Intel® Xeon® D-1500/D-1500 NS and Intel® Xeon® D-1600 NS Processor Families

Specification Update

May 2020

Reference Number: 332054-022US

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.

Contents

§

Revision History

Revision	Description	Date
022	Added BDE121	May 2020
021	Added Note 5 to documentation changes.	March 2020
020	Added PCH27. Updated PCH25 Updated wording for BDE21.	November 2019
019	Added BDE118, BDE119 and BDE120	August 2019
018	Added BDE117 updated Table 2	May 2019
017	Added BDE115 and BDE116.	March 2019
016	Corrected impacted steppings of BDE41. Added BDE114. Added SKU R2DH to table 1.	February 2018
015	Corrected workaround for BDE94 and BDE113.	November 2017
014	Added Intel® Xeon® D-1500 NS Processor Family. Added Erratum BDE111 through BDE113.	September 2017
013	Added Intel® Xeon® D-1500 NS Processor Family.	July 2017
012	Corrected Erratum BDE109. Added Erratum BDE110.	February 2017
011	Deleted Errata BDE30 and BDE42 since they do not apply to this product. Deleted erratum BDE64 as it was a duplicate of BDE55. Fixed title of Erratum BDE94 in Table 2. Corrected status of BDE49 in Table 2.	February 2017
010	Added Erratum:BDE109	January 2017
009	Added Errata: BDE104, BDE105, BDE106, BDE107, BDE108	December 2016
008	Added Erratum: BDE94, BDE95, BDE96, BDE97, BDE98, BDE99, BDE100, BDE101, BDE102, BDE103, PCH26 Modified Erratum: BDE71, PCH25	November 2016
007	Added Errata: BDE92, BDE93	October 2016
006	Added Errata: BDE72, BDE73, BDE74, BDE75, BDE76, BDE77, BDE78, BDE79, BDE80, BDE81, BDE82, BDE83, BDE84, BDE85, BDE86, BDE87, BDE88, BDE89, BDE90, BDE91, LAN5, LAN6 Erratum with title "SVM Doorbells Are Not Correctly Preserved Across Package C-States" was deleted. Table 1 Updated with V2 and Y0 stepping information	August 2016
005	Added Errata BDE71	February 2016
004	Added Errata BDE66, BDE67, BDE68, BDE69, BDE70, PCH1, PCH2, PCH3, PCH4, PCH5, PCH6, PCH7, PCH8, PCH9, PCH10, PCH11, PCH12, PCH13, PCH14, PCH15, PCH16, PCH17, PCH18, PCH19, PCH20, PCH21, PCH22, PCH23, PCH24 Updated LAN1	December 2015
003	Added Errata BDE54, BDE55, BDE56, BDE57, BDE58, BDE59, BDE60, BDE61, BDE62, BDE63, BDE64, BDE65, LAN1, LAN2, LAN3, LAN4, Modified Errata: BDE58, BDE59	November 2015
002	Added Errata BDE60,BDE61, BDE62 and BDE63	July 2015
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 D-1500 Product Family Datasheet Volume One	332050
Intel® Xeon® Processor D-1500 Product Family Datasheet Volume Two	332051
Intel® Xeon® Processor D-1500 Product Family Datasheet Volume Three	332052
Intel® Xeon® Processor D-1500 Product Family Datasheet Volume Four	332053
Intel® Xeon® Processor D-1500 Product Family Thermal Mechanical Specification Design Guide	332055
Intel® Xeon® Processor D-1500 Product Family Boundary Scan Description Language (BSDL)	332056

Nomenclature

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

Specification Changes are modifications to the current published specifications. These changes will be incorporated in 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® D-1500 Processor 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[7]
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0101b		00b	0110b	0110b	V1=0010b
	00000000b	0101b		00b	0110b	0110b	V2=0011b
	00000000b	0101b		00b	0110b	0110b	Y0=0100b
	00000000b	0101b		00b	0110b	0110b	A1[6]=0101b

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.
6. This is a D-1500N Family stepping.

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.

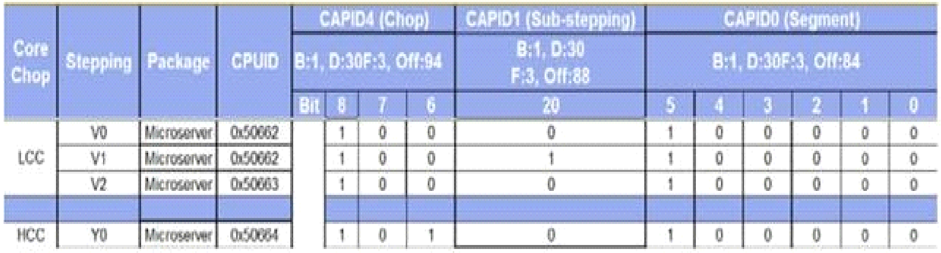

Core Chop	Stepping	Package	CPUID	CAPID4 (Chop)				CAPID1 (Sub-stepping)	CAPID0 (Segment)					
				B:1, D:30F:3, Off:94				B:1, D:30 F:3, Off:88	B:1, D:30F:3, Off:84					
				Bit	8	7	6	20	5	4	3	2	1	0
LCC	V0	Microserver	0x50662		1	0	0	0	1	0	0	0	0	0
	V1	Microserver	0x50662		1	0	0	1	1	0	0	0	0	0
	V2	Microserver	0x50663		1	0	0	0	1	0	0	0	0	0
HCC	Y0	Microserver	0x50664		1	0	1	0	1	0	0	0	0	0

Component Marking Information

Figure 1. Intel® Xeon® D-1500 Processor Families Top-Side Markings (Example)

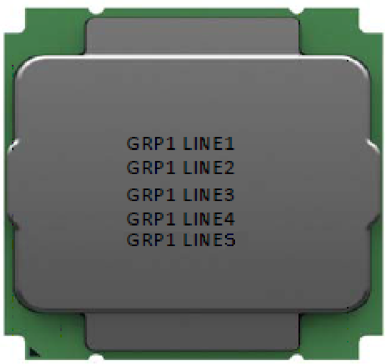

Legend	Mark Text (Production Mark)
GRP1 LINE1:	i{M}{C}YY
GRP1 LINE2:	SUB-BRAND PROC#
GRP1 LINE3:	SSPEC SPEED
GRP1 LINE4:	XXXXX
GRP1 LINE5:	{FPO} {E1}

For the Intel® Xeon® D-1500 / D-1500 NS / D-1600 NS Processor Families SKUs, see https://ark.intel.com/content/www/us/en/ark/products/series/87041/intel-xeon-d-processor.html

Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the Intel® Xeon® D-1500 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 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.

Intel® Xeon® D-1500 Product Family

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

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
BDE1	X	X	X	X	No Fix	Intel® Quick Data Technology DMA Incorrectly Indicates Support For RAID5 And RAID6
BDE2	X	X	X	X	No Fix	Intel® QuickData Technology End Point Does Not Implement Link Capabilities, Link Control, or Link Status Registers
BDE3	X	X	X	X	No Fix	PCIe* Header of a Malformed TLP is Logged Incorrectly
BDE4	X	X	X	X	No Fix	A Malformed TLP May Block ECRC Error Logging
BDE5	X	X			Fixed	Integrated COM Port Interrupt Signaling Not Compatible With The 8259 Implementation
BDE6	X	X	X	X	No Fix	Intel® QuickData Technology DMA Devices Do Not Support Uncorrectable and Correctable Error Detect Mask CSRs
BDE7	X	X	X	X	No Fix	Attempting to Enter ADR May Lead to Unpredictable System Behavior
BDE8	X	X	X		No Fix	Exiting From PC3 or PC6 With DDR4-2133 May Lead to Unpredictable System Behavior
BDE9	X				Fixed	The System May Shut Down Unexpectedly During a Warm Reset
BDE10	X	X	X	X	No Fix	CAT May Not Behave as Expected
BDE11	X	X	X	X	No Fix	Intel® QuickData Technology DMA Devices Do Not Support Uncorrectable and Correctable Error Detect Mask CSRs
BDE12	X	X	X	X	No Fix	EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change
BDE13	X	X	X	X	No Fix	MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error
BDE14	X	X	X	X	No Fix	LER MSRs May Be Unreliable
BDE15	X	X	X	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang
BDE16	X	X	X	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
BDE17	X	X	X	X	No Fix	FREEZE_WHILE_SMM Does Not Prevent Event From Pending PEBS During SMM
BDE18	X	X	X	X	No Fix	Advanced Programmable Interrupt Controller (APIC) Error "Received Illegal Vector" May be Lost
BDE19	X	X	X	X	No Fix	Performance Monitor Precise Instruction Retired Event May Present Wrong Indications
BDE20	X	X	X	X	No Fix	CR0.CD Is Ignored in VMX Operation
BDE21	X	X	X	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation
BDE22	X	X	X	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception
BDE23	X	X	X	X	No Fix	Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered
BDE24	X	X	X	X	No Fix	Pending x87 FPU Exceptions (#MF) May be Signaled Earlier Than Expected
BDE25	X	X	X	X	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction

