

Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series

Specification Update

February 2015

Document Number: 320767-028US



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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

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 information here is subject to change without notice. Do not finalize a design with this information.

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

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

Code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user.

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 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 for details.

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>.

For information on the Enhanced Intel SpeedStep® Technology, see the Processor Spec Finder at <http://ark.intel.com> or contact your Intel representative.

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.

Intel, Intel Core, Celeron, Pentium, 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 – 2015, Intel Corporation. All Rights Reserved.



Contents

| | |
|------------------------------------|----|
| Revision History | 4 |
| Preface | 6 |
| Summary Tables of Changes | 8 |
| Identification Information | 15 |
| Errata | 18 |
| Specification Changes | 56 |
| Specification Clarifications | 60 |
| Documentation Changes | 61 |


§ §



Revision History

| Revision | Description | Date |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| 001 | Initial Release | September 2009 |
| 002 | January Release <ul style="list-style-type: none">• Addition of Errata AAP99-104• Updated AAP84 and AAP90 | January 2010 |
| 003 | February Release <ul style="list-style-type: none">• Added AAP105 and AAP106 | February 2010 |
| 004 | March Release <ul style="list-style-type: none">• Added AAP107 | March 2010 |
| 005 | <ul style="list-style-type: none">• Updated processor identification table with new SKU info• Added AAP108 and AAP109 | April 2010 |
| 006 | <ul style="list-style-type: none">• Added AAP1 in Specification Changes | May 2010 |
| 007 | <ul style="list-style-type: none">• Added AAP110 and AAP111• Added AAP2 and AAP3 in Specification Changes | June 2010 |
| 008 | <ul style="list-style-type: none">• Removed Item Numbering• Added AAP112, AAP113, AAP114 and AAP115 | July 2010 |
| 009 | <ul style="list-style-type: none">• Added AAP116, AAP117 and AAP118• Added AAP4 in Specification Changes | September 2010 |
| 010 | <ul style="list-style-type: none">• Updated processor identification table with new SKU information• Updated AAP41• Added AAP119 and AAP120 | October 2010 |
| 011 | <ul style="list-style-type: none">• Added AAP121 | December 2010 |
| 012 | <ul style="list-style-type: none">• Added AAP122, AAP123 and AAP124 | January 2011 |
| 013 | <ul style="list-style-type: none">• Added AAP125 to AAP130 | February 2011 |
| 014 | <ul style="list-style-type: none">• Added AAP131 | June 2011 |
| 015 | <ul style="list-style-type: none">• Added AAP132 | August 2011 |
| 016 | <ul style="list-style-type: none">• Added AAP133 | September 2011 |
| 017 | <ul style="list-style-type: none">• Updated Erratum AAP122 | October 2011 |
| 018 | <ul style="list-style-type: none">• Added Erratum AAP134 | December 2011 |
| 019 | <ul style="list-style-type: none">• Added Erratum AAP135 | March 2012 |
| 020 | <ul style="list-style-type: none">• Added Documentation Change AAP1 | January 2013 |
| 021 | <ul style="list-style-type: none">• Added Erratum AAP136 | May 2013 |
| 022 | <ul style="list-style-type: none">• Added errata AAP137 and AAP138 | June 2013 |
| 023 | <ul style="list-style-type: none">• Added errata AAP139 and AAP140 | August 2013 |
| 024 | <ul style="list-style-type: none">• No errata added or deleted• Document standardization | October 2013 |
| 025 | <ul style="list-style-type: none">• No errata added or deleted• Document standardization | December 2013 |

Revision History



| Revision | Description | Date |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| 026 | <ul style="list-style-type: none">Updated link to access Intel® 64 and IA-32 Architecture Software Developer's Manual Documentation Changes | July 2014 |
| 027 | <ul style="list-style-type: none">Removed Erratum AAP94 | November 2014 |
| 028 | <ul style="list-style-type: none">Updated Erratum AAP136 | February 2015 |

§ §



Preface

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

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

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

Affected Documents

| Document Title | Document Number / Location |
|-------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| <i>Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 1</i> | 320765 |
| <i>Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2</i> | 320766 |

Related Documents

| Document Title | Document Number / Location |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>AP-485, Intel® Processor Identification and the CPUID Instruction</i> | http://www.intel.com/design/processor/aplnots/241618.htm |
| <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1: Basic Architecture</i> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2A: Instruction Set Reference 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> <i>Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes (see note 1)</i> | http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html |
| <i>ACPI Specifications</i> | www.acpi.info |

Notes:

1. Documentation changes for Intel® 64 and IA-32 Architecture Software Developer's Manual volumes 1, 2A, 2B, 3A, and 3B, and bug fixes are posted in the *Intel® 64 and IA-32 Architecture Software Developer's Manual Documentation Changes*.



Nomenclature

Errata are design defects or errors. These may cause the Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series 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, L3 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 table row indicates this erratum is either new or modified from the previous version of the document.



Errata (Sheet 1 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP1 | X | No Fix | The Processor May Report a #TS Instead of a #GP Fault |
| AAP2 | X | No Fix | REP MOVSB/STOSB 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 |
| AAP3 | 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 |
| AAP4 | X | No Fix | Performance Monitor SSE Retired Instructions May Return Incorrect Values |
| AAP5 | X | No Fix | Premature Execution of a Load Operation Prior to Exception Handler Invocation |
| AAP6 | X | No Fix | MOV To/From Debug Registers Causes Debug Exception |
| AAP7 | X | No Fix | Incorrect Address Computed for Last Byte of FXSAVE/FXRSTOR Image Leads to Partial Memory Update |
| AAP8 | X | No Fix | Values for LBR/BTS/BTM Will Be Incorrect after an Exit from SMM |
| AAP9 | X | No Fix | Single Step Interrupts with Floating Point Exception Pending May Be Mishandled |
| AAP10 | X | No Fix | Fault on ENTER Instruction May Result in Unexpected Values on Stack Frame |
| AAP11 | X | No Fix | IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception |
| AAP12 | X | No Fix | General Protection Fault (#GP) for Instructions Greater Than 15 Bytes May Be Preempted |
| AAP13 | X | No Fix | General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit |
| AAP14 | X | No Fix | LBR, BTS, BTM May Report a Wrong Address When an Exception/Interrupt Occurs in 64-bit Mode |
| AAP15 | X | No Fix | MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang |
| AAP16 | X | No Fix | Corruption of CS Segment Register during RSM While Transitioning From Real Mode to Protected Mode |
| AAP17 | X | No Fix | Performance Monitoring Events for Read Miss to Level 3 Cache Fill Occupancy Counter May Be Incorrect |
| AAP18 | X | No Fix | A VM Exit on MWAIT May Incorrectly Report the Monitoring Hardware As Armed |
| AAP19 | X | No Fix | Delivery Status of the LINT0 Register of the Local Vector Table May Be Lost |
| AAP20 | X | No Fix | Performance Monitor Event SEGMENT_REG_LOADS Counts Inaccurately |
| AAP21 | X | No Fix | #GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code |
| AAP22 | X | No Fix | Improper Parity Error Signaled in the IQ Following Reset When a Code Breakpoint Is Set on a #GP Instruction |
| AAP23 | 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 |
| AAP24 | X | No Fix | IA32_MPERF Counter Stops Counting during On-Demand TM1 |



