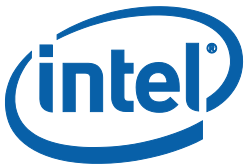

8th Generation Intel® Processor Family for S-Processor Platforms

Specification Update

October 2017

Revision 001

Document Number: 336466-001

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit www.intel.com/design/literature.htm.

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

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

Copyright © 2017, Intel Corporation. All Rights Reserved.

Contents

Revision History

Revision Number	Description	Date
001	• Initial Release	October 2017

§

Preface

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

Information types defined in the Nomenclature section of this document are consolidated into this update document and are no longer published in other documents.

This document may also contain information that has not been previously published.

Affected Documents

Document Title	Document Number
8th Generation Intel® Processor Family for S-Processor Platforms Datasheet	336464 (Vol 1) 336465 (Vol 2)

Related Documents

Document Title	Document Number/Location
AP-485, Intel® Processor Identification and the CPUID Instruction	http://www.intel.com/design/processor/applnots/241618.htm
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 Intel® 64 and IA-32 Intel Architecture Optimization Reference Manual	http://www.intel.com/products/processor/manuals/index.htm
Intel® 64 and IA-32 Architectures Software Developer's Manual Documentation Changes	http://www.intel.com/content/www/us/en/processors/architectures-software-developer-manuals.html
ACPI Specifications	www.acpi.info

Nomenclature

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

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

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

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

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

§

Identification Information

Component Identification via Programming Interface

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

Table 1. S -Processor Lines Component Identification

Reserved	Extended Family	Extended Model	Reserved	Processor Type	Family Code	Model Number	Stepping ID
31:28	27:20	19:16	15:14	13:12	11:8	7:4	3:0
	0000000b	1001b		00b	0110b	1110b	xxxxb

Notes:

1. The Extended Family, Bits [27:20] are used in conjunction with the Family Code, specified in Bits[11:8], to indicate whether the processor belongs to the Celeron® ,Pentium®, 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. See Table 1 for the processor stepping ID number in the CPUID information.
6. When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

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

Component Marking Information

Figure 1. S-Processor Line LGA Top-Side Markings

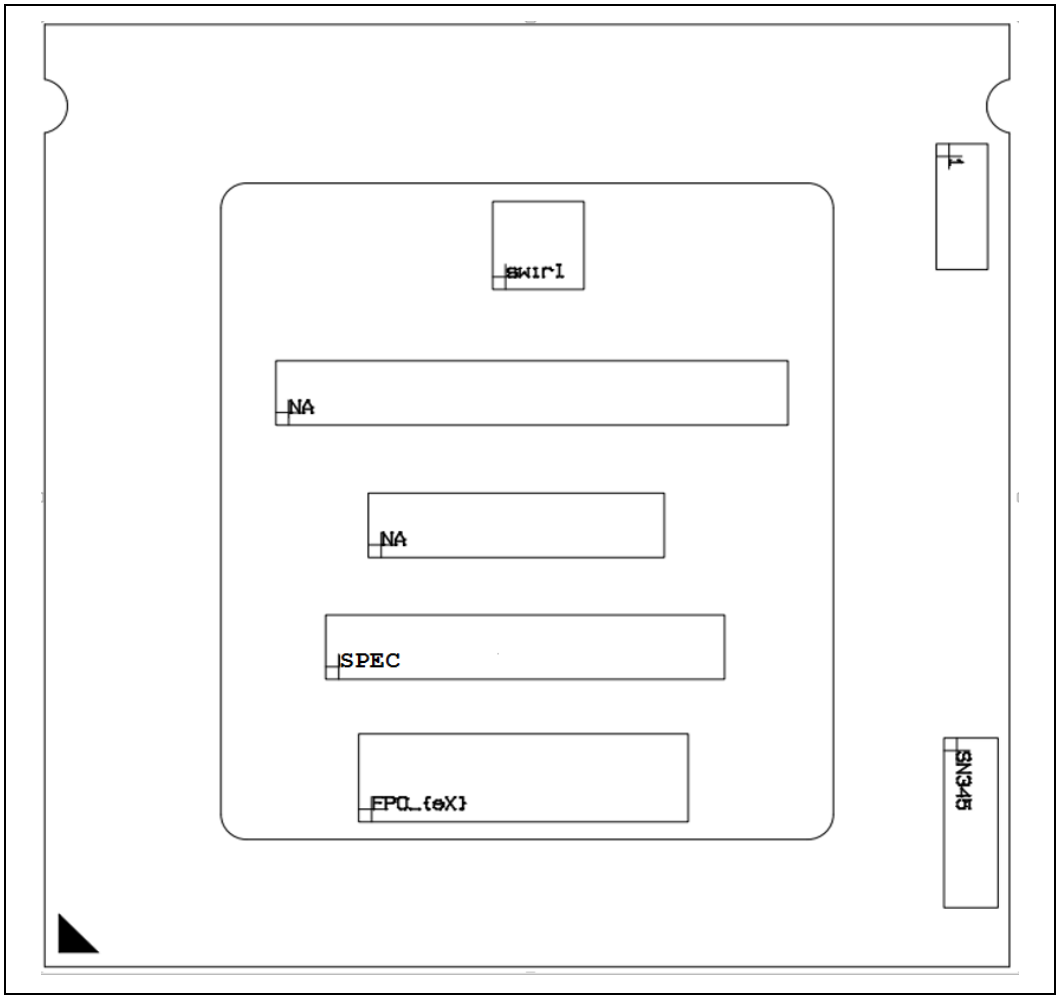

Pin Count: 1151 Package Size: 37.5 mm x 37.5 mm

Production (SSPEC):
Intel logo
BRAND
PROC#
SPEC SPEED
{FPO} {eX}
Note: "1" is used to extract the unit visual ID (2D ID).

Table 2. S-Processor Line

S-Spec #	Processor Number	Step-ping	Cache Size	Functional Core	Processor Graphics Cores	Processor Graphics Freq.	Processor Graphics Maximum Dynamic Freq.	DDR3L Mem. (MT/s)	LPDDR3 Mem. (MT/s)	DDR4 Mem. (MT/s)	Core Freq.	Turbo 1 Core Freq. Rate	Thermal Design Power	Slot / Socket Type
R3N5	I3-8100	B-0	6MB	4	2	0.35GHz	1.1GHz	N/A	N/A	2400	3.6GHz	3.6GHz	65W	LGA1151
R3N4	I3-8350K	B-0	8MB	4	2	0.35GHz	1.15GHz	N/A	N/A	2400	4GHz	4GHz	95W	LGA1151
R3QT	I5-8400	U-0	9MB	6	2	0.35GHz	1.05GHz	N/A	N/A	2666	2.8GHz	4GHz	65W	LGA1151
R3QU	I5-8600K	U-0	9MB	6	2	0.35GHz	1.15GHz	N/A	N/A	2666	3.6GHz	4.3GHz	95W	LGA1151
R3QS	I7-8700	U-0	12MB	6	2	0.35GHz	1.2GHz	N/A	N/A	2666	3.2GHz	4.6GHz	65W	LGA1151
R3QR	I7-8700K	U-0	12MB	6	2	0.35GHz	1.2GHz	N/A	N/A	2666	3.7GHz	4.7GHz	95W	LGA1151

§

Summary Tables of Changes

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

Codes Used in Summary Table

Stepping

X:	Erratum, Specification Change or Clarification that applies to this stepping.
(No mark) or (Blank Box):	This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Status

Doc:	Document change or update that will be implemented.
Planned 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.

