Be Reported**

**Problem:** In rare microarchitectural 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 Tables of Changes".

**CFW121. Processor May Hang at High Temperature With a High-Throughput Graphics Workload**

**Problem:** In systems that support Single Core Turbo frequencies up to 4.7 GHz, the Ring frequency may operate at 4.4GHz. In this condition, if the processor is operating with a



high-throughput graphics workload and the processor is operating near maximum junction temperature, the processor may hang.

**Implication:** Due to this erratum the processor may hang. Intel has only observed this erratum under synthetic test conditions

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

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

#### **CFW122. Gen9 Graphics Intel® VT-d Hardware May Cache Invalid Entries**

**Problem:** The Gen9 graphics subsystem may cache invalid Intel® VT-d context entries.

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

**Workaround:** Software should flush the Gfx Intel® VT-d context cache after any update of context table entries.

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

#### **CFW123. Queued Invalidation Is Prevented When Intel® VT-d is Disabled**

**Problem:** While both the DMA-Remapping and Interrupt-Remapping capabilities are disabled in the Default Intel® VT-d Engine, then Queued Invalidation is incorrectly disabled.

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

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

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

#### **CFW124. Processor May Hang During PKG-C8/C9/C10 Exit**


**Problem:** Due to marginal voltage configuration of internal restore SRAM, the processor may hang.

**Implication:** When this erratum occurs, the processor May Hang. Intel has only observed this erratum in synthetic test conditions.

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

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



**CFW125. 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 Intel Reuse Repository [IRR] and Intel Strategic Research [ISR] APIC registers), the processor may dispatch the second interrupt (from the IRR bit) before the first interrupt has completed and written to the End-Of-Interrupt (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 Tables of Changes"](#).

**CFW126. Executing Some Instructions May Cause Unpredictable Behavior**

**Problem:** Under complex microarchitectural conditions, executing an X87, Intel® AVX, or integer divide instruction may result in unpredictable system behavior.

**Implication:** When this erratum occurs, the system may behave unpredictably. Intel has not observed this erratum with any 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 Tables of Changes"](#).

**CFW127. Incorrect Execution of Internal Branch Instructions May Lead to Unpredictable System Behavior**

**Problem:** Under complex microarchitectural conditions, incorrect execution of internal branch instructions that span multiple 64-byte boundaries (cross cache line), may result in unpredictable system behavior including unexpected page faults (#PF) or invalid-opcode exceptions (#UD) due to incorrect execution of internal branch operations.

**Implication:** When this erratum occurs, the system may exhibit unpredictable system behavior including unexpected #PF or #UD exceptions.

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

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW128. Unexpected Page Faults in Guest Virtualization Environment**

**Problem:** Under complex microarchitectural conditions, a virtualized guest could observe unpredictable system behavior.

**Implication:** When this erratum occurs, systems operating in a virtualization environment may exhibit unexpected page faults (double faults) leading to guest OS shutdown.

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

**Status:** For the steppings affected, see the ["Summary Tables of Changes"](#).

**CFW129. Intel® SGX Key Confidentiality May be Compromised**

**Problem:** Under complex microarchitectural conditions, it may be possible for the value of the Intel® SGX keys to be inferred using speculative execution side channel methods.

**Implication:** If exposed, such keys could allow an attacker to access Intel® SGX enclave data. Processors that do not support Intel® Hyper-Threading Technology (Intel® HT Technology) are not affected by this issue.

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

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

**CFW130. System May Hang Under Complex Conditions**

**Problem:** Under complex conditions, insufficient access control in graphics subsystem may lead to a system hang or crash upon a register read.

**Implication:** When this erratum occurs, a system hang or crash may occur.

**Workaround:** It is possible for a combination of BIOS and a graphics driver to contain a workaround for this erratum.

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

**CFW131. PEG PCIe\* Link May Fail to Link After Resuming from PKG-C8**

**Problem:** The PCI Express Graphics (PEG) IO registers may not be restored after resuming from PKG-C8.

**Implication:** The PEG PCIe\* may fail to link after resuming from PKG-C8.

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

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

**CFW132. Incorrect Error Correcting Code (ECC) Reporting Following the Entry to PKG-C7**

**Problem:** The correctable and uncorrectable ECC errors reported in ECCERRLOG0/1 (MCHBAR Offset 4048h/404Ch) may be overwritten after the entry to PKG-C7.

**Implication:** The DDR4 correctable and uncorrectable ECC errors reported in ECCERRLOG0/1 (MCHBAR Offset 4048h/404Ch) may be unreported after resuming from PKG-C7. Intel has only observed this erratum in a synthetic test environment.

**Workaround:** None identified.

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


**CFW133. PMU MSR\_UNC\_PERF\_FIXED\_CTR is Cleared After Pkg C7 or Deeper**

**Problem:** The Performance Monitoring Unit Uncore Performance Fixed Counter (MSR\_UNC\_PERF\_FIXED\_CTR [MSR 395h]) is cleared after pkg C7 or deeper.

**Implication:** Due to this erratum, once the system enters pkg C7 or deeper, the uncore fixed counter does not reflect the actual count.

**Workaround:** None identified.

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

**CFW134. Performance Monitoring General Counter 2 May Have Invalid Value Written When Intel® TSX Is Enabled**

**Problem:** When Intel® TSX is enabled and there are aborts (HLE or RTM) overlapping with the access or manipulation of the IA32\_PMC2 general-purpose performance counter (Offset: C3h), it may return invalid values.

**Implication:** Software may read invalid value from IA32\_PMC2.

**Workaround:** None Identified.

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

**CFW135. Overflow Flag in IA32\_MC0\_STATUS MSR May be Incorrectly Set**

**Problem:** Under complex micro-architectural conditions, a single internal parity error seen in IA32\_MC0\_STATUS MSR (401h) with MCACOD (bits 15:0) value of 5h and MSCOD (bits 31:16) value of 7h, may set the overflow flag (bit 62) in the same MSR.

**Implication:** Due to this erratum, the IA32\_MC0\_STATUS overflow flag may be set after a single parity error. Intel has not observed this erratum with any commercially available software.

**Workaround:** None identified.

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

**CFW136. Rare Internal Timing Conditions May Lead to Sporadic Hangs During Graphics VTd Flows**

**Problem:** When both Intel® SGX and Graphics RC6 features are enabled, under complex microarchitectural conditions, Intel® VT-d operations towards the Gfx IOMMU may lead to unexpected system behavior.

**Implication:** Due to this erratum, unexpected system behavior will occur.

**Workaround:** None identified.

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



# Specification Changes

---

There are no specification changes in this specification update revision.

§



# Specification Clarifications

---

There are no specification clarifications in this specification update revision.

§



# Documentation Changes

---

There are no documentation changes in this specification update revision.

§