Intel® Xeon® E5-2600 v4 Processor Product Family

Specification Update

April 2022

Reference Number: 333811-008US

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

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

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

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

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

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

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

Intel, the Intel logo, Intel Core, Pentium, Enhanced Intel SpeedStep Technology and Xeon 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.

Contents

Revision History

Revision	Description	Date
008	• Added BDF109	April 2022
007	• Added BDF108 • Added Nomenclature S-Spec Number	March 2022
006	• Added BDF105 - BDF107	November 2020
005	• Added BDF102 - BDF104	September 2019
004	Added BDF100 and BDF101	March 2019
003	• Added BDF86 - BDF99	November 2018
002	• Updated Table 1 and Table 2 • Added Table 5 • Removed BDF42 and BDF61 • Added BDF64 - BDF85	December 2016
001	• Initial Release	May 2015

§

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/Location
Intel® Xeon® Processor E5-2600 v4 Product Family Datasheet, Volume One: Electrical Volume 1 of 2	333809
Intel® Xeon® Processor E5 v4 Product Family Datasheet Volume 2: Registers	333810
Intel® Xeon® Processor E5 v4 Product Family Thermal Mechanical Specification and Design Guide	333812

Document Title	Document Number/ Location
AP-485, Intel® Processor Identification Utility and the CPUID Instruction	Note 1
Intel® Advanced Vector Extensions Programming Reference	Note 1
Intel® Trusted Execution Technology Server Extensions (LT-SX) BIOS Specification	Note 1
Intel® 64 and IA-32 Architecture Software Developer's Manual[1] • Volume 1: Basic Architecture • Volume 2A: Instruction Set Reference Manual A-M • Volume 2B: Instruction Set Reference Manual N-Z • Volume 3A: System Programming Guide • Volume 3B: System Programming Guide • A-32 Intel® Architecture Optimization Reference Manual	http://www.intel.com/products/processor/manuals/index.htm

Notes:

1. Contact your Intel representative for the latest revision and order number of this document.

Nomenclature

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

S-Spec Number is a five-digit code used to identify products. Products are differentiated by their unique characteristics, such as, core speed, L2 cache size, all notes associated with each S-Spec number.

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

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

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

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

Identification Information

Component Identification via Programming Interface

The Intel® Xeon® E5-2600 v4 Processor Product Family Stepping can be identified by the following register contents:

Reserved	Extended Family[1]	Extended Model[2]	Reserved	Processor Type[3]	Family Code[4]	Model Number[5]	Stepping ID[6]
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0100b		00b	0110b	1111b	varies per stepping

Notes:

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

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

Component Marking Information

Figure 1. Intel® Xeon® E5-2600 v4 Processor Product FamilyTop-side Markings (Example)

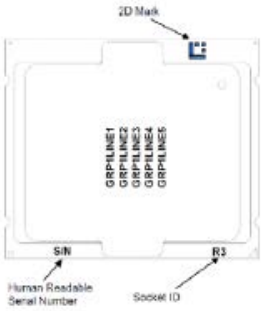

Legend:	Mark Text (Engineering Mark):
GRP1LINE1:	i{M}{C}YY
GRP1LINE2:	INTEL CONFIDENTIAL
GRP1LINE3:	QDF ES SPEED
GRP1LINE4:	XXXXX
GRP1LINE5:	{FPO} {e4}

Legend:	Mark Text (Production Mark):
GRP1LINE1:	i{M}{C}YY
GRP1LINE2:	SUB-BRAND PROC#
GRP1LINE3:	SSPEC SPEED
GRP1LINE4:	XXXXX
GRP1LINE5:	{FPO} {e4}

For the Intel® Xeon® E5-2600 Processor Product family SKUs see https://ark.intel.com/content/www/us/en/ark/products/series/91287/intel-xeon-processor-e5-v4-family.html.

§

Summary Tables of Changes

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

Codes Used in Summary Tables

Stepping

X:	Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.
(No mark) or (Blank box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Page

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

Status

Doc:	Document change or update will be implemented.
Plan Fix:	This erratum may be fixed in a future stepping of the product.
Fixed:	This erratum has been previously fixed.
No Fix:	There are no plans to fix this erratum.

Row

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

Table 1. Integrated Core/Uncore Errata (Sheet 1 of 5)

Number	Steppings	Status	Errata
	B0/M0/R0		
BDF1	X	No Fix	Enabling ISOCH Mode May Cause The System to Hang
BDF2	X	No Fix	PCI BARs in the Home Agent Will Return Non-Zero Values During Enumeration
BDF3	X	No Fix	PCIe* Header of a Malformed TLP is Logged Incorrectly
BDF4	X	No Fix	A Malformed TLP May Block ECRC Error Logging
BDF5	X	No Fix	The System May Hang During an Intel® QuickPath Interconnect (Intel® QPI) Slow to Fast Mode Transition
BDF6	X	No Fix	Unexpected Performance Loss When Turbo Disabled
BDF7	X	No Fix	Attempting to Enter ADR May Lead to Unpredictable System Behavior
BDF8	X	No Fix	Exiting From Package C3 or Package C6 With DDR4-2133 May Lead to Unpredictable System Behavior

Table 1. Integrated Core/Uncore Errata (Sheet 2 of 5)