Errata (Sheet 2 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|-----------------------------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP25 | X | No Fix | The Memory Controller tTHROT_OPREF Timings May Be Violated during Self Refresh Entry |
| AAP26 | X | No Fix | Processor May Over Count Correctable Cache MESI State Errors |
| AAP27 | X | No Fix | Synchronous Reset of IA32_APERF/IA32_MPERF Counters on Overflow Does Not Work |
| AAP28 | X | No Fix | Disabling Thermal Monitor While Processor Is Hot, Then Re-enabling, May Result in Stuck Core Operating Ratio |
| AAP29 | X | No Fix | OVER Bit for IA32_MCI_STATUS Register May Get Set on Specific Internal Error |
| AAP30 | X | No Fix | Writing the Local Vector Table (LVT) When an Interrupt Is Pending May Cause an Unexpected Interrupt |
| AAP31 | X | No Fix | Faulting Intel® MMX™ Technology Instruction May Incorrectly Update x87 FPU Tag Word |
| AAP32 | X | No Fix | xAPIC Timer May Decrement Too Quickly following an Automatic Reload While in Periodic Mode |
| AAP33 | X | No Fix | Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures |
| AAP34 | X | No Fix | B0-B3 Bits in DR6 for Non-Enabled Breakpoints May Be Incorrectly Set |
| AAP35 | X | No Fix | Core C6 May Clear Previously Logged TLB Errors |
| AAP36 | X | No Fix | Performance Monitor Event MISALIGN_MEM_REF May Over Count |
| AAP37 | X | No Fix | Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations |
| AAP38 | X | No Fix | Running with Write Major Mode Disabled May Lead to a System Hang |
| AAP39 | X | No Fix | Infinite Stream of Interrupts May Occur If an ExtINT Delivery Mode Interrupt Is Received While All Cores in C6 |
| AAP40 | X | No Fix | Two xAPIC Timer Event Interrupts May Unexpectedly Occur |
| AAP41 | X | No Fix | EOI Transaction May Not Be Sent If Software Enters Core C6 during an Interrupt Service Routine |
| AAP42 | X | No Fix | FREEZE_WHILE_SMM Does Not Prevent Event from Pending PEBS during SMM |
| AAP43 | X | No Fix | APIC Error "Received Illegal Vector" May Be Lost |
| AAP44 | X | No Fix | DR6 May Contain Incorrect Information When the First Instruction after a MOV SS,r/m or POP SS Is a Store |
| AAP45 | X | No Fix | An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang |
| AAP46 | X | No Fix | IA32_PERF_GLOBAL_CTRL MSR May Be Incorrectly Initialized |
| AAP47 | X | No Fix | Performance Monitor Interrupts Generated from Uncore Fixed Counters (394H) May Be Ignored |
| AAP48 | X | No Fix | Performance Monitor Counter INST_RETIRED.STORES May Count Higher Than Expected |
| AAP49 | X | No Fix | Sleeping Cores May Not Be Woken Up on Logical Cluster Mode Broadcast IPI Using Destination Field Instead of Shorthand |
| AAP50 | X | No Fix | Faulting Executions of FXRSTOR May Update State Inconsistently |



Errata (Sheet 3 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|------------------------------------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP51 | X | No Fix | Performance Monitor Event EPT.EPDPE_MISS May Be Counted While EPT Is Disabled |
| AAP52 | X | No Fix | Memory Aliasing of Code Pages May Cause Unpredictable System Behavior |
| AAP53 | X | No Fix | Performance Monitor Counters May Count Incorrectly |
| AAP54 | X | No Fix | Processor Forward Progress Mechanism Interacting with Certain MSR/CSR Writes May Cause Unpredictable System Behavior |
| AAP55 | X | No Fix | Performance Monitor Event Offcore_response_0 (B7H) Does Not Count NT Stores to Local DRAM Correctly |
| AAP56 | X | No Fix | EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change |
| AAP57 | 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 |
| AAP58 | X | No Fix | Package C3/C6 Transitions When Memory 2x Refresh Is Enabled May Result in a System Hang |
| AAP59 | X | No Fix | Back to Back Uncorrected Machine Check Errors May Overwrite IA32_MC3_STATUS.MSCOD |
| AAP60 | X | No Fix | Memory Intensive Workloads with Core C6 Transitions May Cause System Hang |
| AAP61 | X | No Fix | Corrected Errors with a Yellow Error Indication May Be Overwritten by Other Corrected Errors |
| AAP62 | X | No Fix | PSI# Signal May Incorrectly Be Left Asserted |
| AAP63 | X | No Fix | Performance Monitor Events DCACHE_CACHE_LD and DCACHE_CACHE_ST May Overcount |
| AAP64 | X | No Fix | Rapid Core C3/C6 Transitions May Cause Unpredictable System Behavior |
| AAP65 | X | No Fix | Performance Monitor Events INSTR_RETIRED and MEM_INST_RETIRED May Count Inaccurately |
| AAP66 | X | No Fix | A Page Fault May Not Be Generated When the PS Bit Is Set to "1" in a PML4E or PDPTE |
| AAP67 | X | No Fix | CPURESET Bit Does Not Get Cleared |
| AAP68 | X | No Fix | PHOLD Disable in MISCCTRLSTS Register Does Not Work |
| AAP69 | X | No Fix | PCIe PMCSR Power State Field Incorrectly Allows Requesting of the D1 and D2 Power States |
| AAP70 | X | No Fix | Concurrent Updates to a Segment Descriptor May Be Lost |
| AAP71 | X | No Fix | PMIs May Be Lost during Core C6 Transitions |
| AAP72 | X | No Fix | Uncacheable Access to a Monitored Address Range May Prevent Future Triggering of the Monitor Hardware |
| AAP73 | X | No Fix | BIST Results May Be Additionally Reported after a GETSEC[WAKEUP] or INIT-SIPI Sequence |
| AAP74 | X | No Fix | Pending x87 FPU Exceptions (#MF) May Be Signaled Earlier Than Expected |
| AAP75 | X | No Fix | VM Exits Due to "NMI-Window Exiting" May Be Delayed by One Instruction |
| AAP76 | X | No Fix | Malformed PCIe Packet Generated under Heavy Outbound Load |
| AAP77 | X | No Fix | PCIe* Operation in x16 Mode with Inbound Posted Writes May Be Unreliable |



Errata (Sheet 4 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP78 | X | No Fix | Unpredictable PCI Behavior Accessing Non-existent Memory Space |
| AAP79 | X | No Fix | VM Exits Due to EPT Violations Do Not Record Information about Pre-IRET NMI Blocking |
| AAP80 | 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 |
| AAP81 | X | No Fix | S1 Entry May Cause Cores to Exit C3 or C6 C-State |
| AAP82 | X | No Fix | Multiple Performance Monitor Interrupts Are Possible on Overflow of IA32_FIXED_CTR2 |
| AAP83 | X | No Fix | LBRs May Not Be Initialized during Power-On Reset of the Processor |
| AAP84 | X | No Fix | Unexpected Interrupts May Occur on C6 Exit If Using APIC Timer to Generate Interrupts |
| AAP85 | X | No Fix | LBR, BTM or BTS Records May Have Incorrect Branch from Information after an Enhanced Intel SpeedStep® Technology Transition, T-states, C1E, or Adaptive Thermal Throttling |
| AAP86 | X | No Fix | VMX-Preemption Timer Does Not Count Down at the Rate Specified |
| AAP87 | X | No Fix | Multiple Performance Monitor Interrupts Are Possible on Overflow of Fixed Counter 0 |
| AAP88 | X | No Fix | SVID and SID of Devices 8 and 16 Only Implement Bits [7:0] |
| AAP89 | X | No Fix | No_Soft_Reset Bit in the PMCSR Does Not Operate As Expected |
| AAP90 | X | No Fix | VM Exits Due to LIDT/LGDT/SIDT/SGDT Do Not Report Correct Operand Size |
| AAP91 | X | No Fix | DPRSLPVR Signal May Be Incorrectly Asserted on Transition Between Low Power C-states |
| AAP92 | X | No Fix | Performance Monitoring Events STORE_BLOCKS.NOT_STA and STORE_BLOCKS.STA May Not Count Events Correctly |
| AAP93 | X | No Fix | Storage of PEBS Record Delayed Following Execution of MOV SS or STI |
| AAP94 | X | No Fix | <Erratum Removed> |
| AAP95 | X | No Fix | INVLPG Following INVEPT or INVVPID May Fail to Flush All Translations for a Large Page |
| AAP96 | X | No Fix | LER MSRs May Be Unreliable |
| AAP97 | X | No Fix | MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error |
| AAP98 | X | No Fix | Debug Exception Flags DR6.B0-B3 Flags May Be Incorrect for Disabled Breakpoints |
| AAP99 | X | No Fix | An Exit From the Core C6-state May Result in the Dropping of an Interrupt |
| AAP100 | X | No Fix | PCIe* Extended Capability Structures May Be Incorrect |
| AAP101 | X | No Fix | PMIs during Core C6 Transitions May Cause the System to Hang |
| AAP102 | X | No Fix | 2MB Page Split Lock Accesses Combined with Complex Internal Events May Cause Unpredictable System Behavior |
| AAP103 | X | No Fix | IA32_MC8_CTL2 MSR Is Not Cleared on Processor Warm Reset |
| AAP104 | X | No Fix | The TPM's Locality 1 Address Space Cannot Be Opened |



