

# Intel<sup>®</sup> Xeon<sup>®</sup> Processor 3400 Series

Specification Update

---

*November 2009*



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining, critical control or safety systems, or in nuclear facility applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The processor 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. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

<sup>A</sup>Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See [http://www.intel.com/products/processor\\_number](http://www.intel.com/products/processor_number) for details. Over time processor numbers will increment based on changes in clock, speed, cache, or other features, and increments are not intended to represent proportional or quantitative increases in any particular feature. Current roadmap processor number progression is not necessarily representative of future roadmaps. See [www.intel.com/products/processor\\_number](http://www.intel.com/products/processor_number) for details.

Intel® Active Management Technology requires the computer system to have an Intel(R) AMT-enabled chipset, network hardware and software, as well as connection with a power source and a corporate network connection. Setup requires configuration by the purchaser and may require scripting with the management console or further integration into existing security frameworks to enable certain functionality. It may also require modifications of implementation of new business processes. With regard to notebooks, Intel AMT may not be available or certain capabilities may be limited over a host OS-based VPN or when connecting wirelessly, on battery power, sleeping, hibernating or powered off. For more information, see [www.intel.com/technology/platform-technology/intel-amt/](http://www.intel.com/technology/platform-technology/intel-amt/)

Intel® Trusted Execution Technology (Intel® TXT) requires a computer system with Intel® Virtualization Technology (Intel® Virtualization Technology (Intel® VT-x) and Intel® Virtualization Technology for Directed I/O (Intel® VT-d)), a Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules and an Intel TXT-compatible measured launched environment (MLE). The MLE could consist of a virtual machine monitor, an OS or an application. In addition, Intel TXT requires the system to contain a TPM v1.2, as defined by the Trusted Computing Group and specific software for some uses. For more information, see <http://www.intel.com/technology/security>

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain computer system software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

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

Intel® Hyper-threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology-enabled chipset, BIOS, and operating system. Performance will vary depending on the specific hardware and software you use. For more information including details on which processors support HT Technology, see <http://www.intel.com/info/hyperthreading>.


64-bit computing on Intel architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers and applications enabled for Intel® 64 architecture. Performance will vary depending on your hardware and software configurations. Consult with your system vendor for more information.

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

Intel, Intel Core, Celeron, Pentium, Intel Xeon, Intel Atom, Intel SpeedStep, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2009, Intel Corporation. All Rights Reserved.



## Contents

---

<b>Revision History</b> .....	5
<b>Preface</b> .....	6
<b>Summary Tables of Changes</b> .....	8
<b>Identification Information</b> .....	16
<b>Errata</b> .....	19
<b>Specification Changes</b> .....	53
<b>Specification Clarifications</b> .....	54
<b>Documentation Changes</b> .....	55

§





# Revision History

---

Revision	Description	Date
-001	Initial Release	September 2009
-002	Added Errata AAO101-AAO109.	October 2009
-003	Updated the Processor Identification Table to include two additional SKUs: <ul style="list-style-type: none"><li>• Intel® Xeon® Processor X3430 (S-Spec Number: SLBLJ).</li><li>• Intel® Xeon® Processor L3426 (S-Spec Number: SLBN3).</li></ul> Added Errata AAO110- AAO113.	November 2009



## Preface

---

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

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


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

### Affected Documents

Document Title	Document Number
<i>Intel® Xeon® Processor 3400 Series Datasheet - Volume 1</i>	322371-001
<i>Intel® Xeon® Processor 3400 Series Datasheet - Volume 2</i>	322372-001

### Related Documents

Document Title	Document Number/ Location
<i>AP-485, Intel® Processor Identification and the CPUID Instruction</i>	<a href="http://www.intel.com/design/processor/aplnots/241618.htm">http://www.intel.com/design/processor/aplnots/241618.htm</a>
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide</i> <i>Intel® 64 and IA-32 Intel Architecture Optimization Reference Manual</i>	<a href="http://www.intel.com/products/processor/manuals/index.htm">http://www.intel.com/products/processor/manuals/index.htm</a>
<i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes</i>	<a href="http://www.intel.com/design/processor/specupdt/252046.htm">http://www.intel.com/design/processor/specupdt/252046.htm</a>
<i>ACPI Specifications</i>	<a href="http://www.acpi.info">www.acpi.info</a>



## Nomenclature

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


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

**Specification 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, etc.).



# Summary Tables of Changes

---

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

## Codes Used in Summary Tables

### Stepping

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


### Page


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

### Status

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

### Row

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



Each Specification Update item is prefixed with a capital letter to distinguish the product. The key below details the letters that are used in Intel's microprocessor Specification Updates:

- A = Intel® Xeon® processor 7000 sequence
- C = Intel® Celeron® processor
- D = Intel® Xeon® processor 2.80 GHz
- E = Intel® Pentium® III processor
- F = Intel® Pentium® processor Extreme Edition and Intel® Pentium® D processor
- I = Intel® Xeon® processor 5000 series
- J = 64-bit Intel® Xeon® processor MP with 1MB L2 cache
- K = Mobile Intel® Pentium® III processor
- L = Intel® Celeron® D processor
- M = Mobile Intel® Celeron® processor
- N = Intel® Pentium® 4 processor
- O = Intel® Xeon® processor MP
- P = Intel® Xeon® processor
- Q = Mobile Intel® Pentium® 4 processor supporting Intel® Hyper-Threading technology on 90-nm process technology
- R = Intel® Pentium® 4 processor on 90 nm process
- S = 64-bit Intel® Xeon® processor with 800 MHz system bus (1 MB and 2 MB L2 cache versions)
- T = Mobile Intel® Pentium® 4 processor-M
- U = 64-bit Intel® Xeon® processor MP with up to 8MB L3 cache
- V = Mobile Intel® Celeron® processor on .13 micron process in Micro-FCPGA package
- W = Intel® Celeron® M processor
- X = Intel® Pentium® M processor on 90nm process with 2-MB L2 cache and Intel® processor A100 and A110 with 512-KB L2 cache
- Y = Intel® Pentium® M processor
- Z = Mobile Intel® Pentium® 4 processor with 533 MHz system bus
- AA = Intel® Pentium® D processor 900 sequence and Intel® Pentium® processor Extreme Edition 955, 965
- AB = Intel® Pentium® 4 processor 6x1 sequence
- AC = Intel® Celeron® processor in 478 pin package
- AD = Intel® Celeron® D processor on 65nm process
- AE = Intel® Core™ Duo processor and Intel® Core™ Solo processor on 65nm process
- AF = Intel® Xeon® processor LV
- AG = Intel® Xeon® processor 5100 series
- AH = Intel® Core™2 Duo/Solo Processor for Intel® Centrino® Duo Processor Technology
- AI = Intel® Core™2 Extreme processor X6800 and Intel® Core™2 Duo desktop processor E6000 and E4000 sequence



- AJ = Intel® Xeon® processor 5300 series
- AK = Intel® Core™2 Extreme quad-core processor QX6000 sequence and Intel® Core™2 Quad processor Q6000 sequence
- AL = Intel® Xeon® processor 7100 series
- AM = Intel® Celeron® processor 400 sequence
- AN = Intel® Pentium® dual-core processor
- AO = Intel® Xeon® processor 3200 series
- AP = Intel® Xeon® processor 3000 series
- AQ = Intel® Pentium® dual-core desktop processor E2000 sequence
- AR = Intel® Celeron® processor 500 series
- AS = Intel® Xeon® processor 7200, 7300 series
- AU = Intel® Celeron® Dual Core processor T1400
- AV = Intel® Core™2 Extreme processor QX9650 and Intel® Core™2 Quad processor Q9000 series
- AW = Intel® Core™ 2 Duo processor E8000 series
- AX = Intel® Xeon® processor 5400 series
- AY = Intel® Xeon® processor 5200 series
- AZ = Intel® Core™2 Duo processor and Intel® Core™2 Extreme processor on 45-nm process
- AAA = Intel® Xeon® processor 3300 series
- AAB = Intel® Xeon® E3110 processor
- AAC = Intel® Celeron® dual-core processor E1000 series
- AAD = Intel® Core™2 Extreme processor QX9775
- AAE = Intel® Atom™ processor Z5xx series
- AAF = Intel® Atom™ processor 200 series
- AAG = Intel® Atom™ processor N series
- AAH = Intel® Atom™ processor 300 series
- AAI = Intel® Xeon® processor 7400 series
- AAJ = Intel® Core™ i7 processor and Intel® Core™ i7 Extreme Edition processor
- AAK = Intel® Xeon® processor 5500 series
- AAL = Intel® Pentium Dual-Core processor E5000 series
- AAN = Intel Core i7-800 and i5-700 desktop processor series
- AAO = Intel Xeon processor 3400 series
- AAP = Intel® Core i7-900 Mobile Processor Extreme Edition Series, Intel Core i7-800 and i7-700 Mobile Processor Series




## Errata (Sheet 1 of 5)

Number	Steppings	Status	ERRATA
	B-1		
AAO1	X	No Fix	The Processor May Report a #TS Instead of a #GP Fault
AAO2	X	No Fix	REP MOVS/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may use an Incorrect Data Size or Lead to Memory-Ordering Violations
AAO3	X	No Fix	Code Segment Limit/Canonical Faults on RSM May be Serviced before Higher Priority Interrupts/Exceptions and May Push the Wrong Address Onto the Stack
AAO4	X	No Fix	Performance Monitor SSE Retired Instructions May Return Incorrect Values
AAO5	X	No Fix	Premature Execution of a Load Operation Prior to Exception Handler Invocation
AAO6	X	No Fix	MOV To/From Debug Registers Causes Debug Exception
AAO7	X	No Fix	Incorrect Address Computed For Last Byte of FXSAVE/FXRSTOR Image Leads to Partial Memory Update
AAO8	X	No Fix	Values for LBR/BTS/BTM will be Incorrect after an Exit from SMM
AAO9	X	No Fix	Single Step Interrupts with Floating Point Exception Pending May Be Mishandled
AAO10	X	No Fix	Fault on ENTER Instruction May Result in Unexpected Values on Stack Frame
AAO11	X	No Fix	IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception
AAO12	X	No Fix	General Protection Fault (#GP) for Instructions Greater than 15 Bytes May be Preempted
AAO13	X	No Fix	General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit
AAO14	X	No Fix	LBR, BTS, BTM May Report a Wrong Address when an Exception/Interrupt Occurs in 64-bit Mode
AAO15	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang
AAO16	X	No Fix	Corruption of CS Segment Register During RSM While Transitioning From Real Mode to Protected Mode
AAO17	X	No Fix	Performance Monitoring Events for Read Miss to Level 3 Cache Fill Occupancy Counter may be Incorrect
AAO18	X	No Fix	A VM Exit on MWAIT May Incorrectly Report the Monitoring Hardware as Armed
AAO19	X	No Fix	Delivery Status of the LINT0 Register of the Local Vector Table May be Lost
AAO20	X	No Fix	Performance Monitor Event SEGMENT_REG_LOADS Counts Inaccurately
AAO21	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
AAO22	X	No Fix	Improper Parity Error Signaled in the IQ Following Reset When a Code Breakpoint is Set on a #GP Instruction
AAO23	X	No Fix	An Enabled Debug Breakpoint or Single Step Trap May Be Taken after MOV SS/POP SS Instruction if it is Followed by an Instruction That Signals a Floating Point Exception
AAO24	X	No Fix	IA32_MPERF Counter Stops Counting During On-Demand TM1
AAO25	X	No Fix	The Memory Controller tTHROT_OPREF Timings May be Violated During Self Refresh Entry
AAO26	X	No Fix	Processor May Over Count Correctable Cache MESI State Errors



