



Intel[®] Xeon Phi[™] Processor x200 Product Family

Specification Update

October 2017



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.

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.

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.

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

Intel, Xeon Phi, Xeon, 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 | 4 |
| Preface | 5 |
| Affected Documents | 5 |
| Related Documents | 6 |
| Nomenclature..... | 6 |
| Identification Information | 7 |
| Component Marking Information..... | 8 |
| Summary Tables of Changes | 9 |
| Codes Used in Summary Tables | 9 |
| Stepping..... | 9 |
| Page | 9 |
| Status..... | 9 |
| Row | 9 |
| Errata | 10 |
| Errata Summary | 12 |
| Specification Changes..... | 19 |
| Specification Clarifications..... | 19 |
| Documentation Changes | 19 |



Revision History

| Date | Revision | Description |
|---------------|----------|---|
| October 2017 | 012 | Added Errata KNL31 and KNL32. |
| April 2017 | 011 | Added Errata KNL24 through KNL30. |
| January 2017 | 010 | Added Errata KNL22 and KNL23. |
| December 2016 | 009 | Updated Affected Documents table titles. Updated Table 1. Added Errata KNL19 through KNL21. |
| November 2016 | 008 | Added Errata KNL15 to KNL18. Updated Table 1. |
| October 2016 | 007 | Updated Errata KNL#3 and added Errata KNL#6 to 14. |
| June 2016 | 006 | Added Errata#5. |
| June 2016 | 005 | Initial release |



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 this document and are no longer published in other documents.

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

Affected Documents

| Document Title | Document Number/Location |
|--|--------------------------|
| <i>Knights Landing Processor External Design Specification (EDS), Volume One: Architecture</i> | 547141 |
| <i>Knights Landing Processor External Design Specification (EDS), Volume Two: Registers Part A</i> | 546888 |
| <i>Knights Landing Processor External Design Specification (EDS), Volume Two: Registers Part B</i> | 552048 |
| <i>Intel® Xeon Phi™ x200 - External Design Specification Volume 3 Electrical</i> | 540624 |
| <i>Knights Landing Processor Best Known Configuration (BKC)</i> | 557302 |
| <i>Groveport Platform Design Guide</i> | 525616 |
| <i>Knights Landing Processors Thermal / Mechanical Specification and Design Guide</i> | 538144 |
| <i>Knights Landing Processor BIOS Writer's Guide (BWG), Volume: 1</i> | 545121 |
| <i>Knights Landing Processor BIOS Writer's Guide (BWG), Volume: 2</i> | 545122 |



Related Documents

| Document Title | Document Number/ Location |
|--|---|
| AP-485, Intel® Processor Identification and the CPUID Instruction | Note 1 |
| Processor VR Module and Enterprise VR Down 12.5 [VRM/ERVD] Design Guide | 495853 |
| Intel® Advanced Vector Extensions Programming Reference | Note 1 |
| Intel® Virtualization Technology for Directed I/O Architecture | http:// download.intel.com/ technology/computing/ vptech/Intel(r)_VT_for_ Direct_IO.pdf |
| <i>Intel® 64 and IA-32 Architecture Software Developer's Manual</i> <ul style="list-style-type: none">• <i>Volume 1: Basic Architecture</i>• <i>Volume 2A: Instruction Set Reference Manual A-M</i>• <i>Volume 2B: Instruction Set Reference Manual N-Z</i>• <i>Volume 3A: System Programming Guide</i>• <i>Volume 3B: System Programming Guide</i>• <i>IA-32 Intel® Architecture Optimization Reference Manual</i> | http://www.intel.com/ products/processor/ manuals/index.htm |

Note: Contact your Intel representative for the latest revision and order number of this document.

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.

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

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

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

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

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



Identification Information

Component Identification via Programming Interface

The Intel® Xeon Phi™ Processor x200 stepping can be identified by the following register contents:


| Reserved | Extended Family ¹ | Extended Model ² | Reserved | Processor Type | Family Code ³ | Model Number ⁴ | Stepping ID ^{5,6} |
|----------|------------------------------|-----------------------------|----------|----------------|--------------------------|---------------------------|----------------------------|
| 31:28 | 27:20 | 19:16 | 15:14 | 13:12 | 11:8 | 7:4 | 3:0 |
| | 00000000b | 0101b | | 00b | 0110b | 0111b | varies per stepping |

Notes:

1. The Extended Family, Bits [27:20] are used in conjunction with the Family Code, specified in Bits [11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium 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. See Table 1 for the processor stepping ID number in the CPUID information.
6. Refer to the *Processor BIOS Writer's Guide (BWG), Combined Volumes 1-3* for additional information.

