

# Quad-Core Intel® Xeon® Processor 5400 Series


Specification Update

---

*December 2007*

Order Number : 318585-002

**Notice:** The Quad-Core Intel® Xeon® Processor 5400 Series 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.



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

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

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

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See [http://www.intel.com/products/processor\\_number](http://www.intel.com/products/processor_number) for details.

The Quad-Core Intel® Xeon® Processor 5400 Series 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.

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

Intel, Intel Core, Celeron, Pentium, Intel Xeon, Intel SpeedStep and the Intel logo are trademarks of Intel Corporation in the U. S. and other countries.

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

Copyright © 2007, Intel Corporation. All Rights Reserved.



## Contents

---

<b>Revision History .....</b>	<b>4</b>
<b>Preface .....</b>	<b>5</b>
<b>Summary Tables of Changes .....</b>	<b>7</b>
<b>Identification Information .....</b>	<b>13</b>
<b>Errata .....</b>	<b>15</b>
<b>Specification Changes.....</b>	<b>30</b>
<b>Specification Clarifications .....</b>	<b>31</b>
<b>Documentation Changes .....</b>	<b>32</b>



## Revision History

---

Revision	Description	Date
-001	Initial Release	November 2007
-002	Added Errata AX43 through AX50 Updated Erratum AX26	December 2007



# Preface

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

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

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

## Affected Documents

Document Title	Document Number/ Location
<i>Quad-Core Intel® Xeon® Processor 5400 Series Datasheet</i>	318589

## Related Documents

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



## Nomenclature

**Errata** are design defects or errors. These may cause the Quad-Core Intel® Xeon® Processor 5400 Series behavior to deviate from published specifications. Hardware and software designed to be used with any given stepping must assume that all errata documented for that stepping are present on all devices.

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

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

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

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

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



# Summary Tables of Changes

---

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes which apply to the Quad-Core Intel® Xeon® Processor 5400 Series product. Intel may fix some of the errata in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations:

## Codes Used in Summary Tables

### Stepping

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

### Page

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

### Status

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

### Row

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

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® Celeron® 4 processor 6x1 sequence
AC =	Intel® Celeron® processor in 478 pin package
AD =	Intel® Celeron® D processor on 65nm process
AE =	Intel® Core™ Duo processor and Intel® Core™ Solo processor on 65nm process
AF =	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





AK =	Intel® Core™2 Extreme quad-core processor QX6000 sequence and Intel® Core™2 Quad processor Q6000 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 E2000 sequence
AR =	Intel® Celeron® processor 500 series
AS =	Intel® Xeon® processor 7200, 7300 series
AT =	Intel® Celeron® processor 200 series
AV =	Intel® Core™2 Extreme Processor QX9000 Sequence and Intel® Core™2 Quad Processor Q9000 Sequence processor
AX =	Quad-Core Intel® Xeon® Processor 5400 Series
AY =	Dual-Core Intel® Xeon® Processor 5200 Series