Errata (Sheet 5 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP105 | X | No Fix | PCIe* Link Bit Errors Present during L0s Entry May Cause the System to Hang during L0s Exit |
| AAP106 | X | No Fix | The Combination of a Page-Split Lock Access and Data Accesses That Are Split across Cacheline Boundaries May Lead to Processor Livelock |
| AAP107 | X | No Fix | FP Data Operand Pointer May Be Incorrectly Calculated after an FP Access Which Wraps a 4-Gbyte Boundary in Code That Uses 32-Bit Address Size in 64-bit Mode |
| AAP108 | X | No Fix | IOTLB Invalidations Not Completing on Intel® VT-d Engine for Integrated High Definition Audio |
| AAP109 | X | No Fix | IO_SMI Indication in SMRAM State Save Area May Be Lost |
| AAP110 | X | No Fix | PCIe* Squelch Detect May be Slow to Respond During L0s Entry and May Cause a Surprise Link Down Condition |
| AAP111 | X | No Fix | TR Corruption Due to Save/Restore x87 FPU Pointers in SMRAM |
| AAP112 | X | No Fix | PCIe* Lanes Returning to The Active Power State May Cause The System to Hang |
| AAP113 | X | No Fix | Performance Monitor Events for Hardware Prefetches Which Miss The L1 Data Cache May be Over Counted |
| AAP114 | X | No Fix | Poisoned Write Caused by an Internal Parity Error Targeting IIO PCI Configuration Registers or MMIO Space will Not be Suppressed |
| AAP115 | X | No Fix | VM Exit May Incorrectly Clear IA32_PERF_GLOBAL_CTRL [34:32] |
| AAP116 | X | No Fix | PCIe* Port's LTSSM May Not Transition Properly in the Presence of TS1 or TS2 Ordered Sets That Have Unexpected Symbols Within those Sets |
| AAP117 | X | No Fix | NTB/RP Link Will Send Extra TS2 Ordered Set During Link Training |
| AAP118 | X | No Fix | PCIe* Ports May Not Enter Slave Loopback Mode From the Configuration LTSSM State |
| AAP119 | X | No Fix | Unexpected DMI and PCIe* Link Retraining and Correctable Errors Reported |
| AAP120 | X | No Fix | QPI Lane May Be Dropped During Full Frequency Deskew Phase of Training |
| AAP121 | X | No Fix | PerfMon Overflow Status Can Not be Cleared After Certain Conditions Have Occurred |
| AAP122 | X | No Fix | An Unexpected Page Fault or EPT Violation May Occur After Another Logical Processor Creates a Valid Translation for a Page |
| AAP123 | X | No Fix | L1 Data Cache Errors May be Logged With Level Set to 1 Instead of 0 |
| AAP124 | X | No Fix | Stack Pushes May Not Occur Properly for Events Delivered Immediately After VM Entry to 16-Bit Software |
| AAP125 | X | No Fix | Executing The GETSEC Instruction While Throttling May Result in a Processor Hang |
| AAP126 | X | No Fix | PerfMon Event LOAD_HIT_PRE.SW_PREFETCH May Overcount |
| AAP127 | X | No Fix | Successive Fixed Counter Overflows May be Discarded |
| AAP128 | X | No Fix | #GP May be Signaled When Invalid VEX Prefix Precedes Conditional Branch Instructions |
| AAP129 | X | No Fix | A Logical Processor May Wake From Shutdown State When Branch-Trace Messages or Branch-Trace Stores Are Enabled |
| AAP130 | X | No Fix | Task Switch to a TSS With an Inaccessible LDTR Descriptor May Cause Unexpected Faults |



Errata (Sheet 6 of 6)

| Number | Steppings | Status | ERRATA |
|--------|-----------|--------|--------------------------------------------------------------------------------------------------|
| | B-1 | | |
| AAP131 | X | No Fix | VM Entries That Return From SMM Using VMLAUNCH May Not Update The Launch State of the VMCS |
| AAP132 | X | No Fix | VM Entry May Clear Bytes 81H-83H on Virtual-APIC Page When "Use TPR Shadow" Is 0 |
| AAP133 | X | No Fix | A First Level Data Cache Parity Error May Result in Unexpected Behavior |
| AAP134 | X | No Fix | Intel® Trusted Execution Technology ACM Revocation |
| AAP135 | X | No Fix | An Event May Intervene Before a System Management Interrupt That Results from IN or INS |
| AAP136 | X | No Fix | The Corrected Error Count Overflow Bit in IA32_ MCO_STATUS is Not Updated When the UC Bit is Set |
| AAP137 | X | No Fix | The Upper 32 Bits of CR3 May be Incorrectly Used With 32-Bit Paging |
| AAP138 | X | No Fix | EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly |
| AAP139 | X | No Fix | SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior |
| AAP140 | X | No Fix | Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash |

Specification Changes

| Number | SPECIFICATION CHANGES |
|--------|---------------------------------------------------------------------------|
| AAP1 | Update to Datasheet - Volume 2 to Uncore Revision Identification Register |
| AAP2 | Update to Datasheet - Volume 2 to PCI Express Device Control Register 2 |
| AAP3 | Update to Datasheet - Volume 2 to Completion Timeout Control Register |
| AAP4 | Update to Datasheet - Volume 1 to Table 35 and Table 41 |

Specification Clarifications

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

Documentation Changes

| Number | DOCUMENTATION CHANGES |
|--------|--------------------------------------------------|
| AAP1 | On-Demand Clock Modulation Feature Clarification |

§ §



Identification Information

Component Identification using Programming Interface

The Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series stepping can be identified by the following processor signatures:

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

Notes:

1. The Extended Family, bits [27:20] are used in conjunction with the Family Code, specified in bits [11:8], to indicate whether the processor belongs to the Intel386®, Intel486®, Pentium®, Pentium 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® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series can be identified by the following register contents:

| Processor Stepping | Vendor ID ¹ | Device ID ² | Revision ID ³ |
|--------------------|------------------------|------------------------|--------------------------|
| B-1 | 8086h | D132h | 11h |

Notes:

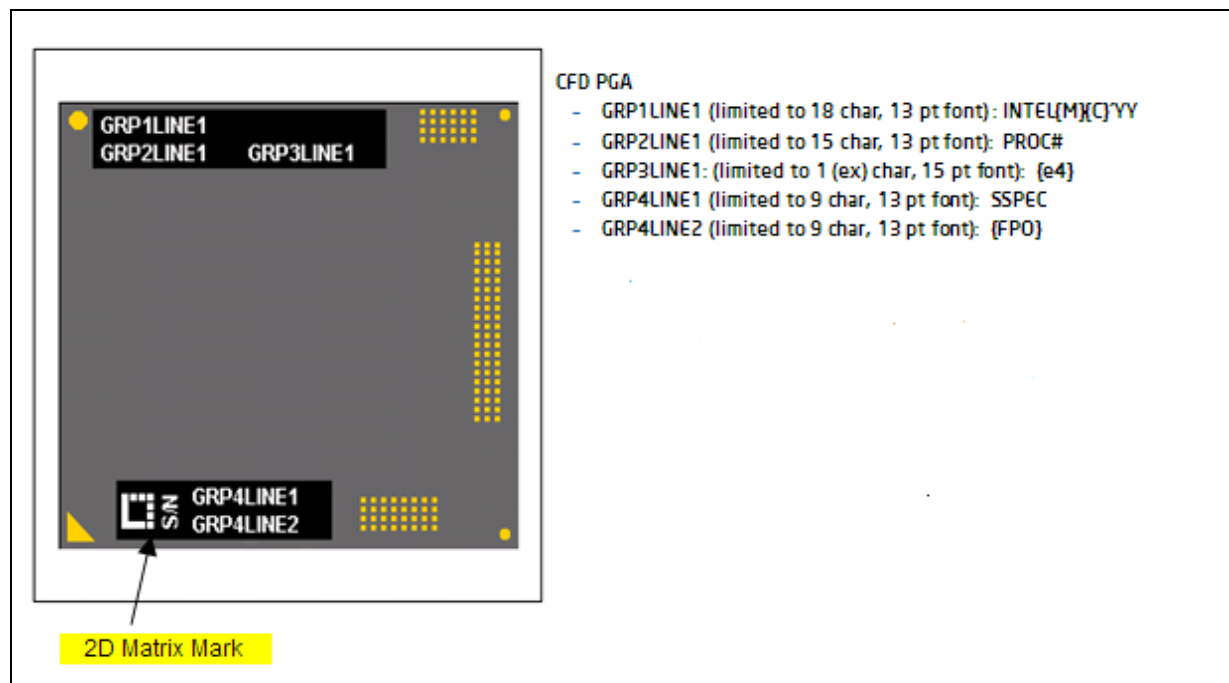
1. The Vendor ID corresponds to Bits 15:0 of the Vendor ID Register located at Offset 00–01h in the PCI Function 0 configuration space.
2. The Device ID corresponds to Bits 15:0 of the Device ID Register located at Device 0 Offset 02–03h in the PCI Function 0 configuration space.
3. The Revision Number 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. Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Component Markings






Table 1. Processor Identification

| S-Spec Number | Processor Number | Stepping | Processor Signature | Core Frequency (GHz) / DDR3 (MHz) | Max Intel® Turbo Boost Technology Frequency (GHz) ¹ | LFM Frequency (GHz) | Shared L3 Cache Size (MB) | Notes |
|---------------|------------------|----------|---------------------|-----------------------------------|----------------------------------------------------------------|---------------------|---------------------------|---------------|
| SLBLW | i7-920XM | B-1 | 106E5h | 2.00 / 1333 | 4 core: 2.26 3 core: 2.26 2 core: 3.06 1 core: 3.20 | 1.200 | 8 | 2, 3, 4, 5 |
| SLBLX | i7-820QM | B-1 | 106E5h | 1.73 / 1333 | 4 core: 2.00 3 core: 2.00 2 core: 2.80 1 core: 3.06 | 1.200 | 8 | 2, 3, 4, 5 |
| SLBLY | i7-720QM | B-1 | 106E5h | 1.60 / 1333 | 4 core: 1.73 3 core: 1.73 2 core: 2.40 1 core: 2.80 | 0.933 | 6 | 2, 3, 4, 5 |
| SLBSC | i7-940XM | B-1 | 106E5h | 2.13 / 1333 | 4 core: 2.40 3 core: 2.40 2 core: 3.20 1 core: 3.33 | 1.200 | 8 | 2, 3, 4, 5, 6 |
| SLBMP | i7-840QM | B-1 | 106E5h | 1.86 / 1333 | 4 core: 2.00 3 core: 2.00 2 core: 2.93 1 core: 3.20 | 1.200 | 8 | 2, 3, 4, 5, 6 |
| SLBQG | i7-740QM | B-1 | 106E5h | 1.73 / 1333 | 4 core: 1.86 3 core: 1.86 2 core: 2.53 1 core: 2.93 | 0.933 | 6 | 2, 3, 4, 5, 6 |

Notes:

1. This column indicates maximum Intel® Turbo Boost Technology frequency (GHz) for 4, 3, 2, or 1 cores active respectively.
2. Intel® Hyper-Threading Technology enabled.
3. Intel® Trusted Execution Technology (Intel® TXT) enabled.
4. Intel® Virtualization Technology (Intel® VT) for IA-32, Intel® 64 and Intel® Architecture (Intel® VT-x) enabled. Intel® Virtualization Technology for Directed I/O (Intel® VT-d) enabled.
5. The core frequency reported in the processor brand string is rounded to 2 decimal digits. (For example, core frequency of 3.06666, repeating 6, is reported as @3.07 in brand string. Core frequency of 1.7333, is reported as @1.73 in brand string.)

§ §



Errata

AAP1. 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.

AAP2. REP MOVSB/STOSB 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 MOVSB or REP STOSB as fast strings. Due to this erratum fast string REP MOVSB/REP STOSB 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 MOVSB or REP STOSB instruction that will execute with fast strings enabled.

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

AAP3. 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.



AAP4. 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.

AAP5. 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/MOVB 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.

AAP6. 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.

**AAP7. 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.

AAP8. 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.

AAP9. 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.

AAP10. 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. 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.



AAP11. 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.

AAP12. 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.

AAP13. 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 4-G 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 4-G limit (0fffffffh).

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

AAP14. 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.

**AAP15. 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.

AAP16. 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 CR0 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 CR0 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.

AAP17. 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_GQ_ALLOC.RT_LLC_MISS (Event 02H), will provide erroneous results.

Implication: The Performance Monitoring UNC_GQ_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.



AAP18. 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.

AAP19. 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.

AAP20. 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.

AAP21. #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.

**AAP22. 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.

AAP23. 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.

AAP24. 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.



AAP25. 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.

AAP26. 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.

AAP27. 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.

AAP28. 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.

AAP29. 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.

**AAP30. 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.

AAP31. Faulting Intel® MMX™ Technology Instruction May Incorrectly Update x87 FPU Tag Word

Problem: Under a specific set of conditions, Intel MMX technology 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 Intel MMX technology store instruction must be the first Intel MMX technology 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.

AAP32. 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.

**AAP33. 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.

AAP34. 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.

AAP35. 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.

AAP36. 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.

**AAP37. 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.

AAP38. 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.

AAP39. 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.

AAP40. 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.

AAP41. 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 may be lost and interrupts may be blocked when core C6 is used during interrupt service routines.

Workaround: Software should check the ISR register and if any interrupts are in service only enter C1.

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



AAP42. 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.

AAP43. 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.

AAP44. 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.

**AAP45. 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.

AAP46. 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.

AAP47. 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.

AAP48. Performance Monitor Counter INST_RETIRED.STORES May Count Higher Than Expected

Problem: Performance Monitoring counter INST_RETIRED.STORES (Event: COH) 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_RETIRED.STORES may report counts higher than expected.

Workaround:None identified.

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



AAP49. 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.

AAP50. 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.

AAP51. Performance Monitor Event EPT.EPDPE_MISS May Be Counted While EPT Is Disabled

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.