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

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



Component Marking Information

Figure 1. Intel® Xeon Phi™ Processor x200 Top-side Markings (Example)

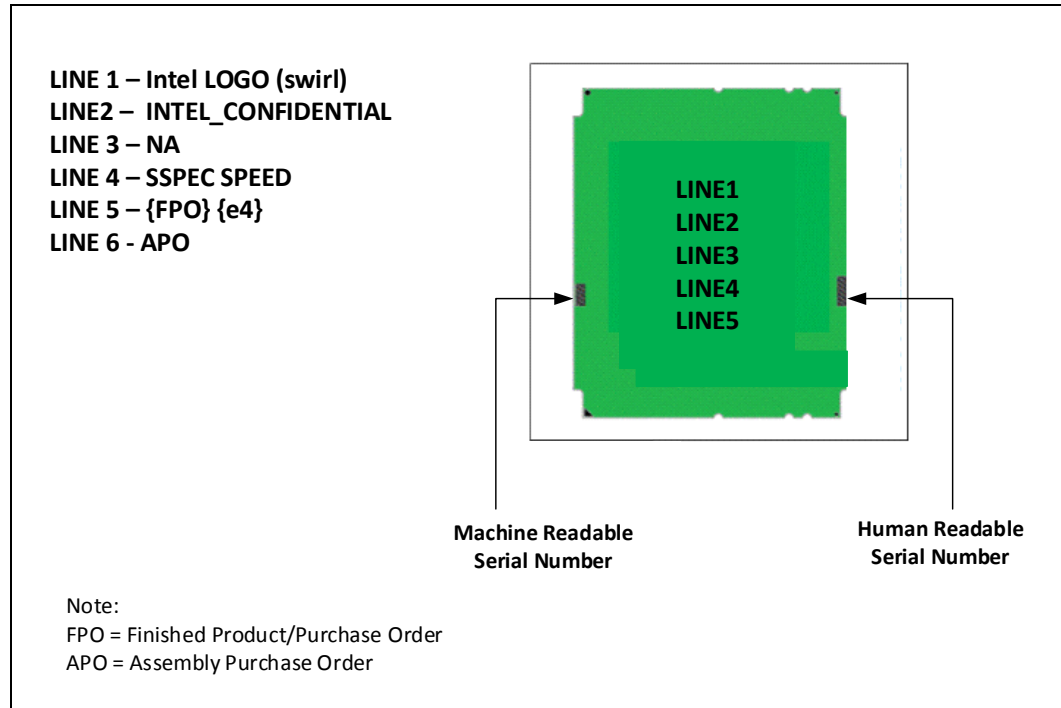



Table 1. Intel® Xeon Phi™ Processor x200 Product Family Identification

| SSPEC# | PRQ Version | Steppings | Max Speed | | | | TDP (W) | # Cores | MCDRAM (GB) |
|--------|-------------|-----------|---------------------------|-------------------------|------------------|------------------|---------|---------|-------------|
| | | | Core / 1 core Turbo (GHz) | Intel® Mesh Freq. (GHz) | DDR Freq. (MT/s) | OPIO Freq (GT/s) | | | |
| R2MD | Early Ship | B0 | 1.4/1.6 | 1.7 | 2400 | 7.2 | 215 | 68 | 16 |
| R2MF | Early Ship | B0 | 1.3/1.5 | 1.7 | 2400 | 7.2 | 215 | 64 | 16 |
| R2ME | Early Ship | B0 | 1.3/1.5 | 1.6 | 2133 | 6.4 | 215 | 64 | 16 |
| R2WY | Production | B0 | 1.5/1.7 | 1.7 | 2400 | 7.2 | 245 | 72 | 16 |
| R2X1 | Production | B0 | 1.4/1.6 | 1.7 | 2400 | 7.2 | 215 | 68 | 16 |
| R2X3 | Production | B0 | 1.3/1.5 | 1.7 | 2400 | 7.2 | 215 | 64 | 16 |
| R2X4 | Production | B0 | 1.3/1.5 | 1.6 | 2133 | 6.4 | 215 | 64 | 16 |

Notes: These processors do not have support for Intel® Virtualization Technology (Intel® VT) and Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d).



