

Intel® Xeon® E3-1200 v3 Processor Family

Specification Update

August 2020



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 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.

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/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-technology.html>

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>.

Intel, Pentium, Intel Core, Intel SpeedStep, 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.



Contents


Revision History	4
Preface	6
Summary Tables of Changes	8
Identification Information	15
Errata	17
Specification Changes	61
Specification Clarifications	62
Documentation Changes	63

§



Revision History

Revision	Description	Date
025	<ul style="list-style-type: none">Added erratum HSW182	August 2020
024	<ul style="list-style-type: none">Added errata HSW180 and HSW181	May 2020
023	<ul style="list-style-type: none">Added errata HSW178 and HSW179	January 2020
022	<ul style="list-style-type: none">ErrataAdded HSW177	August 2018
021	<ul style="list-style-type: none">ErrataAdded HSW176	July 2018
020	<ul style="list-style-type: none">ErrataUpdated HSW173	April 2018
019	<ul style="list-style-type: none">ErrataAdded HSW175	February 2018
018	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW173-HSW174Updated HSW135Removed HSW13, HSW150	March 2017
017	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW171-HSW172Updated HSW150	November 2016
016	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW163-HSW170Updated HSW55, HSW156Removed HSW63	October 2016
015	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW156-HSW162Updated HSW149, HSW155	April 2016
014	<ul style="list-style-type: none">ErrataAdded HSW154-HSW155	January 2016
013	<ul style="list-style-type: none">ErrataAdded HSW153	October 2015
012	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Removed HSW61Added HSW149-HSW152	August 2015
011	<ul style="list-style-type: none">ErrataAdded HSW148	April 2015
010	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW145-HSW147Updated HSW33	February 2015
009	<ul style="list-style-type: none">Added errata HSW141-HSW144	December 2014
008	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Removed HSW140Updated HSW131	September 2014
007	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW133-140Removed HSW75, HSW87, HSW114, and HSW130Added Specification Change HSW1	August 2014



Revision	Description	Date
006	<ul style="list-style-type: none">Added erratum HSW132Added Documentation Change HSW 2Updated link to Intel® 64 and IA-32 Architecture Software Developer's Manual Documentation	July 2014
005	<ul style="list-style-type: none">Updated erratum HSW29Added errata HSW109-131	June 2014
004	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Moved HSW100 to HSW108Added HSW100-107Updated Identification Information	October 2013
003	<ul style="list-style-type: none">Errata<ul style="list-style-type: none">Added HSW59-100Updated Identification Information	August 2013
002	<ul style="list-style-type: none">N/A. No Updates. Revision number added to Revision History to maintain consistency with NDA Specification Update numbering.	N/A
001	<ul style="list-style-type: none">Initial Release.	June 2013



Preface

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

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

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

Affected Documents

Document Title	Document Number
Intel® Xeon® Processor E3-1200 v3 Product Family Datasheet - Volume 1 of 2	328907
Intel® Xeon® Processor E3-1200 v3 Product Family Datasheet - Volume 2 of 2	329000

Related Documents

Document Title	Document Number/ Location
Intel® Architecture Instruction Set Extensions Programming Reference	https://software.intel.com/sites/default/files/m/9/2/3/41604
Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide Intel® 64 and IA-32 Intel Architecture Optimization Reference Manual	http://www.intel.com/products/processor/manuals/index.htm
Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes	https://software.intel.com/content/www/us/en/develop/download/intel-64-and-ia-32-architectures-software-developers-manual-documentation-changes.html
ACPI Specifications	www.acpi.info

Nomenclature

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

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, such as core speed, L2 cache size, package type, and so on, as described in the processor identification information table. Read all notes associated with each S-Spec number.



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

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

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

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



Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the processor. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations.

Codes Used in Summary Tables

Stepping

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 (Sheet 1 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW1	X	No Fix	Last Branch Record (LBR), Branch Trace Store (BTS), Branch Trace Message (BTM) May Report a Wrong Address when an Exception/Interrupt Occurs in 64-bit Mode
HSW2	X	No Fix	EFLAGS Discrepancy on Page Faults and on the Extended Page Table (EPT)-Induced VM (Virtual Machine) Exits after a Translation Change
HSW3	X	No Fix	MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error
HSW4	X	No Fix	Last Exception Record (LER) Model-Specific Register (MSR) May Be Unreliable
HSW5	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang
HSW6	X	No Fix	An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang
HSW7	X	No Fix	General Protection Exception (#GP) on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
HSW8	X	No Fix	FREEZE_WHILE_SMM Does Not Prevent Event From Pending Precise Event Based Sampling (PEBS) During SMM
HSW9	X	No Fix	Advanced Programmable Interrupt Controller (APIC) Error "Received Illegal Vector" May Be Lost
HSW10	X	No Fix	Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations
HSW11	X	No Fix	Performance Monitor Precise Instruction Retired Event May Present Wrong Indications
HSW12	X	No Fix	CR0.CD Is Ignored CR0. in the Virtual Machine Extensions (VMX) Operation
HSW13	X	No Fix	Erratum has been Removed.
HSW14	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a Device-Not-Available (#NM) Exception
HSW15	X	No Fix	Processor May Fail to Acknowledge a TLP Request
HSW16	X	No Fix	Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered
HSW17	X	No Fix	PCIe* Root-port Initiated Compliance State Transmitter Equalization Settings May Be Incorrect
HSW18	X	No Fix	PCIe* Controller May Incorrectly Log Errors on Transition to RxL0s
HSW19	X	No Fix	Unused PCIe* Lanes May Report Correctable Errors
HSW20	X	No Fix	Accessing Physical Memory Space 0-640K through the Graphics Aperture May Cause Unpredictable System Behavior
HSW21	X	No Fix	PCIe* Root Port May Not Initiate Link Speed Change
HSW22	X	No Fix	Pending x87 FPU Exceptions (#MF) May Be Signaled Earlier Than Expected
HSW23	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
HSW24	X	No Fix	VEX.L is Not Ignored with VCVT*2SI Instructions
HSW25	X	No Fix	N/A. Erratum has been removed
HSW26	X	No Fix	Specific Graphics Blitter Instructions May Result in Unpredictable Graphics Controller Behavior
HSW27	X	No Fix	Processor May Enter Shutdown Unexpectedly on a Second Uncorrectable Error
HSW28	X	No Fix	Modified Compliance Patterns for 2.5 GT/s and 5 GT/s Transfer Rates Do Not Follow PCIe* Specification
HSW29	X	No Fix	Performance Monitor Counters May Produce Incorrect Results
HSW30	X	No Fix	Performance Monitor UOPS_EXECUTED Event May Undercount
HSW31	X	No Fix	MSR_PERF_STATUS May Report an Incorrect Core Voltage
HSW32	X	No Fix	PCIe* Atomic Transactions From Two or More PCIe* Controllers May Cause Starvation



Errata (Sheet 2 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW33	X	No Fix	The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When the UC Bit is Set
HSW34	X	No Fix	An Intel® Advanced Vector Extensions (Intel® AVX) Gather Instruction that Causes an EPT Violation May Not Update Previous Elements
HSW35	X	No Fix	PLATFORM_POWER_LIMIT MSR Not Visible
HSW36	X	No Fix	LPDDR Memory May Report Incorrect Temperature
HSW37	X	No Fix	PCIe* Host Bridge DID May Be Incorrect
HSW38	X	No Fix	Time Stamp Counter (TSC) May Be Incorrect After a Deep C-State Exit
HSW39	X	No Fix	PCIe* Controller May Initiate Speed Change While in DL_Init State Causing Certain PCIe* Devices to Fail to Train
HSW40	X	No Fix	Spurious Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Interrupts May Occur When the Primary Fault Overflow (PFO) Bit is Set
HSW41	X	No Fix	N/A. Erratum has been removed
HSW42	X	No Fix	Intel® AVX Gather Instruction That Causes a Fault or VM Exit May Incorrectly Modify Its Destination Register
HSW43	X	No Fix	Inconsistent NaN Propagation May Occur When Executing (V)DPPS Instruction
HSW44	X	No Fix	Display May Flicker When Package C-States Are Enabled
HSW45	X	No Fix	Certain Combinations of Intel® AVX Instructions May Cause Unpredictable System Behavior
HSW46	X	No Fix	Processor May Incorrectly Estimate Peak Power Delivery Requirements
HSW47	X	No Fix	IA32_PERF_CTL MSR is Incorrectly Reset
HSW48	X	No Fix	Processor May Hang During a Function Level Reset of the Display
HSW49	X	No Fix	Intel® AVX Gather Instruction that Should Result in Double Fault (#DF) May Cause Unexpected System Behavior
HSW50	X	No Fix	Throttling and Refresh Rate May Be Incorrect After Exiting Package C-State
HSW51	X	No Fix	Processor May Livelock During On Demand Clock Modulation
HSW52	X	No Fix	IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI is Incorrectly Cleared by SMI
HSW53	X	No Fix	The From-IP for Branch Tracing May Be Incorrect
HSW54	X	No Fix	Thermal Monitor 1 (TM1) Throttling May Continue Indefinitely
HSW55	X	No Fix	Internal Parity Errors May Incorrectly Report Overflow in The IA32_MC2_STATUS MSR
HSW56	X	No Fix	Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX May Over Count
HSW57	X	No Fix	Processor May Run at Incorrect P-State
HSW58	X	No Fix	Performance Monitor Event DSB2MITE_SWITCHES.COUNT May Over Count
HSW59	X	No Fix	Performance Monitor Register UNC_PERF_GLOBAL_STATUS Not Restored on Package C7 Exit
HSW60	X	No Fix	Processor May Not Enter Package C6 or Deeper C-states when the PCIe* Links Are Disabled
HSW61	X	No Fix	<Erratum removed>
HSW62	X	No Fix	Some Performance Monitor Event Counts May Be Inaccurate During the SMT Mode
HSW63	X	No Fix	<Erratum removed>
HSW64	X	No Fix	The Upper 32 Bits of CR3 May Be Incorrectly Used With 32-Bit Paging
HSW65	X	No Fix	Performance Monitor Events HLE_RETIRED.ABORTED_MISC4 And RTM_RETIRED.ABORTED_MISC4 May Over Count
HSW66	X	No Fix	A PCIe* Latency Tolerance Report (LTR) Update Message May Cause The Processor to Hang



Errata (Sheet 3 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW67	X	No Fix	GETSEC Does Not Report Support For the Static Core Root of Trust for Measurement (S-CTRM)
HSW68	X	No Fix	EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly
HSW69	X	No Fix	APIC Timer Might Not Signal an Interrupt While in TSC-Deadline Mode
HSW70	X	No Fix	IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For the VMCS Encoding
HSW71	X	No Fix	Incorrect FROM_IP Value For a Restricted Transactional Memory (RTM) Abort in a BTM or a BTS May Be Observed
HSW72	X	No Fix	Intel® VT-d Hardware May Perform Set Root Table Pointer (SRTP) And Set Interrupt Remapping Table Pointer (SIRTP) Operations on a Package C7 Exit
HSW73	X	No Fix	General-Purpose Performance Counters Can Unexpectedly Increment
HSW74	X	No Fix	Performance Monitoring Events May Report Incorrect Number of Load Hits or Misses to the Last Level Cache (LLC)
HSW75	X	No Fix	N/A. Erratum has been removed
HSW76	X	No Fix	Locked Load Performance Monitoring Events May Under Count
HSW77	X	No Fix	Graphics Processor Ratio And C-State Transitions May Cause a System Hang
HSW78	X	No Fix	Certain Performance Monitoring Events May Over Count Software Demand Loads
HSW79	X	No Fix	Accessing Nonexistent Uncore Performance Monitoring MSRs May Not Signal a #GP
HSW80	X	No Fix	Call Stack Profiling May Produce Extra Call Records
HSW81	X	No Fix	Warm Reset May Fail or Lead to Incorrect Power Regulation
HSW82	X	No Fix	PCIe* Host Bridge DID May Be Incorrect
HSW83	X	No Fix	Transactional Abort May Produce an Incorrect Branch Record
HSW84	X	No Fix	System Management RAM (SMRAM) State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
HSW85	X	No Fix	DMA Remapping Faults for the Graphics Intel® VT-d Unit May not Properly Report Type of Faulted Request
HSW86	X	No Fix	Intel® AVX Gather Instructions Page Faults May Report an Incorrect Faulting Address
HSW87	X	No Fix	N/A. Erratum has been removed
HSW88	X	No Fix	Event Injection by VM Entry May Use an Incorrect B Flag for Stack Segment (SS)
HSW89	X	No Fix	A Fault in the SMM May Result in Unpredictable System Behavior
HSW90	X	No Fix	Processor Frequency is Unexpectedly Limited Below Nominal P1 When cTDP Down is Enabled
HSW91	X	No Fix	The PMI May Be Signaled More Than Once for Performance Monitor Counter Overflow
HSW92	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
HSW93	X	No Fix	DRAM Reported Power Consumption May Be Incorrect After a Warm Reset
HSW94	X	No Fix	RDRAND Execution in a Transactional Region May Cause a System Hang
HSW95	X	No Fix	Uncore Clock Frequency Changes May Cause Audio/Video Glitches
HSW96	X	No Fix	Processor May Experience a Spurious LLC-Related Machine Check During Periods of High Activity
HSW97	X	No Fix	The Processor May Not Enter Package C7 When Using a Panel Self Refresh (PSR) Display
HSW98	X	No Fix	Video or Audio Distortion May Occur
HSW99	X	No Fix	System May Hang When Audio is Enabled During Package C3
HSW100	X	No Fix	INVPCID May Not Cause an #UD in VMX Non-Root Operation
HSW101	X	No Fix	Non-Compliant PFAT Module Base Address May Cause Unpredictable System Behavior