## Errata (Sheet 2 of 5)

Number	Steppings	Status	ERRATA
	B-1		
AAO27	X	No Fix	Synchronous Reset of IA32_APERF/IA32_MPERF Counters on Overflow Does Not Work
AAO28	X	No Fix	Disabling Thermal Monitor While Processor is Hot, Then Re-enabling, May Result in Stuck Core Operating Ratio
AAO29	X	No Fix	PECI Does Not Support PCI Configuration Reads/Writes to Misaligned Addresses
AAO30	X	No Fix	OVER Bit for IA32_MCi_STATUS Register May Get Set on Specific Internal Error
AAO31	X	No Fix	Writing the Local Vector Table (LVT) when an Interrupt is Pending May Cause an Unexpected Interrupt
AAO32	X	No Fix	Faulting MMX Instruction May Incorrectly Update x87 FPU Tag Word
AAO33	X	No Fix	xAPIC Timer May Decrement Too Quickly Following an Automatic Reload While in Periodic Mode
AAO34	X	No Fix	Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures
AAO35	X	No Fix	B0-B3 Bits in DR6 For Non-Enabled Breakpoints May be Incorrectly Set
AAO36	X	No Fix	Core C6 May Clear Previously Logged TLB Errors
AAO37	X	No Fix	Performance Monitor Event MISALIGN_MEM_REF May Over Count
AAO38	X	No Fix	Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations
AAO39	X	No Fix	Running with Write Major Mode Disabled May Lead to a System Hang
AAO40	X	No Fix	Infinite Stream of Interrupts May Occur if an ExtINT Delivery Mode Interrupt is Received while All Cores in C6
AAO41	X	No Fix	Two xAPIC Timer Event Interrupts May Unexpectedly Occur
AAO42	X	No Fix	EOI Transaction May Not be Sent if Software Enters Core C6 During an Interrupt Service Routine
AAO43	X	No Fix	FREEZE_WHILE_SMM Does Not Prevent Event From Pending PEBS During SMM
AAO44	X	No Fix	APIC Error "Received Illegal Vector" May be Lost
AAO45	X	No Fix	DR6 May Contain Incorrect Information When the First Instruction After a MOV SS,r/m or POP SS is a Store
AAO46	X	No Fix	An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang
AAO47	X	No Fix	IA32_PERF_GLOBAL_CTRL MSR May be Incorrectly Initialized
AAO48	X	No Fix	ECC Errors Can Not be Injected on Back-to-Back Writes
AAO49	X	No Fix	Performance Monitor Interrupts Generated From Uncore Fixed Counters (394H) May be Ignored
AAO50	X	No Fix	Performance Monitor Counter INST_RETIRED.STORES May Count Higher than Expected
AAO51	X	No Fix	Sleeping Cores May Not be Woken Up on Logical Cluster Mode Broadcast IPI Using Destination Field Instead of Shorthand
AAO52	X	No Fix	Faulting Executions of FXRSTOR May Update State Inconsistently
AAO53	X	No Fix	Performance Monitor Event EPT.EPDPE_MISS May be Counted While EPT is Disable
AAO54	X	No Fix	Memory Aliasing of Code Pages May Cause Unpredictable System Behavior



## Errata (Sheet 3 of 5)

Number	Steppings	Status	ERRATA
	B-1		
AA055	X	No Fix	Performance Monitor Counters May Count Incorrectly
AA056	X	No Fix	Processor Forward Progress Mechanism Interacting With Certain MSR/CSR Writes May Cause Unpredictable System Behavior
AA057	X	No Fix	Performance Monitor Event Offcore_response_0 (B7H) Does Not Count NT Stores to Local DRAM Correctly
AA058	X	No Fix	EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change
AA059	X	No Fix	System May Hang if MC_CHANNEL_{0,1}_MC_DIMM_INIT_CMD.DO_ZQCL Commands Are Not Issued in Increasing Populated DDR3 Rank Order
AA060	X	No Fix	Package C3/C6 Transitions When Memory 2x Refresh is Enabled May Result in a System Hang
AA061	X	No Fix	Back to Back Uncorrected Machine Check Errors May Overwrite IA32_MC3_STATUS.MSCOD
AA062	X	No Fix	Memory Intensive Workloads with Core C6 Transitions May Cause System Hang
AA063	X	No Fix	Corrected Errors With a Yellow Error Indication May be Overwritten by Other Corrected Errors
AA064	X	No Fix	PSI# Signal May Incorrectly be Left Asserted
AA065	X	No Fix	Memory ECC Errors May be Observed When a UC Partial Write is Followed by a UC Read to the Same Location
AA066	X	No Fix	Performance Monitor Events DCACHE_CACHE_LD and DCACHE_CACHE_ST May Overcount
AA067	X	No Fix	Rapid Core C3/C6 Transitions May Cause Unpredictable System Behavior
AA068	X	No Fix	Performance Monitor Events INSTR_RETIRED and MEM_INST_RETIRED May Count Inaccurately
AA069	X	No Fix	A Page Fault May Not be Generated When the PS bit is set to "1" in a PML4E or PDPTe
AA070	X	No Fix	CPURESET Bit Does Not Get Cleared
AA071	X	No Fix	PHOLD Disable in MISCCTRLSTS Register Does Not Work
AA072	X	No Fix	PCIe PMCSR Power State Field Incorrectly Allows Requesting of the D1 and D2 Power States
AA073	X	No Fix	PECI Accesses to Registers May Fail When Processor is Transitioning to/from Package C6 Power State
AA074	X	No Fix	Concurrent Updates to a Segment Descriptor May be Lost
AA075	X	No Fix	PMIs May be Lost During Core C6 Transitions
AA076	X	No Fix	Uncacheable Access to a Monitored Address Range May Prevent Future Triggering of the Monitor Hardware
AA077	X	No Fix	BIST Results May be Additionally Reported After a GETSEC[WAKEUP] or INIT-SIPI Sequence
AA078	X	No Fix	Pending x87 FPU Exceptions (#MF) May be Signaled Earlier Than Expected
AA079	X	No Fix	VM Exits Due to "NMI-Window Exiting" May Be Delayed by One Instruction
AA080	X	No Fix	Malformed PCIe Packet Generated Under Heavy Outbound Load
AA081	X	No Fix	PCIe Operation in x16 Mode With Inbound Posted Writes May be Unreliable



## Errata (Sheet 4 of 5)

Number	Steppings	Status	ERRATA
	B-1		
AAO82	X	No Fix	Unpredictable PCI Behavior Accessing Non-existent Memory Space
AAO83	X	No Fix	PECI MbxGet() Commands May Fail Several Times Before Passing When Issued During Package C6
AAO84	X	No Fix	VM Exits Due to EPT Violations Do Not Record Information About Pre-IRET NMI Blocking
AAO85	X	No Fix	Intel® VT-d Receiving Two Identical Interrupt Requests May Corrupt Attributes of Remapped Interrupt or Hang a Subsequent Interrupt-Remap-Cache Invalidation Command
AAO86	X	No Fix	S1 Entry May Cause Cores to Exit C3 or C6 C-State
AAO87	X	No Fix	Multiple Performance Monitor Interrupts are Possible on Overflow of IA32_FIXED_CTR2
AAO88	X	No Fix	LBRs May Not be Initialized During Power-On Reset of the Processor
AAO89	X	No Fix	Unexpected Interrupts May Occur on C6 Exit If Using APIC Timer to Generate Interrupts
AAO90	X	No Fix	Package C6 Exit with Memory in Self-Refresh When Using DDR3 RDIMM Memory May Lead to a System Hang
AAO91	X	No Fix	LBR, BTM or BTS Records May have Incorrect Branch From Information After an EIST Transition, T-states, C1E, or Adaptive Thermal Throttling
AAO92	X	No Fix	PECI GetTemp() Reads May Return Invalid Temperature Data in Package C6 State
AAO93	X	No Fix	PECI PCIConfigRd() Followed by a GetTemp() May Cause System Hang in Package C6 State
AAO94	X	No Fix	PECI Mailbox Commands During Package C6 Idle State Transitions May Result in Unpredictable Processor Behavior
AAO95	X	No Fix	VMX-Preemption Timer Does Not Count Down at the Rate Specified
AAO96	X	No Fix	Multiple Performance Monitor Interrupts are Possible on Overflow of Fixed Counter 0
AAO97	X	No Fix	SVID and SID of Devices 8 and 16 only implement bits [7:0]
AAO98	X	No Fix	No_Soft_Reset Bit in the PMCSR Does Not Operate as Expected
AAO99	X	No Fix	VM Exits Due to LIDT/LGDT/SIDT/SGDT Do Not Report Operand Size
AAO100	X	No Fix	PCIConfigRd() and PCIConfigWr() PECI Commands May Silently Fail During Package C6 Exit Events
AAO101	X	No Fix	Performance Monitoring Events STORE_BLOCKS.NOT_STA and STORE_BLOCKS.STA May Not Count Events Correctly
AAO102	X	No Fix	Storage of PEBS Record Delayed Following Execution of MOV SS or STI
AAO103	X	No Fix	Performance Monitoring Event FP_MMX_TRANS_TO_MMX May Not Count Some Transitions
AAO104	X	No Fix	INVLPG Following INVEPT or INVVPID May Fail to Flush All Translations for a Large Page
AAO105	X	No Fix	The PECI Bus May be Tri-stated After System Reset
AAO106	X	No Fix	LER MSRs May Be Unreliable
AAO107	X	No Fix	Multiple ECC Errors May Result in Incorrect Syndrome Being Logged
AAO108	X	No Fix	MCI_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error



## Errata (Sheet 5 of 5)

Number	Steppings	Status	ERRATA
	B-1		
AAO109	X	No Fix	Debug Exception Flags DR6.B0-B3 Flags May be Incorrect for Disabled Breakpoints
AAO110	X	No Fix	An Exit From the Core C6-state May Result in the Dropping of an Interrupt
AAO111	X	No Fix	PCIe Extended Capability Structures May be Incorrect
AAO112	X	No Fix	PMIs During Core C6 Transitions May Cause the System to Hang
AAO113	X	No Fix	IA32_MC8_CTL2 MSR is Not Cleared on Processor Warm Reset

## Specification Changes

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

## Specification Clarifications

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



# Identification Information

## Component Identification via Programming Interface

The Intel Xeon processor 3400 series stepping can be identified by the following register contents:

Reserved	Extended Family <sup>1</sup>	Extended Model <sup>2</sup>	Reserved	Processor Type <sup>3</sup>	Family Code <sup>4</sup>	Model Number <sup>5</sup>	Stepping ID <sup>6</sup>
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0001b		00b	0110	1110b	xxxxb

**Note:**

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 Pro, 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 Processor Type, specified in bits [13:12] indicates whether the processor is an original OEM processor, an OverDrive processor, or a dual processor (capable of being used in a dual processor system).
4. 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.
5. 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.
6. The Stepping ID in bits [3:0] indicates the revision number of that model. See [Table 1](#) for the processor stepping ID number in the CPUID information.