Summary Tables of Changes

The following tables indicate the sighting, specification changes, specification clarifications, or documentation changes which apply to the processor. Intel may fix some of the sightings in a future stepping of the component, and account for the other outstanding issues through documentation or specification changes as noted. These tables use the following notations:

Codes Used in Summary Tables

Stepping

X:Errata exists in the stepping indicated. Specification Change or Clarification that applies to this stepping.

(No mark)

or (Blank box):This erratum is fixed in listed stepping or specification change does not apply to listed stepping.

Page

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

Status

Doc:Document change or update will be implemented.


Plan Fix:This erratum may be fixed in a future stepping of the product.

Fixed:This erratum has been previously fixed.

No Fix planned:There are no plans to fix this erratum.


Row

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




Errata

| Number | Stepping | Status | Errata |
|--------|----------|--------|--|
| | B0 | | |
| KNL1. | X | No Fix | Cacheable MMIO requests may block PCIe* memory requests in SMM mode |
| KNL2. | X | No Fix | DdrCrkTraining CSR Cannot Be Read |
| KNL3. | X | No Fix | Performance Monitoring Event NO_ALLOC_CYCLES Counts Incorrectly on Logical Processors 2 and 3 |
| KNL4. | X | No Fix | A and/or D Bit May Be Set on a Paging-Structure Entry That is Not Present or Has Reserved Bit Set |
| KNL5. | X | No Fix | PCI Express Inbound Memory Write With RO and NS Attributes Set Will Cause the Processor to Hang |
| KNL6. | X | No Fix | CLTT Peci Pass-Through Mode Does Not Work |
| KNL7. | X | No Fix | Uncorrectable Memory Machine Check May Not Set IA32_MC3_STATUS.EN |
| KNL8. | X | No Fix | Package C6 May Cause Incorrect APIC Timer Value |
| KNL9. | X | No Fix | Unsynchronized Cross-Modifying Code Operations Can Cause Unexpected Instruction Execution Results |
| KNL10. | X | No Fix | Upper 48 Bits of the Scatter/Gather Mask Register do Not Function as Expected |
| KNL11. | X | No Fix | DR6 May be Zero After Data Breakpoint on Gather, Scatter or VRSQRT14 Instructions |
| KNL12. | X | No Fix | Performance Monitoring Event CPU_CLK_UNHALTED.THREAD_P Counts Incorrectly on Counter 1 If AnyThread Bit is Set For Counter 0 |
| KNL13. | X | No Fix | An Instruction With 7 or More Prefixes May Cause a Spurious #PF or Spuriously Read UC Memory |
| KNL14. | X | No Fix | Machine Check Exception MSCOD is Incorrect For Poisoned Case |
| KNL15. | X | No Fix | POC_RESET_STRAPS CSR Does Not Report Correct Values |
| KNL16. | X | No Fix | Invalid VEX Instructions May Not Signal a #GP |
| KNL17. | X | No Fix | Performance Monitoring OvfUncore Capability is Not Functional |
| KNL18. | X | No Fix | PECI PCS (Package Configuration Space) Read For Max Thread ID is Incorrect |
| KNL19. | X | No Fix | BIST Results Always Indicate BIST Failure |
| KNL20. | X | No Fix | Incorrect Linear-Address Translation May Be Used for Instruction Fetch |
| KNL21. | X | No Fix | Accesses Between TOHM and $2^{45} - 1$ May Lead to Unpredictable System Behavior |
| KNL22. | X | No Fix | System May Hang When Loading a Second Microcode Update |
| KNL23. | X | No Fix | Programmatic and Peci SMBus Access May Not Work as Intended |
| KNL24. | X | No Fix | System May Hang During Warm Reset |
| KNL25. | X | No Fix | Operating With DDR4-2400 Memory May Cause Unpredictable System Behavior |



| Number | Stepping | Status | Errata |
|--------|----------|--------|--|
| | B0 | | |
| KNL26. | X | No Fix | Enabling DDR Opportunistic Self Refresh May Lead to Memory Errors |
| KNL27. | X | No Fix | Certain Memory Controller Uncorrectable Errors Do Not Signal a Machine Check |
| KNL28. | X | No Fix | PCC is Not Set For Certain Memory Controller Uncorrectable Errors When Poison is Enabled |
| KNL29. | X | No Fix | Memory Controller Machine Check Errors May be Incorrectly Logged |
| KNL30. | X | No Fix | Complex Set of Conditions May Result in Unpredictable System Behavior |
| KNL31. | X | No Fix | Processor May Hang and Machine Check |
| KNL32. | X | No Fix | Unpredictable System Behavior May Occur With MCDRAM Scrubbing Enabled |



Errata Summary

KNL1. Cacheable MMIO requests may block PCIe* memory requests in SMM mode

Problem: When operating in SMM mode, cacheable MMIO requests may prevent PCIe memory requests from completing.