Table 3. Integrated Core/Uncore Errata (Sheet 2 of 5)

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
BDE26	X	X	X	X	No Fix	VEX.L is Not Ignored with VCVT*2SI Instructions
BDE27	X	X	X	X	No Fix	Processor May Livelock During On Demand Clock Modulation
BDE28	X	X	X	X	No Fix	Performance Monitor Events OTHER_ASSISTS.AVX_TO_SSE And OTHER_ASSISTS.SSE_TO_AVX May Over Count
BDE29	X	X	X	X	No Fix	Performance Monitor Event DSB2MITE_SWITCHES.COUNT May Over Count
BDE30						Erratum BDE30 Removed Due to Inapplicability
BDE31	X	X	X	X	No Fix	IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report The Highest Index Value Used For VMCS Encoding
BDE32	X	X	X	X	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
BDE33	X	X	X	X	No Fix	Locked Load Performance Monitoring Events May Under Count
BDE34	X	X	X	X	No Fix	Transactional Abort May Cause an Incorrect Branch Record
BDE35	X	X	X	X	No Fix	PMI May be Signaled More Than Once For Performance Monitor Counter Overflow
BDE36	X	X	X	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
BDE37	X	X	X	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
BDE38	X	X	X	X	No Fix	A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
BDE39	X	X	X	X	No Fix	Intel® Processor Trace Packet Generation May Stop Sooner Than Expected
BDE40	X	X	X	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
BDE41	X				Fixed	Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value
BDE42						Erratum removed due to inapplicability
BDE43	X	X	X	X	No Fix	Performance Monitoring Event INSTR_RETIRED.ALL May Generate Redundant PEBS Records For an Overflow
BDE44	X	X	X	X	No Fix	Intel® Processor Trace (Intel® PT) MODE.Exec, PIP, and CBR Packets Are Not Generated as Expected
BDE45	X	X	X	X	No Fix	Performance Monitor Instructions Retired Event May Not Count Consistently
BDE46	X	X	X	X	No Fix	General-Purpose Performance Counters May be Inaccurate with Any Thread
BDE47	X	X	X	X	No Fix	An Invalid LBR May Be Recorded Following a Transactional Abort
BDE48	X	X	X	X	No Fix	Executing an RSM Instruction With Intel® Processor Trace (Intel® PT) Enabled Will Signal a #GP
BDE49	X				Fixed	Intel® Processor Trace PIP May be Unexpectedly Generated
BDE50	X	X	X	X	No Fix	Processor Core Ratio Changes While in Probe Mode May Result in a Hang
BDE51	X	X	X	X	No Fix	Processor Does Not Check IRTE Reserved Bits
BDE52	X	X	X	X	No Fix	PCIe* TPH Request Capability Structure Incorrectly Advertises Device Specific Mode as Supported
BDE53	X	X			No Fix	Package C3 State or Deeper May Lead to a Reset
BDE54	X	X	X		No Fix	Using Intel® QuickData Technology With ADR May Result in Unpredictable System Behavior
BDE55	X	X	X	X	No Fix	Intel® Advanced Vector Extensions (Intel® AVX) Workloads May Exceed ICCMAX Limits

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

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
BDE56	X	X	X		No Fix	PECI RdPkgConfig Returns Incorrect Maximum Number of Threads
BDE57	X				Fixed	DRAM tCL Improperly Limited to 14
BDE58	X	X	X	X	No Fix	Disabling UFS May Cause a System Hang
BDE59	X	X	X	X	No Fix	Writing MSR_ERROR_CONTROL May Cause a #GP
BDE60	X	X	X	X	No Fix	Isoch and Non-Isoch Traffic Concurrently Accessing The Same Cache Line May Result in System Instability
BDE61	X				Fixed	Intel® QuickPath Interconnect (Intel® QPI) Quiescence Flow And IODC Incorrectly Advertised as Supported
BDE62	X	X	X	X	No Fix	Intel® QuickData Technology Version 3.3 Source Application Tag Mask May Behave Incorrectly
BDE63	X	X	X	X	No Fix	Enabling ACC in VMX Non-Root Operation May Cause System Instability
BDE64	X	X	X	X	No Fix	Deleted. Duplicate with BDE55.
BDE65	X	X	X	X	No Fix	PAGE_WALKER_LOADS Performance Monitoring Event May Count Incorrectly
BDE66	X	X	X	X	No Fix	A spurious Patrol Scrub Error May be Logged
BDE67	X	X	X	X	No Fix	Incorrect NC-SI Flow Control Timing May Lead to BMC Packet Loss
BDE68	X	X	X	X	No Fix	Concurrent 10 GbE Port Operation May Result in Packet Loss or Reduced Bandwidth
BDE69	X				Fixed	MDIO Interface Not Available in S5 State
BDE70	X				Fixed	The Ethernet Flow Director May Incorrectly Route Fragmented Packets
BDE71			X	X	No Fix	The System May Hang When Executing a Complex Sequence of Locked Instructions
BDE72	X	X	X	X	No Fix	PCS11 May be Inaccurate When HWPM is Disabled
BDE73	X	X			Fixed	A Second Power Loss Indication May Not Initiate a Memory ADR Flow
BDE74	X	X	X	X	No Fix	Unexpected Performance Loss When Turbo Disabled
BDE75	X	X	X		No Fix	Reset During PECI Transaction May Cause a Machine Check Exception
BDE76	X	X	X	X	No Fix	Package C-state Transitions While Inband PECI Accesses Are in Progress May Cause Performance Degradation
BDE77	X	X	X	X	No Fix	Data Breakpoint Coincident With a Machine Check Exception May be Lost
BDE78	X	X	X	X	No Fix	Populating Both Memory Channels May Lead to Memory Errors
BDE79	X	X	X	X	No Fix	PECI May Not be Responsive After a Warm Reset Resulting From an IERR
BDE80	X	X	X	X	No Fix	Memory Data Scrambling is Not Compatible with CAP Error Flow
BDE81		X	X	X	No Fix	CAP Error May Cause System Hang
BDE82	X	X	X	X	No Fix	PEBS Record May Be Generated After Being Disabled
BDE83			X	X	No Fix	Software Using Intel® TSX May Behave Unpredictably
BDE84	X	X	X		No Fix	PECI RDPCICONFIGLOCAL and RDPCICONFIG Commands to Certain CSRs Always Return 0
BDE85	X	X	X	X	No Fix	MOVNTDQA From WC Memory May Pass Earlier Locked Instructions
BDE86	X	X	X	X	No Fix	Back-to-Back Page Walks Due to Instruction Fetches May Cause a System Hang

Table 3. Integrated Core/Uncore Errata (Sheet 4 of 5)

