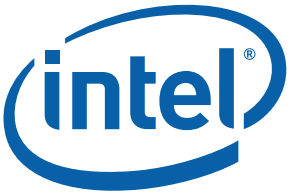

Intel® Core™ i7 Processor Family for LGA2011 Socket

Specification Update

Supporting Desktop Intel® Core™ i7-4960X Extreme Edition Processor Series for the LGA2011 Socket

Supporting Desktop Intel® Core™ i7-49xx and i7-48xx Processor Series for the LGA2011 Socket

April 2015

Revision 015

Reference Number: 329368-015

Contents

Revision History 4
Preface 5
Identification Information 7
Summary Tables of Changes 9
Errata 15
Specification Changes 52
Specification Clarifications 53
Documentation Changes 54

Revision History

Version	Description	Date
001	• Initial release.	September 2013
002	• Added errata CG90-CG100	November 2013
003	• Added errata CG101 and CG102	December 2013
004	• Added errata CG103	February 2014
005	• Revised erratum CG79 to show only the supported performance monitor counters • Added erratum CG104	March 2014
006	• Added errata CG105, CG106, CG107, CG108	April 2014
007	• Added errata CG109, CG110, CG111, CG112, CG113, CG114	May 2014
008	• Added errata CG115, CG116, CG117, CG118, CG119	June 2014
009	• Added erratum CG120	July 2014
010	• Added erratum CG121	August 2014
011	• Modified errata CG115 and CG120 • Added errata CG122, CG123, CG124, CG125 and CG126	September 2014
012	• Added errata CG127 - CG138	October 2014
013	• Added errata CG139 and CG140	November 2014
014	• Added errata CG141 - CG144 • Modified errata CG129	February 2015
015	• Added errata CG145 and CG146	April 2015

§ §

Preface

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

Information types defined in Nomenclature are consolidated into the specification update and are no longer published in other documents.

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

Note: Throughout this document, the Intel® Core™ i7 processor family may be referred to as "processor".

Affected documents

Document Title	Document Number / Location
Intel® Core™ i7 Processor Family for LGA2011 Socket Datasheet – Volume 1 of 2	329366
Intel® Core™ i7 Processor Family for LGA2011 Socket Datasheet – Volume 2 of 2	329367

Note: Document subtitles are:
- Supporting Desktop Intel® Core™ i7-4960X Extreme Edition Processor Series for the LGA2011 Socket
- Supporting Desktop Intel® Core™ i7-49xx and i7-48xx Processor Series for the LGA2011 Socket

Related documents

Document Title	Document Number / Location
AP-485, Intel® Processor Identification and the CPUID Instruction	241618
Intel® 64 and IA-32 Architecture Software Developer's Manual • Volume 1: Basic Architecture • Volume 2A: Instruction Set Reference Manual A-M • Volume 2B: Instruction Set Reference Manual N-Z • Volume 3A: System Programming Guide • Volume 3B: System Programming Guide • IA-32 Intel® Architecture Optimization Reference Manual	325462

Nomenclature

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

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

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

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

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

Identification Information

Component Identification using Programming Interface

The processor stepping can be identified by the following register contents.

Table 1. Processor Signature / Version

Reserved	Extended family[1]	Extended model[2]	Reserved	Processor type[3]	Family code[4]	Model number[5]	Stepping ID
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	00000000b	0011b		00b	0110b	1110b	S1 =0100

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 Intel® 386, Intel® 486, Pentium®, Pentium 4, or Intel® Core™ processor family.
2. The Extended Model, Bits [19:16] in conjunction with the Model Number, specified in Bits [7:4], are used to identify the model of the processor within the processor's family.
3. The Family Code corresponds to Bits [11:8] of the EDX register after RESET, Bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
4. The Model Number corresponds to Bits [7:4] of the EDX register after RESET, Bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
5. The Stepping ID in Bits [3:0] indicates the revision number of that model.

Component Marking

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

Figure 1. Processor Top-side Markings (example)

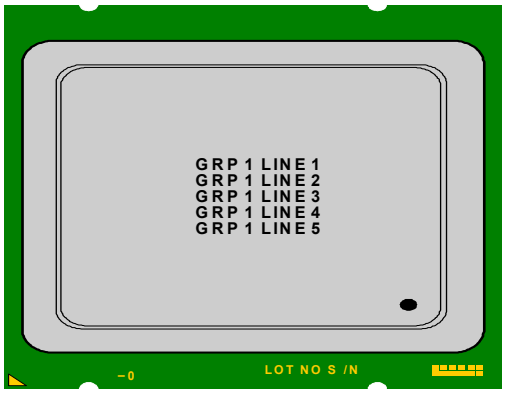

Legend	Mark Text (Production Mark):
GRP1 LINE1:	i{M}{C}YY
GRP1 LINE2:	SUB- BRAND PROC #
GRP1 LINE3:	SSPEC SPEED
GRP1 LINE4:	XXXXX
GRP1 LINE5:	{FPO} {e4}

Table 2. Processor Family Identification

S-spec number	Processor Number	Stepping	CPUID	Core Frequency (GHz)/DDR3 (MHz)/ Intel® QPI (GHz)	TDP (W)	# Cores	Cache size (MB)	Notes
SR1AS	i7-4960X	S-1	306E4h	3.6/1866	130	6	15	
SR1AU	i7-4820K	S-1	306E4h	3.7/1866	130	4	10	
SR1AT	i7-4930K	S-1	306E4h	3.4/1866	130	6	12	

Summary Tables of Changes

The following tables indicate the errata, specification changes, specification clarifications, or documentation changes that apply to the processor 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 uses the following notations:

Codes used in summary tables

Stepping

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

Page

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

Status

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

Row

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

Table 3. Errata (Sheet 1 of 5)

Number	Stepping	Status	ERRATA
	S-1		
CG1	X	No Fix	Core Frequencies at or Below the DRAM DDR Frequency May Result in Unpredictable System Behavior.
CG2	X	No Fix	Quad Rank DIMMs May Not be Properly Refreshed During IBT_OFF Mode.
CG3	X	No Fix	PCIe* TPH Attributes May Result in Unpredictable System Behavior.
CG4	X	No Fix	PCIe* Rx Common Mode Return Loss is Not Meeting the Specification.
CG5	X	No Fix	PCIe* Rx DC Common Mode Impedance is Not Meeting the Specification.
CG6	X	No Fix	QPILS Reports the VNA/VN0 Credits Available for the Processor Rx Rather Than Tx.
CG7	X	No Fix	A PECI RdPciConfigLocal Command Referencing a Non-Existent Device May Return an Unexpected Value.
CG8	X	No Fix	The Vswing of the PCIe* Transmitter Exceeds the Specification.
CG9	X	No Fix	PECI Write Requests That Require a Retry Will Always Time Out.
CG10	X	No Fix	Enabling Opportunistic Self-Refresh and Pkg C2 State Can Severely Degrade PCIe* Bandwidth.
CG11	X	No Fix	Functionally Benign PCIe* Electrical Specification Violation Compendium.
CG12	X	No Fix	Warm Resets May be Converted To Power-On Resets When Recovering From an IERR.
CG13	X	No Fix	A Modification To The Multiple Message Enable Field Does Not Affect The AER Interrupt Message Number Field.
CG14	X	No Fix	Long latency Transactions Can Cause I/O Devices On The Same Link to Time Out.
CG15	X	No Fix	Combining ROL Transactions With Non-ROL Transactions or Marker Skipping Operations May Result in a System Hang.
CG16	X	No Fix	Excessive DRAM RAPL Power Throttling May Lead to a System Hang or USB device Offlining.
CG17	X	No Fix	TSOD-Related SMBus Transactions May Not Complete When Package C-States Are Enabled.
CG18	X	No Fix	The Integrated Memory Controller Does Not Enforce CKE High For tXSDLL DCLKs After Self-Refresh.
CG19	X	No Fix	Intel® QuickData Technology DMA Suspend Does Not Transition From ARMED to HALT State.
CG20	X	No Fix	Routing Intel® High Definition Audio Traffic Through VC1 May Result in System Hang.
CG21	X	No Fix	NTB Operating in NTB/RP Mode With MSI/MSI-X Interrupts May Cause System Hang.
CG22	X	No Fix	Writes to SDOORBELL or B2BDOORBELL In Conjunction With Inbound Access to NTB MMIO Space May Hang System.
CG23	X	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a REP MOVSB or STOSB
CG24	X	No Fix	64-bit REP MOVSB/STOSB May Clear The Upper 32-bits of RCX, RDI And RSI Before Any Data is Transferred
CG25	X	No Fix	An Interrupt Recognized Prior to First Iteration of REP MOVSB/STOSB May Result EFLAGS.RF Being Incorrectly Set
CG26	X	No Fix	Instructions Retired Event May Over Count Execution of IRET Instructions
CG27	X	No Fix	An Event May Intervene Before a System Management Interrupt That Results from IN or INS
CG28	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception
CG29	X	No Fix	Unexpected #UD on VZEROALL/VZEROUPPER
CG30	X	No Fix	Successive Fixed Counter Overflows May be Discarded
CG31	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
CG32	X	No Fix	VM Exits Due to "NMI-Window Exiting" May Not Occur Following a VM Entry to the Shutdown State
CG33	X	No Fix	Execution of INVVPID Outside 64-Bit Mode Cannot Invalidate Translations For 64-Bit Linear Addresses

Table 3. Errata (Sheet 2 of 5)

Number	Stepping	Status	ERRATA
	S-1		
CG34	X	No Fix	REP MOVSB May Incorrectly Update ECX, ESI, and EDI
CG35	X	No Fix	Performance-Counter Overflow Indication May Cause Undesired Behavior
CG36	X	No Fix	VEX.L is Not Ignored with VCVT*2SI Instructions
CG37	X	No Fix	Concurrently Changing the Memory Type and Page Size May Lead to a System Hang
CG38	X	No Fix	MCI_ADDR May be Incorrect For Cache Parity Errors
CG39	X	No Fix	Instruction Fetches Page-Table Walks May be Made Speculatively to Uncacheable Memory
CG40	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
CG41	X	No Fix	The Processor May Not Properly Execute Code Modified Using A Floating-Point Store
CG42	X	No Fix	VM Exits Due to GETSEC May Save an Incorrect Value for "Blocking by STI" in the Context of Probe-Mode Redirection
CG43	X	No Fix	IA32_MC5_CTL2 is Not Cleared by a Warm Reset
CG44	X	No Fix	The Processor May Report a #TS Instead of a #GP Fault
CG45	X	No Fix	IO_SMI Indication in SMRAM State Save Area May be Set Incorrectly
CG46	X	No Fix	Performance Monitor SSE Retired Instructions May Return Incorrect Values
CG47	X	No Fix	IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception
CG48	X	No Fix	Performance Monitoring Event FP_MMX_TRANS_TO_MMX May Not Count Some Transitions
CG49	X	No Fix	General Protection Fault (#GP) for Instructions Greater than 15 Bytes May be Preempted
CG50	X	No Fix	LBR, BTS, BTM May Report a Wrong Address When an Exception/Interrupt Occurs in 64-bit Mode
CG51	X	No Fix	Incorrect Address Computed For Last Byte of FXSAVE/FXRSTOR or XSAVE/XRSTOR Image Leads to Partial Memory Update
CG52	X	No Fix	Values for LBR/BTS/BTM Will be Incorrect after an Exit from SMM
CG53	X	No Fix	EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change
CG54	X	No Fix	B0-B3 Bits in DR6 For Non-Enabled Breakpoints May Be Incorrectly Set
CG55	X	No Fix	MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error
CG56	X	No Fix	Debug Exception Flags DR6.B0-B3 Flags May Be Incorrect for Disabled Breakpoints
CG57	X	No Fix	LER MSRs May Be Unreliable
CG58	X	No Fix	Storage of PEBS Record Delayed Following Execution of MOV SS or STI
CG59	X	No Fix	PEBS Record Not Updated When in Probe Mode
CG60	X	No Fix	Faulting MMX Instruction May Incorrectly Update x87 FPU Tag Word
CG61	X	No Fix	#GP on Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code
CG62	X	No Fix	APIC Error "Received Illegal Vector" May Be Lost
CG63	X	No Fix	Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations
CG64	X	No Fix	Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures
CG65	X	No Fix	LBR, BTM or BTS Records May have Incorrect Branch From Information After an EIST/T-state/S-state/C1E Transition or Adaptive Thermal Throttling
CG66	X	No Fix	FP Data Operand Pointer May Be Incorrectly Calculated After an FP Access Which Wraps a 4-Gbyte Boundary in Code That Uses 32-Bit Address Size in 64-bit Mode
CG67	X	No Fix	VMREAD/VMWRITE Instruction May Not Fail When Accessing an Unsupported Field in VMCS
CG68	X	No Fix	An Unexpected PMI May Occur After Writing a Large Value to IA32_FIXED_CTR2