## Errata (Sheet 1 of 2)

Number	Steppings	Status	ERRATA
	C-0		
AX1	X	No Fix	EFLAGS Discrepancy on a Page Fault After a Multiprocessor TLB Shutdown
AX2	X	No Fix	INVLPG Operation for Large (2M/4M) Pages May be Incomplete under Certain Conditions
AX3	X	No Fix	Store to WT Memory Data May be Seen in Wrong Order by Two Subsequent Loads
AX4	X	No Fix	Non-Temporal Data Store May be Observed in Wrong Program Order
AX5	X	No Fix	Page Access Bit May be Set Prior to Signaling a Code Segment Limit Fault
AX6	X	No Fix	Updating Code Page Directory Attributes without TLB Invalidation May Result in Improper Handling of Code #PF
AX7	X	No Fix	Storage of PEBS Record Delayed Following Execution of MOV SS or STI
AX8	X	No Fix	Performance Monitoring Event FP_MMX_TRANS_TO_MMX May Not Count Some Transitions
AX9	X	No Fix	A REP STOS/MOVS to a MONITOR/MWAIT Address Range May Prevent Triggering of the Monitoring Hardware
AX10	X	No Fix	Performance Monitoring Event MISALIGN_MEM_REF May Over Count
AX11	X	No Fix	The Processor May Report a #TS Instead of a #GP Fault
AX12	X	No Fix	Code Segment Limit Violation May Occur on 4 Gigabyte Limit Check
AX13	X	No Fix	A Write to an APIC Register Sometimes May Appear to Have Not Occurred
AX14	X	No Fix	Last Branch Records (LBR) Updates May be Incorrect after a Task Switch
AX15	X	No Fix	REP MOVS/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may use an Incorrect Data Size or Lead to Memory-Ordering Violations.
AX16	X	No Fix	Upper 32 bits of 'From' Address Reported through BTMs or BTSs May be Incorrect
AX17	X	No Fix	Address Reported by Machine-Check Architecture (MCA) on Single-bit L2 ECC Errors May be Incorrect
AX18	X	No Fix	Code Segment Limit/Canonical Faults on RSM May be Serviced before Higher Priority Interrupts/Exceptions
AX19	X	No Fix	Store Ordering May be Incorrect between WC and WP Memory Type
AX20	X	No Fix	EFLAGS, CR0, CR4 and the EXF4 Signal May be Incorrect after Shutdown
AX21	X	No Fix	Premature Execution of a Load Operation Prior to Exception Handler Invocation
AX22	X	No Fix	Performance Monitoring Events for Retired Instructions (C0H) May Not Be Accurate
AX23	X	No Fix	Returning to Real Mode from SMM with EFLAGS.VM Set May Result in Unpredictable System Behavior
AX24	X	No Fix	CMPSB, LODSB, or SCASB in 64-bit Mode with Count Greater or Equal to 248 May Terminate Early
AX25	X	No Fix	Writing the Local Vector Table (LVT) when an Interrupt is Pending May Cause an Unexpected Interrupt
AX26	X	No Fix	Pending x87 FPU Exceptions (#MF) Following STI May Be Serviced Before Higher Priority Interrupts
AX27	X	No Fix	VERW/VERR/LSL/LAR Instructions May Unexpectedly Update the Last Exception Record (LER) MSR
AX28	X	No Fix	INIT Does Not Clear Global Entries in the TLB



## Errata (Sheet 2 of 2)

Number	Steppings	Status	ERRATA
	C-0		
AX29	X	No Fix	Split Locked Stores May not Trigger the Monitoring Hardware
AX30	X	No Fix	Programming the Digital Thermal Sensor (DTS) Threshold May Cause Unexpected Thermal Interrupts
AX31	X	No Fix	Writing Shared Unaligned Data that Crosses a Cache Line without Proper Semaphores or Barriers May Expose a Memory Ordering Issue
AX32	X	No Fix	General Protection (#GP) Fault May Not Be Signaled on Data Segment Limit Violation above 4-G Limit
AX33	X	No Fix	An Asynchronous MCE During a Far Transfer May Corrupt ESP
AX34	X	Plan Fix	CPUID Reports Architectural Performance Monitoring Version 2 is Supported, When Only Version 1 Capabilities are Available
AX35	X	No Fix	B0-B3 Bits in DR6 May Not be Properly Cleared After Code Breakpoint
AX36	X	No Fix	An xTPR Update Transaction Cycle, if Enabled, May be Issued to the FSB after the Processor has Issued a Stop-Grant Special Cycle
AX37	X	Plan Fix	Performance Monitoring Event IA32_FIXED_CTR2 May Not Function Properly when Max Ratio is a Non-Integer Core-to-Bus Ratio
AX38	X	No Fix	Instruction Fetch May Cause a Livelock During Snoops of the L1 Data Cache
AX39	X	No Fix	Use of Memory Aliasing with Inconsistent Memory Type may Cause a System Hang or a Machine Check Exception
AX40	X	No Fix	A WB Store Following a REP STOS/MOVS or FXSAVE May Lead to Memory-Ordering Violations
AX41	X	Plan 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
AX42	X	No Fix	Using Memory Type Aliasing with Cacheable and WC Memory Types May Lead to Memory Ordering Violations
AX43	X	No Fix	VM Exit Caused by a SIPI Results in Zero Being Saved to the Guest RIP Field in the VMCS
AX44	X	Plan Fix	NMIs May Not Be Blocked by a VM-Entry Failure
AX45	X	Plan Fix	Partial Streaming Load Instruction Sequence May Cause the Processor to Hang
AX46	X	Plan Fix	Self/Cross Modifying Code May Not be Detected or May Cause a Machine Check Exception
AX47	X	Plan Fix	Data TLB Eviction Condition in the Middle of a Cacheline Split Load Operation May Cause the Processor to Hang
AX48	X	Plan Fix	Update of Read/Write (R/W) or User/Supervisor (U/S) or Present (P) Bits without TLB Shutdown May Cause Unexpected Processor Behavior
AX49	X	Plan Fix	RSM Instruction Execution under Certain Conditions May Cause Processor Hang or Unexpected Instruction Execution Results
AX50	X	No Fix	Benign Exception after a Double Fault May Not Cause a Triple Fault Shutdown