Errata (Sheet 4 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW102	X	No Fix	Incorrect LBR Source Address May Be Reported For a Transactional Abort
HSW103	X	No Fix	Address Translation Faults for Intel® Virtualization Technology for Directed I/O (Intel® VT-d) May Not Be Reported for Display Engine Memory Accesses
HSW104	X	No Fix	L3 Cache Corrected Error Count May be Inaccurate After Package C7 Exit
HSW105	X	No Fix	PCIe* Device's SVID is Not Preserved Across The Package C7 C-State
HSW106	X	No Fix	Warm Reset Does Not Stop the GT Power Draw
HSW107	X	No Fix	Unused PCIe* Lanes May Remain Powered After Package C7
HSW108	X	No Fix	Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash
HSW109	X	No Fix	Processor Energy Policy Selection May Not Work as Expected
HSW110	X	No Fix	A PEBS Record May Contain Processor State for an Unexpected Instruction
HSW111	X	No Fix	MSR_PP1_ENERGY_STATUS Reports Incorrect Energy Data
HSW112	X	No Fix	x87 Floating Point Unit Data Pointer (DP) May Be Incorrect After Instructions That Save FP State to Memory
HSW113	X	No Fix	Processor May Hang During Package C7 Exit
HSW114	X	No Fix	N/A. Erratum has been removed
HSW115	X	No Fix	Spurious LLC Machine Check May Occur
HSW116	X	No Fix	Page Fault May Report Incorrect Fault Information
HSW117	X	No Fix	CATERR# Pin Assertion is Not Cleared on a Warm Reset
HSW118	X	No Fix	Uncorrectable Machine Check Error During Core C6 Entry May Not Be Signaled
HSW119	X	No Fix	The SAMPLE/PRELOAD JTAG Command Does Not Sample The Display Transmit Signals
HSW120	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests And Snoop Requests May Be Incorrect
HSW121	X	No Fix	Processor Energy Policy Selection May Not Work as Expected
HSW122	X	No Fix	PCIe* Link May Incorrectly Train to 8.0 GT/s
HSW123	X	No Fix	PCIe* Tx Voltage Reference Cannot Be Changed
HSW124	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
HSW125	X	No Fix	Re-Enabling eDRAM May Log a Machine Check and Hang
HSW126	X	No Fix	Warm Reset Does Not Stop EDRAM Power Draw
HSW127	X	No Fix	Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID
HSW128	X	No Fix	Intel® Smart 2D Display Technology (Intel® S2DDT) May not Function Correctly with Certain High Resolution Displays
HSW129	X	No Fix	Stateless GPGPU A32 Byte Scattered ReadWrite Message Operations May Result in Unpredictable System Behavior
HSW130	X	No Fix	N/A. Erratum has been removed
HSW131	X	No Fix	Spurious Corrected Errors May Be Reported
HSW132	X	No Fix	A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
HSW133	X	No Fix	Package C7 Power Consumption Has Been Observed to Be Higher Than Package C6
HSW134	X	No Fix	An Intel® Hyper-Threading Technology-enabled Processor May Exhibit Unpredictable Behavior During Power or Thermal Management Operations
HSW135	X	No Fix	Certain Perfmon Events May Be Counted Incorrectly When The Processor is Not in C0 State
HSW136	X	No Fix	Software Using Intel® TSX May Result in Unpredictable System Behavior



Errata (Sheet 5 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW137	X	No Fix	A Transient High Temperature Event May Cause Persistent Frequency Restrictions
HSW138	X	No Fix	Running All Cores May Incorrectly Limit the Processor Frequency
HSW139	X	No Fix	Concurrent Core and Graphics Operation at Turbo Ratios May Lead to System Hang
HSW140	X	No Fix	N/A. Erratum has been removed
HSW141	X	No Fix	Performance Monitor Instructions Retired Event May Not Count Consistently
HSW142	X	No Fix	Interactions Between Multiple Unaligned Memory Accesses And Locked Instructions May Lead to a Machine Check
HSW143	X	No Fix	Fixed-Function Performance Counter May Over Count Instructions Retired by 32 When Intel® Hyper-Threading Technology is Enabled
HSW144	X	No Fix	Performance Monitor UOPS_EXECUTED Event May Be Inaccurate When Using Intel® Hyper-Threading Technology
HSW145	X	No Fix	Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a Problem
HSW146	X	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
HSW147	X	No Fix	System May Hang or Video May Be Distorted After Graphics RC6 Exit
HSW148	X	No Fix	Certain eDP* Displays May Not Function as Expected
HSW149	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events May Undercount
HSW150	X	No Fix	Erratum has been Removed.
HSW151	X	No Fix	Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations
HSW152	X	No Fix	An Interrupt Return (IRET) Instruction That Results in a Task Switch Does Not Serialize The Processor
HSW153	X	No Fix	Attempting to Disable Turbo Mode May Cause a #GP
HSW154	X	No Fix	PECI Frequency Limited to 1 MHz
HSW155	X	No Fix	VGATHERQPS That Loads an Element From The APIC-Access Page May Load Other Elements From Incorrect Addresses
HSW156	X	No Fix	An APIC Timer Interrupt During Core C6 Entry May Be Lost
HSW157	X	No Fix	MTF VM Exit on XBEGIN Instruction May Save State Incorrectly
HSW158	X	No Fix	Uncore Performance Monitoring Counters May Be Disabled or Cleared After Package C7
HSW159	X	No Fix	PEBS Record May Be Generated After Being Disabled
HSW160	X	No Fix	PCIe* Ports Do Not Support Data Link Layer (DLL) Link Active Reporting
HSW161	X	No Fix	PCIe* Link Speed Negotiation May Fail After Link is Re-enabled
HSW162	X	No Fix	MOVNTDQA from WC Memory May Pass Earlier Locked Instructions
HSW163	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May Be Lost
HSW164	X	No Fix	A Corrected Internal Parity Error May Result in a System Hang
HSW165	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May Be Lost
HSW166	X	No Fix	Internal Power State Transitions May Cause the Graphics Device to Hang
HSW167	X	No Fix	SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
HSW168	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions
HSW169	X	No Fix	RF May Be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
HSW170	X	No Fix	Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes



Errata (Sheet 6 of 6)

Number	Steppings	Status	ERRATA
	C-0		
HSW171	X	No Fix	An x87 Store Instruction Which Pends #PE While EPT is Enabled May Lead to an Unexpected Machine Check and/or Incorrect x87 State Information
HSW172	X	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
HSW173	X	No Fix	Writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP May #GP When Intel® TSX is Not Supported
HSW174	X	No Fix	APIC Timer Interrupt May Not Be Generated at The Correct Time In TSC-Deadline Mode
HSW175	X	No Fix	Precise Performance Monitoring May Generate Redundant PEBS Records
HSW176	X	No Fix	In eMCA2 Mode, When the Retirement Watchdog Timeout Occurs CATERR# May Be Asserted
HSW177	X	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on the Store
HSW178	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation
HSW179	X	No Fix	System May Hang Under Complex Conditions
HSW180	X	No Fix	PMU MSR_UNC_PERF_FIXED_CTR Is Cleared After Pkg C7 or Deeper
HSW181	X	No Fix	Performance Monitoring General Counter 2 May Have Invalid Value Written When Intel® TSX Is Enabled
HSW182	X	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May Be Incorrectly Set

Specification Changes

Number	SPECIFICATION CHANGES
HSW1	Intel® Transactional Synchronization Extensions (Intel®TSX) Instruction

Specification Clarifications

Number	SPECIFICATION CLARIFICATIONS
	None for this revision of this specification update.

Documentation Changes

Number	DOCUMENTATION CHANGES
HSW1	On-Demand Clock Modulation Feature Clarification
HSW2	Intel® Virtualization Technology (Intel® VT) Clarification



Identification Information

Component Identification using Programming Interface

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

Table 1. Intel® Xeon® E3-1200 v3 Processor Product Family 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
	00000000b	0011b		00b	0110b	1100b	xxxxb

Notes:

1. The Extended Family Bits [27:20] are used in conjunction with the Family Code, specified in Bits [11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium 4, or Intel® Core™ processor family.
2. The Extended Model Bits [19:16], in conjunction with the Model Number specified in Bits [7:4], are used to identify the model of the processor within the processor's family.
3. The Family Code corresponds to Bits [11:8] of the Extended Data Register (EDX) after RESET, Bits [11:8] of the Extended Accumulator Register (EAX) after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
4. The Model Number corresponds to Bits [7:4] of the Extended Data Register (EDX) after RESET, Bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
5. The Stepping ID in Bits [3:0] indicates the revision number of that model. See the processor identification table for the processor stepping ID number in the CPUID information.

When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number, and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

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

The processor can be identified by the following register contents.

Stepping	Vendor ID ¹	Host Device ID ²	Processor Graphics Device ID ³	Revision ID ⁴	CRID
C-0	8086h	0C04h	GT2 = 0416h	06h	06h

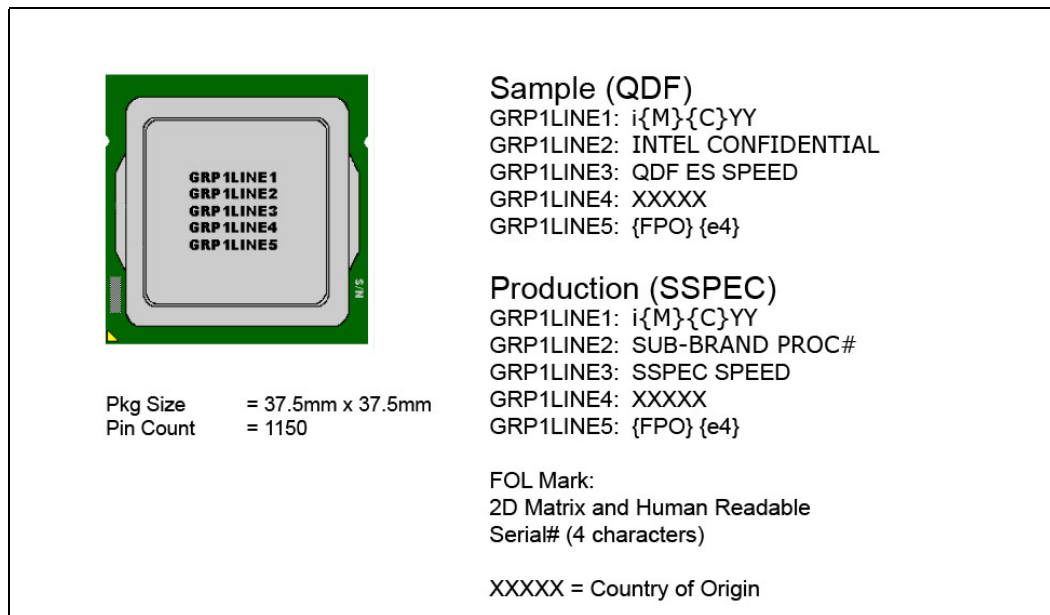
Notes:

1. The Vendor ID corresponds to bits [15:0] of the Vendor ID Register located at offset 00h–01h in the Process Capability Index (PCI) function 0 configuration space.
2. The Host Device ID corresponds to bits [15:0] of the Device ID Register located at Device 0 offset 02h–03h in the PCI function 0 configuration space.
3. The Processor Graphics Device ID (DID2) corresponds to bits [15:0] of the Device ID Register located at Device 2 offset 02h–03h in the PCI function 0 configuration space.
4. The Revision Number corresponds to bits [7:0] of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

Component Marking Information


The processor stepping can be identified by the following component markings.

Figure 1. Intel® Xeon® E3-1200 V3 Processor Product Family-Side Markings



For Intel® Xeon® Processor E3-1200 v3 Product Family SKUs, see:

<https://ark.intel.com/content/www/us/en/ark/products/series/78581/intel-xeon-processor-e3-v3-family.html>



Errata

HSW1. Last Branch Record (LBR), Branch Trace Store (BTS), Branch Trace Message (BTM) May Report a Wrong Address When an Exception/Interrupt Occurs in 64-bit Mode

Problem: An exception or interrupt event should be transparent to the LBR, the BTS and the BTM mechanisms. However, during a specific boundary condition where the exception or interrupt occurs right after the execution of an instruction at the lower canonical boundary (0x00007FFFFFFFFF) in 64-bit mode, the LBR return registers will save a wrong return address with bits 63 to 48 incorrectly sign extended to all 1's. Subsequent the BTS and the BTM operations which report the LBR will also be incorrect.

Implication: The LBR, the BTS and the BTM may report incorrect information in the event of an exception or interrupt.

Workaround: None identified.

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

HSW2. EFLAGS Discrepancy on Page Faults and on the Extended Page Table (EPT)-Induced VM (Virtual Machine) Exits after a Translation Change

Problem: This erratum is regarding the case where paging structures are modified to change a linear address from writable to non-writable without software performing an appropriate Translation Lookaside Buffer (TLB) invalidation. When a subsequent access to that address by a specific instruction (ADD, AND, BTC, BTR, BTS, CMPXCHG, DEC, INC, NEG, NOT, OR, ROL/ROR, SAL/SAR/SHL/SHR, SHLD, SHRD, SUB, XOR, and XADD) causes a page fault or an EPT-induced VM exit, the value saved for EFLAGS may incorrectly contain the arithmetic flag values that the EFLAGS register would have held had the instruction completed without fault or VM exit. For page faults, this can occur even if the fault causes a VM exit or if its delivery causes a nested fault.