Errata Summary Table

ID	Processor Line / Stepping		Status	Title
	S			
	B0	U0		
001	X	X	No Fix	Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures
002	X	X	No Fix	Instruction Fetch May Cause Machine Check if Page Size and Memory Type Was Changed Without Invalidation
003	X	X	No Fix	Execution of VAESIMC or VAESKEYGENASSIST With An Illegal Value for VEX.vvvv May Produce a #NM Exception
004	X	X	No Fix	The Corrected Error Count Overflow Bit in IA32_ MC0_STATUS is Not Updated When The UC Bit is Set
005	X	X	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
006	X	X	No Fix	SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior
007	X	X	No Fix	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
008	X	X	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
009	X	X	No Fix	DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction
010	X	X	No Fix	Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID
011	X	X	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
012	X	X	No Fix	The SMSW Instruction May Execute Within an Enclave
013	X	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG Concurrent With an SMX SENTER/SEXIT May Result in a System Hang
014	X	X	No Fix	Intel® PT TIP.PGD May Not Have Target IP Payload
015	X	X	No Fix	Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a #UD
016	X	X	No Fix	Execution of FXSAVE or FXRSTOR With the VEX Prefix May Produce a #NM Exception
017	X	X	No Fix	WRMSR May Not Clear The Sticky Count Overflow Bit in The IA32_MCi_STATUS MSRs' Corrected Error Count Field
018	X	X	No Fix	PEBS Eventing IP Field May be Incorrect After Not-Taken Branch
019	X	X	No Fix	Debug Exceptions May Be Lost or Misreported Following WRMSR to IA32_BIOS_UPDT_TRIG
020	X	X	No Fix	Complex Interactions With Internal Graphics May Impact Processor Responsiveness
021	X	X	No Fix	Intel® Processor Trace PSB+ Packets May Contain Unexpected Packets

ID	Processor Line / Stepping		Status	Title
	S			
	B0	U0		
022	X	X	No Fix	Placing an Intel® PT ToPA in Non-WB Memory or Writing It Within a Transactional Region May Lead to System Instability
023	X	X	No Fix	VM Entry That Clears TraceEn May Generate a FUP
024	X	X	No Fix	Performance Monitor Event For Outstanding Offcore Requests And Snoop Requests May be Incorrect
025	X	X	No Fix	ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK
026	X	X	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
027	X	X	No Fix	ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero
028	X	X	No Fix	A VMX Transition Attempting to Load a Non-Existent MSR May Result in a Shutdown
029	X	X	No Fix	Transitions Out of 64-bit Mode May Lead to an Incorrect FDP And FIP
030	X	X	No Fix	Intel® PT FUP May be Dropped After OVF
031	X	X	No Fix	ENCLS[ECREATE] Causes #GP if Enclave Base Address is Not Canonical
032	X	X	No Fix	Processor Graphics IOMMU Unit May Report Spurious Faults
033	X	X	No Fix	Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State
034	X	X	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction
035	X	X	No Fix	ENCLS[EINIT] Instruction May Unexpectedly #GP
036	X	X	No Fix	Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop
037	X	X	No Fix	WRMSR to IA32_BIOS_UPDT_TRIG May be Counted as Multiple Instructions
038	X	X	No Fix	Branch Instructions May Initialize MPX Bound Registers Incorrectly
039	X	X	No Fix	Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled
040	X	X	No Fix	Processor May Run Intel® AVX Code Much Slower Than Expected
041	X	X	No Fix	Intel® PT Buffer Overflow May Result in Incorrect Packets
042	X	X	No Fix	Last Level Cache Performance Monitoring Events May Be Inaccurate
043	X	X	No Fix	#GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave

ID	Processor Line / Stepping		Status	Title
	S			
	B0	U0		
044	X	X	No Fix	Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception
045	X	X	No Fix	Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits
046	X	X	No Fix	CR3 Filtering Does Not Compare Bits [11:5] of CR3 and IA32_RTIT_CR3_MATCH in PAE Paging Mode
047	X	X	No Fix	x87 FDP Value May be Saved Incorrectly
048	X	X	No Fix	PECI Frequency Limited to 1 MHz
049	X	X	No Fix	Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults
050	X	X	No Fix	Intel® PT CYCThresh Value of 13 is Not Supported
051	X	X	No Fix	Enabling VMX-Preemption Timer Blocks HDC Operation
052	X	X	No Fix	Integrated Audio Codec May Not be Detected
053	X	X	No Fix	Display Flickering May be Observed with Specific eDP Panels
054	X	X	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
055	X	X	No Fix	MACHINE_CLEARS.MEMORY ORDERING Performance Monitoring Event May Undercount
056	X	X	No Fix	Some Counters May Not Freeze On Performance Monitoring Interrupts
057	X	X	No Fix	Instructions And Branches Retired Performance Monitoring Events May Overcount
058	X	X	No Fix	Some OFFCORE_RESPONSE Performance Monitoring Events May Overcount
059	X	X	No Fix	#GP After RSM May Push Incorrect RFLAGS Value When Intel® PT is Enabled
060	X	X	No Fix	Access to SGX EPC Page in BLOCKED State is Not 062 Reported as an SGX-Induced Page Fault
061	X	X	No Fix	MTF VM Exit on XBEGIN Instruction May Save State Incorrectly
062	X	X	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
063	X	X	No Fix	Certain Non-Canonical IA32_BNDCFGS Values Will Not Cause VM-Entry Failures
064	X	X	No Fix	PEBS EventingIP Field May Be Incorrect Under Certain Conditions
065	X	X	No Fix	HWP's Guaranteed_Performance Updated Only on Configurable TDP Changes

ID	Processor Line / Stepping		Status	Title
	S			
	B0	U0		
066	X	X	No Fix	RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
067	X	X	No Fix	Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters
068	X	X	No Fix	HWP's Maximum_Performance Value is Reset to 0xFF
069	X	X	No Fix	HWP's Guaranteed_Performance and Relevant Status/Interrupt May be Updated More Than Once Per Second
070	X	X	No Fix	Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes
071	X	X	No Fix	HWP May Generate Thermal Interrupt While Not Enabled
072	X	X	No Fix	Camera Device Does Not Issue an MSI When INTx is Enabled
073	X	X	No Fix	Attempts to Retrain a PCIe* Link May be Ignored
074	X	X	No Fix	PCIe* Port Does Not Support DLL Link Activity Reporting
075	X	X	No Fix	BNDLDX And BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access
076	X	X	No Fix	RING_PERF_LIMIT_REASONS May be Incorrect
077	X	X	No Fix	Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions
078			No Fix	EDRAM Corrected Error Events May Not be Properly Logged After a Warm Reset
079	X	X	No Fix	Some Bits in MSR_MISC_PWR_MGMT May be Updated on Writing Illegal Values to This MSR
080	X	X	No Fix	Violations of Intel® Software Guard Extensions (Intel® SGX) Access-Control Requirements Produce #GP Instead of #PF
081	X	X	No Fix	IA32_RTIT_CR3_MATCH MSR Bits[11:5] Are Treated As Reserved
082	X	X	No Fix	The Intel PT CR3 Filter is Not Re-evaluated on VM Entry
083	X	X	No Fix	Display Slowness May be Observed Under Certain Display Commands Scenario
084	X	X	No Fix	CPUID TLB Associativity Information is Inaccurate
085	X		No Fix	Using Different Vendors For 2400 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang

ID	Processor Line / Stepping		Status	Title
	S			
	B0	U0		
086	X	X	No Fix	Two DIMMs Per Channel 2133MHz DDR4 SODIMM Daisy-Chain Systems With Different Vendors May Hang
087	X	X	No Fix	Unpredictable System Behavior May Occur in DDR4 Multi-Rank System
088	X	X	No Fix	Processor May Hang on Complex Sequence of Conditions
089	X	X	No Fix	Potential Partial Trace Data Loss in Intel® Trace Hub ODLA When Storing to Memory.
090		X	No Fix	Using Different Vendors For 2666 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang

Errata

001	**Reported Memory Type May Not Be Used to Access the VMCS and Referenced Data Structures**
Problem	Bits 53:50 of the IA32_VMX_BASIC MSR report the memory type that the processor uses to access the VMCS and data structures referenced by pointers in the VMCS. Due to this erratum, a VMX access to the VMCS or referenced data structures will instead use the memory type that the MTRRs (memory-type range registers) specify for the physical address of the access.
Implication	Bits 53:50 of the IA32_VMX_BASIC MSR report that the WB (write-back) memory type will be used but the processor may use a different memory type.
Workaround	Software should ensure that the VMCS and referenced data structures are located at physical addresses that are mapped to WB memory type by the MTRRs.
Status	For the steppings affected, see the Summary Table of Changes.

002	**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 steppings affected, see the Summary Table of Changes.

003	**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 steppings affected, see the Summary Table of Changes.

004	**The Corrected Error Count Overflow Bit in IA32_ MC0_STATUS is Not Updated When The UC Bit is Set**
Problem	After a UC (uncorrected) error is logged in the IA32_MC0_STATUS MSR (401H), corrected errors will continue to be counted in the lower 14 bits (bits 51:38) of the Corrected Error Count. Due to this erratum, the sticky count overflow bit (bit 52) of the Corrected Error Count will not get updated when the UC bit (bit 61) is set to 1.
Implication	The Corrected Error Count Overflow indication will be lost if the overflow occurs after an uncorrectable error has been logged.
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

005	**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 steppings affected, see the Summary Table of Changes.