AAP52. 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.

**AAP53. 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_PERFEVTSELx 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.

AAP54. 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.

AAP55. Performance Monitor Event Offcore_response_0 (B7H) Does Not Count NT Stores to Local DRAM Correctly

Problem: When a IA32_PERFEVTSELx 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.



AAP56. 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.

AAP57. 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.

AAP58. 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.

**AAP59. 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.

AAP60. 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.

AAP61. 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.

AAP62. 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.



AAP63. 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.

AAP64. 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.

AAP65. 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.

AAP66. 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.

**AAP67. 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.

AAP68. 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.

AAP69. 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.

AAP70. 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.

AAP71. 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.

**AAP72. 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.

AAP73. 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.

AAP74. 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.

**AAP75. 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.

AAP76. 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.

AAP77. 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.

AAP78. 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.

**AAP79. 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.

AAP80. 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.

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

AAP81. 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.

AAP82. 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.

**AAP83. 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.

AAP84. Unexpected Interrupts May Occur on C6 Exit If Using APIC Timer to Generate Interrupts

Problem: 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. This erratum may occur when the APIC timer is near expiration when entering the core C6 state.

Implication: Due to this erratum, unexpected interrupt vectors could be sent from the APIC to a logical processor.

Workaround: Software should stop the APIC timer (by writing 0 to the Initial Count Register) before allowing the core to enter the C6 state.

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

AAP85. LBR, BTM or BTS Records May Have Incorrect Branch from Information after an Enhanced Intel SpeedStep® Technology 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 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 Enhanced Intel SpeedStep Technology transition, T-states, C1E, or Adaptive Thermal Throttling.

Workaround: None identified.

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

AAP86. 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.



AAP87. 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.

AAP88. 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.

AAP89. 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_{hot} to D0. Due to this erratum the device does not perform an internal reset upon transitioning from D3_{hot} 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 preformed when transitioning from D3_{hot} to D0.

Workaround: None identified.

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

**AAP90. VM Exits Due to LIDT/LGDT/SIDT/SGDT Do Not Report Correct Operand Size**

Problem: When a VM exit occurs due to a LIDT, LGDT, SIDT, or SGDT instruction with a 32-bit operand, Bit 11 of the VM-exit instruction information field should be set to 1. Due to this erratum, this bit is instead cleared to 0 (indicating a 16-bit operand).

Implication: Virtual-machine monitors cannot rely on Bit 11 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.

AAP91. DPRSLPVR Signal May Be Incorrectly Asserted on Transition Between Low Power C-states

Problem: On entry to or exit from package C6 states, DPRSLPVR (Deeper Sleep Voltage Regulator) signal may be incorrectly asserted.

Implication: Due to this erratum, platform voltage regulator may shutdown

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.

AAP92. 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.

AAP93. 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.



AAP94. <Erratum Removed>

AAP95. 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 4 KBytes 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.

AAP96. 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.

AAP97. 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.

AAP98. 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.

AAP99. 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.

AAP100. 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.

AAP101. 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.

**AAP102. 2MB Page Split Lock Accesses Combined with Complex Internal Events May Cause Unpredictable System Behavior**

Problem: A 2-MB Page Split Lock (a locked access that spans two 2-MB large pages) coincident with additional requests that have particular address relationships in combination with a timing sensitive sequence of complex internal conditions may cause unpredictable system behavior.

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

Workaround: None identified.

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

AAP103. 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.

AAP104. The TPM's Locality 1 Address Space Cannot Be Opened

Problem: Due to this erratum, writing to TXT.CMD.OPEN.LOCALITY1 (FED2_0380H) does not open the Locality 1 address space to the TPM (Trusted Platform Module).

Implication: Software that uses the TPM's Locality 1 address space will not be able to gain access to it.

Workaround: All operations for the TPM should be done using Locality 0 or Locality 2 instead of Locality 1.

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

AAP105. PCIe* Link Bit Errors Present during L0s Entry May Cause the System to Hang during L0s Exit

Problem: During L0s entry PCIe* link bit errors may be generated due to a slow shutdown response from the PCIe analog circuits. As a result, the PCIe analog circuits may now take longer to establish bit lock during the L0s exit sequence. In some cases bit lock may not be achieved and may result in a system hang.

Implication: While exiting from L0s the PCIe* bus may go into recovery mode. At the 5 GB/s rate system hangs may occur while exiting from L0s; however the hangs have not been seen on commercially available systems.

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.

**AAP106. The Combination of a Page-Split Lock Access and Data Accesses That Are Split across Cacheline Boundaries May Lead to Processor Livelock**

Problem: Under certain complex micro-architectural conditions, the simultaneous occurrence of a page-split lock and several data accesses that are split across cacheline boundaries may lead to processor livelock.

Implication: Due to this erratum, a livelock may occur that can only be terminated by a processor reset. 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.

AAP107. FP Data Operand Pointer May Be Incorrectly Calculated after an FP Access Which Wraps a 4-Gbyte Boundary in Code That Uses 32-Bit Address Size in 64-bit Mode

Problem: The FP (Floating Point) Data Operand Pointer is the effective address of the operand associated with the last non-control FP instruction executed by the processor. If an 80-bit FP access (load or store) uses a 32-bit address size in 64-bit mode and the memory access wraps a 4-Gbyte boundary and the FP environment is subsequently saved, the value contained in the FP Data Operand Pointer may be incorrect.

Implication: Due to this erratum, the FP Data Operand Pointer may be incorrect. Wrapping an 80-bit FP load around a 4-Gbyte boundary in this way is not a normal programming practice. Intel has not observed this erratum with any commercially available software.

Workaround:If the FP Data Operand Pointer is used in a 64-bit operating system which may run code accessing 32-bit addresses, care must be taken to ensure that no 80-bit FP accesses are wrapped around a 4-Gbyte boundary.

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

AAP108. IOTLB Invalidations Not Completing on Intel® VT-d Engine for Integrated High Definition Audio

Problem: IOTLB invalidation in the Intel® VT-d engine for integrated High Definition Audio device may not complete and cause IVT field, bit [63] of IOTLBINV register (Offset 0x1208 in Memory Mapped IO region described by VTBAR {device 8, function 0, offset 0x180}), to not be cleared as expected. As a result, software may continue to poll this bit and not detect successful invalidation completion.

Implication: When Intel® VT-d engine for integrated High Definition Audio device is enabled and software requests for IOTLB invalidation while audio traffic is active, the request may not complete and may result in a software hang. Intel has not observed this erratum with any commercially available software.

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.

AAP109. IO_SMI Indication in SMRAM State Save Area May Be Lost

Problem: The IO_SMI bit (bit 0) in the IO state field at SMRAM offset 7FA4H is set to "1" by the processor to indicate a System Management Interrupt (SMI) is either taken immediately after a successful I/O instruction or is taken after a successful iteration of a REP I/O instruction. Due to this erratum, the setting of the IO_SMI bit may be lost. This may happen under a complex set of internal conditions with Intel® Hyper-Threading Technology enabled and has not been observed with commercially available software.

Implication: Due to this erratum, SMI handlers may not be able to identify the occurrence of I/O SMIs.

Workaround:None identified.

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

**AAP110. PCIe* Squelch Detect May be Slow to Respond During L0s Entry and May Cause a Surprise Link Down Condition**

Problem: When entering the L0s idle state the PCIe* squelch detect response may be slower than expected. This slow response can cause the PCIe* interface at the downstream port to unexpectedly enter the L0s.FTS (Fast Training Sequence) state instead of the normal operation which is staying in the L0s.idle state until the Tx side of the upstream port exits squelch. This unexpected state transition may cause a recovery entry leading to a Surprise Link Down condition.

