

Dual-Core Intel[®] Xeon[®] Processor 3000 Series

Specification Update

February 2008

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

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

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

Enabling Execute Disable Bit functionality requires a PC with a processor with Execute Disable Bit capability and a supporting operating system. Check with your PC manufacturer on whether your system delivers Execute Disable Bit functionality.

The Dual-Core Intel® Xeon® Processor may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

Φ Intel® 64 Architecture requires a computer system with a processor, chipset, BIOS, operating system, device drivers, and applications enabled for Intel 64. Processor will not operate (including 32-bit operation) without an Intel 64-enabled BIOS. Performance will vary depending on your hardware and software configurations. See <http://www.intel.com/technology/intel64/> for more information including details on which processors support Intel 64, or consult with your system vendor for more information.

± Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations. Intel Virtualization Technology-enabled BIOS and VMM applications are currently in development.

Intel, Celeron, Pentium, Xeon, Intel SpeedStep, Intel Core, Core Inside 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 © 2007, Intel Corporation


Contents

Preface	6
Summary Tables of Changes	8
Identification Information	17
Component Identification Information.....	19
Errata	21
Specification Changes	72
Specification Clarifications	73
Documentation Changes	74

Revision History

Revision	Date	Description
001	September 2006	Initial release of the Dual-Core Intel® Xeon® Processor 3000 Series Specification Update
002	November 2006	<ul style="list-style-type: none"> • Changed references from "AI" to "AP" • Updated Status for Erratum AP55 in Errata Table • Replaced Erratum AP10 with a new erratum • Updated Errata AP46, AP53, AP72 • Added Errata AP78-AP85 • Added Specification Change AP1 • Corrected Plan information in Summary Table of Changes for Errata AP16, AP69, AP70, AP72 and AP75 • Updated Processor Identification information
003	December 2006	<ul style="list-style-type: none"> • Replaced Erratum AP61 with a new erratum • Updated Erratum AP75 and AP83 • Added Errata AP86 – AP90 • Corrected Plan Information in Summary Table of Changes for Erratum AP8 • Reset numbering for Notes applicable to Table 1, 2 in Component Identification Section
004	January 2007	<ul style="list-style-type: none"> • Added Errata AP91 - AP94
005	February 2007	<ul style="list-style-type: none"> • Updated Erratum AP70 • Added Errata AP95-AP97
006	April 2007	<ul style="list-style-type: none"> • Added Errata AP98-AP104 • Updated Erratum AP33 in Summary Table of Changes • Added L step info
007	April 2007 out of cycle	<ul style="list-style-type: none"> • Added Erratum AP105 • Added Specification Clarification AP1
008	May 2007	<ul style="list-style-type: none"> • Updated Errata AP14, AP25 and AP26 • Updated Component Identification Information
009	July 2007	<ul style="list-style-type: none"> • Added Erratum AP106 to AP111 • Included G0 stepping information (updated summary tables of change and updated processor identification information)
010	August 2007	<ul style="list-style-type: none"> • Added Erratum AP112 • Added new S-Spec numbers per PCN107455

011	October 2007	<ul style="list-style-type: none"> • Added Errata AP113 - AP123 • Updated Plan Status for AP6, AP20, AP21, AP22, AP23, AP31, AP38 AP39, AP40, AP41, AP42, AP44, AP50, AP55, AP56, AP57 AP61, AP66, AP69, AP70, AP72, AP75 AP79, AP91, AP94, AP101, AP102, AP109
012	November 2007	<ul style="list-style-type: none"> • Added Erratum AP124 • Deleted typo "M0 Column" from Errata Table
013	December 2007	<ul style="list-style-type: none"> • Updated Erratum AP8 • Updated Plan Status for AP121, AP122, AP123 • Added Erratum AP125
014	January 2008	<ul style="list-style-type: none"> • Added Erratum AP126
015	February 2008	<ul style="list-style-type: none"> • Updated Erratum AP48 • Updated Erratum AP51 • Deleted Erratum AP123 (because it is repeat of AP108) and replaced with a new Erratum



Preface

This document is an update to the specifications contained in the documents listed in the following Affected Documents/Related Documents table. It is a compilation of device and document errata and specification clarifications and changes, and is intended for hardware system manufacturers and for software developers of applications, operating system, and tools.

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents. This document may also contain information that has not been previously published.

Affected Documents

Document Title	Document Number
<i>Dual-Core Intel® Xeon® Processor 3000 Series Datasheet</i>	314915

Related Documents

Document Title	Document Location
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual Volume 1: Basic Architecture</i>	http://www.intel.com/products/processor/manuals/index.htm
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual Volume 2A: Instruction Set Reference Manual A-M</i>	
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual Volume 2B: Instruction Set Reference Manual, N-Z</i>	
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual Volume 3A: System Programming Guide</i>	
<i>Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual Volume 3B: System Programming Guide</i>	

Nomenclature

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

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

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

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 the next release of the specifications.

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

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 table indicates the Specification Changes, Errata, Specification Clarifications or Documentation Changes, which apply to the listed MCH steppings. Intel intends to fix some of the errata in a future stepping of the component, and to account for the other outstanding issues through documentation or Specification Changes as noted. This table uses the following notations:

Codes Used in Summary Table

Stepping

X:	Erratum, 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.

Status

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

Row

Shaded:	This item is either new or modified from the previous version of the document.
---------	--

Item Numbering

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

A =	Dual-Core Intel® Xeon® processor 7000 sequence
C =	Intel® Celeron® processor
D =	Dual-Core Intel® Xeon® processor 2.80 GHz
E =	Intel® Pentium® III processor
F =	Intel® Pentium® processor Extreme Edition and Intel® Pentium® D processor
I =	Dual-Core Intel® Xeon® processor 5000 series
J =	64-bit Intel® Xeon® processor MP with 1MB L2 cache
K =	Mobile Intel® Pentium® III processor
L =	Intel® Celeron® D processor
M =	Mobile Intel® Celeron® processor
N =	Intel® Pentium® 4 processor
O =	Intel® Xeon® processor MP
P =	Intel® Xeon® processor
Q =	Mobile Intel® Pentium® 4 processor supporting Hyper-Threading technology on 90-nm process technology
R =	Intel® Pentium® 4 processor on 90 nm process
S =	64-bit Intel® Xeon® processor with 800 MHz system bus (1 MB and 2 MB L2 cache versions)
T =	Mobile Intel® Pentium® 4 processor-M
U =	64-bit Intel® Xeon® processor MP with up to 8MB L3 cache
V =	Mobile Intel® Celeron® processor on .13 micron process in Micro-FCPGA package
W =	Intel® Celeron® M processor
X =	Intel® Pentium® M processor on 90nm process with 2-MB L2 cache and Intel® processor A100 and A110 with 512-KB L2 cache
Y =	Intel® Pentium® M processor
Z =	Mobile Intel® Pentium® 4 processor with 533 MHz system bus
AA =	Intel® Pentium® D Processor 900 sequence and Intel® Pentium® processor Extreme Edition 955, 965
AB =	Intel® Pentium® 4 Processor 6x1 sequence
AC =	Intel(R) Celeron(R) processor in 478 pin package
AD =	Intel(R) Celeron(R) D processor on 65nm process
AE =	Intel® Core™ Duo processor and Intel® Core™ Solo processor on 65nm process
AF =	Dual-Core Intel® Xeon® processor LV
AG =	Dual-Core Intel® Xeon® processor 5100 series
AH =	Intel® Core™2 Duo/Solo Processor for Intel® Centrino® Duo processor technology
AI =	Intel® Core™2 Extreme processor X6800 and Intel® Core™2 Duo desktop processor E6000 and E4000 sequence
AJ =	Quad-Core Intel® Xeon® processor 5300 series

Summary Tables of Changes

AK =	Intel® Core™2 Extreme quad-core processor QX6000 sequence and Intel® Core™2 Quad processor Q6000Δ sequence
AL =	Dual-Core Intel® Xeon® processor 7100 series
AM =	Intel® Celeron® processor 400 sequence
AN =	Intel® Pentium® dual-core processor
AO =	Quad-Core Intel® Xeon® processor 3200 series
AP =	Dual-Core Intel® Xeon® processor 3000 series
AQ =	Intel® Pentium® dual-core desktop processor E2000Δ sequence
AR =	Intel® Celeron processor 500 series
AS =	Intel® Xeon® processor 7200, 7300 series
AV =	Intel® Core™2 Extreme processor QX9000 series and Intel® Core™2 Quad processor Q9000 series
AX =	Quad-Core Intel® Xeon® processor 5400 series
AW =	Intel® Core™ 2 Duo desktop processor E8000 series
AY =	Dual-Core Intel® Xeon® processor 5200 series
AAA =	Quad-Core Intel® Xeon® processor 3300 series
AAB =	Dual-Core Intel® Xeon® E3110 processor
AAC =	Intel® Celeron® dual-core processor E1000 series

The Specification Updates for the Pentium® processor, Pentium® Pro processor, and other Intel products do not use this convention

NO	B1	B2	L2	GO	Plan	ERRATA
AP1	X	X	X	X	No Fix	Writing the Local Vector Table (LVT) when an Interrupt is Pending May Cause an Unexpected Interrupt
AP2	X	X	X	X	No Fix	LOCK# Asserted During a Special Cycle Shutdown Transaction May Unexpectedly De-assert
AP3	X	X	X	X	No Fix	Address Reported by Machine-Check Architecture (MCA) on Single-bit L2 ECC Errors May be Incorrect
AP4	X	X	X	X	No Fix	VERW/VERR/LSL/LAR Instructions May Unexpectedly Update the Last Exception Record (LER) MSR
AP5	X	X	X	X	No Fix	DR3 Address Match on MOVD/MOVQ/MOVRTQ Memory Store Instruction May Incorrectly Increment Performance Monitoring Count for Saturating SIMD Instructions Retired (Event CFH)
AP6	X	X	X		Fixed	SYSRET May Incorrectly Clear RF (Resume Flag) in the RFLAGS Register
AP7	X	X	X	X	No Fix	General Protection Fault (#GP) for Instructions Greater than 15 Bytes May be Preempted
AP8	X	X	X	X	No Fix	Pending x87 FPU Exceptions (#MF) Following STI May Be Serviced Before Higher Priority Interrupts
AP9	X	X	X	X	No Fix	The Processor May Report a #TS Instead of a #GP Fault
AP10	X	X	X	X	No Fix	Single Step Interrupts with Floating Point Exception Pending May Be Mishandled
AP11	X	X	X	X	No Fix	A Write to an APIC Register Sometimes May Appear to Have Not Occurred

NO	B1	B2	L2	GO	Plan	ERRATA
AP12	X	X	X	X	No Fix	Programming the Digital Thermal Sensor (DTS) Threshold May Cause Unexpected Thermal Interrupts
AP13	X	X	X	X	No Fix	Count Value for Performance-Monitoring Counter PMH_PAGE_WALK May be Incorrect
AP14	X	X	X	X	No Fix	LER MSRs May be Incorrectly Updated
AP15	X	X	X	X	No Fix	Performance Monitoring Events for Retired Instructions (COH) May Not Be Accurate
AP16	X	X	X	X	No Fix	Performance Monitoring Event For Number Of Reference Cycles When The Processor Is Not Halted (3CH) Does Not Count According To The Specification
AP17	X	X	X	X	No Fix	Using 2M/4M Pages When A20M# Is Asserted May Result in Incorrect Address Translations
AP18	X	X	X	X	No Fix	Writing Shared Unaligned Data that Crosses a Cache Line without Proper Semaphores or Barriers May Expose a Memory Ordering Issue
AP19	X	X	X	X	No Fix	Code Segment limit violation may occur on 4 Gigabyte limit check
AP20	X	X	X		Fixed	FP Inexact-Result Exception Flag May Not Be Set
AP21	X	X	X		Fixed	Global Pages in the Data Translation Look-Aside Buffer (DTLB) May Not Be Flushed by RSM instruction before Restoring the Architectural State from SMRAM
AP22	X	X	X		Fixed	Sequential Code Fetch to Non-canonical Address May have Non-deterministic Results
AP23	X	X	X		Fixed	VMCALL to Activate Dual-monitor Treatment of SMIs and SMM Ignores Reserved Bit settings in VM-exit Control Field
AP24	X	X	X	X	No Fix	The PECI Controller Resets to the Idle State
AP25	X	X	X	X	No Fix	Some Bus Performance Monitoring Events May Not Count Local Events under Certain Conditions
AP26	X	X	X	X	No Fix	Premature Execution of a Load Operation Prior to Exception Handler Invocation
AP27	X	X	X	X	No Fix	General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit
AP28	X	X	X	X	No Fix	EIP May be Incorrect after Shutdown in IA-32e Mode
AP29	X	X	X	X	No Fix	#GP Fault is Not Generated on Writing IA32_MISC_ENABLE [34] When Execute Disable Bit is Not Supported
AP30	X	X			Fixed	(E)CX May Get Incorrectly Updated When Performing Fast String REP MOVS or Fast String REP STOS With Large Data Structures
AP31	X	X	X		Fixed	Performance Monitoring Events for Retired Loads (CBH) and Instructions Retired (COH) May Not Be Accurate
AP32	X	X	X	X	No Fix	Upper 32 bits of 'From' Address Reported through BTMs or BTSSs May be Incorrect
AP33	X	X	X		Plan Fix	Unsynchronized Cross-Modifying Code Operations Can Cause Unexpected Instruction Execution Results