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

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

The Intel Xeon processor 3400 series can be identified by the following register contents:

Stepping	Vendor ID <sup>1</sup>	Device ID <sup>2</sup>	Revision ID <sup>3</sup>
B-1	8086h	D130h	11h

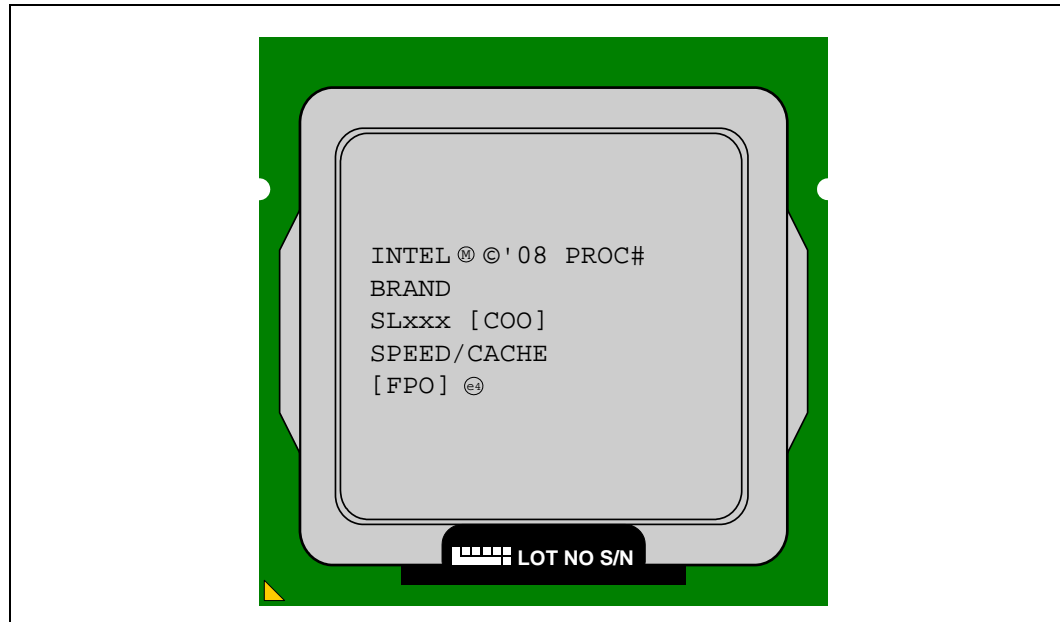
**Notes:**

1. The Vendor ID corresponds to bits 15:0 of the Vendor ID Register located at offset 00–01h in the PCI function 0 configuration space.
2. The Device ID corresponds to bits 15:0 of the Device ID Register located at Device 0 offset 02–03h in the PCI function 0 configuration space.
3. The Revision Number corresponds to bits 7:0 of the Revision ID Register located at offset 08h in the PCI function 0 configuration space.

## Component Marking Information

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

**Figure 1. Processor Production Top-side Markings (Example)**



**Table 1. Processor Identification (Sheet 1 of 2)**

S-Spec Number	Processor Number	Stepping	Processor Signature	Core Frequency (GHz) / DDR3 (MHz)	Max Intel® Turbo Boost Technology Frequency (GHz) <sup>2</sup>	Shared L3 Cache Size (MB)	Notes
SLBJH	X3470	B-1	106E5h	2.93 / 1333	4 core: 3.20 3 core: 3.20 2 core: 3.46 1 core: 3.60	8	1, 3, 4, 5, 6
SLBJK	X3460	B-1	106E5h	2.80 / 1333	4 core: 2.93 3 core: 2.93 2 core: 3.33 1 core: 3.46	8	1, 3, 4, 5, 6
SLBLD	X3450	B-1	106E5h	2.66 / 1333	4 core: 2.80 3 core: 2.80 2 core: 3.20 1 core: 3.20	8	1, 3, 4, 5, 6
SLBLF	X3440	B-1	106E5h	2.53 / 1333	4 core: 2.66 3 core: 2.66 2 core: 2.80 1 core: 2.93	8	1, 3, 4, 5, 6





Table 1. Processor Identification (Sheet 2 of 2)

S-Spec Number	Processor Number	Stepping	Processor Signature	Core Frequency (GHz) / DDR3 (MHz)	Max Intel® Turbo Boost Technology Frequency (GHz) <sup>2</sup>	Shared L3 Cache Size (MB)	Notes
SLBLJ	X3430	B-1	106E5h	2.40 / 1333	4 core: 2.53 3 core: 2.53 2 core: 2.66 1 core: 2.80	8	1, 4, 5, 6
SLBN3	L3426	B-1	106E5h	1.86 / 1333	4 core: 2.13 3 core: 2.13 2 core: 3.06 1 core: 3.20	8	3, 4, 5, 6, 7

**Notes:**

1. This processor has TDP of 95W and meets the 1156\_VR\_CONF\_09B VR Configuration.
2. This column indicates maximum Intel® Turbo Boost Technology frequency (GHz) for 4, 3, 2, or 1 cores active respectively.
3. Intel® Hyper-Threading Technology enabled.
4. Intel® Trusted Execution Technology (Intel® TXT) enabled.
5. Intel® Virtualization Technology for IA-32, Intel® 64 and Intel® Architecture (Intel® VT-x) enabled.
6. Intel® Virtualization Technology for Directed I/O (Intel® VT-d) enabled.
7. This processor has TDP of 45 W.



# Errata

---

## **AAO1. The Processor May Report a #TS Instead of a #GP Fault**

**Problem:** A jump to a busy TSS (Task-State Segment) may cause a #TS (invalid TSS exception) instead of a #GP fault (general protection exception).

**Implication:** Operation systems that access a busy TSS may get invalid TSS fault instead of a #GP fault. 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.

## **AAO2. REP MOVS/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may use an Incorrect Data Size or Lead to Memory-Ordering Violations**

**Problem:** Under certain conditions as described in the Software Developers Manual section "Out-of-Order Stores For String Operations in Pentium 4, Intel Xeon, and P6 Family Processors" the processor performs REP MOVS or REP STOS as fast strings. Due to this erratum fast string REP MOVS/REP STOS instructions that cross page boundaries from WB/WC memory types to UC/WP/WT memory types, may start using an incorrect data size or may observe memory ordering violations.

**Implication:** Upon crossing the page boundary the following may occur, dependent on the new page memory type:

- UC the data size of each write will now always be 8 bytes, as opposed to the original data size.
- WP the data size of each write will now always be 8 bytes, as opposed to the original data size and there may be a memory ordering violation.
- WT there may be a memory ordering violation.

**Workaround:** Software should avoid crossing page boundaries from WB or WC memory type to UC, WP or WT memory type within a single REP MOVS or REP STOS instruction that will execute with fast strings enabled.

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



**AAO3. Code Segment Limit/Canonical Faults on RSM May be Serviced before Higher Priority Interrupts/Exceptions and May Push the Wrong Address Onto the Stack**

**Problem:** Normally, when the processor encounters a Segment Limit or Canonical Fault due to code execution, a #GP (General Protection Exception) fault is generated after all higher priority Interrupts and exceptions are serviced. Due to this erratum, if RSM (Resume from System Management Mode) returns to execution flow that results in a Code Segment Limit or Canonical Fault, the #GP fault may be serviced before a higher priority Interrupt or Exception (e.g. NMI (Non-Maskable Interrupt), Debug break(#DB), Machine Check (#MC), etc.). If the RSM attempts to return to a non-canonical address, the address pushed onto the stack for this #GP fault may not match the non-canonical address that caused the fault.

**Implication:** Operating systems may observe a #GP fault being serviced before higher priority Interrupts and Exceptions. Intel has not observed this erratum on any commercially available software.

**Workaround:** None identified.

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


**AAO4. Performance Monitor SSE Retired Instructions May Return Incorrect Values**

**Problem:** Performance Monitoring counter SIMD\_INST\_RETIRED (Event: C7H) is used to track retired SSE instructions. Due to this erratum, the processor may also count other types of instructions resulting in higher than expected values.

**Implication:** Performance Monitoring counter SIMD\_INST\_RETIRED may report count higher than expected.

**Workaround:** None identified.

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



### **AAO5. Premature Execution of a Load Operation Prior to Exception Handler Invocation**

**Problem:** If any of the below circumstances occur, it is possible that the load portion of the instruction will have executed before the exception handler is entered.

- If an instruction that performs a memory load causes a code segment limit violation.
- If a waiting X87 floating-point (FP) instruction or MMX™ technology (MMX) instruction that performs a memory load has a floating-point exception pending.
- If an MMX or SSE/SSE2/SSE3/SSSE3 extensions (SSE) instruction that performs a memory load and has either CR0.EM=1 (Emulation bit set), or a floating-point Top-of-Stack (FP TOS) not equal to 0, or a DNA exception pending.

**Implication:** In normal code execution where the target of the load operation is to write back memory there is no impact from the load being prematurely executed, or from the restart and subsequent re-execution of that instruction by the exception handler. If the target of the load is to uncached memory that has a system side-effect, restarting the instruction may cause unexpected system behavior due to the repetition of the side-effect. Particularly, while CR0.TS [bit 3] is set, a MOVD/MOVQ with MMX/XMM register operands may issue a memory load before getting the DNA exception.

**Workaround:** Code which performs loads from memory that has side-effects can effectively workaround this behavior by using simple integer-based load instructions when accessing side-effect memory and by ensuring that all code is written such that a code segment limit violation cannot occur as a part of reading from side-effect memory.

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

### **AAO6. MOV To/From Debug Registers Causes Debug Exception**

**Problem:** When in V86 mode, if a MOV instruction is executed to/from a debug registers, a general-protection exception (#GP) should be generated. However, in the case when the general detect enable flag (GD) bit is set, the observed behavior is that a debug exception (#DB) is generated instead.

**Implication:** With debug-register protection enabled (i.e., the GD bit set), when attempting to execute a MOV on debug registers in V86 mode, a debug exception will be generated instead of the expected general-protection fault.

**Workaround:** In general, operating systems do not set the GD bit when they are in V86 mode. The GD bit is generally set and used by debuggers. The debug exception handler should check that the exception did not occur in V86 mode before continuing. If the exception did occur in V86 mode, the exception may be directed to the general-protection exception handler.

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



### **AAO7. Incorrect Address Computed For Last Byte of FXSAVE/FXRSTOR Image Leads to Partial Memory Update**

**Problem:** A partial memory state save of the 512-byte FXSAVE image or a partial memory state restore of the FXRSTOR image may occur if a memory address exceeds the 64KB limit while the processor is operating in 16-bit mode or if a memory address exceeds the 4GB limit while the processor is operating in 32-bit mode.

**Implication:** FXSAVE/FXRSTOR will incur a #GP fault due to the memory limit violation as expected but the memory state may be only partially saved or restored.

**Workaround:** Software should avoid memory accesses that wrap around the respective 16-bit and 32-bit mode memory limits.

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

### **AAO8. Values for LBR/BTS/BTM will be Incorrect after an Exit from SMM**

**Problem:** After a return from SMM (System Management Mode), the CPU will incorrectly update the LBR (Last Branch Record) and the BTS (Branch Trace Store), hence rendering their data invalid. The corresponding data if sent out as a BTM on the system bus will also be incorrect.

**Note:** This issue would only occur when one of the 3 above mentioned debug support facilities are used.

**Implication:** The value of the LBR, BTS, and BTM immediately after an RSM operation should not be used.

**Workaround:** None identified.

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

### **AAO9. Single Step Interrupts with Floating Point Exception Pending May Be Mishandled**

**Problem:** In certain circumstances, when a floating point exception (#MF) is pending during single-step execution, processing of the single-step debug exception (#DB) may be mishandled.

**Implication:** When this erratum occurs, #DB will be incorrectly handled as follows:

- #DB is signaled before the pending higher priority #MF (Interrupt 16)
- #DB is generated twice on the same instruction

**Workaround:** None identified.

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



### **AAO10. Fault on ENTER Instruction May Result in Unexpected Values on Stack Frame**

**Problem:** The ENTER instruction is used to create a procedure stack frame. Due to this erratum, if execution of the ENTER instruction results in a fault, the dynamic storage area of the resultant stack frame may contain unexpected values (i.e. residual stack data as a result of processing the fault).

**Implication:** Data in the created stack frame may be altered following a fault on the ENTER instruction. Please refer to "Procedure Calls For Block-Structured Languages" in *IA-32 Intel® Architecture Software Developer's Manual, Vol. 1, Basic Architecture*, for information on the usage of the ENTER instructions. This erratum is not expected to occur in ring 3. Faults are usually processed in ring 0 and stack switch occurs when transferring to ring 0. Intel has not observed this erratum on any commercially available software.

**Workaround:** None identified.

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

### **AAO11. IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception**

**Problem:** In IA-32e mode, it is possible to get an Alignment Check Exception (#AC) on the IRET instruction even though alignment checks were disabled at the start of the IRET. This can only occur if the IRET instruction is returning from CPL3 code to CPL3 code. IRETs from CPL0/1/2 are not affected. This erratum can occur if the EFLAGS value on the stack has the AC flag set, and the interrupt handler's stack is misaligned. In IA-32e mode, RSP is aligned to a 16-byte boundary before pushing the stack frame.

**Implication:** In IA-32e mode, under the conditions given above, an IRET can get a #AC even if alignment checks are disabled at the start of the IRET. This erratum can only be observed with a software generated stack frame.