## Specification Changes

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

## Specification Clarifications

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

## Documentation Changes

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



# Identification Information

## Component Identification via Programming Interface

The Quad-Core Intel® Xeon® Processor 5400 Series stepping can be identified by the following register contents:

Reserved	Extended Family <sup>1</sup>	Extended Model <sup>2</sup>	Reserved	Processor Type <sup>3</sup>	Family Code <sup>4</sup>	Model Number <sup>5</sup>	Stepping ID <sup>6</sup>
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	0000000b	0001b		00b	0110b	0111b	XXXXb

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

**Notes:**

1. The Extended Family, bits [27:20] are used in conjunction with the Family Code, specified in bits [11:8], to indicate whether the processor belongs to the Intel386, Intel486, Pentium, Pentium Pro, Pentium 4, or Intel® Core™ processor family.
2. The Extended Model, bits [19:16] in conjunction with the Model Number, specified in bits [7:4], are used to identify the model of the processor within the processor's family.
3. The Processor Type, specified in bits [13:12] indicates whether the processor is an original OEM processor, an OverDrive processor, or a dual processor (capable of being used in a dual processor system).
4. The Family Code corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
5. The Model Number corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
6. The Stepping ID in bits [3:0] indicates the revision number of that model. See Table 2 for the processor stepping ID number in the CPUID information.


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



## Component Marking Information

Quad-Core Intel® Xeon® Processor 5400 Series stepping can be identified by the following component markings:

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



**Table 1. Quad-Core Intel® Xeon® Processor 5400 Series Identification Information**

S-Spec	Processor Number	Core Stepping	CPUID <sup>1</sup>	Core Freq (GHz)	Data Bus Freq (MHz)	L2 Cache Size	Processor Package Revision	Notes
SLANZ	X5482	C-0	10676h	3.20	1600	12 (2x6MB)	01	4
SLANP	X5460	C-0	10676h	3.16	1333	12 (2x6MB)	01	3
SLASA	X5472	C-0	10676h	3	1600	12 (2x6MB)	01	3
SLASB	X5450	C-0	10676h	3	1333	12 (2x6MB)	01	3
SLANR	E5472	C-0	10676h	3	1600	12 (2x6MB)	01	2
SLANQ	E5450	C-0	10676h	3	1333	12 (2x6MB)	01	2
SLANS	E5440	C-0	10676h	2.83	1333	12 (2x6MB)	01	2
SLANT	E5462	C-0	10676h	2.80	1600	12 (2x6MB)	01	2
SLANU	E5430	C-0	10676h	2.66	1333	12 (2x6MB)	01	2
SLANV	E5420	C-0	10676h	2.50	1333	12 (2x6MB)	01	2
SLANW	E5410	C-0	10676h	2.33	1333	12 (2x6MB)	01	2
SLAP2	E5405	C-0	10676h	2	1333	12 (2x6MB)	01	2, 5, 6

### Notes:

1. CPUID is 00001067sh, where 's' is the stepping number.
2. This is a Quad-Core Intel® Xeon® Processor 5400 Series with a 80W TDP (Thermal Design Power).
3. This is a Quad-Core Intel® Xeon® Processor 5400 Series with a 120W TDP (Thermal Design Power).
4. This is a Quad-Core Intel® Xeon® Processor 5400 Series with a 150W TDP (Thermal Design Power).
5. Like other processors that have Max Ratio equal to the Min Ratio, these processors will report support for C1E and TM2 but will not change frequency or voltage in these states.
6. These parts do not support Enhanced Intel SpeedStep(R) Technology.