Number	Steppings	Status	Errata
	B0/M0/R0		
BDF9	X	No Fix	The System May Shut Down Unexpectedly During a Warm Reset
BDF10	X	No Fix	CAT May Not Behave as Expected
BDF11	X	No Fix	LBR, BTS, BTM May Report a Wrong Address when an Exception/Interrupt Occurs in 64-bit Mode
BDF12	X	No Fix	EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change
BDF13	X	No Fix	MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error
BDF14	X	No Fix	LER MSRs May Be Unreliable
BDF15	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang
BDF16	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
BDF17	X	No Fix	FREEZE_WHILE_SMM Does Not Prevent Event From Pending PEBS During SMM
BDF18	X	No Fix	APIC Error "Received Illegal Vector" May be Lost
BDF19	X	No Fix	Performance Monitor Precise Instruction Retired Event May Present Wrong Indications
BDF20	X	No Fix	CR0.CD Is Ignored in VMX Operation
BDF21	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation
BDF22	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception
BDF23	X	No Fix	Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered
BDF24	X	No Fix	Pending x87 FPU Exceptions (#MF) May be Signaled Earlier Than Expected
BDF25	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
BDF26	X	No Fix	VEX.L is Not Ignored with VCVT*2SI Instructions
BDF27	X	No Fix	Processor May Livelock During On Demand Clock Modulation
BDF28	X	No Fix	Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX May Over Count
BDF29	X	No Fix	Performance Monitor Event DSB2MITE_SWITCHES.COUNT May Over Count
BDF30	X	No Fix	Timed MWAIT May Use Deadline of a Previous Execution
BDF31	X	No Fix	IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For VMCS Encoding
BDF32	X	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
BDF33	X	No Fix	Locked Load Performance Monitoring Events May Under Count
BDF34	X	No Fix	Transactional Abort May Cause an Incorrect Branch Record
BDF35	X	No Fix	PMI May be Signaled More Than Once For Performance Monitor Counter Overflow
BDF36	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
BDF37	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
BDF38	X	No Fix	A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
BDF39	X	No Fix	Intel® Processor Trace Packet Generation May Stop Sooner Than Expected
BDF40	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
BDF41	X	No Fix	Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value

Table 1. Integrated Core/Uncore Errata (Sheet 3 of 5)

Number	Steppings	Status	Errata
	B0/M0/R0		
BDF42	X	Fixed	Removed
BDF43	X	No Fix	Performance Monitoring Event INSTR_RETIRED.ALL May Generate Redundant PEBS Records For an Overflow
BDF44	X	No Fix	Reset During PECI Transaction May Cause a Machine Check Exception
BDF45	X	No Fix	Intel® Processor Trace (Intel® PT) MODE.Exec, PIP, and CBR Packets Are Not Generated as Expected
BDF46	X	No Fix	Performance Monitor Instructions Retired Event May Not Count Consistently
BDF47	X	No Fix	General-Purpose Performance Counters May be Inaccurate with Any Thread
BDF48	X	No Fix	An Invalid LBR May Be Recorded Following a Transactional Abort
BDF49	X	No Fix	Executing an RSM Instruction With Intel® Processor Trace Enabled Will Signal a #GP
BDF50	X	No Fix	Intel® Processor Trace PIP May be Unexpectedly Generated
BDF51	X	No Fix	Processor Core Ratio Changes While in Probe Mode May Result in a Hang
BDF52	X	No Fix	Processor Does Not Check IRTE Reserved Bits
BDF53	X	No Fix	PCIe* TPH Request Capability Structure Incorrectly Advertises Device Specific Mode as Supported
BDF54	X	No Fix	Package C3 State or Deeper May Lead to a Reset
BDF55	X	No Fix	VMX-Preemption Timer May Stop Operating When ACC is Enabled
BDF56	X	No Fix	Intel® Advanced Vector Extensions (Intel® AVX) Workloads May Exceed ICCMAX Limits
BDF57	X	No Fix	Writing MSR_ERROR_CONTROL May Cause a #GP
BDF58	X	No Fix	Enabling ACC in VMX Non-Root Operation May Cause System Instability
BDF59	X	No Fix	A Spurious Patrol Scrub Error May be Logged
BDF60	X	No Fix	Performance Monitoring Counters May Produce Incorrect Results for BR_INST_RETIRED Event on Logical Processor.
BDF61	X	No Fix	Removed
BDF62	X	No Fix	Processor Instability May Occur When Using The PECI RdIAMSR Command
BDF63	X	No Fix	A #VE May Not Invalidate Cached Translation Information
BDF64	X	No Fix	Package C-state Transitions While Inband PECI Accesses Are in Progress May Cause Performance Degradation
BDF65	X	No Fix	Attempting Concurrent Enabling of Intel® Processor Trace (Intel® PT) With LBR, BTS, or BTM Results in a #GP
BDF66	X	No Fix	A DDR4 C/A Parity Error in Lockstep Mode May Result in a Spurious Uncorrectable Error
BDF67	X	No Fix	Cores May be Unable to Reach Maximum Turbo Frequency
BDF68	X	No Fix	PEBS Record May Be Generated After Being Disabled
BDF69	X	No Fix	Software Using Intel® TSX May Behave Unpredictably
BDF70	X	No Fix	Some E5-1607V4 And E5-1603V4 Parts Will Incorrectly Report Support For DDR4-2400
BDF71	X	No Fix	PROCHOT# Assertion During Warm Reset May Cause Persistent Performance Reduction
BDF72	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May be Lost
BDF73	X	No Fix	Internal Parity Errors May Incorrectly Report Overflow in the IA32_MC0_STATUS MSR
BDF74	X	No Fix	Incorrect VMCS Used for PML-Index field on VMX Transitions Into and Out of SMM

Table 1. Integrated Core/Uncore Errata (Sheet 4 of 5)

