

Intel[®] Xeon[®] E3-1200 v4 Processor Family

Specification Update

July 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® 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 may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, Pentium, Celeron, Intel Core, Intel Xeon, Enhanced Intel SpeedStep Technology, Intel Software Guard Extensions (Intel SGX), Intel Processor Trace (Intel PT), Intel TSX (Intel Transactional Synchronization Extensions), Intel Memory Protection Extensions (Intel MPX), Intel Advanced Vector Extensions (Intel AVX), Intel Virtualization Technology for Directed I/O (Intel VT-d), Thermal Design Power (TDP), Intel Smart Sound Technology (Intel SST), Intel Graphics Driver and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2020, Intel Corporation. All rights reserved.




Content

Revision History	4
Preface	5
Summary Tables of Changes	7
Identification Information	13
Errata	15
Specification Changes	49
Specification Clarifications	50
Documentation Changes	51



Revision History

Date	Revision	Description
July 2020	019	Added erratum BDW129
April 2020	018	Added erratum BDW128
March 2020	017	Added erratum BDW127
February 2020	016	Added erratum BDW126 Revised erratum BDW72
August 2019	015	Added errata BDW124 and BDW125
February 2019	014	Added erratum BDW123
August 2018	013	Added erratum BDW122
July 2018	012	Added erratum BDW121
June 2018	011	Added erratum BDW120
April 2018	010	Updated erratum BDW117
February 2018	009	Added erratum BDW119
March 2017	008	Updated erratum BDW109 Added errata BDW116 - BDW118 Removed errata BDW13, BDW51
November 2016	007	Updated errata BDW85, BDW100 Added errata BDW114 - BDW115
October 2016	006	Updated errata BDW29, BDW90 Added errata BDW102 - BDW113
April 2016	005	Updated erratum BDW64 Added errata BDW96 - BDW101
January 2016	004	Updated Table 3. Product Family SKUs Added errata BDW94 - BDW95
October 2015	003	Added errata BDW89 - BDW93
August 2015	002	Added errata BDW84 - BDW88
June 2015	001	Initial Release



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
<i>Intel® Xeon® Processor E3-1200 v4 Product Family Datasheet, Volume 1 of 2</i>	332374
<i>Intel® Xeon® Processor E3-1200 v4 Product Family Datasheet, Volume 2 of 2</i>	332375

Related Documents

Document Title	Document Number / Location
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 2A: Instruction Set Reference, A-L -- Refer to section "AP-485, Intel® Processor Identification and the CPUID Instruction"</i>	http://www.intel.com/design/processor/aplnots/241618.htm
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction set reference, A-L</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction set reference, M-U</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C: Instruction set reference, V-Z</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System programming guide, part 1</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System programming guide, part 2</i> <i>Intel® 64 and IA-32 Intel Architecture Optimization Reference Manual</i>	http://www.intel.com/products/processor/manuals/index.htm
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html
<i>Advanced Configuration Power Interface (ACPI) Specifications</i>	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, etc. as described in the processor identification information table. Read all notes associated with each S-Spec number.

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 5)

Number	Steppings	Status	Errata
	G-0		
BDW1	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
BDW2	X	No Fix	EFLAGS Discrepancy on Page Faults and on Extended Page Table (EPT)-Induced Virtual Machine (VM) Exits after a Translation Change
BDW3	X	No Fix	MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a Data Translation Look Aside Buffer (DTLB) ErrorLast Exception Record (LER) Model-Specific Registers (MSRs) May Be Unreliable
BDW4	X	No Fix	Last Exception Record (LER) Model-Specific Registers (MSRs) May Be Unreliable
BDW5	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang



Errata (Sheet 2 of 5)

Number	Steppings	Status	Errata
	G-0		
BDW6	X	No Fix	An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang
BDW7	X	No Fix	General Protection Exception (#GP) on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
BDW8	X	No Fix	FREEZE_WHILE_SMM Does Not Prevent Event From Pending Precise Event Based Sampling (PEBS) During SMM
BDW9	X	No Fix	Advanced Programmable Interrupt Controller (APIC) Error "Received Illegal Vector" May be Lost
BDW10	X	No Fix	Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations
BDW11	X	No Fix	Performance Monitor Precise Instruction Retired Event May Present Wrong Indications
BDW12	X	No Fix	CR0.CD Is Ignored in Virtual Machine Extensions (VMX) Operation
BDW13	X	No Fix	Erratum has been Removed
BDW14	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a Device-Not-Available (#NM) Exception
BDW15	X	No Fix	Processor May Fail to Acknowledge a Transaction Layer Packet (TLP) Request
BDW16	X	No Fix	Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered
BDW17	X	No Fix	PCIe* Root-port Initiated Compliance State Transmitter Equalization Settings May be Incorrect
BDW18	X	No Fix	PCIe* Controller May Incorrectly Log Errors on Transition to RxL0s
BDW19	X	No Fix	Unused PCIe* Lanes May Report Correctable Errors
BDW20	X	No Fix	PCIe* Root Port May Not Initiate Link Speed Change
BDW21	X	No Fix	Pending x87 FPU Exceptions (#MF) May be Signaled Earlier Than Expected
BDW22	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
BDW23	X	No Fix	VEX.L is Not Ignored with VCVT*2SI Instructions
BDW24	X	No Fix	PCIe* Atomic Transactions From Two or More PCIe* Controllers May Cause Starvation
BDW25	X	No Fix	The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When The Uncorrected (UC) Bit is Set
BDW26	X	No Fix	PCIe* Controller May Initiate Speed Change While in DL_Init State Causing Certain PCIe* Devices to Fail to Train
BDW27	X	No Fix	Spurious Intel® Virtualization Technology (Intel® VT) Interrupts May Occur When the PFO Bit is Set
BDW28	X	No Fix	Processor May Livelock During On Demand Clock Modulation
BDW29	X	No Fix	Internal Parity Errors May Incorrectly Report Overflow in The IA32_MC2_STATUS MSR
BDW30	X	No Fix	Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX May Over Count
BDW31	X	No Fix	Performance Monitor Event DSB2MITE_SWITCHES.COUNT May Over Count
BDW32	X	No Fix	Timed MWAIT May Use Deadline of a Previous Execution
BDW33	X	No Fix	IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For Virtual-Machine Control Structure (VMCS) Encoding
BDW34	X	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
BDW35	X	No Fix	Locked Load Performance Monitoring Events May Under Count
BDW36	X	No Fix	Transactional Abort May Produce an Incorrect Branch Record
BDW37	X	No Fix	SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior



Errata (Sheet 3 of 5)

Number	Steppings	Status	Errata
	G-0		
BDW38	X	No Fix	Performance Monitoring Interrupt (PMI) May be Signaled More Than Once For Performance Monitor Counter Overflow
BDW39	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
BDW40	X	No Fix	Intel® Turbo Boost Technology May be Incorrectly Reported as Supported on 5th Generation Intel® Core™ i3 U-series, and select Mobile Intel Pentium® processors and Mobile Intel Celeron® processors
BDW41	X	No Fix	The SAMPLE/PRELOAD JTAG Command Does Not Sample The Display Transmit Signals
BDW42	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
BDW43	X	No Fix	Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID
BDW44	X	No Fix	Back to Back Updates of the Intel® Virtualization Technology for Directed I/O (Intel® VT-d) Root Table Pointer May Lead to an Unexpected DMA Remapping Fault
BDW45	X	No Fix	A MOV to CR3 When the EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
BDW46	X	No Fix	Peer I/O Device Writes to The Graphics Memory Address range (GMADR) May Lead to a System Hang
BDW47	X	No Fix	Spurious Corrected Errors May be Reported
BDW48	X	No Fix	Intel® Processor Trace (Intel® PT) Packet Generation May Stop Sooner Than Expected
BDW49	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
BDW50	X	No Fix	Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value
BDW51	X	Fixed	Erratum has been Removed
BDW52	X	No Fix	Spurious Corrected Errors May be Reported
BDW53	X	No Fix	Performance Monitoring Event INSTR_RETIRED.ALL May Generate Redundant PEBS Records For an Overflow
BDW54	X	No Fix	Concurrent Core And Graphics Operation at Turbo Ratios May Lead to System Hang
BDW55	X	No Fix	The System May Hang on First Package C6 or deeper C-State
BDW56	X	No Fix	Using The FIVR Spread Spectrum Control Mailbox May Not Produce The Requested Range
BDW57	X	No Fix	Intel® PT MODE.Exec, PIP, and CBR Packets Are Not Generated as Expected
BDW58	X	No Fix	Performance Monitor Instructions Retired Event May Not Count Consistently
BDW59	X	No Fix	General-Purpose Performance Counters May be Inaccurate with Any Thread
BDW60		No Fix	An Incorrect LBR or Intel® PT Packet May Be Recorded Following a Transactional Abort
BDW61	X	No Fix	Executing an RSM Instruction With Intel® PT Enabled Will Signal a #GP
BDW62	X	Fixed	Intel® Processor Trace PIP May be Unexpectedly Generated
BDW63	X	Fixed	A Virtualization Exception (#VE) May Not Invalidate Cached Translation Information
BDW64	X	Fixed	Some Performance Monitor Events May Overcount During TLB Misses
BDW65	X	No Fix	Intel® PT PSB+ Packets May Contain Unexpected Packets
BDW66	X	No Fix	Writing Non-Zero Value to IA32_RTIT_CR3_MATCH [63:48] Will Cause #GP
BDW67	X	Fixed	Core C6 May Cause Interrupts to be Serviced Out of Order
BDW68	X	Fixed	LPDDR3 Memory Training May Cause Platform Boot Failure
BDW69	X	No Fix	Aggressive Ramp Down of Voltage May Result in Unpredictable Behavior
BDW70	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests And Snoop Requests May be Incorrect



Errata (Sheet 4 of 5)

Number	Steppings	Status	Errata
	G-0		
BDW71	X	No Fix	DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction
BDW72	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation
BDW73	X	No Fix	The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated After a UC Error is Logged
BDW74	X	No Fix	Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a #UD
BDW75	X	No Fix	Processor Operation at Turbo Frequencies Above 3.2 GHz May Cause The Processor to Hang
BDW76	X	No Fix	DDR-1600 With a Reference Clock of 100Mhz May Cause S3 Entry Failure
BDW77	X	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
BDW78	X	No Fix	System May Hang or Video May be Distorted After Graphics RC6 Exit
BDW79	X	No Fix	Certain Embedded DisplayPort* (eDP*) Displays May Not Function as Expected
BDW80	X	No Fix	Instruction Fetch Power Saving Feature May Cause Unexpected Instruction Execution
BDW81	X	No Fix	IA Core Ratio Change Coincident With Outstanding Read to the Display Engine (DE) May Cause a System Hang
BDW82	X	No Fix	PL3 Power Limit Control Mechanism May Not Release Frequency Restrictions
BDW83	X	No Fix	I/O Subsystem Clock Gating May Cause a System Hang
BDW84	X	No Fix	PAGE_WALKER_LOADS Performance Monitoring Event May Count Incorrectly
BDW85	X	No Fix	Some DRAM And L3 Cache Performance Monitoring Events May Undercount
BDW86	X	No Fix	The System May Hang When Executing a Complex Sequence of Locked Instructions
BDW87	X	No Fix	Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations
BDW88	X	No Fix	An IRET Instruction That Results in a Task Switch Does Not Serialize The Processor
BDW89	X	No Fix	Attempts to Retrain a PCIe* Link May be Ignored
BDW90	X	No Fix	Attempting Concurrent Enabling of Intel® PT With LBR, BTS, or BTM Results in a #GP
BDW91	X	No Fix	Setting TraceEn While Clearing BranchEn in IA32_RTIT_CTL Causes a #GP
BDW92	X	No Fix	Processor Graphics I/O Memory Management Unit (IOMMU) Unit May Not Mask DMA Remapping Faults
BDW93	X	No Fix	Processor Graphics IOMMU Unit May Report Spurious Faults
BDW94	X	No Fix	PECI Frequency Limited to 1 MHz
BDW95	X	No Fix	Reads or Writes to LBRs with Intel® PT Enabled Will Result in a #GP
BDW96	X	No Fix	Frequency Difference Between Intel Architecture (IA) Core(s) and Ring Domains May Cause Unpredictable System Behavior
BDW97	X	No Fix	Back-to-Back Page Walks Due to Instruction Fetches May Cause a System Hang
BDW98	X	No Fix	Performance Monitoring Counters May Produce Incorrect Results for BR_INST_RETIRED Event on Logical Processor
BDW99	X	No Fix	PEBS Record May Be Generated After Being Disabled
BDW100	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events Related to RFO Request Types May Count Incorrectly
BDW101	X	No Fix	MOVNTDQA From Write Combining (WC) Memory May Pass Earlier Locked Instructions
BDW102	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May be Lost
BDW103	X	No Fix	Internal Parity Errors May Incorrectly Report Overflow in the IA32_MC0_STATUS MSR
BDW104	X	No Fix	Software Using Intel® TSX May Behave Unpredictably