Table 3. Errata (Sheet 3 of 5)

Number	Stepping	Status	ERRATA
	S-1		
CG69	X	No Fix	A Write to the IA32_FIXED_CTR1 MSR May Result in Incorrect Value in Certain Conditions
CG70	X	No Fix	#GP May be Signaled When Invalid VEX Prefix Precedes Conditional Branch Instructions
CG71	X	No Fix	Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered
CG72	X	No Fix	PCMPESTRI, PCMPESTRM, VPCMPESTRI and VPCMPESTRM Always Operate with 32-bit Length Registers
CG73	X	No Fix	MSR_PKG_Cx_RESIDENCY MSRs May Not Be Accurate
CG74	X	No Fix	During Package Power States Repeated PCIe* and/or DMI L1 Transitions May Cause a System Hang
CG75	X	No Fix	MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang
CG76	X	No Fix	PCMPESTRI, PCMPESTRM, VPCMPESTRI and VPCMPESTRM Always Operate With 32-bit Length Registers
CG77	X	No Fix	Clock Modulation Duty Cycle Cannot Be Programmed to 6.25%
CG78	X	No Fix	Processor May Livelock During On Demand Clock Modulation
CG79	X	No Fix	Performance Monitor Counters May Produce Incorrect Results
CG80	X	No Fix	Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash
CG81	X	No Fix	IA32_FEATURE_CONTROL MSR May be Un-Initialized on a Cold Reset
CG82	X	No Fix	PEBS May Unexpectedly Signal a PMI After the PEBS Buffer is Full
CG83	X	No Fix	Execution of GETSEC[SEXIT] May Cause a Debug Exception to Be Lost
CG84	X	No Fix	An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang
CG85	X	No Fix	The Corrected Error Count Overflow Bit in IA32_ MC0_STATUS is Not Updated After a UC Error is Logged
CG86	X	No Fix	IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report the Highest Index Value Used for VMCS Encoding
CG87	X	No Fix	The Upper 32 Bits of CR3 May be Incorrectly Used With 32-Bit Paging
CG88	X	No Fix	EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly
CG89	X	No Fix	CPUID Faulting is Not Enumerated Properly
CG90	X	X	Incorrect Size Reported by PCIe* NTB BAR Registers
CG91	X	X	The Vswing of the PCIe* Transmitter Exceeds The Specification
CG92	X	X	PECI_WAKE_MODE is Always Reported as Disabled
CG93	X	X	Poisoned PCIe* AtomicOp Completions May Return an Incorrect Byte Count
CG94	X	X	PCIe* Device 3 Does Not Log an Error in UNCERRSTS When an Invalid Sequence Number in an Ack DLLP is Received
CG95	X	X	Programmable Ratio Limits For Turbo Mode is Reported as Disabled
CG96	X	X	PCIe* TLPs in Disabled VC Are Not Reported as Malformed
CG97	X	X	PCIe* Header of a Malformed TLP is Logged Incorrectly
CG98	X	X	PCIe* Link May Fail to Train to 8.0 GT/s
CG99	X	X	PCIe* May Associate Lanes That Are Not Part of Initial Link Training to L0 During Upconfiguration
CG100	X	X	Incorrect Speed and De-emphasis Level Selection During DMI Compliance Testing
CG101	X	X	Single PCIe* ACS Violation or UR Response May Result in Multiple Correctable Errors Logged
CG102	X	X	PCIe* Extended Tag Field May be Improperly Set
CG103	X	X	Power Meter May Under-Estimate Package Power
CG104	X	X	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1

Table 3. Errata (Sheet 4 of 5)

Number	Stepping	Status	ERRATA
	S-1		
CG105	X	X	PMON Counters Overflow May Not Trigger PMON Global Freeze
CG106	X	X	Processor May Log a Machine Check when MSI is signaled by a device
CG107	X	X	PECI May Not be Able to Access IIO CSRs
CG108	X	X	Spurious Patrol Scrub Errors Observed During a Warm Reset
CG109	X	X	Receiver Termination Impedance On PCIe* 3.0 Does Not Comply With The Specification
CG110	X	X	Platform Recovery After a Machine Check May Fail
CG111	X	X	PECI May be Non-responsive When System is in BMC Init Mode
CG112	X	X	Processor May Issue Unexpected NAK DLLP Upon PCIe* L1 Exit
CG113	X	X	Core C-state Residency Counters May Return Stale Data
CG114	X	X	PCIe* Ports Operating at 8 GT/s May Issue an Additional Packet After Stop and Scream Occurs
CG115	X	X	Reading DDRIO Broadcast CSRs May Return Incorrect Data
CG116	X	X	Corrected Filtering Indication May be Incorrect in LLC Machine Check Bank Status Register
CG117	X	X	A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation
CG118	X	X	RTID_POOL_CONFIG Registers Incorrectly Behave as a Read-Write Registers
CG119	X	X	Catastrophic Trip Triggered at Lower Than Expected Temperatures
CG120	X	X	PCIe* Hot Plug Slot Status Register May Not Indicate Command Completed
CG121	X	X	Misaligned PCIe* Configuration Register Writes May Not be Dropped
CG122	X	X	PCIe* Correctable Error Status Register May Not Log Receiver Error at 8.0 GT/s
CG123	X	X	Configuring PCIe* Port 3a as an NTB Disables EOI Forwarding to Port 2a
CG124	X	X	PCIe* LBMS Bit Incorrectly Set
CG125	X	X	PCIe* DLW is Not Supported When Operating at 8GT/s
CG126	X	X	Spurious Patrol Scrub Errors May Be Reported During Exit From Deep Package C-States
CG127	X	X	Local PCIe* P2P Traffic on x4 Ports May Cause a System Hang
CG128	X	X	NTB Operating In NTB/RP Mode May Complete Transactions With Incorrect ReqID
CG129	X	X	Warm Reset May Cause PCIe Hot-Plug Sequencing Failure
CG130	X	X	Performance Monitoring IA32_PERF_GLOBAL_STATUS.CondChgd Bit Not Cleared by Reset
CG131	X	X	PCIe* TLP Translation Request Errors Are Not Properly Logged For Invalid Memory Writes
CG132	X	X	Threshold-Based Status Indicator Not Updated After a UC or UCR Occurs
CG133	X	X	PCIe* Slave Loopback May Transmit Incorrect Sync Headers
CG134	X	X	PCIe* Type 1 VDMs May be Silently Dropped
CG135	X	X	Writing PCIe* Port 2A DEVCTRL May Have Side Effects When Port 2 is Bifurcated
CG136	X	X	Performance Monitor Instructions Retired Event May Not Count Consistently
CG137	X	X	Accessing SB01BASE MMIO Space in The Presence of Bi-directional NTB Traffic May Result in a System Hang
CG138	X	X	Patrol Scrubbing Doesn’t Skip Ranks Disabled After DDR Training
CG139	X	X	Multiple Intel® VT-d Translation Requests From a Single Agent May Result in a Hang
CG140	X	X	The System May Shut Down Unexpectedly During a Warm Reset
CG141	X	X	Intel® VT-d Memory Check Error on an Intel® QuickData Technology Channel May Cause All Other Channels to Master Abort
CG142	X	X	Writes To Some Control Register Bits Ignore Byte Enable

Table 3. Errata (Sheet 5 of 5)

Number	Stepping	Status	ERRATA
	S-1		
CG143	X	X	Invalid Intel® QuickData Technology XOR Descriptor Source Addressing May Lead to Unpredictable System Behavior
CG144	X	X	DDRIO SVID And SDID CSR Writes Do Not Behave as Expected
CG145	X	X	Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation
CG146	X	X	High Frequency Noise on DDR SMBus Signals May Prevent Proper Detection of Memory

Table 4. Specification Clarifications

No.	Specification clarifications
1	None

Table 5. Specification Changes

No.	Specification changes
1	None

Table 6. Documentation Changes

No.	Documentation changes
CG1	SDM, Volume 3B: On-Demand Clock Modulation Feature Clarification

§ §

Errata

CG1 Core Frequencies at or Below the DRAM DDR Frequency May Result in Unpredictable System Behavior.

Problem: The Enhanced Intel SpeedStep® Technology can dynamically adjust the core operating frequency to as low as 1200 MHz. Due to this erratum, under complex conditions and when the cores are operating at or below the DRAM DDR frequency, unpredictable system behavior may result.

Implication: Systems using Enhanced Intel SpeedStep Technology with DDR3-1333 or DDR3-1600 memory devices are subject to unpredictable system behavior.

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

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

CG2 Quad Rank DIMMs May Not be Properly Refreshed During IBT_OFF Mode.

Problem: The Integrated Memory Controller incorporates a power savings mode known as IBT_OFF (Input Buffer Termination disabled). Due to this erratum, Quad Rank DIMMs may not be properly refreshed during IBT_OFF mode.

Implication: Use of IBT_OFF mode with Quad Rank DIMMs may result in unpredictable system behavior.

Workaround: A BIOS workaround has been identified.

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

CG3 PCIe* TPH Attributes May Result in Unpredictable System Behavior.

Problem: TPH (Transactions Processing Hints) are optional aids to optimize internal processing of PCIe* transactions. Due to this erratum, certain transactions with TPH attributes may be misdirected, resulting in unpredictable system behavior.

Implication: Use of the TPH feature may affect system stability.

Workaround: A BIOS workaround has been identified.

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

CG4 PCIe* Rx Common Mode Return Loss is Not Meeting the Specification.

Problem: The PCIe* specification requires that the Rx Common Mode Return Loss in the range of 0.05 to 2.5 GHz must be limited to -6 dB. The processor's PCIe* Rx do not meet this requirement. The PCIe* Rx Common Mode Return at 500 MHz has been found to be between -3.5 and -4 dB on a limited number of samples.

Implication: Intel has not observed any functional failures due to this erratum with any commercially available PCIe* devices.

Workaround: None identified.

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

CG5 PCIe* Rx DC Common Mode Impedance is Not Meeting the Specification.

Problem: When the PCIe* Rx termination is not powered, the DC Common Mode impedance has the following requirement: ≥10 kohm over 0 to 200 mV range with respect to ground and ≥20 kohm for voltages ≥200 mV with respect to ground. The processor's PCIe* Rx do not meet this requirement at 85°C or greater. In a limited number of samples Intel has measured an impedance as low as 9.85 kohm at 50 mV.

Implication: Intel has not observed any functional impact due to this violation with any commercially available system.

Workaround: None identified.

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

CG6 QPILS Reports the VNA/VN0 Credits Available for the Processor Rx Rather Than Tx.

Problem: The QPILS register (CPUBUS(1); Devices 8,9; Function 0; Offset 0x48), according to the Intel® QuickPath Interconnect Specification at revision 1.1 and later, should report the VNA/VN0 credits available for the processor Tx (Transmit port). Due to this erratum, the QPILS register reports the VNA/VN0 credits available for the processor Rx (Receive port).

Implication: This is a violation of the specification but no functional failures have been observed due to this erratum.

Workaround: None

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

CG7 A PECI RdPciConfigLocal Command Referencing a Non-Existent Device May Return an Unexpected Value.

Problem: Configuration reads to nonexistent PCI configuration registers should return 0FFFF_FFFFH. Due to this erratum, when the PECI RdPciConfigLocal command references a nonexistent PCI configuration register, the value 0000_0000H may be returned instead of the expected 0FFFF_FFFFH.

Implication: A PECI RdPciConfigLocal command referencing a nonexistent device may observe a return value of 0000_0000H. Software expecting a return value of 0FFFF_FFFFH to identify nonexistent devices may not work as expected.