Implication: None identified. Although the EFLAGS value saved by an affected event (a page fault or an EPT-induced VM exit) may contain incorrect arithmetic flag values, Intel has not identified software that is affected by this erratum. This erratum will have no further effects once the original instruction is restarted because the instruction will produce the same results as if it had initially completed without fault or VM exit.

Workaround: If the handler of the affected events inspects the arithmetic portion of the saved EFLAGS value, then system software should perform a synchronized paging structure modification and the TLB invalidation.

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


HSW3. MCI_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error

Problem: A single Data Translation Look Aside Buffer (DTLB) error can incorrectly set the Overflow (bit [62]) in the MCI_Status register. A DTLB error is indicated by MCA error code (bits [15:0]) appearing as binary value, 000x 0000 0001 0100, in the MCI_Status register.

Implication: Due to this erratum, the Overflow bit in the MCI_Status register may not be an accurate indication of multiple occurrences of DTLB errors. There is no other impact to normal processor functionality.

Workaround: None identified.

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

**HSW4. Last Exception Record (LER) Model-Specific Register (MSR) May Be Unreliable**

Problem: Due to certain internal processor events, updates to the LER, MSRs, MSR_LER_FROM_LIP (1DDH) and MSR_LER_TO_LIP (1DEH), may happen when no update was expected.

Implication: The values of the LER MSRs may be unreliable.

Workaround: None identified.

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

HSW5. MONITOR or CLFLUSH on the Local xAPIC's Address Space Results in Hang

Problem: If the target linear address range for a MONITOR or CLFLUSH is mapped to the local xAPIC's address space, the processor will hang.

Implication: When this erratum occurs, the processor will hang. The local xAPIC's address space must be uncached. The MONITOR instruction only functions correctly if the specified linear address range is of the type write-back. CLFLUSH flushes data from the cache. Intel has not observed this erratum with any commercially available software.

Workaround: Do not execute MONITOR or CLFLUSH instructions on the local xAPIC address space.

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

HSW6. An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang

Problem: Uncorrectable errors logged in IA32_CR_MC2_STATUS MSR (409H) may also result in a system hang causing an Internal Timer Error (MCACOD = 0x0400h) to be logged in another machine check bank (IA32_MCI_STATUS).

Implication: Uncorrectable errors logged in IA32_CR_MC2_STATUS can further cause a system hang and an Internal Timer Error to be logged.

Workaround: None identified.

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

HSW7. General Protection Exception (#GP) on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code

Problem: During a #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](#).

**HSW8. FREEZE_WHILE_SMM Does Not Prevent Event From Pending Precise Event Based Sampling (PEBS) During SMM**

Problem: In general, a PEBS record should be generated on the first count of the event after the counter has overflowed. However, IA32_DEBUGCTL_MSR.FREEZE_WHILE_SMM (MSR 1D9H, bit [14]) prevents performance counters from counting during the System Management Model (SMM). Due to this erratum, if

1. A performance counter overflowed before a System Management Interrupt (SMI)
2. A PEBS record has not yet been generated because another count of the event has not occurred
3. The monitored event occurs during SMM

then a PEBS record will be saved after the next Response Surface Method (RSM) instruction.

When FREEZE_WHILE_SMM is set, a PEBS should not be generated until the event occurs outside of SMM.

Implication: A PEBS record may be saved after an RSM instruction due to the associated performance counter detecting the monitored event during SMM; even when FREEZE_WHILE_SMM is set.

Workaround: None identified.

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

HSW9. Advanced Programmable Interrupt Controller (APIC) Error "Received Illegal Vector" May Be Lost

Problem: APIC may not update the Error Status Register (ESR) flag Received Illegal Vector bit [6] properly when an illegal vector error is received on the same internal clock that the ESR is being written (as part of the write-read ESR access flow). The corresponding error interrupt will also not be generated for this case.

Implication: Due to this erratum, an incoming illegal vector error may not be logged into ESR properly and may not generate an error interrupt.

Workaround: None identified.

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

HSW10. Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations

Problem: Under complex microarchitectural conditions, if software changes the memory type for data being actively used and shared by multiple threads without the use of semaphores or barriers, software may see load operations execute out of order.

Implication: Memory ordering may be violated. Intel has not observed this erratum with any commercially available software.

Workaround: Software should ensure pages are not being actively used before requesting their memory type be changed.

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

HSW11. Performance Monitor Precise Instruction Retired Event May Present Wrong Indications

Problem: When the Precise Distribution for Instructions Retired (PDIR) mechanism is activated (INST_RETIRE.ALL (event C0H, umask value 00H) on Counter 1 programmed in PEBS mode), the processor may return wrong PEBS or Performance Monitoring Interrupt (PMI) interrupts and/or incorrect counter values if the counter is reset with a Sample-After-Value (SAV) below 100 (the SAV is the counter reset value software programs in the MSR IA32_PMC1[47:0] in order to control interrupt frequency).

Implication: Due to this erratum, when using low SAV values, the program may get incorrect PEBS or PMI interrupts and/or an invalid counter state.

Workaround: The sampling driver should avoid using SAV<100.

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

HSW12. CR0.CD Is Ignored CR0. in the Virtual Machine Extensions (VMX) Operation

Problem: If CR0.CD=1, the Memory Type Range Register (MTRR) and the PAT should be ignored and the Un-cachable (UC) memory type should be used for all memory accesses. Due to this erratum, a logical processor in the VMX operation will operate as if CR0.CD=0 even if that bit is set to 1.

Implication: Algorithms that rely on cache disabling may not function properly in the VMX operation.

Workaround: Algorithms that rely on cache disabling should not be executed in the VMX root operation.

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

HSW13. Erratum has been Removed.

HSW14. Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a Device-Not-Available (#NM) Exception

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

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

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

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

HSW15. Processor May Fail to Acknowledge a TLP Request

Problem: When a PCI Express* (PCIe*) root port's receiver is in Receiver L0s power state and the port initiates a Recovery event, it will issue Training Sets to the link partner. The link partner will respond by initiating an L0s exit sequence. Prior to transmitting its own Training Sets, the link partner may transmit a TLP request. Due to this erratum, the root port may not acknowledge the TLP request.

Implication: After completing the Recovery event, the PCIe* link partner will replay the TLP request. The link partner may set a Correctable Error status bit, which has no functional effect.

Workaround: None identified.

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



HSW16. Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered

Problem: If the local-APIC timer's Current-Count Register (CCR) is 0, software should be able to determine whether a previously generated timer interrupt is being delivered by first reading the delivery-status bit in the Logic and Validation Technology (LVT) timer register and then reading the bit in the Interrupt-Request Register (IRR) corresponding to the vector in the LVT timer register. If both values are read as 0, no timer interrupt should be in the process of being delivered. Due to this erratum, a timer interrupt may be delivered even if the CCR is 0 and the LVT and IRR bits are read as 0. This can occur only if the Divide Configuration Register (DCR) is greater than or equal to 4. The erratum does not occur if software writes zero to the Initial Count Register before reading the LVT and IRR bits.

Implication: Software that relies on reads of the LVT and IRR bits to determine whether a timer interrupt is being delivered may not operate properly.

Workaround: Software that uses the local-APIC timer must be prepared to handle the timer interrupts, even those that would not be expected based on reading CCR and the LVT and IRR bits; alternatively, software can avoid the problem by writing zero to the Initial Count Register before reading the LVT and IRR bits.

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

HSW17. PCIe* Root-port Initiated Compliance State Transmitter Equalization Settings May Be Incorrect

Problem: If the processor is directed to enter PCIe* Polling.Compliance at 5.0 GT/s or 8.0 GT/s transfer rates, it should use the Link Control 2 Compliance Preset/De-emphasis field (bits [15:12]) to determine the correct de-emphasis level. Due to this erratum, when the processor is directed to enter Polling.Compliance from 2.5 GT/s transfer rate, it retains 2.5 GT/s de-emphasis values.

Implication: The processor may operate in Polling.Compliance mode with an incorrect transmitter de-emphasis level.

Workaround: None identified.

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

HSW18. PCIe* Controller May Incorrectly Log Errors on Transition to RxL0s

Problem: Due to this erratum, if a link partner transitions to RxL0s state within 20 ns of entering L0 state, the PCIe* controller may incorrectly log an error in "Correctable Error Status.Receiver Error Status" field (Bus 0, Device 2, Function 0, 1, 2 and Device 6, Function 0, offset 1D0H, bit 0).

Implication: Correctable receiver errors may be incorrectly logged. Intel has not observed any functional impact due to this erratum with any commercially available add-in cards.

Workaround: None identified.

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

HSW19. Unused PCIe* Lanes May Report Correctable Errors

Problem: Due to this erratum, during PCIe* link down configuration, unused lanes may report a Correctable Error Detected in Bus 0, Device 1, Function 0-2, and Device 6, Function 0, Offset 158H, Bit 0.

Implication: Correctable Errors may be reported by a PCIe* controller for unused lanes.

Workaround: None identified.

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

**HSW20. Accessing Physical Memory Space 0-640K through the Graphics Aperture May Cause Unpredictable System Behavior**

- Problem:** The physical memory space 0-640K, when accessed through the graphics aperture, may result in a failure for writes to complete or reads to return incorrect results.
- Implication:** A hang or functional failure may occur during graphics operation such as OGL or OCL conformance tests, 2D/3D games and graphics intensive application.
- 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](#).

HSW21. PCIe* Root Port May Not Initiate Link Speed Change

- Problem:** The PCIe* Base specification requires the upstream component to maintain the PCIe* link at the target link speed or the highest speed supported by both components on the link, whichever is lower. PCIe* root port will not initiate the link speed change without being triggered by the software, when the root port maximum link speed is configured to be 5.0 GT/s. System BIOS will trigger the link speed change under normal boot scenarios. However, the BIOS is not involved in some scenarios such as link disable/re-enable or secondary bus reset, and, therefore, the speed change may not occur unless initiated by the downstream component. This erratum does not affect the ability of the downstream component to initiate a link speed change. All known 5.0Gb/s-capable PCIe* downstream components have been observed to initiate the link speed change without relying on the root port to do so.
- Implication:** Due to this erratum, the PCIe* root port may not initiate a link speed change during some hardware scenarios, causing the PCIe* link to operate at a lower than expected speed. Intel has not observed this erratum with any commercially available platform.
- Workaround:** None identified.
- Status:** For the steppings affected, see the [Summary Tables of Changes](#).

HSW22. Pending x87 FPU Exceptions (#MF) May Be Signaled Earlier Than Expected

- Problem:** x87 instructions that trigger the #MF normally service interrupts before the #MF. Due to this erratum, if an instruction that triggers the #MF is executed while Enhanced Intel SpeedStep® Technology transitions, Intel® Turbo Boost Technology transitions, or Thermal Monitor events occur, the pending #MF may be signaled before pending interrupts are serviced.
- Implication:** Software may observe the #MF being-signalized before pending interrupts are serviced.
- Workaround:** None identified.
- Status:** For the steppings affected, see the [Summary Tables of Changes](#).



HSW23. 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 (that is, 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](#).

HSW24. VEX.L is Not Ignored with VCVT*2SI Instructions

Problem: The VEX.L bit should be ignored for the VCVTSS2SI, VCVTSD2SI, VCVTTSS2SI, and VCVTTSD2SI instructions; however, due to this erratum, the VEX.L bit is not ignored and will cause a #UD.

Implication: Unexpected #UDs will be seen when the VEX.L bit is set to 1 with VCVTSS2SI, VCVTSD2SI, VCVTTSS2SI, and VCVTTSD2SI instructions.

Workaround: Software should ensure that the VEX.L bit is set to 0 for all scalar instructions.

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

HSW25. N/A. Erratum has been removed

HSW26. Specific Graphics Blitter Instructions May Result in Unpredictable Graphics Controller Behavior

Problem: Specific source-copy blitter instructions in Intel® HD Graphics 4600 Processor may result in unpredictable behavior when a blit source and destination overlap.

Implication: Due to this erratum, the processor may exhibit unpredictable graphics controller behavior. 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](#).

HSW27. Processor May Enter Shutdown Unexpectedly on a Second Uncorrectable Error

Problem: If an IA32_MCI_STATUS MSR contains an uncorrectable error with MCACOD=0x406 and a second uncorrectable error occurs after warm reset but before the first error is cleared by zeroing the IA32_MCI_STATUS MSR, a shutdown will occur.

Implication: When this erratum occurs, the processor will unexpectedly shut down instead of executing the machine check handler.

Workaround: None identified. Software should clear IA32_MCI_STATUS MSRs as early as possible to minimize the possibility of this erratum occurring.

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

**HSW28. Modified Compliance Patterns for 2.5 GT/s and 5 GT/s Transfer Rates Do Not Follow PCIe* Specification**

Problem: The PCIe* controller does not produce the PCIe* specification defined sequence for the Modified Compliance Pattern at 2.5 GT/s and 5 GT/s transfer rates. This erratum is not seen at 8 GT/s transfer rates.

Implication: Normal PCIe* operation is unaffected by this erratum.

Workaround: None identified.

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

HSW29. Performance Monitor Counters May Produce Incorrect Results

Problem: When operating with the Simultaneous Multi Threading (SMT) enabled, a memory at-retirement performance monitoring event (from the list below) may be dropped or may increment an enabled event on the corresponding counter with the same number on the physical core's other thread rather than the thread experiencing the event. Processors with the SMT disabled in the BIOS are not affected by this erratum.