Implication: This erratum may cause a system hang while trying reach the L0s state.

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.

AAP111. TR Corruption Due to Save/Restore x87 FPU Pointers in SMRAM

Problem: If x87 FPU instruction and data pointers are saved in SMRAM, the TR (Task Register) selector may be restored incorrectly on the exit from SMM.

Implication: The TR selector containing incorrect data 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.

AAP112. PCIe* Lanes Returning to The Active Power State May Cause The System to Hang

Problem: Under certain conditions, when the PCIe lanes come out of the S0 power savings state, the clocks may change asynchronously leading to a system hang.

Implication: A System hang may occur when coming out of the S0 power saving state.

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.

AAP113. Performance Monitor Events for Hardware Prefetches Which Miss The L1 Data Cache May be Over Counted

Problem: Hardware prefetches that miss the L1 data cache but cannot be processed immediately due to resource conflicts will count and then retry. This may lead to incorrectly incrementing the L1D_PREFETCH.MISS (event 4EH, umask 02H) event multiple times for a single miss.

Implication: The count reported by the L1D_PREFETCH.MISS event may be higher than expected.

Workaround: None identified.

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

**AAP114. Poisoned Write Caused by an Internal Parity Error Targeting IIO PCI Configuration Registers or MMIO Space will Not be Suppressed**

Problem: When due to an internal parity error, a processor attempts to write poisoned data to a PCI configuration register in the IIO (Integrated I/O) module (internal PCI devices on bus IIOBUSNO) or to the MMIO space decoded by a BAR in the IIO module, the poisoned data will not be dropped. However, even though the poisoned data will not be dropped the internal Intel® QuickPath Interconnect logic will log and report an error in the IA32_MCO_STATUS MSR (401H) with MCACOD equal to 0000 1110 xxxx xx11 and bit 16 or 17 set.

Implication: Poisoned data may be written to PCI configuration registers or MMIO space causing a machine check exception. It is possible for these writes to lead to unpredictable system behavior.

Workaround: None identified.

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

AAP115. VM Exit May Incorrectly Clear IA32_PERF_GLOBAL_CTRL [34:32]

Problem: If the "load IA32_PERF_GLOBAL_CTRL" VM-exit control is 1, a VM exit should load the IA32_PERF_GLOBAL_CTRL MSR (38FH) from the IA32_PERF_GLOBAL_CTRL field in the guest-state area of the VMCS. Due to this erratum, such a VM exit may instead clear bits 34:32 of the MSR, loading only bits 31:0 from the VMCS.

Implication: All fixed-function performance counters will be disabled after an affected VM exit, even if the VM exit should have enabled them based on the IA32_PERF_GLOBAL_CTRL field in the guest-state area of the VMCS.

Workaround: A VM monitor that wants the fixed-function performance counters to be enabled after a VM exit may do one of two things: (1) clear the "load IA32_PERF_GLOBAL_CTRL" VM-exit control; or (2) include an entry for the IA32_PERF_GLOBAL_CTRL MSR in the VM-exit MSR-load list.

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

AAP116. PCIe* Port's LTSSM May Not Transition Properly in the Presence of TS1 or TS2 Ordered Sets That Have Unexpected Symbols Within those Sets

Problem: When a PCIe* port receives TS1 and/or TS2 ordered sets with unexpected symbols (per the PCIe* Base Specification), the port's LTSSM (Link Training State Machine) might not transition according to the PCIe* Base Specification requirements. The LTSSM may incorrectly stay in its current state, or transition to an incorrect state. If the unexpected symbols are sporadic in nature the link will recover and go to the proper state.

Implication: PCIe* Port's LTSSM may not transition according to PCIe* Base Specification as described above. This problem has not been seen in real system testing, but was discovered by synthetic tests designed to check for illegal conditions.

Workaround: None identified.

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



AAP117. NTB/RP Link Will Send Extra TS2 Ordered Set During Link Training

Problem: The NTB (Non-Transparent Bridge) when operating in NTB/RP (Root Port) mode will send a superfluous TS2 ordered set after transitioning to the CONFIGURATION.IDLE state during link training. This TS2 ordered set may contain invalid capability data.

Implication: NTB/RP Link will transmit a TS2 ordered set after transitioning to the CONFIGURATION.IDLE state. No impact expected for specification compliant PCIe partners. Specification compliant PCIe link partners will have transitioned to CONFIGURATION.IDLE before this ordered set is sent and will ignore it.

Workaround: None identified.

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

AAP118. PCIe* Ports May Not Enter Slave Loopback Mode From the Configuration LTSSM State

Problem: If a PCIe* port's LTSSM (Link Training State Machine) is in the CONFIG.LINK_WIDTH_START state, it may not enter slave loopback mode when requested to do so by the link partner. If the request is missed the link will continue to train and enter the Slave loopback mode after it first transitions through the L0 and RECOVERY LTSSM states.

Implication: Due to this erratum, PCIe* ports may be delayed in entering the slave loopback mode.

Workaround: None identified.

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

AAP119. Unexpected DMI and PCIe* Link Retraining and Correctable Errors Reported

Problem: When the processor exits the package C6 power state, the PCIe* and DMI ports may enter a state where they will NAK all packets for a short time. If this condition persists long enough so that the same packet is NAKed four times, the link will retrain and a correctable error may be signaled by the PCIe end point. Overall performance of the link is not impacted.

Implication: Due to this erratum, unexpected link retraining and correctable errors may be reported.

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.

AAP120. QPI Lane May Be Dropped During Full Frequency Deskew Phase of Training

Problem: A random QPI Lane may be dropped during the lane deskew phase while the QPI Bus is training at full frequency.

Implication: When there are multiple resets after the QPI Bus has reached full speed operation there is a small chance that a lane could be dropped during the deskew phase of training. In the case of a lane being dropped this will be detected and a retry will be done until the link is established and the lane is re-trained.

Workaround: None identified.

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

**AAP121. PerfMon Overflow Status Can Not be Cleared After Certain Conditions Have Occurred**

Problem: Under very specific timing conditions, if software tries to disable a PerfMon counter through MSR IA32_PERF_GLOBAL_CTRL (0x38F) or through the per-counter event-select (e.g. MSR 0x186) and the counter reached its overflow state very close to that time, then due to this erratum the overflow status indication in MSR IA32_PERF_GLOBAL_STAT (0x38E) may be left set with no way for software to clear it.

Implication: Due to this erratum, software may be unable to clear the PerfMon counter overflow status indication.

Workaround: Software may avoid this erratum by clearing the PerfMon counter value prior to disabling it and then clearing the overflow status indication bit.

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

AAP122. An Unexpected Page Fault or EPT Violation May Occur After Another Logical Processor Creates a Valid Translation for a Page

Problem: An unexpected page fault (#PF) or EPT violation may occur for a page under the following conditions:

- The paging structures initially specify no valid translation for the page.
- Software on one logical processor modifies the paging structures so that there is a valid translation for the page (e.g., by setting to 1 the present bit in one of the paging-structure entries used to translate the page).
- Software on another logical processor observes this modification (e.g., by accessing a linear address on the page or by reading the modified paging-structure entry and seeing value 1 for the present bit).
- Shortly thereafter, software on that other logical processor performs a store to a linear address on the page.

In this case, the store may cause a page fault or EPT violation that indicates that there is no translation for the page (e.g., with bit 0 clear in the page-fault error code, indicating that the fault was caused by a not-present page). Intel has not observed this erratum with any commercially available software.

Implication: An unexpected page fault may be reported. There are no other side effects due to this erratum.

Workaround: System software can be constructed to tolerate these unexpected page faults. See Section "Propagation of Paging-Structure Changes to Multiple Processors" of Volume 3A of IA-32 Intel® Architecture Software Developer's Manual, for recommendations for software treatment of asynchronous paging-structure updates.