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
BDE87	X	X	X	X	No Fix	Embedded LAN Controller PCIe* AER May Not Work
BDE88	X	X	X	X	No Fix	Performance Monitoring Counters May Produce Incorrect Results for BR_INST_RETIRED Event on Logical Processor
BDE89	X	X			No Fix	Some Performance Monitor Events May Overcount During TLB Misses
BDE90	X	X	X	X	No Fix	An Intel® Hyper-Threading Technology (Intel® HT Technology) Enabled Processor May Exhibit Internal Parity Errors or Unpredictable System Behavior
BDE91			X	X	No Fix	Writing The IIO_LLC_WAYS MSR Results in an Incorrect Value
BDE92		X	X	X	No Fix	Spurious PCIe* End-to-End Parity Errors May be Logged in R2PINGERRLOG0
BDE93	X	X	X	X	No Fix	Support For Intel® QuickData Technology XOR With GF Multiply Operation is Incorrectly Advertised
BDE94		X	X	X	No Fix	Warm Reset Leads to Spurious PCIe* Malformed TLP Errors
BDE95		X	X	X	No Fix	Improper PROCHOT# Assertion When Package C-States Are Enabled
BDE96		X	X	X	No Fix	Platforms With PCIe* Hot-Plug May Hang During a Warm Reset
BDE97		X	X	X	No Fix	A #VE May Not Invalidate Cached Translation Information
BDE98		X	X	X	No Fix	Processor Instability May Occur When Using The PECI RdIAMSR Command
BDE99	X	X	X	X	No Fix	Internal Parity Errors May Incorrectly Report Overflow in the IA32_MC0_STATUS MSR
BDE100		X	X	X	No Fix	Inband PECI Concurrent With OS Patch Load May Result in Incorrect Throttling Causing Reduced System Performance
BDE101	X	X	X	X	No Fix	RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
BDE102	X	X	X	X	No Fix	Interrupt Remapping May Lead to a System Hang
BDE103	X	X	X	X	No Fix	Some DRAM And L3 Cache Performance Monitoring Events May Undercount
BDE104	X	X	X	X	No Fix	General-Purpose Performance Monitoring Counters 4-7 Will Not Increment Do Not Count With USR Mode Only Filtering
BDE105	X	X	X	X	No Fix	Writing MSR_LASTBRANCH_x_FROM_IP May #GP When Intel® TSX is Not Supported
BDE106	X	X	X	X	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
BDE107	X	X	X	X	No Fix	JTAG Boundary Scan For Intel QPI And PCIe* Lanes May Report Incorrect Stuck at 1 Errors
BDE108	X	X	X	X	No Fix	Bi-Directional PCIe* Posted Transactions May Lead to System Hang
BDE109		X	X	X	No Fix	DDR Chip Select Signal May Toggle After Entry to Self Refresh
BDE110	X	X	X	X	No Fix	APIC Timer Interrupt May Not be Generated at The Correct Time In TSC-Deadline Mode
BDE111	X	X	X	X	No Fix	LAN Controller Does Not Send PM_PME Message After Wake Event From D3Hot
BDE112	X	X	X	X	No Fix	Link Down Events Behind PCIe Device Connected to CPU Root Ports Can Cause CTO > 50ms on Other Root Ports
BDE113	X	X	X	X	No Fix	NVDIMM Data May Not be Preserved Correctly on Power Loss or ADR Activation
BDE114	X	X	X	X	No Fix	Debug Exceptions May Be Lost in The Case Of Machine Check Exception

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

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
BDE115	X	X	X	X	No Fix	The System May Hang When Exiting package C6 State
BDE116	X	X	X	X	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
BDE117				X	No Fix	Integrated 10 Gigabit Ethernet Controllers May Not be Visible in PCIe Configuration Space
BDE118	x	x	x	x	No Fix	MSR_TURBO_ACTIVATION_RATIO MSR Cannot be Locked
BDE119	x	x	x	x	No Fix	Package C3 or C6 Transition May Lead to a Machine Check Exception Logging Uncorrected DDR ECC Errors
BDE120	x	x	x	x	No Fix	Intel® MBM Counters May Report System Memory Bandwidth Incorrectly
BDE121	x	x	x	x	No Fix	Performance Monitoring General Counter 2 May Have Invalid Value Written When Transactional Synchronization Extensions (TSX) Is Enabled

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	Clarification on Operation of LT_LOCK_MEMORY MSR.

Integrated Core/Uncore Errata

BDE1 Intel® Quick Data Technology DMA Incorrectly Indicates Support For RAID5 And RAID6

Problem: The processor incorrectly advertises support for Galois Field multiply with a '1' in bit 9 of DMACAPABILITY (CBBAR{0-3}, Offset 10H).

Implication: RAID5 and RAID6 operations require descriptors using Galois Field multiply operations (OpTypes in the range of the 089H-08BH). Using these OpTypes will signal CHANERR.DCERR (bit 10) and will halt the corresponding channel.

Workaround: None identified.

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

BDE2 Intel® QuickData Technology End Point Does Not Implement Link Capabilities, Link Control, or Link Status Registers

Problem: The PCI Express* Base Specification, revision 3.1, requires a PCIe* endpoint to implement the Link Capabilities, Link Control, and Link Status registers (at PCI Express Capability Structure offsets 0CH, 10H, and 12H, respectively). However, the processor's Intel QuickData Technology DMA end points (Device 0; Function 0, 1, 2, 3) do not implement these registers. The Intel QuickData Technology DMA endpoints' PCI bus number is specified in CPUBUSNO2[15:8] (Bus 1; Device 16; Function 7; Offset 0D4H).

Implication: Software that expects these registers to be implemented by the end points may not behave as expected

Workaround: None identified.

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

BDE3 PCIe* Header of a Malformed TLP is Logged Incorrectly

Problem: If a PCIe port receives a malformed Transaction Layer Packet (TLP), 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*.

BDE4 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*.

BDE5 Integrated COM Port Interrupt Signaling Not Compatible With The 8259 Implementation

Problem: The Integrated COM ports generate active-low interrupts. In PIC or Virtual Wire mode, the 8259 implementation can’t sense active-low interrupts.

Implication: Any software using PIC or Virtual Wire mode is unable to service interrupts from the Integrated COM ports.

Workaround: Use symmetric I/O mode when Integrated COM port interrupts are necessary.

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

BDE6 Intel® QuickData Technology DMA Devices Do Not Support Uncorrectable and Correctable Error Detect Mask CSRs

Problem: Intel QuickData Technology DMA devices include ERRUNCDETMSK and ERRCORDETMSK CSRs (Device 0; Function 0-3; offsets 140H and 144H respectively). Due to this erratum, there is no error masking functionality associated with these CSRs.

Implication: Writing these CSRs does not affect the operation of the processor.

Workaround: None identified.

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

BDE7 Attempting to Enter ADR May Lead to Unpredictable System Behavior

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

BDE8 Exiting From PC3 or PC6 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*.

BDE9 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*.

BDE10 CAT May Not Behave as Expected

Problem: Due to this erratum, Cache Allocation Technology (CAT) 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*.

BDE11 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 Last Branch Record (LBR), Branch Trace Store (BTS) and Branch Trace Message (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 (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*.

BDE12 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*.

BDE13 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*.

BDE14 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*.

BDE15 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*.

BDE16 #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*.

BDE17 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*.

BDE18 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*.

BDE19 Performance Monitor Precise Instruction Retired Event May Present Wrong Indications

Problem: When the Precise Distribution for Instructions Retired (PDIR) 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*.

BDE20 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.*

BDE21 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 (e.g., PDE); then invalidate any translations for the affected linear addresses; and then modify the relevant paging-structure entry to set the P flag and establish the new page size. 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.*

BDE22 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.*

BDE23 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 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 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.*

BDE24 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.*

BDE25 DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction

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

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

Workaround: None identified.

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

BDE26 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.*

BDE27 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.*

BDE28 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.*

BDE29 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.*

BDE30 Erratum BDE30 Removed Due to Inapplicability

BDE31 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.*

BDE32 Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

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

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

Workaround: None identified.

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

BDE33 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.*

BDE34 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 Last Branch Record (LBR) 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.*

BDE35 PMI May be Signaled More Than Once For Performance Monitor Counter Overflow

Problem: Due to this erratum, Performance Monitoring Interrupt (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.*

BDE36 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.*

BDE37 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.*

BDE38 A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation

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

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

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

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

BDE39 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.*

BDE40 PEBS Eventing IP Field May be Incorrect After Not-Taken Branch

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

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

Workaround: None identified.

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

BDE41 Reading The Memory Destination of an Instruction That Begins an HLE Transaction May Return The Original Value

Problem: An Hardware Lock Elision (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.*

BDE42 Erratum BDE42 Removed Due to Inapplicability

BDE43 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.*

BDE44 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), Paging Information Packet (PIP), and Core: Bus Ratio (CBR) 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.*

BDE45 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.*

BDE46 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.*

BDE47 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 Last Branch Record (LBR) 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.*

BDE48 Executing an RSM Instruction With Intel® Processor Trace (Intel® PT) Enabled Will Signal a #GP

Problem: Upon delivery of an System Management Interrupt (SMI), 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.*

BDE49 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.*

BDE50 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.*

BDE51 Processor Does Not Check IRTE Reserved Bits

Problem: As per the Intel® Virtualization Technology for Directed I/O (Intel® VT-d) specification, bits 63:Host Address Width (HAW) of the Posted Interrupt Descriptor Upper Address field in the (Interrupt Remapping Table Entry (IRTE) 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.*

BDE52 PCIe* TPH Request Capability Structure Incorrectly Advertises Device Specific Mode as Supported

Problem: The TPH Transaction layer packet Processing Hints (TPH) 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.*

BDE53 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.