Summary Tables of Changes

NO	B1	B2	L2	GO	Plan	ERRATA
AP34	X	X	X	X	No Fix	MSRs Actual Frequency Clock Count (IA32_APERF) or Maximum Frequency Clock Count (IA32_MPERF) May Contain Incorrect Data after a Machine Check Exception (MCE)
AP35	X	X	X	X	No Fix	Incorrect Address Computed For Last Byte of FXSAVE/FXRSTOR Image Leads to Partial Memory Update
AP36	X	X	X	X	No Fix	Split Locked Stores May not Trigger the Monitoring Hardware
AP37	X	X			Fixed	REP CMPS/SCAS Operations May Terminate Early in 64-bit Mode when RCX >= 0X100000000
AP38	X	X	X		Fixed	FXSAVE/FXRSTOR Instructions which Store to the End of the Segment and Cause a Wrap to a Misaligned Base Address (Alignment <= 0x10h) May Cause FPU Instruction or Operand Pointer Corruption
AP39	X	X	X		Fixed	Cache Data Access Request from One Core Hitting a Modified Line in the L1 Data Cache of the Other Core May Cause Unpredictable System Behavior
AP40	X	X	X		Fixed	PREFETCHH Instruction Execution under Some Conditions May Lead to Processor Livelock
AP41	X	X	X		Fixed	PREFETCHH Instructions May Not be Executed when Alignment Check (AC) is Enabled
AP42	X	X	X		Fixed	Upper 32 Bits of the FPU Data (Operand) Pointer in the FXSAVE Memory Image May Be Unexpectedly All 1's after FXSAVE
AP43	X	X	X		Plan Fix	Concurrent Multi-processor Writes to Non-dirty Page May Result in Unpredictable Behavior
AP44	X	X	X		Fixed	Performance Monitor IDLE_DURING_DIV (18h) Count May Not be Accurate
AP45	X	X	X	X	No Fix	Values for LBR/BTS/BTM will be Incorrect after an Exit from SMM
AP46	X	X	X	X	No Fix	Shutdown/VMX-Abort Condition May Disable Non-Bootstrap Processors
AP47	X	X			Fixed	SYSCALL Immediately after Changing EFLAGS.TF May Not Behave According to the New EFLAGS.TF
AP48	X	X	X	X	No Fix	Code Segment Limit/Canonical Faults on RSM May be Serviced before Higher Priority Interrupts/Exceptions
AP49	X	X	X	X	No Fix	VM Bit is Cleared on Second Fault Handled by Task Switch from Virtual-8086 (VM86)
AP50	X	X	X		Fixed	IA32_FMASK is Reset during an INIT
AP51	X	X	X	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
AP52	X	X	X	X	No Fix	Last Branch Records (LBR) Updates May be Incorrect after a Task Switch
AP53	X	X	X	X	No Fix	IO_SMI Indication in SMRAM State Save Area May Be Set Incorrectly
AP54	X	X	X	X	No Fix	INIT Does Not Clear Global Entries in the TLB
AP55	X	X	X		Fixed	Using Memory Type Aliasing with Memory Types WB/WT May Lead to Unpredictable Behavior

NO	B1	B2	L2	GO	Plan	ERRATA
AP56	X	X	X		Fixed	Update of Read/Write (R/W) or User/Supervisor (U/S) or Present (P) Bits without TLB Shutdown May Cause Unexpected Processor Behavior
AP57	X	X	X		Fixed	BTS Message May Be Lost When the STPCLK# Signal is Active
AP58	X	X	X	X	No Fix	CMPSB, LODSB, or SCASB in 64-bit Mode with Count Greater or Equal to 2 ⁴⁸ May Terminate Early
AP59	X	X	X	X	No Fix	REP MOVSS/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may use an Incorrect Data Size or Lead to Memory-Ordering Violations.
AP60	X	X	X	X	No Fix	MOV To/From Debug Registers Causes Debug Exception
AP61	X	X	X		Fixed	Debug Register May Contain Incorrect Information on a MOVSS or POPSS Instruction Followed by SYSRET
AP62	X	X	X	X	No Fix	EFLAGS Discrepancy on a Page Fault After a Multiprocessor TLB Shutdown
AP63	X	X	X	X	No Fix	LBR, BTS, BTM May Report a Wrong Address when an Exception/Interrupt Occurs in 64-bit Mode
AP64	X	X	X	X	No Fix	Returning to Real Mode from SMM with EFLAGS.VM Set May Result in Unpredictable System Behavior
AP65	X	X	X	X	No Fix	A Thermal Interrupt is Not Generated when the Current Temperature is Invalid
AP66	X	X	X		Fixed	VMLAUNCH/VMRESUME May Not Fail when VMCS is Programmed to Cause VM Exit to Return to a Different Mode
AP67	X	X	X	X	No Fix	IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception
AP68	X	X	X	X	No Fix	Performance Monitoring Event FP_ASSIST May Not be Accurate
AP69	X	X	X		Fixed	CPL-Qualified BTS May Report Incorrect Branch-From Instruction Address
AP70	X	X	X		Fixed	PEBS Does Not Always Differentiate Between CPL-Qualified Events
AP71	X	X	X	X	No Fix	PMI May Be Delayed to Next PEBS Event
AP72	X	X	X		Fixed	A PMI Due to a PEBS Buffer Overflow Will Not be Signaled Unless IA32_DEBUGCTL[12] is Set
AP73	X	X	X	X	No Fix	The BS Flag in DR6 May be Set for Non-Single-Step #DB Exception
AP74	X	X	X	X	No Fix	An Asynchronous MCE During a Far Transfer May Corrupt ESP
AP75	X	X	X		Fixed	In Single-Stepping on Branches Mode, the BS Bit in the Pending-Debug-Exceptions Field of the Guest State Area will be Incorrectly Set by VM Exit on a MOV to CR8 Instruction
AP76	X	X	X	X	No Fix	B0-B3 Bits in DR6 May Not be Properly Cleared After Code Breakpoint
AP77	X	X	X	X	No Fix	BTM/BTS Branch-From Instruction Address May be Incorrect for Software Interrupts
AP78	X	X	X	X	No Fix	Last Branch Records (LBR) Updates May be Incorrect After a Task Switch

Summary Tables of Changes

NO	B1	B2	L2	GO	Plan	ERRATA
AP79	X	X	X		Fixed	REP Store Instructions in a Specific Situation may cause the Processor to Hang
AP80	X	X	X	X	No Fix	Performance Monitoring Events for L1 and L2 Miss May Not be Accurate
AP81	X	X	X	X	No Fix	Store to WT Memory Data May be Seen in Wrong Order by Two Subsequent Loads
AP82	X	X	X	X	No Fix	A MOV Instruction from CR8 Register with 16 Bit Operand Size Will Leave Bits 63:16 of the Destination Register Unmodified
AP83	X	X	X	X	No Fix	Non-Temporal Data Store May be Observed in Wrong Program Order
AP84	X	X	X	X	No Fix	Performance Monitor SSE Retired Instructions May Return Incorrect Values
AP85	X	X	X	X	No Fix	Fault on ENTER Instruction May Result in Unexpected Values on Stack Frame
AP86	X	X			Fixed	CPUID Reports Architectural Performance Monitoring Version 2 is Supported, When Only Version 1 Capabilities are Available
AP87	X	X	X	X	No Fix	Unaligned Accesses to Paging Structures May Cause the Processor to Hang
AP88	X	X	X	X	No Fix	Microcode Updates Performed During VMX Non-root Operation Could Result in Unexpected Behavior
AP89	X	X	X	X	No Fix	INVLPG Operation for Large (2M/4M) Pages May be Incomplete under Certain Conditions
AP90	X	X	X	X	No Fix	Page Access Bit May be Set Prior to Signaling a Code Segment Limit Fault
AP91	X	X	X		Fixed	Update of Attribute Bits on Page Directories without Immediate TLB Shutdown May Cause Unexpected Processor Behavior
AP92	X	X	X		Plan Fix	Invalid Instructions May Lead to Unexpected Behavior
AP93	X	X	X	X	No Fix	EFLAGS, CR0, CR4 and the EXF4 Signal May be Incorrect after Shutdown
AP94	X	X	X		Fixed	Performance Monitoring Counter MACRO_INSTS.DECODED May Not Count Some Decoded Instructions
AP95	X	X	X		Fixed	The Stack Size May be Incorrect as a Result of VIP/VIF Check on SYSEXIT and SYSRET
AP96	X	X	X	X	No Fix	Performance Monitoring Event SIMD_UOP_TYPE_EXEC.MUL is Counted Incorrectly for PMULUDQ Instruction
AP97	X	X	X	X	No Fix	Storage of PEBS Record Delayed Following Execution of MOV SS or STI
AP98	X	X	X	X	No Fix	Store Ordering May be Incorrect between WC and WP Memory Types
AP99	X	X	X	X	No Fix	Updating Code Page Directory Attributes without TLB Invalidation May Result in Improper Handling of Code #PF
AP100	X	X	X		Fixed	Performance Monitoring Event CPU_CLK_UNHALTED.REF May Not Count Clock Cycles According to the Processors Operating Frequency
AP101			X		Fixed	(E)CX May Get Incorrectly Updated When Performing Fast String REP STOS With Large Data Structures

NO	B1	B2	L2	GO	Plan	ERRATA
AP102	X	X	X		Fixed	Performance Monitoring Event BR_INST_RETIRED May Count CPUID Instructions as Branches
AP103	X	X	X	X	No Fix	Performance Monitoring Event MISALIGN_MEM_REF May Over Count
AP104	X	X	X	X	No Fix	A REP STOS/MOVS to a MONITOR/MWAIT Address Range May Prevent Triggering of the Monitoring Hardware
AP105	X	X	X		Fixed	False Level One Data Cache Parity Machine-Check Exceptions May be Signaled
AP106	X	X	X	X	No Fix	A Memory Access May Get a Wrong Memory Type Following a #GP due to WRMSR to an MTRR Mask
AP107	X	X	X	X	No Fix	PMI While LBR Freeze Enabled May Result in Old/Out-of-date LBR Information
AP108					Fixed	VMCALL failure due to corrupt MSEG location may cause VM Exit to load the machine state incorrectly
AP109	X	X	X		Fixed	Overlap of an Intel® VT APIC Access Page in a Guest with the DS Save Area May Lead to Unpredictable Behavior
AP110	X	X	X	X	No Fix	VTPR Write Access During Event Delivery May Cause an APIC-Access VM Exit
AP111				X	No Fix	BIST Failure After Reset
AP112	X	X	X	X	No Fix	Performance Monitoring Event FP_MMX_TRANS_TO_MMX May Not Count Some Transitions
AP113				X	No Fix	When One Core Executes SEXIT the Other Core's Last Branch Recording May be Incorrect
AP114				X	No Fix	A GETSEC[ENTERACCS] Instruction Executed Immediately after GETSEC[WAKEUP] Instruction May Result in a Processor Hang
AP115	X	X	X	X	No Fix	Instruction Fetch May Cause a Livelock During Snoops of the L1 Data Cache
AP116	X	X	X	X	No Fix	Use of Memory Aliasing with Inconsistent Memory Type may Cause a System Hang or a Machine Check Exception
AP117	X	X	X	X	No Fix	A WB Store Following a REP STOS/MOVS or FXSAVE May Lead to Memory-Ordering Violations
AP118	X	X	X	X	No Fix	VM Exit with Exit Reason "TPR Below Threshold" Can Cause the Blocking by MOV/POP SS and Blocking by STI Bits to be Cleared in the Guest Interruptibility-State Field
AP119	X	X	X	X	No Fix	Using Memory Type Aliasing with cacheable and WC Memory Types May Lead to Memory Ordering Violations
AP120	X	X	X	X	No Fix	VM Exit due to Virtual APIC-Access May Clear RF
AP121					Fixed	Fixed Function Performance Counters MSR_PERF_FIXED_CTR1 (30AH) and MSR_PERF_FIXED_CTR2 (30BH) are Not Cleared When the Processor is Reset
AP122					Fixed	VTPR Access May Lead to System Hang
AP123	X	X	X	X	No Fix	IA32_MC1_STATUS MSR Bit[60] Does Not Reflect Machine Check Error Reporting Enable Correctly

Summary Tables of Changes

NO	B1	B2	L2	GO	Plan	ERRATA
AP124	X	X	X	X	No Fix	RSM Instruction Execution under Certain Conditions May Cause Processor Hang or Unexpected Instruction Execution Results
AP125	X	X	X		Fixed	NMIs May Not Be Blocked by a VM-Entry Failure
AP126	X	X	X	X	No Fix	Benign Exception after a Double Fault May Not Cause a Triple Fault Shutdown

Number	SPECIFICATION CHANGES
-	There are no Specification Changes in this Specification Update revision.