006	**SMRAM State-Save Area Above the 4GB Boundary May Cause Unpredictable System Behavior**
Problem	If BIOS uses the RSM instruction to load the SMBASE register with a value that would cause any part of the SMRAM state-save area to have an address above 4-GBytes, subsequent transitions into and out of SMM (system-management mode) might save and restore processor state from incorrect addresses.
Implication	This erratum may cause unpredictable system behavior. Intel has not observed this erratum with any commercially available system.
Workaround	Ensure that the SMRAM state-save area is located entirely below the 4GB address boundary.
Status	For the steppings affected, see the Summary Table of Changes.

007	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
Problem	x87 instructions that trigger #MF normally service interrupts before the #MF. Due to this erratum, if an instruction that triggers #MF is executing when an Enhanced Intel SpeedStep® Technology transitions, an Intel® Turbo Boost Technology transitions, or a Thermal Monitor events occurs, the #MF may be taken before pending interrupts are serviced.
Implication	Software may observe #MF being signaled before pending interrupts are serviced.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

008	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
Problem	During RTM (Restricted Transactional Memory) operation when branch tracing is enabled using BTM (Branch Trace Message) or BTS (Branch Trace Store), the incorrect EIP value (From_IP pointer) may be observed for an RTM abort.
Implication	Due to this erratum, the From_IP pointer may be the same as that of the immediately preceding taken branch.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

009	DR6 Register May Contain an Incorrect Value When a MOV to SS or POP SS Instruction is Followed by an XBEGIN Instruction
Problem	If XBEGIN is executed immediately after an execution of MOV to SS or POP SS, a transactional abort occurs and the logical processor restarts execution from the fallback instruction address. If execution of the instruction at that address causes a debug exception, bits [3:0] of the DR6 register may contain an incorrect value.
Implication	When the instruction at the fallback instruction address causes a debug exception, DR6 may report a breakpoint that was not triggered by that instruction, or it may fail to report a breakpoint that was triggered by the instruction.
Workaround	Avoid following a MOV SS or POP SS instruction immediately with an XBEGIN instruction.
Status	For the steppings affected, see the Summary Table of Changes.

010	Opcode Bytes F3 0F BC May Execute As TZCNT Even When TZCNT Not Enumerated by CPUID
Problem	If CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 1 then opcode bytes F3 0F BC should be interpreted as TZCNT otherwise they will be interpreted as REP BSF. Due to this erratum, opcode bytes F3 0F BC may execute as TZCNT even if CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 0.
Implication	Software that expects REP prefix before a BSF instruction to be ignored may not operate correctly since there are cases in which BSF and TZCNT differ with regard to the flags that are set and how the destination operand is established.
Workaround	Software should use the opcode bytes F3 0F BC only if CPUID.(EAX=07H, ECX=0):EBX.BMI1 (bit 3) is 1 and only if the functionality of TZCNT (and not BSF) is desired.

Status	For the steppings affected, see the Summary Table of Changes.

011	**#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code**
Problem	During a #GP (General Protection Exception), the processor pushes an error code on to the exception handler's stack. If the segment selector descriptor straddles the canonical boundary, the error code pushed onto the stack may be incorrect.
Implication	An incorrect error code may be pushed onto the stack. Intel has not observed this erratum with any commercially available software.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

012	**The SMSW Instruction May Execute Within an Enclave**
Problem	The SMSW instruction is illegal within an SGX (Software Guard Extensions) enclave, and an attempt to execute it within an enclave should result in a #UD (invalid-opcode exception). Due to this erratum, the instruction executes normally within an enclave and does not cause a #UD.
Implication	The SMSW instruction provides access to CR0 bits 15:0 and will provide that information inside an enclave. These bits include NE, ET, TS, EM, MP and PE.
Workaround	None identified. If SMSW execution inside an enclave is unacceptable, system software should not enable SGX.
Status	For the steppings affected, see the Summary Table of Changes.

013	**WRMSR to IA32_BIOS_UPDT_TRIG Concurrent With an SMX SENTER/SEXIT May Result in a System Hang**
Problem	Performing WRMSR to IA32_BIOS_UPDT_TRIG (MSR 79H) on a logical processor while another logical processor is executing an SMX (Safer Mode Extensions) SENTER/SEXIT operation (GETSEC[SENTER] or GETSEC[SEXIT] instruction) may cause the processor to hang.
Implication	When this erratum occurs, the system will hang. Intel has not observed this erratum with any commercially available system.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

014	**Intel® PT TIP.PGD May Not Have Target IP Payload**
Problem	When Intel PT (Intel Processor Trace) is enabled and a direct unconditional branch clears IA32_RTIT_STATUS.FilterEn (MSR 571H, bit 0), due to this erratum, the resulting TIP.PGD (Target IP Packet, Packet Generation Disable) may not have an IP payload with the target IP.
Implication	It may not be possible to tell which instruction in the flow caused the TIP.PGD using only the information in trace packets when this erratum occurs.
Workaround	The Intel PT trace decoder can compare direct unconditional branch targets in the source with the FilterEn address range(s) to determine which branch cleared FilterEn.

Status	For the steppings affected, see the Summary Table of Changes.

015	**Operand-Size Override Prefix Causes 64-bit Operand Form of MOVBE Instruction to Cause a #UD**
Problem	Execution of a 64 bit operand MOVBE instruction with an operand-size override instruction prefix (66H) may incorrectly cause an invalid-opcode exception (#UD).
Implication	A MOVBE instruction with both REX.W=1 and a 66H prefix will unexpectedly cause an #UD (invalid-opcode exception). Intel has not observed this erratum with any commercially available software.
Workaround	Do not use a 66H instruction prefix with a 64-bit operand MOVBE instruction.
Status	For the steppings affected, see the Summary Table of Changes.

016	**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 steppings affected, see the Summary Table of Changes.

017	**WRMSR May Not Clear The Sticky Count Overflow Bit in The IA32_MCi_STATUS MSRs' Corrected Error Count Field**
Problem	The sticky count overflow bit is the most significant bit (bit 52) of the Corrected Error Count Field (bits[52:38]) in IA32_MCi_STATUS MSRs. Once set, the sticky count overflow bit may not be cleared by a WRMSR instruction. When this occurs, that bit can only be cleared by power-on reset.
Implication	Software that uses the Corrected Error Count field and expects to be able to clear the sticky count overflow bit may misinterpret the number of corrected errors when the sticky count overflow bit is set. This erratum does not affect threshold-based CMCI (Corrected Machine Check Error Interrupt) signaling.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

018	**PEBS Eventing IP Field May be Incorrect After Not-Taken Branch**
Problem	When a PEBS (Precise-Event-Based-Sampling) record is logged immediately after a not-taken conditional branch (Jcc instruction), the Eventing IP field should contain the address of the first byte of the Jcc instruction. Due to this erratum, it may instead contain the address of the instruction preceding the Jcc instruction.
Implication	Performance monitoring software using PEBS may incorrectly attribute PEBS events that occur on a Jcc to the preceding instruction.
Workaround	None identified.

Status	For the steppings affected, see the Summary Table of Changes.

019	**Debug Exceptions May Be Lost or Misreported Following WRMSR to IA32_BIOS_UPDT_TRIG**
Problem	If the WRMSR instruction writes to the IA32_BIOS_UPDT_TRIG MSR (79H) immediately after an execution of MOV SS or POP SS that generated a debug exception, the processor may fail to deliver the debug exception or, if it does, the DR6 register contents may not correctly reflect the causes of the debug exception.
Implication	Debugging software may fail to operate properly if a debug exception is lost or does not report complete information.
Workaround	Software should avoid using WRMSR instruction immediately after executing MOV SS or POP SS
Status	For the steppings affected, see the Summary Table of Changes.

020	**Complex Interactions With Internal Graphics May Impact Processor Responsiveness**
Problem	Under complex conditions associated with the use of internal graphics, the processor may exceed the MAX_LAT CSR values (PCI configuration space, offset 03FH, bits[7:0]).
Implication	When this erratum occurs, the processor responsiveness is affected. Intel has not observed this erratum with any commercially available software.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

021	**Intel® Processor Trace PSB+ Packets May Contain Unexpected Packets**
Problem	Some Intel Processor Trace packets should be issued only between TIP.PGE (Target IP Packet.Packet Generation Enable) and TIP.PGD (Target IP Packet.Packet Generation Disable) packets. Due to this erratum, when a TIP.PGE packet is generated it may be preceded by a PSB+ (Packet Stream Boundary) that incorrectly includes FUP (Flow Update Packet) and MODE.Exec packets.
Implication	Due to this erratum, FUP and MODE.Exec may be generated unexpectedly.
Workaround	Decoders should ignore FUP and MODE.Exec packets that are not between TIP.PGE and TIP.PGD packets.
Status	For the steppings affected, see the Summary Table of Changes.