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

BDE54 Using Intel® QuickData Technology With ADR May Result in Unpredictable System Behavior

Problem: Using Intel QuickData with Asynchronous DRAM Refresh (ADR) may result in unpredictable system behavior.

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.

BDE55 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.

BDE56 PECI RdPkgConfig Returns Incorrect Maximum Number of Threads

Problem: PECI RdPkgConfig command with index 0 may not return the correct maximum thread ID in parameter 3.

Implication: Software that depends on maximum thread ID 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.

BDE57 DRAM tCL Improperly Limited to 14

Problem: The processor incorrectly limits the DRAM tCL parameter to 14.

Implication: DRAM operating at 2133 MT/s and higher that requires tCL greater than 14 may not operate correctly.

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

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

BDE58 Disabling UFS May Cause a System Hang

Problem: Due to this erratum, disabling Uncore Frequency Scaling (UFS) via the BIOS Mailbox WRITE_PCU_MISC_CONFIG (0x06) command with UFS Disable (bit 28) set in the MISC_BIT_VECTOR also disables periodic RCOMP.

Implication: The UFS disable side effect of disabling periodic RCOMP will lead to unreliable memory subsystem operation resulting in a system hang.

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

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

BDE59 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.

BDE60 Isoch and Non-Isoch Traffic Concurrently Accessing The Same Cache Line May Result in System Instability

Problem: Isoch traffic and non-Isoch traffic concurrently accessing the same cache line may lead to system instability.

Implication: When this erratum occurs, the system may hang or signal a machine check.

Workaround: None identified. The memory accessed by Isoch traffic and non-Isoch traffic must not overlap.

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

BDE61 Intel® QuickPath Interconnect (Intel® QPI) Quiescence Flow And IODC Incorrectly Advertised as Supported

Problem: The Power Control Unit CAPID3 (Bus 1; Device 30; Function 3; Offset 90H) bit 11 incorrectly advertises IODC (I/O Directory Cache) support. Further, Intel QPI quiescence flow is incorrectly advertised as supported by VIRTUAL_MSR_CR_QUIESCE_CTL1 (MSR 50H) bit 7.

Implication: Due to this erratum, software may invoke the Intel QPI quiescence flow and/or enable IODC. Either action may lead to unpredictable system behavior.

Workaround: None identified. Software should not invoke the Intel QPI quiescence flow and should not enable IODC.

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

BDE62 Intel® QuickData Technology Version 3.3 Source Application Tag Mask May Behave Incorrectly

Problem: Upon receipt of a Data Integrity Field (DIF) update with strip descriptor containing a non-zero source application tag mask, a spurious DIF error may be logged.

Implication: Intel QuickData Technology may incorrectly log a DIF error and halt operations.

Workaround: For a software workaround, contact your Intel representative.

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

BDE63 Enabling ACC in VMX Non-Root Operation May Cause System Instability

Problem: Autonomous C-State Control (ACC) 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.

BDE64 Deleted. Duplicate with BDE55.

BDE65 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.

BDE66 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.

BDE67 Incorrect NC-SI Flow Control Timing May Lead to BMC Packet Loss

Problem: As described in the Network Controller Sideband Interface (NC-SI) Specification (version 1.0.1), the processor should periodically transmit an XOFF flow control frame so long as an underlying traffic congestion condition persists. Due to this erratum, the interval between XOFF message transmissions by the processor exceeds the specification’s limit.

Implication: Exceeding the XOFF transmission interval may cause the platform’s Baseboard Management Controller (BMC) to incorrectly conclude that the traffic congestion condition has been resolved. The BMC then may resume sending packets, creating the possibility of packet loss.

Workaround: None identified. BMC software must retry lost packets.

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

BDE68 Concurrent 10 GbE Port Operation May Result in Packet Loss or Reduced Bandwidth

Problem: Due to this erratum, the processor is not able to sustain concurrent, high bandwidth operation of both 10 GbE ports using small packets (about 112 bytes or less).

Implication: If flow control is enabled, this erratum may result in reduced 10 GbE port bandwidth. If flow control is not enabled, this erratum may result in packet loss. This erratum may be observed during the L3 forwarding performance test under RFC2544 using the Intel Data Plan Development Kit. This erratum has not been observed when exactly one 10 GbE port is active.

Workaround: None identified.

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

BDE69 MDIO Interface Not Available in S5 State

Problem: The Management Data I/O (MDIO) interface is not available in S5 state after a G3 to S5 transition.

Implication: Without MDIO, firmware cannot enable certain I/O interfaces resulting in lack of support for 10 GbE connection in S5 state and no Wake on LAN support after G3 to S5 transition.

Workaround: It is possible for BIOS to contain a workaround to enable Wake on LAN, utilizing LAN NVM image in WW12 or later LEK releases. There is no workaround available for 10 GbE; the link will be limited to 1 Gb/S transfer speeds.

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

BDE70 The Ethernet Flow Director May Incorrectly Route Fragmented Packets

Problem: The Flow Director in the processor’s Ethernet controller should route packets to the optimal core based on the packet’s outer header. Due to this erratum, for fragmented packets, the Flow Director may route the packets to the default queue.

Implication: The performance gains expected from Flow Director may not be realized for fragmented packets. There are no functional implications.

Workaround: None identified.

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

BDE71 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*.

BDE72 PCS11 May be Inaccurate When HWPM is Disabled

Problem: Package Configuration Space 11 (PCS11) accumulates time during which Running Average Power Limit (RAPL) determines the processor frequency. If the processor cannot reduce its power sufficiently to meet the RAPL-defined power limit, it should still update PCS11. Under these conditions, if Hardware Power Management (HWPM) is disabled, the processor will not update PCS11

Implication: When HWPM is disabled, PCS11 may not be accurate.

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

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

BDE73 A Second Power Loss Indication May Not Initiate a Memory ADR Flow

Problem: In response to a power loss indication, the processor will initiate an Asynchronous DRAM Refresh (ADR) flow. Upon power restoration, an ADR resume event will be initiated. Due to this erratum, if the platform signals an additional power loss event prior to completion of the ADR resume flow, an additional ADR flow is not initiated.

Implication: The occurrence of this erratum can result in unpredictable system behavior.

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

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

BDE74 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.

BDE75 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.

BDE76 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 1.2GHz 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.

BDE77 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.

BDE78 Populating Both Memory Channels May Lead to Memory Errors

Problem: When both memory channels are populated with SODIMMs correctable and uncorrectable memory errors may occur.

Implication: When this erratum occurs, a higher than expected rate of correctable and uncorrectable memory errors will be logged.

Workaround: A BIOS workaround has been identified. Please refer to memory reference code version 18R00 or later and release notes.

BDE79 PECI May Not be Responsive After a Warm Reset Resulting From an IERR

Problem: Due to this erratum, Platform Environmental Control Interface (PECI) may become non-responsive and always return a 0x91 completion code after a warm reset that follows a Power Controller Unit (PCU) fatal machine check. The 0x91 completion code indicates that a hardware, firmware, or associated logic error blocked processing of the PECI request.

Implication: When this erratum occurs, using PECI to configure the processor or diagnose processor errors is not feasible. In some cases, a cold reset may be needed to regain functionality.

Workaround: None identified.

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

BDE80 Memory Data Scrambling is Not Compatible with CAP Error Flow

Problem: Command/Address Parity (CAP) error flow and DDR4 WrCRC error flow for DIMM isolation fail when data scrambling is enabled.

Implication: Error isolation software may not behave as expected when memory data scrambling is enabled.

Workaround: None identified. Do not enable memory data scrambling in the presence of error isolation software.

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

BDE81 CAP Error May Cause System Hang

Problem: A Command/Address Parity (CAP) error may result in a system hang that logs an Internal Timeout Error Machine Check (IA32_MCi_STATUS.MCACOD=0400H and IA32_MCi_STATUS.MSCOD=0080H).