Number	Steppings	Status	Errata
	B0/M0/R0		
BDF75	X	No Fix	Certain Microcode Updates May Result in Incorrect Throttling Causing Reduced System Performance
BDF76	X	No Fix	An Intel® Hyper-Threading Technology Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior
BDF77	X	No Fix	Inband PECI Concurrent With OS Patch Load May Result in Incorrect Throttling Causing Reduced System Performance
BDF78	X	No Fix	Writing The IIO_LLC_WAYS MSR Results in an Incorrect Value
BDF79	X	No Fix	Turbo May Be Delayed After Exiting C6 When Using HWP
BDF80	X	No Fix	IA32_MC4_STATUS.VAL May be Incorrectly Cleared by Warm Reset
BDF81	X	No Fix	Interrupt Remapping May Lead to a System Hang
BDF82	R0	No Fix	MEM_HOT_C23_N DIMM Temperature Reporting Does Not Function Correctly
BDF83	X	No Fix	Bi-Directional PCIe* Posted Transactions May Lead to System Hang
BDF84	X	No Fix	Excessive Uncorrected and Corrected Memory Errors May Occur Following S3 Resume or Warm Reset
BDF85	X	No Fix	Writing MSR_LASTBRANCH_x_FROM_IP May #GP When Intel® TSX is Not Supported
BDF86	X	No Fix	PCIe* Ports May Incorrectly Detect a Receiver
BDF87	X	No Fix	Some DRAM And L3 Cache Performance Monitoring Events May Undercount
BDF88	X	No Fix	APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode
BDF89	X	No Fix	Loading Microcode Updates or Executing an Authenticated Code Module May Result in a System Hang
BDF90	X	No Fix	NVDIMM Data May Not be Preserved Correctly on Power Loss or ADR Activation
BDF91	X	No Fix	Link Down Events behind PCIe Device connected to CPU Root Ports Can Cause CTO > 50ms on other Root Ports
BDF92	X	No Fix	Removed
BDF93	X	No Fix	Reads From MSR_LER_TO_LIP May Not Return a Canonical Address
BDF94	X	No Fix	Processor May Hang After Multiple Microcode Updates Loaded
BDF95	X	No Fix	In eMCA2 Mode, When The Retirement Watchdog Timeout Occurs CATERR# May be Asserted
BDF96	X	No Fix	Systems That Enable Both OSB And IODC May Exhibit Unexpected System Behavior
BDF97	X	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store
BDF98	X	No Fix	Writing the CMCI_DISABLE bit in the ERROR_CONTROL MSR will #GP Fault
BDF99	X	No Fix	System May Hang When Operating at High Temperatures
BDF100	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
BDF101	X	No Fix	The System May Hang When Exiting package C6 State
BDF102	X	No Fix	Intel® MBM Counters May Report System Memory Bandwidth Incorrectly
BDF103	X	No Fix	When Operating at Maximum Uncore Frequency, The Processor May Hang
BDF104	X	No Fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes
BDF105	X	No Fix	Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set
BDF106	X	No Fix	PECI RdAMSR() to Non-existent MSR Might Result in IERR

Table 1. Integrated Core/Uncore Errata (Sheet 5 of 5)

Number	Steppings	Status	Errata
	B0/M0/R0		
BDF107	X	No Fix	System logs a message channel timeout in MCA Bank
BDF108.	X	No Fix	Retried PECI PCIConfigLocal Register Accesses May Not Operate Correctly
BDF109.	X	No Fix	Debug Exceptions May Be Lost in The Case Of Machine Check Exception

Specification Changes

Number	SPECIFICATION CHANGES
1	None for this revision of this specification update.

Specification Clarifications

No.	SPECIFICATION CLARIFICATIONS
1	None for this revision of this specification update.

Documentation Changes

No.	DOCUMENTATION CHANGES
1	None for this revision of this specification update.

Integrated Core/Uncore Errata

BDF1 Enabling ISOCH Mode May Cause The System to Hang

Problem: When ISOCH (Isochronous) operation is enabled within BIOS, the system may hang and fail to boot.

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

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*.

BDF2 PCI BARs in the Home Agent Will Return Non-Zero Values During Enumeration

Problem: During system initialization the Operating System may access the standard PCI BARs (Base Address Registers). Due to this erratum, accesses to the Home Agent BAR registers (Bus 1; Device 18; Function 0,4; Offsets 0x14-0x24) will return non-zero values.

Implication: The operating system may issue a warning. Intel has not observed any functional failures due to this erratum.

Workaround: None identified.

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

BDF3 PCIe* Header of a Malformed TLP is Logged Incorrectly

Problem: If a PCIe port receives a malformed TLP (Transaction Layer Packet), an error is logged in the UNCERRSTS register (Device 0; Function 0; Offset 14CH and Device 2-3; Function 0-3; Offset 14CH). Due to this erratum, the header of the malformed TLP is logged incorrectly in the HDRLOG register (Device 0; Function 0; Offset 164H and Device 2-3; Function 0-3; Offset 164H).

Implication: The PCIe header of a malformed TLP is not logged correctly.

Workaround: None identified.

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

BDF4 A Malformed TLP May Block ECRC Error Logging

Problem: If a PCIe* port receives a Malformed TLP that also would generate an ECRC Check Failed error, it should report a Malformed TLP error. When Malformed TLP errors are masked, the processor should report the lower-precedence ECRC Check Failed error but, due to this erratum, it does not.

Implication: Software that relies upon ECRC Check Failed error indication may not behave as expected.

Workaround: None identified.

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

BDF5 The System May Hang During an Intel® QuickPath Interconnect (Intel® QPI) Slow to Fast Mode Transition

Problem: During an Intel QPI slow mode to fast mode transition, the LL_STATUS field of the QPIPCSTS register (Bus 0; Device 8,9,10; Function 0; Offset 0xc0) may not be correctly updated to reflect link readiness.

Implication: The system may hang waiting for the QPIPCSTS.LL_STATUS to update.

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*.

BDF6 Unexpected Performance Loss When Turbo Disabled

Problem: When Intel Turbo Boost Technology is disabled by IA32_MISC_ENABLES MSR (416H) TURBO_MODE_DISABLE bit 38, the Ring operating frequency may be below P1 operating frequency.