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

MEM_UOP_RETIREDD.LOADS
MEM_UOP_RETIREDD.STORES
MEM_UOP_RETIREDD.LOCK
MEM_UOP_RETIREDD.SPLIT
MEM_UOP_RETIREDD.STLB_MISS
MEM_LOAD_UOPS_RETIREDD.HIT_LFB
MEM_LOAD_UOPS_RETIREDD.L1_HIT
MEM_LOAD_UOPS_RETIREDD.L2_HIT
MEM_LOAD_UOPS_RETIREDD.L3_HIT
MEM_LOAD_UOPS_L3_HIT_RETIREDD.XSNP_HIT
MEM_LOAD_UOPS_L3_HIT_RETIREDD.XSNP_HITM
MEM_LOAD_UOPS_L3_HIT_RETIREDD.XSNP_MISS
MEM_LOAD_UOPS_L3_HIT_RETIREDD.XSNP_NONE
MEM_LOAD_UOPS_RETIREDD.L3_MISS
MEM_LOAD_UOPS_L3_MISS_RETIREDD.LOCAL_DRAM
MEM_LOAD_UOPS_L3_MISS_RETIREDD.REMOTE_DRAM
MEM_LOAD_UOPS_RETIREDD.L2_MISS

Implication: Due to this erratum, certain performance monitoring event will produce unreliable results during hyper-threaded operation.

Workaround: None identified.

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

**HSW30. Performance Monitor UOPS_EXECUTED Event May Undercount**

Problem: The performance monitor event UOPS_EXECUTED (Event B1H, any Unmask) should count the number of UOPs executed each cycle. However, due to this erratum, when eight UOPs execute in one cycle, these UOPs will not be counted.

Implication: The performance monitor event UOPS_EXECUTED may reflect a count lower than the actual number of events.

Workaround: None identified.

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

HSW31. MSR_PERF_STATUS May Report an Incorrect Core Voltage

Problem: The core operating voltage can be determined by dividing MSR_PERF_STATUS MSR (198H) bits [47:32] by 2^{13} . However, due to this erratum, this calculation may report half the actual core voltage.

Implication: The core operating voltage may be reported incorrectly.

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](#).

HSW32. PCIe* Atomic Transactions From Two or More PCIe* Controllers May Cause Starvation

Problem: On a Processor PCIe* controller configuration in which two or more controllers receive concurrent atomic transactions, a PCIe* controller may experience starvation, which eventually can lead to a completion timeout.

Implication: Atomic transactions from two or more PCIe* controllers may lead to a completion timeout. Atomic transactions from only one controller will not be affected by this erratum. Intel has not observed this erratum with any commercially available device.

Workaround: None identified.

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

HSW33. The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When the UC Bit is Set

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

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

Workaround: None identified.

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

HSW34. An Intel® Advanced Vector Extensions (Intel® AVX) Gather Instruction that Causes an EPT Violation May Not Update Previous Elements

Problem: When execution of an Intel® AVX gather instruction causes an Extended Page Table (EPT) violation due to a specific element, all previous elements should be complete. Due to this erratum, such an execution may fail to complete previous elements. In addition, the instruction's mask operand is not updated. This erratum applies only if the EPT violation occurs while updating an accessed or dirty flag in a paging-structure entry. Instructions impacted by this erratum are: VGATHERDPS, VGATHERDPD, VGATHERQPS, VGATHERQPD, VPGATHERDD, VPGATHERDQ, VPGATHERQD, and VPGATHERQQ.

Implication: This erratum may prevent a gather instruction from making forward progress.

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](#).

HSW35. PLATFORM_POWER_LIMIT MSR Not Visible

Problem: The PLATFORM_POWER_LIMIT MSR (615H) is used to control the Power limit 3 (PL3) mechanism of the processor. Due to this erratum, this MSR is not visible to software.

Implication: Software is unable to read or write the PLATFORM_POWER_LIMIT MSR. If software attempts to access this MSR, a general protection fault will 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](#).

HSW36. LPDDR Memory May Report Incorrect Temperature

Problem: When any of the four possible LPDDR ranks are not populated, the unpopulated ranks will report a default temperature of 85C as a three bit value of 011b. If the system has unpopulated ranks, the temperature of memory will be reported as 85C in PCU_CR_DDR_DIMM_HOTTEST_ABSOLUTE (MCHBAR Bus 0; Device 0; Function 0; offset 58B8H) in bits [5:7], until any of the populated ranks report a higher temperature than this.

Implication: When the memory temperature is less than or equal to 85C it may be reported as 85C. This erratum does not affect DDR3 and DDR3L memory types.

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](#).

HSW37. PCIe* Host Bridge DID May Be Incorrect

Problem: The PCIe* Host Bridge DID register (Bus 0; Device 0; Offset 2H) contents may be incorrect after a Package C7 exit.

Implication: Software that depends on the Host Bridge DID value may not behave as expected after a Package C7 exit.

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](#).

HSW38. Time Stamp Counter (TSC) May Be Incorrect After a Deep C-State Exit

Problem: On exiting from Package C6 or deeper, the processor may incorrectly restore the TSC.

Implication: Software using the TSC may produce incorrect result and/or 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](#).

HSW39. PCIe* Controller May Initiate Speed Change While in DL_Init State Causing Certain PCIe* Devices to Fail to Train

Problem: The PCIe* controller supports hardware autonomous speed change capabilities. Due to this erratum, the PCIe* controller may initiate speed change while in the DL_Init state which may prevent link training for certain PCIe* devices.

Implication: Certain PCIe* devices may fail to complete DL_Init causing the PCIe* link to fail to train.

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](#).

HSW40. Spurious Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) Interrupts May Occur When the Primary Fault Overflow (PFO) Bit is Set

Problem: When the PFO field (bit [0] in the Intel® VT for Directed I/O (Intel® VT-d) FSTS [Fault Status] register) is set to 1, further faults should not generate an interrupt. Due to this erratum, further interrupts may still occur.

Implication: Unexpected Invalidation Queue Error interrupts may occur. Intel has not observed this erratum with any commercially available software.

Workaround: Software should be written to handle spurious Intel® VT-d fault interrupts.

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

HSW41. N/A. Erratum has been removed

HSW42. Intel® AVX Gather Instruction That Causes a Fault or VM Exit May Incorrectly Modify Its Destination Register

Problem: An execution of a 128-bit Intel® AVX gather instruction zeros the upper 128 bits of the instruction's destination register, unless access to the first unmasked element causes a fault or VM exit. Due to this erratum, these bits may be cleared even when accessing the first unmasked element causes a fault or VM exit. Instructions impacted by this erratum are: VGATHERDPS, VGATHERDPD, VGATHERQPS, VGATHERQPD, VPGATHERDD, VPGATHERDQ, VPGATHERQD, and VPGATHERQQ.

Implication: Software that depends on the destination register of a 128-bit Intel® AVX gather instruction to remain unchanged after access of the first unmasked element results in fault or VM exit 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](#).

HSW43. Inconsistent NaN Propagation May Occur When Executing (V)DPPS Instruction

Problem: Upon completion of the (V)DPPS instruction with multiple different NaN encodings in the input elements, software may observe different NaN encodings in the destination elements.

Implication: Inconsistent NaN encodings in the destination elements for the (V) DPPS instruction may be observed.

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](#).

**HSW44. Display May Flicker When Package C-States Are Enabled**

Problem: When package C-States are enabled, the display may not be refreshed at the correct rate.

Implication: When this erratum occurs, the user may observe flickering on the display.

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](#).

HSW45. Certain Combinations of Intel® AVX Instructions May Cause Unpredictable System Behavior

Problem: Execution of certain combinations of Intel® AVX instructions may lead to unpredictable system behavior.

Implication: When this erratum occurs, unpredictable system behaviors, including system hang or incorrect results can 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](#).

HSW46. Processor May Incorrectly Estimate Peak Power Delivery Requirements

Problem: Under certain conditions, the processor may incorrectly calculate the frequency at which the cores and graphics engine can operate, while still meeting voltage regulator and power supply peak power delivery capabilities. When this occurs, combined with high power workloads, system shutdown may be observed.

Implication: When this erratum occurs, system shutdown may be observed under high power workloads.

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](#).

HSW47. IA32_PERF_CTL MSR is Incorrectly Reset

Problem: The IA32_PERF_CTL MSR (199H) is not initialized correctly after a processor reset.

Implication: If software reads the IA32_PERF_CTL MSR before writing it, software can observe an incorrect reset value. Although incorrect values are reported to software, the correct default values for this register are still used by the processor. No performance or power impact occurs due to this erratum.

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](#).


HSW48. Processor May Hang During a Function Level Reset of the Display

Problem: When package C-States are enabled, it is possible that the processor may hang when software performs a Function Level Reset of the display via bit 1 of the Advanced Features Control Register (Bus 0; Device 2; Function 0; Offset 0A8H).

Implication: When this erratum occurs, the processor may hang.

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](#).



HSW49. Intel® AVX Gather Instruction that Should Result in Double Fault (#DF) May Cause Unexpected System Behavior

Problem: Due to this erratum, an execution of a 128-bit Intel® AVX gather instruction may fail to generate a #DF when expected. Instructions impacted by this erratum are: VGATHERDPS, VGATHERDPD, VGATHERQPS, VGATHERQPD, VPGATHERDD, VPGATHERDQ, VPGATHERQD, and VPGATHERQQ.

Implication: When this erratum occurs, an operation which should cause a #DF, may result in unexpected system behavior.

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](#).

HSW50. Throttling and Refresh Rate May Be Incorrect After Exiting Package C-State

Problem: When the Open Loop Thermal Management (OLMT) feature is enabled, the DIMM thermal status reported in DDR_THERM_PERDIMM_STATUS (MCHBAR Offset 588CH) may be incorrect following an exit from Package C3 or deeper.

Implication: The incorrect DIMM thermal status may result in degraded performance from unneeded memory throttling and excessive DIMM refresh rates.

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](#).

HSW51. Processor May Livelock During On Demand Clock Modulation

Problem: The processor may livelock when: (1) A processor thread has enabled on demand clock modulation via bit 4 of the IA32_CLOCK_MODULATION MSR (19AH) and the clock modulation duty cycle is set to 12.5% (02H in bits 3:0 of the same MSR). (2) The other processor thread does not have on demand clock modulation enabled and that thread is executing a stream of instructions with the lock prefix that either split a cacheline or access UC memory.

Implication: Program execution may stall on both threads of the core subject to this erratum.

Workaround: This erratum will not occur if clock modulation is enabled on all threads when using on demand clock modulation or if the duty cycle programmed in the IA32_CLOCK_MODULATION MSR is 18.75% or higher.

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

HSW52. IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI is Incorrectly Cleared by SMI

Problem: FREEZE_PERFMON_ON_PMI (bit 12) in the IA32_DEBUGCTL MSR (1D9H) is erroneously cleared during delivery of an SMI.

Implication: As a result of this erratum, the performance monitoring counters will continue to count after a PMI occurs in SMM.

Workaround: None identified.

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

HSW53. The From-IP for Branch Tracing May Be Incorrect

Problem: The BTM and the BTS report the "From-IP" indicating the source address of the branch instruction. Due to this erratum, the BTM and the BTS may repeat the "From-IP" value previously reported. The "To-IP" value is not affected.

Implication: Using the BTM or the BTS reports to reconstruct program execution may be unreliable.

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](#).

HSW54. Thermal Monitor 1 (TM1) Throttling May Continue Indefinitely

Problem: TM1 throttling may continue when the processor's temperature decreases below the throttling point while the processor is in Package C3 or deeper.

Implication: The processor will continue thermal throttling but does not indicate if it is hot.

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](#).

HSW55. Internal Parity Errors May Incorrectly Report Overflow in The IA32_MC2_STATUS MSR

Problem: Due to this erratum, uncorrectable internal parity error reports with an IA32_MC2_STATUS.MCACOD (bits [15:0]) value of 0005H and an IA32_MC2_STATUS.MSCOD (bits [31:16]) value of 0004H may incorrectly set the IA32_MC2_STATUS.OVER flag (bit 62) indicating an overflow even when only a single error has been observed.

Implication: IA32_MC2_STATUS.OVER may not accurately indicate multiple occurrences of uncorrectable internal parity errors. There is no other impact to normal processor functionality.

Workaround: None identified.

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

HSW56. Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX May Over Count

Problem: The Performance Monitor events OTHER_ASSISTS.AVX_TO_SSE (Event C1H; Umask 08H) and OTHER_ASSISTS.SSE_TO_AVX (Event C1H; Umask 10H) incorrectly increment and over count when a Hardware Lock Elision (HLE) abort occurs.

Implication: The Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX may over count.

Workaround: None identified.

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

HSW57. Processor May Run at Incorrect P-State

Problem: The processor package may use stale software Performance State (P-State) requests when one or more logical processors are idle.

Implication: The processor package may run at a higher or lower than expected P-State. This issue may persist as long as any logical processor is idle.

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](#).



HSW58. Performance Monitor Event DSB2MITE_SWITCHES.COUNT May Over Count

Problem: The Performance Monitor Event DSB2MITE_SWITCHES.COUNT (Event ABH; Umask 01H) should count the number of the Decode Stream Buffer (DSB) to the Macro Instruction Translation Engine (MITE) switches. Due to this erratum, the DSB2MITE_SWITCHES.COUNT event will count speculative switches and cause the count to be higher than expected.

Implication: The Performance Monitor Event DSB2MITE_SWITCHES.COUNT may report count higher than expected.

Workaround: None identified.