Workaround: Software that performs enumeration via the PECI "RdPciConfigLocal" command should interpret 0FFFF_FFFFH and 0000_0000H values for the Vendor Identification and Device Identification Register as indicating a nonexistent device.

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

CG8 The Vswing of the PCIe* Transmitter Exceeds the Specification.

Problem: The PCIe* specification defines a limit for the Vswing (voltage swing) of the differential lines that make up a lane to be 1200 mV peak-to-peak when operating at 2.5 GT/s and 5 GT/s. Intel has found that the processor's PCIe* transmitter may exceed this specification. Peak-to-peak swings on a limited number of samples have been observed up to 1450 mV.

Implication: For those taking direct measurements of the PCIe* transmit traffic coming from the processor may detect that the Vswing exceeds the PCIe* specification. Intel has not observed any functional failures due to this erratum.

Workaround: None identified.

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

CG9 PECI Write Requests That Require a Retry Will Always Time Out.

Problem: PECI 3.0 introduces a 'Host Identification' field as a way for the PECI host device to identify itself to the PECI client. This is intended for use in future PECI systems that may support more than one PECI originator. Since PECI 3.0 systems do not support the use of multiple originators, PECI 3.0 host devices should zero out the unused Host ID field. PECI 3.0 also introduces a 'retry' bit as a way for the PECI host to indicate to the client that the current request is a 'retry' of a previous read or write operation. Unless the PECI 3.0 host device zeroes out the byte containing the 'Host ID & Retry bit' information, PECI write requests that require a retry will never complete successfully.

Implication: PECI write requests that require a retry may never complete successfully. Instead, they will return a timeout completion code of 81H for a period ranging from 1 ms to 30 ms if the 'RETRY' bit is asserted.

Workaround: PECI 3.0 host devices should zero out the byte that contains the Host ID and Retry bit information for all PECI requests at all times including retries.

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

CG10 Enabling Opportunistic Self-Refresh and Pkg C2 State Can Severely Degrade PCIe* Bandwidth.

Problem: Due to this erratum, enabling opportunistic self-refresh can lead to the memory controller over-aggressively transitioning DRAM to self-refresh mode when the processor is in Pkg C2 state.

Implication: The PCIe* interface peak bandwidth can be degraded by as much as 90%

Workaround: A BIOS workaround has been identified.

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

CG11 Functionally Benign PCIe* Electrical Specification Violation Compendium.

Problem: Violations of PCIe* electrical specifications listed in the table below have been observed.

Specification	Violation Description
Deemphasis ratio limit: -3.5±0.5 dB	Ave: -3.8 dB, Min: -4.09 dB
At 5 GT/s operation, the receiver must tolerate AC common mode voltage of 300 mV (peak-to-peak) and must tolerate 78.1 ps jitter.	Simultaneous worst case AC common mode voltage and worst case jitter during 5 GT/s operation may result in intermittent failures leading to subsequent recovery events.
TTX-UPW-TJ (uncorrelated total pulse width jitter) maximum of 24 ps.	Samples have measured as high as 25 ps.
The Transmitter PLL bandwidth and peaking for PCIe* at 5 GT/s is either 8 to 16 MHz with 3 dB of peaking or 5 to 16 MHz with 1 dB of peaking.	Samples have measured 7.8-16 MHz with 1.3 dB of peaking.
During the LTSSM Receiver Detect State, common-mode resistance to ground is 40 to 60 ohms.	Samples have measured up to 100 ohms.
8 GT/s Receiver Stressed Eye	Samples marginally pass or fail the 10-12 BER target under stressed eye conditions.
8 GT/s PLL Bandwidth: 2 to 4 MHz with 2 dB peaking.	Samples have a measured bandwidth of up to 4.1 MHz

Implication: Intel has not observed failures from the violations listed in this erratum on any commercially available platforms and/or using commercially available PCIe* devices.

Workaround: None identified.

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

CG12 Warm Resets May be Converted To Power-On Resets When Recovering From an IERR.

Problem: When a warm reset is attempted and an IERR (Internal Error) happens as indicated by the IA32_MCi_STATUS.MCOD of 0000_0100_0000_0000, a power-on reset occurs instead.

Implication: The values in the machine check bank will be lost as a result of the power-on reset. This prevents a OS, BIOS, or the BMC (Baseboard Management Controller) from logging the content of the error registers or taking any post-reset actions that are dependent on the machine check information.

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

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

CG13 A Modification To The Multiple Message Enable Field Does Not Affect The AER Interrupt Message Number Field.

Problem: The (Advanced Error Interrupt) Message Number field (RPERRSTS Devices 0-3; Functions 0-3; Offset 178H; bits[31:27]) should be updated when the number of messages allocated to the root port is changed by writing the Multiple Message Enable field (MSIMSGCTL Device 3; Function 0; Offset 62H; bits[6:4]). However, writing the Multiple Message Enable in the root port does not update the Advanced Error Interrupt Message Number field.

Implication: Due to this erratum, software can allocate only one MSI (Message Signaled Interrupt) to the root port.

Workaround: None identified

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

CG14 Long latency Transactions Can Cause I/O Devices On The Same Link to Time Out.

Problem: Certain long latency transactions - for example, master aborts on inbound traffic, locked transactions, peer-to-peer transactions, or vendor defined messages - conveyed over the PCIe* and DMI2 interfaces can block the progress of subsequent transactions for extended periods. In certain cases, these delays may lead to I/O device timeout that can result in device error reports and/or device off-lining.

Implication: Due to this erratum, devices that generate PCIe* or DMI2 traffic characterized by long latencies can interfere with other traffic types on the same link. This may result in reduced I/O performance and device timeout errors. USB traffic can be particularly sensitive to these delays.

Workaround: Avoid the contributing conditions. This can be accomplished by separating traffic types to be conveyed on different links and/or reducing or eliminating long latency transactions.

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

CG15 Combining ROL Transactions With Non-ROL Transactions or Marker Skipping Operations May Result in a System Hang.

Problem: When Intel® QuickData Technology DMA ROL (Raid On Load) transactions and non-ROL transactions are simultaneously active, and the non-ROL address offsets are not cacheline boundary aligned, the non-ROL transaction's last partial cacheline data write may be lost leading to a system hang. In addition, when Intel® QuickData Technology DMA ROL transactions are active, marker skipping operations may lead to a system hang.

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

Workaround: None identified. When ROL transactions and non-ROL transactions are simultaneously active, all non-ROL address offsets must be aligned on cacheline boundaries. Further, marker skipping operations may not be used on any DMA channel when ROL transactions are active.

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

CG16 Excessive DRAM RAPL Power Throttling May Lead to a System Hang or USB device Offlining.

Problem: DRAM RAPL (Running Average Power Limit) is a facility for limiting the maximum power consumption of the memory subsystem. DRAM RAPL's control mechanism constrains the number of memory transactions during a particular time period. Due to this erratum, a very low power limit can throttle certain memory subsystem configurations to an extent that system failure, ranging from permanent loss of USB devices to system hangs, may result.

Implication: Using DRAM RAPL to regulate the memory subsystem power to a very low level may cause platform instability.

Workaround: It is possible for the BIOS to contain processor configuration data and code changes as a workaround for this erratum.

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

CG17 TSOD-Related SMBus Transactions May Not Complete When Package C-States Are Enabled.

Problem: The processor may not complete SMBus (System Management Bus) transactions targeting the TSOD (Temperature Sensor On DIMM) when Package C-States are enabled. Due to this erratum, if the processor transitions into a Package C-State while an SMBus transaction with the TSOD is in process, the processor will suspend receipt of the transaction. The transaction completes while the processor is in a Package C-State. Upon exiting Package C-State, the processor will attempt to resume the SMBus transaction, detect a protocol violation, and log an error.

Implication: When Package C-States are enabled, the SMBus communication error rate between the processor and the TSOD may be higher than expected.

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

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

CG18 The Integrated Memory Controller Does Not Enforce CKE High For tXSDLL DCLKs After Self-Refresh.

Problem: The JEDEC STANDARD DDR3 SDRAM Specification (No. 79-3E) requires that the CKE signal be held high for tXSDLL DCLKs after exiting self-refresh before issuing commands that require a locked DLL (Delay-Locked Loop). Due to this erratum, the Integrated Memory Controller may not meet this requirement with 512 Mb, 1 Gb, and 2 Gb devices in single rank per channel configurations.

Implication: Violating tXSDLL may result in DIMM clocking issues and may lead to unpredictable system behavior.

Workaround: A BIOS workaround has been identified.

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

CG19 Intel® QuickData Technology DMA Suspend Does Not Transition From ARMED to HALT State.

Problem: Suspending an Intel® QuickData Technology DMA channel while in the ARMED state should transition the channel to the HALT state. Due to this erratum, suspending a DMA channel while in the ARMED state does not change the state to HALT and will cause the DMA engine, when subsequently activated, to ignore the first descriptor's fence control bit and may cause the DMA engine to prematurely discard the first descriptor during the copy stage.

Implication: Suspending a DMA channel while in the ARMED state will cause the DMA engine to ignore descriptor fencing, possibly issue completion status without actually completing all descriptors, and may be subject to unexpected activation of DMA transfers.

Workaround: Check the DMA_trans_state (CHANSTS_0; Bus 0; MMIO BAR: CB_BAR [0:7]; Offset 88H; bits[2:0]) to ensure the channel state is either IDLE (001b) or ACTIVE (000b) before setting Susp_DMA (CHANCMD; Bus 0; MMIO BAR: CB_BAR [0:7]; Offset 84H;bit 2).

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

CG20 Routing Intel® High Definition Audio Traffic Through VC1 May Result in System Hang.

Problem: When bit 9 in the IIOMISCCTRL CSR (Bus 0; Device 5; Function 0; Offset 1C0H) is set, VCp inbound traffic (Intel® HD Audio) is routed through VC1 to optimize isochronous traffic performance. Due to this erratum, VC1 may not have sufficient bandwidth for all traffic routed through it; overflows may occur.

Implication: This erratum can result in lost completions that may cause a system hang.

Workaround: A BIOS workaround has been identified.

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

CG21 NTB Operating in NTB/RP Mode With MSI/MSI-X Interrupts May Cause System Hang.

Problem: The NTB (nontransparent bridge) operating in NTB/RP (NTB to Root Port mode) using Message Signaled Interrupts (MSI or MSI-X) in the presence of locks may result in a system hang.

Implication: Due to this erratum, system may hang under the condition described above.

Workaround: A BIOS workaround has been identified.

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

CG22 Writes to SDOORBELL or B2BDOORBELL In Conjunction With Inbound Access to NTB MMIO Space May Hang System.

Problem: A posted write targeting the SDOORBELL (Offset 64H) or B2BDOORBELL (Offset 140H) MMIO registers in the region define by Base Address Register PB01BASE (Bus 0; Device 3; Function 0: Offset 10H) or SB01BASE (Bus M; Device 0; Function 0; Offset 10H) may hang the system. This system hang may occur if the NTB (Non-Transparent Bridge) is processing a transaction from the secondary side of the NTB that is targeting the NTB shared MMIO registers or targeting the secondary side configuration registers when the write arrives.

Implication: The system may hang if the processor writes to the local SDOORBELL or B2BDOORBELL register at the same time that the NTB is processing an inbound transaction.

Workaround: In NTB/NTB (back-to-back) mode, do not use the B2BDOORBELL to send interrupts from the local to remote host. Instead, configure one of the following local register pairs to point to the remote SB01BASE region:

- PB23BASE (Device: 3; Function: 0; Offset: 18H) and PBAR2XLAT (Offset 10H) from PB01BASE or SB01BASE regions;
- PB45BASE (Device: 3; Function: 0; Offset: 20H) and PBAR4XLAT (Offset 18H) from PB01BASE, or SB01BASE regions;

The local host may then write directly to the PDOORBELL (Offset 60H) from the PB23BASE/PB45BASE region defined above.

In NTB/RP (bridge to root port) mode, the SDOORBELL register cannot be used by the processor on the primary side of the NTB to interrupt the processor on the secondary side. Instead, dedicate a BAR and XLAT pair, either PB23BASE/PBAR2XLAT or PB45BASE/PBAR4XLAT, to generate an interrupt directed directly into the MSI/MSIx (Message Signaled Interrupt) interrupt range on the remote processor. The device driver or client on the remote host must point the appropriate PBARnXLAT register to its MSI/MSIx interrupt range. The processor on the primary side can then write the MSI/MSIx interrupt to the dedicated BAR which will be translated by the NTB to the MSI/MSIx region of the secondary side's processor.