Implication: When this erratum occurs, a correctable CAP error results in an uncorrectable 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.

BDE82 PEBS Record May Be Generated After Being Disabled

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

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

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

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

BDE83 Software Using Intel® TSX May Behave Unpredictably

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

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

BDE84 PECI RDPCICONFIGLOCAL and RDPCICONFIG Commands to Certain CSRs Always Return 0

Problem: Due to this erratum, the Platform Environment Control Interface (PECI) RDPCICFGLOCAL and RDPCICFG commands attempting to access the following CSRs always return 0:

Due to this erratum, the PECI RDPCICFGLOCAL and RDPCICFG commands attempting to access the following CSRs always return 0:

Table 4. CSRs affected

Bud	Device	Function	Offset
0xFF	8 to 31	(all)	0x00 to 0x0F
0xFF	(all)	(all)	0x08 to 0x0B
0xFF	8 to 31	(all)	0x30 to 0x3F

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.

BDE85 MOVNTDQA From WC Memory May Pass Earlier Locked Instructions

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

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

Workaround: None identified. Software that relies on a locked instruction to fence subsequent executions of (V)MOVNTDQA should insert an MFENCE instruction between the locked instruction and subsequent (V)MOVNTDQA instruction.

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

BDE86 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.

BDE87 Embedded LAN Controller PCIe* AER May Not Work

Problem: When Virtualization is enabled on the embedded LAN controller's PF1 (Physical Function 1) and not on PF0, PCIe Advanced Error Reporting (AER) for Virtual Function (VF) does not work.

Implication: This erratum results in a PCI Express Base Specification compliance violation.

Workaround: None identified.

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

BDE88 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: V1, V2, A1: None identified; Y0: It is possible for the BIOS to contain a workaround for this erratum.

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

BDE89 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.*

BDE90 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 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*.

BDE91 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*.

BDE92 Spurious PCIe* End-to-End Parity Errors May be Logged in R2PINGERRLOG0

Problem: A warm reset after PCIe End-to-End parity has been enabled by setting HaSysDefeature3.EnableE2EPAR (Bus 1; Device 18; Function 0; Offset 9CH; bit 2) to 1 may result in the logging of spurious parity errors in R2PINGERRLOG0.ParErrE2E1 and/or R2PINGERRLOG0.ParErrE2E0 (Bus 1; Device 16; Function 0; Offset 4CH; bits 13 and 12, respectively).

Implication: When this erratum occurs, spurious PCIe End-to-End parity errors are logged.

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*.

BDE93 Support For Intel® QuickData Technology XOR With GF Multiply Operation is Incorrectly Advertised

Problem: Intel QuickData Technology DMACAPABILITY register bit 9 is asserted, incorrectly indicating support for the XOR with GF Multiply operation (OpTypes 089H - 08BH). This erratum applies to channels 0 and 1.

Implication: Due to this erratum, software reading bit 9 of the DMACAPABILITY register may incorrectly expect Intel QuickData Technology channels 0 and 1 to support RAID 5 and RAID 6.

Workaround: Software can avoid this erratum by ignoring Intel QuickData Technology DMACAPABILITY register bit 9.

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

BDE94 Warm Reset Leads to Spurious PCIe* Malformed TLP Errors

Problem: A warm reset after the normal completion of BIOS results in PCIe Malformed TLP errors being logged with IIO global fatal error status bits for the Intel® QuickData Technology and GBE root ports.

Implication: When this erratum occurs, a fatal error is reported after warm reset. Intel has not observed this erratum to impact the operation of any commercially available software.

Workaround: None identified.

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

BDE95 Improper PROCHOT# Assertion When Package C-States Are Enabled

Problem: PROCHOT# may be improperly asserted in certain conditions when Package C-States are enabled, resulting in unexpected throttling.

Implication: Unexpected throttling 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*.

BDE96 Platforms With PCIe* Hot-Plug May Hang During a Warm Reset

Problem: During a warm reset, the processor may incorrectly operate the PCIe Hot-Plug SMBUS I/O Expander leading to a system hang.

Implication: When this erratum occurs, the system will hang during a warm reset.

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

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

BDE97 A #VE May Not Invalidate Cached Translation Information

Problem: An Extended Page Table (EPT) violation that causes a #Virtualization Exception (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*.

BDE98 Processor Instability May Occur When Using The PECI RdIAMSR Command

Problem: Under certain circumstances, reading a machine check register using the Platform Environmental Control Interface (PECI) 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*.

BDE99 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*.

BDE100 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*.

BDE101 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 (Processor Event Based Sampling) or BTS (Branch Trace Store) address translation, the RF (resume flag) 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*.

BDE102 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: It is possible for the BIOS to contain a workaround for this erratum.

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

BDE103 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*.

BDE104 General-Purpose Performance Monitoring Counters 4-7 Will Not Increment 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_PERFEVTSELx 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*.

BDE105 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® Transactional Synchronization Extensions (Intel® TSX) (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*.

BDE106 Performance Monitoring Counters May Undercount When Using CPL Filtering

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

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

Workaround: None identified.

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

BDE107 JTAG Boundary Scan For Intel QPI And PCIe* Lanes May Report Incorrect Stuck at 1 Errors

Problem: Boundary Scan testing of the Intel QPI and PCIe interfaces may incorrectly report a recurring stuck at 1 failure on Intel QPI and PCIe receiver lanes. This erratum only affects Boundary Scan testing and does not affect functional operation of the Intel QPI and PCIe interfaces.

Implication: This erratum may result in Boundary Scan test failures reported on one or more of the Intel QPI and PCIe lanes.

Workaround: None identified.

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

BDE108 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: 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*.

BDE109 DDR Chip Select Signal May Toggle After Entry to Self Refresh

Problem: The processor may assert the DDR Chip Select signal can assert immediately after the Self Refresh Entry DDR command is issued during Package C6 entry.

Implication: When this erratum occurs, the JEDEC DDR4 t_CPDED timing specification is violated possibly leading to uncorrectable machine check errors and/or the logging of patrol scrub errors.

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*.

BDE110 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 Advanced Programmable Interrupt Controller (APIC) timer is configured for TSC-Deadline mode, a timer interrupt may be generated much earlier than expected or much later 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*.

BDE111 LAN Controller Does Not Send PM_PME Message After Wake Event From D3Hot

Problem: When the LAN controller is in D3Hot and detects a wake event, it wakes the system but does not send a PM_PME message to the root port.

Implication: Software may not be able to determine the LAN controller caused the wake event from D3Hot.

Workaround: None Identified.

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

BDE112 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 >50 ms 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 50 ms) 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*.

BDE113 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, memory data may be incorrect and may lead to an ECC error on next access.

Workaround: None Identified.

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

BDE114 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 Processor Context Corrupted (PCC) 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: For the Steppings affected, see the *Summary Tables of Changes*.

BDE115 The System May Hang When Exiting package C6 State

Problem: Under complex micro-architectural 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.

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

BDE116 Using Intel® TSX Instructions May Lead to Unpredictable System Behavior

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

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

Status: It is possible for BIOS to contain a workaround for this erratum. For the Steppings affected, see the *Summary Tables of Changes*.

BDE117 Integrated 10 Gigabit Ethernet Controllers May Not be Visible in PCIe Configuration Space

Problem: After reset events that do not send a RESET_WARN, the integrated 10 Gigabit Ethernet controllers may not be visible in PCIe configuration space. This only occurs when one or both of the LAN controllers are configured to be disabled in S5.

Implication: Due to this erratum, the 10 Gigabit Ethernet controllers may not be visible to software after a reset.

Workaround: None identified.

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

BDE118 MSR_TURBO_ACTIVATION_RATIO MSR Cannot be Locked

Problem: Setting the TURBO_ACTIVATION_RATIO_LOCK field (bit 31) of the MSR_TURBO_ACTIVATION_RATIO MSR (64H) has no effect; it does not block future writes to the MSR_TURBO_ACTIVATION_RATIO MSR.

Implication: Software cannot rely on locking the MSR_TURBO_ACTIVATION_RATIO MSR.

Workaround: None identified.