022	**Placing an Intel® PT ToPA in Non-WB Memory or Writing It Within a Transactional Region May Lead to System Instability**
Problem	If an Intel PT (Intel® Processor Trace) ToPA (Table of Physical Addresses) is not placed in WB (writeback) memory or is written by software executing within an Intel® TSX (Intel® Transactional Synchronization Extension) transactional region, the system may become unstable.
Implication	Unusual treatment of the ToPA may lead to system instability.
Workaround	None identified. Intel PT ToPA should reside in WB memory and should not be written within a Transactional Region.

Status	For the steppings affected, see the Summary Table of Changes.

023	**VM Entry That Clears TraceEn May Generate a FUP**
Problem	If VM entry clears Intel® PT (Intel Processor Trace) IA32_RTIT_CTL.TraceEn (MSR 570H, bit 0) while PacketEn is 1 then a FUP (Flow Update Packet) will precede the TIP.PGD (Target IP Packet, Packet Generation Disable). VM entry can clear TraceEn if the VM-entry MSR-load area includes an entry for the IA32_RTIT_CTL MSR.
Implication	When this erratum occurs, an unexpected FUP may be generated that creates the appearance of an asynchronous event taking place immediately before or during the VM entry.
Workaround	The Intel PT trace decoder may opt to ignore any FUP whose IP matches that of a VM entry instruction.
Status	For the steppings affected, see the Summary Table of Changes.

024	**Performance Monitor Event For Outstanding Offcore Requests And Snoop Requests May be Incorrect**
Problem	The performance monitor event OFFCORE_REQUESTS_OUTSTANDING (Event 60H, any Umask Value) should count the number of offcore outstanding transactions each cycle. Due to this erratum, the counts may be higher or lower than expected.
Implication	The performance monitor event OFFCORE_REQUESTS_OUTSTANDING may reflect an incorrect count.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

025	**ENCLU[EGETKEY] Ignores KEYREQUEST.MISCMASK**
Problem	The Intel® SGX (Software Guard Extensions) ENCLU[EGETKEY] instruction ignores the MISCMASK field in KEYREQUEST structure when computing a provisioning key, a provisioning seal key, or a seal key.
Implication	ENCLU[EGETKEY] will return the same key in response to two requests that differ only in the value of KEYREQUEST.MISCMASK. Intel has not observed this erratum with any commercially available software.
Workaround	When executing the ENCLU[EGETKEY] instruction, software should ensure the bits set in KEYREQUEST.MISCMASK are a subset of the bits set in the current SECS's MISCSELECT field.
Status	For the steppings affected, see the Summary Table of Changes.

026	**POPCNT Instruction May Take Longer to Execute Than Expected**
Problem	POPCNT instruction execution with a 32 or 64 bit operand may be delayed until previous non-dependent instructions have executed.
Implication	Software using the POPCNT instruction may experience lower performance than expected.
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

027	ENCLU[EREPORT] May Cause a #GP When TARGETINFO.MISCSELECT is Non-Zero
Problem	The Intel® SGX (Software Guard extensions) ENCLU[EREPORT] instruction may cause a #GP (general protection fault) if any bit is set in TARGETINFO structure's MISCSELECT field.
Implication	This erratum may cause unexpected general-protection exceptions inside enclaves.
Workaround	When executing the ENCLU[EREPORT] instruction, software should ensure the bits set in TARGETINFO.MISCSELECT are a subset of the bits set in the current SECS's MISCSELECT field.
Status	For the steppings affected, see the Summary Table of Changes.

028	A VMX Transition Attempting to Load a Non-Existent MSR May Result in a Shutdown
Problem	A VMX transition may result in a shutdown (without generating a machine-check event) if a non-existent MSR is included in the associated MSR-load area. When such a shutdown occurs, a machine check error will be logged with IA32_MCi_STATUS.MCACOD (bits [15:0]) of 406H, but the processor does not issue the special shutdown cycle. A hardware reset must be used to restart the processor.
Implication	Due to this erratum, the hypervisor may experience an unexpected shutdown.
Workaround	Software should not configure VMX transitions to load non-existent MSRs.
Status	For the steppings affected, see the Summary Table of Changes.

029	Transitions Out of 64-bit Mode May Lead to an Incorrect FDP And FIP
Problem	A transition from 64-bit mode to compatibility or legacy modes may result in cause a subsequent x87 FPU state save to zeroing bits [63:32] of the FDP (x87 FPU Data Pointer Offset) and the FIP (x87 FPU Instruction Pointer Offset).
Implication	Leaving 64-bit mode may result in incorrect FDP and FIP values when x87 FPU state is saved.
Workaround	None identified. 64-bit software should save x87 FPU state before leaving 64-bit mode if it needs to access the FDP and/or FIP values.
Status	For the steppings affected, see the Summary Table of Changes.

030	Intel® PT FUP May be Dropped After OVF
Problem	Some Intel PT (Intel Processor Trace) OVF (Overflow) packets may not be followed by a FUP (Flow Update Packet) or TIP.PGE (Target IP Packet, Packet Generation Enable).
Implication	When this erratum occurs, an unexpected packet sequence is generated.
Workaround	When it encounters an OVF without a following FUP or TIP.PGE, the Intel PT trace decoder should scan for the next TIP, TIP.PGE, or PSB+ to resume operation.
Status	For the steppings affected, see the Summary Table of Changes.

031	ENCLS[ECREATE] Causes #GP if Enclave Base Address is Not Canonical
Problem	The ENCLS[ECREATE] instruction uses an SECS (SGX enclave control structure) referenced by the SRCPAGE pointer in the PAGEINFO structure, which is referenced by the RBX register. Due to this erratum, the instruction causes a #GP (general-protection fault) if the SECS attributes indicate that the enclave should operate in 64-bit mode and the enclave base linear address in the SECS is not canonical.
Implication	System software will incur a general-protection fault if it mistakenly programs the SECS with a non-canonical address. Intel has not observed this erratum with any commercially available software.
Workaround	System software should always specify a canonical address as the base address of the 64-bit mode enclave.
Status	For the steppings affected, see the Summary Table of Changes.

032	Processor Graphics IOMMU Unit May Report Spurious Faults
Problem	The IOMMU unit for Processor Graphics pre-fetches context (or extended-context) entries to improve performance. Due to the erratum, the IOMMU unit may report spurious DMA remapping faults if prefetching encounters a context (or extended-context) entry which is not marked present.
Implication	Software may observe spurious DMA remapping faults when the present bit for the context (or extended-context) entry corresponding to the Processor Graphics device (Bus: 0; Device: 2; Function: 0) is cleared. These faults may be reported when the Processor Graphics device is quiescent.
Workaround	None identified. Instead of marking a context not present, software should mark the context (or extended-context) entry present while using the page table to indicate all the memory pages referenced by the context entry is not present.
Status	For the steppings affected, see the Summary Table of Changes.

033	Processor DDR VREF Signals May Briefly Exceed JEDEC Spec When Entering S3 State
Problem	Voltage glitch of up to 200mV on the VREF signal lasting for about 1mS may be observed when entering System S3 state. This violates the JEDEC DDR specifications.
Implication	Intel has not observed this erratum to impact the operation of any commercially available system.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

034	**DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction**
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 MMX instruction that uses a memory addressing mode with an index or a store instruction.
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 (i.e., following them only with an instruction that writes (E/R)SP).
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

035	**ENCLS[EINIT] Instruction May Unexpectedly #GP**
Problem	When using Intel® SGX (Software Guard Extensions), the ENCLS[EINIT] instruction will incorrectly cause a #GP (general protection fault) if the MISCSELECT field of the SIGSTRUCT structure is not zero.
Implication	This erratum may cause an unexpected #GP, but only if software has set bits in the MISCSELECT field in SIGSTRUCT structure that do not correspond to extended features that can be written to the MISC region of the SSA (State Save Area). Intel has not observed this erratum with any commercially available software.
Workaround	When executing the ENCLS[EINIT] instruction, software should only set bits in the MISCSELECT field in the SIGSTRUCT structure that are enumerated as 1 by CPUID.(EAX=12H,ECX=0):EBX (the bit vector of extended features that can be written to the MISC region of the SSA).
Status	For the steppings affected, see the Summary Table of Changes.