**Workaround:** Software should not generate misaligned stack frames for use with IRET.

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

### **AAO12. General Protection Fault (#GP) for Instructions Greater than 15 Bytes May be Preempted**

**Problem:** When the processor encounters an instruction that is greater than 15 bytes in length, a #GP is signaled when the instruction is decoded. Under some circumstances, the #GP fault may be preempted by another lower priority fault (e.g. Page Fault (#PF)). However, if the preempting lower priority faults are resolved by the operating system and the instruction retried, a #GP fault will occur.

**Implication:** Software may observe a lower-priority fault occurring before or in lieu of a #GP fault. Instructions of greater than 15 bytes in length can only occur if redundant prefixes are placed before the instruction.

**Workaround:** None identified.

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



**AAO13. General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit**

**Problem:** In 32-bit mode, memory accesses to flat data segments (base = 00000000h) that occur above the 4G limit (0fffffffh) may not signal a #GP fault.

**Implication:** When such memory accesses occur in 32-bit mode, the system may not issue a #GP fault.

**Workaround:** Software should ensure that memory accesses in 32-bit mode do not occur above the 4G limit (0fffffffh).

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

**AAO14. LBR, BTS, BTM May Report a Wrong Address when an Exception/Interrupt Occurs in 64-bit Mode**

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

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

**Workaround:** None identified.

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

**AAO15. 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.



#### **AAO16. Corruption of CS Segment Register During RSM While Transitioning From Real Mode to Protected Mode**

**Problem:** During the transition from real mode to protected mode, if an SMI (System Management Interrupt) occurs between the MOV to CRO that sets PE (Protection Enable, bit 0) and the first FAR JMP, the subsequent RSM (Resume from System Management Mode) may cause the lower two bits of CS segment register to be corrupted.

**Implication:** The corruption of the bottom two bits of the CS segment register will have no impact unless software explicitly examines the CS segment register between enabling protected mode and the first FAR JMP. *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3A: System Programming Guide, Part 1*, in the section titled "Switching to Protected Mode" recommends the FAR JMP immediately follows the write to CRO to enable protected mode. 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.

#### **AAO17. Performance Monitoring Events for Read Miss to Level 3 Cache Fill Occupancy Counter may be Incorrect**

**Problem:** Whenever an Level 3 cache fill conflicts with another request's address, the miss to fill occupancy counter, UNC\_GO\_ALLOC.RT\_LLC\_MISS (Event 02H), will provide erroneous results.

**Implication:** The Performance Monitoring UNC\_GO\_ALLOC.RT\_LLC\_MISS event may count a value higher than expected. The extent to which the value is higher than expected is determined by the frequency of the L3 address conflict.

**Workaround:** None identified.

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

#### **AAO18. A VM Exit on MWAIT May Incorrectly Report the Monitoring Hardware as Armed**

**Problem:** A processor write to the address range armed by the MONITOR instruction may not immediately trigger the monitoring hardware. Consequently, a VM exit on a later MWAIT may incorrectly report the monitoring hardware as armed, when it should be reported as unarmed due to the write occurring prior to the MWAIT.

**Implication:** If a write to the range armed by the MONITOR instruction occurs between the MONITOR and the MWAIT, the MWAIT instruction may start executing before the monitoring hardware is triggered. If the MWAIT instruction causes a VM exit, this could cause its exit qualification to incorrectly report 0x1. In the recommended usage model for MONITOR/MWAIT, there is no write to the range armed by the MONITOR instruction between the MONITOR and the MWAIT.

**Workaround:** Software should never write to the address range armed by the MONITOR instruction between the MONITOR and the subsequent MWAIT.

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



### **AAO19. Delivery Status of the LINT0 Register of the Local Vector Table May be Lost**

**Problem:** The Delivery Status bit of the LINT0 Register of the Local Vector Table will not be restored after a transition out of C6 under the following conditions

- LINT0 is programmed as level-triggered
- The delivery mode is set to either Fixed or ExtINT
- There is a pending interrupt which is masked with the interrupt enable flag (IF)

**Implication:** Due to this erratum, the Delivery Status bit of the LINT0 Register will unexpectedly not be set. Intel has not observed this erratum with any commercially available software or system.

**Workaround:** None identified.

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

### **AAO20. Performance Monitor Event SEGMENT\_REG\_LOADS Counts Inaccurately**

**Problem:** The performance monitor event SEGMENT\_REG\_LOADS (Event 06H) counts instructions that load new values into segment registers. The value of the count may be inaccurate.

**Implication:** The performance monitor event SEGMENT\_REG\_LOADS may reflect a count higher or lower than the actual number of events.

**Workaround:** None identified.

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

### **AAO21. #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.


### **AAO22. Improper Parity Error Signaled in the IQ Following Reset When a Code Breakpoint is Set on a #GP Instruction**

**Problem:** While coming out of cold reset or exiting from C6, if the processor encounters an instruction longer than 15 bytes (which causes a #GP) and a code breakpoint is enabled on that instruction, an IQ (Instruction Queue) parity error may be incorrectly logged resulting in an MCE (Machine Check Exception).

**Implication:** When this erratum occurs, an MCE may be incorrectly signaled.

**Workaround:** None identified.

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



**AAO23. An Enabled Debug Breakpoint or Single Step Trap May Be Taken after MOV SS/POP SS Instruction if it is Followed by an Instruction That Signals a Floating Point Exception**

**Problem:** A MOV SS/POP SS instruction should inhibit all interrupts including debug breakpoints until after execution of the following instruction. This is intended to allow the sequential execution of MOV SS/POP SS and MOV [r/e]SP, [r/e]BP instructions without having an invalid stack during interrupt handling. However, an enabled debug breakpoint or single step trap may be taken after MOV SS/POP SS if this instruction is followed by an instruction that signals a floating point exception rather than a MOV [r/e]SP, [r/e]BP instruction. This results in a debug exception being signaled on an unexpected instruction boundary since the MOV SS/POP SS and the following instruction should be executed atomically.

**Implication:** This can result in incorrect signaling of a debug exception and possibly a mismatched Stack Segment and Stack Pointer. If MOV SS/POP SS is not followed by a MOV [r/e]SP, [r/e]BP, there may be a mismatched Stack Segment and Stack Pointer on any exception. Intel has not observed this erratum with any commercially available software or system.

**Workaround:** As recommended in the *IA32 Intel® Architecture Software Developer's Manual*, the use of MOV SS/POP SS in conjunction with MOV [r/e]SP, [r/e]BP will avoid the failure since the MOV [r/e]SP, [r/e]BP will not generate a floating point exception. Developers of debug tools should be aware of the potential incorrect debug event signaling created by this erratum.

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

**AAO24. IA32\_MPERF Counter Stops Counting During On-Demand TM1**

**Problem:** According to the *Intel® 64 and IA-32 Architectures Software Developer's Manual* Volume 3A: System Programming Guide, the ratio of IA32\_MPERF (MSR E7H) to IA32\_APERF (MSR E8H) should reflect actual performance while TM1 or on-demand throttling is activated. Due to this erratum, IA32\_MPERF MSR stops counting while TM1 or on-demand throttling is activated, and the ratio of the two will indicate higher processor performance than actual.

**Implication:** The incorrect ratio of IA32\_APERF/IA32\_MPERF can mislead software P-state (performance state) management algorithms under the conditions described above. It is possible for the Operating System to observe higher processor utilization than actual, which could lead the OS into raising the P-state. During TM1 activation, the OS P-state request is irrelevant and while on-demand throttling is enabled, it is expected that the OS will not be changing the P-state. This erratum should result in no practical implication to software.

**Workaround:** None identified.

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

**AAO25. The Memory Controller tTHROT\_OPREF Timings May be Violated During Self Refresh Entry**

**Problem:** During self refresh entry, the memory controller may issue more refreshes than permitted by tTHROT\_OPREF (bits 29:19 in MC\_CHANNEL\_{0,1}\_REFRESH\_TIMING CSR).

**Implication:** The intention of tTHROT\_OPREF is to limit current. Since current supply conditions near self refresh entry are not critical, there is no measurable impact due to this erratum.

**Workaround:** None identified.

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



#### **AAO26. Processor May Over Count Correctable Cache MESI State Errors**

**Problem:** Under a specific set of conditions, correctable Level 2 cache hierarchy MESI state errors may be counted more than once per occurrence of a correctable error.

**Implication:** Correctable Level 2 cache hierarchy MESI state errors may be reported in the MCI\_STATUS register at a rate higher than their actual occurrence.

**Workaround:** None identified.

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

#### **AAO27. Synchronous Reset of IA32\_APERF/IA32\_MPERF Counters on Overflow Does Not Work**

**Problem:** When either the IA32\_MPERF or IA32\_APERF MSR (E7H, E8H) increments to its maximum value of 0xFFFF\_FFFF\_FFFF\_FFFF, both MSRs are supposed to synchronously reset to 0x0 on the next clock. This synchronous reset does not work. Instead, both MSRs increment and overflow independently.

**Implication:** Software can not rely on synchronous reset of the IA32\_APERF/IA32\_MPERF registers.

**Workaround:** None identified.

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

#### **AAO28. Disabling Thermal Monitor While Processor is Hot, Then Re-enabling, May Result in Stuck Core Operating Ratio**

**Problem:** If a processor is at its TCC (Thermal Control Circuit) activation temperature and then Thermal Monitor is disabled by a write to IA32\_MISC\_ENABLE MSR (1A0H) bit [3], a subsequent re-enable of Thermal Monitor will result in an artificial ceiling on the maximum core P-state. The ceiling is based on the core frequency at the time of Thermal Monitor disable. This condition will only correct itself once the processor reaches its TCC activation temperature again.

**Implication:** Since Intel requires that Thermal Monitor be enabled in order to be operating within specification, this erratum should never be seen during normal operation.

**Workaround:** Software should not disable Thermal Monitor during processor operation.

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


#### **AAO29. PECC Does Not Support PCI Configuration Reads/Writes to Misaligned Addresses**

**Problem:** The PECC (Platform Environment Control Interface) specification allows for partial reads from or writes to misaligned addresses within the PCI configuration space. However, the PECC client does not properly interpret addresses that are Dword (4 byte) misaligned and may read or write incorrect data.

**Implication:** Due to this erratum, writes to or reads from Dword misaligned addresses could result in unintended side effects and unpredictable behavior.

**Workaround:** PECC host controllers may issue byte, word and Dword reads and writes as long as they are aligned to Dword addresses.

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



### **AAO30. OVER Bit for IA32\_MCi\_STATUS Register May Get Set on Specific Internal Error**

**Problem:** If a specific type of internal unclassified error is detected, as identified by IA32\_MCi\_STATUS.MCACOD=0x0405, the IA32\_MCi\_STATUS.OVER (overflow) bit [62] may be erroneously set.

**Implication:** The OVER bit of the MCI\_STATUS register may be incorrectly set for a specific internal unclassified error.

**Workaround:** None identified.

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

### **AAO31. Writing the Local Vector Table (LVT) when an Interrupt is Pending May Cause an Unexpected Interrupt**

**Problem:** If a local interrupt is pending when the LVT entry is written, an interrupt may be taken on the new interrupt vector even if the mask bit is set.

**Implication:** An interrupt may immediately be generated with the new vector when a LVT entry is written, even if the new LVT entry has the mask bit set. If there is no Interrupt Service Routine (ISR) set up for that vector the system will GP fault. If the ISR does not do an End of Interrupt (EOI) the bit for the vector will be left set in the in-service register and mask all interrupts at the same or lower priority.

**Workaround:** Any vector programmed into an LVT entry must have an ISR associated with it, even if that vector was programmed as masked. This ISR routine must do an EOI to clear any unexpected interrupts that may occur. The ISR associated with the spurious vector does not generate an EOI, therefore the spurious vector should not be used when writing the LVT.

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

### **AAO32. Faulting MMX Instruction May Incorrectly Update x87 FPU Tag Word**

**Problem:** Under a specific set of conditions, MMX stores (MOVD, MOVQ, MOVNTQ, MASKMOVQ) which cause memory access faults (#GP, #SS, #PF, or #AC), may incorrectly update the x87 FPU tag word register.