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

BDE119 Package C3 or C6 Transition May Lead to a Machine Check Exception Logging Uncorrected DDR ECC Errors

Problem: With certain processor MCUs (Microcode Updates M1050662_00000019, M1050663_07000016, M1050664_0F000014, M1050665_0E00000C), a Package C3 or Package C6 state transition may lead to a machine check exception with IA32_MC9_STATUS.MSCOD=0x0010 or IA32_MC10_STATUS.MSCOD=0x0010 (Patrol Scrub Error) or IA32_MC7_STATUS.MSCOD=0x0001 (Read Data Error) with Uncorrected Error also set.

Implication: A system hang may occur due to an uncorrected machine check exception.

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

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

BDE120 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*.

BDE121 Performance Monitoring General Counter 2 May Have Invalid Value Written When Intel® Transactional Synchronization Extensions (Intel® TSX) Is Enabled

Problem: When a Intel® TSX is enabled and there are aborts (Hardware Lock Elision [HLE] or Restricted Transactional Memory [RTM]) overlapping with access or manipulation of the IA32_PMC2 general-purpose performance counter (Offset: C3h) it may return invalid value.

Implication: Implication: Software may read invalid value from IA32_PMC2.

Workaround: Workaround: None identified

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

LAN Specific Sightings

Table 5. LAN Errata

Number	Steppings				Status	ERRATA
	V1	V2	Y0	A1		
LAN1	X				Fixed	SR-IOV is Not Supported on The Integrated 10 GbE LAN Controller
LAN2	X	X	X	X	No Fix	Excessive CRC Error Rate May be Observed on 10GBase-KR Links
LAN3	X	X	X	X	No Fix	Incorrect Link Fault Errors May be Logged at Power-On Reset
LAN4	X	X	X	X	No Fix	A CRC Error May Occur Upon Link Reset in Ethernet 1000Base-KX Mode
LAN5	X	X	X	X	No Fix	Embedded LAN Device May Report Incorrect PCIe* Completion Status When Virtual Function is Enabled
LAN6	X	X	X	X	No Fix	Enabling LPLU May Prevent Establishing a Link in Low Power System States

LAN1 SR-IOV is Not Supported on The Integrated 10 GbE LAN Controller

Problem: When Single Root IO Virtualization (SR-IOV) is enabled, a Virtual Function Level Reset (VFLR) or Bus Master Enable (BME) clear on a Virtual Function (VF) of one 10GbE Physical Function (PF) also affects the corresponding VF on the other PF of the same device.

Implication: A VF may appear to reset or hang unexpectedly.

Workaround: None identified. Either utilize VF on only one PF or assign VFs from both PFs such that they do not overlap.

Status: Fix.

LAN2 Excessive CRC Error Rate May be Observed on 10GBase-KR Links

Problem: After a 10GBase-KR link has been successfully established, data transfers may experience a CRC error rate above the IEEE 802.3 specification limit.

Implication: Intel has observed this erratum to occur approximately once in every 30 local PHY resets or Auto-Negotiate restarts initiated by the link partner. This erratum does not affect the reliability of the link because packets with CRC errors are automatically retransmitted.

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.

LAN3 Incorrect Link Fault Errors May be Logged at Power-On Reset

Problem: Due to this erratum, during a power-on reset, LAN link errors may be incorrectly logged in the LINKS register bit 3 and the HLREG1 register bits (7:8).

Implication: Error reporting may indicate false local fault and idle errors.

Workaround: None identified. The Link Status register and the MAC local fault Count register should be ignored before the link is established.

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

LAN4 A CRC Error May Occur Upon Link Reset in Ethernet 1000Base-KX Mode

Problem: In 1000Base-KX mode (on 10GBASE-KR PHY) at reset, a malformed packet may be generated leading to the link partner's receiver logging a CRC error.

Implication: Link partner's receiver could log a CRC error upon link reset in 1000Base-KX mode. After the first CRC error, no further CRC errors will be observed due to this erratum. Intel has not observed this erratum to impact the functionality of any commercially available system.

Workaround: None identified.

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

LAN5 Embedded LAN Device May Report Incorrect PCIe* Completion Status When Virtual Function is Enabled

Problem: When an embedded LAN device receives a PCIe Unsupported Request (UR) or INVALID completion status targeted to a Virtual Function (VF), it is expected that the Received_Master_Abort bit (PCISTS, offset 6, bit 13) will be set in the VF. Due to this erratum, the Received_Master_Abort bit is set in the PCISTS CSR associated with the Physical Function (PF) instead.

Implication: When this erratum occurs, the associated device driver will see an incorrect error report.

Workaround: None identified.

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

LAN6 Enabling LPLU May Prevent Establishing a Link in Low Power System States

Problem: In a low power state (S3, S4, or S5) and with GbE Low Power Link Up (LPLU) enabled, the processor will attempt to establish a LAN link at the 1G LAN link speed regardless of link partner capabilities.

Implication: If the link partner does not support the 1G LAN link speed and LPLU is enabled, a link will not be established when the system is in a low power state. Wake on LAN and other network traffic will not function until the system exits the low power state.

Workaround: Software can avoid this erratum by enabling LPLU only when the link partner is known to support the 1G LAN link speed.

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

Integrated PCH Sightings

The following table provides a list of PCH issues and the steppings affected

Table 6. Integrated PCH Specific Sightings

Number	Steppings				Status	Errata
	V1	V2	Y0	A1		
PCH1	X	X	X	X	No Fix	USB Isoch In Transfer Error Issue
PCH2	X	X	X	X	No Fix	USB Babble Detected with SW Over-Scheduling
PCH3	X	X	X	X	No Fix	USB Full-/Low-Speed EOP Issue
PCH4	X	X	X	X	No Fix	Asynchronous Retries Prioritized Over Periodic Transfers
PCH5	X	X	X	X	No Fix	USB FS/LS Incorrect Number of Retries
PCH6	X	X	X	X	No Fix	USB Full-/Low-speed Port Reset or Clear TT Buffer Request
PCH7	X	X	X	X	No Fix	USB RMH Think Time Issue
PCH8	X	X	X	X	No Fix	USB Full-/low-speed Device Removal Issue
PCH9	X	x	X	X	No Fix	xHC Data Packet Header and Payload Mismatch Error Condition
PCH10	X	X	X	X	No Fix	USB SuperSpeed Packet with Invalid Type Field Issue
PCH11	X	X	X	X	No Fix	xHC Behavior with Three Consecutive Failed U3 Entry Attempts
PCH12	X	X	X	X	No Fix	Incorrect IRQ(x) Vector Returned for 8259 Interrupts With RAEOI Enabled
PCH13	X	x	X	X	No Fix	Max Packet Size and Transfer Descriptor Length Mismatch
PCH14	X	X	X	X	No Fix	PCIe* Root Ports Unsupported Request Completion
PCH15	X	X	X	X	No Fix	SATA Signal Voltage Level Violation
PCH16	X	X	X	X	No Fix	Super-speed Device Re-Enumeration
PCH17	X	x	X	X	No Fix	Set Latency Tolerance Value Command Completion Event Issue
PCH18	X	X	X	X	No Fix	LFPS Detect Threshold
PCH19	X	X	X	X	No Fix	SMBus Hold Time
PCH20	X	X	X	X	No Fix	RMH Port Disabled Due to Device Initiated Remote Wake
PCH21	X	x	X	X	No Fix	Enumeration Issue when Resuming for Sx
PCH22	X	X	X	X	No Fix	SATA Lock Lost with During Link Negotiation
PCH23	X	X	X	X	No Fix	PCIe* Clocking Mode Switch Issue
PCH24	X	X	X	X	No Fix	USB xHCI may Execute a Stale Transfer Request Block (TRB)
PCH25	X	X	X	X	No Fix	xHCI Host Controller Reset May Cause a System Hang
PCH26	X				Fix	Integrated COM Ports Baud Rates Are Not Generated as Expected
PCH27	X	X	X	X	No Fix	xHCI Short Packet Event Using Non-Event Data TRB

PCH1 USB Isoch In Transfer Error Issue

Problem: If a USB full-speed inbound isochronous transaction with a packet length 190 bytes or greater is started near the end of a micro-frame the PCH may see more than 189 bytes in the next micro-frame. Implication: If the PCH sees more than 189 bytes for a micro-frame an error will be sent to software and the isochronous transfer will be lost. If a single data packet is lost no perceptible impact for the end user is expected.