036	**Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop**
Problem	If an Intel PT (Intel® Processor Trace) internal buffer overflow occurs immediately before software executes a taken branch or event that enters an Intel PT TraceStop region, the OVF (Overflow) packet may be lost.
Implication	The trace decoder will not see the OVF packet, nor any subsequent packets (e.g., TraceStop) that were lost due to overflow.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

037	WRMSR to IA32_BIOS_UPDT_TRIG May be Counted as Multiple Instructions
Problem	When software loads a microcode update by writing to MSR IA32_BIOS_UPDT_TRIG (79H) on multiple logical processors in parallel, a logical processor may, due to this erratum, count the WRMSR instruction as multiple instruction-retired events.
Implication	Performance monitoring with the instruction-retired event may over count by up to four extra events per instance of WRMSR which targets the IA32_BIOS_UPDT_TRIG register.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

038	Branch Instructions May Initialize MPX Bound Registers Incorrectly
Problem	Depending on the current Intel® MPX (Memory Protection Extensions) configuration, execution of certain branch instructions (near CALL, near RET, near JMP, and Jcc instructions) without a BND prefix (F2H) initialize the MPX bound registers. Due to this erratum, execution of such a branch instruction on a user-mode page may not use the MPX configuration register appropriate to the current privilege level (BNDCFGU for CPL 3 or BNDCFGS otherwise) for determining whether to initialize the bound registers; it may thus initialize the bound registers when it should not, or fail to initialize them when it should.
Implication	After a branch instruction on a user-mode page has executed, a #BR (bound-range) exception may occur when it should not have or a #BR may not occur when one should have.
Workaround	If supervisor software is not expected to execute instructions on user-mode pages, software can avoid this erratum by setting CR4.SMEP[bit 20] to enable supervisor-mode execution prevention (SMEP). If SMEP is not available or if supervisor software is expected to execute instructions on user-mode pages, no workaround is identified.
Status	For the steppings affected, see the Summary Table of Changes.

039	Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled
Problem	If Intel PT (Intel Processor Trace) is enabled, WRMSR will not cause a general-protection exception (#GP) on an attempt to write a non-canonical value to any of the following MSRs: • MSR_LASTBRANCH_{0 - 31}_FROM_IP (680H – 69FH) • MSR_LASTBRANCH__{0 - 31}_TO_IP (6C0H – 6DFH) • MSR_LASTBRANCH_FROM_IP (1DBH) • MSR_LASTBRANCH_TO_IP (1DCH) • MSR_LASTINT_FROM_IP (1DDH) • MSR_LASTINT_TO_IP (1DEH)Instead the same behavior will occur as if a canonical value had been written. Specifically, the WRMSR will be dropped and the MSR value will not be changed.
Implication	Due to this erratum, an expected #GP may not be signaled.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

040	Processor May Run Intel® AVX Code Much Slower Than Expected
Problem	After a C6 state exit, the execution rate of AVX instructions may be reduced.
Implication	Applications using AVX instructions may run slower than expected.
Workaround	It is possible for the BIOS to contain a workaround
Status	For the steppings affected, see the Summary Table of Changes.

041	Intel® PT Buffer Overflow May Result in Incorrect Packets
Problem	Under complex micro-architectural conditions, an Intel PT (Processor Trace) OVF (Overflow) packet may be issued after the first byte of a multi-byte CYC (Cycle Count) packet, instead of any remaining bytes of the CYC.
Implication	When this erratum occurs, the splicing of the CYC and OVF packets may prevent the Intel PT decoder from recognizing the overflow. The Intel PT decoder may then encounter subsequent packets that are not consistent with expected behavior.
Workaround	None Identified. The decoder may be able to recognize that this erratum has occurred when a two-byte CYC packet is followed by a single byte CYC, where the latter 2 bytes are 0xf302, and where the CYC packets are followed by a FUP (Flow Update Packet) and a PSB+ (Packet Stream Boundary+). It should then treat the two CYC packets as indicating an overflow.
Status	For the steppings affected, see the Summary Table of Changes.

042	Last Level Cache Performance Monitoring Events May be Inaccurate
Problem	The performance monitoring events LONGEST_LAT_CACHE.REFERENCE (Event 2EH; Umask 4FH) and LONGEST_LAT_CACHE.MISS (Event 2EH; Umask 41H) count requests that reference or miss in the last level cache. However, due to this erratum, the count may be incorrect.
Implication	LONGEST_LAT_CACHE events may be incorrect.
Workaround	None identified. Software may use the following OFFCORE_REQUESTS model-specific sub events that provide related performance monitoring data: DEMAND_DATA_RD, DEMAND_CODE_RD, DEMAND_RFO, ALL_DATA_RD, L3_MISS_DEMAND_DATA_RD, ALL_REQUESTS.
Status	For the steppings affected, see the Summary Table of Changes.

043	#GP Occurs Rather Than #DB on Code Page Split Inside an Intel® SGX Enclave
Problem	When executing within an Intel® SGX (Software Guard Extensions) enclave, a #GP (general-protection exception) may be delivered instead of a #DB (debug exception) when an instruction breakpoint is detected. This occurs when the instruction to be executed spans two pages, the second of which has an entry in the EPCM (enclave page cache map) that is not valid.
Implication	Debugging software may not be invoked when an instruction breakpoint is detected.
Workaround	Software should ensure that all pages containing enclave instructions have valid EPCM entries.
Status	For the steppings affected, see the Summary Table of Changes.

044	Execution of VAESENCLAST Instruction May Produce a #NM Exception Instead of a #UD Exception
Problem	Execution of VAESENCLAST with VEX.L= 1 should signal a #UD (Invalid Opcode) exception, however, due to the erratum, a #NM (Device Not Available) exception may be signaled.
Implication	As a result of this erratum, an operating system may restore AVX and other state unnecessarily.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

045	Intel® SGX Enclave Accesses to the APIC-Access Page May Cause APIC-Access VM Exits
Problem	In VMX non-root operation, Intel SGX (Software Guard Extensions) enclave accesses to the APIC-access page may cause APIC-access VM exits instead of page faults.
Implication	A VMM (virtual-machine monitor) may receive a VM exit due to an access that should have caused a page fault, which would be handled by the guest OS (operating system).
Workaround	A VMM avoids this erratum if it does not map any part of the EPC (Enclave Page Cache) to the guest's APIC-access address; an operating system avoids this erratum if it does not attempt indirect enclave accesses to the APIC.
Status	For the steppings affected, see the Summary Table of Changes.

046	CR3 Filtering Does Not Compare Bits [11:5] of CR3 and IA32_RTIT_CR3_MATCH in PAE Paging Mode
Problem	In PAE paging mode, the CR3[11:5] are used to locate the page-directory-pointer table. Due to this erratum, those bits of CR3 are not compared to IA32_RTIT_CR3_MATCH (MSR 572H) when IA32_RTIT_CTL.CR3Filter (MSR 570H, bit 7) is set.
Implication	If multiple page-directory-pointer tables are co-located within a 4KB region, CR3 filtering will not be able to distinguish between them so additional processes may be traced.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

047	x87 FDP Value May be Saved Incorrectly
Problem	Execution of the FSAVE, FNSAVE, FSTENV, or FNSTENV instructions in real-address mode or virtual-8086 mode may save an incorrect value for the x87 FDP (FPU data pointer). This erratum does not apply if the last non-control x87 instruction had an unmasked exception.
Implication	Software operating in real-address mode or virtual-8086 mode that depends on the FDP value for non-control x87 instructions without unmasked exceptions may not operate properly.

Workaround	None identified. Software should use the FDP value saved by the listed instructions only when the most recent non-control x87 instruction incurred an unmasked exception.
Status	For the steppings affected, see the Summary Table of Changes.

048	**PECI Frequency Limited to 1 MHz**
Problem	The PECI (Platform Environmental Control Interface) 3.1 specification's operating frequency range is 0.2 MHz to 2 MHz. Due to this erratum, PECI may be unreliable when operated above 1 MHz.
Implication	Platforms attempting to run PECI above 1 MHz may not behave as expected.
Workaround	None identified. Platforms should limit PECI operating frequency to 1 MHz.
Status	For the steppings affected, see the Summary Table of Changes.