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

AAP123. L1 Data Cache Errors May be Logged With Level Set to 1 Instead of 0

Problem: When an L1 Data Cache error is logged in IA32_MCI_STATUS[15:0], which is the MCA Error Code Field, with a cache error type of the format 0000 0001 RRRR TTLL, the LL field may be incorrectly encoded as 01b instead of 00b.

Implication: An error in the L1 Data Cache may report the same LL value as the L2 Cache. Software should not assume that an LL value of 01b is the L2 Cache.

Workaround: None identified.

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



AAP124. Stack Pushes May Not Occur Properly for Events Delivered Immediately After VM Entry to 16-Bit Software

Problem: The stack pushes for an event delivered after VM entry and before execution of an instruction in VMX non-root operation may not occur properly. The erratum applies only if the VM entry establishes IA32_EFER.LMA = 0 and CS.D = 0 and only if the event handler is also invoked with CS.D = 0.

Implication: This erratum affects events that are pending upon completion of VM entry and that do not cause VM exits. Examples include debug exceptions, interrupts, and general-protection faults generated in virtual-8086 mode by the mode's virtual interrupt mechanism. The erratum applies only if the VM entry is not to IA-32e mode and is to 16-bit operation, and only if the relevant handler uses 16-bit operation. The incorrect stack pushes resulting from the erratum may cause incorrect guest operation. 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.

AAP125. Executing The GETSEC Instruction While Throttling May Result in a Processor Hang

Problem: If the processor throttles due to either high temperature thermal conditions or due to an explicit operating system throttling request (TT1) while executing GETSEC[SENTER] or GETSEC[SEXIT] instructions, then under certain circumstances, the processor may hang. Intel has not been observed this erratum with any commercially available software.

Implication: Possible hang during execution of GETSEC instruction.

Workaround: None identified.

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

AAP126. PerfMon Event LOAD_HIT_PRE.SW_PREFETCH May Overcount

Problem: PerfMon event LOAD_HIT_PRE.SW_PREFETCH (event 4CH, umask 01H) should count load instructions hitting an ongoing software cache fill request initiated by a preceding software prefetch instruction. Due to this erratum, this event may also count when there is a preceding ongoing cache fill request initiated by a locking instruction.

Implication: PerfMon event LOAD_HIT_PRE.SW_PREFETCH may overcount.

Workaround: None identified.

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

AAP127. Successive Fixed Counter Overflows May be Discarded

Problem: Under specific internal conditions, when using Freeze PerfMon on PMI feature (bit 12 in IA32_DEBUGCTL.Freeze_PerfMon_on_PMI, MSR 1D9H), if two or more PerfMon Fixed Counters overflow very closely to each other, the overflow may be mishandled for some of them. This means that the counter's overflow status bit (in MSR_PERF_GLOBAL_STATUS, MSR 38EH) may not be updated properly; additionally, PMI interrupt may be missed if software programs a counter in Sampling-Mode (PMI bit is set on counter configuration).

Implication: Successive Fixed Counter overflows may be discarded when Freeze PerfMon on PMI is used.

Workaround: Software can avoid this by:

- Avoid using Freeze PerfMon on PMI bit
- Enable only one fixed counter at a time when using Freeze PerfMon on PMI

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

**AAP128. #GP May be Signaled When Invalid VEX Prefix Precedes Conditional Branch Instructions**

Problem: When a 2-byte opcode of conditional branch (opcodes 0F8xH, for any value of x) instruction resides in 16-bit code-segment and is associated with invalid VEX prefix, it may sometimes signal a #GP fault (illegal instruction length > 15-bytes) instead of #UD (illegal opcode).

Implication: #GP fault instead of a #UD signaled on an illegal instruction.

Workaround: None identified.

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

AAP129. A Logical Processor May Wake From Shutdown State When Branch-Trace Messages or Branch-Trace Stores Are Enabled

Problem: Normally, a logical processor that entered the shutdown state will remain in that state until a break event (NMI, SMI, INIT) occurs. Due to this erratum, if CR4.MCE (Machine Check Enable) is 0 and a branch-trace message or branch-trace store is pending at the time of a machine check, the processor may not remain in shutdown state. In addition, if the processor was in VMX non-root operation when it improperly woke from shutdown state, a subsequent VM exit may save a value of 2 into the activity-state field in the VMCS (indicating shutdown) even though the VM exit did not occur while in shutdown state.

Implication: This erratum may result in unexpected system behavior. If a VM exit saved a value of 2 into the activity-state field in the VMCS, the next VM entry will take the processor to shutdown state.

Workaround: Software should ensure that CR4.MCE is set whenever IA32_DEBUGCTL MSR (60EH) TR bit [6] is set.

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

AAP130. Task Switch to a TSS With an Inaccessible LDTR Descriptor May Cause Unexpected Faults

Problem: A task switch may load the LDTR (Local Descriptor Table Register) with an incorrect segment descriptor if the LDT (Local Descriptor Table) segment selector in the new TSS specifies an inaccessible location in the GDT (Global Descriptor Table).

Implication: Future accesses to the LDT may result in unpredictable system behavior.

Workaround: Operating system code should ensure that segment selectors used during task switches to the GDT specify offsets within the limit of the GDT and that the GDT is fully paged into memory.

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

AAP131. VM Entries That Return From SMM Using VMLAUNCH May Not Update The Launch State of the VMCS

Problem: Successful VM entries using the VMLAUNCH instruction should set the launch state of the VMCS to "launched". Due to this erratum, such a VM entry may not update the launch state of the current VMCS if the VM entry is returning from SMM.

Implication: Subsequent VM entries using the VMRESUME instruction with this VMCS will fail. RFLAGS.ZF is set to 1 and the value 5 (indicating VMRESUME with non-launched VMCS) is stored in the VM-instruction error field. This erratum applies only if dual monitor treatment of SMI and SMM is active.

Workaround: None identified.

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



AAP132. VM Entry May Clear Bytes 81H-83H on Virtual-APIC Page When "Use TPR Shadow" Is 0

Problem: VM entry should not clear bytes 81H-83H on the virtual-APIC page if the "use TPR shadow" VM-execution control is 0. Due to this erratum, VM entry will do so if the "virtualize x2APIC mode" VM-execution control is 1.

Implication: VM entries with the 0-setting of the "use TPR shadow" VM-execution control and the 1-setting of the "virtualize x2APIC mode" VM-execution control cause any non-zero data at bytes 81H-83H on the virtual-APIC page to be lost. Note that this combination of settings is not allowed; any such VM entry will fail after clearing these bytes.

Workaround: Software should always set the "use TPR shadow" VM-execution control to 1 whenever it sets that "virtualize x2APIC mode" VM-execution control to 1.

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

AAP133. A First Level Data Cache Parity Error May Result in Unexpected Behavior

Problem: When a load occurs to a first level data cache line resulting in a parity error in close proximity to other software accesses to the same cache line and other locked accesses the processor may exhibit unexpected behavior.

Implication: Due to this erratum unpredictable system behavior may occur. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

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

AAP134. Intel® Trusted Execution Technology ACM Revocation

Problem: SINIT ACM i7_QUAD_SINIT_20.BIN or earlier are revoked and will not launch with new processor configuration information.

Implication: Due to this erratum, SINIT ACM i7_QUAD_SINIT_20.BIN and earlier will be revoked.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. All Intel® TXT enabled software must use SINIT ACM i7_QUAD_SINIT_51.BIN or later.