This erratum will occur when the following additional conditions are also met.

- The MMX store instruction must be the first MMX instruction to operate on x87 FPU state (i.e. the x87 FP tag word is not already set to 0x0000).
- For MOVD, MOVQ, MOVNTQ stores, the instruction must use an addressing mode that uses an index register (this condition does not apply to MASKMOVQ).

**Implication:** If the erratum conditions are met, the x87 FPU tag word register may be incorrectly set to a 0x0000 value when it should not have been modified.

**Workaround:** None identified.

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



**AAO33. xAPIC Timer May Decrement Too Quickly Following an Automatic Reload While in Periodic Mode**

**Problem:** When the xAPIC Timer is automatically reloaded by counting down to zero in periodic mode, the xAPIC Timer may slip in its synchronization with the external clock. The xAPIC timer may be shortened by up to one xAPIC timer tick.

**Implication:** When the xAPIC Timer is automatically reloaded by counting down to zero in periodic mode, the xAPIC Timer may slip in its synchronization with the external clock. The xAPIC timer may be shortened by up to one xAPIC timer tick.

**Workaround:** None identified.

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

**AAO34. Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures**

**Problem:** Bits 53:50 of the IA32\_VMX\_BASIC MSR report the memory type that the processor uses to access the VMCS and data structures referenced by pointers in the VMCS. Due to this erratum, a VMX access to the VMCS or referenced data structures will instead use the memory type that the MTRRs (memory-type range registers) specify for the physical address of the access.

**Implication:** Bits 53:50 of the IA32\_VMX\_BASIC MSR report that the WB (write-back) memory type will be used but the processor may use a different memory type.

**Workaround:** Software should ensure that the VMCS and referenced data structures are located at physical addresses that are mapped to WB memory type by the MTRRs.

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

**AAO35. B0-B3 Bits in DR6 For Non-Enabled Breakpoints May be Incorrectly Set**

**Problem:** Some of the B0-B3 bits (breakpoint conditions detect flags, bits [3:0]) in DR6 may be incorrectly set for non-enabled breakpoints when the following sequence happens:


1. MOV or POP instruction to SS (Stack Segment) selector;
2. Next instruction is FP (Floating Point) that gets FP assist
3. Another instruction after the FP instruction completes successfully
4. A breakpoint occurs due to either a data breakpoint on the preceding instruction or a code breakpoint on the next instruction.

Due to this erratum a non-enabled breakpoint triggered on step 1 or step 2 may be reported in B0-B3 after the breakpoint occurs in step 4.

**Implication:** Due to this erratum, B0-B3 bits in DR6 may be incorrectly set for non-enabled breakpoints.

**Workaround:** Software should not execute a floating point instruction directly after a MOV SS or POP SS instruction.

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



### **AAO36. Core C6 May Clear Previously Logged TLB Errors**

**Problem:** Following an exit from core C6, previously logged TLB (Translation Lookaside Buffer) errors in IA32\_MCI\_STATUS may be cleared.

**Implication:** Due to this erratum, TLB errors logged in the associated machine check bank prior to core C6 entry may be cleared. Provided machine check exceptions are enabled, the machine check exception handler can log any uncorrectable TLB errors prior to core C6 entry. The TLB marks all detected errors as uncorrectable.

**Workaround:** As long as machine check exceptions are enabled, the machine check exception handler can log the TLB error prior to core C6 entry. This will ensure the error is logged before it is cleared.

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

### **AAO37. Performance Monitor Event MISALIGN\_MEM\_REF May Over Count**

**Problem:** The MISALIGN\_MEM\_REF Performance Monitoring (Event 05H) may over count memory misalignment events, possibly by orders of magnitude.

**Implication:** Software relying on MISALIGN\_MEM\_REF to count cache line splits for optimization purposes may read excessive number of memory misalignment events.

**Workaround:** None identified.

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

### **AAO38. Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations**

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

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

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

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

### **AAO39. Running with Write Major Mode Disabled May Lead to a System Hang**

**Problem:** With write major mode disabled, reads will be favored over writes and under certain circumstances this can lead to a system hang.

**Implication:** Due to this erratum a system hang 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.



#### **AAO40. Infinite Stream of Interrupts May Occur if an ExtINT Delivery Mode Interrupt is Received while All Cores in C6**

**Problem:** If all logical processors in a core are in C6, an ExtINT delivery mode interrupt is pending in the xAPIC and interrupts are blocked with EFLAGS.IF=0, the interrupt will be processed after C6 wakeup and after interrupts are re-enabled (EFLAGS.IF=1). However, the pending interrupt event will not be cleared.

**Implication:** Due to this erratum, an infinite stream of interrupts will occur on the core servicing the external interrupt. Intel has not observed this erratum with any commercially available software/system.

**Workaround:** None identified.

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

#### **AAO41. Two xAPIC Timer Event Interrupts May Unexpectedly Occur**

**Problem:** If an xAPIC timer event is enabled and while counting down the current count reaches 1 at the same time that the processor thread begins a transition to a low power C-state, the xAPIC may generate two interrupts instead of the expected one when the processor returns to C0.

**Implication:** Due to this erratum, two interrupts may unexpectedly be generated by an xAPIC timer event.

**Workaround:** None identified.

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


#### **AAO42. EOI Transaction May Not be Sent if Software Enters Core C6 During an Interrupt Service Routine**

**Problem:** If core C6 is entered after the start of an interrupt service routine but before a write to the APIC EOI register, the core may not send an EOI transaction (if needed) and further interrupts from the same priority level or lower may be blocked.

**Implication:** EOI transactions and interrupts may be blocked when core C6 is used during interrupt service routines. 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.



#### **AAO43. 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.

#### **AAO44. APIC Error “Received Illegal Vector” May be Lost**

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

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

**Workaround:** None identified.

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

#### **AAO45. DR6 May Contain Incorrect Information When the First Instruction After a MOV SS,r/m or POP SS is a Store**

**Problem:** Normally, each instruction clears the changes in DR6 (Debug Status Register) caused by the previous instruction. However, the instruction following a MOV SS,r/m (MOV to the stack segment selector) or POP SS (POP stack segment selector) instruction will not clear the changes in DR6 because data breakpoints are not taken immediately after a MOV SS,r/m or POP SS instruction. Due to this erratum, any DR6 changes caused by a MOV SS,r/m or POP SS instruction may be cleared if the following instruction is a store.

**Implication:** When this erratum occurs, incorrect information may exist in DR6. This erratum will not be observed under normal usage of the MOV SS,r/m or POP SS instructions (i.e., following them with an instruction that writes [e/r]SP). When debugging or when developing debuggers, this behavior should be noted.

**Workaround:** None identified.

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



#### **AAO46. An Uncorrectable Error Logged in IA32\_CR\_MC2\_STATUS May also Result in a System Hang**

**Problem:** Uncorrectable errors logged in IA32\_CR\_MC2\_STATUS MSR (409H) may also result in a system hang causing an Internal Timer Error (MCACOD = 0x0400h) to be logged in another machine check bank (IA32\_MCI\_STATUS).

**Implication:** Uncorrectable errors logged in IA32\_CR\_MC2\_STATUS can further cause a system hang and an Internal Timer Error to be logged.

**Workaround:** None identified.

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

#### **AAO47. IA32\_PERF\_GLOBAL\_CTRL MSR May be Incorrectly Initialized**

**Problem:** The IA32\_PERF\_GLOBAL\_CTRL MSR (38FH) bits [34:32] may be incorrectly set to 7H after reset; the correct value should be 0H.

**Implication:** The IA32\_PERF\_GLOBAL\_CTRL MSR bits [34:32] may be incorrect after reset (EN\_FIXED\_CTR{0, 1, 2} may be enabled).

**Workaround:** None identified.

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

#### **AAO48. ECC Errors Can Not be Injected on Back-to-Back Writes**

**Problem:** ECC errors should be injected on every write that matches the address set in the MC\_CHANNEL\_{0,1}\_ADDR\_MATCH CSRs. Due to this erratum if there are two back-to-back writes that match MC\_CHANNEL\_{0,1}\_ADDR\_MATCH, the 2nd write will not have the error injected.

**Implication:** The 2nd back-to-back write that matches MC\_CHANNEL\_{0,1}\_ADDR\_MATCH will not have the ECC error properly injected. Setting MC\_CHANNEL\_{0,1}\_ADDR\_MATCH to a specific address will reduce the chance of being impacted by this erratum.

**Workaround:** Only injecting errors to specific address should reduce the chance on being impacted by this erratum.

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


#### **AAO49. Performance Monitor Interrupts Generated From Uncore Fixed Counters (394H) May be Ignored**

**Problem:** Performance monitor interrupts (PMI's) from Uncore fixed counters are ignored when Uncore general performance monitor counters 3B0H-3BFH are not programmed.

**Implication:** This erratum blocks a usage model in which each of the cores can sample its own performance monitor events synchronously based on single interrupt from the Uncore.

**Workaround:** Program any one of the Uncore general performance monitor counters with a valid performance monitor event and enable the event by setting the local enable bit in the corresponding performance monitor event select MSR. For the usage model where no counting is desired, program that Uncore general performance counter's global enable bit to be zero.

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



#### **AAO50. Performance Monitor Counter INST\_RETIREDD.STORES May Count Higher than Expected**

**Problem:** Performance Monitoring counter INST\_RETIREDD.STORES (Event: C0H) is used to track retired instructions which contain a store operation. Due to this erratum, the processor may also count other types of instructions including WRMSR and MFENCE.

**Implication:** Performance Monitoring counter INST\_RETIREDD.STORES may report counts higher than expected.

**Workaround:** None identified.

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

#### **AAO51. Sleeping Cores May Not be Woken Up on Logical Cluster Mode Broadcast IPI Using Destination Field Instead of Shorthand**

**Problem:** If software sends a logical cluster broadcast IPI using a destination shorthand of 00B (No Shorthand) and writes the cluster portion of the Destination Field of the Interrupt Command Register to all ones while not using all 1s in the mask portion of the Destination Field, target cores in a sleep state that are identified by the mask portion of the Destination Field may not be woken up. This erratum does not occur if the destination shorthand is set to 10B (All Including Self) or 11B (All Excluding Self).

**Implication:** When this erratum occurs, cores which are in a sleep state may not wake up to handle the broadcast IPI. Intel has not observed this erratum with any commercially available software.

**Workaround:** Use destination shorthand of 10B or 11B to send broadcast IPIs.

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

#### **AAO52. Faulting Executions of FXRSTOR May Update State Inconsistently**

**Problem:** The state updated by a faulting FXRSTOR instruction may vary from one execution to another.

**Implication:** Software that relies on x87 state or SSE state following a faulting execution of FXRSTOR may behave inconsistently.

**Workaround:** Software handling a fault on an execution of FXRSTOR can compensate for execution variability by correcting the cause of the fault and executing FXRSTOR again.

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

#### **AAO53. Performance Monitor Event EPT.EPDPE\_MISS May be Counted While EPT is Disable**

**Problem:** Performance monitor event EPT.EPDPE\_MISS (Event: 4FH, Umask: 08H) is used to count Page Directory Pointer table misses while EPT (extended page tables) is enabled. Due to this erratum, the processor will count Page Directory Pointer table misses regardless of whether EPT is enabled or not.

**Implication:** Due to this erratum, performance monitor event EPT.EPDPE\_MISS may report counts higher than expected.

**Workaround:** Software should ensure this event is only enabled while in EPT mode.

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



#### **AAO54. Memory Aliasing of Code Pages May Cause Unpredictable System Behavior**

**Problem:** The type of memory aliasing contributing to this erratum is the case where two different logical processors have the same code page mapped with two different memory types. Specifically, if one code page is mapped by one logical processor as write-back and by another as uncacheable and certain instruction fetch timing conditions occur, the system may experience unpredictable behavior.