Errata (Sheet 5 of 5)

Number	Steppings	Status	Errata
	G-0		
BDW105	X	No Fix	Incorrect VMCS Used for Page Modification Log (PML)-Index field on VMX Transitions Into and Out of SMM
BDW106	X	No Fix	An APIC Timer Interrupt During Core C6 Entry May be Lost
BDW107	X	No Fix	An Intel® Hyper-Threading Technology (Intel® HT Technology) Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior
BDW108	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May be Lost
BDW109	X	No Fix	Performance Monitoring Counters May Undercount When Using a CPL Filtering
BDW110	X	No Fix	An Intel® HT Technology Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior
BDW111	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions
BDW112	X	No Fix	RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
BDW113	X	No Fix	Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes
BDW114	X	No Fix	An x87 Store Instruction Which Pends Precision Exception (#PE) While EPT is Enabled May Lead to an Unexpected Machine Check and/or Incorrect x87 State Information
BDW115	X	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
BDW116	X	No Fix	General-Purpose Performance Monitoring Counters 4-7 Do Not Count With USR Mode Only Filtering
BDW117	X	No Fix	Writing MSR_LASTBRANCH_x_FROM_IP and MSR_LER_FROM_LIP May #GP When Intel® TSX is Not Supported
BDW118	X	No Fix	APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode
BDW119	X	No Fix	Precise Performance Monitoring May Generate Redundant PEBS Records
BDW120	X	No Fix	Reads From MSR_LER_TO_LIP May Not Return a Canonical Address
BDW121	X	No Fix	In eMCA2 Mode, When The Retirement Watchdog Timeout Occurs CATERR# May be Asserted
BDW122	X	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store
BDW123	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
BDW124	X	No Fix	Graphics Intel® VT-d Hardware May Cache Invalid Entries
BDW125	X	No Fix	Gen9 Graphics Intel® VT-d Hardware May Cache Invalid Entries
BDW126	X	No Fix	System May Hang Under Complex Conditions
BDW127	X	Fixed	Incorrect CPUID Reporting For Intel® Turbo Boost Max Technology 3.0 Capability
BDW128	X	No Fix	PMU MSR_UNC_PERF_FIXED_CTR is Cleared After Pkg C7 or Deeper
BDW129	X	No Fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes



Specification Changes


Number	Specification Changes
N/A	None for this revision of this specification update.

Specification Clarifications

Number	Specification Clarifications
N/A	None for this revision of this specification update.

Documentation Changes

Number	Documentation Changes
N/A	None for this revision of this specification update.



Identification Information

Component Identification using Programming Interface

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

Table 1. 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	0100b		00b	0110b	0111b	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) register after RESET, Bits [11:8] of the Extended Accumulator Register (EAX) register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
4. The Model Number corresponds to Bits [7:4] of the EDX register after RESET, Bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
5. The Stepping ID in Bits [3:0] indicates the revision number of that model. 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 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.

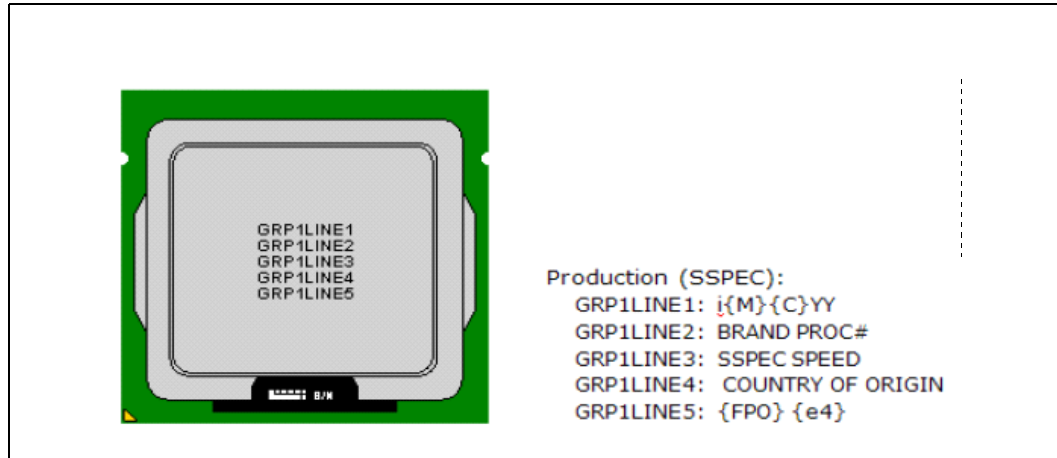
Table 2. Processor Identification by Register Contents

Processor Line	Stepping	Vendor ID	Host Device ID	Processor Graphics Device ID	Revision ID	Compatibility Revision ID
H-Processor Line	G-0	8086h	1618h	GT3 = 162Ah	0Ah	0Ah

Component Marking Information

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

Figure 1. Intel® Xeon® Processor E3-1200 v4 Product Family LGA Top-Side Markings



Pin Count: 1150

Package Size: 37.5 mm x 37.5 mm

For Intel® Xeon® Processor E3-1200 v4 Product Family SKUs, see: <https://ark.intel.com/content/www/us/en/ark/products/series/87722/intel-xeon-processor-e3-v4-family.html>



Errata

BDW1. 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/interrupt event should be transparent to the LBR, BTS and BTM mechanisms. However, during a specific boundary condition where the exception/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 BTS and BTM operations which report the LBR will also be incorrect.