Implication: Processor performance may be below expectations for P1 operating frequency.

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*.

BDF7 Attempting to Enter ADR May Lead to Unpredictable System Behavior

Problem: Due to this erratum, an attempt to transition the memory subsystem to ADR (Asynchronous DRAM Self Refresh) mode may fail.

Implication: This erratum may lead to 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*.

BDF8 Exiting From Package C3 or Package C6 With DDR4-2133 May Lead to Unpredictable System Behavior

Problem: Due to this erratum, with DDR4-2133 memory, exiting from PC3 (package C3) or PC6 (package C6) state may lead to unpredictable system behavior.

Implication: This erratum may lead to 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*.

BDF9 The System May Shut Down Unexpectedly During a Warm Reset

Problem: Certain complex internal timing conditions present when a warm reset is requested can prevent the orderly completion of in-flight transactions. It is possible under these conditions that the warm reset will fail and trigger a full system shutdown.

Implication: When this erratum occurs, the system will shut down and all machine check error logs will be lost.

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*.

BDF10 CAT May Not Behave as Expected

Problem: Due to this erratum, CAT (Cache Allocation Technology) way enforcement may not behave as configured.

Implication: When this erratum occurs, cache quality of service guarantees may not be met.

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*.

BDF11 LBR, BTS, 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 (Last Branch Record), BTS (Branch Trace Store) and BTM (Branch Trace Message) mechanisms. However, during a specific boundary condition where the exception/interrupt occurs right after the execution of an instruction at the lower canonical boundary (0x00007FFFFFFFFFFF) 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*.

BDF12 EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change

Problem: EFLAGS Discrepancy on Page Faults and on EPT-Induced 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 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*.

BDF13 MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error

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

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

Workaround: None identified.

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

BDF14 LER MSRs May Be Unreliable

Problem: Due to certain internal processor events, updates to the LER (Last Exception Record) 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*.

BDF15 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*.

BDF16 #GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code

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

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

Workaround: None identified.

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

BDF17 FREEZE_WHILE_SMM Does Not Prevent Event From Pending 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 SMM (System Management Mode). Due to this erratum, if

1. A performance counter overflowed before an 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*.

BDF18 APIC Error "Received Illegal Vector" May be Lost

Problem: APIC (Advanced Programmable Interrupt Controller) may not update the ESR (Error Status Register) 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*.

BDF19 Performance Monitor Precise Instruction Retired Event May Present Wrong Indications

Problem: When the PDIR (Precise Distribution for Instructions Retired) mechanism is activated (INST_RETIRED.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*.

BDF20 CR0.CD Is Ignored in VMX Operation

Problem: If CR0.CD=1, the 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.*

BDF21 Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation

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

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

Workaround: Software should not write to a paging-structure entry in a way that would change, for any linear address, both the page size and the memory type. It can instead use the following algorithm: first clear the P flag in the relevant paging-structure entry (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 and memory type.

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

BDF22 Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception

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

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

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

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

BDF23 Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered

Problem: If the local-APIC timer's CCR (current-count register) 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 LVT timer register and then reading the bit in the IRR (interrupt-request register) 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 DCR (Divide Configuration Register) 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.*

BDF24 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-signaled before pending interrupts are serviced.

Workaround: None identified.

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

BDF25 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.*

BDF26 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.*

BDF27 Processor May Livelock During On Demand Clock Modulation

Problem: The processor may livelock when (1) a processor thread has enabled on demand clock modulation via bit 4 of the IA32_CLOCK_MODULATION MSR (19AH) and the clock modulation duty cycle is set to 12.5% (02H in bits 3:0 of the same MSR), 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 cacheline or access UC memory.

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

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

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

BDF28 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 HLE (Hardware Lock Elision) 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.*

BDF29 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 DSB (Decode Stream Buffer) to MITE (Macro Instruction Translation Engine) 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.*

BDF30 Timed MWAIT May Use Deadline of a Previous Execution

Problem: A timed MWAIT instruction specifies a 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.*

BDF31 IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For 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.*

BDF32 Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

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

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

Workaround: None identified.

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

BDF33 Locked Load Performance Monitoring Events May Under Count

Problem: The performance monitoring events MEM_TRANS_RETIRED.LOAD_LATENCY (Event CDH; Umask 01H), MEM_LOAD_RETIRED.L2_HIT (Event D1H; Umask 02H), and MEM_UOPS_RETIRED.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.*

BDF34 Transactional Abort May Cause 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 (Last Branch Record) is not correctly reported.

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

Workaround: None identified.

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

BDF35 PMI May be Signaled More Than Once For Performance Monitor Counter Overflow

Problem: Due to this erratum, PMI (Performance Monitoring Interrupt) 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 EOI (end-of- Interrupt) register write for the interrupt until after the overflow indications have been cleared.

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

BDF36 Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception

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

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

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

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

BDF37 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.*

BDF38 A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation

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

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

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

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

BDF39 Intel® Processor Trace 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 (Processor Trace) 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.*

BDF40 PEBS Eventing IP Field May be Incorrect After Not-Taken Branch

Problem: When a PEBS (Precise-Event-Based-Sampling) 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.*

BDF41 Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value

Problem: An HLE (Hardware Lock Elision) 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.*

BDF42 Removed

BDF43 Performance Monitoring Event INSTR_RETIRED.ALL May Generate Redundant PEBS Records For an Overflow

Problem: Due to this erratum, the performance monitoring feature PDIR (Precise Distribution of Instructions Retired) for INSTR_RETIRED.ALL (Event C0H; Umask 01H) will generate redundant PEBS (Precise Event Based Sample) 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.*

BDF44 Reset During PECI Transaction May Cause a Machine Check Exception

Problem: If a PECI transaction is interrupted by a warm reset, it may result in a machine check exception with MCACOD of 0x402.