**Implication:** If this erratum occurs the system may have unpredictable behavior including a system hang. The aliasing of memory regions, a condition necessary for this erratum to occur, is documented as being unsupported in the *Intel 64 and IA-32 Intel® Architecture Software Developer's Manual, Volume 3A*, in the section titled *Programming the PAT*. Intel has not observed this erratum with any commercially available software or system.

**Workaround:** Code pages should not be mapped with uncacheable and cacheable memory types at the same time.

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

#### **AAO55. Performance Monitor Counters May Count Incorrectly**

**Problem:** Under certain circumstances, a general purpose performance counter, IA32\_PMC0-4 (C1H - C4H), may count at core frequency or not count at all instead of counting the programmed event.

**Implication:** The Performance Monitor Counter IA32\_PMCx may not properly count the programmed event. Due to the requirements of the workaround there may be an interruption in the counting of a previously programmed event during the programming of a new event.

**Workaround:** Before programming the performance event select registers, IA32\_PERFEVTSELx MSR (186H - 189H), the internal monitoring hardware must be cleared. This is accomplished by first disabling, saving valid events and clearing from the select registers, then programming three event values 0x4300D2, 0x4300B1 and 0x4300B5 into the IA32\_PERFEVTSELx MSRs, and finally continuing with new event programming and restoring previous programming if necessary. Each performance counter, IA32\_PMCx, must have its corresponding IA32\_PFEVTSELx MSR programmed with at least one of the event values and must be enabled in IA32\_PERF\_GLOBAL\_CTRL MSR (38FH) bits [3:0]. All three values must be written to either the same or different IA32\_PERFEVTSELx MSRs before programming the performance counters. Note that the performance counter will not increment when its IA32\_PERFEVTSELx MSR has a value of 0x4300D2, 0x4300B1 or 0x4300B5 because those values have a zero UMASK field (bits [15:8]).

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



**AAO56. Processor Forward Progress Mechanism Interacting With Certain MSR/CSR Writes May Cause Unpredictable System Behavior**

**Problem:** Under specific internal conditions, a mechanism within the processor to ensure forward progress may interact with writes to a limited set of MSRs/CSRs and consequently may lead to unpredictable system behavior.

**Implication:** This erratum may cause 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.

**AAO57. Performance Monitor Event Offcore\_response\_0 (B7H) Does Not Count NT Stores to Local DRAM Correctly**

**Problem:** When a IA32\_PERFVTSELx MSR is programmed to count the Offcore\_response\_0 event (Event:B7H), selections in the OFFCORE\_RSP\_0 MSR (1A6H) determine what is counted. The following two selections do not provide accurate counts when counting NT (Non-Temporal) Stores:

- OFFCORE\_RSP\_0 MSR bit [14] is set to 1 (LOCAL\_DRAM) and bit [7] is set to 1 (OTHER): NT Stores to Local DRAM are not counted when they should have been.
- OFFCORE\_RSP\_0 MSR bit [9] is set to (OTHER\_CORE\_HIT\_SNOOP) and bit [7] is set to 1 (OTHER): NT Stores to Local DRAM are counted when they should not have been.

**Implication:** The counter for the Offcore\_response\_0 event may be incorrect for NT stores.

**Workaround:** None identified.

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


**AAO58. 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.



**AAO59. System May Hang if MC\_CHANNEL\_{0,1}\_MC\_DIMM\_INIT\_CMD.DO\_ZQCL Commands Are Not Issued in Increasing Populated DDR3 Rank Order**

**Problem:** ZQCL commands are used during initialization to calibrate DDR3 termination. A ZQCL command can be issued by writing 1 to the MC\_CHANNEL\_{0,1}\_MC\_DIMM\_INIT\_CMD.DO\_ZQCL (Device 4,5,6, Function 0, Offset 15, bit[15]) field and it targets the DDR3 rank specified in the RANK field (bits[7:5]) of the same register. If the ZQCL commands are not issued in increasing populated rank order then ZQ calibration may not complete, causing the system to hang.

**Implication:** Due to this erratum the system may hang if writes to the MC\_CHANNEL\_{0,1}\_MC\_DIMM\_INIT\_CMD.DO\_ZQCL field are not in increasing populated DDR3 rank order.

**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.

**AAO60. Package C3/C6 Transitions When Memory 2x Refresh is Enabled May Result in a System Hang**

**Problem:** If ASR\_PRESENT (MC\_CHANNEL\_{0,1}\_REFRESH\_THROTTLE\_SUPPORT CSR function 0, offset 68H, bit [0], Auto Self Refresh Present) is clear which indicates that high temperature operation is not supported on the DRAM, the memory controller will not enter self-refresh if software has REF\_2X\_NOW (bit 4 of the MC\_CLOSED\_LOOP CSR, function 3, offset 84H) set. This scenario may cause the system to hang during C3/C6 entry.

**Implication:** Failure to enter self-refresh can delay C3/C6 power state transitions to the point that a system hang may result with CATERR being asserted. REF\_2X\_NOW is used to double the refresh rate when the DRAM is operating in extended temperature range. The ASR\_PRESENT was intended to allow low power self refresh with DRAM that does not support automatic self refresh.

**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.


**AAO61. Back to Back Uncorrected Machine Check Errors May Overwrite IA32\_MC3\_STATUS.MSCOD**

**Problem:** When back-to-back uncorrected machine check errors occur that would both be logged in the IA32\_MC3\_STATUS MSR (40CH), the IA32\_MC3\_STATUS.MSCOD (bits [31:16]) field may reflect the status of the most recent error and not the first error. The rest of the IA32\_MC3\_STATUS MSR contains the information from the first error.

**Implication:** Software should not rely on the value of IA32\_MC3\_STATUS.MSCOD if IA32\_MC3\_STATUS.OVER (bit [62]) is set.

**Workaround:** None identified.

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



### **AAO62. Memory Intensive Workloads with Core C6 Transitions May Cause System Hang**

**Problem:** Under a complex set of internal conditions, a system running a high cache stress and I/O workload combined with the presence of frequent core C6 transitions may result in a system hang.

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

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

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

### **AAO63. Corrected Errors With a Yellow Error Indication May be Overwritten by Other Corrected Errors**

**Problem:** A corrected cache hierarchy data or tag error that is reported with IA32\_MCi\_STATUS.MCACOD (bits [15:0]) with value of 000x\_0001\_xxxx\_xx01 (where x stands for zero or one) and a yellow threshold-based error status indication (bits [54:53] equal to 10B) may be overwritten by a corrected error with a no tracking indication (00B) or green indication (01B).

**Implication:** Corrected errors with a yellow threshold-based error status indication may be overwritten by a corrected error without a yellow indication.

**Workaround:** None identified.

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

### **AAO64. PSI # Signal May Incorrectly be Left Asserted**

**Problem:** When some of the cores in the processor are in C3/C6 state, the PSI# (Power Status Indicator) signal may incorrectly be left asserted when another core makes a frequency change request without changing the operating voltage. Since this erratum results in a possible maximum core current greater than the PSI# threshold of 15A, PSI# should have been de-asserted.

**Implication:** Due to this erratum, platform voltage regulator tolerances may be exceeded and a subsequent system reset 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.

### **AAO65. Memory ECC Errors May be Observed When a UC Partial Write is Followed by a UC Read to the Same Location**

**Problem:** Memory ECC errors may be observed when an uncachable partial write is closely followed by an uncacheable read to the same location. This erratum will only occur if the BIOS is configured to use Closed\_Page memory mode (Bit 0 of MC\_CONTROL register Device 3 Function 0 Offset 48H set to 1).

**Implication:** Correctable and Uncorrectable ECC errors may be logged in the IA32\_MCi\_Status registers.

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



#### **AAO66. Performance Monitor Events DCACHE\_CACHE\_LD and DCACHE\_CACHE\_ST May Overcount**

**Problem:** The performance monitor events DCACHE\_CACHE\_LD (Event 40H) and DCACHE\_CACHE\_ST (Event 41H) count cacheable loads and stores that hit the L1 cache. Due to this erratum, in addition to counting the completed loads and stores, the counter will incorrectly count speculative loads and stores that were aborted prior to completion.

**Implication:** The performance monitor events DCACHE\_CACHE\_LD and DCACHE\_CACHE\_ST may reflect a count higher than the actual number of events.

**Workaround:** None identified.

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

#### **AAO67. Rapid Core C3/C6 Transitions May Cause Unpredictable System Behavior**

**Problem:** Under a complex set of internal conditions, cores rapidly performing C3/C6 transitions in a system with Intel® Hyper-Threading Technology enabled may cause a machine check error (IA32\_MCI\_STATUS.MCACOD = 0x0106), system hang or unpredictable system behavior.

**Implication:** This erratum may cause a machine check error, system hang or 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.

#### **AAO68. Performance Monitor Events INSTR\_RETIRED and MEM\_INST\_RETIRED May Count Inaccurately**

**Problem:** The performance monitor event INSTR\_RETIRED (Event C0H) should count the number of instructions retired, and MEM\_INST\_RETIRED (Event 0BH) should count the number of load or store instructions retired. However, due to this erratum, they may undercount.

**Implication:** The performance monitor event INSTR\_RETIRED and MEM\_INST\_RETIRED may reflect a count lower than the actual number of events.

**Workaround:** None identified.

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


#### **AAO69. A Page Fault May Not be Generated When the PS bit is set to "1" in a PML4E or PDPTE**

**Problem:** On processors supporting Intel® 64 architecture, the PS bit (Page Size, bit 7) is reserved in PML4Es and PDPTEs. If the translation of the linear address of a memory access encounters a PML4E or a PDPTE with PS set to 1, a page fault should occur. Due to this erratum, PS of such an entry is ignored and no page fault will occur due to its being set.

**Implication:** Software may not operate properly if it relies on the processor to deliver page faults when reserved bits are set in paging-structure entries.

**Workaround:** Software should not set bit 7 in any PML4E or PDPTE that has Present Bit (Bit 0) set to "1".

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



#### **AAO70. CPURESET Bit Does Not Get Cleared**

**Problem:** CPURESET (bit 10 of SYRE Device 8; Function 2; Offset 0CCH) allows the processor to be independently reset without assertion of the PLTRST# signal upon a 0 to 1 transition. The CPURESET bit does not get cleared and must be cleared by software.

**Implication:** The processor will not be reset if a 1 is written to this bit while it is already a one.

**Workaround:** The CPURESET bit must be cleared by software prior to setting it.

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

#### **AAO71. PHOLD Disable in MISCCTRLSTS Register Does Not Work**

**Problem:** PHOLD Disable (PCI Hold Disable, bit [23] in MISCCTRLSTS Device 0; Function 0; Offset 188H) does not function as described. Setting this bit will not cause the processor to respond with Unsupported Request and log a fatal error upon receiving an Assert\_PHOLD message from the PCH (Platform Controller Hub).

**Implication:** Due to this erratum, it is not possible to disable PHOLD requests from the PCH.

**Workaround:** None identified.

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

#### **AAO72. PCIe PMCSR Power State Field Incorrectly Allows Requesting of the D1 and D2 Power States**

**Problem:** The PCIe PMCSR (Power Management Control and Status Register, Device 3,4,5,6; Function 0; Offset E4H) incorrectly allows the writing/requesting of the D1 and D2 Power States in the Power State field (bits[1:0] of PMCSR) when these states are not supported.

**Implication:** Given that the device does not support the D1 and D2 states, attempts to write those states should have been ignored. The PCIe port does not change power state from D0 or D3hot when the Power State bits are written to D1 or D2, so there is no functional impact to the PCIe port. However, the Power State field is incorrectly modified.

**Workaround:** None identified.

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

#### **AAO73. PECl Accesses to Registers May Fail When Processor is Transitioning to/from Package C6 Power State**

**Problem:** A PECl (Platform Environment Control Interface) access to PCI configuration registers while the device is transitioning to or from package C6 may fail. Writes may not update the target register and reads may return incorrect data. The PECl bus will not show any indication the transaction failed.

**Implication:** PECl accesses to PCI configuration registers may not be processed 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.