Number	SPECIFICATION CLARIFICATIONS
-	There are no Specification Clarifications in this Specification Update revision.

Number	DOCUMENTATION CHANGES
-	There are no Documentation Changes in this Specification Update revision.

Identification Information

Figure 1. Processor Top-Side Markings Example for the Dual-Core Intel® Xeon® Processor 3000 Series with 2MB L2 Cache with 1066 MHz FSB

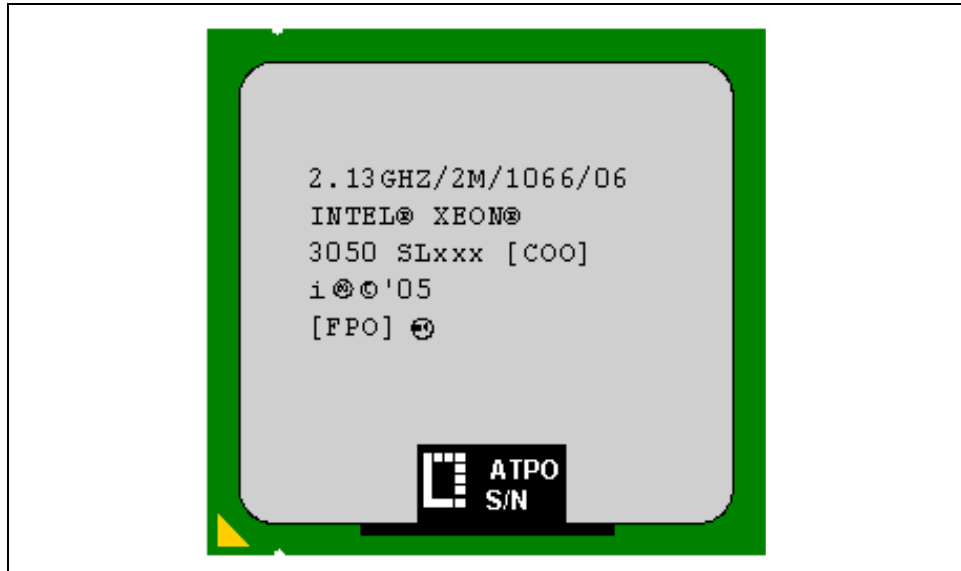
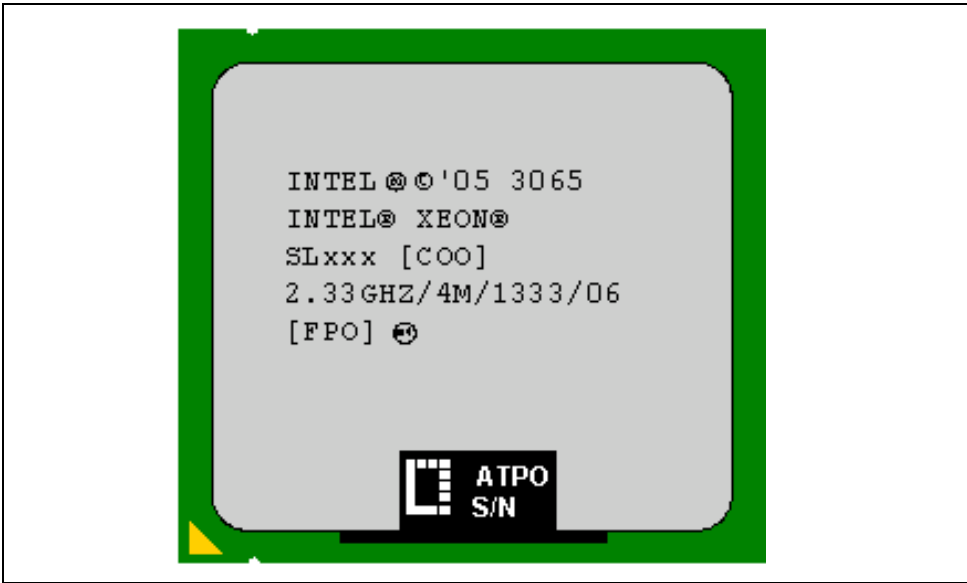


Figure 2. Processor Top-Side Markings Example for the Dual-Core Intel® Xeon® Processor 3000 Series with 4MB L2 Cache with 1066 MHz FSB



Figure 3. Processor Top-Side Markings Example for the Dual-Core Intel® Xeon® Processor 3000 Series with 4MB L2 Cache with 1333 MHz FSB



Component Identification Information

The Dual-Core Intel® Xeon® processor 3000 series can be identified by the following values:

Family ¹	Model ²	Brand
0110b	1111b	10110011b

NOTES:

1. The Family 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.
2. The Model 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.

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. Refer to the *Intel® Processor Identification and the CPUID Instruction Application Note 485* and the *Conroe Processor Family BIOS Writer's Guide (BWG)* for further information on the CPUID instruction.

The following notes are applicable to Table 1 through 2.

1. These processors support the 775_VR_CONFIG_06 specifications
2. These parts support Intel® 64 Architecture
3. These parts support Intel® Virtualization Technology (Intel® VT)
4. These parts support Execute Disable Bit Feature
5. These parts have PROCHOT# enabled
6. These parts have THERMTRIP# enabled
7. These parts support Intel® Thermal Monitor 2 (TM2) feature
8. These parts have PECI enabled
9. These parts have Tdiode enabled
10. These parts have Enhanced Intel SpeedStep® Technology (EIST) enabled
11. These parts have Extended HALT State (C1E) enabled
12. These parts have Extended HALT power of 22W
13. These parts have Extended HALT power of 12W

Component Identification Information

Table 1. Dual-Core Intel® Xeon® Processor 3000 Series 2M SKU Identification Information

S-Spec	Core Stepping	L2 Cache Size (bytes)	Processor Signature	Processor Number	Speed Core/Bus	Package	Notes
SL9TW	B2	2M	06F6h	3040	1.86 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
SL9TY	B2	2M	06F6h	3050	2.13 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
SL9VT	L2	2M	06F2h	3040	1.86 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SLAC2	L2	2M	06F2h	3040	1.86 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SL9VS	L2	2M	06F2h	3050	2.13 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SLABZ	L2	2M	06F2h	3050	2.13 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13

Table 2. Dual-Core Intel® Xeon® Processor 3000 Series 4M SKU Identification Information

S-Spec	Core Stepping	L2 Cache Size (bytes)	Processor Signature	Processor Number	Speed Core/Bus	Package	Notes
SL9TZ	B2	4M	06F6h	3060	2.40 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
SL9U2	B2	4M	06F6h	3070	2.66 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12
SL9ZH	B2	4M	06F6h	3060	2.40 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SLACD	B2	4M	06F6h	3060	2.40 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SL9ZC	B2	4M	06F6h	3070	2.66 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SLACC	B2	4M	06F6h	3070	2.66 GHz / 1066 MHz	775-land LGA	1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 13
SLAA9	G0	4M	06FBh	3065	2.33 GHz / 1333 MHz	775-land LGA	1, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16
SLAA3	G0	4M	06FBh	3075	2.66 GHz / 1333 MHz	775-land LGA	1, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16
SLAA2	G0	4M	06FBh	3085	3.00 GHz / 1333 MHz	775-land LGA	1, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16

Errata

AP1. 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.

AP2. LOCK# Asserted During a Special Cycle Shutdown Transaction May Unexpectedly De-assert

Problem: During a processor shutdown transaction, when LOCK# is asserted and if a DEFER# is received during a snoop phase and the Locked transaction is pipelined on the front side bus (FSB), LOCK# may unexpectedly de-assert.

Implication: When this erratum occurs, the system may hang during shutdown. Intel has not observed this erratum with any commercially available systems or software.

Workaround: None identified.

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

AP3. Address Reported by Machine-Check Architecture (MCA) on Single-bit L2 ECC Errors May be Incorrect

Problem: When correctable Single-bit ECC errors occur in the L2 cache, the address is logged in the MCA address register (MCI_ADDR). Under some scenarios, the address reported may be incorrect.

Implication: Software should not rely on the value reported in MCI_ADDR, for Single-bit L2 ECC errors.

Workaround: None identified.

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

AP4. VERW/VERR/LSL/LAR Instructions May Unexpectedly Update the Last Exception Record (LER) MSR

Problem: The LER MSR may be unexpectedly updated, if the resultant value of the Zero Flag (ZF) is zero after executing the following instructions

- 1) VERR (ZF=0 indicates unsuccessful segment read verification)
- 2) VERW (ZF=0 indicates unsuccessful segment write verification)
- 3) LAR (ZF=0 indicates unsuccessful access rights load)
- 4) LSL (ZF=0 indicates unsuccessful segment limit load)

Implication: The value of the LER MSR may be inaccurate if VERW/VERR/LSL/LAR instructions are executed after the occurrence of an exception.

Workaround: Software exception handlers that rely on the LER MSR value should read the LER MSR before executing VERW/VERR/LSL/LAR instructions.

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

AP5. DR3 Address Match on MOVD/MOVQ/MOVNTQ Memory Store Instruction May Incorrectly Increment Performance Monitoring Count for Saturating SIMD Instructions Retired (Event CFH)

Problem: Performance monitoring for Event CFH normally increments on saturating SIMD instruction retired. Regardless of DR7 programming, if the linear address of a retiring memory store MOVD/MOVQ/MOVNTQ instruction executed matches the address in DR3, the CFH counter may be incorrectly incremented.

Implication: The value observed for performance monitoring count for saturating SIMD instructions retired may be too high. The size of the error is dependent on the number of occurrences of the conditions described above, while the counter is active.

Workaround: None Identified.

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

AP6. SYSRET May Incorrectly Clear RF (Resume Flag) in the RFLAGS Register

Problem: In normal operation, SYSRET will restore the value of RFLAGS from R11 (the value previously saved upon execution of the SYSCALL instruction). Due to this erratum, the RFLAGS.RF bit will be unconditionally cleared after execution of the SYSRET instruction.

Implication: The SYSRET instruction can not be used if the RF flag needs to be set after returning from a system call. Intel has not observed this erratum with any commercially available software.

Workaround: Use the IRET instruction to return from a system call, if RF flag has to be set after the return.

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

AP7. 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.

AP8. Pending x87 FPU Exceptions (#MF) Following STI May Be Serviced Before Higher Priority Interrupts

Problem: Interrupts that are pending prior to the execution of the STI (Set Interrupt Flag) instruction are normally serviced immediately after the instruction following the STI. An exception to this is if the following instruction triggers a #MF. In this situation, the interrupt should be serviced before the #MF. Because of this erratum, if following STI, an instruction that triggers a #MF is executed while STPCLK#, Enhanced Intel SpeedStep Technology transitions or Thermal Monitor events occur, the pending #MF may be serviced before higher priority interrupts.

Implication: Software may observe #MF being serviced before higher priority interrupts.

Workaround: None Identified.

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

AP9. 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.

AP10. 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.

AP11. A Write to an APIC Register Sometimes May Appear to Have Not Occurred

Problem: With respect to the retirement of instructions, stores to the uncacheable memory-based APIC register space are handled in a non-synchronized way. For example if an instruction that masks the interrupt flag, e.g. CLI, is executed soon after an uncacheable write to the Task Priority Register (TPR) that lowers the APIC priority, the interrupt masking operation may take effect before the actual priority has been lowered. This may cause interrupts whose priority is lower than the initial TPR, but higher than the final TPR, to not be serviced until the interrupt enabled flag is finally set, i.e. by STI instruction. Interrupts will remain pending and are not lost.

Implication: In this example the processor may allow interrupts to be accepted but may delay their service.

Workaround: This non-synchronization can be avoided by issuing an APIC register read after the APIC register write. This will force the store to the APIC register before any subsequent instructions are executed. No commercial operating system is known to be impacted by this erratum.

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

AP12. Programming the Digital Thermal Sensor (DTS) Threshold May Cause Unexpected Thermal Interrupts

Problem: Software can enable DTS thermal interrupts by programming the thermal threshold and setting the respective thermal interrupt enable bit. When programming DTS value, the previous DTS threshold may be crossed. This will generate an unexpected thermal interrupt.

Implication: Software may observe an unexpected thermal interrupt occur after reprogramming the thermal threshold.

Workaround: In the ACPI/OS implement a workaround by temporarily disabling the DTS threshold interrupt before updating the DTS threshold value.

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

AP13. Count Value for Performance-Monitoring Counter PMH_PAGE_WALK May be Incorrect

Problem: Performance-Monitoring Counter PMH_PAGE_WALK is used to count the number of page walks resulting from Data Translation Look-Aside Buffer (DTLB) and Instruction Translation Look-Aside (ITLB) misses. Under certain conditions, this counter may be incorrect.

Implication: There may be small errors in the accuracy of the counter.

Workaround: None identified.