049	**Processor Graphics IOMMU Unit May Not Mask DMA Remapping Faults**
Problem	Intel® Virtualization Technology for Directed I/O specification specifies setting the FPD (Fault Processing Disable) field in the context (or extended-context) entry of IOMMU to mask recording of qualified DMA remapping faults for DMA requests processed through that context entry. Due to this erratum, the IOMMU unit for Processor Graphics device may record DMA remapping faults from Processor Graphics device (Bus: 0; Device: 2; Function: 0) even when the FPD field is set to 1.
Implication	Software may continue to observe DMA remapping faults recorded in the IOMMU Fault Recording Register even after setting the FPD field.
Workaround	None identified. Software may mask the fault reporting event by setting the IM (Interrupt Mask) field in the IOMMU Fault Event Control register (Offset 038H in GFXVTBAR).
Status	For the steppings affected, see the Summary Table of Changes.

050	**Intel® PT CYCThresh Value of 13 is Not Supported**
Problem	Intel PT (Intel® Processor Trace) CYC (Cycle Count) threshold is configured through CYCThresh field in bits [22:19] of IA32_RTIT_CTL MSR (570H). A value of 13 is advertised as supported by CPUID (leaf 14H, sub-lead 1H). Due to this erratum, if CYCThresh is set to 13 then the CYC threshold will be 0 cycles instead of 4096 (213-1) cycles.
Implication	CYC packets may be issued in higher rate than expected if threshold value of 13 is used.
Workaround	None identified. Software should not use value of 13 for CYC threshold.
Status	For the steppings affected, see the Summary Table of Changes.

051	**Enabling VMX-Preemption Timer Blocks HDC Operation**
Problem	HDC (Hardware Duty Cycling) will not put the physical package into the forced idle state while any logical processor is in VMX non-root operation and the "activate VMX-preemption timer" VM-execution control is 1.

Implication	HDC will not provide the desired power reduction when the VMX-preemption timer is active in VMX non-root operation.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

052	Integrated Audio Codec May Not be Detected
Problem	Integrated Audio Codec may lose power when LPSP (Low-Power Single Pipe) mode is enabled for an eDP* (embedded DisplayPort) or DP/HDMI ports. Platforms with Intel® SST (Intel® Smart Sound Technology) enabled are not affected.
Implication	The Audio Bus driver may attempt to do enumeration of Codecs when eDP or DP/HDMI port enters LPSP mode, due to this erratum, the Integrated Audio Codec will not be detected and audio maybe be lost.
Workaround	Intel® Graphics Driver 15.40.11.4312 or later will prevent the Integrated Audio Codec from losing power when LPSP mode is enabled.
Status	For the steppings affected, see the Summary Table of Changes.

053	Display Flickering May be Observed with Specific eDP Panels
Problem	The processor may incorrectly configure transmitter buffer characteristics if the associated eDP panel requests VESA equalization preset 3, 5, 6, or 8.
Implication	Display flickering or display loss maybe observed.
Workaround	Intel® Graphics Driver version 15.40.12.4326 or later contains a workaround for this erratum.
Status	For the steppings affected, see the Summary Table of Changes.

054	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
Problem	BTS (Branch Trace Store) and BTM (Branch Trace Message) send branch records to the Debug Store management area and system bus respectively. The Branch Predicted bit (bit 4 of eighth byte in BTS/BTM records) should report whether the most recent branch was predicted correctly. Due to this erratum, the Branch Predicted bit may be incorrect.
Implication	BTS and BTM cannot be used to determine the accuracy of branch prediction.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

055	MACHINE_CLEARS.MEMORY_ORDERING Performance Monitoring Event May Undercount
Problem	The performance monitoring event MACHINE_CLEARS.MEMORY_ORDERING (Event C3H; Umask 02H) counts the number of machine clears caused by memory ordering conflicts. However due to this erratum, this event may undercount for VGATHER*/VPGATHER* instructions of four or more elements.
Implication	MACHINE_CLEARS.MEMORY_ORDERING performance monitoring event may undercount.

Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

056	**CTR_FRZ May Not Freeze Some Counters**
Problem	IA32_PERF_GLOBAL_STATUS.CTR_FRZ (MSR 38EH, bit 59) is set when either (1) IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI (MSR 1D9H, bit 12) is set and a PMI is triggered, or (2) software sets bit 59 of IA32_PERF_GLOBAL_STATUS_SET (MSR 391H). When set, CTR_FRZ should stop all core performance monitoring counters from counting. However, due to this erratum, IA32_PMC4-7 (MSR C5-C8H) may not stop counting. IA32_PMC4-7 are only available when a processor core is not shared by two logical processors.
Implication	General performance monitoring counters 4-7 may not freeze when IA32_PERF_GLOBAL_STATUS.CTR_FRZ is set.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

057	**Instructions And Branches Retired Performance Monitoring Events May Overcount**
Problem	The performance monitoring events INST_RETIRED (Event C0H; any Umask value) and BR_INST_RETIRED (Event C4H; any Umask value) count instructions retired and branches retired, respectively. However, due to this erratum, these events may overcount in certain conditions when: - Executing VMASKMOV* instructions with at least one masked vector element - Executing REP MOVS or REP STOS with Fast Strings enabled (IA32_MISC_ENABLES MSR (1A0H), bit 0 set) - An MPX #BR exception occurred on BNDLDX/BNDSTX instructions and the BR_INST_RETIRED (Event C4H; Umask is 00H or 04H) is used.
Implication	INST_RETIRED and BR_INST_RETIRED performance monitoring events may overcount.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

058	**Some OFFCORE_RESPONSE Performance Monitoring Events May Overcount**
Problem	The performance monitoring events OFFCORE_RESPONSE (Events B7H and BBH) should count off-core responses matching the request-response configuration specified in MSR_OFFCORE_RSP_0 and MSR_OFFCORE_RSP_1 (1A6H and 1A7H, respectively) for core-originated requests. However, due to this erratum, DMND_RFO (bit 1), DMND_IFETCH (bit 2) and OTHER (bit 15) request types may overcount.
Implication	Some OFFCORE_RESPONSE events may overcount.
Workaround	None identified. Software may use the following model-specific events that provide related performance monitoring data: OFFCORE_REQUESTS (all sub-events), L2_TRANS.L2_WB and L2_RQSTS.PF_MISS.
Status	For the steppings affected, see the Summary Table of Changes.

059	Instructions Fetch #GP After RSM During Inter® PT May Push Incorrect RFLAGS Value on Stack
Problem	If Intel PT (Processor Trace) is enabled, a #GP (General Protection Fault) caused by the instruction fetch immediately following execution of an RSM instruction may push an incorrect value for RFLAGS onto the stack.
Implication	Software that relies on RFLAGS value pushed on the stack under the conditions described may not work properly.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

060	Access to SGX EPC Page in BLOCKED State is Not Reported as an SGX-Induced Page Fault
Problem	If a page fault results from attempting to access a page in the SGX (Intel® Software Guard Extensions) EPC (Enclave Page Cache) that is in the BLOCKED state, the processor does not set bit 15 of the error code and thus fails to indicate that the page fault was SGX-induced.
Implication	Due to this erratum, software may not recognize these page faults as being SGX-induced.
Workaround	Before using the EBLOCK instruction to marking a page as BLOCKED, software should use paging to mark the page not present.
Status	For the steppings affected, see the Summary Table of Changes.

061	MTF VM Exit on XBEGIN Instruction May Save State Incorrectly
Problem	Execution of an XBEGIN instruction while the monitor trap flag VM-execution control is 1 will be immediately followed by an MTF VM exit. If advanced debugging of RTM transactional regions has been enabled, the VM exit will erroneously save as instruction pointer the address of the XBEGIN instruction instead of the fallback instruction address specified by the XBEGIN instruction. In addition, it will erroneously set bit 16 of the pending-debug-exceptions field in the VMCS indicating that a debug exception or a breakpoint exception occurred.
Implication	Software using the monitor trap flag to debug or trace transactional regions may not operate properly. Intel has not observed this erratum with any commercially available software.
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

062	Performance Monitoring Counters May Undercount When Using CPL Filtering
Problem	Performance Monitoring counters configured to count only OS or only USR events by setting exactly one of bits 16 or 17 in IA32_PERFEVTSELx MSRs (186H-18DH) may not count for a brief period during the transition to a new CPL.