#### **AAO74. Concurrent Updates to a Segment Descriptor May be Lost**

**Problem:** If a logical processor attempts to set the accessed bit in a code or data segment descriptor while another logical processor is modifying the same descriptor, both modifications of the descriptor may be lost.

**Implication:** Due to this erratum, updates to segment descriptors may not be preserved. Intel has not observed this erratum with any commercially available software or system.

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

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

#### **AAO75. PMIs May be Lost During Core C6 Transitions**

**Problem:** If a performance monitoring counter overflows and causes a PMI (Performance Monitoring Interrupt) at the same time that the core is entering C6, then the PMI may be lost.

**Implication:** PMIs may be lost during a C6 transition.

**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.

#### **AAO76. Uncacheable Access to a Monitored Address Range May Prevent Future Triggering of the Monitor Hardware**

**Problem:** It is possible that an address range which is being monitored via the MONITOR instruction could be written without triggering the monitor hardware. A read from the monitored address range which is issued as uncacheable (for example having the CRO.CD bit set) may prevent subsequent writes from triggering the monitor hardware. A write to the monitored address range which is issued as uncacheable, may not trigger the monitor hardware and may prevent subsequent writes from triggering the monitor hardware.

**Implication:** The MWAIT instruction will not exit the optimized power state and resume program flow if the monitor hardware is not triggered.

**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.


#### **AAO77. BIST Results May be Additionally Reported After a GETSEC[WAKEUP] or INIT-SIPI Sequence**

**Problem:** BIST results should only be reported in EAX the first time a logical processor wakes up from the Wait-For-SIPI state. Due to this erratum, BIST results may be additionally reported after INIT-SIPI sequences and when waking up RLP's from the SENTER sleep state using the GETSEC[WAKEUP] command.

**Implication:** An INIT-SIPI sequence may show a non-zero value in EAX upon wakeup when a zero value is expected. RLP's waking up for the SENTER sleep state using the GETSEC[WAKEUP] command may show a different value in EAX upon wakeup than before going into the SENTER sleep state.

**Workaround:** If necessary software may save the value in EAX prior to launching into the secure environment and restore upon wakeup and/or clear EAX after the INIT-SIPI sequence.

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



### **AAO78. 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.

### **AAO79. VM Exits Due to "NMI-Window Exiting" May Be Delayed by One Instruction**

**Problem:** If VM entry is executed with the "NMI-window exiting" VM-execution control set to 1, a VM exit with exit reason "NMI window" should occur before execution of any instruction if there is no virtual-NMI blocking, no blocking of events by MOV SS, and no blocking of events by STI. If VM entry is made with no virtual-NMI blocking but with blocking of events by either MOV SS or STI, such a VM exit should occur after execution of one instruction in VMX non-root operation. Due to this erratum, the VM exit may be delayed by one additional instruction.

**Implication:** VMM software using "NMI-window exiting" for NMI virtualization should generally be unaffected, as the erratum causes at most a one-instruction delay in the injection of a virtual NMI, which is virtually asynchronous. The erratum may affect VMMs relying on deterministic delivery of the affected VM exits.

**Workaround:** None identified.

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

### **AAO80. Malformed PCIe Packet Generated Under Heavy Outbound Load**

**Problem:** When running the PCIe ports in a 2x8 configuration at 5.0GT/S speed with heavy outbound write traffic, malformed packets could be generated. The length in the header field will not match the actual payload size.

**Implication:** Due to this erratum, malformed PCIe packets could be transmitted.

**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.

### **AAO81. PCIe Operation in x16 Mode With Inbound Posted Writes May be Unreliable**

**Problem:** Under a complex set of conditions, it is possible that with PCIe configured for x16 operation inbound writes may store incorrect data.

**Implication:** PCIe operation with inbound writes in x16 mode may be unreliable.

**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.



#### **AAO82. Unpredictable PCI Behavior Accessing Non-existent Memory Space**

**Problem:** Locked instructions whose memory reference is split across cache line boundaries and are aborted on PCI behind Intel® 5 Series Chipset and Intel® 3400 Series Chipset may cause subsequent PCI writes to be unpredictable.

**Implication:** Aborted split lock accesses to non-existent PCI memory space behind Intel 5 Series Chipset and Intel 3400 Series Chipset may cause PCI devices to subsequently become inoperable until a platform reset. Intel has not observed this erratum with commercially available software and has only observed this in a synthetic test environment.

**Workaround:** None identified.

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

#### **AAO83. PECI MbxGet() Commands May Fail Several Times Before Passing When Issued During Package C6**

**Problem:** PECI (Platform Environment Control Interface) MbxSend() requests may become blocked when the processor is in package C6. This temporary blocking may cause subsequent MbxGet() commands to result in the receipt of a bad write FCS (frame checksum).

**Implication:** Due to this erratum, as long as the host retries the MbxGet() command the results will be delivered upon the subsequent exit from package C6, but this may take several milliseconds depending on the platform or operating system.

**Workaround:** PECI MbxGet() commands may need to be retried several times before successful completion.

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

#### **AAO84. VM Exits Due to EPT Violations Do Not Record Information About Pre-IRET NMI Blocking**

**Problem:** With certain settings of the VM-execution controls VM exits due to EPT violations set bit 12 of the exit qualification if the EPT violation was a result of an execution of the IRET instruction that commenced with non-maskable interrupts (NMIs) blocked. Due to this erratum, such VM exits will instead clear this bit.

**Implication:** Due to this erratum, a virtual-machine monitor that relies on the proper setting of bit 12 of the exit qualification may deliver NMIs to guest software prematurely.

**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.


#### **AAO85. Intel® VT-d Receiving Two Identical Interrupt Requests May Corrupt Attributes of Remapped Interrupt or Hang a Subsequent Interrupt-Remap-Cache Invalidation Command**

**Problem:** If the Intel® VT-d (Intel® Virtualization Technology for Directed I/O) interrupt-remapping hardware receives two identical back-to-back interrupt requests, then the attributes of the remapped interrupt returned may be corrupted. This interrupt sequence may also hang the system if the software executes a subsequent interrupt-remap-cache invalidation command.

**Implication:** This scenario may lead to unpredictable external interrupt behavior; or a subsequent interrupt-remap-cache invalidation command submitted by software may hang.

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

**Workaround:**



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

**AAO86. S1 Entry May Cause Cores to Exit C3 or C6 C-State**

**Problem:** Under specific circumstances, S1 entry may cause a logical processor to spuriously wake up from C3 or C6 and transition to a C0/S1 state. Upon S1 exit, these logical processors will be operating in C0.

**Implication:** In systems where S1 is used for power savings, customers may observe higher S1 power than expected and software may observe a different C-state on S1 exit than on S1 entry.

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

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

**AAO87. Multiple Performance Monitor Interrupts are Possible on Overflow of IA32\_FIXED\_CTR2**

**Problem:** When multiple performance counters are set to generate interrupts on an overflow and more than one counter overflows at the same time, only one interrupt should be generated. However, if one of the counters set to generate an interrupt on overflow is the IA32\_FIXED\_CTR2 (MSR 30BH) counter, multiple interrupts may be generated when the IA32\_FIXED\_CTR2 overflows at the same time as any of the other performance counters.

**Implication:** Multiple counter overflow interrupts may be unexpectedly generated.

**Workaround:** None identified.

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

**AAO88. LBRs May Not be Initialized During Power-On Reset of the Processor**

**Problem:** If a second reset is initiated during the power-on processor reset cycle, the LBRs (Last Branch Records) may not be properly initialized.

**Implication:** Due to this erratum, debug software may not be able to rely on the LBRs out of power-on reset.

**Workaround:** Ensure that the processor has completed its power-on reset cycle prior to initiating a second reset.

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


**AAO89. Unexpected Interrupts May Occur on C6 Exit If Using APIC Timer to Generate Interrupts**

**Problem:** During a complex set of conditions, if the APIC timer is being used to generate interrupts, unexpected interrupts not related to the APIC timer may be signaled when a core exits the C6 power state. The APIC timer stops counting in C6 and as such isn't typically used to generate interrupts when the C6 core power state is enabled.

**Implication:** Unexpected interrupt vectors could be sent from the APIC to a logical processor. 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.



#### **AAO90. Package C6 Exit with Memory in Self-Refresh When Using DDR3 RDIMM Memory May Lead to a System Hang**

**Problem:** When using DDR3 RDIMM memory and exiting from the C6 low power state with memory in self-refresh the CS (Chip Select) signals may remain in tri-state during tSTAB (CLK Stabilization time) thus violating the JEDEC Standard: *Definition of the SSTE32882 Registering Clock Driver with Parity and Quad Chip Selects for DDR3 RDIMM Applications*. As detailed in the JEDEC specification the CS signals should transition from tri-state to high to exit the Clock Stopped Power Down Mode.

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

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

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

#### **AAO91. LBR, BTM or BTS Records May have Incorrect Branch From Information After an EIST Transition, T-states, C1E, or Adaptive Thermal Throttling**

**Problem:** The "From" address associated with the LBR (Last Branch Record), BTM (Branch Trace Message) or BTS (Branch Trace Store) may be incorrect for the first branch after an EIST (Enhanced Intel® SpeedStep Technology) transition, T-states, C1E (C1 Enhanced), or Adaptive Thermal Throttling.

**Implication:** When the LBRs, BTM or BTS are enabled, some records may have incorrect branch "From" addresses for the first branch after an EIST transition, T-states, C1E, or Adaptive Thermal Throttling.

**Workaround:** None identified.

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


#### **AAO92. PECI GetTemp() Reads May Return Invalid Temperature Data in Package C6 State**

**Problem:** The PECI (Platform Environment Control Interface) GetTemp() command may occasionally return incorrect temperature data.

**Implication:** The temperature data reported over PECI should always be a negative value and represents a delta below the onset of TCC (thermal control circuit) activation, as indicated by PROCHOT#. The PECI GetTemp() command may occasionally return incorrect temperature data when the processor is in the package C6 state. The error occurrence rate and returned processor temperature values are random including both hot and cold readings. Note that this error may cause the processor to return positive PECI temperature values that may not necessarily be indicative of a thermal event requiring an immediate shutdown.

**Workaround:** Intel recommends discarding processor temperature values less than -100 or greater than 0, and the use of appropriate temperature smoothing filters in the range -100 to 0 to minimize fan speed fluctuations, if any, due to these errors. Intel does not recommend initiating system shutdown solely based on PECI readings. For systems using the PECI temperature data to facilitate system shutdown, Intel recommends initiating a shutdown only if a PECI value of 0 is returned over three consecutive PECI temperature reads.

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



**AAO93. PECE PCIConfigRd() Followed by a GetTemp() May Cause System Hang in Package C6 State**

**Problem:** The PECE (Platform Environment Control Interface) PCIConfigRd() command immediately followed by a PECE GetTemp() command may result in a system hang.

**Implication:** When the processor is in the package C6 state, a PECE PCIConfigRd() command immediately followed by a GetTemp() command may result in a system hang. If PCIConfigRd() is never used, then this erratum will not be observed.

**Workaround:** A PCIConfigWr() command should be issued in between PCIConfigRd() and GetTemp() commands. The PCIConfigWr() command may be issued to any valid PECE writable CSR address, including a benign CSR address such as 0x23058.

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

**AAO94. PECE Mailbox Commands During Package C6 Idle State Transitions May Result in Unpredictable Processor Behavior**

**Problem:** If a PECE (Platform Environment Control Interface) mailbox command is executed at the same time that the processor is entering or exiting the package C6 idle state, unpredictable processor behavior or an incorrect mailbox response may result.

**Implication:** The PECE mailbox commands are not reliable during processor package C6 idle state and may result in unpredictable processor behavior or incorrect PECE responses.

**Workaround:** It is possible for the BIOS to contain a workaround for this erratum. This workaround involves disabling PECE mailbox functions during package C6 idle state causing the processor to miss responding to requests during this time. The workaround may also result in PECE mailbox completion code responses of 0x85 ("Mailbox is Idle - no data available") when executed during normal active operating conditions.