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

Workaround: None identified.

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

BDW2. EFLAGS Discrepancy on Page Faults and on Extended Page Table (EPT)-Induced Virtual Machine (VM) 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 TLB invalidation.

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

BDW3. MCI_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a Data Translation Look Aside Buffer (DTLB) Error

Problem: A single 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*.

**BDW4. Last Exception Record (LER) Model-Specific Registers (MSRs) 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](#).

BDW5. 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](#).

BDW6. 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](#).


BDW7. 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](#).

**BDW8. 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 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](#).

BDW9. 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](#).


BDW10. 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](#).

**BDW11. 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/PMI interrupts and/or incorrect counter values if the counter is reset with a SAV below 100 (Sample-After-Value is the counter reset value software programs in 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](#).

BDW12. CR0.CD Is Ignored in Virtual Machine Extensions (VMX) Operation

Problem: If CR0.CD=1, the Memory Type Range Registers (MTRRs) and PAT should be ignored and the UC memory type should be used for all memory accesses. Due to this erratum, a logical processor in 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 VMX operation.

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

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

BDW13. Erratum has been Removed**BDW14. 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](#).


BDW15. Processor May Fail to Acknowledge a Transaction Layer Packet (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](#).

**BDW16. 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](#).

BDW17. 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](#).

BDW18. 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](#).

**BDW19. 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](#).

BDW20. 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](#).

BDW21. Pending x87 FPU Exceptions (#MF) May be Signaled Earlier Than Expected

Problem: x87 instructions that trigger #MF normally service interrupts before the #MF. Due to this erratum, if an instruction that triggers #MF is 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 #MF being-signalized before pending interrupts are serviced.

Workaround: None identified.

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



BDW22. 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](#).

BDW23. 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](#).

BDW24. 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](#).

BDW25. The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated When The Uncorrected (UC) Bit is Set

Problem: After a 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](#).

**BDW26. 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](#).

BDW27. Spurious Intel® Virtualization Technology (Intel® VT) Interrupts May Occur When the PFO Bit is Set

Problem: When the Primary Fault Overflow (PFO) field (bit [0] in the Intel® VT 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](#).


BDW28. Processor May Livelock During On Demand Clock Modulation

Problem: The processor may live-lock 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), and (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 cache line or access the 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](#).

**BDW29. 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](#).

BDW30. 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 an 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](#).

BDW31. 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 Decode Stream Buffer (DSB) to 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](#).


BDW32. Timed MWAIT May Use Deadline of a Previous Execution

Problem: A timed MWAIT instruction specifies a Time Stamp Counter (TSC) deadline for execution resumption. If a wake event causes execution to resume before the deadline is reached, a subsequent timed MWAIT instruction may incorrectly use the deadline of the previous timed MWAIT when that previous deadline is earlier than the new one.

Implication: A timed MWAIT may end earlier 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](#).

**BDW33. IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For Virtual-Machine Control Structure (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](#).

BDW34. Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

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

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

Workaround: None identified.

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

BDW35. 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](#).


BDW36. 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](#).

**BDW37. SMRAM State-Save Area Above the 4GB 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 4GB address boundary.

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

BDW38. Performance Monitoring Interrupt (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 an 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*.

BDW39. 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*.

BDW40. Intel® Turbo Boost Technology May be Incorrectly Reported as Supported on 5th Generation Intel® Core™ i3 U-series, and select Mobile Intel Pentium® processors and Mobile Intel Celeron® processors

Problem: The 5th Generation Intel Core™ i3 U-series, and select Mobile Pentium® and Celeron® processors may incorrectly report support for Intel Turbo Boost Technology via CPUID.06H.EAX bit 1.

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

Workaround: None identified.

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

**BDW41. 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](#).

BDW42. 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](#).


BDW43. Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID

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

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

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

**BDW44. Back to Back Updates of the Intel® Virtualization Technology for Directed I/O (Intel® VT-d) Root Table Pointer May Lead to an Unexpected DMA Remapping Fault**

Problem: An Intel® VT-d Root Table Pointer update that completes followed by a second Root Table Pointer update that also completes, without performing a global invalidation of either the context-cache or the IOTLB between the two updates, may lead to an unexpected DMA remapping fault.

Implication: Back to back Root Table Pointer updates may cause an unexpected DMA remapping fault. Intel has not observed this erratum with any commercially available software.

Workaround: Software must not perform a second Root Table Pointer update before doing a global invalidation of either the context-cache or the IOTLB.

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

BDW45. A MOV to CR3 When the 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 Table of Changes*.

BDW46. Peer I/O Device Writes to The Graphics Memory Address range (GMADR) May Lead to a System Hang

Problem: The system may hang when a peer I/O device uses the peer aperture to directly write into the GMADR.

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

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

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

BDW47. 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 Machine Check Architecture (MCA) Error Code (bits [15:0]) of 0x0005. If the 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*.

**BDW48. Intel® Processor Trace (Intel® PT) Packet Generation May Stop Sooner Than Expected**

Problem: Setting the STOP bit (bit 4) in a Table of Physical Addresses entry directs the processor to stop Intel® PT packet generation when the associated output region is filled. The processor indicates this has occurred by setting the Stopped bit (bit 5) of IA32_RTIT_STATUS MSR (571H). Due to this erratum, packet generation may stop earlier than expected.

Implication: When this erratum occurs, the OutputOffset field (bits [62:32]) of the IA32_RTIT_OUTPUT_MASK_PTRS MSR (561H) holds a value that is less than the size of the output region which triggered the STOP condition; Intel® PT analysis software should not attempt to decode packet data bytes beyond the OutputOffset.