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

AP14. LER MSRs May be Incorrectly Updated

Problem: The LER (Last Exception Record) MSRs, MSR_LER_FROM_LIP (1DDH) and MSR_LER_TO_LIP (1DEH) may contain incorrect values after any of the following:

- Either STPCLK#, NMI (NonMaskable Interrupt) or external interrupts
- CMP or TEST instructions with an uncacheable memory operand followed by a conditional jump
- STI/POP SS/MOV SS instructions followed by CMP or TEST instructions and then by a conditional jump

Implication: When the conditions for this erratum occur, the value of the LER MSRs may be incorrectly updated.

Workaround: None identified.

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

AP15. Performance Monitoring Events for Retired Instructions (COH) May Not Be Accurate

Problem: The INST_RETIRED performance monitor may miscount retired instructions as follows:

- Repeat string and repeat I/O operations are not counted when a hardware interrupt is received during or after the last iteration of the repeat flow.
- VMLAUNCH and VMRESUME instructions are not counted.
- HLT and MWAIT instructions are not counted. The following instructions, if executed during HLT or MWAIT events, are also not counted:
 - a) RSM from a C-state SMI during an MWAIT instruction.
 - b) RSM from an SMI during a HLT instruction.

Implication: There may be a smaller than expected value in the INST_RETIRED performance monitoring counter. The extent to which this value is smaller than expected is determined by the frequency of the above cases.

Workaround: None Identified.

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

AP16. Performance Monitoring Event For Number Of Reference Cycles When The Processor Is Not Halted (3CH) Does Not Count According To The Specification

Problem: The CPU_CLK_UNHALTED performance monitor with mask 1 counts bus clock cycles instead of counting the core clock cycles at the maximum possible ratio. The maximum possible ratio is computed by dividing the maximum possible core frequency by the bus frequency.

Implication: The CPU_CLK_UNHALTED performance monitor with mask 1 counts a value lower than expected. The value is lower by exactly one multiple of the maximum possible ratio.

Workaround: Multiply the performance monitor value by the maximum possible ratio.

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

AP17. Using 2M/4M Pages When A20M# Is Asserted May Result in Incorrect Address Translations

Problem: An external A20M# pin if enabled forces address bit 20 to be masked (forced to zero) to emulate real-address mode address wraparound at 1 megabyte. However, if all of the following conditions are met, address bit 20 may not be masked.

- Paging is enabled
- A linear address has bit 20 set
- The address references a large page
- A20M# is enabled

Implication: When A20M# is enabled and an address references a large page the resulting translated physical address may be incorrect. This erratum has not been observed with any commercially available operating system.

Workaround: Operating systems should not allow A20M# to be enabled if the masking of address bit 20 could be applied to an address that references a large page. A20M# is normally only used with the first megabyte of memory.

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

AP18. Writing Shared Unaligned Data that Crosses a Cache Line without Proper Semaphores or Barriers May Expose a Memory Ordering Issue

Problem: Software which is written so that multiple agents can modify the same shared unaligned memory location at the same time may experience a memory ordering issue if multiple loads access this shared data shortly

thereafter. Exposure to this problem requires the use of a data write which spans a cache line boundary.

Implication: This erratum may cause loads to be observed out of order. Intel has not observed this erratum with any commercially available software or system.

Workaround: Software should ensure at least one of the following is true when modifying shared data by multiple agents:

- The shared data is aligned
- Proper semaphores or barriers are used in order to prevent concurrent data accesses.

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

AP19. Code Segment limit violation may occur on 4 Gigabyte limit check

Problem: Code Segment limit violation may occur on 4 Gigabyte limit check when the code stream wraps around in a way that one instruction ends at the last byte of the segment and the next instruction begins at 0x0.

Implication: This is a rare condition that may result in a system hang. Intel has not observed this erratum with any commercially available software, or system.

Workaround: Avoid code that wraps around segment limit.

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

AP20. FP Inexact-Result Exception Flag May Not Be Set

Problem: When the result of a floating-point operation is not exactly representable in the destination format (1/3 in binary form, for example), an inexact-result (precision) exception occurs. When this occurs, the PE bit (bit 5 of the FPU status word) is normally set by the processor. Under certain rare conditions, this bit may not be set when this rounding occurs. However, other actions taken by the processor (invoking the software exception handler if the exception is unmasked) are not affected. This erratum can only occur if one of the following FST instructions is one or two instructions after the floating-point operation which causes the precision exception:

- FST m32real
- FST m64real
- FSTP m32real
- FSTP m64real
- FSTP m80real
- FIST m16int

- FIST m32int
- FISTP m16int
- FISTP m32int
- FISTP m64int
- FISTTP m16int
- FISTTP m32int
- FISTTP m64int

Note that even if this combination of instructions is encountered, there is also a dependency on the internal pipelining and execution state of both instructions in the processor.

Implication: Inexact-result exceptions are commonly masked or ignored by applications, as it happens frequently, and produces a rounded result acceptable to most applications. The PE bit of the FPU status word may not always be set upon receiving an inexact-result exception. Thus, if these exceptions are unmasked, a floating-point error exception handler may not recognize that a precision exception occurred. Note that this is a "sticky" bit, i.e., once set by an inexact-result condition, it remains set until cleared by software.

Workaround: This condition can be avoided by inserting either three NOPs or three non-floating-point non-Jcc instructions between the two floating-point instructions.

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

AP21. Global Pages in the Data Translation Look-Aside Buffer (DTLB) May Not Be Flushed by RSM instruction before Restoring the Architectural State from SMRAM

Problem: The Resume from System Management Mode (RSM) instruction does not flush global pages from the Data Translation Look-Aside Buffer (DTLB) prior to reloading the saved architectural state.

Implication: If SMM turns on paging with global paging enabled and then maps any of linear addresses of SMRAM using global pages, RSM load may load data from the wrong location.

Workaround: Do not use global pages in system management mode.

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

AP22. Sequential Code Fetch to Non-canonical Address May have Non-deterministic Results

Problem: If code sequentially executes off the end of the positive canonical address space (falling through from address 00007fffffffffff to non-canonical address 0000800000000000), under some circumstances the

code fetch will be converted to a canonical fetch at address ffff800000000000.

Implication: Due to this erratum, the processor may transfer control to an unintended address. The result of fetching code at that address is unpredictable and may include an unexpected trap or fault, or execution of the instructions found there.

Workaround: If the last page of the positive canonical address space is not allocated for code (4K page at 00007ffffff000 or 2M page at 00007ffffe0000) then the problem cannot occur.

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

AP23. **VMCALL to Activate Dual-monitor Treatment of SMIs and SMM Ignores Reserved Bit settings in VM-exit Control Field**

Problem: Processors supporting Intel® Virtualization Technology can execute VMCALL from within the Virtual Machine Monitor (VMM) to activate dual-monitor treatment of SMIs and SMM. Due to this erratum, if reserved bits are set to values inconsistent with VMX Capability MSRs, VMCALL may not VMFail.

Implication: VMCALL executed to activate dual-monitor treatment of SMIs and SMM may not VMFail due to incorrect reserved bit settings in VM-Exit control field.

Workaround: Software should ensure that all VMCS reserved bits are set to values consistent with VMX Capability MSRs.

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

AP24. **The PECI Controller Resets to the Idle State**

Problem: After reset, the Platform Environment Control Interface (PECI) client controller should first identify a PECI bus idle condition and only then search for the first rising edge. Due to this erratum, the processor PECI controller resets into the "Idle Detected" state upon processor reset. If another PECI device on the platform is attempting to send a message as the processor PECI controller comes out of reset, the processor PECI controller will typically experience a Frame Check Sequence error and move to the idle state. Rarely, the processor PECI controller may interpret that the message was intended for it and try to reply. In this case a message may be corrupted but this situation will be caught and handled by the PECI error handling protocol.

Implication: The processor PECI controller resets to an incorrect state but the error handling capability of PECI will resolve the situation so that the processor will be able to respond to an incoming message immediately after reset and will not disregard an incoming message that arrives before an idle bus is formally detected.

Workaround: No workaround is necessary due to the PECI error handling protocol.

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

AP25. Some Bus Performance Monitoring Events May Not Count Local Events under Certain Conditions

Problem: Many Performance Monitoring Events require core-specificity, which specifies which core's events are to be counted (local core, other core or both cores). Due to this erratum, some Bus Performance Monitoring events may not count when the core-specificity is set to the local core.

The following Bus Performance Monitoring events will not count power management related events for local core-specificity:

- BUS_TRANS_IO (Event: 6CH) – Will not count I/O level reads resulting from package-resolved C-state
- BUS_TRANS_ANY (Event: 70H) – Will not count Stop-Grants

Implication: The count values for the affected events may be lower than expected. The degree of undercount depends on the occurrence of erratum conditions while the affected events are active.

Workaround: None identified.

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

AP26. Premature Execution of a Load Operation Prior to Exception Handler Invocation

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

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

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

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

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

AP27. General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit

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

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

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

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

AP28. EIP May be Incorrect after Shutdown in IA-32e Mode

Problem: When the processor is going into shutdown state the upper 32 bits of the instruction pointer may be incorrect. This may be observed if the processor is taken out of shutdown state by NMI#.

Implication: A processor that has been taken out of the shutdown state may have an incorrect EIP. The only software which would be affected is diagnostic software that relies on a valid EIP.

Workaround: None identified.

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

AP29. #GP Fault is Not Generated on Writing IA32_MISC_ENABLE [34] When Execute Disable Bit is Not Supported

Problem: A #GP fault is not generated on writing to IA32_MISC_ENABLE [34] bit in a processor which does not support Execute Disable Bit functionality.

Implication: Writing to IA32_MISC_ENABLE [34] bit is silently ignored without generating a fault.

Workaround: None identified.

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

AP30. (E)CX May Get Incorrectly Updated When Performing Fast String REP MOVS or Fast String REP STOS With Large Data Structures

Problem: When performing Fast String REP MOVS or REP STOS commands with data structures [(E)CX*Data Size] larger than the supported address size structure (64K for 16-bit address size and 4G for 32-bit address size) some addresses may be processed more than once. After an amount of data greater than or equal to the address size structure has been processed, external events (such as interrupts) will cause the (E)CX registers to be incremented by a value that corresponds to 64K bytes for 16 bit address size and 4G bytes for 32 bit address size.

Implication: (E)CX may contain an incorrect count which may cause some of the MOVS or STOS operations to re-execute. Intel has not observed this erratum with any commercially available software.

Workaround: Do not use values in (E)CX that when multiplied by the data size give values larger than the address space size (64K for 16-bit address size and 4G for 32-bit address size).

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

AP31. Performance Monitoring Events for Retired Loads (CBH) and Instructions Retired (COH) May Not Be Accurate

Problem: The following events may be counted as instructions that contain a load by the MEM_LOAD_RETIRED performance monitor events and may be counted as loads by the INST_RETIRED (mask 01H) performance monitor event:

- Prefetch instructions
- x87 exceptions on FST* and FBSTP instructions
- Breakpoint matches on loads, stores, and I/O instructions
- Stores which update the A and D bits
- Stores that split across a cache line
- VMX transitions
- Any instruction fetch that misses in the ITLB

Implication: The MEM_LOAD_RETIRED and INST_RETIRED (mask 01H) performance monitor events may count a value higher than expected. The extent to which the values are higher than expected is determined by the frequency of the above events.

Workaround: None identified.

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

AP32. Upper 32 bits of 'From' Address Reported through BTMs or BTSs May be Incorrect

Problem: When a far transfer switches the processor from 32-bit mode to IA-32e mode, the upper 32 bits of the 'From' (source) addresses reported through the BTMs (Branch Trace Messages) or BTSs (Branch Trace Stores) may be incorrect.

Implication: The upper 32 bits of the 'From' address debug information reported through BTMs or BTSs may be incorrect during this transition

Workaround: None identified.

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

AP33. Unsynchronized Cross-Modifying Code Operations Can Cause Unexpected Instruction Execution Results

Problem: The act of one processor, or system bus master, writing data into a currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called cross-modifying code (XMC). XMC that does not force the second processor to execute a synchronizing instruction, prior to execution of the new code, is called unsynchronized XMC.

Software using unsynchronized XMC to modify the instruction byte stream of a processor can see unexpected or unpredictable execution behavior from the processor that is executing the modified code.

Implication: In this case, the phrase "unexpected or unpredictable execution behavior" encompasses the generation of most of the exceptions listed in the Intel Architectures Software Developer's Manual Volume 3: System Programming Guide, including a General Protection Fault (GPF) or other unexpected behaviors. In the event that unpredictable execution causes a GPF the application executing the unsynchronized XMC operation would be terminated by the operating system.

Workaround: In order to avoid this erratum, programmers should use the XMC synchronization algorithm as detailed in the Intel Architectures Software Developer's Manual Volume 3: System Programming Guide, Section: Handling Self- and Cross-Modifying Code.

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

AP34. MSRs Actual Frequency Clock Count (IA32_APERF) or Maximum Frequency Clock Count (IA32_MPERF) May Contain Incorrect Data after a Machine Check Exception (MCE)

Problem: When an MCE occurs during execution of a RDMSR instruction for MSRs Actual Frequency Clock Count (IA32_APERF) or Maximum Frequency Clock Count (IA32_MPERF), the current and subsequent RDMSR instructions for these MSRs may contain incorrect data.