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

CG23 DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a REP MOVSB or STOSB

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

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

Workaround: None identified.

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

CG24 64-bit REP MOVSB/STOSB May Clear The Upper 32-bits of RCX, RDI And RSI Before Any Data is Transferred

Problem: If a REP MOVSB/STOSB is executed in 64-bit mode with an address size of 32 bits, and if an interrupt is being recognized at the start of the instruction operation, the upper 32-bits of RCX, RDI and RSI may be cleared, even though no data has yet been copied or written.

Implication: Due to this erratum, the upper 32-bits of RCX, RDI and RSI may be prematurely cleared.

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

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

CG25 An Interrupt Recognized Prior to First Iteration of REP MOVSB/STOSB May Result EFLAGS.RF Being Incorrectly Set

Problem: If a REP MOVSB/STOSB is executed and an interrupt is recognized prior to completion of the first iteration of the string operation, EFLAGS may be saved with RF=1 even though no data has been copied or stored. The Software Developer's Manual states that RF will be set to 1 for such interrupt conditions only after the first iteration is complete.

Implication: Software may not operate correctly if it relies on the value saved for EFLAGS.RF when an interrupt is recognized prior to the first iteration of a string instruction. Debug exceptions due to instruction breakpoints are delivered correctly despite this erratum; this is because the erratum occurs only after the processor has evaluated instruction-breakpoint conditions.

Workaround: Software whose correctness depends on value saved for EFLAGS.RF by delivery of the affected interrupts can disable fast-string operation by clearing Fast-String Enable in bit 0 in the IA32_MISC_ENABLE MSR (1A0H).

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

CG26 Instructions Retired Event May Over Count Execution of IRET Instructions

Problem: Under certain conditions, the performance monitoring event Instructions Retired (Event C0H, Unmask 00H) may over count the execution of IRET instruction.

Implication: Due to this erratum, performance monitoring event Instructions Retired may over count.

Workaround: None identified.

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

CG27 An Event May Intervene Before a System Management Interrupt That Results from IN or INS

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

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

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

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

CG28 Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception

Problem: The VAESIMC and VAESKEYGENASSIST instructions should produce a #UD (Invalid-Opcode) exception if the value of the vvvv field in the VEX prefix is not 1111b. Due to this erratum, if CR0.TS is "1", the processor may instead produce a #NM (Device-Not-Available) exception.

Implication: Due to this erratum, some undefined instruction encodings may produce a #NM instead of a #UD exception.

Workaround: Software should always set the vvvv field of the VEX prefix to 1111b for instances of the VAESIMC and VAESKEYGENASSIST instructions.

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

CG29 Unexpected #UD on VZEROALL/VZEROUPPER

Problem: Execution of the VZEROALL or VZEROUPPER instructions in 64-bit mode with VEX.W set to 1 may erroneously cause a #UD (invalid-opcode exception).

Implication: The affected instructions may produce unexpected invalid-opcode exceptions in 64-bit mode.

Workaround: Compilers should encode VEX.W = 0 for the VZEROALL and VZEROUPPER instructions.

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

CG30 Successive Fixed Counter Overflows May be Discarded

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

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

Workaround: Software can avoid this by:

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

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

CG31 Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception

Problem: Attempt to use FXSAVE or FXRSTOR with a VEX prefix should produce a #UD (Invalid-Opcode) exception. If either the TS or EM flag bits in CR0 are set, a #NM (device-not-available) exception will be raised instead of #UD exception.

Implication: Due to this erratum a #NM exception may be signaled instead of a #UD exception on an FXSAVE or an FXRSTOR with a VEX prefix.

Workaround: Software should not use FXSAVE or FXRSTOR with the VEX prefix.

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

CG32 VM Exits Due to "NMI-Window Exiting" May Not Occur Following a VM Entry to the Shutdown State

Problem: If VM entry is made with the "virtual NMIs" and "NMI-window exiting", VM-execution controls set to 1, and if there is no virtual-NMI blocking after VM entry, a VM exit with exit reason "NMI window" should occur immediately after VM entry unless the VM entry put the logical processor in the wait-for SIPI state. Due to this erratum, such VM exits do not occur if the VM entry put the processor in the shutdown state.

Implication: A VMM may fail to deliver a virtual NMI to a virtual machine in the shutdown state.

Workaround: Before performing a VM entry to the shutdown state, software should check whether the "virtual NMIs" and "NMI-window exiting" VM-execution controls are both 1. If they are, software should clear "NMI-window exiting" and inject an NMI as part of VM entry.

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

CG33 Execution of INVVPID Outside 64-Bit Mode Cannot Invalidate Translations For 64-Bit Linear Addresses

Problem: Executions of the INVVPID instruction outside 64-bit mode with the INVVPID type "individual-address invalidation" ignore bits 63:32 of the linear address in the INVVPID descriptor and invalidate translations for bits 31:0 of the linear address.

Implication: The INVVPID instruction may fail to invalidate translations for linear addresses that set bits in the range 63:32. Because this erratum applies only to executions outside 64-bit mode, it applies only to attempts by a 32-bit virtual-machine monitor (VMM) to invalidate translations for a 64-bit guest. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

CG34 REP MOVSB May Incorrectly Update ECX, ESI, and EDI

Problem: Under certain conditions, if the execution of a REP MOVSB instruction is interrupted, the values of ECX, ESI and EDI may contain values that represent a later point in the execution of the instruction than the actual interruption point.

Implication: Due to this erratum ECX, ESI, and EDI may be incorrectly advanced, resulting in unpredictable system behavior.

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

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

CG35 Performance-Counter Overflow Indication May Cause Undesired Behavior

Problem: Under certain conditions (listed below) when a performance counter overflows, its overflow indication may remain set indefinitely. This erratum affects the general-purpose performance counters IA32_PMC{0-7} and the fixed-function performance counters IA32_FIXED_CTR{0-2}. The erratum may occur if any of the following conditions are applied concurrent to when an actual counter overflow condition is reached:

1. Software disables the counter either globally through the IA32_PERF_GLOBAL_CTRL MSR (38FH), or locally through the IA32_PERFEVTSEL{0-7} MSRs (186H-18DH), or the IA32_FIXED_CTR_CTRL MSR (38DH).
2. Software sets the IA32_DEBUGCTL MSR (1D9H) FREEZE_PERFMON_ON_PMI bit [12].
3. The processor attempts to disable the counters by updating the state of the IA32_PERF_GLOBAL_CTRL MSR (38FH) as part of transitions such as VM exit, VM entry, SMI, RSM, or processor C-state.

Implication: Due to this erratum, the corresponding overflow status bit in IA32_PERF_GLOBAL_STATUS MSR (38DH) for an affected counter may not get cleared

when expected. If a corresponding counter is configured to issue a PMI (performance monitor interrupt), multiple PMIs may be signaled from the same overflow condition. Likewise, if a corresponding counter is configured in PEBS mode (applies to only the general purpose counters), multiple PEBS events may be signaled.

Workaround: None identified.

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

CG36 VEX.L is Not Ignored with VCVT*2SI Instructions

Problem: The VEX.L bit should be ignored for the VCVTSS2SI, VCVTSD2SI, VCVTTSS2SI, and VCVTTSD2SI instructions, however due to this erratum the VEX.L bit is not ignored and will cause a #UD.

Implication: Unexpected #UDs will be seen when the VEX.L bit is set to 1 with VCVTSS2SI, VCVTSD2SI, VCVTTSS2SI, and VCVTTSD2SI instructions.

Workaround: Software should ensure that the VEX.L bit is set to 0 for all scalar instructions.

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

CG37 Concurrently Changing the Memory Type and Page Size May Lead to a System Hang

Problem: Under a complex set of micro architectural conditions, the system may hang if software changes the memory type and page size used to translate a linear address while a TLB (Translation Lookaside Buffer) holds a valid translation for that linear address.

Implication: Due to this erratum, the system may hang. Intel has not observed this erratum with any commercially available software.

Workaround: None identified. Please refer to Software Developer's Manual, volume 3, section "Recommended Invalidation" for the proper procedure for concurrently changing page attributes and page size.

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

CG38 MCI_ADDR May be Incorrect For Cache Parity Errors

Problem: In cases when a WBINVD instruction evicts a line containing an address or data parity error (MCOD of 0x124, and MSCOD of 0x10), the address of this error should be logged in the MCi_ADDR register. Due to this erratum, the logged address may be incorrect, even though MCi_Status.ADDRV (bit 63) is set.

Implication: The address reported in MCi_ADDR may not be correct for cases of a parity error found during WBINVD execution.

Workaround: None identified.

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

CG39 Instruction Fetches Page-Table Walks May be Made Speculatively to Uncacheable Memory

Problem: Page-table walks on behalf of instruction fetches may be made speculatively to uncacheable (UC) memory.

Implication: If any paging structures are located at addresses in uncacheable memory that are used for memory-mapped I/O, such I/O operations may be invoked as a result of speculative execution that would never actually occur in the executed code path. Intel has not observed this erratum with any commercially available software.

Workaround: Software should avoid locating paging structures at addresses in uncacheable memory that are used for memory-mapped I/O.

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

CG40 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 affected steppings, see the Summary Tables of Changes.

CG41 The Processor May Not Properly Execute Code Modified Using A Floating-Point Store

Problem: Under complex internal conditions, a floating-point store used to modify the next sequential instruction may result in the old instruction being executed instead of the new instruction.

Implication: Self- or cross-modifying code may not execute as expected. Intel has not observed this erratum with any commercially available software.

Workaround: None identified. Do not use floating-point stores to modify code.

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

CG42 VM Exits Due to GETSEC May Save an Incorrect Value for "Blocking by STI" in the Context of Probe-Mode Redirection

Problem: The GETSEC instruction causes a VM exit when executed in VMX non-root operation. Such a VM exit should set bit 0 in the interruptibility-state field in the virtual-machine control structure (VMCS) if the STI instruction was blocking interrupts at the time GETSEC commenced execution. Due to this erratum, a VM exit executed in VMX non-root operation may erroneously clear bit 0 if redirection to probe mode occurs on the GETSEC instruction.

Implication: After returning from probe mode, a virtual interrupt may be incorrectly delivered prior to GETSEC instruction. Intel has not observed this erratum with any commercially software.

Workaround: None identified.

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

CG43 IA32_MC5_CTL2 is Not Cleared by a Warm Reset

Problem: IA32_MC5_CTL2 MSR (285H) is documented to be cleared on any reset. Due to this erratum this MSR is only cleared upon a cold reset.

Implication: The algorithm documented in Software Developer's Manual, Volume 3, section titled "CMCI Initialization" or any other algorithm that counts the IA32_MC5_CTL2 MSR being cleared on reset will not function as expected after a warm reset.

Workaround: None identified.

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

CG44 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 affected steppings, see the Summary Tables of Changes.

CG45 IO_SMI Indication in SMRAM State Save Area May be Set Incorrectly

Problem: he 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
- A 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: or the affected steppings, see the Summary Tables of Changes.

CG46 Performance Monitor SSE Retired Instructions May Return Incorrect Values

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

Implication: Performance Monitoring counter SIMD_INST_RETIRED may report count higher than expected.

Workaround: None identified.

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

CG47 IRET under Certain Conditions May Cause an Unexpected Alignment Check Exception

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

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

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

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

CG48 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 under counting 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 affected steppings, see the Summary Tables of Changes.

CG49 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 affected steppings, see the Summary Tables of Changes.

CG50 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 (0x00007FFFFFFFFFFF) in 64-bit mode, the LBR return registers will save a wrong return address with bits 63 to 48 incorrectly sign extended to all 1's. Subsequent BTS and BTM operations which report the LBR will also be incorrect.

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

Workaround: None identified.

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

CG51 Incorrect Address Computed For Last Byte of FXSAVE/FXRSTOR or XSAVE/XRSTOR Image Leads to Partial Memory Update

Problem: A partial memory state save of the FXSAVE or XSAVE image or a partial memory state restore of the FXRSTOR or XRSTOR 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 or XSAVE/XRSTOR 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 affected steppings, see the Summary Tables of Changes.