Workaround: None identified.

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

BDW49. PEBS Eventing IP Field May be Incorrect After Not-Taken Branch

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

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

Workaround: None identified.

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

BDW50. Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value

Problem: An HLE transactional region begins with an instruction with the XACQUIRE prefix. Due to this erratum, reads from within the transactional region of the memory destination of that instruction may return the value that was in memory before the transactional region began.

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

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

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

BDW51. Erratum has been Removed**BDW52. Spurious Corrected Errors May be Reported**

Problem: Due this erratum, spurious corrected errors may be logged in the MC0_STATUS register with the valid (bit 63) set, the uncorrected error (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 the 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: It is possible for the BIOS to contain a workaround for this erratum.

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

**BDW53. Performance Monitoring Event INSTR_RETIRE.ALL May Generate Redundant PEBS Records For an Overflow**

Problem: Due to this erratum, the performance monitoring feature Precise Distribution of Instructions Retired (PDIR) for INSTR_RETIRE.ALL (Event C0H; Umask 01H) will generate redundant PEBS records for a counter overflow. This can occur if the lower 6 bits of the performance monitoring counter are not initialized or reset to 0, in the PEBS counter reset field of the DS Buffer Management Area.

Implication: The performance monitor feature PDIR, may generate redundant PEBS records for an overflow.

Workaround: Initialize or reset the counters such that lower 6 bits are 0.

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

BDW54. 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](#).

BDW55. The System May Hang on First Package C6 or deeper C-State

Problem: Under certain conditions following a cold boot, exiting the first package C6 or deeper C-State may hang the system.

Implication: Due to this erratum, the system may hang exiting a package C6 or deeper C-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](#).

BDW56. Using The FIVR Spread Spectrum Control Mailbox May Not Produce The Requested Range

Problem: Values programmed into the Fully Integrated Voltage Regulator Spread Spectrum Control (FIVR SSC) Mailbox may not result in the expected spread spectrum range.

Implication: The actual FIVR spread spectrum range may not be the same as the programmed values affecting the usefulness of FIVR SSC Mailbox as a means to reduce Electromagnetic Interference (EMI).

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](#).

**BDW57. Intel® PT MODE.Exec, PIP, and CBR Packets Are Not Generated as Expected**

Problem: The Intel® PT MODE.Exec (MODE packet – Execution mode leaf), Paging Information Packet (PIP), and Core:Bus Ratio (CBR) packets are generated at the following Packet Stream Boundary (PSB+) event rather than at the time of the originating event as expected.

Implication: The decoder may not be able to properly disassemble portions of the binary or interpret portions of the trace because many packets may be generated between the MODE.Exec, PIP, and CBR events and the following PSB+ event.

Workaround: The processor inserts these packets as status packets in the PSB+ block. The decoder may have to skip forward to the next PSB+ block in the trace to obtain the proper updated information to continue decoding.

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

BDW58. 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](#).

BDW59. General-Purpose Performance Counters May be Inaccurate with Any Thread

Problem: The IA32_PMCx MSR (C1H - C8H) general-purpose performance counters may report inaccurate counts when the associated event selection IA32_PERFEVTSELx MSR's (186H - 18DH) AnyThread field (bit 21) is set and either the OS field (bit 17) or USR field (bit 16) is set (but not both set).

Implication: Due to this erratum, IA32_PMCx counters may be inaccurate.

Workaround: None identified

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

BDW60. An Incorrect LBR or Intel® PT Packet May Be Recorded Following a Transactional Abort

Problem: Use of Intel® TSX may result in a transactional abort. If an abort occurs immediately following a branch instruction, an incorrect branch target may be logged in an LBR or in an Intel® PT packet before the LBR or the Intel® PT packet produced by the abort.

Implication: The LBR or the Intel® PT packet immediately preceding a transactional abort may indicate an unexpected branch target.

Workaround: None identified.

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



BDW61. Executing an RSM Instruction With Intel® PT Enabled Will Signal a #GP

Problem: Upon delivery of a SMI, the processor saves and then clears TraceEn in the IA32_RTIT_CTL MSR (570H), thus disabling Intel® PT. If the SMI handler enables Intel® PT and it remains enabled when an RSM instruction is executed, a shutdown event should occur. Due to this erratum, the processor does not shutdown but instead generates a #GP.

Implication: When this erratum occurs, a #GP will be signaled.

Workaround: If software enables the Intel® PT in system-management mode, it should disable Intel® PT before executing RSM.

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

BDW62. Intel® Processor Trace PIP May be Unexpectedly Generated

Problem: When Intel® PT is enabled, Packet Stream Boundary packets may include a PIP even though the OS field (bit 2) of IA32_RTIT_CTL MSR (570H) is 0.

Implication: When this erratum occurs, user-mode tracing (indicated by IA32_RTIT_CTL.OS = 0) may include CR3 address information. This may be an undesirable leakage of kernel information.

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](#).

BDW63. A Virtualization Exception (#VE) May Not Invalidate Cached Translation Information

Problem: An EPT violation that causes a #VE may not invalidate the guest-physical mappings that were used to translate the guest-physical address that caused the EPT violation.

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

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

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

BDW64. Some Performance Monitor Events May Overcount During TLB Misses


Problem: The following Performance Monitor Events may significantly overcount when multiple TLB misses happen nearly concurrently:

1. ITLB_MISSES (Event 85H, Umask 01H, 02H, 04H, 08H, 10H)
2. DTLB_LOAD_MISSES (Event 08H, Umask 01H, 02H, 04H, 08H, 10H)
3. DTLB_STORE_MISSES (Event 49H, Umask 01H, 02H, 04H, 08H, 10H)
4. PAGE_WALKER_LOADS (Event BCH, all Umasks)

Implication: When this erratum occurs, counts accumulated for the listed events may significantly exceed the correct counts.

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](#).