Implication: After an MCE event, accesses to the IA32_APERF and IA32_MPERF MSRs may return incorrect data. A subsequent reset will clear this condition.

Workaround: None identified.

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

AP35. 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.

AP36. Split Locked Stores May not Trigger the Monitoring Hardware

Problem: Logical processors normally resume program execution following the MWAIT, when another logical processor performs a write access to a WB cacheable address within the address range used to perform the MONITOR operation. Due to this erratum, a logical processor may not resume execution until the next targeted interrupt event or O/S timer tick following a locked store that spans across cache lines within the monitored address range.

Implication: The logical processor that executed the MWAIT instruction may not resume execution until the next targeted interrupt event or O/S timer tick in the case where the monitored address is written by a locked store which is split across cache lines.

Workaround: Do not use locked stores that span cache lines in the monitored address range.

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

AP37. REP CMPS/SCAS Operations May Terminate Early in 64-bit Mode when RCX >= 0X10000000

Problem: REP CMPS (Compare String) and SCAS (Scan String) instructions in 64-bit mode may terminate before the count in RCX reaches zero if the initial value of RCX is greater than or equal to 0X10000000.

Implication: Early termination of REP CMPS/SCAS operation may be observed and RFLAGS may be incorrectly updated.

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.

AP38. FXSAVE/FXRSTOR Instructions which Store to the End of the Segment and Cause a Wrap to a Misaligned Base Address (Alignment \leq 0x10h) May Cause FPU Instruction or Operand Pointer Corruption

Problem: If a FXSAVE/FXRSTOR instruction stores to the end of the segment causing a wrap to a misaligned base address (alignment \leq 0x10h), and one of the following conditions is satisfied:

- 1) 32-bit addressing, obtained by using address-size override, when in 64-bit mode
- 2) 16-bit addressing in legacy or compatibility mode

Then, depending on the wrap-around point, one of the below saved values may be corrupted:

- FPU Instruction Pointer Offset
- FPU Instruction Pointer Selector
- FPU Operand Pointer Selector
- FPU Operand Pointer Offset

Implication: This erratum could cause FPU Instruction or Operand pointer corruption and may lead to unexpected operations in the floating point exception handler.

Workaround: Avoid segment base mis-alignment and address wrap-around at the segment boundary.

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

AP39. Cache Data Access Request from One Core Hitting a Modified Line in the L1 Data Cache of the Other Core May Cause Unpredictable System Behavior

Problem: When request for data from Core 1 results in a L1 cache miss, the request is sent to the L2 cache. If this request hits a modified line in the L1 data cache of Core 2, certain internal conditions may cause incorrect data to be returned to the Core 1.

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.

AP40. PREFETCHH Instruction Execution under Some Conditions May Lead to Processor Livelock

Problem: PREFETCHH instruction execution after a split load and dependent upon ongoing store operations may lead to processor livelock.

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

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.

AP41. PREFETCHH Instructions May Not be Executed when Alignment Check (AC) is Enabled

Problem: PREFETCHT0, PREFETCHT1, PREFETCHT2 and PREFETCHNTA instructions may not be executed when Alignment Check is enabled.

Implication: PREFETCHH instructions may not perform the data prefetch if Alignment Check is enabled.

Workaround: Clear the AC flag (bit 18) in the EFLAGS register and/or the AM bit (bit 18) of Control Register CR0 to disable alignment checking.

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

AP42. Upper 32 Bits of the FPU Data (Operand) Pointer in the FXSAVE Memory Image May Be Unexpectedly All 1's after FXSAVE

Problem: The upper 32 bits of the FPU Data (Operand) Pointer may incorrectly be set to all 1's instead of the expected value of all 0's in the FXSAVE memory image if all of the following conditions are true:

- The processor is in 64-bit mode.
- The last floating point operation was in compatibility mode
- Bit 31 of the FPU Data (Operand) Pointer is set.
- An FXSAVE instruction is executed

Implication: Software depending on the full FPU Data (Operand) Pointer may behave unpredictably.

Workaround: None identified.

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

AP43. Concurrent Multi-processor Writes to Non-dirty Page May Result in Unpredictable Behavior

Problem: When a logical processor writes to a non-dirty page, and another logical-processor either writes to the same non-dirty page or explicitly sets the dirty bit in the corresponding page table entry, complex interaction with internal processor activity may cause unpredictable system behavior.

Implication: This erratum may result in unpredictable system behavior and hang.

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

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

AP44. Performance Monitor IDLE_DURING_DIV (18h) Count May Not be Accurate

Problem: Performance monitoring events that count the number of cycles the divider is busy and no other execution unit operation or load operation is in progress may not be accurate.

Implication: The counter may reflect a value higher or lower than the actual number of events.

Workaround: None identified.

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

AP45. 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.

AP46. Shutdown/VMX-Abort Condition May Disable Non-Bootstrap Processors

Problem: When a logical processor encounters an error resulting in shutdown or VMX-Abort, non-bootstrap processors in the package may be unexpectedly disabled.

Implication: Non-bootstrap logical processors in the package that have not observed the error condition may be disabled and may not respond to INIT#, SMI#, NMI#, SIPI or other events.

Workaround: When this erratum occurs, RESET# must be asserted to restore multi-core functionality.

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

AP47. SYSCALL Immediately after Changing EFLAGS.TF May Not Behave According to the New EFLAGS.TF

Problem: If a SYSCALL instruction follows immediately after EFLAGS.TF was updated and IA32_FMASK.TF (bit 8) is cleared, then under certain circumstances SYSCALL may behave according to the previous EFLAGS.TF.

Implication: When the problem occurs, SYSCALL may generate an unexpected debug exception, or may skip an expected debug exception.

Workaround: Mask EFLAGS.TF by setting IA32_FMASK.TF (bit 8).

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

AP48. 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.

AP49. VM Bit is Cleared on Second Fault Handled by Task Switch from Virtual-8086 (VM86)

Problem: Following a task switch to any fault handler that was initiated while the processor was in VM86 mode, if there is an additional fault while servicing the original task switch then the VM bit will be incorrectly cleared in EFLAGS, data segments will not be pushed and the processor will not return to the correct mode upon completion of the second fault handler via IRET.

Implication: When the OS recovers from the second fault handler, the processor will no longer be in VM86 mode. Normally, operating systems should prevent interrupt task switches from faulting, thus the scenario should not occur under normal circumstances.¹

Workaround: None Identified

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

AP50. IA32_FMASK is Reset during an INIT

Problem: IA32_FMASK MSR (0xC0000084) is reset during INIT.

Implication: Implication: If an INIT takes place after IA32_FMASK is programmed, the processor will overwrite the value back to the default value.

Workaround: Operating system software should initialize IA32_FMASK after INIT.

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

AP51. 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: Status: For the steppings affected, see the Summary Tables of Changes.

AP52. Last Branch Records (LBR) Updates May be Incorrect after a Task Switch

Problem: A Task-State Segment (TSS) task switch may incorrectly set the LBR_FROM value to the LBR_TO value.

Implication: The LBR_FROM will have the incorrect address of the Branch Instruction.

Workaround: None Identified.

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

AP53. IO_SMI Indication in SMRAM State Save Area May Be Set Incorrectly

Problem: The IO_SMI bit in SMRAM's location 7FA4H is set to "1" by the CPU to indicate a System Management Interrupt (SMI) occurred as the result of executing an instruction that reads from an I/O port. Due to this erratum, the IO_SMI bit may be incorrectly set by

- A non-I/O instruction.
- SMI is pending while a lower priority event interrupts
- A REP I/O read
- An I/O read that redirects to MWAIT

Implication: SMM handlers may get false IO_SMI indication.

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

AP54. INIT Does Not Clear Global Entries in the TLB

Problem: INIT may not flush a TLB entry when:

- The processor is in protected mode with paging enabled and the page global enable flag is set (PGE bit of CR4 register)
- G bit for the page table entry is set
- TLB entry is present in TLB when INIT occurs

Implication: Implication: Software may encounter unexpected page fault or incorrect address translation due to a TLB entry erroneously left in TLB after INIT.

Workaround: Write to CR3, CR4 (setting bits PSE, PGE or PAE) or CR0 (setting bits PG or PE) registers before writing to memory early in BIOS code to clear all the global entries from TLB.

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

AP55. Using Memory Type Aliasing with Memory Types WB/WT May Lead to Unpredictable Behavior

Problem: Memory type aliasing occurs when a single physical page is mapped to two or more different linear addresses, each with different memory type. Memory type aliasing with the memory types WB and WT may cause the processor to perform incorrect operations leading to unpredictable behavior.

Implication: Software that uses aliasing of WB and WT memory types may observe unpredictable 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.

AP56. Update of Read/Write (R/W) or User/Supervisor (U/S) or Present (P) Bits without TLB Shutdown May Cause Unexpected Processor Behavior

Problem: Updating a page table entry by changing R/W, U/S or P bits without TLB shutdown (as defined by the 4 step procedure in "Propagation of Page Table and Page Directory Entry Changes to Multiple Processors" In volume 3A of the IA-32 Intel® Architectures Software Developer's

Manual), in conjunction with a complex sequence of internal processor micro-architectural events, may lead to unexpected processor behavior.

Implication: This erratum may lead to livelock, shutdown or other unexpected processor behavior. 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.

AP57. BTS Message May Be Lost When the STPCLK# Signal is Active.

Problem: STPCLK# is asserted to enable the processor to enter a low-power state. Under some circumstances, when STPCLK# becomes active, the BTS (Branch Trace Store) message may be either lost and not written or written with corrupted branch address to the Debug Store area

Implication: BTS messages may be lost or be corrupted in the presence of STPCLK# assertions.

Workaround: None Identified.

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

AP58. CMPSB, LODSB, or SCASB in 64-bit Mode with Count Greater or Equal to 2^{48} May Terminate Early

Problem: In 64-bit Mode CMPSB, LODSB, or SCASB executed with a repeat prefix and count greater than or equal to 2^{48} may terminate early. Early termination may result in one of the following.

- The last iteration not being executed
- Signaling of a canonical limit fault (#GP) on the last iteration

Implication: While in 64-bit mode, with count greater or equal to 2^{48} , repeat string operations CMPSB, LODSB or SCASB may terminate without completing the last iteration. Intel has not observed this erratum with any commercially available software.

Workaround: Do not use repeated string operations with RCX greater than or equal to 2^{48} .

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

AP59. REP MOVSB/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may

use an Incorrect Data Size or Lead to Memory-Ordering Violations.

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

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

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

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

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

AP60. MOV To/From Debug Registers Causes Debug Exception

Problem: When in V86 mode, if a MOV instruction is executed to/from on debug register, 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.

AP61. Debug Register May Contain Incorrect Information on a MOVSS or POPSS Instruction Followed by SYSRET

Problem: In IA-32e mode, if a MOVSS or POPSS instruction with a debug breakpoint is followed by the SYSRET instruction; incorrect information may exist in the Debug Status Register (DR6).

Implication: When debugging or when developing debuggers, this behavior should be noted. This erratum will not occur under normal usage of the MOVSS or POPSS instructions (i.e., following them with a MOV ESP instruction).

Workaround: Do not attempt to put a breakpoint on MOVSS and POPSS instructions that are followed by a SYSRET.

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

AP62. EFLAGS Discrepancy on a Page Fault After a Multiprocessor TLB Shutdown

Problem: This erratum may occur when the processor executes one of the following read-modify-write arithmetic instructions and a page fault occurs during the store of the memory operand: ADD, AND, BTC, BTR, BTS, CMPXCHG, DEC, INC, NEG, NOT, OR, ROL/ROR, SAL/SAR/SHL/SHR, SHLD, SHRD, SUB, XOR, and XADD. In this case, the EFLAGS value pushed onto the stack of the page fault handler may reflect the status of the register after the instruction would have completed execution rather than before it. The following conditions are required for the store to generate a page fault and call the operating system page fault handler:

- The store address entry must be evicted from the DTLB by speculative loads from other instructions that hit the same way of the DTLB before the store has completed. DTLB eviction requires at least three-load operations that have linear address bits 15:12 equal to each other and address bits 31:16 different from each other in close physical proximity to the arithmetic operation.
- The page table entry for the store address must have its permissions tightened during the very small window of time between the DTLB eviction and execution of the store. Examples of page permission tightening include from Present to Not Present or from Read/Write to Read Only, etc.
- Another processor, without corresponding synchronization and TLB flush, must cause the permission change.

Implication: This scenario may only occur on a multiprocessor platform running an operating system that performs "lazy" TLB shutdowns. The memory image of the EFLAGS register on the page fault handler's stack prematurely contains the final arithmetic flag values although the instruction has not yet completed. Intel has not identified any operating systems that inspect the arithmetic portion of the EFLAGS

register during a page fault nor observed this erratum in laboratory testing of software applications.