Implication	A measurement of ring transitions (using the edge-detect bit 18 in IA32_PERFEVTSELx) may undercount, such as CPL_CYCLES.RING0_TRANS (Event 5CH, Umask 01H). Additionally, the sum of an OS-only event and a USR-only event may not exactly equal an event counting both OS and USR. Intel has not observed any other software-visible impact
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

063	**Certain Non-Canonical IA32_BNDCFGS Values Will Not Cause VM-Entry Failures**
Problem	If the VM-entry controls Load IA32_BNDCFGS field (bit 16) is 1, VM-entry should fail when the value of the guest IA32_BNDCFGS field in the VMCS is not canonical (that is, when bits 63:47 are not identical). Due to this erratum, VM-entry does not fail if bits 63:48 are identical but differ from bit 47. In this case, VM-entry loads the IA32_BNDCFGS MSR with a value in which bits 63:48 are identical to the value of bit 47 in the VMCS field.
Implication	If the value of the guest IA32_BNDCFGS field in the VMCS is not canonical, VM-entry may load the IA32_BNDCFGS MSR with a value different from that of the VMCS field.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

064	**PEBS EventingIP Field May Be Incorrect Under Certain Conditions**
Problem	The EventingIP field in the PEBS (Processor Event-Based Sampling) record reports the address of the instruction that triggered the PEBS event. Under certain complex microarchitectural conditions, the EventingIP field may be incorrect.
Implication	When this erratum occurs, performance monitoring software may not attribute the PEBS events to the correct instruction.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

065	**HWP's Guaranteed_Performance Updated Only on Configurable TDP Changes**
Problem	According to HWP (Hardware P-states) specification, the Guaranteed_Performance field (bits[15:8]) in the IA32_HWP_CAPABILITIES MSR (771H) should be updated as a result of changes in the configuration of TDP, RAPL (Running Average Power Limit), and other platform tuning options that may have dynamic effects on the actual guaranteed performance support level. Due to this erratum, the processor will update the Guaranteed_Performance field only as a result of configurable TDP dynamic changes.
Implication	Software may read a stale value of the Guaranteed _Performance field.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

066	RF May be Incorrectly Set in The EFLAGS That is Saved on a Fault in PEBS or BTS
Problem	After a fault due to a failed PEBS (Processor Event Based Sampling) or BTS (Branch Trace Store) address translation, the RF (resume flag) may be incorrectly set in the EFLAGS image that is saved.
Implication	When this erratum occurs, a code breakpoint on the instruction following the return from handling the fault will not be detected. This erratum only happens when the user does not prevent faults on PEBS or BTS.
Workaround	Software should always prevent faults on PEBS or BTS.
Status	For the steppings affected, see the Summary Table of Changes.

067	Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters
Problem	Due to this erratum, if IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI (MSR 1D9H, bit 12) is set to 1 when Intel PT (Processor Trace) triggers a ToPA (Table of Physical Addresses) PMI (PerfMon Interrupt), performance monitoring counters are not frozen as expected.
Implication	Performance monitoring counters will continue to count for events that occur during PMI handler execution.
Workaround	PMI handler software can programmatically stop performance monitoring counters upon entry.
Status	For the steppings affected, see the Summary Table of Changes.

068	HWP's Maximum_Performance Value is Reset to 0xFF
Problem	According to HWP (Hardware P-states) specification, the reset value of the Maximum_Performance field (bits [15:8]) in IA32_HWP_REQUEST MSR (774h) should be set to the value of IA32_HWP_CAPABILITIES MSR (771H) Highest_Performance field (bits[7:0]) after reset. Due to this erratum, the reset value of Maximum_Performance is always set to 0xFF.
Implication	Software may see an unexpected value in Maximum Performance field. Hardware clipping will prevent invalid performance states.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

069	HWP's Guaranteed_Performance and Relevant Status/Interrupt May be Updated More Than Once Per Second
Problem	According to HWP (Hardware P-states) specification, the Guaranteed_Performance field (bits[15:8]) in the IA32_HWP_CAPABILITIES MSR (771H) and the Guaranteed_Performance_Change (bit 0) bit in IA32_HWP_STATUS MSR (777H) should not be changed more than once per second nor should the thermal interrupt associated with the change to these fields be signaled more than once per second. Due to this erratum, the processor may change these fields and generate the associated interrupt more than once per second
Implication	HWP interrupt rate due to Guaranteed_Performance field change can be higher than specified

Workaround	Clearing the Guaranteed_Performance_Change status bit no more than once per second will ensure that interrupts are not generated at too fast a rate
Status	For the steppings affected, see the Summary Table of Changes.

070	**Some Memory Performance Monitoring Events May Produce Incorrect Results When Filtering on Either OS or USR Modes**
Problem	The memory at-retirement performance monitoring events (listed below) may produce incorrect results when a performance counter is configured in OS-only or USR-only modes (bits 17 or 16 in IA32_PERFEVTSELx MSR). Counters with both OS and USR bits set are not affected by this erratum. The list of affected memory at-retirement events is as follows: MEM_INST_RETIRED.STLB_MISS_LOADS event D0H, umask 11H MEM_INST_RETIRED.STLB_MISS_STORES event D0H, umask 12H MEM_INST_RETIRED.LOCK_LOADS event D0H, umask 21H MEM_INST_RETIRED.SPLIT_LOADS event D0H, umask 41H MEM_INST_RETIRED.SPLIT_STORES event D0H, umask 42H MEM_LOAD_RETIRED.L2_HIT event D1H, umask 02H MEM_LOAD_RETIRED.L3_HIT event D1H, umask 04H MEM_LOAD_RETIRED.L4_HIT event D1H, umask 80H MEM_LOAD_RETIRED.L1_MISS event D1H, umask 08H MEM_LOAD_RETIRED.L2_MISS event D1H, umask 10H MEM_LOAD_RETIRED.L3_MISS event D1H, umask 20H MEM_LOAD_RETIRED.FB_HIT event D1H, umask 40H MEM_LOAD_L3_HIT_RETIRED.XSNP_MISS event D2H, umask 01H MEM_LOAD_L3_HIT_RETIRED.XSNP_HIT event D2H, umask 02H MEM_LOAD_L3_HIT_RETIRED.XSNP_HITM event D2H, umask 04H MEM_LOAD_L3_HIT_RETIRED.XSNP_NONE event D2H, umask 08H
Implication	The listed performance monitoring events may produce incorrect results including PEBS records generated at an incorrect point
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

071	**HWP May Generate Thermal Interrupt While Not Enabled**
Problem	Due to this erratum, the conditions for HWP (Hardware P-states) to generate a thermal interrupt on a logical processor may generate thermal interrupts on both logical processors of that core.
Implication	If two logical processors of a core have different configurations of HWP (e.g. only enabled on one), an unexpected thermal interrupt may occur on one logical processor due to the HWP settings of the other logical processor.
Workaround	Software should configure HWP consistently on all logical processors of a core.
Status	For the steppings affected, see the Summary Table of Changes.

072	Camera Device Does Not Issue an MSI When INTx is Enabled
Problem	When both MSI (Message Signaled Interrupts) and legacy INTx are enabled by the camera device, INTx is asserted rather than issuing the MSI, in violation of the PCI Local Bus Specification.
Implication	Due to this erratum, camera device interrupts can be lost leading to device failure.
Workaround	The camera device must disable legacy INTx by setting bit 10 of PCICMD (Bus 0; Device 5; Function 0; Offset 04H) before MSI is enabled
Status	For the steppings affected, see the Summary Table of Changes.

073	Attempts to Retrain a PCIe* Link May be Ignored
Problem	A PCIe link should retrain when Retrain Link (bit 5) in the Link Control register (Bus 0; Device 1; Functions 0,1,2; Offset 0xB0) is set. Due to this erratum, if the link is in the L1 state, it may ignore the retrain request
Implication	The PCIe link may not behave as expected.
Workaround	It is possible for the BIOS to contain a workaround for this erratum.
Status	For the steppings affected, see the Summary Table of Changes.

074	PCIe* Port Does Not Support DLL Link Activity Reporting
Problem	The PCIe Base specification requires DLL (Data Link Layer) Link Activity Reporting when 8 GT/s link speed is supported. Due to this erratum, link activity reporting is not supported
Implication	Due to this erratum, PCIe port does not support DLL Link Activity Reporting when 8 GT/s is supported.
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

075	BNDLDX And BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access
Problem	BNDLDX and BNDSTX instructions access the bound's directory and table to load or store bounds. These accesses should signal #GP (general protection exception) when the address is not canonical (i.e. bits 48 to 63 are not the sign extension of bit 47). Due to this erratum, #GP may not be generated by the processor when a non-canonical address is used by BNDLDX or BNDSTX for their bound directory memory access.
Implication	Intel has not observed this erratum with any commercially available software.
Workaround	Software should use canonical addresses for bound directory accesses.
Status	For the steppings affected, see the Summary Table of Changes.

076	RING_PERF_LIMIT_REASONS May be Incorrect
Problem	Under certain conditions, RING_PERF_LIMIT_REASONS (MSR 6B1H) may incorrectly assert the OTHER status bit (bit 8) as well as the OTHER log bit (bit 24).