Note: Intel has only observed the issue in a synthetic test environment where precise control of packet scheduling is available, and has not observed this failure in its compatibility validation testing.

- Isochronous traffic is periodic and cannot be retried thus it is considered good practice for software to schedule isochronous transactions to start at the beginning of a micro-frame. Known software solutions follow this practice.
- •To sensitize the system to the issue additional traffic such as other isochronous transactions or retries of asynchronous transactions would be required to push the inbound isochronous transaction to the end of the micro-frame.

Workaround: None

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

PCH2 USB Babble Detected with SW Over-Scheduling

Problem: If software violates USB periodic scheduling rules for full-speed isochronous traffic by over-scheduling, the RMH may not handle the error condition properly and return a completion split with more data than the length expected.

Implication: If the RMH returns more data than expected, the endpoint will detect packet babble for that transaction and the packet will be dropped. Since over-scheduling occurred to create the error condition, the packet would be dropped regardless of RMH behavior. If a single isochronous data packet is lost, no perceptible impact to the end user is expected.

Note: USB software over-scheduling occurs when the amount of data scheduled for a micro-frame exceeds the maximum budget. This is an error condition that violates the USB periodic scheduling rule.

Note: This failure has only been recreated synthetically with USB software intentionally over-scheduling traffic to hit the error condition.

Workaround: None

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

PCH3 USB Full-/Low-Speed EOP Issue

Problem: If the EOP of the last packet in a USB Isochronous split transaction (Transaction >189 bytes) is dropped or delayed 3 ms or longer the following may occur:

- If there are no other pending low-speed or full-speed transactions the RMH will not send SOF, or Keep-Alive. Devices connected to the RMH will interpret this condition as idle and will enter suspend.
- If there is other pending low-speed or full-speed transactions, the RMH will drop the isochronous transaction and resume normal operation.

Implication: If there are no other transactions pending, the RMH is unaware a device entered suspend and may starting sending a transaction without waking the device. The implication is device dependent, but a device may stall and require a reset to resume functionality.

If there are other transactions present, only the initial isochronous transaction may be lost. The loss of a single isochronous transaction may not result in end user perceptible impact.

Note: Intel has only observed this failure when using software that does not comply with the USB specification and violates the hardware isochronous scheduling threshold by terminating transactions that are already in progress

Workaround: None

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

PCH4 Asynchronous Retries Prioritized Over Periodic Transfers

Problem: The integrated USB RMH incorrectly prioritizes full-speed and low-speed asynchronous retries over dispatchable periodic transfers.

Implication: Periodic transfers may be delayed or aborted. If the asynchronous retry latency causes the periodic transfer to be aborted, the impact varies depending on the nature of periodic transfer:

- If a periodic interrupt transfer is aborted, the data may be recovered by the next instance of the interrupt or the data could be dropped.
- If a periodic isochronous transfer is aborted, the data will be dropped. A single dropped periodic transaction should not be noticeable by end user.

Note: This issue has only been seen in a synthetic environment. The USB spec does not consider the occasional loss of periodic traffic a violation.

Workaround: None

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

PCH5 USB FS/LS Incorrect Number of Retries

Problem: A USB low-speed transaction may be retried more than three times, and a USB full speed transaction may be retried less than three times if all of the following conditions are met:

- A USB low-speed transaction with errors, or the first retry of the transaction occurs near the end of a micro-frame, and there is not enough time to complete another retry of the low-speed transaction in the same micro-frame.
- There is pending USB full-speed traffic and there is enough time left in the micro-frame to complete one or more attempts of the full-speed transaction.
- Both the low-speed and full-speed transactions must be asynchronous (Bulk/Control) and must have the same direction either in or out.

Note: Per the USB EHCI Specification a transaction with errors should be attempted a maximum of 3 times if it continues to fail.

Implication: For low-speed transactions the extra retry(s) allow a transaction additional chance(s) to recover regardless of if the full-speed transaction has errors or not.

If the full-speed transactions also have errors, the PCH may retry the transaction fewer times than required, stalling the device prematurely. Once stalled, the implication is software dependent, but the device may be reset by software.

Workaround: None

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

PCH6 USB Full-/Low-speed Port Reset or Clear TT Buffer Request

Problem: One or more full-/low-speed USB devices on the same RMH controller may be affected if the devices are not suspended and either (a) software issues a Port Reset OR (b) software issues a Clear TT Buffer request to a port executing a split full-/low-speed Asynchronous Out command.

- The small window of exposure for full-speed device is around 1.5 microseconds and around 12 microseconds for a low-speed device.

Implication: The affected port may stall or receive stale data for a newly arrived split transfer.

Implication: occurring at the time of the Port Reset or Clear TT Buffer request.

Note: This issue has only been observed in a synthetic test environment.

Workaround: None

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

PCH7 USB RMH Think Time Issue

Problem: The USB RMH Think Time may exceed its declared value in the RMH hub descriptor register of 8 full-speed bit times.

Implication: If the USB driver fully subscribes a USB micro-frame, LS/FS transactions may exceed the micro-frame boundary.

Note: No functional failures have been observed.

Workaround: None

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

PCH8 USB Full-/low-speed Device Removal Issue

Problem: If two or more USB full-/low-speed devices are connected to the same USB controller, the devices are not suspended, and one device is removed, one or more of the devices remaining in the system may be affected by the disconnect.

Implication: The implication is device dependent. A device may experience a delayed transaction, stall and be recovered via software, or stall and require a reset such as a hot plug to resume normal functionality.

Workaround: None

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

PCH9 xHC Data Packet Header and Payload Mismatch Error Condition

Problem: If a Super Speed device sends a Data Packet Header (DPH) to the xHC with a data length field that specifies less data than is actually sent in the Data Packet Payload (DPP), the xHC will accept the packet instead of discarding the packet as invalid.

Note: The USB 3.0 specification requires a device to send a DPP matching the amount of data specified by the DPH.

Implication: The amount of data specified in the DPH will be accepted by the xHC and the remaining data will be discarded and may result in anomalous system behavior.

Note: This issue has only been observed in a synthetic test environment with a synthetic device.

Workaround: None

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

PCH10 USB SuperSpeed Packet with Invalid Type Field Issue

Problem: If the encoding for the "type" field for a SuperSpeed packet is set to a reserved value and the encoding for the "subtype" field is set to "ACK", the xHC may accept the packet as a valid acknowledgment transaction packet instead of ignoring the packet.

Note: The USB 3.0 specification requires that a device never set any defined fields to reserved values.

Implication: System implication is dependent on the misbehaving device and may result in anomalous system behavior.

Note: This issue has only been observed in a synthetic test environment with a synthetic device.

Workaround: None

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

PCH11 xHC Behavior with Three Consecutive Failed U3 Entry Attempts

Problem: The xHC does not transition to the SS.Inactive USB 3.0 Link Training and Status State Machine (LTSSM) state after a SuperSpeed device fails to enter U3 upon three consecutive attempts.

Note: The USB 3.0 specification requires a SuperSpeed device to enter U3 when directed.

Implication: The xHC will continue to try to initiate U3. The implication is driver and operating system dependent.

Workaround: None.

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

PCH12 Incorrect IRQ(x) Vector Returned for 8259 Interrupts With RAEOI Enabled

Problem: If multiple interrupts are active prior to an interrupt acknowledge cycle with Rotating Automatic End of Interrupt (RAEOI) mode of operation enabled for 8259 interrupts (0-7), an incorrect IRQ(x) vector may be returned to the processor.

Implication: Implications of an incorrect IRQ(x) vector being returned to the CPU are SW implementation dependent.

Note: This issue has only been observed in a synthetic test environment.

Workaround: None.

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

PCH13 Max Packet Size and Transfer Descriptor Length Mismatch

Problem: The xHC may incorrectly handle a request from a low-speed or full-speed device when all the following conditions are true:

- The sum of the packet fragments equals the length specified by the Transfer Descriptor (TD)
- The TD length is less than the Max Packet Size (MPS) for the device
- The last packet received in the transfer is “0” or babble bytes

Implication: The xHC will halt the endpoint if all the above conditions are met. All functions associated with the endpoint will stop functioning until the device is unplugged and reinserted.

Workaround: None.