Workaround: No workaround is needed upon normal restart of the instruction, since this erratum is transparent to the faulting code and results in correct instruction behavior. Operating systems may ensure that no processor is currently accessing a page that is scheduled to have its page permissions tightened or have a page fault handler that ignores any incorrect state.

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

AP63. 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.

AP64. Returning to Real Mode from SMM with EFLAGS.VM Set May Result in Unpredictable System Behavior

Problem: Returning back from SMM mode into real mode while EFLAGS.VM is set in SMRAM may result in unpredictable system behavior.

Implication: If SMM software changes the values of the EFLAGS.VM in SMRAM, it may result in unpredictable system behavior. Intel has not observed this behavior in commercially available software.

Workaround: SMM software should not change the value of EFLAGS.VM in SMRAM.

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

AP65. A Thermal Interrupt is Not Generated when the Current Temperature is Invalid

Problem: When the DTS (Digital Thermal Sensor) crosses one of its programmed thresholds it generates an interrupt and logs the event (IA32_THERM_STATUS MSR (019Ch) bits [9,7]). Due to this erratum, if the DTS reaches an invalid temperature (as indicated IA32_THERM_STATUS MSR bit[31]) it does not generate an interrupt even if one of the programmed thresholds is crossed and the corresponding log bits become set.

Implication: When the temperature reaches an invalid temperature the CPU does not generate a Thermal interrupt even if a programmed threshold is crossed.

Workaround: None identified.

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

AP66. VMLAUNCH/VMRESUME May Not Fail when VMCS is Programmed to Cause VM Exit to Return to a Different Mode

Problem: VMLAUNCH/VMRESUME instructions may not fail if the value of the "host address-space size" VM-exit control differs from the setting of IA32_EFER.LMA.

Implication: Programming the VMCS to allow the monitor to be in different modes prior to VMLAUNCH/VMRESUME and after VM-exit may result in undefined behavior

Workaround: Software should ensure that "host address-space size" VM-exit control has the same value as IA32_EFER.LMA at the time of VMLAUNCH/VMRESUME.

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

AP67. 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: Workaround: Software should not generate misaligned stack frames for use with IRET.

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

AP68. Performance Monitoring Event FP_ASSIST May Not be Accurate

Problem: Performance monitoring event FP_ASSIST (11H) may be inaccurate as assist events will be counted twice per actual assist in the following specific cases:

- FADD and FMUL instructions with a NaN(Not a Number) operand and a memory operand
- FDIV instruction with zero operand value in memory

In addition, an assist event may be counted when DAZ (Denormals-Are-Zeros) and FTZ (Flush-To-Zero) flags are turned on even though no actual assist occurs.

Implication: The counter value for the performance monitoring event FP_ASSIST (11H) may be larger than expected. The size of the error is dependent on the number of occurrences of the above conditions while the event is active.

Workaround: None identified.

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

AP69. CPL-Qualified BTS May Report Incorrect Branch-From Instruction Address

Problem: CPL (Current Privilege Level)-qualified BTS (Branch Trace Store) may report incorrect branch-from instruction address under the following conditions:

- Either BTS_OFF_OS[9] or BTS_OFF_USR[10] is selected in IA32_DEBUGCTL MSR (1D9H)
- Privilege-level transitions occur between CPL > 0 and CPL 0 or vice versa.

Implication: Due to this erratum, the From address reported by BTS may be incorrect for the described conditions.

Workaround: None identified.

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

AP70. PEBS Does Not Always Differentiate Between CPL-Qualified Events

- Problem:** Performance monitoring counter configured to sample PEBS (Precise Event Based Sampling) events at a certain privilege level may count samples at the wrong privilege level.
- Implication:** Performance monitoring counter may be higher than expected for CPL-qualified events. Do not use performance monitoring counters for precise event sampling when the precise event is dependent on the CPL value.
- Workaround:** Do not use performance monitoring counters for precise event sampling when the precise event is dependent on the CPL value.
- Status:** For the steppings affected, see the Summary Tables of Changes.

AP71. PMI May Be Delayed to Next PEBS Event

- Problem:** After a PEBS (Precise Event-Based Sampling) event, the PEBS index is compared with the PEBS threshold, and the index is incremented with every event. If PEBS index is equal to the PEBS threshold, a PMI (Performance Monitoring Interrupt) should be issued. Due to this erratum, the PMI may be delayed by one PEBS event.
- Implication:** Debug Store Interrupt Service Routines may observe delay of PMI occurrence by one PEBS event.
- Workaround:** None identified.
- Status:** For the steppings affected, see the Summary Tables of Changes.

AP72. A PMI Due to a PEBS Buffer Overflow Will Not be Signaled Unless IA32_DEBUGCTL[12] is Set

- Problem:** IA32_PERF_GLOBAL_STATUS MSR (38EH) bit [62] when set, indicates that a PEBS (Precise Event-Based Sampling) overflow has occurred and a PMI (Performance Monitor Interrupt) has been sent. Due to this erratum, this bit will not be set unless IA32_DEBUGCTL MSR (1D9H) bit [12] (which stops all Performance Monitor Counters upon a PMI) is also set.
- Implication:** Unless IA32_DEBUGCTL[12] is set, IA32_PERF_GLOBAL_STATUS[62] will not signal that a PMI was generated due to a PEBS Overflow.
- 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.

AP73. The BS Flag in DR6 May be Set for Non-Single-Step #DB Exception

Problem: DR6 BS (Single Step, bit 14) flag may be incorrectly set when the TF (Trap Flag, bit 8) of the EFLAGS Register is set, and a #DB (Debug Exception) occurs due to one of the following:

- DR7 GD (General Detect, bit 13) being bit set;
- INT1 instruction;
- Code breakpoint

Implication: The BS flag may be incorrectly set for non-single-step #DB exception.

Workaround: None identified.

AP74. An Asynchronous MCE During a Far Transfer May Corrupt ESP

Problem: If an asynchronous machine check occurs during an interrupt, call through gate, FAR RET or IRET and in the presence of certain internal conditions, ESP may be corrupted.

Implication: If the MCE (Machine Check Exception) handler is called without a stack switch, then a triple fault will occur due to the corrupted stack pointer, resulting in a processor shutdown. If the MCE is called with a stack switch, e.g. when the CPL (Current Privilege Level) was changed or when going through an interrupt task gate, then the corrupted ESP will be saved on the stack or in the TSS (Task State Segment), and will not be used.

Workaround: Use an interrupt task gate for the machine check handler.

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

AP75. In Single-Stepping on Branches Mode, the BS Bit in the Pending-Debug-Exceptions Field of the Guest State Area will be Incorrectly Set by VM Exit on a MOV to CR8 Instruction

Problem: In a system supporting Intel® Virtualization Technology, the BS bit (bit 14 of the Pending-Debug-Exceptions field) in the guest state area will be incorrectly set when all of the following conditions occur:

- The processor is running in VMX non-root as a 64 bit mode guest;
- The "CR8-load existing" VM-execution control is 0 and the "use TPR shadow" VM-execution is 1;
- Both BTF (Single-Step On Branches, bit 1) of the IA32_DEBUGCTL MSR (1D9H) Register and the TF (Trap Flag, bit 8) of the RFLAGS Register are set;
- "MOV CR8, reg" attempts to program a TPR (Task Priority Register) value that is below the TPR threshold and causes a VM exit.

Implication: A Virtual-Machine will sample the BS bit and will incorrectly inject a Single-Step trap to the guest.

Workaround: A Virtual-Machine Monitor must manually disregard the BS bit in the Guest State Area in case of a VM exit due to a TPR value below the TPR threshold.

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

AP76. B0-B3 Bits in DR6 May Not be Properly Cleared After Code Breakpoint

Problem: B0-B3 bits (breakpoint conditions detect flags, bits [3:0]) in DR6 may not be properly cleared when the following sequence happens:

- 1) POP instruction to SS (Stack Segment) selector;
- 2) Next instruction is FP (Floating Point) that gets FP assist followed by code breakpoint.

Implication: B0-B3 bits in DR6 may not be properly cleared.

Workaround: None identified.

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

AP77. BTM/BTS Branch-From Instruction Address May be Incorrect for Software Interrupts.

Problem: When BTM (Branch Trace Message) or BTS (Branch Trace Store) is enabled, a software interrupt may result in the overwriting of BTM/BTS branch-from instruction address by the LBR (Last Branch Record) branch-from instruction address.

Implication: A BTM/BTS branch-from instruction address may get corrupted for software interrupts.

Workaround: None identified.

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

AP78. Last Branch Records (LBR) Updates May be Incorrect After a Task Switch

Problem: A Task-State Segment (TSS) task switch may incorrectly set the LBR_FROM value to the LBR_TO value.

Implication: The LBR_FROM will have the incorrect address of the Branch Instruction.

Workaround: None identified.

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

AP79. REP Store Instructions in a Specific Situation may cause the Processor to Hang

Problem: During a series of REP (repeat) store instructions a store may try to dispatch to memory prior to the actual completion of the instruction. This behavior depends on the execution order of the instructions, the timing of a speculative jump and the timing of an uncacheable memory store. All types of REP store instructions are affected by this erratum.

Implication: When this erratum occurs, the processor may live lock and/or result in a system hang.

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

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

AP80. Performance Monitoring Events for L1 and L2 Miss May Not be Accurate

Problem: Performance monitoring events 0CBh with an event mask value of 02h or 08h (MEM_LOAD_RETIRED.L1_LINE_MISS or MEM_LOAD_RETIRED.L2_LINE_MISS) may under count the cache miss events.

Implication: Performance monitoring events 0CBh with an event mask value of 02h or 08h may show a count which is lower than expected; the amount by which the count is lower is dependent on other conditions occurring on the same load that missed the cache.

Workaround: None Identified.

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

AP81. Store to WT Memory Data May be Seen in Wrong Order by Two Subsequent Loads

Problem: When data of Store to WT memory is used by two subsequent loads of one thread and another thread performs cacheable write to the same address the first load may get the data from external memory or L2 written by another core, while the second load will get the data straight from the WT Store.

Implication: Software that uses WB to WT memory aliasing may violate proper store ordering.

Workaround: Do not use WB to WT aliasing.

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

AP82. A MOV Instruction from CR8 Register with 16 Bit Operand Size Will Leave Bits 63:16 of the Destination Register Unmodified

Problem: Moves to/from control registers are supposed to ignore REX.W and the 66H (operand size) prefix. In systems supporting Intel® Virtualization Technology, when the processor is operating in VMX non-root operation and "use TPR shadow" VM-execution control is set to 1, a MOV instruction from CR8 with a 16 bit operand size (REX.W =0 and 66H prefix) will only store 16 bits and leave bits 63:16 at the destination register unmodified, instead of storing zeros in them.

Implication: 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.

AP83. Non-Temporal Data Store May be Observed in Wrong Program Order

Problem: When non-temporal data is accessed by multiple read operations in one thread while another thread performs a cacheable write operation to the same address, the data stored may be observed in wrong program order (i.e. later load operations may read older data).

Implication: Software that uses non-temporal data without proper serialization before accessing the non-temporal data may observe data in wrong program order.

Workaround: Software that conforms to the Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A, section "Buffering of Write Combining Memory Locations" will operate correctly.

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

AP84. 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 inaccurately also count certain 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.

AP85. Fault on ENTER Instruction May Result in Unexpected Values on Stack Frame

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

Implication: Data in the created stack frame may be altered following a fault on the ENTER instruction. Please refer to "Procedure Calls For Block-Structured Languages" in Intel® 64 and IA-32 Intel® Architectures 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.

AP86. CPUID Reports Architectural Performance Monitoring Version 2 is Supported, When Only Version 1 Capabilities are Available

Problem: CPUID leaf 0Ah reports the architectural performance monitoring version that is available in EAX[7:0]. Due to this erratum CPUID reports the supported version as 2 instead of 1.

Implication: Software will observe an incorrect version number in CPUID.0Ah.EAX [7:0] in comparison to which features are actually supported.

Workaround: Software should use the recommended enumeration mechanism described in the Architectural Performance Monitoring section of the Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3: System Programming Guide.

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

AP87. Unaligned Accesses to Paging Structures May Cause the Processor to Hang

Problem: When an unaligned access is performed on paging structure entries, accessing a portion of two different entries simultaneously, the processor may live lock.

Implication: When this erratum occurs, the processor may live lock causing a system hang.

Workaround: Do not perform unaligned accesses on paging structure entries.

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

AP88. Microcode Updates Performed During VMX Non-root Operation Could Result in Unexpected Behavior

Problem: When Intel® Virtualization Technology is enabled, microcode updates are allowed only during VMX root operations. Attempts to apply microcode updates while in VMX non-root operation should be silently ignored. Due to this erratum, the processor may allow microcode updates during VMX non-root operations if not explicitly prevented by the host software.

Implication: Microcode updates performed in non-root operation may result in unexpected system behavior.

Workaround: Host software should intercept and prevent loads to IA32_BIOS_UPDT_TRIG MSR (79H) during VMX non-root operations. There are two mechanism that can be used (1) Enabling MSR access protection in the VM-execution controls or (2) Enabling selective MSR protection of IA32_BIOS_UPDT_TRIG MSR.