Implication	When this erratum occurs, software using this register will incorrectly report clipping because of the OTHER reason.
Workaround	It is possible for the BIOS to contain a workaround for this erratum.
Status	For the steppings affected, see the Summary Table of Changes.

077	**Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions**
Problem	The performance monitoring events MEM_TRANS_RETIRED.LOAD_LATENCY_* (Event CDH; UMask 01H; any latency) count load instructions whose latency exceed a predefined threshold, where the loads are randomly selected using the load latency facility (an extension of PEBS). However due to this erratum, these events may count incorrectly for VGATHER*/VPGATHER* instructions.
Implication	The Load Latency Performance Monitoring events may be Inaccurate for Gather instructions.
Workaround	None identified
Status	For the steppings affected, see the Summary Table of Changes.

078	**EDRAM Corrected Error Events May Not be Properly Logged After a Warm Reset**
Problem	After a warm reset, an EDRAM corrected error may not be logged correctly until the associated machine check register is initialized. This erratum may affect IA32_MC8_STATUS or IA32_MC10_STATUS.
Implication	The EDRAM corrected error information may be lost when this erratum occurs.
Workaround	Data from the affected machine check registers should be read and the registers initialized as soon as practical after a warm reset.
Status	For the steppings affected, see the Summary Table of Changes.

079	**Some Bits in MSR_MISC_PWR_MGMT May be Updated on Writing Illegal Values to This MSR**
Problem	Attempts to write illegal values to MSR_MISC_PWR_MGMT (MSR 0x1AA) result in #GP (General Protection Fault) and should not change the MSR value. Due to this erratum, some bits in the MSR may be updated on writing an illegal value.
Implication	Certain fields may be updated with allowed values when writing illegal values to MSR_MISC_PWR_MGMT. Such writes will always result in #GP as expected.
Workaround	None identified. Software should not attempt to write illegal values to this MSR.
Status	For the steppings affected, see the Summary Table of Changes.

080	**Violations of Intel® Software Guard Extensions (Intel® SGX) Access-Control Requirements Produce #GP Instead of #PF**
Problem	Intel® Software Guard Extensions (Intel® SGX) define new access-control requirements on memory accesses. A violation of any of these requirements causes a page fault (#PF) that sets bit 15 (SGX) in the page-fault error code. Due to this erratum, these violations instead cause general-protection exceptions (#GP).

Implication	Software resuming from system sleep states S3 or S4 and relying on receiving a page fault from the above enclave accesses may not operate properly.
Workaround	Software can monitor #GP faults to detect that an enclave has been destroyed and needs to be rebuilt after resuming from S3 or S4
Status	For the steppings affected, see the Summary Table of Changes.

081	**IA32_RTIT_CR3_MATCH MSR Bits[11:5] Are Treated As Reserved**
Problem	Due to this erratum, bits[11:5] in IA32_RTIT_CR3_MATCH (MSR 572H) are reserved; an MSR write that attempts to set that field to a non-zero value will result in a #GP fault.
Implication	The inability to write the identified bit field does not affect the functioning of Intel® PT (Intel® Processor Trace) operation because, as described in erratum SKL061, the bit field that is the subject of this erratum is not used during Intel PT CR3 filtering.
Workaround	Ensure that bits 11:5 of the value written to IA32_RTIT_CR3_MATCH are zero, including cases where the selected page-directory-pointer-table base address has non-zero bits in this range.
Status	For the steppings affected, see the Summary Table of Changes.

082	**The Intel PT CR3 Filter is Not Re-evaluated on VM Entry**
Problem	On a VMRESUME or VMLAUNCH with both TraceEn[0] and CR3Filter[7] in IA32_RTIT_CTL (MSR 0570H) set to 1 both before the VM Entry and after, the new value of CR3 is not compared with IA32_RTIT_CR3_MATCH (MSR 0572H).
Implication	The Intel PT (Processor Trace) CR3 filtering mechanism may continue to generate packets despite a mismatching CR3 value, or may fail to generate packets despite a matching CR3, as a result of an incorrect value of IA32_RTIT_STATUS.ContextEn[1] (MSR 0571H) that results from the failure to re-evaluate the CR3 match on VM entry.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

083	**Display Slowness May be Observed Under Certain Display Commands Scenario**
Problem	Back to back access to the VGA register ports (I/O addresses 0x3C2, 0x3CE, 0x3CF) will experience higher than expected latency.
Implication	Due to this erratum, the processor may redraw the slowly when in VGA mode.
Workaround	None identified.
Status	For the steppings affected, see the Summary Table of Changes.

084	**CPUID TLB Associativity Information is Inaccurate**
Problem	CPUID leaf 2 (EAX=02H) TLB information inaccurately reports that the shared 2nd-Level TLB is 6-way set associative (value C3H), although it is 12-way set associative. Other information reported by CPUID leaf 2 is accurate.

Implication	Software that uses CPUID shared 2nd-level TLB associativity information for value C3H may operate incorrectly. Intel has not observed this erratum to impact the operation of any commercially available software
Workaround	None identified. Software should ignore the shared 2nd-Level TLB associativity information reported by CPUID for the affected processors.
Status	For the steppings affected, see the Summary Table of Changes.

085	**Using Different Vendors For 2400 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang**
Problem	When using 2400 MHz DDR4 UDIMMs from different vendors within the same channel, a higher rate of correctable errors may occur or the system may hang.
Implication	Due to this erratum, reported correctable error counts may increase or the system may hang.
Workaround	None identified. Use a single vendor for 2400 MHz UDIMMs.
Status	For the steppings affected, see the Summary Table of Changes.

086	**Two DIMMs Per Channel 2133MHz DDR4 SODIMM Daisy-Chain Systems With Different Vendors May Hang**
Problem	When, on a single memory channel with 2133 MHz DDR4 SODIMMs, mixing different vendors or mixing single rank and dual rank DIMMs, may lead to a higher rate of correctable error to system hangs.
Implication	Due to this erratum, reported correctable error counts may increase or system may hang.
Workaround	Use a single vendor for and do not mix single rank and dual rank 2133 MHz DDR4 SODIMM.
Status	For the steppings affected, see the Summary Table of Changes.

087	**Unpredictable System Behavior May Occur in DDR4 Multi-Rank System**
Problem	Due to incorrect configuration of DDR4 ODT by BIOS, it is possible for a multi-rank system to violate section 4.27 of the DDR4 JEDEC spec revision JESED79-4A.
Implication	Due to this erratum, complex microarchitectural conditions may result in unpredictable system behavior.
Workaround	A BIOS workaround has been identified.
Status	For the steppings affected, see the Summary Table of Changes.

088	**Processor May Hang on Complex Sequence of Conditions**
Problem	A complex set of architectural and micro-architectural conditions may lead to a processor hang with an internal timeout error (MCACOD 0400H) logged into IA32_MCi_STATUS. When both logical processors in a core are active, this erratum will not occur unless there is no store on one of the logical processors for more than 10 seconds.
Implication	This erratum may result in a processor hang. Intel has not observed this erratum with any commercially available software.

Workaround	None Identified.
Status	For the steppings affected, see the Summary Table of Changes.

089	**Potential Partial Trace Data Loss in Intel® Trace Hub ODLA When Storing to Memory**
Problem	When Intel® Trace Hub's ODLA (On-Die Logic Analyzer) is configured to trace to memory, under complex microarchitectural conditions, the trace may lose a timestamp.
Implication	Some ODLA trace data may be lost. This erratum does not affect other trace data sources. Typically, lost trace data will be displayed as "OVERFLOW." Subsequent timestamps will allow the trace decoder to resume tracing. Intel has not observed this erratum in commercially available software.
Workaround	None identified. For a particular workload, changing the memory buffer size or disabling deep compression may eliminate the microarchitectural condition that causes the erratum.
Status	For the steppings affected, see the Summary Table of Changes.

090	**Using Different Vendors For 2666 MHz DDR4 UDIMMs May Cause Correctable Errors or a System Hang**
Problem	When using 2666 MHz DDR4 UDIMMs from different vendors or mixing single rank and dual rank DIMMs, within the same channel, a higher rate of correctable errors may occur or the system may hang.
Implication	Due to this erratum, reported correctable error counts may increase or the system may hang.
Workaround	None identified. Use a single vendor and do not mix single rank and dual rank for 2666 MHz UDIMMs.
Status	For the steppings affected, see the Summary Table of Changes.

§

Specification Changes

There are no Specification Changes in this Specification Update revision.

§