# Errata

---

## **AX1. 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:

1. 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.
2. 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.
3. 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](#).

## **AX2. 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](#).

### **AX3. 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](#).

### **AX4. 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](#).

### **AX5. 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](#).

### **AX6. 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:

1. A PDE (Page Directory Entry) is modified without invalidating the corresponding TLB (Translation Look-aside Buffer) entry
2. 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
3. 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](#).

#### **AX7. 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](#).

#### **AX8. 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](#).

#### **AX9. 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](#).

#### **AX10. 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 the 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](#).

#### **AX11. 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](#).

#### **AX12. 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](#).

#### **AX13. 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](#).



#### **AX14. 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](#).

#### **AX15. REP MOVS/STOS Executing with Fast Strings Enabled and Crossing Page Boundaries with Inconsistent Memory Types may use an Incorrect Data Size or Lead to Memory-Ordering Violations.**

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

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

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

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

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX16. 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](#).

#### **AX17. 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](#).

#### **AX18. Code Segment Limit/Canonical Faults on RSM May be Serviced before Higher Priority Interrupts/Exceptions**

**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.)

**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](#).

#### **AX19. Store Ordering May be Incorrect between WC and WP Memory Type**

**Problem:** According to Intel® 64 and IA-32 Intel Architecture 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](#).

#### **AX20. 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](#).

#### **AX21. 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.

1. If an instruction that performs a memory load causes a code segment limit violation.
2. If a waiting X87 floating-point (FP) instruction or MMX™ technology (MMX) instruction that performs a memory load has a floating-point exception pending.
3. 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](#).

## **AX22. Performance Monitoring Events for Retired Instructions (C0H) 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](#).

## **AX23. 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](#).

## **AX24. 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 248, 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](#).

#### **AX25. 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](#).

#### **AX26. 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](#).

#### **AX27. 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](#).

#### **AX28. 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:** 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](#).

#### **AX29. 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](#).

#### **AX30. 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](#).

#### **AX31. 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](#).

#### **AX32. 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 (0xffffffffh) 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 (0xffffffffh).

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX33. 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 new 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](#).

#### **AX34. 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](#).

#### **AX35. 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](#).

#### **AX36. An xTPR Update Transaction Cycle, if Enabled, May be Issued to the FSB after the Processor has Issued a Stop-Grant Special Cycle**

**Problem:** According to the FSB (Front Side Bus) protocol specification, no FSB cycles should be issued by the processor once a Stop-Grant special cycle has been issued to the bus. If xTPR update transactions are enabled by clearing the IA32\_MISC\_ENABLES[bit 23] at the time of Stop-Clock assertion, an xTPR update transaction cycle may be issued to the FSB after the processor has issued a Stop Grant Acknowledge transaction.

**Implication:** When this erratum occurs in systems using C-states C2 (Stop-Grant State) and higher the result could be a system hang.

**Workaround:** BIOS must leave the xTPR update transactions disabled (default).

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX37. Performance Monitoring Event IA32\_FIXED\_CTR2 May Not Function Properly when Max Ratio is a Non-Integer Core-to-Bus Ratio**

**Problem:** Performance Counter IA32\_FIXED\_CTR2 (MSR 30BH) event counts CPU reference clocks when the core is not in a halt state. This event is not affected by core frequency changes (e.g., P states, TM2 transitions) but counts at the same frequency as the Time-Stamp Counter IA32\_TIME\_STAMP\_COUNTER (MSR 10H). Due to this erratum, the IA32\_FIXED\_CTR2 will not function properly when the non-integer core-to-bus ratio multiplier feature is used and when a non-zero value is written to IA32\_FIXED\_CTR2. Non-integer core-to-bus ratio enables additional operating frequencies. This feature can be detected by IA32\_PLATFORM\_ID (MSR 17H) bit [23].

**Implication:** The Performance Monitoring Event IA32\_FIXED\_CTR2 may result in an inaccurate count when the non-integer core-to-bus multiplier feature is used.

**Workaround:** If writing to IA32\_FIXED\_CTR2 and using a non-integer core-to-bus ratio multiplier, always write a zero.

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX38. 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](#).