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

AP89. INVLPG Operation for Large (2M/4M) Pages May be Incomplete under Certain Conditions

Problem: The INVLPG instruction may not completely invalidate Translation Look-aside Buffer (TLB) entries for large pages (2M/4M) when both of the following conditions exist:

- Address range of the page being invalidated spans several Memory Type Range Registers (MTRRs) with different memory types specified
- INVLPG operation is preceded by a Page Assist Event (Page Fault (#PF) or an access that results in either A or D bits being set in a Page Table Entry (PTE))

Implication: Stale translations may remain valid in TLB after a PTE update resulting in unpredictable system behavior. Intel has not observed this erratum with any commercially available software.

Workaround: Software should ensure that the memory type specified in the MTRRs is the same for the entire address range of the large page.

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

AP90. Page Access Bit May be Set Prior to Signaling a Code Segment Limit Fault

Problem: If code segment limit is set close to the end of a code page, then due to this erratum the memory page Access bit (A bit) may be set for the

subsequent page prior to general protection fault on code segment limit.

Implication: When this erratum occurs, a non-accessed page which is present in memory and follows a page that contains the code segment limit may be tagged as accessed.

Workaround: Erratum can be avoided by placing a guard page (non-present or non-executable page) as the last page of the segment or after the page that includes the code segment limit.

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

AP91. Update of Attribute Bits on Page Directories without Immediate TLB Shutdown May Cause Unexpected Processor Behavior

Problem: Updating a page directory entry (or page map level 4 table entry or page directory pointer table entry in IA-32e mode) by changing R/W, U/S or P bits without immediate TLB shutdown (as described by the 4 step procedure in "*Propagation of Page Table and Page Directory Entry Changes to Multiple Processors*" In volume 3A of the *Intel® 64 and IA-32 Architectures Software Developer's Manual*), in conjunction with a complex sequence of internal processor micro-architectural events, may lead to unexpected processor behavior.

Implication: This erratum may lead to livelock, shutdown or other unexpected processor 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.

AP92. Invalid Instructions May Lead to Unexpected Behavior

Implication: Invalid instructions due to undefined opcodes or instructions exceeding the maximum instruction length (due to redundant prefixes placed before the instruction) may lead, under complex circumstances, to unexpected behavior.

Implication: The processor may behave unexpectedly due to invalid instructions. 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.

AP93. EFLAGS, CR0, CR4 and the EXF4 Signal May be Incorrect after Shutdown

Problem: When the processor is going into shutdown due to an RSM inconsistency failure, EFLAGS, CR0 and CR4 may be incorrect. In addition the EXF4 signal may still be asserted. This may be observed if the processor is taken out of shutdown by NMI#.

Implication: A processor that has been taken out of shutdown may have an incorrect EFLAGS, CR0 and CR4. In addition the EXF4 signal may still be asserted.

Workaround: None identified.

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

AP94. Performance Monitoring Counter MACRO_INSTS.DECODED May Not Count Some Decoded Instructions

Problem: MACRO_INSTS.DECODED performance monitoring counter (Event 0AAH, Umask 01H) counts the number of macro instructions decoded, but not necessarily retired. The event is undercounted when the decoded instructions are a complete loop iteration that is decoded in one cycle and the loop is streamed by the LSD (Loop Stream Detector), as described in the *Optimizing the Front End* section of the *Intel® 64 and IA-32 Architectures Optimization Reference Manual*.

Implication: The count value returned by the performance monitoring counter MACRO_INST.DECODED may be lower than expected. The degree of undercounting is dependent on the occurrence of loop iterations that are decoded in one cycle and whether the loop is streamed by the LSD while the counter is active.

Workaround: None identified.

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

AP95. The Stack Size May be Incorrect as a Result of VIP/VIF Check on SYSEXIT and SYSRET

Problem: The stack size may be incorrect under the following scenario:

- The stack size was changed due to a SYSEXIT or SYSRET
- PVI (Protected Mode Virtual Interrupts) mode was enabled (CR4.PVI == 1)
- Both the VIF (Virtual Interrupt Flag) and VIP (Virtual Interrupt Pending) flags of the EFLAGS register are set

Implication: If this erratum occurs the stack size may be incorrect, consequently this may result in 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.

AP96. Performance Monitoring Event SIMD_UOP_TYPE_EXEC.MUL is Counted Incorrectly for PMULUDQ Instruction

Problem: Performance Monitoring Event SIMD_UOP_TYPE_EXEC.MUL (Event select 0B3H, Umask 01H) counts the number of SIMD packed multiply micro-ops executed. The count for PMULUDQ micro-ops may be lower than expected. No other instruction is affected.

Implication: The count value returned by the performance monitoring event SIMD_UOP_TYPE_EXEC.MUL may be lower than expected. The degree of undercount depends on actual occurrences of PMULUDQ instructions, while the counter is active.

Workaround: None identified.

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

AP97. 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.

AP98. Store Ordering May be Incorrect between WC and WP Memory Types

Problem: According to *Intel® 64 and IA-32 Architectures Software Developer's Manual*, Volume 3A "Methods of Caching Available", WP (Write Protected) stores should drain the WC (Write Combining) buffers in the same way as UC (Uncacheable) memory type stores do. Due to this erratum, WP stores may not drain the WC buffers.

Implication: Memory ordering may be violated between WC and WP stores.

Workaround: None identified.

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

AP99. Updating Code Page Directory Attributes without TLB Invalidation May Result in Improper Handling of Code #PF

Problem: Code #PF (Page Fault exception) is normally handled in lower priority order relative to both code #DB (Debug Exception) and code Segment Limit Violation #GP (General Protection Fault). Due to this erratum, code #PF may be handled incorrectly, if all of the following conditions are met:

- A PDE (Page Directory Entry) is modified without invalidating the corresponding TLB (Translation Look-aside Buffer) entry
- Code execution transitions to a different code page such that both
 - The target linear address corresponds to the modified PDE
 - The PTE (Page Table Entry) for the target linear address has an A (Accessed) bit that is clear
- One of the following simultaneous exception conditions is present following the code transition
 - Code #DB and code #PF
 - Code Segment Limit Violation #GP and code #PF

Implication: Software may observe either incorrect processing of code #PF before code Segment Limit Violation #GP or processing of code #PF in lieu of code #DB.

Workaround: None identified.

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

AP100. Performance Monitoring Event CPU_CLK_UNHALTED.REF May Not Count Clock Cycles According to the Processors Operating Frequency

Problem: Performance Counter MSR_PERF_FIXED_CTR2 (MSR 30BH) that counts CPU_CLK_UNHALTED.REF clocks, should count these clock cycles at a constant rate that is determined by the maximum resolved boot frequency, as programmed by BIOS. Due to this erratum, the rate is instead set by the maximum core-clock to bus-clock ratio of the processor, as indicated by hardware.

Implication: No functional impact as a result of this erratum. If the maximum resolved boot frequency as programmed by BIOS is different from the frequency implied by the maximum core-clock to bus-clock ratio of the processor as indicated by hardware, then the following effects may be observed:

- Performance Monitoring Event CPU_CLK_UNHALTED.REF will count at a rate different than the TSC (Time Stamp Counter)
- When running a system with several processors that have different maximum core-clock to bus-clock ratios, CPU_CLK_UNHALTED.REF

monitoring events at each processor will be counted at different rates and therefore will not be comparable.

Workaround: Calculate the ratio of the rates at which the TSC and the CPU_CLK_UNHALTED.REF performance monitoring event count (this can be done by measuring simultaneously their counted value while executing code) and adjust the CPU_CLK_UNHALTED.REF event count to the maximum resolved boot frequency using this ratio.

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

AP101. (E)CX May Get Incorrectly Updated When Performing Fast String REP STOS With Large Data Structures

Problem: When performing Fast String REP STOS commands with data structures [(E)CX*Data Size] larger than the supported address size structure (64K for 16-bit address size and 4G for 32-bit address size) some addresses may be processed more than once. After an amount of data greater than or equal to the address size structure has been processed, external events (such as interrupts) will cause the (E)CX registers to be incremented by a value that corresponds to 64K bytes for 16 bit address size and 4G bytes for 32 bit address size.

Implication: (E)CX may contain an incorrect count which may cause some of the STOS operations to re-execute. Intel has not observed this erratum with any commercially available software.

Workaround: Do not use values in (E)CX that when multiplied by the data size give values larger than the address space size (64K for 16-bit address size and 4G for 32-bit address size).

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

AP102. Performance Monitoring Event BR_INST_RETIRED May Count CPUID Instructions as Branches

Problem: Performance monitoring event BR_INST_RETIRED (C4H) counts retired branch instructions. Due to this erratum, two of its sub-events mistakenly count for CPUID instructions as well. Those sub events are: BR_INST_RETIRED.PRED_NOT_TAKEN (Umask 01H) and BR_INST_RETIRED.ANY (Umask 00H).

Implication: The count value returned by the performance monitoring event BR_INST_RETIRED.PRED_NOT_TAKEN or BR_INST_RETIRED.ANY may be higher than expected. The extent of over counting depends on the occurrence of CPUID instructions, while the counter is active.

Workaround: None identified.

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

AP103. Performance Monitoring Event MISALIGN_MEM_REF May Over Count

Problem: Performance monitoring event MISALIGN_MEM_REF (05H) is used to count the number of memory accesses that cross an 8-byte boundary and are blocked until retirement. Due to this erratum, the performance monitoring event MISALIGN_MEM_REF also counts other memory accesses.

Implication: The performance monitoring event MISALIGN_MEM_REF may over count. The extent of over counting depends on the number of memory accesses retiring while the counter is active.

Workaround: None identified.

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

AP104. A REP STOS/MOVS to a MONITOR/MWAIT Address Range May Prevent Triggering of the Monitoring Hardware

Problem: The MONITOR instruction is used to arm the address monitoring hardware for the subsequent MWAIT instruction. The hardware is triggered on subsequent memory store operations to the monitored address range. Due to this erratum, REP STOS/MOVS fast string operations to the monitored address range may prevent the actual triggering store to be propagated to the monitoring hardware.

Implication: A logical processor executing an MWAIT instruction may not immediately continue program execution if a REP STOS/MOVS targets the monitored address range.

Workaround: Software can avoid this erratum by not using REP STOS/MOVS store operations within the monitored address range.

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

AP105. False Level One Data Cache Parity Machine-Check Exceptions May be Signaled

Problem: Executing an instruction stream containing invalid instructions/data may generate a false Level One Data Cache parity machine-check exception.

Implication: The false Level One Data Cache parity machine-check exception is reported as an uncorrected machine-check error. An uncorrected machine-check error is treated as a fatal exception by the operating system and may cause a shutdown and/or reboot.

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

Status: For the steppings affected, see the Summary

AP106. A Memory Access May Get a Wrong Memory Type Following a #GP due to WRMSR to an MTRR Mask

Problem: The TLB (Translation Lookaside Buffer) may indicate a wrong memory type on a memory access to a large page (2M/4M Byte) following the recovery from a #GP (General Protection Fault) due to a WRMSR to one of the IA32_MTRR_PHYSMASKn MSRs with reserved bits set.

Implication: When this erratum occurs, a memory access may get an incorrect memory type leading to unexpected system operation. As an example, an access to a memory mapped I/O device may be incorrectly marked as cacheable, become cached, and never make it to the I/O device. Intel has not observed this erratum with any commercially available software.

Workaround: Software should not attempt to set reserved bits of IA32_MTRR_PHYSMASKn MSRs.

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

AP107. PMI While LBR Freeze Enabled May Result in Old/Out-of-date LBR Information

Problem: When Precise Event-Based Sampling (PEBS) is configured with Performance Monitoring Interrupt (PMI) on PEBS buffer overflow enabled and Last Branch Record (LBR) Freeze on PMI enabled by setting FREEZE_LBRS_ON_PMI flag (bit 11) to 1 in IA32_DEBUGCTL (MSR 1D9H), the LBR stack is frozen upon the occurrence of a hardware PMI request. Due to this erratum, the LBR freeze may occur too soon (i.e. before the hardware PMI request).

Implication: Following a PMI occurrence, the PMI handler may observe old/out-of-date LBR information that does not describe the last few branches before the PEBS sample that triggered the PMI.

Workaround: None identified.

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

AP108. VMCALL failure due to corrupt MSEG location may cause VM Exit to load the machine state incorrectly

Problem: In systems supporting Intel® Virtualization Technology, if a VMCALL failure occurs due to a corrupt Monitor Segment (MSEG), subsequent VM Exits may load machine state incorrectly.

Implication: Occurrence of this erratum may result in a VMX abort.

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.

AP109. Overlap of an Intel® VT APIC Access Page in a Guest with the DS Save Area May Lead to Unpredictable Behavior

Problem: Logging of a branch record or a PEBS (precise-event-based-sampling) record to the DS (debug store) save area that overlaps with the APIC access page may lead to unpredictable behavior.