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

HSW59. Performance Monitor Register UNC_PERF_GLOBAL_STATUS Not Restored on Package C7 Exit

Problem: MSR_UNC_PERF_GLOBAL_STATUS (392H) is a global status register which indicates the overflow of uncore performance monitor counters. The content of this register is lost in package C7 state.

Implication: If any uncore performance monitor counter has overflowed before entering the package C7 state, the MSR_UNC_PERF_GLOBAL_STATUS register will no longer reflect the overflow after exiting C7 state.

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](#).

HSW60. Processor May Not Enter Package C6 or Deeper C-states when the PCIe* Links Are Disabled

Problem: If the PCIe* links are disabled via Link Disable (Bus 0, Device 1, Functions [2:1], Offset B0h, bit 4) and the PCIe* controller is enabled (Bus 0, Device 0, Function 0, Offset 54h, bits [2:1] = 11), then the processor will be unable to enter Package C6 or deeper C-states.

Implication: Due to this erratum, the process will not enter Package C6 or deeper C-states.

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](#).

HSW61. <Erratum removed>

HSW62. Some Performance Monitor Event Counts May Be Inaccurate During the SMT Mode

Problem: The performance monitor event OFFCORE_REQUESTS_OUTSTANDING (Event 60H, any Umask Value) should count the number of occurrences that loads or stores stay in the super queue each cycle. The performance monitor event CYCLE_ACTIVITY.CYCLES_L2_PENDING (Event A3H, Umask 01H) should count the number of cycles that demand loads stay in the super queue. However, due to this erratum, these events may count inaccurately during the SMT mode.

Implication: The performance monitor events OFFCORE_REQUESTS_OUTSTANDING and CYCLE_ACTIVITY.L2_PENDING may be unreliable during SMT Mode.

Workaround: None identified.

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

HSW63. <Erratum removed>
HSW64. The Upper 32 Bits of CR3 May Be Incorrectly Used With 32-Bit Paging

Problem: When 32-bit paging is in use, the processor should use a page directory located at the 32-bit physical address specified in bits 31:12 of CR3; the upper 32 bits of CR3 should be ignored. Due to this erratum, the processor will use a page directory located at the 64-bit physical address specified in bits 63:12 of CR3.

Implication: The processor may use an unexpected page directory or, if the EPT is in use, cause an unexpected EPT violation. This erratum applies only if software enters 64-bit mode, loads CR3 with a 64-bit value, and then returns to 32-bit paging without changing CR3. Intel has not observed this erratum with any commercially available software.

Workaround: Software that has executed in 64-bit mode should reload CR3 with a 32-bit value before returning to 32-bit paging.

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

HSW65. Performance Monitor Events HLE_RETIREDA.BORTED_MISC4 And RTM_RETIREDA.BORTED_MISC4 May Over Count

Problem: The Performance Monitor Events HLE_RETIREDA.BORTED_MISC4 (Event C8H; Umask 40H) and RTM_RETIREDA.BORTED_MISC4 (Event C9H; Umask 40H) are defined to count the number of transactional aborts due to incompatible memory types. Due to this erratum, they may count additional unrelated transactional aborts.

Implication: The Performance Monitor Events HLE_RETIREDA.BORTED_MISC4 and RTM_RETIREDA.BORTED_MISC4 counts may be greater than the number of aborts due to incompatible memory types. This can result in nonzero counts when all memory types are compatible.

Workaround: None identified.

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

HSW66. A PCIe* Latency Tolerance Report (LTR) Update Message May Cause The Processor to Hang

Problem: If a PCIe* device sends an LTR update message while the processor is in a package C6 or deeper, the processor may hang.

Implication: Due to this Erratum, the processor may hang if a PCIe* LTR update message is received while in a Package C6 or deeper.

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](#).

HSW67. GETSEC Does Not Report Support For the Static Core Root of Trust for Measurement (S-CTRM)

Problem: Processors with Intel® Boot Guard Technology that has GETSEC[PARAMETERS] leaf 5 Extended Accumulator Register (EAX) bit 5 set indicates support for processor rooted S-CTRM. Due to this erratum, that bit will not be set even though processor rooted S-CTRM is supported.

Implication: Software may be unaware of support for processor rooted S-CTRM.

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](#).



HSW68. EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly

Problem: If a memory access to a linear address requires the processor to update an accessed or dirty flag in a paging-structure entry, and if that update causes an EPT violation, the processor should store the linear address into the "guest linear address" field in the Virtual-Machine Control Structure (VMCS). Due to this erratum, the processor may store an incorrect value into bits 11:0 of this field. (The processor correctly stores the guest-physical address of the paging-structure entry into the "guest-physical address" field in the VMCS.)

Implication: Software may not be easily able to determine the page offset of the original memory access that caused the EPT violation. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: Software requiring the page offset of the original memory access address can derive it by simulating the effective address computation of the instruction that caused the EPT violation.

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

HSW69. APIC Timer Might Not Signal an Interrupt While in TSC-Deadline Mode

Problem: If the APIC timer is in TSC-deadline mode and is armed when a timed MWAIT instruction is executed, the timer expiration might not cause an interrupt.

Implication: Software depending on APIC timer TSC-deadline mode interrupts 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](#).

HSW70. IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For the VMCS Encoding

Problem: IA32_VMX_VMCS_ENUM MSR (48AH) bits 9:1 report the highest index value used for any VMCS encoding. Due to this erratum, the value 21 is returned in bits 9:1 although there is a VMCS field whose encoding uses the index value 23.

Implication: Software that uses the value reported in IA32_VMX_VMCS_ENUM[9:1] to read and write all VMCS fields may omit one field.

Workaround: None identified.

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

HSW71. Incorrect FROM_IP Value For a Restricted Transactional Memory (RTM) Abort in a BTM or a BTS May Be Observed

Problem: During the RTM operation when branch tracing is enabled using the BTM or 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](#).

HSW72. Intel® VT-d Hardware May Perform Set Root Table Pointer (SRTP) And Set Interrupt Remapping Table Pointer (SIRTP) Operations on a Package C7 Exit

Problem: On a package C7 exit, Intel® VT-d hardware may spuriously perform the SRTP and the SIRTP operations. A package C7 exit can cause the value programmed by software in the RTA_REG (IRTA_REG) to be visible to hardware before software executes a GCMD.SRTP command. This will result in hardware using the new values for the Direct Media Access (DMA) and interrupt translation page-walks, possibly before they are intended to be used by software.

Implication: If software has updated the root table pointer but has not executed the SRTP command, then the root table pointer update will happen unexpectedly, causing the VMM to walk incorrect or non-existent tables. Intel has not observed this erratum with any commercially available software.

Workaround: Privileged software should not execute a MWAIT (because it can trigger a package C7 entry/exit) between writing to RTA_REG (IRTA_REG) and GCMD_REG.SRTP (GCMD_REG.SIRTP) registers.

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

HSW73. General-Purpose Performance Counters Can Unexpectedly Increment

Problem: A performance monitor event programmed in a general-purpose performance counter should count the number of occurrences of the event selected in IA32_PERFVTSEL{0-7} MSR (186H-18DH). If the Invert (INV, bit 23) is set to 1 and a non-zero Counter Mask (CMASK) bits [31:24] value is used, due to this erratum, the event may over count in the case that either of Operating System (OS mode, bit 17) or User mode (USR, bit 16) is selected. Over counting will occur for the cycles spent in the non-matching Current Privilege Level (CPL).

Implication: General-purpose performance counters may reflect counts higher than the actual number of events when the INV bit is set, the CMASK is a non-zero value and either the OS or USR bit is set.

Workaround: None identified.

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

HSW74. Performance Monitoring Events May Report Incorrect Number of Load Hits or Misses to the Last Level Cache (LLC)

Problem: The following performance monitor events should count the numbers of loads hitting or missing the LLC. However due to this erratum, The L3_hit related events may over count and the L3_miss related events may undercount.

MEM_LOAD_RETIRED.L3_HIT (Event D1H, Umask 40H)

MEM_LOAD_RETIRED.L3_MISS (Event D1H, Umask 20H)


MEM_LOAD_L3_HIT_RETIRED.XSNP_NONE (Event D2H, Umask 08H)

MEM_LOAD_LLC_MISS_RETIRED.LOCAL_DRAM (Event D3H, Umask 01H)

Implication: The listed performance monitoring events may be inaccurate.

Workaround: None identified.

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



HSW75. N/A. Erratum has been removed

HSW76. Locked Load Performance Monitoring Events May Under Count

Problem: The performance monitoring events MEM_TRANS_RETIRE.LOAD_LATENCY (Event CDH; Umask 01H), MEM_LOAD_RETIRE.L2_HIT (Event D1H; Umask 02H), and MEM_UOPS_RETIRE.LOCKED (Event DOH; Umask 20H) should count the number of locked loads. Due to this erratum, these events may under count for locked transactions that hit the L2 cache.

Implication: The above event count will under count on locked loads hitting the L2 cache.

Workaround: None identified.

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

HSW77. Graphics Processor Ratio And C-State Transitions May Cause a System Hang

Problem: If ratio or C-state changes involving the processor core and processor graphics occur at the same time or while processor graphics are active under certain internal conditions, the ratio change may not complete.

Implication: The system may hang during C-state or ratio changes.

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](#).

HSW78. Certain Performance Monitoring Events May Over Count Software Demand Loads

Problem: The following performance monitor events should count the number of software demand loads. However due to this erratum, they may also include requests from the Next Page Prefetcher and over count.

OFFCORE_REQUESTS_OUTSTANDING.DEMAND_DATA (Event 60H; Umask 01H)

OFFCORE_REQUESTS.DEMAND_DATA (Event B0H; Umask 01H)

CYCLE_ACTIVITY.L2_Pending (Event A3H; Umask 01H)

L2_HIT_MISS.LOAD (Event 24H; Umask 01H)

Implication: The listed performance monitoring events may reflect a count higher than the actual number of events.

Workaround: None identified.

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

HSW79. Accessing Nonexistent Uncore Performance Monitoring MSRs May Not Signal a #GP

Problem: An access to an uncore Performance Monitor MSR beyond the number reported in the MSR_UNC_CBO_CONFIG MSR (396H) bits[3:0] should signal a #GP. Due to this erratum, the processor may hang instead of signaling a #GP.

Implication: When software accesses nonexistent uncore performance monitoring MSRs, the logical processor may hang instead of signaling a #GP.

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](#).

HSW80. Call Stack Profiling May Produce Extra Call Records

Problem: The performance monitoring Call Stack Profiling function should not generate call records for “zero length calls” (call instructions targeting the location following the instruction). However, due to this erratum, the processor will produce call records for zero length calls.

Implication: The performance monitoring the LBR call stack MSRs are incorrect in the presence of “zero length calls” because calls and returns do not match.

Workaround: None identified.

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

HSW81. Warm Reset May Fail or Lead to Incorrect Power Regulation

Problem: Due to this erratum, after a warm reset, the processor may fail to boot properly or may cause power to be regulated to an incorrect level.

Implication: The processor may not be able to control the Voltage Regulator (VR) to advertised specifications, leading to in a system hang, a machine check, or improper power regulation.

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](#).

HSW82. PCIe* Host Bridge DID May Be Incorrect

Problem: The PCIe* Host Bridge DID register (Bus 0; Device 0; Function 0; Offset 2H) contents may be incorrect.

Implication: Software that depends on the Host Bridge DID value 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](#).

HSW83. Transactional Abort May Produce an Incorrect Branch Record

Problem: If an Intel® Transactional Synchronization Extensions (Intel® TSX) transactional abort event occurs during a string instruction, the From-IP in the LBR is not correctly reported.

Implication: Due to this erratum, an incorrect From-IP on the LBR stack may be observed.

Workaround: None identified.

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

HSW84. System Management RAM (SMRAM) State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior

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

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

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

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

**HSW85. DMA Remapping Faults for the Graphics Intel® VT-d Unit May not Properly Report Type of Faulted Request**

Problem: When a fault occurs during DMA remapping of Graphics accesses at the Graphics Intel® VT-d unit, the type of faulted request (read or write) should be reported in bit 126 of the FRCD_REG register in the remapping hardware memory map register set. Due to this erratum, the request type may not be reported correctly.

Implication: Software processing the DMA remapping faults may not be able to determine the type of faulting graphics device DMA request.

Workaround: None identified.

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

HSW86. Intel® AVX Gather Instructions Page Faults May Report an Incorrect Faulting Address

Problem: If the software modifies a paging-structure entry to relax the access rights for a linear address and does not perform a TLB invalidation, a subsequent execution of an Intel® AVX gather instruction that accesses that address may generate a page fault that loads CR2 (which should contain the faulting linear address) with an incorrect value.

Implication: Software handling an affected page fault may not operate correctly.

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](#).

HSW87. N/A. Erratum has been removed**HSW88. Event Injection by VM Entry May Use an Incorrect B Flag for Stack Segment (SS)**

Problem: The stack accesses made by VM-entry event injection may use an incorrect value for the B flag (default stack-pointer size and upper bound) for the SS.

Implication: An affected stack access may use an incorrect address or an incorrect segment upper bound. This may result in unpredictable system behavior.

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

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

HSW89. A Fault in the SMM May Result in Unpredictable System Behavior

Problem: The value of the SS register as well as the CPL may be incorrect following a fault in SMM. The erratum can occur only if a fault occurs following an SMI and before the software has loaded the SS register (for example, with the MOV SS instruction).

Implication: This erratum may cause unpredictable system behavior. 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](#).