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

AAP135. An Event May Intervene Before a System Management Interrupt That Results from IN or INS

Problem: If an I/O instruction (IN, INS, OUT, or OUTS) results in an SMI (system-management interrupt), the processor will set the IO_SMI bit at offset 7FA4H in SMRAM. This interrupt should be delivered immediately after execution of the I/O instruction so that the software handling the SMI can cause the I/O instruction to be re-executed. Due to this erratum, it is possible for another event (e.g., a nonmaskable interrupt) to be delivered before the SMI that follows the execution of an IN or INS instruction.

Implication: If software handling an affected SMI uses I/O instruction restart, the handler for the intervening event will not be executed.

Workaround: The SMM handler has to evaluate the saved context to determine if the SMI was triggered by an instruction that read from an I/O port. The SMM handler must not restart an I/O instruction if the platform has not been configured to generate a synchronous SMI for the recorded I/O port address.

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

**AAP136. The Corrected Error Count Overflow Bit in IA32_MCO_STATUS is Not Updated When the UC Bit is Set**

Problem: After a UC (uncorrected) error is logged in the IA32_MCO_STATUS MSR (401H), corrected errors will continue to be counted in the lower 14 bits (bits 51:38) of the Corrected Error Count. Due to this erratum, the sticky count overflow bit (bit 52) of the Corrected Error Count will not get updated when the UC bit (bit 61) is set to 1.

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

Workaround: None identified.

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

AAP137. The Upper 32 Bits of CR3 May be Incorrectly Used With 32-Bit Paging

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

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

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

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

AAP138. EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly

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

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

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

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

AAP139. SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior

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

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

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

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



AAP140. Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash

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

Implication: This erratum may result in unexpected faults, an uncorrectable TLB error logged in IA32_MCI_STATUS.MCACOD (bits [15:0]) with a value of 0000_0000_0001_xxxx (where x stands for 0 or 1), a guest or hypervisor crash, or other unpredictable system behavior.

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

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

§ §



Specification Changes

AAP1. Update to Datasheet - Volume 2 to Uncore Revision Identification Register

Issue: The Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2 Section 4.4.3 will be updated in red text below.

Affected Docs: Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2

| Device: 0 Function: 0-1 Offset:08h | | | |
|---------------------------------------------|------|---------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Device: 2 Function: 0 Offset:08h | | | |
| Device: 3 Function: 0-1, 4 Offset:08h | | | |
| Device: 4-5 Function: 0-3 Offset:08h | | | |
| Bit | Attr | Default | Description |
| 7:4 | RO | 0h | RID Major Steppings which required all masks be regenerated. B1 stepping: 0h |
| 3:0 | RO | 4h | RID Minor Revision Identification Number Increment for each steppings which don't require all masks to be regenerated. B1 stepping: 4h |



AAP2. Update to Datasheet - Volume 2 to PCI Express Device Control Register 2

Issue: The Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2 Section 3.3.4.31 will be updated in red text below.

Affected Docs: Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2

| Register: DEVCTRL2 Device: 0 (DMI), 3, 5 (PCIe) Function: 0 Offset: B8h | | | |
|--------------------------------------------------------------------------------------------------------------|------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit | Attr | Default | Description |
| 15:6 | RO | 0h | Reserved |
| 5 | RW | 0 | Alternative RID Interpretation (ARI) Enable When set to 1b, ARI is enabled for the Root Port. |
| 4 | RW | 0 | Completion Timeout Disable When set to 1b, this bit disables the Completion Timeout mechanism for all NP tx that I/O issues on the PCIe/DMI link. When 0b, completion timeout is enabled. Software can change this field while there is active traffic in the root port. |
| 3:0 | RW | 0000b | Completion Timeout Value on NP Tx that Integrated I/O Issues on PCIe/DMI – In Devices that support Completion Timeout programmability, this field allows system software to modify the Completion Timeout range. The following encodings and corresponding timeout ranges are defined: 0000b = 2 ms 0001b = Reserved (Integrated I/O aliases to 0000b) 0010b = Reserved (Integrated I/O aliases to 0000b) 0101b = 4 ms 0110b = 10 ms 1001b = 40 ms 1010b = 210 ms 1101b = 800 ms 1110b = 2 s - 6 s When OS selects 2 s to 6 s range, the CTOCTRL register further controls the timeout value within that range. For all other ranges selected by OS, the timeout value within that range is fixed in Integrated I/O hardware. Software can change this field while there is active traffic in the root port. |

**AAP3. Update to Datasheet - Volume 2 to Completion Timeout Control Register**

Issue: *x=The Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2* Section 3.3.5.8 will be updated in red text below.

Affected Docs: *Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 2*

| Register:CTOCTRL Device:0 (DMI), 3, 5 (PCIe) Function:0 Offset:1E0h | | | |
|----------------------------------------------------------------------------------------------------------|------|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Bit | Attr | Default | Description |
| 31:10 | RV | 00 | <i>Reserved</i> |
| 9:8 | RW | 00 | XP-to-PCIe Timeout Select within 2 s to 6 s Range When OS selects a timeout range of 2 s to 6 s for Windows* XP (that affect NP tx issued to the PCIe/DMI) using the root port's DEVCTRL2 register, this field selects the sub-range within that larger range, for additional controllability. 00: 2 s 01: 4 s 10: 6 s 11: Reserved Note: this field is subject to redefinition based on design feedback |
| 7:0 | RV | 00 | <i>Reserved</i> |



AAP4. Update to Datasheet - Volume 1 to Table 35 and Table 41

Issue: The Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 1 Table 35 in Section 7.6 and Table 41 in Section 7.10.1 will be updated in red text below.

Affected Docs: Intel® Core™ i7-900 Mobile Processor Extreme Edition Series, Intel® Core™ i7-800 and i7-700 Mobile Processor Series Datasheet - Volume 1

Table 35

| | | | |
|--------------|------|------------|-----------------------------------------|
| Single Ended | (qa) | CMOS Input | PM_EXT_TS#[0], PM_EXT_TS#[1], CFG[17:0] |
| Single Ended | (qb) | CMOS Input | RSTIN# |
| Single Ended | (qc) | CMOS Input | PM_SYNC |

Table 41

| Symbol | Alpha Group | Parameter | Min | Typ | Max | Units | Notes ^{1,8} |
|-----------------|--------------------------------------|--------------------|------------------------|-----|------------------------|-------|----------------------|
| V _{IL} | (m), (n), (p), (qa), (qb), (s) | Input Low Voltage | | | 0.64 * V _{TT} | V | 2,3 |
| V _{IH} | (m), (n), (p), (qa), (qb), (qc), (s) | Input High Voltage | 0.76 * V _{TT} | | | V | 2,3,5 |
| V _{IL} | (g), (qc) | Input Low Voltage | | | 0.40 * V _{TT} | V | 2,3 |

§ §



Specification Clarifications

There are no, new Specification Clarifications in this Specification Update revision.

§ §



Documentation Changes

AAP1. On-Demand Clock Modulation Feature Clarification

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

For Hyper-Threading Technology enabled processors, the IA32_CLOCK_MODULATION register is duplicated for each logical processor. In order for the On-demand clock modulation feature to work properly, the feature must be enabled on all the logical processors within a physical processor. If the programmed duty cycle is not identical for all the logical processors, the processor clock will modulate to the highest duty cycle programmed for processors if the CPUID DisplayFamily_DisplayModel signatures is listed in Table 14-2. For all other processors, if the programmed duty cycle is not identical for all logical processors in the same core, the processor will modulate at the lowest programmed duty cycle.

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

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

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

| DisplayFamily_Display Model | DisplayFamily_Display Model | DisplayFamily_Display Model | DisplayFamily_Display Model |
|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| 0F_xx | 06_1C | 06_1A | 06_1E |
| 06_1F | 06_25 | 06_26 | 06_27 |
| 06_2C | 06_2E | 06_2F | 06_35 |
| 06_36 | — | — | — |

§ §



Documentation Changes