**BDW65. Intel® PT PSB+ Packets May Contain Unexpected Packets**

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

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

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

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

BDW66. Writing Non-Zero Value to IA32_RTIT_CR3_MATCH [63:48] Will Cause #GP

Problem: Bits [63:48] of the IA32_RTIT_CR3_MATCH MSR (0572H) are incorrectly treated as reserved and therefore writing non-zero values to them will cause a #GP

Implication: Due to this erratum, a #GP fault will occur if a non-zero value is written to IA32_RTIT_CR3_MATCH[63:48].

Workaround: Software should avoid writing non-zero values to bits [63:48] of the IA32_RTIT_CR3_MATCH MSR.

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

BDW67. Core C6 May Cause Interrupts to be Serviced Out of Order

Problem: If the APIC In-Service Register (ISR) indicates in-progress interrupt(s) at Core C6 entry, a lower priority interrupt pending in the Interrupt Request Register (IRR) may be executed after Core C6 exit, delaying completion of the higher priority interrupt's service routine.

Implication: An interrupt may be processed out of its intended priority order immediately after Core C6 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](#)

BDW68. LPDDR3 Memory Training May Cause Platform Boot Failure

Problem: Due to this erratum, LPDDR3 memory sub-systems may not successfully complete training.

Implication: When this erratum occurs, the platform may fail to boot successfully

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

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


BDW69. Aggressive Ramp Down of Voltage May Result in Unpredictable Behavior

Problem: Aggressive ramp down of Vcc voltage may result in insufficient voltage to meet power demand.

Implication: Due to this erratum, unpredictable system behavior or hangs 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](#).



BDW70. 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*.

BDW71. DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction

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

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

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

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

BDW72. 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*.

**BDW73. The Corrected Error Count Overflow Bit in IA32_MC0_STATUS is Not Updated After a UC Error is Logged**

Problem: When a UC error is logged in the IA32_MC0_STATUS MSR (401H), corrected errors will continue to update 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 after a UC error is logged.

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*.

BDW74. Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a #UD

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

Implication: A MOVBE instruction with both REX.W=1 and a 66H prefix will unexpectedly cause a #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*.

BDW75. Processor Operation at Turbo Frequencies Above 3.2 GHz May Cause The Processor to Hang

Problem: The processor may not run reliably when operating at turbo frequencies above 3.2 GHz.

Implication: Due to this erratum, 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*.

BDW76. DDR-1600 With a Reference Clock of 100Mhz May Cause S3 Entry Failure

Problem: Due to this erratum, Platform State S3 entry with a DDR-1600 memory subsystem may cause the DDR reference clock, when configured at 100 MHz, to briefly switch to 133 MHz resulting in unpredictable system behavior.

Implication: When this erratum occurs, the system may experience unpredictable system behavior.

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

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


BDW77. 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*.

**BDW78. 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: It is possible for the BIOS to contain a workaround for this erratum.

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

BDW79. Certain Embedded DisplayPort* (eDP*) Displays May Not Function as Expected

Problem: When the processor attempts to receive data on the 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(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*.

BDW80. Instruction Fetch Power Saving Feature May Cause Unexpected Instruction Execution

Problem: Under a complex set of micro-architectural conditions, an instruction fetch dynamic power savings feature may cause the processor to execute unexpected instructions.

Implication: When this erratum occurs, instances of unexpected #GP or Page fault (#PF) have been observed. Unexpected faults may lead to an application or operating system crash.

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*.

BDW81. IA Core Ratio Change Coincident With Outstanding Read to the Display Engine (DE) May Cause a System Hang

Problem: An outstanding read from an IA core to the DE that is coincident with an IA core ratio change may result in a system hang.

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

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

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

BDW82. PL3 Power Limit Control Mechanism May Not Release Frequency Restrictions

Problem: The PL3 mechanism imposes peak frequency constraints on all domains (Core, Graphics, and Ring) when a current spike that might cause accelerated battery aging is detected. Due to this erratum, these constraints may not be released when the current spike has ended.

Implication: The processor clock frequencies may be unnecessarily limited.

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*.

**BDW83. I/O Subsystem Clock Gating May Cause a System Hang**

Problem: Certain complex internal conditions and timing relationships during clock gating of the I/O subsystem may cause a system hang and may lead to a timeout machine check with an IA32_MCI_STATUS.MCACOD of 0400H.

Implication: Due to this erratum, the processor may hang and may report 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](#).

BDW84. PAGE_WALKER_LOADS Performance Monitoring Event May Count Incorrectly

Problem: Due to this erratum, the PAGE_WALKER_LOADS (Event BCH) performance monitoring event may overcount or may undercount.

Implication: These performance monitoring events may not produce reliable results.

Workaround: None identified.

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

BDW85. Some DRAM And L3 Cache Performance Monitoring Events May Undercount

Problem: Due to this erratum, the supplier may be mis-attributed to unknown, and the following events may undercount:

MEM_LOAD_UOPS_RETIRED.L3_HIT (Event D1H Umask 04H)

MEM_LOAD_UOPS_RETIRED.L3_MISS (Event D1H Umask 20H)

MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_MISS (Event D2H Umask 01H)

MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HIT (Event D2H Umask 02H)

MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_HITM (Event D2H Umask 04H)

MEM_LOAD_UOPS_L3_HIT_RETIRED.XSNP_NONE (Event D2H Umask 08H)

MEM_LOAD_UOPS_L3_MISS_RETIRED.LOCAL_DRAM (Event D3H Umask 01H)

MEM_TRANS_RETIRED.LOAD_LATENCY (Event CDH Umask 01H)

Implication: The affected events may undercount, resulting in inaccurate memory profiles. For the affected events that are precise, PEBS records may be generated at incorrect points. Intel has observed incorrect counts by as much as 20%.

Workaround: None identified.

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

BDW86. The System May Hang When Executing a Complex Sequence of Locked Instructions

Problem: Under certain internal timing conditions while executing a complex sequence of locked instructions, the system may hang.

Implication: The system may hang while executing a complex sequence of locked instructions and cause an Internal Timeout Error Machine Check (IA32_MCI_STATUS.MCACOD=0400H).

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

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



BDW87. Certain Settings of VM-Execution Controls May Result in Incorrect Linear-Address Translations

Problem: If 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](#).