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


**AAO95. VMX-Preemption Timer Does Not Count Down at the Rate Specified**

**Problem:** The VMX-preemption timer should count down by 1 every time a specific bit in the TSC (Time Stamp Counter) changes. (This specific bit is indicated by IA32\_VMX\_MISC bits [4:0] (0x485h) and has a value of 5 on the affected processors.) Due to this erratum, the VMX-preemption timer may instead count down at a different rate and may do so only intermittently.

**Implication:** The VMX-preemption timer may cause VM exits at a rate different from that expected by software.

**Workaround:** None identified.

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



#### **AAO96. Multiple Performance Monitor Interrupts are Possible on Overflow of Fixed Counter 0**

**Problem:** The processor can be configured to issue a PMI (performance monitor interrupt) upon overflow of the IA32\_FIXED\_CTR0 MSR (309H). A single PMI should be observed on overflow of IA32\_FIXED\_CTR0, however multiple PMIs are observed when this erratum occurs.

This erratum only occurs when IA32\_FIXED\_CTR0 overflows and the processor and counter are configured as follows:

- Intel® Hyper-Threading Technology is enabled
- IA32\_FIXED\_CTR0 local and global controls are enabled
- IA32\_FIXED\_CTR0 is set to count events only on its own thread (IA32\_FIXED\_CTR\_CTRL MSR (38DH) bit [2] = '0')
- PMIs are enabled on IA32\_FIXED\_CTR0 (IA32\_FIXED\_CTR\_CTRL MSR bit [3] = '1')
- Freeze\_on\_PMI feature is enabled (IA32\_DEBUGCTL MSR (1D9H) bit [12] = '1')

**Implication:** When this erratum occurs there may be multiple PMIs observed when IA32\_FIXED\_CTR0 overflows

**Workaround:** Disable the FREEZE\_PERFMON\_ON\_PMI feature in IA32\_DEBUGCTL MSR (1D9H) bit [12].

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

#### **AAO97. SVID and SID of Devices 8 and 16 only implement bits [7:0]**

**Problem:** Bits [15:8] of SVID (Subsystem Vendor ID, Offset 2CH) and the SID (Subsystem Device ID, Offset 2EH) of devices 8 and 16 are not implemented. Only the lower bits [7:0] of these registers can be written to, though the PCI-e specification indicates that these are 16-bit registers.

**Implication:** Only bits [7:0] of SVID and SID can be written. Bits [15:8] will always be read as 0.

**Workaround:** None identified.

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


#### **AAO98. No\_Soft\_Reset Bit in the PMCSR Does Not Operate as Expected**

**Problem:** When the No\_Soft\_Reset bit in the Power Management Control and Status Register (PMCSR; Bus 0; Devices 0, 3, 4, 5; Function 0; Offset 0xE4; Bit 3) is cleared the device should perform an internal reset upon transitioning from D3<sub>hot</sub> to D0. Due to this erratum the device does not perform an internal reset upon transitioning from D3<sub>hot</sub> to D0.

**Implication:** When the No\_Soft\_reset bit in the PMCSR register is set or cleared no internal reset of the device will be performed when transitioning from D3<sub>hot</sub> to D0.

**Workaround:** None identified.

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



### **AAO99. VM Exits Due to LIDT/LGDT/SIDT/SGDT Do Not Report Operand Size**

**Problem:** When a VM exit occurs outside 64-bit mode due to a LIDT, LGDT, SIDT, or SGDT instruction, the operand size should be reported in bits 13:12 of the VM-Exit Instruction-Information Field (value 1 indicates 16-bit; value 2 indicates 32-bit). Due to this erratum, these bits are always reported as zero.

**Implication:** Virtual Machine Monitors cannot rely on bits 13:12 of the VM-Exit Instruction-Information Field to determine the operand size of the instruction causing the VM exit.

**Workaround:** Virtual Machine Monitor software may decode the instruction to determine operand size.

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

### **AAO100. PCIConfigRd() and PCIConfigWr() PECl Commands May Silently Fail During Package C6 Exit Events**

**Problem:** When PCIConfigRd() or PCIConfigWr() commands coincide with processor package C6 exits under the right timing conditions, they may fail to execute but still produce 'passing' responses.

**Implication:** When the timing conditions of this erratum are met, reads will return a value of "all zeroes" for the return data and writes will have no effect while both commands will return a passing completion code. The rate of occurrence of this issue is dependent on frequency and duration of C6 entry/exit events and PECl polling rate.

**Workaround:** None identified.

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

### **AAO101. Performance Monitoring Events STORE\_BLOCKS.NOT\_STA and STORE\_BLOCKS.STA May Not Count Events Correctly**

**Problem:** Performance Monitor Events STORE\_BLOCKS.NOT\_STA and STORE\_BLOCKS.STA should only increment the count when a load is blocked by a store. Due to this erratum, the count will be incremented whenever a load hits a store, whether it is blocked or can forward. In addition this event does not count for specific threads correctly.

**Implication:** If Intel® Hyper-Threading Technology is disabled, the Performance Monitor events STORE\_BLOCKS.NOT\_STA and STORE\_BLOCKS.STA may indicate a higher occurrence of loads blocked by stores than have actually occurred. If Intel Hyper-Threading Technology is enabled, the counts of loads blocked by stores may be unpredictable and they could be higher or lower than the correct count.

**Workaround:** None identified.


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

### **AAO102. Storage of PEBS Record Delayed Following Execution of MOV SS or STI**

**Problem:** When a performance monitoring counter is configured for PEBS (Precise Event Based Sampling), overflow of the counter results in storage of a PEBS record in the PEBS buffer. The information in the PEBS record represents the state of the next instruction to be executed following the counter overflow. Due to this erratum, if the counter overflow occurs after execution of either MOV SS or STI, storage of the PEBS record is delayed by one instruction.

**Implication:** When this erratum occurs, software may observe storage of the PEBS record being delayed by one instruction following execution of MOV SS or STI. The state information in the PEBS record will also reflect the one instruction delay.

**Workaround:** None identified.



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

**AAO103. Performance Monitoring Event FP\_MMX\_TRANS\_TO\_MMX May Not Count Some Transitions**

**Problem:** Performance Monitor Event FP\_MMX\_TRANS\_TO\_MMX (Event CCH, Umask 01H) counts transitions from x87 Floating Point (FP) to MMX™ instructions. Due to this erratum, if only a small number of MMX instructions (including EMMS) are executed immediately after the last FP instruction, a FP to MMX transition may not be counted.

**Implication:** The count value for Performance Monitoring Event FP\_MMX\_TRANS\_TO\_MMX may be lower than expected. The degree of undercounting is dependent on the occurrences of the erratum condition while the counter is active. 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.

**AAO104. INVLPG Following INVEPT or INVVPID May Fail to Flush All Translations for a Large Page**

**Problem:** This erratum applies if the address of the memory operand of an INVEPT or INVVPID instruction resides on a page larger than 4KBytes and either (1) that page includes the low 1 MBytes of physical memory; or (2) the physical address of the memory operand matches an MTRR that covers less than 4 MBytes. A subsequent execution of INVLPG that targets the large page and that occurs before the next VM-entry instruction may fail to flush all TLB entries for the page. Such entries may persist in the TLB until the next VM-entry instruction.

**Implication:** Accesses to the large page between INVLPG and the next VM-entry instruction may incorrectly use translations that are inconsistent with the in-memory page tables.

**Workaround:** None Identified.

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

**AAO105. The PECCI Bus May be Tri-stated After System Reset**

**Problem:** During power-up, the processor may improperly assert the PECCI (Platform Environment Control Interface) pin. This condition is cleared as soon as Bus Clock starts toggling. However, if the PECCI host (also referred to as the master or originator) incorrectly determines this asserted state as another PECCI host initiating a transaction, it may release control of the bus resulting in a permanent tri-state condition.

**Implication:** Due to this erratum, the PECCI host may incorrectly determine that it is not the bus master and consequently PECCI commands initiated by the PECCI software layer may receive incorrect/invalid responses.

**Workaround:** To workaround this erratum the PECCI host should pull the PECCI bus low to initiate a PECCI transaction.


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

**AAO106. 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.

#### **AAO107. Multiple ECC Errors May Result in Incorrect Syndrome Being Logged**

**Problem:** When multiple correctable DRAM ECC errors occur the processor may log the syndrome of a previous error or may log an unknown value.

**Implication:** Due to this erratum, the value logged in the IA32\_MCI\_MISC MSR will not correspond to the most recent error.

**Workaround:** None identified.

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

#### **AAO108. 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.

#### **AAO109. Debug Exception Flags DR6.B0-B3 Flags May be Incorrect for Disabled Breakpoints**

**Problem:** When a debug exception is signaled on a load that crosses cache lines with data forwarded from a store and whose corresponding breakpoint enable flags are disabled (DR7.G0-G3 and DR7.L0-L3), the DR6.B0-B3 flags may be incorrect.

**Implication:** The debug exception DR6.B0-B3 flags may be incorrect for the load if the corresponding breakpoint enable flag in DR7 is disabled.

**Workaround:** None identified.

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

#### **AAO110. An Exit From the Core C6-state May Result in the Dropping of an Interrupt**

**Problem:** In a complex set of internal conditions when the processor exits from Core C6 state, it is possible that an interrupt may be dropped.

**Implication:** Due to this erratum, an interrupt may be dropped. Intel has not observed this erratum with any commercially available software.

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

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

#### **AAO111. PCIe Extended Capability Structures May be Incorrect**

**Problem:** The PCIe Extended Capability structure at Offset 0x100 of Bus 0; Devices 0, 3, 4, 5 and 6 contains a Capability ID of AER (Advanced Error Reporting), but these devices do not support AER. The Next Capability Offset field of this Extended Capability structure contains 0x150 which is the offset of the next Extended Capability structure. For Bus 0; Devices 4, 5, and 6, the Next Capability Offset field of the Extended Capability structure at offset 0x150 should contain 0 to indicate the end of the capability chain but



instead contains 0x160. All fields of the Extended Capability structure at offset 0x160 are 0x0. A Capability ID of 0x0 is a reserved Capability ID.

**Implication:** Software that enables features based upon the existence of the AER may not observe the expected behavior associated with this capability.

**Workaround:** None identified.

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

### **AAO112. PMIs During Core C6 Transitions May Cause the System to Hang**

**Problem:** If a performance monitoring counter overflows and causes a PMI (Performance Monitoring Interrupt) at the same time that the core enters C6, then this may cause the system to hang.

**Implication:** Due to this erratum, the processor may hang when a PMI coincides with core C6 entry.

**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.


### **AAO113. IA32\_MC8\_CTL2 MSR is Not Cleared on Processor Warm Reset**

**Problem:** After processor warm reset the IA32\_MC8\_CTL2 MSR (288H) should be zero. Due to this erratum the IA32\_MC8\_CTL2 MSR is not zeroed on processor warm reset.

**Implication:** When this erratum occurs, the IA32\_MC8\_CTL2 MSR will not be zeroed by warm reset. Software that expects the values to be 0 coming out of warm reset may not behave as expected.

**Workaround:** BIOS should zero the IA32\_MC8\_CTL2 MSR after a warm reset.

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




# Specification Changes

---

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

- *Intel® Xeon® Processor 3400 Series Datasheet – Volumes 1 and 2*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

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



## Specification Clarifications

---

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

- *Intel® Xeon® Processor 3400 Series Datasheet – Volumes 1 and 2*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

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



## Documentation Changes

---

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

- *Intel® Xeon® Processor 3400 Series Datasheet – Volumes 1 and 2*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference Manual A-M*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B: Instruction Set Reference Manual N-Z*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A: System Programming Guide*
- *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3B: System Programming Guide*

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

Note: Documentation changes for Intel® 64 and IA-32 Architecture Software Developer's Manual volumes 1, 2A, 2B, 3A, and 3B will be posted in a separate document, Intel® 64 and IA-32 Architecture Software Developer's Manual Documentation Changes. Follow the link below to become familiar with this file.

<http://developer.intel.com/products/processor/manuals/index.htm>

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

### §