CG52 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 affected steppings, see the Summary Tables of Changes.

CG53 EFLAGS Discrepancy on Page Faults and on EPT-Induced VM Exits after a Translation Change

Problem: This erratum is regarding the case where paging structures are modified to change a linear address from writable to non-writable without software performing an appropriate TLB invalidation. When a subsequent access to that address by a specific instruction (ADD, AND, BTC, BTR, BTS, CMPXCHG, DEC, INC, NEG, NOT, OR, ROL/ROR, SAL/SAR/SHL/SHR, SHLD, SHRD, SUB, XOR, and XADD) causes a page fault or an EPT-induced VM exit, the value saved for EFLAGS may incorrectly contain the arithmetic flag values that the EFLAGS register would have held had the instruction completed without fault or VM exit. For page faults, this can occur even if the fault causes a VM exit or if its delivery causes a nested fault.

Implication: None identified. Although the EFLAGS value saved by an affected event (a page fault or an EPT-induced VM exit) may contain incorrect arithmetic flag values, Intel has not identified software that is affected by this erratum. This erratum will have no further effects once the original instruction is restarted because the instruction will produce the same results as if it had initially completed without fault or VM exit.

Workaround: If the handler of the affected events inspects the arithmetic portion of the saved EFLAGS value, then system software should perform a synchronized paging structure modification and TLB invalidation.

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

CG54 B0-B3 Bits in DR6 For Non-Enabled Breakpoints May Be Incorrectly Set

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

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

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

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

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

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

CG55 MCi_Status Overflow Bit May Be Incorrectly Set on a Single Instance of a DTLB Error

Problem: A single Data Translation Look Aside Buffer (DTLB) error can incorrectly set the Overflow (bit [62]) in the MCi_Status register. A DTLB error is indicated by M error code (bits [15:0]) appearing as binary value, 000x 0000 0001 0100, in the MCi_Status register.

Implication: Due to this erratum, the Overflow bit in the MCi_Status register may not be an accurate indication of multiple occurrences of DTLB errors. There is no other impact to normal processor functionality.

Workaround: None identified.

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

CG56 Debug Exception Flags DR6.B0-B3 Flags May Be Incorrect for Disabled Breakpoints

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

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

Workaround: None identified.

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

CG57 LER MSRs May Be Unreliable

Problem: Due to certain internal processor events, updates to the LER (Last Exception Record) MSRs, MSR_LER_FROM_LIP (1DDH) and MSR_LER_TO_LIP (1DEH), may happen when no update was expected.

Implication: The values of the LER MSRs may be unreliable.

Workaround: None Identified.

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

CG58 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 affected steppings, see the Summary Tables of Changes.

CG59 PEBS Record Not Updated When in Probe Mode

Problem: When a performance monitoring counter is configured for PEBS (Precise Event Based Sampling), overflows of the counter can result in storage of a PEBS record in the PEBS buffer. Due to this erratum, if the overflow occurs during probe mode, it may be ignored and a new PEBS record may not be added to the PEBS buffer.

Implication: Due to this erratum, the PEBS buffer may not be updated by overflows that occur during probe mode.

Workaround: None identified.

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

CG60 Faulting MMX Instruction May Incorrectly Update x87 FPU Tag Word

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

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

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

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

Workaround: None identified.

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

CG61 #GP on Segment Selector Descriptor That Straddles Canonical Boundary May Not Provide Correct Exception Error Code

Problem: During a #GP (General Protection Exception), the processor pushes an error code on to the exception handler's stack. If the segment selector descriptor straddles the canonical boundary, the error code pushed onto the stack may be incorrect.

Implication: An incorrect error code may be pushed onto the stack. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

CG62 APIC Error "Received Illegal Vector" May Be Lost

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

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

Workaround: None identified.

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

CG63 Changing the Memory Type for an In-Use Page Translation May Lead to Memory-Ordering Violations

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

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

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

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

CG64 Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures

Problem: Bits 53:50 of the IA32_VMX_BASIC MSR report the memory type that the processor uses to access the VMCS and data structures referenced by pointers in the VMCS. Due to this erratum, a VMX access to the VMCS or referenced data structures will instead use the memory type that the MTRRs (memory-type range registers) specify for the physical address of the access.

Implication: Bits 53:50 of the IA32_VMX_BASIC MSR report that the WB (write-back) memory type will be used but the processor may use a different memory type.

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

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

CG65 LBR, BTM or BTS Records May have Incorrect Branch From Information After an EIST/T-state/S-state/C1E Transition or Adaptive Thermal Throttling

Problem: The "From" address associated with the LBR (Last Branch Record), BTM (Branch Trace Message) or BTS (Branch Trace Store) may be incorrect for the first branch after a transition of:

- EIST (Enhanced Intel® SpeedStep Technology)
- T-state (Thermal Monitor states)
- S1-state (ACPI package sleep state)
- C1E (Enhanced C1 Low Power state)
- Adaptive Thermal Throttling

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

Workaround: None identified.

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

CG66 FP Data Operand Pointer May Be Incorrectly Calculated After an FP Access Which Wraps a 4-Gbyte Boundary in Code That Uses 32-Bit Address Size in 64-bit Mode

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

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

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

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

CG67 VMREAD/VMWRITE Instruction May Not Fail When Accessing an Unsupported Field in VMCS

Problem: The Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 2B states that execution of VMREAD or VMWRITE should fail if the value of the instruction's register source operand corresponds to an unsupported field in the VMCS (Virtual Machine Control Structure). The correct operation is that the logical processor will set the ZF (Zero Flag), write 0CH into the VM-instruction error field and for VMREAD leave the instruction's destination operand unmodified. Due to this erratum, the instruction may instead clear the ZF, leave the VM-instruction error field unmodified and for VMREAD modify the contents of its destination operand.

Implication: Accessing an unsupported field in VMCS will fail to properly report an error. In addition, VMREAD from an unsupported VMCS field may unexpectedly change its destination operand. Intel has not observed this erratum with any commercially available software.

Workaround: Software should avoid accessing unsupported fields in a VMCS.

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

CG68 An Unexpected PMI May Occur After Writing a Large Value to IA32_FIXED_CTR2

Problem: If the fixed-function performance counter IA32_FIXED_CTR2 MSR (30BH) is configured to generate a performance-monitor interrupt (PMI) on overflow and the counter's value is greater than FFFFFFFFFFC0H, then this erratum may incorrectly cause a PMI if software performs a write to this counter.

Implication: A PMI may be generated unexpectedly when programming IA32_FIXED_CTR2. Other than the PMI, the counter programming is not affected by this erratum as the attempted write operation does succeed.

Workaround: None identified.

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

CG69 A Write to the IA32_FIXED_CTR1 MSR May Result in Incorrect Value in Certain Conditions

Problem: Under specific internal conditions, if software tries to write the IA32_FIXED_CTR1 MSR (30AH) a value that has all bits [31:1] set while the counter was just about to overflow when the write is attempted (i.e. its value was 0xFFFF FFFF FFFF), then due to this erratum the new value in the MSR may be corrupted.

Implication: Due to this erratum, IA32_FIXED_CTR1 MSR may be written with a corrupted value.

Workaround: Software may avoid this erratum by writing zeros to the IA32_FIXED_CTR1 MSR, before the desired write operation.

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

CG70 #GP May be Signaled When Invalid VEX Prefix Precedes Conditional Branch Instructions

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

Implication: Due to this erratum, #GP fault instead of a #UD may be signaled on an illegal instruction.

Workaround: None identified.

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

CG71 Interrupt From Local APIC Timer May Not Be Detectable While Being Delivered

Problem: If the local-APIC timer's CCR (current-count register) is 0, software should be able to determine whether a previously generated timer interrupt is being delivered by first reading the delivery-status bit in the LVT timer register and then reading the bit in the IRR (interrupt-request register) corresponding to the vector in the LVT timer register. If both values are read as 0, no timer interrupt should be in the process of being delivered. Due to this erratum, a timer interrupt may be delivered even if the CCR is 0 and the LVT and IRR bits are read as 0. This can occur only if the DCR (Divide Configuration Register) is greater than or equal to 4. The erratum does not occur if software writes zero to the Initial Count Register before reading the LVT and IRR bits.

Implication: Software that relies on reads of the LVT and IRR bits to determine whether a timer interrupt is being delivered may not operate properly.

Workaround: Software that uses the local-APIC timer must be prepared to handle the timer interrupts, even those that would not be expected based on reading CCR and the LVT and IRR bits; alternatively, software can avoid the problem by writing zero to the Initial Count Register before reading the LVT and IRR bits.

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

CG72 PCMPESTRI, PCMPESTRM, VPCMPESTRI and VPCMPESTRM Always Operate with 32-bit Length Registers

Problem: In 64-bit mode, using REX.W=1 with PCMPESTRI and PCMPESTRM or VEX.W=1 with VPCMPESTRI and VPCMPESTRM should support a 64-bit length operation with RAX/RDX. Due to this erratum, the length registers are incorrectly interpreted as 32-bit values.

Implication: Due to this erratum, using REX.W=1 with PCMPESTRI and PCMPESTRM as well as VEX.W=1 with VPCMPESTRI and VPCMPESTRM do not result in promotion to 64-bit length registers.

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

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

CG73 MSR_PKG_Cx_RESIDENCY MSRs May Not Be Accurate

Problem: If the processor is in a package C-state for an extended period of time (greater than 40 seconds) with no wake events, the value in the MSR_PKG_C{2,3,6,7}_RESIDENCY MSRs (60DH and 3F8H–3FAH) will not be accurate.

Implication: Utilities that report C-state residency times will report incorrect data in cases of long duration package C-states.

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

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

CG74 During Package Power States Repeated PCIe* and/or DMI L1 Transitions May Cause a System Hang

Problem: Under a complex set of internal conditions and operating temperature, when the processor is in a deep power state (package C3, C6 or C7) and the PCIe and/or DMI links are toggling in and out of L1 state, the system may hang.

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

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

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

CG75 MONITOR or CLFLUSH on the Local XAPIC's Address Space Results in Hang

Problem: If the target linear address range for a MONITOR or CLFLUSH is mapped to the local xAPIC's address space, the processor will hang.

Implication: When this erratum occurs, the processor will hang. The local xAPIC's address space must be uncached. The MONITOR instruction only functions correctly if the specified linear address range is of the type write-back. CLFLUSH flushes data from the cache. Intel has not observed this erratum with any commercially available software.

Workaround: Do not execute MONITOR or CLFLUSH instructions on the local xAPIC address space.

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

CG76 PCMPESTRI, PCMPESTRM, VPCMPESTRI and VPCMPESTRM Always Operate With 32-bit Length Registers

Problem: In 64-bit mode, using REX.W=1 with PCMPESTRI and PCMPESTRM or VEX.W=1 with VPCMPESTRI and VPCMPESTRM should support a 64-bit length operation with RAX/RDX. Due to this erratum, the length registers are incorrectly interpreted as 32-bit values.

Implication: Due to this erratum, using REX.W=1 with PCMPESTRI and PCMPESTRM as well as VEX.W=1 with VPCMPESTRI and VPCMPESTRM do not result in promotion to 64-bit length registers.

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

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

CG77 Clock Modulation Duty Cycle Cannot Be Programmed to 6.25%

Problem: When programming field T_STATE_REQ of the IA32_CLOCK_MODULATION MSR (19AH) bits [3:0] to '0001, the actual clock modulation duty cycle will be 12.5% instead of the expected 6.25% ratio.

Implication: Due to this erratum, it is not possible to program the clock modulation to a 6.25% duty cycle.

Workaround: None Identified.

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

CG78 Processor May Livelock During On Demand Clock Modulation

Problem: The processor may livelock when (1) a processor thread has enabled on demand clock modulation via bit 4 of the IA32_CLOCK_MODULATION MSR (19AH) and the clock modulation duty cycle is set to 12.5% (02H in bits 3:0 of the same MSR), and (2) the other processor thread does not have on demand clock modulation enabled and that thread is executing a stream of instructions with the lock prefix that either split a cache line or access UC memory.

Implication: Program execution may stall on both threads of the core subject to this erratum.

Workaround: This erratum will not occur if clock modulation is enabled on all threads when using on demand clock modulation or if the duty cycle programmed in the IA32_CLOCK_MODULATION MSR is 18.75% or higher.