**HSW90. Processor Frequency is Unexpectedly Limited Below Nominal P1 When cTDP Down is Enabled**

Problem: When the Configurable Thermal Design Power (cTDP) Down is enabled on a processor branded as Core[®] i3 or Pentium[®], the processor frequency will be limited to the cTDP Down P1 frequency (Max Non-Turbo Frequency) when it should be able to operate between the cTDP Down frequency P1 and the nominal P1 frequency.

Implication: When cTDP is enabled, the processor cannot achieve expected frequencies.

Workaround: None identified.

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

HSW91. The PMI May Be Signaled More Than Once for Performance Monitor Counter Overflow

Problem: Due to this erratum, the PMI may be repeatedly issued until the counter overflow bit is cleared in the overflowing counter.

Implication: Multiple PMIs may be received when a performance monitor counter overflows.

Workaround: None identified. If the PMI is programmed to generate a Non-Maskable Interrupt (NMI), software may delay the End-of-Interrupt (EOI) register write for the interrupt until after the overflow indications have been cleared.

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

HSW92. Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception

Problem: Attempt to use FXSAVE or FXRSTOR with a VEX prefix should produce a #UD exception. If either the TS or EM flag bits in CR0 are set, a #NM exception will be raised instead of #UD exception.

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

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

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

HSW93. DRAM Reported Power Consumption May Be Incorrect After a Warm Reset

Problem: The DRAM power consumption reported by MSR_DRAM_Energy_Status (619H) bits [31:0] is reset to zero a short time after RESET# is de-asserted rather than during RESET# assertion. Due to this erratum, the MSR_DRAM_ENERGY_STATUS register will not increase monotonically after reset.

Implication: PECI reads to the Package Configuration Space (PCS) for DDR Energy Status (Index 4) may appear to change abruptly shortly after a warm reset.

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](#).

HSW94. RDRAND Execution in a Transactional Region May Cause a System Hang

Problem: Execution of the Random Number Generator (RDRAND) instruction inside an Intel[®] TSX transactional region may cause the logical processor to hang.

Implication: A system hang may occur as a result of this erratum.

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](#).

**HSW95. Uncore Clock Frequency Changes May Cause Audio/Video Glitches**

Problem: On some processors, the time required to change the uncore clock frequency may be large enough to significantly lengthen the latency of I/O requests to memory, possibly resulting in audio or video glitches.

Implication: Audio/Video glitches may occur during uncore ratio changes.

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](#).

HSW96. Processor May Experience a Spurious LLC-Related Machine Check During Periods of High Activity

Problem: Due to certain internal conditions while running core and memory intensive operations, some processors may incorrectly report an LLC related machine check with a IA32_MCI_STATUS.MCACOD value of 110AH.

Implication: Due to this erratum, the processor may experience a machine check.

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](#).

HSW97. The Processor May Not Enter Package C7 When Using a Panel Self Refresh (PSR) Display

Problem: The processor datasheet specifies that entering package C7 requires enabling the PSR for certain display resolutions, along with other conditions. Due to this erratum, the processor may not enter package C7 when connected to a PSR-enabled display even if all of the required conditions are met.

Implication: Due to this erratum, the processor may not enter package C7.

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](#).

HSW98. Video or Audio Distortion May Occur

Problem: Due to this erratum, internal processor operations can occasionally delay the completion of memory read requests enough to cause video or audio streaming underrun.

Implication: Visible artifacts such as flickering on a video device or glitches on audio 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](#).

HSW99. System May Hang When Audio is Enabled During Package C3

Problem: When audio is enabled while in package C3 state or deeper, audio memory traffic continues to be generated. Due to this erratum, the processor logic required for memory traffic may be powered down.

Implication: When this erratum occurs, the processor logic required for audio memory traffic may not be operational, resulting in a system hang.

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](#).

**HSW100. INVPCID May Not Cause a #UD in VMX Non-Root Operation**

Problem: The INVPCID instruction should cause a #UD in VMX non-root operation if either bit 31 of the primary processor-based VM-execution controls (activate secondary controls) or bit 12 of the secondary processor-based VM-execution controls (enable INVPCID) is 0. Due to this erratum, the INVPCID instruction will not cause #UD if “activate secondary controls” is 0 and “enable INVPCID” is 1. Instead, the instruction will either execute normally or cause a VM exit if the “INVLPG exiting” VM-execution control is 1.

Implication: The processor may cause a VM exit that software does not expect. 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](#).

HSW101. Non-Compliant PFAT Module Base Address May Cause Unpredictable System Behavior

Problem: The Platform Firmware Armoring Technology (PFAT) requires the PFAT module base address be 256 KB aligned and reside in the first 4 GB of memory. If the BIOS does not comply with these requirements when setting up the PFAT module, the processor should GP# at PFAT launch. Due to this erratum, a #GP fault may not be generated.

Implication: A PFAT module that does not follow the PFAT module base address requirements may result in unpredictable system behavior.

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

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

HSW102. Incorrect LBR Source Address May Be Reported For a Transactional Abort

Problem: If the fetch of an instruction in a transactional region causes a fault, a transactional abort occurs. If the LBRs are enabled, the source address recorded for such a transactional abort is the address of the instruction being fetched. If that instruction was itself the target of an earlier branch instruction, this erratum may erroneously record the address of the branch instruction as the source address for the transactional abort.

Implication: Trace reconstruction software that uses the LBR information may fail when this erratum occurs.

Workaround: None identified

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


HSW103. Address Translation Faults for Intel® Virtualization Technology for Directed I/O (Intel® VT-d) May Not Be Reported for Display Engine Memory Accesses

Problem: The Intel® VT-d hardware unit supporting the Processor Graphics device (Bus 0; Device 2; Function 0) may not report address translation faults detected on Display Engine memory accesses when the Context Cache is disabled or during time periods when Context Cache is being invalidated.

Implication: Due to this erratum, Display Engine accesses that fault are correctly aborted but may not be reported in the FSTS_REG fault reporting register (GFXVTDBAR offset 034H).

Workaround: None identified

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

**HSW104. L3 Cache Corrected Error Count May Be Inaccurate After Package C7 Exit**

Problem: The corrected error count for L3 cache errors reported in IA32_MCi_STATUS. Corrected Error Count (bits [52:38]) with an MCACOD of 0001 0001 xxxx xxxx (x can be 0 or 1) may be incorrectly restored to a smaller value during exit from Package C7.

Implication: The corrected error count for L3 cache errors in IA32_MCi_STATUS may be inaccurate after Package C7 exit.

Workaround: None identified.

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

HSW105. PCIe* Device's SVID is Not Preserved Across The Package C7 C-State

Problem: Bus 0, Device 7, Function 0's Subsystem Vendor Identification (SVID) register (Subsystem Vendor Identification, Offset 2CH) is not preserved across package C7 C-State transitions.

Implication: This may cause the operating system to think the device has been replaced with a different device.

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](#).

HSW106. Warm Reset Does Not Stop the GT Power Draw

Problem: Due to this erratum, if the GT is enabled prior to a warm reset, it will remain powered after the warm reset. The processor will make incorrect power management decisions because it assumes the GT is not drawing power after a warm reset.

Implication: The processor may draw more current than expected from an external VR. The processor may also put the external VR into a low power state where it will be unable to supply the sufficient power resulting in unpredictable system behavior.

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

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

HSW107. Unused PCIe* Lanes May Remain Powered After Package C7

Problem: If a PCIe* controller is enabled and either has unused lanes or no PCIe* device is present, the link or unused lanes should enter a low power state. Due to this erratum, after exiting Package C7, the unused link or unused lanes may remain powered.

Implication: Power consumption may be greater 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](#).

HSW108. Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash

Problem: If a logical processor has EPT enabled, is using 32-bit Physical Address Extension (PAE) paging, and accesses the virtual-APIC page, then a complex sequence of internal processor micro-architectural events may cause an incorrect address translation or machine check on either logical processor.

Implication: This erratum may result in unexpected faults, an uncorrectable TLB error logged in IA32_MCi_STATUS.MCACOD (bits [15:0]) with a value of 0000_0000_0001_xxxx (where x stands for 0 or 1), a guest or hypervisor crash, or other unpredictable system behavior.

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



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

HSW109. Processor Energy Policy Selection May Not Work as Expected

Problem: When the IA32_ENERGY_PERF_BIAS MSR (1B0H) is set to a value of 4 or more, the processor will try to increase the energy efficiency of Turbo mode. However, this functionality is effectively disabled if the software requested P-state exceeds the maximum P-state supported by the processor. This has the effect of decreasing the energy efficiency of the processor while in Turbo mode.

Implication: When this erratum occurs, reduced battery life and reduced energy efficiency may occur.

Workaround: The BIOS should set the max ACPI _PST object to the max supported turbo ratio, ensuring that the software P-state request does not exceed the maximum ratio supported by the processor. Note that this workaround will disable Core Ratio Overclocking.

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

HSW110. A PEBS Record May Contain Processor State for an Unexpected Instruction

Problem: If a performance counter has overflowed and is configured for PEBS, the processor will arm the PEBS hardware within a bounded number of cycles called the skid (see the discussion of skid and related topics in the Precise Distribution of Instructions Retired section of the *Intel® 64 and IA-32 Architectures Software Developer Manual*). Once the PEBS hardware is armed, the processor should capture processor state in a PEBS record following the execution of the next instruction that causes the counter to increment (a “triggering” instruction). Due to this erratum, the capture of processor state may occur at an instruction after the first triggering instruction following the skid but not beyond the second triggering instruction after the skid.

Implication: A PEBS record may contain processor state (including instruction pointer) not associated with the triggering instruction.

Workaround: None identified.

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

HSW111. MSR_PP1_ENERGY_STATUS Reports Incorrect Energy Data

Problem: The MSR_PP1_ENERGY_STATUS MSR (641H) bits [31:0] reports incorrect energy data.

Implication: Due to this erratum, reported Intel Integrated Graphics domain energy consumption may not be accurate.

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](#).


HSW112. x87 Floating Point Unit Data Pointer (DP) May Be Incorrect After Instructions That Save FP State to Memory

Problem: Under certain conditions, the value of the x87 FPU DP saved by the FSAVE/FNSAVE, FSTENV/FNSTENV, FXSAVE, XSAVE, or XSAVEOPT instructions may be incorrect.

Implication: Due to this erratum, the x87 FPU DP may be incorrect.

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](#).



HSW113. Processor May Hang During Package C7 Exit

Problem: Under certain internal timing conditions, the processor might not properly exit package C7 leading to a hang.

Implication: Due to this erratum, the package C7 state may not be reliable. Intel has not observed this erratum with any commercially available system.

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](#).

HSW114. N/A. Erratum has been removed

HSW115. Spurious LLC Machine Check May Occur

Problem: Under certain stressful conditions while running at ring ratios higher than 30, the processor may experience a spurious LLC machine check as indicated by IA32_MCi_STATUS.MCACOD (bits [15:0]) with value 000x 0001 0000 1010 (where x is 0 or 1).

Implication: When this erratum occurs, an uncorrectable LLC error will be logged and the system may hang or restart.

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](#).

HSW116. Page Fault May Report Incorrect Fault Information

Under the following conditions:

1. A read-modify-write instruction's memory source or destination (for example, ADD memory, reg) crossing a cache line boundary.
2. That instruction executing without fault.
3. While the read-modify-write instruction is executing, one or more of the following page table attributes associated with its memory operand are modified:
 - a. the Dirty (D) flag was 0 when the instruction was initiated but was concurrently set to 1, or
 - b. one of the relevant Read/Write (R/W) flags was 0 when the instruction was initiated but was concurrently set to 1, or
 - c. if the read-modify-write instruction executes at CPL = 3 and one of the relevant User/Supervisor (U/S) flags was 0 when the instruction was initiated but was concurrently set to 1.
4. A subsequent instruction executing within a narrow timing window that experiences a page fault.
5. There is no serializing instruction between the read-modify-write instruction and the faulting instruction.

The page fault (in #4) may report an incorrect error code and faulting linear address. These would describe the read-modify-write instruction's memory access instead of that of the faulting instruction. (The address of the faulting instruction is reported correctly.)

Implication: The erratum makes it appear that the page fault resulted from an access that occurred prior to the faulting instruction. Because the earlier access completed without faulting, a page-fault handler may identify the page fault as transient (or spurious) and re-execute the faulting instruction (for example, by executing an Interrupt Return [IRET]). In such cases, the erratum will not recur; the page fault on the later access will recur

and will be reported correctly. If the page-fault handler does not re-execute the faulting instruction, this erratum may result in unpredictable system behavior. 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](#).

HSW117. CATERR# Pin Assertion is Not Cleared on a Warm Reset

Problem: If the CATERR# pin is held asserted to indicate a fatal error, a subsequent warm reset event will not cause the CATERR# pin to de-assert.

Implication: When this erratum occurs, platforms that monitor the CATERR# pin may be unable to detect a fatal error after a warm reset or may incorrectly respond to a CATERR# pin assertion although an error may not have occurred subsequent to the warm reset event.

Workaround: The CATERR# pin can be de-asserted by a cold reset event.

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

HSW118. Uncorrectable Machine Check Error During Core C6 Entry May Not Be Signaled

Problem: Machine Check exceptions occurring during core C6 entry may be ignored.

Implication: When this erratum occurs, incorrect state may be saved during core C6 entry and subsequently restored during core C6 exit, resulting in unpredictable system behavior.

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

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

HSW119. The SAMPLE/PRELOAD JTAG Command Does Not Sample The Display Transmit Signals

Problem: The Display Transmit signals are not correctly sampled by the SAMPLE/PRELOAD JTAG* Command, violating the Boundary Scan specification (IEEE 1149.1).