Implication: If PCIe memory requests are blocked by cacheable MMIO requests, the system may hang.

Workaround: Do not issue cacheable MMIO request during SMM mode.

Status: No Fix

KNL2. DdrCrClkTraining CSR Cannot Be Read

Problem: Due to this erratum, two fields in DdrCrClkTraining CSR (Bus: 2; Device:17; Function: 5; Offset: A1Ch), that should have been read-write, VmseLogicDelay0 (bit 29) and VmseLogicDelay2 (bit 30), are write-only; reading the CSR always returns 0s for these fields.

Implication: Data read from this register is not reliable. Intel has not observed this erratum to affect processor functionality.

Workaround: None identified.

Status: No Fix

KNL3. Performance Monitoring Event NO_ALLOC_CYCLES Counts Incorrectly on Logical Processors 2 and 3

Problem: Due to this erratum, all sub-events of performance monitoring event NO_ALLOC_CYCLES (Event CAH) will be incorrect on logical processors 2 and 3. Logical processors 0 and 1 are not affected by this erratum.

Implication: Using the NO_ALLOC_CYCLES performance monitoring event on logical processors 2 or 3 will not provide reliable results.

Workaround: Use the results of the NO_ALLOC_CYCLES performance monitoring event only when running the workload on logical processors 0 and/or 1.

Status: No Fix

KNL4. A and/or D Bit May Be Set on a Paging-Structure Entry That is Not Present or Has Reserved Bit Set

Problem: The A (Accessed, bit 5) and/or D (Dirty, bit 6) bits in a paging-structure entry (e.g., a Page-Table Entry) may be set to 1 even when that entry has its Present bit cleared or has a reserved bit set. This can only occur when one logical processor has cleared the Present bit or set a reserved bit in a paging-structure entry, while at the same time another logical processor accesses the contents of a linear address mapped by that entry.

Implication: Software that does not expect hardware to modify a paging-structure entry when it is marked not present, or has a reserved bit set, may behave unexpectedly.

Workaround: Operating systems can take steps to ensure they are not exposed to this erratum. Contact your Intel representative for further information.

Status: No Fix



KNL5. PCI Express Inbound Memory Write With RO and NS Attributes Set Will Cause the Processor to Hang

Problem: Inbound memory writes with the RO (Relaxed Ordering) and NS (No Snoop) attributes set in the TLP will cause the processor to hang.

Implication: When this erratum occurs, the system will experience a hang.

Workaround: BIOS should disable relaxed ordering for inbound transactions.

Status: No Fix

KNL6. CLTT PECI Pass-Through Mode Does Not Work

Problem: When CLTT (Closed Loop Thermal Throttling) PECI Pass-through mode is enabled, PECI thermal update messages are dropped.

Implication: Integrated Memory Controller does not receive thermal data from the PECI update messages.

Workaround: Do not use CLTT PECI Pass-through mode for DDR throttling.

Status: No Fix

KNL7. Uncorrectable Memory Machine Check May Not Set IA32_MC3_STATUS.EN

Problem: When an uncorrectable memory error follows a correctable memory error, the error logged may not set the EN field (bit 60) in the IA32_MC3_STATUS MSR (40CH).

Implication: An uncorrectable machine check may be logged with EN incorrectly set to 0, possibly leading to unexpected machine check handler behavior.

Workaround: None Identified.

Status: No Fix

KNL8. Package C6 May Cause Incorrect APIC Timer Value

Problem: The APIC timer CCR (Current Count Register) may be too high after waking from package C6. The next timer interrupt will be delivered at the correct time. However, in Periodic mode, the reload of the APIC timer may be delayed.

Implication: When this erratum occurs, CCR reads may be too high and in Periodic mode an APIC timer interrupt may occur later than expected.

Workaround: None Identified.