BDW88. An 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](#).

BDW89. Attempts to Retrain a PCIe* Link May be Ignored

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

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

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

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

BDW90. Attempting Concurrent Enabling of Intel® PT With LBR, BTS, or BTM Results in a #GP

Problem: If the LBR, the BTS, or the BTM are enabled in the IA32_DEBUGCTL MSR (1D9H), an attempt to enable Intel® PT in IA32_RTIT_CTL MSR (570H) results in a #GP. (Note that the BTM enable bit in IA32_DEBUGCTL MSR is named "TR".) Correspondingly, if Intel® PT was previously enabled when an attempt is made to enable the LBR, the BTS, or the BTM, a #GP will occur.

Implication: An unexpected #GP may occur when concurrently enabling any one of the LBR, the BTS, or the BTM with Intel® PT.

Workaround: None identified.

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

**BDW91. Setting TraceEn While Clearing BranchEn in IA32_RTIT_CTL Causes a #GP**

Problem: A WRMSR to IA32_RTIT_CTL (MSR 0570H) that sets TraceEn (bit 0) and clears BranchEn (bit 13) will cause a #GP.

Implication: Intel® Processor Trace cannot be enabled without enabling control flow trace packets.

Workaround: None identified.

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

BDW92. Processor Graphics I/O Memory Management Unit (IOMMU) Unit May Not Mask DMA Remapping Faults

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

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

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

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

BDW93. Processor Graphics IOMMU Unit May Report Spurious Faults

Problem: The IOMMU unit for Processor Graphics pre-fetches context (or extended-context) entries to improve performance. Due to the erratum, the IOMMU unit may report spurious DMA remapping faults if prefetching encounters a context (or extended-context) entry which is not marked present.

Implication: Software may observe spurious DMA remapping faults when the present bit for the context (or extended-context) entry corresponding to the Processor Graphics device (Bus: 0; Device: 2; Function: 0) is cleared. These faults may be reported when the Processor Graphics device is quiescent.

Workaround: None identified. Instead of marking a context not present, software should mark the context (or extended-context) entry present while using the page table to indicate all the memory pages referenced by the context entry is not present.

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


BDW94. Platform Environmental Control Interface (PECI) Frequency Limited to 1 MHz

Problem: The PEFI 3.1 specification's operating frequency range is 0.2 MHz to 2 MHz. Due to this erratum, PEFI may be unreliable when operated above 1 MHz.

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

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

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

**BDW95. Reads or Writes to LBRs with Intel® PT Enabled Will Result in a #GP**

Problem: On processors where the use of Intel® PT and the LBRs are mutually exclusive, reads of the LBR MSR should return 0s and writes to them should be ignored. Due to this erratum, reads and writes to the LBR MSR while IA32_RTIT_CTL MSR (570H) TraceEn bit 0 is 1 will result in a #GP.

Implication: When this erratum occurs, a #GP will occur. LBRs are not available when Intel® PT is enabled.

Workaround: None identified.

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

BDW96. Frequency Difference Between Intel Architecture (IA) Core(s) and Ring Domains May Cause Unpredictable System Behavior

Problem: Operating one or more of the IA cores at a frequency significantly higher than the ring operating frequency may cause unpredictable system behavior. Intel has observed this erratum to occur when the software explicitly requests the ring and IA core(s) to operate at different frequencies or when IA core(s) are transitioning in and out of C-states with the IA core(s) operating at frequencies much higher than the ring frequency. Exposure to this erratum may be increased when the IA cores run at or close to P0 P-state frequency.

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

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

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

BDW97. Back-to-Back Page Walks Due to Instruction Fetches May Cause a System Hang

Problem: Multiple code fetches in quick succession that generate page walks may result in a system hang causing an Internal Timer Error (an MCACOD value of 0400H) logged into IA32_MCi_STATUS bits [15:0].

Implication: Due to this erratum, the processor may hang and report 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*.


BDW98. Performance Monitoring Counters May Produce Incorrect Results for BR_INST_RETIRED Event on Logical Processor

Problem: Performance monitoring event BR_INST_RETIRED (C4H) counts retired branch instructions. Due to this erratum, when operating on logical processor 1 of any core, BR_INST_RETIRED.FAR_BRANCH (Event C4H; Umask 40H) and BR_INST_RETIRED.ALL_BRANCHES (Event C4H; Umask 04H) may count incorrectly. Logical processor 0 of all cores and cores with SMT disabled are not affected by this erratum.

Implication: Due to this erratum, certain performance monitoring event may produce unreliable results when SMT is enabled.

Workaround: None identified.

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

**BDW99. 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 DS (Debug Store) 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*.

BDW100. Some OFFCORE_RESPONSE Performance Monitoring Events Related to RFO Request Types May Count Incorrectly

Problem: The performance monitoring events OFFCORE_RESPONSE (Events B7H and BBH) should count uncore responses matching the request-response configuration specified in MSR_OFFCORE_RSP_0 (1A6H) and MSR_OFFCORE_RSP_1 (1A7H) for core-originated requests. However, due to this erratum, response type NO_SUPP bit [17] may be reported instead of LOCAL bit [26] for request types DMND_RFO bit [1] and PF_RFO bit [5].

Implication: The specified performance monitoring events may count incorrectly.

Workaround: None identified.

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

BDW101. MOVNTDQA From Write Combining (WC) Memory May Pass Earlier Locked Instructions

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


BDW102. 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*.

**BDW103. Internal Parity Errors May Incorrectly Report Overflow in the IA32_MC0_STATUS MSR**

Problem: Due to this erratum, an uncorrectable internal parity error with an IA32_MC0_STATUS.MCACOD (bits [15:0]) value of 0005H may incorrectly set the IA32_MC0_STATUS.OVER flag (bit 62) indicating an overflow when a single error has been observed

Implication: IA32_MC0_STATUS.OVER may not accurately indicate multiple occurrences of errors. There is no other impact to normal processor functionality.

Workaround: None identified.