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

CG79 Performance Monitor Counters May Produce Incorrect Results

Problem: When operating with SMT enabled, a memory at-retirement performance monitoring event (from the list below) may be dropped or may increment an enabled event on the corresponding counter with the same number on the physical core's other thread rather than the thread experiencing the event. Processors with SMT disabled in BIOS are not affected by this erratum.

The list of affected memory at-retirement events is as follows:

MEM_UOP_RETIRED.LOADS
MEM_UOP_RETIRED.STORES
MEM_UOP_RETIRED.LOCK
MEM_UOP_RETIRED.SPLIT
MEM_UOP_RETIRED.STLB_MISS
MEM_LOAD_UOPS_RETIRED.HIT_LFB
MEM_LOAD_UOPS_RETIRED.L1_HIT
MEM_LOAD_UOPS_RETIRED.L2_HIT
MEM_LOAD_UOPS_RETIRED.LLC_HIT
MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HIT
MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_HITM
MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_MISS
MEM_LOAD_UOPS_LLC_HIT_RETIRED.XSNP_NONE
MEM_LOAD_UOPS_RETIRED.LLC_MISS
MEM_LOAD_UOPS_LLC_MISS_RETIRED.LOCAL_DRAM
MEM_LOAD_UOPS_LLC_MISS_RETIRED.REMOTE_DRAM
MEM_LOAD_UOPS_RETIRED.L2_MISS

Implication: Due to this erratum, certain performance monitoring event may produce unreliable results when SMT is enabled.

Workaround: None Identified.

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

CG80 Virtual-APIC Page Accesses With 32-Bit PAE Paging May Cause a System Crash

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

Implication: This erratum may result in unexpected faults, an uncorrectable TLB error logged in IA32_MCi_STATUS.MCACOD (bits [15:0]) with a value of 0000_0000_0001_xxxxb (where x stands for 0 or 1), a guest or hypervisor crash, or other unpredictable system behavior.

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

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

CG81 IA32_FEATURE_CONTROL MSR May be Un-Initialized on a Cold Reset

Problem: IA32_FEATURE_CONTROL MSR (3Ah) may have random values after RESET (including the reserved and Lock bits), and the read-modify-write of the reserved bits and/or the Lock bit being incorrectly set may cause an unexpected GP fault.

Implication: Due to this erratum, an unexpected GP fault may occur and BIOS may not complete initialization.

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

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

CG82 PEBS May Unexpectedly Signal a PMI After the PEBS Buffer is Full

Problem: The Software Developer's Manual states that no PMI should be generated when PEBS index reaches PEBS Absolute Maximum. Due to this erratum a PMI may be generated even though the PEBS buffer is full.

Implication: PEBS may trigger a PMI even though the PEBS index has reached the PEBS Absolute Maximum.

Workaround: None identified.

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

CG83 Execution of GETSEC[SEXIT] May Cause a Debug Exception to Be Lost

Problem: A debug exception occurring at the same time that GETSEC[SEXIT] is executed or when an SEXIT doorbell event is serviced may be lost.

Implication: Due to this erratum, there may be a loss of a debug exception when it happens concurrently with the execution of GETSEC[SEXIT]. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

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

CG84 An Uncorrectable Error Logged in IA32_CR_MC2_STATUS May also Result in a System Hang

Problem: Uncorrectable errors logged in IA32_CR_MC2_STATUS MSR (409H) may also result in a system hang causing an Internal Timer Error (MCACOD = 0x0400h) to be logged in another machine check bank (IA32_MCi_STATUS).

Implication: Uncorrectable errors logged in IA32_CR_MC2_STATUS can further cause a system hang and an Internal Timer Error to be logged.

Workaround: None identified.

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

CG85 The Corrected Error Count Overflow Bit in IA32_ MC0_STATUS is Not Updated After a UC Error is Logged

Problem: When a UC (uncorrected) error is logged in the IA32_MC0_STATUS MSR (401H), corrected errors will continue to update the lower 14 bits (bits 51:38) of the Corrected Error Count. Due to this erratum, the sticky count overflow bit (bit 52) of the Corrected Error Count will not get updated after a UC error is logged.

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

Workaround: None identified.

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

CG86 IA32_VMX_VMCS_ENUM MSR (48AH) Does Not Properly Report the Highest Index Value Used for VMCS Encoding

Problem: IA32_VMX_VMCS_ENUM MSR (48AH) bits 9:1 report the highest index value used for any VMCS encoding. Due to this erratum, the value 21 is returned in bits 9:1 although there is a VMCS field whose encoding uses the index value 23.

Implication: Software that uses the value reported in IA32_VMX_VMCS_ENUM[9:1] to read and write all VMCS fields may omit one field.

Workaround: None identified.

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

CG87 The Upper 32 Bits of CR3 May be Incorrectly Used With 32-Bit Paging

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

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

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

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

CG88 EPT Violations May Report Bits 11:0 of Guest Linear Address Incorrectly

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

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

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

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

CG89 CPUID Faulting is Not Enumerated Properly

Problem: A processor that supports the CPUID-faulting feature enumerates this capability by setting PLATFORM_INFO MSR (CEH) bit 31. Due to this erratum, the processor will erroneously clear this bit.

Implication: Software that depends upon CPUID faulting will incorrectly determine that the processor does not support the feature.

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

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

CG90 Incorrect Size Reported by PCIe* NTB BAR Registers

Problem: The PCIe NTB (Non-Transparent Bridge) BARs (Base Address Register) at PB23BASE (Bus 0; Device 3; Function 0; Offset 0x18) and PB45BASE ((Bus 0; Device 3; Function 0; Offset 0x20) are used to determine the size of the requested memory region. Due to this erratum, these registers return incorrect values.

Implication: When this erratum occurs, incorrect memory region size(s) can lead to overlapping BAR regions.

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

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

CG91 The Vswing of the PCIe* Transmitter Exceeds The Specification

Problem: The PCIe Specification defines a limit for the Vswing (Voltage Swing) of the differential lines that make up a lane to be 1200 mV peak-to-peak when operating at 2.5 GT/s and 5 GT/s. Intel has found that the processor's PCIe transmitter may exceed this specification. Peak-to-peak swings on a limited number of samples have been observed up to 1450 mV.

Implication: For those taking direct measurements of the PCIe transmit traffic coming from the processor may detect that the Vswing exceeds the PCIe Specification. Intel has not observed any functional failures due to this erratum.

Workaround: None identified.

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

CG92 PECI_WAKE_MODE is Always Reported as Disabled

Problem: Due to this erratum, the state of PECI_WAKE_MODE is always reported as disabled. The PECI (Platform Environment Control Interface) PCS (Package Configuration Service) WRITE_PECI_WAKE_MODE (0x5) command correctly updates the state of PECI_WAKE_MODE, but the PECI PCS READ_PECI_WAKE_MODE (0x5) always reports the PECI_WAKE_MODE as disabled.

Implication: Software depending on the reported value for PECI_WAKE_MODE may not behave as expected.

Workaround: None identified.

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

CG93 Poisoned PCIe* AtomicOp Completions May Return an Incorrect Byte Count

Problem: A poisoned PCIe AtomicOp request completion may have an incorrect byte count.

Implication: When this erratum occurs, PCIe devices which enable byte count checking will log an unexpected completion and issue a CTO (Completion Time Out).

Workaround: None identified.

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

CG94 PCIe* Device 3 Does Not Log an Error in UNCERRSTS When an Invalid Sequence Number in an Ack DLLP is Received

Problem: If the processor‵s PCIe device 3 controller receives an invalid sequence number in an Ack DLLP (Data Link Layer Packet), it is expected to log an uncorrectable error for the affected port in bit [4] of the UNCERRSTS register (Bus 0; Device 3; Function 3:0; Offset 14CH). Due to this erratum, no data link protocol error is logged when an invalid sequence number in an Ack DLLP occurs on PCIe device 3.

Implication: Software that uses this register upon an uncorrectable PCIe error will not be able to identify this specific error type.

Workaround: None identified.

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

CG95 Programmable Ratio Limits For Turbo Mode is Reported as Disabled

Problem: The Programmable Ratio Limits for Turbo Mode in bit 28 of the MSR_PLATFORM_INFO MSR (CEH) should be 1 but, due to this erratum, it is reported as 0.

Implication: Due to this erratum, software will incorrectly assume it cannot dynamically vary the factory configured ratio limit values specified in MSR_TURBO_RATIO_LIMIT MSR (1ADH) and MSR_TURBO_RATIO_LIMIT1 MSR (1AEH).

Workaround: Software should treat the Programmable Ratio Limits for Turbo Mode bit as set.

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

CG96 PCIe* TLPs in Disabled VC Are Not Reported as Malformed

Problem: The PCIe Base Specification requires processors to report a TLP (Transaction Layer Packet) with a TC (Traffic Class) that is not mapped to any enabled VC (Virtual Channel) in an Ingress Port as a Malformed TLP. Due to this erratum, a TLP received on the DMI port that not is mapped to an enabled VC is not reported as a Malformed TLP.

Implication: Receipt of a TLP with an unmapped PCIe TC may lead to completion time out events or other unexpected system behavior. Intel has not observed this erratum with any commercially available software or platform.

Workaround: None identified.

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

CG97 PCIe* Header of a Malformed TLP is Logged Incorrectly

Problem: If a PCIe port receives a malformed TLP (Transaction Layer Packet), an error is logged in the UNCERRSTS register (Device 0; Function 0; Offset 14CH and Device 2-3; Function 0-3; Offset 14CH). Due to this erratum, the header of the malformed TLP is logged incorrectly in the HDRLOG register (Device 0; Function 0; Offset 164H and Device 2-3; Function 0-3; Offset 164H).

Implication: The PCIe header of a malformed TLP is not logged correctly.

Workaround: None identified.

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

CG98 PCIe* Link May Fail to Train to 8.0 GT/s

Problem: Due to this erratum, with certain 8.0 GT/s-capable link partners, the PCIe link may fail to train to 8.0 GT/s as requested.

Implication: When this erratum occurs, the PCIe link will enter an infinite speed-change request loop.

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

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

CG99 PCIe* May Associate Lanes That Are Not Part of Initial Link Training to L0 During Upconfiguration

Problem: The processor should not associate any lanes that were not part of the initial link training in subsequent upconfiguration requests from an endpoint. Due to this erratum, the processor may associate any Lane that has exited Electrical Idle, even if it is beyond the width of the initial Link training.

Implication: Upconfiguration requests may result in a Link wider than the initially-trained Link.

Workaround: Endpoints must ensure that upconfiguration requests do not request a Link width wider than that negotiated during initial Link training

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

CG100 Incorrect Speed and De-emphasis Level Selection During DMI Compliance Testing

Problem: When the DMI port is operating as a PCIe* port, it supports only 2.5GT/s and 5GT/s data rates. According to the PCIe specification, the data rate and de-emphasis level for the compliance patterns should be based on the maximum data rate supported. Due to this erratum, the port may select an 8GT/s data rate and associated de-emphasis level during compliance testing mode.

Implication: When doing PCIe load board compliance testing, the DMI port may transmit using 8GT/s data rate and de-emphasis levels.

Workaround: None identified.

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

CG101 Single PCIe* ACS Violation or UR Response May Result in Multiple Correctable Errors Logged

Problem: An ACS (Access Control Services) error or UR (Unsupported Request) PCIe completion status can trigger a LER (Live Error Recovery) if they are unmasked in the LER_UNCERRMSK (Bus 0; Device 0/1/2/3; Function 0/0-1/0-3/0-3; Offset 28CH) and LER_XPUNCERRMSK (Bus 0; Device 0/1/2/3; Function 0/0-1/0-3/0-3; Offset 290H) CSRs, respectively. Due to this erratum, the Root Port Error Status "multiple_correctable_error_received" bit (RPERRSTS[1], CPUBUSNO(0), Device 0:3, Functions 0/0-1/0-3/0-3, Offset 0x178) may be set upon on a single ACS or UR error.

Implication: PCIe error handling software may not behave as expected after an ACS error or a UR completion status. Intel has not observed this erratum with any commercially available software or system.

Workaround: None identified.