#### **AX39. 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](#).

#### **AX40. 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](#).

#### **AX41. 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](#).

#### **AX42. 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](#).

#### **AX43. VM Exit Caused by a SIPI Results in Zero Being Saved to the Guest RIP Field in the VMCS**

**Problem:** If a logical processor is in VMX non-root operation and in the wait-for-SIPI state, an occurrence of a start-up IPI (SIPI) causes a VM exit. Due to this erratum, such VM exits always save zero into the RIP field of the guest-state area of the virtual-machine control structure (VMCS) instead of the value of RIP before the SIPI was received.

**Implication:** In the absence of virtualization, a SIPI received by a logical processor in the wait-for-SIPI state results in the logical processor starting execution from the vector sent in the SIPI regardless of the value of RIP before the SIPI was received. A virtual-machine monitor (VMM) responding to a SIPI-induced VM exit can emulate this behavior because the SIPI vector is saved in the lower 8 bits of the exit qualification field in the VMCS. Such a VMM should be unaffected by this erratum. A VMM that does not emulate this behavior may need to recover the old value of RIP through alternative means. Intel has not observed this erratum with any commercially available software.

**Workaround:** VMM software that may respond to SIPI-induced VM exits by resuming the interrupt guest context without emulating the non-virtualized SIPI response should (1) save from the VMCS (using VMREAD) the value of RIP before any VM entry to the wait-for-SIPI state; and (2) restore to the VMCS (using VMWRITE) that value before the next VM entry that resumes the guest in any state other than wait-for-SIPI.

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX44. 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](#).



#### **AX45. Partial Streaming Load Instruction Sequence May Cause the Processor to Hang**

**Problem:** Under some rare conditions, when multiple streaming load instructions (MOVNTDQA) are mixed with non-streaming loads that split across cache lines, the processor may hang.

**Implication:** Under the scenario described above, the processor may 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. However, streaming behavior may be re-enabled by setting bit 5 to 1 of the MSR at address 0x21 for software development or testing purposes. If this bit is changed, then a read-modify-write should be performed to preserve other bits of this MSR. When the streaming behavior is enabled and using streaming load instructions, always consume a full cache line worth of data and/or avoid mixing them with non-streaming memory references. If streaming loads are used to read partial cache lines, and mixed with non-streaming memory references, use fences to isolate the streaming load operations from non-streaming memory operations.

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX46. Self/Cross Modifying Code May Not be Detected or May Cause a Machine Check Exception**

**Problem:** If instructions from at least three different ways in the same instruction cache set exist in the pipeline combined with some rare internal state, self-modifying code (SMC) or cross-modifying code may not be detected and/or handled.

**Implication:** An instruction that should be overwritten by another instruction while in the processor pipeline may not be detected/modified, and could retire without detection. Alternatively the instruction may cause a Machine Check Exception. 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](#).

#### **AX47. Data TLB Eviction Condition in the Middle of a Cacheline Split Load Operation May Cause the Processor to Hang**

**Problem:** If the TLB translation gets evicted while completing a cacheline split load operation, under rare scenarios the processor may hang.

**Implication:** The cacheline split load operation may not be able to complete normally, and the machine may hang and generate Machine Check Exception. 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](#).

#### **AX48. 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, even when transitioning these bits from 0 to 1, without keeping the affected linear address range coherent with all TLB (Translation Lookaside Buffers) and paging-structures caches in the processor, in conjunction with a complex sequence of internal processor micro-architectural events and store operations, 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.

**Implication:** None identified.

**Status:** For the steppings affected, see the [Summary Tables of Changes](#).

#### **AX49. RSM Instruction Execution under Certain Conditions May Cause Processor Hang or Unexpected Instruction Execution Results**

**Problem:** 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](#).

#### **AX50. 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. Due to this erratum, any benign faults while attempting to call double-fault handler will not cause a shutdown. However Contributory Exceptions and Page Faults will continue to cause a triple fault shutdown.

**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:

- Quad-Core Intel® Xeon® Processor 5400 Series *Datasheet*

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



# Specification Clarifications

---

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

- Quad-Core Intel® Xeon® Processor 5400 Series *Datasheet*

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



## Documentation Changes

---

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

- Quad-Core Intel® Xeon® Processor 5400 Series *Datasheet*

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

**Note:**

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

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

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