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

BDW104. Software Using Intel® TSX May Behave Unpredictably

Problem: Under a complex set of internal timing conditions and system events, software using the Intel® TSX instructions may behave unpredictably.

Implication: This erratum may result in unpredictable behavior of the software using Intel® TSX.

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*.

BDW105. Incorrect VMCS Used for Page Modification Log (PML)-Index field on VMX Transitions Into and Out of SMM

Problem: The (PML index field is saved to an incorrect VMCS on an SMM VM exit. VM entries that return from SMM restore the PML-index field from that same incorrect VMCS.

Implication: The PML-index field is correctly maintained for expected use cases, in which the SMM-transfer monitor (STM) does not access the PML-index field in the SMM VMCS. If the STM uses VMREAD to read the field, it will get an incorrect value. In addition, the processor will ignore any modification of the field that the STM makes using VMWRITE. Intel has not observed this erratum to impact any commercially available software.

Workaround: None identified. To access the PML-index field, STM software should first load the current-VMCS pointer with a pointer to the executive VMCS.

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


BDW106. 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: It is possible for the BIOS to contain a workaround for this erratum.

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

**BDW107. An Intel® Hyper-Threading Technology (Intel® HT Technology) Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior**

Problem: Under a complex series of microarchitectural events while running Intel® HT Technology, a correctable internal parity error or unpredictable system behavior including may occur.

Implication: A correctable error (IA32_MC0_STATUS.MCACOD=0005H and IA32_MC0_STATUS.MSCOD=0001H) may be logged. The unpredictable system behavior frequently leads to faults (for example #UD, #PF, #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*.

BDW108. 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*.

BDW109. Performance Monitoring Counters May Undercount When Using a CPL Filtering

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

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

Workaround: None identified.

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


BDW110. An Intel® HT Technology Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior

Problem: Under a complex series of microarchitectural events while running Intel® HT Technology, a correctable internal parity error or unpredictable system behavior may occur.

Implication: A correctable error (IA32_MC0_STATUS.MCACOD=0005H and IA32_MC0_STATUS.MSCOD=0001H) may be logged. The unpredictable system behavior frequently leads to faults (for example #UD, #PF, #GP).

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

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



BDW111. 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](#).

BDW112. 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](#).

BDW113. 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_PERFEVTSELx MSR). Counters with both OS and USR bits set are not affected by this erratum.

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

MEM_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*.

BDW114. An x87 Store Instruction Which Pends Precision Exception (#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*.

BDW115. Load Latency Performance Monitoring Facility May Stop Counting

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

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

Workaround: None identified.

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


BDW116. General-Purpose Performance Monitoring Counters 4-7 Do Not Count With USR Mode Only Filtering

Problem: The IA32_PMC4-7 MSR (C5H-C8H) general-purpose performance monitoring counters will not count when the associated CPL filter selection in IA32_PERFVTSELx MSR's (18AH-18DH) USR field (bit 16) is set while OS field (bit 17) is not set.

Implication: Software depending upon IA32_PMC4-7 to count only USR events will not operate as expected. Counting OS only events or OS and USR events together is unaffected by this erratum.

Workaround: None identified.

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

**BDW117. 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](#).

BDW118. 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](#).

BDW119. Precise Performance Monitoring May Generate Redundant PEBS Records

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

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

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

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

**BDW120. Reads From MSR_LER_TO_LIP May Not Return a Canonical Address**

Problem: Due to this erratum, reads from MSR_LER_TO_LIP (MSR 1DEH) may return values for bits [63:61] that are not equal to bit [47].

Implication: Reads from MSR_LER_TO_LIP may return a non-canonical address where bits [63:61] may be incorrect. Using this value as an address, including restoring the MSR value that was read, may cause a #GP.

Workaround: Software should ensure the value read in MSR_LER_TO_LIP bit [47] is replicated in bits [63:61]. This is most easily accomplished by sign extending from bit [47] to bits [63:48].

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

BDW121. 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*.

BDW122. 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*.


BDW123. Using Intel® TSX Instructions May Lead to Unpredictable System Behavior

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

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

Workaround: the *Intel® Transactional Synchronization Extension* It is possible for BIOS to contain a workaround for this erratum.

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

**BDW124. Graphics Intel® VT-d Hardware May Cache Invalid Entries**

Problem: The processor's graphics IOMMU may cache invalid Intel® VT-d context entries. This violates the Intel® VT-d specification for HW Caching Mode where hardware implementations of this architecture must not cache invalid entries.

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

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

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

BDW125. Gen9 Graphics Intel® VT-d Hardware May Cache Invalid Entries

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

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

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

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

BDW126. 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*.

BDW127. Incorrect CPUID Reporting For Intel® Turbo Boost Max Technology 3.0 Capability

Problem: The CPUID instruction in leaf 06h does not report Intel® Turbo Boost Max Technology 3.0 capability in EAX[14].

Implication: Due to this Erratum, software cannot detect if Intel® Turbo Boost Max Technology 3.0 capability is supported.

Workaround: None identified.

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

BDW128. 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*.

BDW129. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

Problem: Resuming from C6 Sleep-State, with Fixed Interrupts of the same priority queued (in the corresponding bits of the Intel Reuse Repository (IRR) and Intel Strategic Research (ISR) Advanced Programmable Interrupt Controller (APIC) registers), the processor



may dispatch the second interrupt (from the IRR bit) before the first interrupt has completed and written to the EOI register, causing the first interrupt to never complete.

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

Workaround: 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, A-L*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction set reference, M-U*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C: Instruction set reference, V-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System programming guide, part 1*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System programming guide, part 2*

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






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, A-L*
- *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2B: Instruction set reference, M-U*
- *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 2C: Instruction set reference, V-Z*
- *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3A: System programming guide, part 1*
- *Intel® 64 and IA-32 Architectures Software Developer’s Manual, Volume 3B: System programming guide, part 2*

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, A-L*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction set reference, M-U*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2C: Instruction set reference, V-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System programming guide, part 1*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System programming guide, part 2*

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, 2C, 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>.

There are no new documentation changes in this specification update revision.