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

CG102 PCIe* Extended Tag Field May be Improperly Set

Problem: The Extended Tag field in the TLP Header will not be zero for TLPs issued by PCIe ports 1a, 1b, 2c, 2d, 3c, and 3d even when the Extended Tag Field Enable bit in the Device Control Register (Offset 08H, bit 8) is 0.

Implication: This erratum does not affect ports 0, 2a, 2b, 3a and 3b. This erratum will not result in any functional issues when using device that properly track and return the full 8 bit Extended Tag value with the affected ports. However, if the Extended Tag field is not returned by a device connected to an affected port then this erratum may result in unexpected completions and completion timeouts.

Workaround: None identified.

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

CG103 Power Meter May Under-Estimate Package Power

Problem: Power Meter provides a real-time power consumption estimate for the processor. Depending on operating conditions and variations in certain component-specific characteristics, the reported power may be below the actual power consumption.

Implication: Due to Intel® Turbo Boost Technology using the Power Meter to compare instantaneous power consumption to the rated TDP, the core frequency in P0 may set to a ratio where the processor exceeds its rated TDP. Further, using the average power limit facility (RAPL) may cause the processor to run at a power consumption level that is higher than expected.

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

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

CG104 VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1

Problem: When "XD Bit Disable" in the IA32_MISC_ENABLE MSR (1A0H) bit 34 is set to 1, it should not be possible to enable the "execute disable" feature by setting IA32_EFER.NXE. Due to this erratum, a VM exit that occurs with the 1-setting of the "load IA32_EFER" VM-exit control may set IA32_EFER.NXE even if IA32_MISC_ENABLE bit 34 is set to 1. This erratum can occur only if IA32_MISC_ENABLE bit 34 was set by guest software in VMX non-root operation.

Implication: Software in VMX root operation may execute with the "execute disable" feature enabled despite the fact that the feature should be disabled by the IA32_MISC_ENABLE MSR. Intel has not observed this erratum with any commercially available software.

Workaround: A virtual-machine monitor should not allow guest software to write to the IA32_MISC_ENABLE MSR.

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

CG105 PMON Counters Overflow May Not Trigger PMON Global Freeze

Problem: In the event of a R2PCIe / R3QPI PMON register PMONCTRSTATUS0 (Device 19; Function 1; Offset 0xF8) and PMONCTRSTATUS{0:1} (Device 18,19; Function 5,6; Offset 0xF8) overflow, the PMON unit sends an overflow message to a global PMON manager. The global PMON manager then asserts the global freeze signal and disables all of its counters. Due to this erratum, the global PMON manager will not assert the global freeze signal and disable all of its counters.

Implication: PMON counters may fail to freeze when either of R2PCIe /R3QPI PMON Counters Overflow. The counts on PMON unit may not be accurate.

Workaround: Upon receipt of an overflow message, software should set bit 8 of the PMONUNITCTRL0 (Device 19; Function 1; Offset 0xF4) and PMONUNITCTRL{0:1} (Device 18,19; Function 5,6; Offset 0xF4) to freeze the counter.

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

CG106 Processor May Log a Machine Check when MSI is signaled by a device

Problem: In certain cases, when MSI (Message Signaled Interrupt) is signaled by a device, a machine check error may be logged in the IA32_MC4_STATUS MSR (411H) with MSCOD field bits [31:24] equal to values of 73H or 74H.

Implication: A machine check may be observed when MSI is signaled by a device.

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

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

CG107 PECI May Not be Able to Access IIO CSRs

Problem: Due to this erratum, when the processor has viral enabled and an uncorrectable error occurs in the core, PECI (Platform Environment Control Interface) may not be able to access IIO (Integrated I/O) CSRs.

Implication: When this erratum occurs, IIO CSR access using a PECI RdPCIConfigLocal() or WrPCIConfigLocal() command will return a status of 91H, indicating that the request could not be processed.

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

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

CG108 Spurious Patrol Scrub Errors Observed During a Warm Reset

Problem: The patrol scrub engine continues to run during a warm reset; this can lead to spurious errors being reported by the Memory Controller while memory is in Self Refresh.

Implication: Due to this erratum, erroneous patrol scrub errors may be observed during a warm reset.

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

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

CG109 Receiver Termination Impedance On PCIe* 3.0 Does Not Comply With The Specification

Problem: The PCIe* Base Specification revision 3.0 defines ZRX-HIGH-IMP-DC-NEG and ZRX-HIGH-IMP-DC-POS for termination impedance of the receiver. The specified impedance for a negative voltage (-150 mV to 0V) is expected to be greater than 1 Kohm. Sampled measurements of this impedance as low as 400 ohms have been seen. The specified impedance for a positive voltage (> 200 mV) is greater than 20 Kohms. Sampled measurements of this impedance as low as 14.6 Kohms have been seen.

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

Workaround: None identified.

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

CG110 Platform Recovery After a Machine Check May Fail

Problem: While attempting platform recovery after a machine check (as indicated by CATERR# signaled from the legacy socket), the original error condition may prevent normal platform recovery which can lead to a second machine check. A remote processor detecting a second Machine Check Event will hang immediately

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

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

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

CG111 PECI May be Non-responsive When System is in BMC Init Mode

Problem: The allow_peci_pcode_error_rsp field in the DYNAMIC_PERF_POWER_CTL CSR (Device 10; Function 2; Offset 0x64H; bit 16) does not retain its value after a warm reset. When the system is in BMC Init mode, this erratum can cause PECI (Platform Environment Control Interface) access to be non-responsive after a warm reset caused by a Machine Check Event.

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

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

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

CG112 Processor May Issue Unexpected NAK DLLP Upon PCIe* L1 Exit

Problem: Upon exiting the L1 link power state, the processor's PCIe port may unexpectedly issue a NAK DLLP (Data Link Layer Packet).

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

Workaround: None identified.

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

CG113 Core C-state Residency Counters May Return Stale Data

Problem: Updates to the Core C-state residency counter MSRs occur after a core wakeup has completed. Reads of these MSRs which occur between a core wakeup and the subsequent Core C-state residency counter updates will return the pre-sleep values. The affected MSRs are MSR_CORE_C3_RESIDENCY (3FCH), MSR_CORE_C6_RESIDENCY (3FDH), and MSR_CORE_C7_RESIDENCY (3FEH).

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

Workaround: None identified.

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

CG114 PCIe* Ports Operating at 8 GT/s May Issue an Additional Packet After Stop and Scream Occurs

Problem: The processor's Stop and Scream mode requires that an attempt to send a poisoned packet results in that poisoned packet and all subsequent packets being dropped. When operating a PCIe link at 8 GT/s, Stop and Scream will drop the attempted poisoned packet but, due to this erratum, one additional good packet may be sent.

Implication: Intel has not observed functional failures from this erratum on any commercially available platforms using any commercially available PCIe device.

Workaround: None identified.

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

CG115 Reading DDRIO Broadcast CSRs May Return Incorrect Data

Problem: Device 31, Functions 6,7 are DDRIO broadcast CSRs; writing one broadcast CSR writes the corresponding CSR in each of the four memory channels (Individual memory channel CSRs are located at Device 17,31; Functions 0,1,4,5). Due to this erratum, reading DDRIO broadcast CSRs may return incorrect data.

Implication: When this erratum occurs, software that reads broadcast CSRs may behave unexpectedly.

Workaround: None identified.

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

CG116 Corrected Filtering Indication May be Incorrect in LLC Machine Check Bank Status Register

Problem: The Corrected Filtering indication in IA32_MC{17-28}_STATUS.MCACOD bit 12 may be calculated incorrectly under some conditions.

Implication: Software using the Corrected Filtering indication in IA32_MC{17-28}_STATUS.MCACOD may not function as expected.

Workaround: None identified.

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

CG117 A MOV to CR3 When EPT is Enabled May Lead to an Unexpected Page Fault or an Incorrect Page Translation

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

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

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

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

CG118 RTID_POOL_CONFIG Registers Incorrectly Behave as a Read-Write Registers

Problem: The RTID_POOL_CONFIG CSRs (Device 12; Function 0-5; Offset ACH and Device 13, Function 0-5; Offset ACH) were intended to be Read-Only. Due to this erratum, these registers behave incorrectly as Read-Write.

Implication: Writes to the RTID_POOL_CONFIG CSRs may lead to unexpected results.

Workaround: None identified. Software should write to RTID_POOL_CONFIG_SHADOW CSRs (Device 12; Function 0-5; Offset B0H and Device 13; Function 0-5; Offset B0H) rather than RTID_POOL_CONFIG CSRs.

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

CG119 Catastrophic Trip Triggered at Lower Than Expected Temperatures

Problem: Catastrophic Trip is intended to provide protection when the temperature of the processor exceeds a critical threshold by immediately shutting down the processor. Due to this erratum, the Catastrophic Trip may be triggered well below the critical threshold.

Implication: When this erratum occurs, the processor improperly issues a catastrophic shutdown causing a system failure.

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

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

CG120 PCIe* Hot Plug Slot Status Register May Not Indicate Command Completed

Problem: The PCIe Base Specification requires a write to the Slot Control register (Offset A8H) to generate a hot plug command when the downstream port is hot plug capable. Due to this erratum, a hot plug command is generated only when one or more of the Slot Control register bits [11:6] are changed.

Implication: Writes to the Slot Control register that leave bits [11:6] unchanged will not generate a hot plug command and will therefore not generate a command completed event. Software that expects a command completed event may not behave as expected.

Workaround: It is possible for software to implement a one-second timeout in lieu of receiving a command completed event.

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

CG121 Misaligned PCIe* Configuration Register Writes May Not be Dropped

Problem: Misaligned PCIe configuration register writes should be dropped and an error should be logged by setting UBOXErrSts.CFGWrAddrMisAligned (Device 11; Function 0, Offset 0x64; bit 5) to 1. Due to this erratum, a misaligned register write may not be dropped.

Implication: Misaligned PCIe configuration register writes may result in unpredictable system behavior. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

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

CG122 PCIe* Correctable Error Status Register May Not Log Receiver Error at 8.0 GT/s

Problem: Due to this erratum, correctable PCIe receiver errors may not be logged in the DPE field (bit 15) of the PCISTS CSR (Bus:0; Device 1,2,3; Function 0-1,0-3,0-3; Offset 6H) when operating at 8.0 GT/s.

Implication: Correctable receiver errors during 8.0 GT/s operation may not be visible to the OS or driver software.

Workaround: None identified.

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

CG123 Configuring PCIe* Port 3a as an NTB Disables EOI Forwarding to Port 2a

Problem: Configuring PCIe Port 3a as an NTB (non-transparent bridge) requires disabling EOI (End Of Interrupt) broadcast forwarding to this port by setting bit 26 of MISCCTRLSTS CSR (Bus 0; Device 3; Function 0; Offset 188H) to 0. Due to this erratum, disabling EOI broadcast forwarding to Port 3a improperly disables EOI broadcast forwarding to Port 2a.

Implication: Some platform configurations will not behave as expected.

Workaround: If Port 3a is configured as an NTB then devices requiring EOI messages (those using Message Signaled Interrupts and those with their own IO APIC) must not be connected to port 2a.

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

CG124 PCIe* LBMS Bit Incorrectly Set

Problem: If a PCIe Link autonomously changes width or speed for reasons other than to attempt to correct unreliable Link operation, the Port should set LABS bit (Link Autonomous Bandwidth Status) (Bus 0; Device 0; Function 0 and Device 1; Function 0-1 and Device 2-3; Function 0-3; Offset 0x1A2; bit 15). Due to this erratum, the processor will not set this bit and will incorrectly set LBMS bit (Link Bandwidth Management Status) (Bus 0; Device 0; Function 0 and Device 1; Function 0-1 and Device 2-3; Function 0-3; Offset 0x1A2; bit14) instead.

Implication: Software that uses the LBMS bit or LABS bit may behave incorrectly.

Workaround: None identified.

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

CG125 PCIe* DLW is Not Supported When Operating at 8GT/s

Problem: Problem: DLW (Dynamic Link Width) is an optional PCIe feature enabling a PCIe device to dynamically change the link width to manage power and bandwidth. When a PCIe device is operating at 8 GT/s, an attempt to change the link width using DLW may result in a Surprise Link Down error.

Implication: Due to this erratum, the processor may experience Surprise Link Down errors.