Implication: Guest software configured to log branch records or PEBS records cannot specify the DS (debug store) save area within the APIC-access page. Under any expected usage model this type of overlap is not expected to exist. One should be aware of the fact that the specified DS address is of linear form while the APIC access page is of a physical form. Any solution that wishes to avoid this condition will need to comprehend the linear-to-physical translation of the DS related address pointers with respect to the mapping of the physical APIC access page to avoid such an overlap. Under normal circumstances for correctly written software, such an overlap is not expected to exist. Intel has not observed this erratum with any commercially available software.

Workaround: For a fully comprehensive workaround, the VMM should not allow the logging of branch or PEBS records while guest software is running if the "virtualize APIC accesses" VM-execution control is 1.

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

AP110. VTPR Write Access During Event Delivery May Cause an APIC-Access VM Exit

Problem: VTPR write accesses should not cause APIC-access VM exits but instead should cause data to be written to the virtual-APIC page. Due to this erratum, a VTPR write access during event delivery may cause an APIC-access VM exit with no data being written to the virtual-APIC page.

Implication: VTPR accesses are accesses to offset 80H on the APIC-access page. VTPR write accesses can occur during event delivery when pushing data on the stack. Because event delivery performs multiple stack pushes, an event delivery that includes a VTPR write access will also include at least one other write to the APIC-access page. That other write will cause an APIC-access VM exit. Thus, even in the presence of this erratum, any event delivery that includes a VTPR write access will cause an APIC-access VM exit. The only difference with respect to correct behavior will be with regard to page offset saved in the exit qualification by the APIC-access VM exit. A VMM should be able to emulate the event delivery correctly even with the incorrect offset.

Workaround: The VMM should emulate any event delivery that causes an APIC-access VM exit in the same way regardless of the offset saved in the exit qualification.

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

AP111. BIST Failure After Reset

Problem: The processor may show an erroneous BIST (built-in self test) result in bit [17] of EAX register when coming out of reset.

Implication: When this erratum occurs, an erroneous BIST failure will be reported in EAX bit [17]. This failure can be ignored since it is not accurate.

Workaround: It is possible for BIOS to workaround this erratum by masking off bit [17] of the EAX register after coming out of reset.

Status: For the steppings affected, see the Summary

AP112. Performance Monitoring Event FP_MMX_TRANS_TO_MMX May Not Count Some Transitions

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

Implication: The count value for Performance Monitoring Event FP_MMX_TRANS_TO_MMX may be lower than expected. The degree of undercounting is dependent on the occurrences of the erratum condition while the counter is active. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

AP113. When One Core Executes SEXIT the Other Core's Last Branch Recording May be Incorrect

Problem: In processors supporting Intel® Trusted Execution Technology when one core is executing SEXIT and the other core is executing a control-transfer instruction, the FROM_IP field contained in the last branch information may be incorrect for the following:

- LBR (Last Branch Record) MSRs

- BTM (Branch Traces Messages) on the bus
- BTS (Branch Trace Store) records written by the debug store mechanism

Implication: Due to this erratum, last branch information may be incorrect after one core executes SEXIT. 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.

AP114. A GETSEC[ENTERACCS] Instruction Executed Immediately after GETSEC[WAKEUP] Instruction May Result in a Processor Hang

Problem: In dual core processor systems supporting Intel® Trusted Execution Technology, a processor hang or unpredictable system behavior may occur if the ILP (Initiating Logical Processor) executes GETSEC[WAKEUP] and then executes GETSEC[ENTERACCS] without making sure that the RLP (Responding Logical Processor) has woken up in between these two instructions.

Implication: This may cause the processor to hang or execute code down an unintended path.

Workaround: Software must be written to ensure that the RLP has woken-up in response to GETSEC[WAKEUP] instruction and then execute GETSEC[ENTERACCS] instruction.

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

AP115. Instruction Fetch May Cause a Livelock During Snoops of the L1 Data Cache

Problem: A livelock may be observed in rare conditions when instruction fetch causes multiple level one data cache snoops.

Implication: Due to this erratum, a livelock may occur. Intel has not observed this erratum with any commercially available software.

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

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

AP116. Use of Memory Aliasing with Inconsistent Memory Type may Cause a System Hang or a Machine Check Exception

Problem: Software that implements memory aliasing by having more than one linear addresses mapped to the same physical page with different

cache types may cause the system to hang or to report a machine check exception (MCE). This would occur if one of the addresses is non-cacheable and used in a code segment and the other is a cacheable address. If the cacheable address finds its way into the instruction cache, and the non-cacheable address is fetched in the IFU, the processor may invalidate the non-cacheable address from the fetch unit. Any micro-architectural event that causes instruction restart will be expecting this instruction to still be in the fetch unit and lack of it will cause a system hang or an MCE.

Implication: This erratum has not been observed with commercially available software.

Workaround: Although it is possible to have a single physical page mapped by two different linear addresses with different memory types, Intel has strongly discouraged this practice as it may lead to undefined results. Software that needs to implement memory aliasing should manage the memory type consistency.

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

AP117. A WB Store Following a REP STOS/MOVS or FXSAVE May 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 may perform REP MOVS or REP STOS as write combining stores (referred to as "fast strings") for optimal performance. FXSAVE may also be internally implemented using write combining stores. Due to this erratum, stores of a WB (write back) memory type to a cache line previously written by a preceding fast string/FXSAVE instruction may be observed before string/FXSAVE stores.

Implication: A write-back store may be observed before a previous string or FXSAVE related store. Intel has not observed this erratum with any commercially available software.

Workaround: Software desiring strict ordering of string/FXSAVE operations relative to subsequent write-back stores should add an MFENCE or SFENCE instruction between the string/FXSAVE operation and following store-order sensitive code such as that used for synchronization.

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

AP118. VM Exit with Exit Reason "TPR Below Threshold" Can Cause the Blocking by MOV/POP SS and Blocking by STI Bits to be Cleared in the Guest Interruptibility-State Field

Problem: As specified in Section, “VM Exits Induced by the TPR Shadow”, in the *Intel® 64 and IA-32 Architectures Software Developer’s Manual*, Volume 3B, a VM exit occurs immediately after any VM entry performed with the “use TPR shadow”, “activate secondary controls”, and “virtualize APIC accesses” VM-execution controls all set to 1 and with the value of the TPR shadow (bits 7:4 in byte 80H of the virtual-APIC page) less than the TPR-threshold VM-execution control field. Due to this erratum, such a VM exit will clear bit 0 (blocking by STI) and bit 1 (blocking by MOV/POP SS) of the interruptibility-state field of the guest-state area of the VMCS (bit 0 - blocking by STI and bit 1 - blocking by MOV/POP SS should be left unmodified).

Implication: Since the STI, MOV SS, and POP SS instructions cannot modify the TPR shadow, bits 1:0 of the interruptibility-state field will usually be zero before any VM entry meeting the preconditions of this erratum; behavior is correct in this case. However, if VMM software raises the value of the TPR-threshold VM-execution control field above that of the TPR shadow while either of those bits is 1, incorrect behavior may result. This may lead to VMM software prematurely injecting an interrupt into a guest. Intel has not observed this erratum with any commercially available software.

Workaround: VMM software raising the value of the TPR-threshold VM-execution control field should compare it to the TPR shadow. If the threshold value is higher, software should not perform a VM entry; instead, it could perform the actions that it would normally take in response to a VM exit with exit reason “TPR below threshold”.

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

AP119. Using Memory Type Aliasing with cacheable and WC Memory Types May Lead to Memory Ordering Violations

Problem: Memory type aliasing occurs when a single physical page is mapped to two or more different linear addresses, each with different memory types. Memory type aliasing with a cacheable memory type and WC (write combining) may cause the processor to perform incorrect operations leading to memory ordering violations for WC operations.

Implication: Software that uses aliasing between cacheable and WC memory types may observe memory ordering errors within WC memory operations. Intel has not observed this erratum with any commercially available software.

Workaround: None identified. Intel does not support the use of cacheable and WC memory type aliasing, and WC operations are defined as weakly ordered.

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

AP120. VM Exit due to Virtual APIC-Access May Clear RF

Problem: RF (Resume Flag), bit 16 of the EFLAGS/RFLAGS register, is used to restart instruction execution without getting an instruction breakpoint on the instruction following a debug breakpoint exception. Due to this erratum, in a system supporting Intel® Virtualization Technology, when a VM Exit occurs due to Virtual APIC-Access (Advanced Programmable Interrupt Controller-Access) the EFLAGS/RFLAGS saved in the VMCS (Virtual-Machine Control Structure) may contain an RF value of 0.

Implication: When this erratum occurs, following a VM Exit due to a Virtual APIC-access, the processor may unintentionally break on the subsequent instruction after VM entry.

Workaround: None identified.

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

AP121. Fixed Function Performance Counters MSR_PERF_FIXED_CTR1 (30AH) and MSR_PERF_FIXED_CTR2 (30BH) are Not Cleared When the Processor is Reset

Problem: The Fixed Function Performance Counters that count the number of core cycles and reference cycles when the core is not in a halt state are not cleared when the processor is reset.

Implication: The MSR_PERF_FIXED_CTR1 and MSR_PERF_FIXED_CTR2 counters may contain unexpected values after reset.

Workaround: BIOS can workaround this erratum by clearing the counters at processor initialization time.

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

AP122. VTPR Access May Lead to System Hang

Problem: The logical processor may hang if an instruction performs a VTPR access and the next instruction to be executed is located on a different code page.

Implication: Software running VMX non-root operation may cause a logical processor to hang if the virtual-machine monitor (VMM) sets both the "use TPR shadow" and "virtualize APIC accesses" VM-execution controls.

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.

AP123. IA32_MC1_STATUS MSR Bit[60] Does Not Reflect Machine Check Error Reporting Enable Correctly

Problem: IA32_MC1_STATUS MSR (405H) bit[60] (EN- Error Enabled) is supposed to indicate whether the enable bit in the IA32_MC1_CTL MSR (404H) was set at the time of the last update to the IA32_MC1_STATUS MSR. Due to this erratum, IA32_MC1_STATUS MSR bit[60] instead reports the current value of the IA32_MC1_CTL MSR enable bit.

Implication: IA32_MC1_STATUS MSR bit [60] may not reflect the correct state of the enable bit in the IA32_MC1_CTL MSR at the time of the last update.

Workaround: None identified.

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

AP124. RSM Instruction Execution under Certain Conditions May Cause Processor Hang or Unexpected Instruction Execution Results

Problem: RSM instruction execution, under certain conditions triggered by a complex sequence of internal processor micro-architectural events, may lead to processor hang, or unexpected instruction execution results.

Implication: In the above sequence, the processor may live lock or hang, or RSM instruction may restart the interrupted processor context through a nondeterministic EIP offset in the code segment, resulting in unexpected instruction execution, unexpected exceptions or system hang. 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.

AP125. NMIs May Not Be Blocked by a VM-Entry Failure

Problem: The *Intel® 64 and IA-32 Architectures Software Developer's Manual Volume 3B: System Programming Guide, Part 2* specifies that, following a VM-entry failure during or after loading guest state, "the state of blocking by NMI is what it was before VM entry." If non-maskable interrupts (NMIs) are blocked and the "virtual NMIs" VM-execution control set to 1, this erratum may result in NMIs not being blocked after a VM-entry failure during or after loading guest state.

Implication: VM-entry failures that cause NMIs to become unblocked may cause the processor to deliver an NMI to software that is not prepared for it.

Workaround: VMM software should configure the virtual-machine control structure (VMCS) so that VM-entry failures do not occur.

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

AP126. Benign Exception after a Double Fault May Not Cause a Triple Fault Shutdown

Problem: According to the *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3A, "Exception and Interrupt Reference"*, if another exception occurs while attempting to call the double-fault handler, the processor enters shutdown mode. However due to this erratum, only Contributory Exceptions and Page Faults will cause a triple fault shutdown, whereas a benign exception may not.

Implication: If a benign exception occurs while attempting to call the double-fault handler, the processor may hang or may handle the benign exception. 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.

Specification Changes

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

- *Dual-Core Intel® Xeon® Processor 3000 Series Electrical, Mechanical, and Thermal Specifications (EMTS)*
- *Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual, volumes 1, 2A, 2B, 3A, and 3B*

All Specification Changes will be incorporated into a future version of the appropriate Dual-Core Intel® Xeon® processor 3000 series documentation.

Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Over time processor numbers will increment based on changes in clock, speed, cache, FSB, 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 http://www.intel.com/products/processor_number for details.

Specification Clarifications

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

- *Dual-Core Intel® Xeon® Processor 3000 Series Electrical, Mechanical, and Thermal Specifications (EMTS)*
- *Intel® 64 and IA-32 Intel® Architectures Software Developer's Manual, volumes 1, 2A, 2B, 3A, and 3B*

All Specification Clarifications will be incorporated into a future version of the appropriate Dual-Core Intel® Xeon® processor documentation.

Documentation Changes

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

- *Dual-Core Intel® Xeon® Processor 3000 Series Electrical, Mechanical, and Thermal Specifications (EMTS)*

All Documentation Changes will be incorporated into a future version of the appropriate Dual-Core Intel® Xeon® processor documentation.

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

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