Implication: When this erratum occurs, the system becomes unresponsive and a machine check will be generated.

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.*

BDF45 Intel® Processor Trace (Intel® PT) MODE.Exec, PIP, and CBR Packets Are Not Generated as Expected

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

BDF46 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.*

BDF47 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.

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.*

BDF48 An Invalid LBR May Be Recorded Following a Transactional Abort

Problem: Use of Intel® Transactional Synchronization Extensions may result in a transactional abort. If an abort occurs immediately following a branch instruction, an invalid LBR (Last Branch Record) may be recorded before the LBR produced by the abort.

Implication: The invalid LBR may interfere with execution path reconstruction prior to the transactional abort.

Workaround: None identified.

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

BDF49 Executing an RSM Instruction With Intel® Processor Trace Enabled Will Signal a #GP

Problem: Upon delivery of an SMI (System Management Interrupt), the processor saves and then clears TraceEn in the IA32_RTIT_CTL MSR (570H), thus disabling Intel® Processor Trace (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 (general-protection exception).

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

Workaround: If software enables 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.*

BDF50 Intel® Processor Trace PIP May be Unexpectedly Generated

Problem: When Intel® Processor Trace is enabled, PSB+ (Packet Stream Boundary) packets may include a PIP (Paging Information Packet) 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 BIOS to contain a workaround for this erratum.

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

BDF51 Processor Core Ratio Changes While in Probe Mode May Result in a Hang

Problem: If a processor core ratio change occurs while the processor is in probe mode, the system may hang.

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

Workaround: None identified. Processor core ratio changes may be disabled to avoid this erratum.

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

BDF52 Processor Does Not Check IRTE Reserved Bits

Problem: As per the Intel® Virtualization Technology for Directed I/O (Intel® VT-d) specification, bits 63:HAW (Host Address Width) of the Posted Interrupt Descriptor Upper Address field in the IRTE (Interrupt Remapping Table Entry) must be checked for a value of 0; violations must be reported as an interrupt-remapping fault. Due to this erratum, hardware does not perform this check and does not signal an interrupt-remapping fault on violations.

Implication: If software improperly programs the reserved address bits of posted interrupt descriptor upper address in the IRTE to a value other than zero, hardware will not detect and report the violation.

Workaround: Software must ensure posted interrupt address bits 63:HAW in the IRTE are zero.

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

BDF53 PCIe* TPH Request Capability Structure Incorrectly Advertises Device Specific Mode as Supported

Problem: The TPH (Transaction layer packet Processing Hints) Requester Capability Structure (PCI Express Extended Capability ID type 0017H) incorrectly reports that Device Specific Mode is supported in its TPH Requester Capability Register (bit 2 at offset 04H in the capability structure).

Implication: The processor supports only No ST (Steering Tag) Mode. The PCI Express Base Specification allows, in this instance, the TPH Requester Capability Structure's TPH Requester Control Register (at offset 08H) bits 2:0 to be hardwired to '000', forcing No ST Mode. Advertising Device Specific Mode but forcing No ST Mode is a violation of the PCI Express Base Specification (and may be reported as a compliance issue). Intel has not observed this erratum to impact the operation of any commercially available system.

Workaround: None identified.

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

BDF54 Package C3 State or Deeper May Lead to a Reset

Problem: Due to this erratum, the processor may reset and signal a Machine Check error with a IA32_MCi_STATUS.MCACOD value of 0400H when in Package C3 state or deeper.

Implication: When this erratum occurs, the processor will reset and report an uncorrectable machine check error.

Workaround: A BIOS code change has been identified and may be implemented as a workaround for this erratum. It is possible for the BIOS to contain a workaround for this erratum

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

BDF55 VMX-Preemption Timer May Stop Operating When ACC is Enabled

Problem: When the MSR_PKG_CST_CONFIG_CONTROL.ACC_Enable bit (MSR E2H, bit 16) is set, the VMX-preemption timer is not decremented in the HLT state.

Implication: When ACC (Autonomous C-State Control) is enabled, the VMX-preemption timer may not cause a VM exit when 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*.

BDF56 Intel® Advanced Vector Extensions (Intel® AVX) Workloads May Exceed ICCMAX Limits

Problem: Intel AVX workloads require a reduced maximum turbo ratio. Due to this erratum, the Intel AVX turbo ratio is higher than expected which may cause the processor to exceed ICCMAX limits and lead to unpredictable system behavior.

Implication: Due to this erratum, the processor may exhibit 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*.

BDF57 Writing MSR_ERROR_CONTROL May Cause a #GP

Problem: A WRMSR that attempts to set MODE1_MEMERROR_REPORT field (bit 1) and/or MEM_CORRERR_LOGGING_DISABLE field (bit 5) of the MSR_ERROR_CONTROL MSR (17FH) may incorrectly cause a #GP (General Protection exception).

Implication: Due to this erratum, if BIOS attempts to change the value of the listed bits, a #GP 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*.

BDF58 Enabling ACC in VMX Non-Root Operation May Cause System Instability

Problem: ACC (Autonomous C-State Control) is enabled by setting ACC_Enable (bit 16) of MSR_PKG_CST_CONFIG_CONTROL (E2H) to '1'. If ACC is enabled while the processor is in VMX non-root operation, an unexpected VM exit, a machine check, or unpredictable system behavior may result.

Implication: Enabling ACC may lead to system instability.

Workaround: None identified. BIOS should not enable ACC.

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

BDF59 A Spurious Patrol Scrub Error May be Logged

Problem: When a memory ECC error occurs, a spurious patrol scrub error may also be logged on another memory channel.

Implication: A patrol scrub correctable error may be incorrectly logged.