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

PCH14 PCIe* Root Ports Unsupported Request Completion

Problem: The PCIe* root ports may return an Unsupported Request (UR) completion with an incorrect lower address field in response to a memory read if any of the following occur:

- Bus Master Enable is disabled in the PCIe Root Port’s Command register (PCICMD bit 2 =0)
- Address Type (AT) field of the Transaction Layer Packet (TLP) header is non-zero
- The requested upstream address falls within the memory range claimed by the secondary side of the bridge
- Requester ID with Bus Number of 0

Implication: The UR Completion with an incorrect lower address field may be handled as a Malformed TLP causing the Requester to send an ERR_NONFATAL or ERR_FATAL message upstream to the root port.

Workaround: None.

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

PCH15 SATA Signal Voltage Level Violation

Problem: SATA transmit buffers have been designed to maximize performance and robustness over a variety of routing scenarios. As a result, the SATA transmit signaling voltage levels may exceed the maximum motherboard TX connector and device RX connector voltage specifications as defined in section 7.2.2.3 of the Serial ATA specification, rev 3.1. This issue applies to Gen 1 (1.5 Gb/s).

Implication: None known.

Workaround: None.

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

PCH16 Super-speed Device Re-Enumeration

Problem: If a Super-speed device is connected to the xHC and an unexpected device pulse occurs on the USB3R{n,p} signals during an exit from U3 low power link state, the xHC may falsely detect a connection event.

Implication: The Super-speed device may re-enumerate when resuming from U3. Implications of re-enumeration are driver, application and operating system dependent.

Note: A Super-speed device may enter the U3 low power link state during S3 or selective suspend. There are no known cases of data loss since the Super-speed device always re-enumerates.

Workaround: None.

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

PCH17 Set Latency Tolerance Value Command Completion Event Issue

Problem: The xHCI controller does not return a value of ‘0’ for slot ID in the command completion event Transfer Request Block (TRB) for a set latency tolerance value command.

Note: This violates the command completion event TRB description in section 6.4.2.2 of the eXtensible Host Controller Interface for Universal Serial Bus (xHCI) specification, revision 1.0.

Implication: There are no known functional failures due to this issue.

Note: Set latency tolerance value command is specific to the controller and not the slot. Software knows which command was issued and which fields are valid to check for the event.

Note: xHCI CV compliance test suite: Test TD4.10: Set Latency Tolerance Value Command Test may issue a warning. No wavier from USBIF required for this warning.

Workaround: None.

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

PCH18 LFPS Detect Threshold

Problem: The xHC Low Frequency Periodic Signal (LFPS) detect threshold of 400 mV is higher than the USB 3.0 specification maximum of 300 mV.

Implication: The xHC may not recognize LFPS from Super-speed devices transmitting at the minimum low power peak-to-peak differential voltage (400 mV) as defined by USB 3.0 specification.

Note: The low power peak-to-peak voltage transmission level is intended for devices soldered down to the motherboard.

Workaround: None.

Note: For optimal interoperability across all implementations, Intel recommends that designs utilize soldered down Super-speed devices that support standard peak-to-peak differential voltage levels (800 mV minimum).

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

PCH19 SMBus Hold Time

Problem: The SMBus data hold time may be less than the 300 ns minimum defined by the Intel® Xeon® Processor D-1500 Product Family External Design Specification (EDS).

Implication: There are no known functional failures due to this issue.

Workaround: None.

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

PCH20 RMH Port Disabled Due to Device Initiated Remote Wake

Problem: During resume from Global Suspend, the RMH controller may not send SOF soon enough to prevent a device from entering suspend again. A collision on the port may occur if a device initiated remote wake occurs before the RMH controller sends SOF.

Note: Intel has only observed this issue when two USB devices on the same RMH controller send remote wake within 30 ms window while RMH controller is resuming from Global Suspend.

Implication: The RMH host controller may detect the collision as babble and disable the port.

Workaround: Intel recommends system software to check bit 3 (Port Enable/Disable Change) together with bit 7 (Suspend) of Port N Status and Control PORTC registers when determining which port(s) have initiated remote wake.

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

PCH21 Enumeration Issue when Resuming for Sx

Problem: If a device is attached while the platform is in S3 or S4 and the device is assigned the highest assignable Slot ID upon resume, the xHC may attempt to access an unassigned main memory address.

Implication: Accessing unassigned main memory address may cause a system software timeout leading to possible system hang.

Workaround: System Software can detect the timeout and perform a host controller reset to avoid the system hang.

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

PCH22 SATA Lock Lost with During Link Negotiation

Problem: During link speed negotiation, if a receiver error occurs after host SATA controller locks on a device's ALIGN primitive, the host SATA controller may be unable to train the link.

Note: This issue only occurs when SSC is disabled on the drive and has only been observed at SATA Gen2 speeds.

Implication: A SATA device connected to the SATA controller may fail to train and become inoperative.

Workaround: A workaround is available in Reference Code.

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

PCH23 PCIe* Clocking Mode Switch Issue

Problem: The PCIe link may become unstable when switching from non-common clock mode to common clock mode with some PCIe devices.

Implication: The PCIe link may report link errors or train to a lower speed. Implication is device dependent.

Workaround: A workaround is available in Reference Code.

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

PCH24 USB xHCI may Execute a Stale Transfer Request Block (TRB)

Problem: When a USB 3.0 or USB 2.0 hub with numerous active Full-Speed (FS) or Low-Speed (LS) periodic endpoints attached is removed and then reconnected to an USB xHCI port, the xHCI controller may fail to fully refresh its cache of TRB records. The controller may read and execute a stale TRB and place a pointer to it in a Transfer Event TRB.

Note: In some cases, the xHCI controller may read de-allocated memory pointed to by a TRB of a disabled slot. The xHCI controller may also place a pointer to that memory in the event ring, causing the xHCI driver to access that memory and process its contents, resulting in system hang, failure to enumerate devices, or other anomalous system behavior.

Note: This issue has only been observed in a stress test environment

Workaround: None

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

PCH25 xHCI Host Controller Reset May Cause a System Hang

Problem: Within 1 ms of setting the Host Controller Reset bit (HCRST bit 1) of the USB Command Register (xHCI BAR + 80h) the xHCI host controller may fail to respond to register accesses.

Implication: The system may hang.

Workaround: None identified.

Note: Note: Software must not make any accesses to the xHCI Host Controller registers for 1 ms after setting the HCRST bit 1 of the USB Command Register (xHCI BAR + 80h) and must add a 120 ms delay in between consecutive xHCI host controller resets.

Status: No Fix.

PCH26 Integrated COM Ports Baud Rates Are Not Generated as Expected

Problem: Due to this erratum, the Integrated COM port UARTs use a reference clock of 2.4576 MHz rather than the expected 1.8432 MHz.

Implication: Divisor ratios used to generate UART baud rate clocks from the reference clock must take into account the actual reference clock frequency. Further, while typical baud rates can be precisely generated with the 2.4576 MHz reference clock (e.g., 38400, 19200, 9600, 4800, 2400, and 1200), the 115200, 57600, and 28800 baud rates cannot be generated and the 14400 baud rate will have a tolerable (-3%) frequency error.

Workaround: It is possible for BIOS to contain processor configuration data and code changes as a workaround for this erratum. Software that does not use BIOS to configure the Integrated COM ports will need to modify its divisor ratios appropriately.

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

PCH27 xHCI Short Packet Event Using Non-Event Data TRB

Problem: The xHCI may generate an unexpected short packet event for the last transfer's Transfer Request Block (TRB) when using Non-Event Data TRB with multiples TRBs.

Implication: Transfer may fail due to the packet size error.

Implication: Note: This issue has only been observed in an synthetic environment. No known implication has been identified with commercial software.

Workaround: None identified. Intel recommends software to use Data Event TRBs for short packet completion.

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

Documentation Changes

1. Clarification on Operation of LT_LOCK_MEMORY MSR.

On V2 and Y0 stepping processors, once BIOS sets bit 0 (CONFIG_LOCK) in LT_LOCK_MEMORY (MSR 2E7h) the following additional registers are locked and can not be changed:

RCBA - Root Complex Base Address (Bus 0; Device 31; Function 0; Offset F0h)

GPIOBASE GPIO - Base Address (Bus 0; Device 31; Function0; Offset 48h)

FD - Function Disable (RCBA + 3418h)

GP_IO_SEL (GPIOBASE+04h) Regardless of setting of GPIO Lockdown Enable in GPIO Control.