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

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

CG126 Spurious Patrol Scrub Errors May Be Reported During Exit From Deep Package C-States

Problem: When exiting from Package C3 or deeper, spurious Memory Scrubbing Errors may be reported with IA32_MC(13-16)_STATUS.MCACOD with a value of 0000_0000_1100_CCCCb (where CCCC is the channel number).

Implication: The patrol scrub errors reported when this erratum occurs are uncorrectable and may result in a system reset.

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

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

CG127 Local PCIe* P2P Traffic on x4 Ports May Cause a System Hang

Problem: Under certain conditions, P2P (Peer-to-Peer) traffic between x4 PCIe ports on the same processor (i.e., local) may cause a system hang.

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

Workaround: None identified.

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

CG128 NTB Operating In NTB/RP Mode May Complete Transactions With Incorrect ReqID

Problem: When the NTB (Non-Transparent Bridge) is operating in NTB/RP (NTB Root Port mode) it is possible for transactions to be completed with the incorrect ReqID (Requester ID). This erratum occurs when an outbound transaction is aborted before a completion for inbound transaction is returned.

Implication: Due to this erratum, a completion timeout and an unexpected completion may be seen by the processor connected to the NTB/RP. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

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

CG129 Warm Reset May Cause PCIe Hot-Plug Sequencing Failure

Problem: The Integrated I/O unit uses the VPP (Virtual Pin Port) to communicate with power controllers, switches, and LEDs associated with PCIe Hot-Plug sequencing. Due to this erratum, a warm reset occurring when a VPP transaction is in progress may result in an extended VPP stall, termination of the in flight VPP transaction, or a transient power down of slots subject to VPP power control.

Implication: During or shortly after a warm reset, when this erratum occurs, PCIe Hot-Plug sequencing may experience transient or persistent failures or slots may experience unexpected transient power down events. In certain instances, a cold reset may be needed to fully restore operation.

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

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

CG130 Performance Monitoring IA32_PERF_GLOBAL_STATUS.CondChgd Bit Not Cleared by Reset

Problem: The IA32_PERF_GLOBAL_STATUS MSR (38EH) should be cleared by reset. Due to this erratum, CondChgd (bit 63) of the IA32_PERF_GLOBAL_STATUS MSR may not be cleared.

Implication: When this erratum occurs, performance monitoring software may behave unexpectedly.

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

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

CG131 PCIe* TLP Translation Request Errors Are Not Properly Logged For Invalid Memory Writes

Problem: A PCIe Memory Write TLP (Transaction Layer Packet) with an AT field value of 01b (address translation request) does not set the UR (Unsupported Request) bit (UNCERRSTS CSR, Bus 0; Device 0; Function 0; Offset 0x14C; Bit 20) as required by the PCIe Base Specification.

Implication: System or software monitoring error status bits may not be notified of an unsupported request. When this erratum occurs, the processor sets the 'advisory_non_fatal_error_status' bit (CORERRSTS CSR, Bus 0; Device 0; Function 0; Offset 0x158; Bit 13) and drops the failing transaction.

Workaround: None identified.

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

CG132 Threshold-Based Status Indicator Not Updated After a UC or UCR Occurs

Problem: In Machine Check Status MSRs for the cache, bits 54:53 are defined as the Threshold-based Error Status Indicator for cache errors. If a UC (uncorrectable error) or UCR (uncorrectable recoverable error) occurs in the cache, additional correctable errors that occur in the cache that exceed the threshold value will not cause the Threshold-based Error Status Indicator to change from Green to Yellow status once UC or UCR is indicated.

Implication: After OS recovery of UCR errors without PCC (processor context corrupt), the Threshold-based Error Status data may not have been updated for CE (correctable errors) during the window of time from the UCR error until the software clears the IA32_MCi_STATUS MSRs.

Workaround: None identified.

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

CG133 PCIe* Slave Loopback May Transmit Incorrect Sync Headers

Problem: The PCIe Base Specification requires that, in the Loopback.Active state, a loopback slave re-transmits the received bit stream bit-for-bit on the corresponding Tx. If the link is directed to enter loopback slave mode at 8 GT/s via TS1 ordered sets with both the Loopback and Compliance Receive bits set, the processor may place sync headers in incorrect locations in the loopback bit stream.

Implication: In PCIe CEM (Card Electromechanical specification) Rx compliance testing directing the link to loopback slave mode, the received data may not be correctly re-transmitted on the Tx, causing the test to fail.

Workaround: None Identified.

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

CG134 PCIe* Type 1 VDMs May be Silently Dropped

Problem: Due to this erratum, a PCIe Type 1 VDMs (Vendor Defined Message) is silently dropped unless the vendor ID is the MCTP (Management Component Transport Protocol) value of 0x1AB4.

Implication: PCIe Type 1 VDMs may be unexpectedly dropped. Intel has not observed this erratum to impact the operation of any commercially available system.

Workaround: None identified.

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

CG135 Writing PCIe* Port 2A DEVCTRL May Have Side Effects When Port 2 is Bifurcated

Problem: When PCIe port 2 is bifurcated, due to this erratum, a write to DEVCTRL (Bus 0x0, Device 0x2, Function 0x0, Offset 0x98) for PCIe port 2A may affect the max_payload_size field (bits[7:5]) in DEVCTRL for port 2B, port 2C, and/or port 2D (Bus 0x0, Device 0x2, Function 0x1-3, Offset 0x98). This erratum applies only when an affected max_payload_size field value for port 2B, port 2C, and/or port 2D is different than the max_payload_size field value written to port 2A's DEVCTRL.

Implication: The max_payload_size values for port 2B, 2C, and 2D may not match the end point device and may result in unpredictable system behavior.

Workaround: None identified.

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

CG136 Performance Monitor Instructions Retired Event May Not Count Consistently

Problem: The Performance Monitor Instructions Retired event (Event C0H; Umask 00H) and the instruction retired fixed counter IA32_FIXED_CTR0 MSR (309H) are used to count the number of instructions retired. Due to this erratum, certain internal conditions may cause the counter(s) to increment when no instruction has retired or to intermittently not increment when instructions have retired.

Implication: A performance counter counting instructions retired may over count or under count. The count may not be consistent between multiple executions of the same code.

Workaround: None identified.

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

CG137 Accessing SB01BASE MMIO Space in The Presence of Bi-directional NTB Traffic May Result in a System Hang

Problem: While transactions are originating from both sides of the NTB, accessing SB01BASE MMIO space may cause the system to hang. For example, the NTB convention of using SB01BASE MMIO doorbell or scratchpad registers to convey interrupt messages across the NTB may result in a system hang. This erratum does not apply if transactions originate from only one side of the NTB.

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

Workaround: Do not reference SB01BASE MMIO space when transactions may originate from both sides of the NTB. For the example given above, this may require substituting a polling method in place of SB01BASE doorbells or scratchpad register accesses.

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

CG138 Patrol Scrubbing Doesn't Skip Ranks Disabled After DDR Training

Problem: If a rank is detected as failed after completing DDR training then BIOS will mark it as disabled. Disabled ranks are omitted from the OS memory map. Due to this erratum, a rank disabled after DDR training completes is not skipped by the Patrol Scrubber. Patrol Scrubbing of the disabled ranks may result in superfluous correctable and uncorrectable memory error reports.

Implication: Disabling ranks after DDR training may result in the over-reporting of memory errors.

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

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

CG139 Multiple Intel® VT-d Translation Requests From a Single Agent May Result in a Hang

Problem: A single agent issuing multiple requests to the same memory page requiring Intel® VT-d address translation may result in a hang.

Implication: When this erratum occurs, the processor will hang. Intel has not observed this erratum with any commercially available system.

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

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

CG140 The System May Shut Down Unexpectedly During a Warm Reset

Problem: Certain complex internal timing conditions present when a warm reset is requested can prevent the orderly completion of in-flight transactions. It is possible under these conditions that the warm reset will fail and trigger a full system shutdown.

Implication: When this erratum occurs, the system will shut down and all machine check error logs will be lost.

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

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

CG141 Intel® VT-d Memory Check Error on an Intel® QuickData Technology Channel May Cause All Other Channels to Master Abort

Problem: An Intel QuickData DMA access to Intel® VT-d protected memory that results in a protected memory check error may cause master abort completions on all other Intel QuickData DMA channels.

Implication: Due to this erratum, an error during Intel QuickData DMA access to an Intel® VT-d protected memory address may cause a master abort on other Intel QuickData DMA channels.

Workaround: None identified.

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

CG142 Writes To Some Control Register Bits Ignore Byte Enable

Problem: Due to this erratum, partial writes to some registers write the full register. The affected registers are:

SADDBGMM2_CFG (Device 12; Function 0-7; Offset 0xA8 and Device 13; Function 0-6; Offset 0xA8) and

LLCERRINJ_CFG (Device 12; Function 0-7; Offset 0xFC and Device 13; Function 0-6; Offset 0xFC)

Implication: Partial writes of the registers listed above may result in changes to register bytes that were intended to be unmodified.

Workaround: None identified. Use aligned, full-width DWORD (32-bit) read-modify-write sequencing to change a portion or portions of the registers listed.

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

CG143 Invalid Intel® QuickData Technology XOR Descriptor Source Addressing May Lead to Unpredictable System Behavior

Problem: Intel QuickData Technology (i.e. Crystal Beach DMA v3.2) does not correctly halt and report aborts on illegal source addresses placed in a CBDMA descriptor regardless of type (Legacy or PQ). This abort condition may cause unpredictable system behavior.

Implication: This erratum may lead to unpredictable system behavior.

Workaround: Ensure XOR DMA descriptor source addresses targets valid DRAM memory locations.

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

CG144 DDRIO SVID And SDID CSR Writes Do Not Behave as Expected

Problem: SVID and SDID (Bus 1; Device 17,31; Functions 0-7; Offsets 0x2C and 0x2E) registers implement write-once registers within their configuration space. Due to this erratum, write accesses to either of these registers individually, with double-word sized accesses, may prevent any further update to the other register.

Implication: Writes to SVID and SDID registers may not work as intended.

Workaround: Writes to the SVID and SDID registers must be made with byte or word writes.

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

CG145 Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation

Problem: This erratum may cause a machine-check error (IA32_MCi_STATUS.MCACOD=0150H) on the fetch of an instruction that crosses a 4-KByte address boundary. It applies only if (1) the 4-KByte linear region on which the instruction begins is originally translated using a 4-KByte page with the WB memory type; (2) the paging structures are later modified so that linear region is translated using a large page (2-MByte, 4-MByte, or 1-GByte) with the UC memory type; and (3) the instruction fetch occurs after the paging-structure modification but before software invalidates any TLB entries for the linear region.

Implication: Due to this erratum an unexpected machine check with error code 0150H may occur, possibly resulting in a shutdown. Intel has not observed this erratum with any commercially available software.

Workaround: Software should not write to a paging-structure entry in a way that would change, for any linear address, both the page size and the memory type. It can instead use the following algorithm: first clear the P flag in the relevant paging-structure entry (e.g., PDE); then invalidate any translations for the affected linear addresses; and then modify the relevant paging-structure entry to set the P flag and establish the new page size and memory type.

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

CG146 High Frequency Noise on DDR SMBus Signals May Prevent Proper Detection of Memory

Problem: During the processor power up sequence, high frequency noise may occur on the DDR SMBus SDA and SCL signals interfering with correct information transfer.

Implication: When this erratum occurs, high frequency noise may cause certain voltage translator components to latch up and, as a result, the system may not be able to detect memory.

Workaround: None identified.

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

§

Specification Changes

There are no specification changes in this specification update revision.

Specification Clarifications

There are no specification clarifications in this specification update revision.

§

Documentation Changes

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

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

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

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

http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html

CG1 SDM, Volume 3B: On-Demand Clock Modulation Feature Clarification

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

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

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

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

Table 14.2 CPUID Signatures for Legacy Processors That Resolve to Higher Performance Setting of Conflicting Duty Cycle Requests

DisplayFamily_ DisplayModel	DisplayFamily_ DisplayModel	DisplayFamily_ DisplayModel	DisplayFamily_ DisplayModel
0F_xx	06_1C	06_1A	06_1E
06_1F	06_25	06_26	06_27
06_2C	06_2E	06_2F	06_35
06_36			

§ §