Workaround: The Home Agent error registers and correctable error count registers (Bus 1; Device 20; Function 2; Offset 104-110) provides accurate error information.

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

BDF60 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*.

BDF61 Removed

BDF62 Processor Instability May Occur When Using The PECI RdIAMSR Command

Problem: Under certain circumstances, reading a machine check register using the PECI (Platform Environmental Control Interface) RdIAMSR command may result in a machine check, processor hang or shutdown.

Implication: Machine check, hang or shutdown may be observed when using the PECI RdIAMSR command.

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*.

BDF63 A #VE May Not Invalidate Cached Translation Information

Problem: An EPT (Extended Page Table) violation that causes a #VE (virtualization exception) 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*.

BDF64 Package C-state Transitions While Inband PECI Accesses Are in Progress May Cause Performance Degradation

Problem: When a Package C-state transition occurs at the same time an inband PECI transaction occurs, PROCHOT# may be incorrectly asserted.

Implication: Incorrect assertion of PROCHOT# reduces the core frequency to the minimum operating frequency of 1.2 GHz resulting in persistent performance degradation.

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*.

BDF65 Attempting Concurrent Enabling of Intel® Processor Trace (Intel® PT) With LBR, BTS, or BTM Results in a #GP

Problem: If LBR (Last Branch Records), BTS (Branch Trace Store), or BTM (Branch Trace Messages) are enabled in the IA32_DEBUGCTL MSR (1D9H), an attempt to enable Intel PT (Intel® Processor Trace) in IA32_RTIT_CTL MSR (570H) results in a #GP (general protection exception). (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 LBR, BTS, or BTM, a #GP will occur.

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

Workaround: None identified.

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

BDF66 A DDR4 C/A Parity Error in Lockstep Mode May Result in a Spurious Uncorrectable Error

Problem: If a memory C/A (Command/Address) parity error occurs while the memory subsystem is configured in lockstep mode then the channel that observed the error will properly log the error but the associated channel in lockstep will incorrectly log an uncorrectable error in its IA32_MCi_STATUS MSR.

Implication: Due to this erratum, incorrect logging of an uncorrectable memory error in IA32_MCi_STATUS may occur.

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*.

BDF67 Cores May be Unable to Reach Maximum Turbo Frequency

Problem: Due to this erratum, processors with more than ten cores may be limited to less than the specified maximum turbo frequency.

Implication: When this erratum occurs, the processor performance is reduced.

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*.

BDF68 PEBS Record May Be Generated After Being Disabled

Problem: A performance monitoring counter may generate a PEBS (Precise Event Based Sampling) 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*.

BDF69 Software Using Intel® TSX May Behave Unpredictably

Problem: Under a complex set of internal timing conditions and system events, software using the Intel TSX (Intel Transactional Synchronization Extensions) 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*.

BDF70 Some E5-1607V4 And E5-1603V4 Parts Will Incorrectly Report Support For DDR4-2400

Problem: Some E5-1607V4 and E5-1603V4 parts will incorrectly report that they support DDR4-2400. Using DDR4-2400 DIMMs may result in unpredictable system behavior.

Implication: System may operate their memory sub-systems at DDR4-2400 rather than DDR4-2133.

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*.

BDF71 PROCHOT# Assertion During Warm Reset May Cause Persistent Performance Reduction

Problem: Assertion of PROCHOT# after RESET# de-assertion but before BIOS has completed reset initialization (indicated by CPL3) may result in persistent processor throttling. Asserting PROCHOT# during and after RESET# assertion for FRB (Fault Resilient Boot) tri-stating of the processor is not affected by this erratum.

Implication: When this erratum occurs, the resultant persistent throttling substantially reduces the processor's performance.

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*.

BDF72 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*.

BDF73 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*.

BDF74 Incorrect VMCS Used for PML-Index field on VMX Transitions Into and Out of SMM

Problem: The PML (Page Modification Log) 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 STM (SMM-transfer monitor) 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*.

BDF75 Certain Microcode Updates May Result in Incorrect Throttling Causing Reduced System Performance

Problem: Microcode updates with signature less than 0B000017 loaded by the operating system may result in excessive and persistent throttling that significantly reduces system performance.

Implication: When this erratum occurs, reduced performance may occur, concurrent with an incorrect assertion of the PROCHOT# signal.

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*.

BDF76 An Intel® Hyper-Threading Technology Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior

Problem: Under a complex series of microarchitectural events while running Intel Hyper-Threading 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 (e.g. #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*.

BDF77 Inband PECI Concurrent With OS Patch Load May Result in Incorrect Throttling Causing Reduced System Performance

Problem: Microcode updates loaded by the operating system may result in excessive and persistent throttling that significantly reduces system performance.

Implication: When this erratum occurs, performance may be reduced, concurrent with an incorrect assertion of the PROCHOT# signal.

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*.

BDF78 Writing The IIO_LLC_WAYS MSR Results in an Incorrect Value

Problem: Writing the IIO_LLC_WAYS MSR (C8Bh) always sets bits [1:0] regardless of the value written.

Implication: IIO cache way allocation may not act as intended. Intel has not seen any functional failure due to this erratum.

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

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

BDF79 Turbo May Be Delayed After Exiting C6 When Using HWP

Problem: Due to this erratum, enabling HWP (Hardware-Controlled Performance States) by setting bit 0 of IA32_PM_ENABLE (MSR 770H) may lead to an unexpected delay in reaching turbo frequencies after a core exits C6 sleep state. This erratum does not occur when HWP is not enabled.

Implication: When this erratum occurs, enabling HWP may lead to a visible reduction of system performance.

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*.