Implication: The SAMPLE/PRELOAD command cannot be used to sample Display Transmit signals.

Workaround: None identified.

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

HSW120. Performance Monitor Event For Outstanding Offcore Requests And Snoop 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](#).

**HSW121. Processor Energy Policy Selection May Not Work as Expected**

Problem: When the IA32_ENERGY_PERF_BIAS MSR (1B0H) is set to a value of 4 or more, the processor will try to increase the energy efficiency of Turbo mode. However, this functionality is effectively disabled if the software requested P-state exceeds the maximum P-state supported by the processor.

Implication: When this erratum occurs, the energy efficiency control 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](#).

HSW122. PCIe* Link May Incorrectly Train to 8.0 GT/s

Problem: During PCIe* 8.0 GT/s Phase 2 Equalization training, the received per-lane transmitter coefficients for physical lanes 8-15 may be incorrectly applied to the PCIe* transmitters.

Implication: Due to this erratum, a PCIe* link may either fail to train to the 8.0 GT/s transfer speed, experience link errors, or periodically retrain (possibly dropping to a lower link speed).

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum. The BIOS may bypass PCIe* Gen 3.0 Phase 2 Equalization training by setting EQPH2BYP in bit 15 of PCIE_CR_EQCFG_0_1_0_MMR (Device 1; Function 0; Offset 0DD8H) to 1.

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

HSW123. PCIe* Tx Voltage Reference Cannot Be Changed

Problem: PCIe* Tx Voltage Reference Select is available via the PCIE_CR_AFEbND[0:7]CFG1 (Device 1; Function 0) registers in field TxVrefSel bits [9:5]. Due to this erratum, changes to these values will have no effect.

Implication: For PCIe*, setting the Tx Voltage Reference Select to non-default values will not produce the reference levels documented in the register description. Tx swing control utilizes Tx Voltage Reference; Tx swing cannot be adjusted from default.

Workaround: None identified.

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

HSW124. 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](#).

HSW125. Re-Enabling eDRAM May Log a Machine Check and Hang

Problem: If the eDRAM was disabled as a result of a package C-State entry of C2 or higher or a software request, the subsequent package C-state exit or software request to re-enable eDRAM may result in a machine check logged in IA32_MCI_STATUS.MCACOD [15:0] with of value 402H and subsequent system hang.

Implication: Due to this erratum, the system may log a machine check and hang.

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](#).

HSW126. Warm Reset Does Not Stop EDRAM Power Draw

Problem: Due to this erratum, if the EDRAM is enabled prior to a warm reset, it will remain powered after the warm reset. The processor will make incorrect power management decisions because it assumes the EDRAM is not drawing power after a warm reset.

Implication: The processor may draw more current than expected from an external VR. The processor may also put the external VR into a low power state where it will be unable to supply the sufficient power resulting in unpredictable system behavior.

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

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

HSW127. 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 Bit Scan Forward (BSF). Due to this erratum, opcode bytes F3 0F BC may execute as TZCNT even if CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 0.

Implication: Software that expects REP prefix before a BSF instruction to be ignored may not operate correctly, since there are cases in which the BSF and the 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 the TZCNT (and not BSF) is desired.

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

HSW128. Intel® Smart 2D Display Technology (Intel® S2DDT) May not Function Correctly with Certain High Resolution Displays

Problem: A limitation in Intel® S2DDT, commonly known as frame buffer compression, may result in pixel data being supplied too slowly to the display.

Implication: Screen flickering or blank screen may be observed on certain high resolution displays.

Workaround: The Intel® Graphics Driver contains a workaround for this erratum.

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

HSW129. Stateless GPGPU A32 Byte Scattered ReadWrite Message Operations May Result in Unpredictable System Behavior

Problem: A stateless General Purpose Computing on Graphics Processing Units (GPGPU) Byte Scattered ReadWrite Message Operation that immediately follows any thread performing GPGPU Byte Scattered Message Operation with state may be incorrectly handled as an access with state leading to unpredictable system behavior.

Implication: This erratum may result in unpredictable system behavior. The stateless model for A32 Byte Scattered Message Operations is not supported.



Workaround: None identified.

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

HSW130. N/A. Erratum has been removed

HSW131. Spurious Corrected Errors May Be Reported

Problem: Due this erratum, spurious corrected errors may be logged in the IA32_MC0_STATUS register with the valid field (bit 63) set, the uncorrected error field (bit 61) not set, a Model Specific Error Code (bits [31:16]) of 0x000F, 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 corrected errors that are benign. These corrected errors may be safely ignored.

Workaround: None identified.

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

HSW132. A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation

Problem: If the EPT is enabled, a MOV to CR3 or VMFUNC may be followed by an unexpected page fault or the use of an incorrect page translation.

Implication: Guest software may crash or experience unpredictable behavior as a result of this erratum.

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](#).

HSW133. Package C7 Power Consumption Has Been Observed to Be Higher Than Package C6

Problem: Package C7 power consumption may be higher than package C6 power consumption.

Implication: When this erratum occurs, power consumption will be higher than expected.

Workaround: It is possible for the BIOS to contain a workaround for this erratum which demotes package C7 to package C6 when power consumption in package C7 is likely to be higher than in package C6.

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

HSW134. An Intel® Hyper-Threading Technology-enabled Processor May Exhibit Unpredictable Behavior During Power or Thermal Management Operations

Problem: When both logical processors in a core are idled due to power or thermal management operations, such as thermal events or C-state entry, under certain circumstances, instruction fetches initiated before entering the idle state may not complete correctly, resulting in unpredictable system behavior.

Implication: Due to this erratum, the processor may exhibit unpredictable system behavior. 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](#).

HSW135. Certain Perfmon Events May Be Counted Incorrectly When The Processor is Not in C0 State

Problem: Due to this erratum, the PerfMon events listed below may be counted when the logical processor is not in C0 State.

IDQ.EMPTY (Event 79H, Umask 02H)

IDQ_UOPS_NOT_DELIVERED.CORE (Event 9CH, Umask 01H)

RESOURCE_STALLS.ANY (Event A2H, Umask 01H)

CYCLE_ACTIVITY.CYCLES_LDM_PENDING (Event A3H, Umask 02H, Cmask 02H)

CYCLE_ACTIVITY.CYCLES_NO_EXECUTE (Event A3H, Umask 04H, Cmask 04H)

CYCLE_ACTIVITY.STALLS_LDM_PENDING (Event A3H, Umask 06H, Cmask 06H)

Implication: The count will be higher than expected.

Workaround: None identified.

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

HSW136. Software Using Intel® TSX May Result in Unpredictable System Behavior

Problem: Under a complex set of internal timing conditions and system events, software using the Intel® TSX instructions may result in unpredictable system behavior.

Implication: This erratum may result in unpredictable system behavior.

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

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

HSW137. A Transient High Temperature Event May Cause Persistent Frequency Restrictions

Problem: If the VR reports a high temperature condition, the processor will limit the ratio on all domains (Core/Graphics/Ring) to their respective maximum non-turbo ratios. When the thermal condition is no longer present, it is expected that the processor should release this constraint and allow the domains operate in their turbo region. Due to this erratum, if the thermal event ends when the processor is in Package C6 or deeper, the constraint will not be removed.

Implication: The processor will not operate at the highest available frequencies and will have a negative impact on performance. This constraint on the ratios are cleared upon a warm or cold reset.

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](#).

HSW138. Running All Cores May Incorrectly Limit the Processor Frequency

Problem: When all IA cores in the processor are running but not executing Intel® AVX instructions, and the processor is not constrained by PL1/PL2 power limits or thermal limits, the cores should be able to operate at the 4C turbo frequency. Due to this erratum, the processor may limit core frequency under these conditions as much as several bins below the 4C turbo frequency.

Implication: When this erratum occurs, the processor will not meet specified performance levels.

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](#).



HSW139. Concurrent Core and Graphics Operation at Turbo Ratios May Lead to System Hang

Problem: Workloads that attempt concurrent operation of cores and graphics in their respective turbo ranges, under certain conditions, may result in a system hang.

Implication: Concurrent core and graphics operation may hang the system.

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](#).

HSW140. N/A. Erratum has been removed

HSW141. Performance Monitor Instructions Retired Event May Not Count Consistently

Problem: The Performance Monitor Instructions Retired event (Event C0H; Umask 00H) and the instruction retired fixed counter IA32_FIXED_CTR0 MSR (309H) are used to count the number of instructions retired. Due to this erratum, certain internal conditions may cause the counter(s) to increment when no instruction has retired or to intermittently not increment when instructions have retired.

Implication: A performance counter counting instructions retired may over count or under count. The count may not be consistent between multiple executions of the same code.

Workaround: None identified.

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

HSW142. Interactions Between Multiple Unaligned Memory Accesses And Locked Instructions May Lead to a Machine Check

Problem: Under a complex set of conditions, interactions between multiple locked operations sharing certain low order address bits and data accesses that span a 4 KByte boundary may result in a processor internal timeout machine check (IA32_MCI_STATUS.MCACOD = 0x0400).

Implication: Due to this erratum, the processor may signal a machine check exception. Intel has not observed this erratum with any commercially available system.

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](#).


HSW143. Fixed-Function Performance Counter May Over Count Instructions Retired by 32 When Intel® Hyper-Threading Technology is Enabled

Problem: If, while Intel® Hyper-Threading Technology is enabled, the IA32_FIXED_CTR0 MSR (309H) is enabled by setting bits 0 and/or 1 in the IA32_PERF_FIXED_CTR_CTRL MSR (38DH) before setting bit 32 in the IA32_PERF_GLOBAL_CTRL MSR (38FH) then IA32_FIXED_CTR0 may over count by up to 32.

Implication: When this erratum occurs, the fixed-function performance counter IA32_FIXED_CTR0 may over count by up to 32.

Workaround: The following sequence avoids this erratum (steps 1 and 2 are needed if the counter was previously enabled):

1. Clear bit 32 in the IA32_PERF_GLOBAL_CTRL MSR (38FH) and clear bits 1 and 0 in the IA32_PERF_FIXED_CTR_CTRL MSR (38DH).
2. Zero the IA32_FIXED_CTR0 MSR.
3. Set bit 32 in the IA32_PERF_GLOBAL_CTRL MSR.
4. Set bits 0 and/or 1 in the IA32_PERF_FIXED_CTR_CTRL MSR as desired.



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

HSW144. Performance Monitor UOPS_EXECUTED Event May Be Inaccurate When Using Intel® Hyper-Threading Technology

Problem: The performance monitor event UOPS_EXECUTED (Event B1H, Umask 01H) counts the number of UOPs executed each cycle. However, due to this erratum, when using Intel® HT Technology, the UOPs may not be assigned to the correct logical processor.

Implication: The total number of UOPs executed by a core will be counted correctly, but the division of UOPs between its logical processors may be incorrect.

Workaround: None identified.

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

HSW145. Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a Problem

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](#).

HSW146. 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](#).

HSW147. System May Hang or Video May Be Distorted After Graphics RC6 Exit

Problem: In a specific scenario, when the processor graphics exits RC6 and a processor core exits C6 at the same time, the system may become unresponsive or the video may become distorted.

Implication: The system may hang or video may be distorted.

Workaround: None identified.

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


HSW148. Certain eDP* Displays May Not Function as Expected

Problem: When the processor attempts to receive data on the Embedded DisplayPort* (eDP*) AUX bus, the impedance seen by the display's AUX bus drivers will be significantly below the VESA* eDP* specification's requirement for the Vaux(Receiver [Rx]) (eDP* Auxiliary Channel) input impedance.

Implication: Certain eDP* displays may not operate as expected.

Workaround: None identified.

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



HSW149. Some OFFCORE_RESPONSE Performance Monitoring Events May Undercount

Problem: The performance monitoring events OFFCORE_RESPONSE (Events B7H and BBH) should count uncore responses matching the request-response configuration specified in MSR_OFFCORE_RSPs (1A6H and 1A7H, respectively) for core-originated requests. However, due to this erratum, COREWB (bit 3), PF_L3_DATA_RD (bit 7), PF_L3_RFO (bit 8), PR_L3_CODE_RD (bit 9), SPLIT_LOCK_UC_LOCK (bit 10), and STREAMING_STORES (bit 11) request types may undercount.

Implication: These performance monitoring events may not produce reliable results for the listed request types.

Workaround: None identified.

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

HSW150. Erratum has been Removed.

HSW151. Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations

Problem: If the VM exit occurs from a guest with primary processor-based VM-execution control "activate secondary controls" set to 0 and the secondary processor-based VM-execution control "enable VPID" set to 1, then after a later VM entry with VPID fully enabled ("activate secondary controls" and "enable VPID" set to 1), the processor may use stale linear address translations.

Implication: The processor may incorrectly translate linear addresses. Intel has not observed this erratum with any commercially available software.

Workaround: Software should not enter a guest with "enable VPID" set to 1 when "activate secondary controls" is set to 0.

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

HSW152. An Interrupt Return (IRET) Instruction That Results in a Task Switch Does Not Serialize The Processor

Problem: An IRET instruction that results in a task switch by returning from a nested task does not serialize the processor (contrary to the *Software Developer's Manual Vol. 3* section titled "Serializing Instructions").

Implication: Software which depends on the serialization property of IRET during task switching may not behave as expected. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified. Software can execute an MFENCE instruction immediately prior to the IRET instruction if serialization is needed.

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

HSW153. Attempting to Disable Turbo Mode May Cause a #GP

Problem: Bit 38 of IA32_MISC_ENABLE MSR (1A0H) is Turbo Mode Disable on processors that support Intel® Dynamic Acceleration. Due to this erratum, that bit may be incorrectly treated as reserved; attempting to set Turbo Mode Disable results in a #GP even when it reads as 1.