BDF80 IA32_MC4_STATUS.VAL May be Incorrectly Cleared by Warm Reset

Problem: Due to this erratum, the IA32_MC4_STATUS. VAL (MSR 411H, bit 63) may be incorrectly cleared by a warm reset.

Implication: Software may be unaware that a machine check occurred before the warm reset.

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

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

BDF81 Interrupt Remapping May Lead to a System Hang

Problem: Under complex micro-architectural conditions, back-to-back interrupt requests when interrupt remapping is enabled may lead to a system hang.

Implication: When this erratum occurs, the system hang may be associated with a queued invalidation of the IOAPIC that does not complete.

Workaround: None Identified.

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

BDF82 MEM_HOT_C23_N DIMM Temperature Reporting Does Not Function Correctly

Problem: On single HA (Home Agent) systems, the MEM_HOT_C23_N signal can be configured as an output signal that is asserted when a DIMM temperature exceeds the throttle threshold (c.f. dimm_temp_th CSRs at Bus: 1; Device: 20; Function: 0,1; Offset: 120H, 124H). Due to this erratum, MEM_HOT_C23_N is not asserted when it should be.

Implication: Platforms that rely on the MEM_HOT_C23_N for DIMM temperature-based throttling will not behave as expected, potentially leading to unpredictable system behavior, excessive DIMM aging, and DIMM failure. This erratum does not affect MEM_HOT_C23_N when configured as an input.

Workaround: Single HA platforms should use Open Loop Thermal Throttling for DIMM temperature control, use MEM_HOT_C01_N as a proxy for MEM_HOT_C23_N, or have the BMC (or other external agent) periodically read the DIMM temperature via PECI then use the

MEM_HOT_C23_N signal as an input to throttle DIMM activity as needed. See Grantley Platform Design Guide Rev. 2.2, IBL ID: 506549 for further details.

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

BDF83 Bi-Directional PCIe* Posted Transactions May Lead to System Hang

Problem: Certain bi-directional PCIe posted traffic patterns between CPU nodes may lead to a loss of flow control credits resulting in a link hang.

Implication: Implication: Deadlock on a PCIe link may result in a system hang.

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*.

BDF84 Excessive Uncorrected and Corrected Memory Errors May Occur Following S3 Resume or Warm Reset

Problem: Following S3 resume or warm reset, uncorrected and corrected memory errors may occur.

Implication: When this erratum occurs, the system will log correctable errors, signal a machine check, or shut down.

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*.

BDF85 Writing MSR_LASTBRANCH_x_FROM_IP May #GP When Intel® TSX is Not Supported

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

Implication: The value read from MSR_LASTBRANCH_x_FROM_IP 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, 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*.

BDF86 PCIe* Ports May Incorrectly Detect a Receiver

Problem: CC_SPARE (Bus 0; Device 6; Function 7; Offset 0x650) and CC_SPARE2 (Bus 0; Device 6 Function 7; Offset 0x64C) registers are not cleared on a warm reset.

Implication: PCIe ports that are configured using CC_SPARE and/or CC_SPARE2 to enable their lowest power state may incorrectly detect a receiver following a warm reset.

Workaround: None identified. Do not set CC_SPARE or CC_SPARE2 registers bits.

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

BDF87 Some DRAM And L3 Cache Performance Monitoring Events May Undercount

Problem: Due to this erratum, the supplier may be misattributed 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*.

BDF88 APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode

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

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

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

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

BDF89 Loading Microcode Updates or Executing an Authenticated Code Module May Result in a System Hang

Problem: An uncorrectable error (IA32_MC3_STATUS.MCACOD=0400 and IA32_MC3_STATUS.MSCOD=0080) may be logged for processors that have more than 2.5MB last-level-cache per core on attempting to load a microcode update or execute an authenticated code module. This issue does not occur with microcode updates with a signature of 0x0b000021 and greater.

Implication: Due to this erratum, the processor may hang when attempting to load a microcode update or execute an authenticated code module.

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.

BDF90 NVDIMM Data May Not be Preserved Correctly on Power Loss or ADR Activation

Problem: When entering ADR (Asynchronous DRAM Self-Refresh), whether through power loss or a specific ADR command, concurrent reads to the NVDIMM may prevent the data from being properly preserved.

Implication: After an ADR event volatile data may be incorrect and may lead to an ECC error on next access.

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.

BDF91 Link Down Events behind PCIe Device connected to CPU Root Ports Can Cause CTO > 50ms on other Root Ports

Problem: When a downstream switch connected to a CPU Root Port experiences a link down it may cause a back pressure event that prevents other CPU root ports from completing transaction for >50ms but less than 100ms

Implication: When intentionally disabling a PCIe link in the system the IIO Arbiter can get stuck for > 50ms causing other endpoints to exceed their CT value (of 50ms) which is reported as a fatal system ERR2 condition

Workaround: Set PCIe CTOs to 100ms or greater if in a vulnerable configuration.

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

BDF92 Removed due to inapplicability

BDF93 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.

BDF94 Processor May Hang After Multiple Microcode Updates Loaded

Problem: Under certain conditions, a microcode update load may hang if another microcode update was already loaded, resulting in an Internal Timer Error Machine Check (IA32_MCi_STATUS.MCACOD=400H; bits 15:0).

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.

BDF95 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.

BDF96 Systems That Enable Both OSB And IODC May Exhibit Unexpected System Behavior

Problem: If a platform with four or more sockets is configured to enable both OSB (Opportunistic Snoop Broadcast) and IODC (Input Output Directory Cache) or if a platform with two or more sockets is configured to enable OSB, IODC, and CoD (Cluster on Die), then the system may exhibit unexpected behavior.

Implication: Due to this erratum, the system may exhibit unexpected system behavior.

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

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

BDF97 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 (e.g., #PF) or VM exit. The value written to the MXCSR exceptions flags is what would have been written if there were no fault.

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

Workaround: None identified.

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

BDF98 Writing the CMCI_DISABLE bit in the ERROR_CONTROL MSR will #GP Fault

Problem: Writing a 1 to bit 4 (CMCI_DISABLE) of MSR_ERROR_CONTROL (MSR 0x17F) will result in a #GP Fault and the bit will not be set.

Implication: When this erratum occurs, the system will #GP Fault and the CMCI (Corrected Machine Check Interrupt) will not be disabled.

Workaround: None identified. Software cannot disable CMCI.

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

BDF99 System May Hang When Operating at High Temperatures

Problem: When operating at high temperatures above 85^{o}C, the processor may hang with the CATERR# pin asserted. There may be a Machine Check Exception (IA32_MC19_STATUS Valid Bit (bit 63) will be set to 0) and hang.

Implication: Due to this erratum, the processor may hang or fail to boot.

Workaround: It is possible for a BIOS to contain a workaround for this erratum. Please refer to Table 7 (microcode table) and/or release notes. This erratum affects the following QDFs: R2NC, R2NE, R2SE, R2SB, R2SJ, R2S4, and R2S5.

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

BDF100 Using Intel® TSX Instructions May Lead to Unpredictable System Behavior

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

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

Workaround: It is possible for BIOS to contain a workaround for this erratum. Please see the Intel® White Paper “Performance Monitoring Impact of TSX Memory Ordering Issue” Doc ID#604224 or contact your Intel® Representative for more information.Please refer to Table 7 (microcode table) and/or release notes.

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

BDF101 The System May Hang When Exiting package C6 State

Problem: Under complex microarchitectural conditions, a package C6 exit may not complete, which will lead to a system hang with a resulting machine check error (IA32_MCi_STATUS.MCACOD=0402H and IA32_MCi_STATUS.MSCOD= (0900H, 7100H, or 7300H)

Implication: When this erratum occurs, the system will hang.

Workaround: It is possible for BIOS to contain a workaround for this erratum. Please refer to Table 7 (microcode table) and/or release notes.

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

BDF102 Intel® MBM Counters May Report System Memory Bandwidth Incorrectly

Problem: Intel® Memory Bandwidth Monitoring (MBM) counters track metrics according to the assigned Resource Monitor ID (RMID) for that logical core. The IA32_QM_CTR register

(MSR 0xC8E), used to report these metrics, may report incorrect system bandwidth for certain RMID values.

Implication: Due to this erratum, system memory bandwidth may not match what is reported.

Workaround: It is possible for software to contain code changes to work around this erratum. Please see the white paper titled Intel® Resource Director Technology (Intel® RDT) Reference Manual found at https://software.intel.com/en-us/intel-resource-director-technology-rdt-reference-manual for more information.

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

BDF103 When Operating at Maximum Uncore Frequency, The Processor May Hang

Problem: Under rare microarchitectural conditions, if the processor is operating at maximum uncore frequency, it may hang without logging any machine check exceptions (MCEs) or other internal errors (IERRs).

Implication: When this erratum occurs, the system will hang with no active machine check exceptions or other internal errors.

Workaround: It is possible for software to limit the maximum ratios at which the processor may operate.

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

BDF104 A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

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

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

Workaround: A workaround for this erratum is available in microcode. Please see Microcode Updates above in this document.

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

BDF105 Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set

Problem: Under complex micro-architectural conditions, a single internal parity error seen in IA32_MC0_STATUS MSR (401h) with MCACOD (bits 15:0) value of 5h and MSCOD (bits 31:16) value of 7h, may set the overflow flag (bit 62) in the same MSR.

Implication: Due to this erratum, the IA32_MC0_STATUS overflow flag may be set after a single parity error. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

BDF106 PECI RdAMSR() to Non-existent MSR Might Result in IERR

Problem: If a PECI RdIAMSR() transaction is initiated to a non-existent MSR before any microcode patch is loaded (either via FIT or by BIOS), the system may signal IERR# with bank 4 logging a MCCOD of 0x0403 and a `MSEC_FW` code of 0x09.

Implication: Due to this erratum, PECI accesses to unimplemented MSRs may result in IERR# assertion.

Workaround: Avoid RdIAMSR() access to unimplemented MSRs or delay all PECI transactions until 500 ms after RESET_N de-assertion.

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

BDF107 System logs a message channel timeout in MCA Bank

Problem: PCU unable to communicate on the message channel logs a message channel timeout.

Implication: System will hang or reboot with CATERR logged in SEL log and machine check bank 4 populated with 0x73 or 0x9. All cores in C6 (voltage off). Ring pll is off (socket had entered pkg-C6). pkgC fsm state is PKGC_UNBLOCK_L1_EXIT.

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

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

BDF108. Retried PECI PCIConfigLocal Register Accesses May Not Operate Correctly

Problem: When the processor requests a PECI PCIConfigLocal Read or Write command to be retried, and the PECI host immediately retries the command (within 150 us), the processor may fail to correctly process the retried PECI command.

Implication: Due to this erratum, the PECI PCIConfigLocal Read command may return incorrect data, and the PECI PCIConfigLocal Write command may incorrectly update the target.

Workaround: None identified.

Status: No Fix.

BDF109. Debug Exceptions May Be Lost in The Case Of Machine Check Exception

Problem: If both a machine check exception and a debug exception are pending on the same instruction boundary, then the machine check exception gets priority and the debug exception may be lost, even if the PCC (processor context corrupted) field is cleared in all of the machine check banks (bit 57=0 in all IA32_MCi_STATUS MSR). This can happen in the case that an instruction triggered a data breakpoint while an unrelated machine check event was received.

Implication: Debugging software may fail to operate as expected if a debug exception is lost.

Workaround: None identified.

Status: No Fix.

§