Implication: When this erratum occurs, a WRMSR to IA32_MISC_ENABLE unexpectedly causes a #GP.

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](#).

**HSW154. PECI Frequency Limited to 1 MHz**

Problem: The 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](#).

HSW155. VGATHERQPS That Loads an Element From The APIC-Access Page May Load Other Elements From Incorrect Addresses

Problem: If the "virtualize APIC accesses" VM-execution control is 1, a 256-bit VGATHERQPS with an element that maps to the APIC-access page may use incorrect addresses to load other elements.

Implication: Loading from an incorrect address can result in unexpected behavior with respect to data, faults, or VM exits. This erratum will occur only if a guest operating system attempts to access the APIC using the VGATHERQPS instruction. 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](#).

HSW156. An APIC Timer Interrupt During Core C6 Entry May Be Lost

Problem: Due to this erratum, an APIC timer interrupt coincident with the core entering C6 state may be lost rather than held for servicing later.

Implication: A lost APIC timer interrupt may lead to missed deadlines or a system hang.

Workaround: None identified.

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

HSW157. Monitor Trap Flag (MTF) VM Exit on XBEGIN Instruction May Save State Incorrectly

Problem: Execution of an XBEGIN instruction while the "monitor trap flag" VM-execution control is 1 will be immediately followed by an MTF VM exit. If advanced debugging of RTM transactional regions has been enabled, the VM exit will erroneously save the address of the XBEGIN instruction as the instruction pointer (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.

Implication: Software 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.

Workaround: None identified.

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

**HSW158. Uncore Performance Monitoring Counters May Be Disabled or Cleared After Package C7**

Problem: Upon exiting Package C7, the following Uncore performance monitoring MSRs may be cleared to zero:

MSR_ UNC _PERF_GLOBAL_CTRL (391H)

MSR_ UNC _PERF_GLOBAL_STATUS (392H)

MSR_ UNC _PERF_FIXED_CTRL (394H)

MSR_ UNC _PERF_FIXED_CTR (395H)

Implication: Uncore performance monitoring counters may be disabled and some counter state may be cleared after Package C7.

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](#).

HSW159. PEBS Record May Be Generated After Being Disabled

Problem: A performance monitoring counter may generate a PEBS record after disabling PEBS or the performance monitoring counter by clearing the corresponding enable bit in IA32_PEBS_ENABLE MSR (3F1H) or IA32_PERF_GLOBAL_CTRL MSR (38FH).

Implication: A PEBS record generated after a VMX transition will store into memory according to the post-transition Debug Store (DS) configuration. These stores may be unexpected if PEBS is not enabled following the transition.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. A software workaround is possible through disallowing PEBS during VMX non-root operation and disabling PEBS prior to VM entry.

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

HSW160. PCIe* Ports Do Not Support Data Link Layer (DLL) Link Active Reporting

Problem: The PCIe* Base Specification requires every "Downstream Port that supports Link speeds greater than 5.0 GT/s" to support DLL Link Active Reporting. However, the PCIe* ports do not support DLL Link Active Reporting.

Implication: Due to this erratum, the PCIe* ports do not support DLL Link Active Reporting. This may be reported by a PCIe* compliance test.

Workaround: None identified.

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

HSW161. PCIe* Link Speed Negotiation May Fail After Link is Re-enabled

Problem: If a PCIe* link is established, then disabled, and the link partner's advertised speeds are changed while the link is disabled, the link may fail to correctly negotiate link speed when it is re-enabled.

Implication: Due to this erratum, the PCIe* link speed negotiation may fail after re-enabling a disabled port.

Workaround: None identified.

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

**HSW162. MOVNTDQA from WC Memory May Pass Earlier Locked Instructions**

- Problem:** An execution of (V)MOVNTDQA (streaming load instruction) that loads from Write Combining (WC) memory may appear to pass an earlier locked instruction that accesses a different cache line.
- Implication:** Software that expects a lock to fence subsequent (V)MOVNTDQA instructions may not operate properly.
- Workaround:** None identified. Software that relies on a locked instruction to fence subsequent executions of (V)MOVNTDQA should insert an MFENCE instruction between the locked instruction and subsequent (V)MOVNTDQA instruction.
- Status:** For the steppings affected, see the [Summary Tables of Changes](#).

HSW163. Data Breakpoint Coincident With a Machine Check Exception May Be Lost

- Problem:** If a data breakpoint occurs coincident with a machine check exception, then the data breakpoint may be lost.
- Implication:** Due to this erratum, a valid data breakpoint may be lost.
- Workaround:** None identified.
- Status:** For the steppings affected, see the [Summary Tables of Changes](#).

HSW164. A Corrected Internal Parity Error May Result in a System Hang

- Problem:** A corrected Internal Parity Error (IA32_MC0_STATUS.MCACOD=0005H and IA32_MC0_STATUS.MSCOD=000FH, MSR 401H bits [15:0] and bits [31:16] respectively) may cause a system hang.
- Implication:** Due to this erratum, a corrected internal parity error may cause a system hang. Reset, SMI, or INIT will end the system hang.
- 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](#).

HSW165. Data Breakpoint Coincident With a Machine Check Exception May Be Lost

- Problem:** If a data breakpoint occurs coincident with a machine check exception, then the data breakpoint may be lost.
- Implication:** Due to this erratum, a valid data breakpoint may be lost.
- Workaround:** None identified.
- Status:** For the steppings affected, see the [Summary Tables of Changes](#).

**HSW166. Internal Power State Transitions May Cause the Graphics Device to Hang**

Problem: On certain processors, when the graphics device transitions among active power states in response to dynamic power demand, the graphics device may become unresponsive.

Implication: When this erratum occurs, the graphics device may hang, resulting in a frozen or blank display. The graphics driver may be able to restart the graphics device.

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

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

HSW167. SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior

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

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

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

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

HSW168. 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](#).

HSW169. 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](#).

HSW170. 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 (listed below) may produce incorrect results when a performance counter is configured in OS-only or USR-only modes (bits 17 or 16 in IA32_PERFECTSELx 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_UOPS_RETIRED.STLB_MISS_LOADS event D0H, umask 11H
MEM_UOPS_RETIRED.STLB_MISS_STORES event D0H, umask 12H
MEM_UOPS_RETIRED.LOCK_LOADS event D0H, umask 21H
MEM_UOPS_RETIRED.SPLIT_LOADS event D0H, umask 41H
MEM_UOPS_RETIRED.SPLIT_STORES event D0H, umask 42H

MEM_LOAD_UOPS_RETIRED.L2_HIT event D1H, umask 02H
MEM_LOAD_UOPS_RETIRED.L3_HIT event D1H, umask 04H
MEM_LOAD_UOPS_RETIRED.L1_MISS event D1H, umask 08H
MEM_LOAD_UOPS_RETIRED.L2_MISS event D1H, umask 10H
MEM_LOAD_UOPS_RETIRED.L3_MISS event D1H, umask 20H
MEM_LOAD_UOPS_RETIRED.HIT_LFB event D1H, umask 40H

MEM_LOAD_L3_HIT_RETIRED.XSNP_MISS event D2H, umask 01H
MEM_LOAD_L3_HIT_RETIRED.XSNP_HIT event D2H, umask 02H
MEM_LOAD_L3_HIT_RETIRED.XSNP_HITM event D2H, umask 04H
MEM_LOAD_L3_HIT_RETIRED.XSNP_NONE event D2H, umask 08H

MEM_LOAD_UOPS_L3_MISS_RETIRED.LOCAL_DRAM event D3H, umask 01H

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](#).



HSW171. An x87 Store Instruction Which Pends #PE While EPT is Enabled May Lead to an Unexpected Machine Check and/or Incorrect x87 State Information

Problem: The execution of an x87 store instruction which causes a #PE to be pended and also causes a VM-exit due to an EPT violation or misconfiguration may lead the VMM logging a machine check exception with a cache hierarchy error (IA32_MCi_STATUS.MCACOD = 0150H and IA32_MCi_STATUS.MSCOD = 000FH). Additionally, FSW.PE and FSW.ES (bits 5 and 7 of the FPU Status Word) may be incorrectly set to 1, and the x87 Last Instruction Opcode (FOP) may be incorrect.

Implication: When this erratum occurs, the VMM may receive an expected machine check exception, and software attempting to handle the #PE may not behave as expected.

Workaround: None identified.

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

HSW172. 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® Hyper-Threading 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](#).

HSW173. Writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP May #GP When Intel® TSX is Not Supported

Problem: Due to this erratum, on processors that do not support Intel® TSX (CPUID.07H.EBX bits 4 and 11 are both zero), writes to MSR_LASTBRANCH_x_FROM_IP (MSR 680H to 68FH) and MSR_LER_FROM_LIP (MSR 1DDH) may #GP unless bits [62:61] are equal to bit [47].

Implication: The value read from MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP is unaffected by this erratum; bits [62:61] contain IN_TSX and TSX_ABORT information respectively. Software restoring these MSRs from saved values are subject to this erratum.

Workaround: Before writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP, ensure the value being written has bit [47] replicated in bits [62:61]. This is most easily accomplished by sign extending from bit [47] to bits [62:48].

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

HSW174. APIC Timer Interrupt May Not Be Generated at The Correct Time In TSC-Deadline Mode

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

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

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

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

HSW175. Precise Performance Monitoring May Generate Redundant PEBS Records

Problem: Processor Event Based Sampling 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](#).

HSW176. In eMCA2 Mode, When the Retirement Watchdog Timeout Occurs CATERR# May Be Asserted

Problem: A Retirement Watchdog Timeout (MCACOD = 0x0400) in Enhanced MCA2 (eMCA2) mode will cause the CATERR# pin to be pulsed in addition to an MSMI# pin assertion. In addition, a Machine Check Abort (#MC) will be pended in the cores along with the MSMI.

Implication: Due to this erratum, systems that expect to only see MSMI# will also see CATERR# pulse when a Retirement Watchdog Timeout occurs. The CATERR# pulse can be safely ignored.

Workaround: None identified.

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


HSW177. 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](#).

**HSW178. 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 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](#).

HSW179. 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: None identified.

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

HSW180. 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](#).

**HSW181. Performance Monitoring General Counter 2 May Have Invalid Value Written When Intel® TSX Is Enabled**

Problem: When Intel® Transactional Synchronization Extensions (Intel® TSX) is enabled, and there are aborts (HLE or RTM) overlapping with access or manipulation of the IA32_PMC2 general-purpose performance counter (Offset: C3h), it may return invalid value.

Implication: Software may read invalid value from IA32_PMC2.

Workaround: None identified.

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

HSW182. 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](#).

§



Specification Changes

The Specification Changes listed in this section apply to the following documents:

- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

HSW1. Intel® Transactional Synchronization Extensions (Intel®TSX) Instruction

- Due to Erratum HSW136, Intel® Transactional Synchronization Extensions (Intel®TSX) instructions are disabled and are only supported for software development. See your Intel representative for details.

§




Specification Clarifications

The Specification Clarifications listed in this section may apply to the following documents:

- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

There are no new Specification Changes in this Specification Update revision.

§



Documentation Changes

The Documentation Changes listed in this section apply to the following documents:

- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

All Documentation Changes will be incorporated into a future version of the appropriate Processor documentation.

Note: Documentation changes for *Intel® 64 and IA-32 Architecture Software Developer's Manual* volumes 1, 2A, 2B, 3A, and 3B will be posted in a separate document, *Intel® 64 and IA-32 Architecture Software Developer's Manual Documentation Changes*. Use the following link to access this file: <http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html>.

On-Demand Clock Modulation Feature Clarification

Software Controlled Clock Modulation section of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*, Volume 3B: System Programming Guide will be modified to differentiate On-demand clock modulation feature on different processors. The clarification will state:

For Intel® Hyper-Threading Technology enabled processors, the IA32_CLOCK_MODULATION register is duplicated for each logical processor. In order for the On-demand clock modulation feature to work properly, the feature must be enabled on all the logical processors within a physical processor. If the programmed duty cycle is not identical for all the logical processors, the processor clock will modulate to the highest duty cycle programmed for processors if the CPUID DisplayFamily_DisplayModel signatures is listed in [Table 14-2. CPUID Signatures for Legacy Processors that Resolve to Higher Performance Setting of Conflicting Duty Cycle Requests](#). For all other processors, if the programmed duty cycle is not identical for all logical processors in the same core, the processor will modulate at the lowest programmed duty cycle.

For multiple processor cores in a physical package, each core can modulate to a programmed duty cycle independently.

For the Intel P6 family processors, on-demand clock modulation was implemented through the chipset, which controlled clock modulation through the processor's STPCLK# pin.

**Table 14-2. CPUID Signatures for Legacy Processors that Resolve to Higher Performance Setting of Conflicting Duty Cycle Requests**

Display Family Display Model	Display Family Display Model	Display Family Display Model	Display Family Display Model
0F_xx	06_1C	06_1A	06_1E
06_1F	06_25	06_26	06_27
06_2C	06_2E	06_2F	06_35
06_36	N/A	N/A	N/A

HSW2. Intel® Virtualization Technology (Intel® VT) Clarification

Section 3.1 will be modified to include the following paragraph:

It is recommended to avoid device direct assignment to Virtual Machines in virtualized environments with this processor, due to the lack of Access Control Services (ACS) support in PCIe* root ports. Some Operating Systems may check for ACS support and potentially disable direct device assignment (that is, affects SR-IOV setup/configuration within the server) as well.

§