Status: No Fix

KNL9. Unsynchronized Cross-Modifying Code Operations Can Cause Unexpected Instruction Execution Results

Problem: The act of one processor or system bus master writing data into a currently executing code segment of a second processor with the intent of having the second processor execute that data as code is called cross-modifying code (XMC). XMC that does not force the second processor to execute a synchronizing instruction prior to execution of the new code is called unsynchronized XMC. Software using unsynchronized XMC to modify the instruction byte stream of a processor can see unexpected or unpredictable execution behavior from the processor that is executing the modified code.

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



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

Status: No Fix

KNL10. Upper 48 Bits of the Scatter/Gather Mask Register do Not Function as Expected

Problem: When one element of a scatter or gather instruction faults, the upper 48 bits of the mask register will be cleared. Ideally, only bits corresponding to completed operations should be cleared.

Implication: Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instructions supported by KNL use only the lower 16 bits of these mask registers. Intel has not observed this erratum to affect commercially available software.

Workaround: None Identified

Status: No Fix

KNL11. DR6 May be Zero After Data Breakpoint on Gather, Scatter or VRSQRT14 Instructions

Problem: If gather, scatter or VRSQRT14 instructions trigger a data breakpoint, the Debug Status Register (DR6) may be cleared.

Implication: Software will see a breakpoint trap but with no indication of which data breakpoint register was triggered.

Workaround: None Identified.

Status: No Fix

KNL12. Performance Monitoring Event CPU_CLK_UNHALTED.THREAD_P Counts Incorrectly on Counter 1 If AnyThread Bit is Set For Counter 0

Problem: Due to this erratum, GP (general-purpose) counter 1 may overcount CPU_CLK_UNHALTED.THREAD_P (Event 3CH Umask 00H) if IA32_PERFEVTSEL0.AnyThread (MSR 186H, bit 21) is set.

Implication: CPU_CLK_UNHALTED.THREAD_P event should not be relied upon on GP counter 1.

Workaround: Use only GP counter 0 or fixed counter 1 for unhalting core cycles.

Status: No Fix

KNL13. An Instruction With 7 or More Prefixes May Cause a Spurious #PF or Spuriously Read UC Memory

Problem: An instruction with 7 or more prefixes can result in a spurious code fetch that may signal a #PF (Page Fault) or read UC (un-cacheable) memory.

Implication: A spurious UC memory access may result in unexpected and undesired side effect(s). The OS may mishandle a spurious #PF due to there being no reason for the #PF

Workaround: Avoid using 7 or more prefixes on an instruction. If limiting the number of prefixes is not feasible then marking MMIO (memory mapped I/O) as XD (execute disable) in the page tables will prevent speculative reads from UC MMIO.

Status: No Fix



KNL14. Machine Check Exception MSCOD is Incorrect For Poisoned Case

Problem: When poisoned data is received at the EDC, MSCOD should be logged as a data error (encoding = 0x2). Instead, the encoding is set for an uncorrectable error (0x40). The effected machine check bank registers are IA32_MC7_STATUS through IA32_MC14_STATUS.

Implication: Error reporting for poisoned data is incorrectly reported as an uncorrected error and not a data error. The error flow will be incorrect for poisoned data.

Workaround: None Identified.

Status: No Fix

KNL15. POC_RESET_STRAPS CSR Does Not Report Correct Values

Problem: The POC_RESET_STRAPS CSR (Bus: 1; Device 30; Function 1; Offset: A0H) does not correctly report the strap settings.

Implication: The register cannot be used to check the strap settings.

Workaround: A BIOS workaround has been identified. Contact Intel representative for more information.

Status: No Fix

KNL16. Invalid VEX Instructions May Not Signal a #GP

Problem: Under certain conditions, invalid VEX instructions with prefixes may not signal a #GP.

Implication: Processor may not operate as expected with invalid VEX instructions.

Workaround: None Identified.

Status: No Fix

KNL17. Performance Monitoring OvfUncore Capability is Not Functional

Problem: Due to this erratum, IA32_PERF_GLOBAL_STATUS.OvfUncore (MSR 38EH, bit[61]) is always 0 and writing 1 to IA32_PERF_GLOBAL_OVF_CTRL.ClrOvfUncore (MSR 390H, bit[61]) signals #GP.

Implication: Software attempting to use OvfUncore capability may not function as expected.

Workaround: None identified.

Status: No Fix

KNL18. PECCI PCS (Package Configuration Space) Read For Max Thread ID is Incorrect

Problem: The PECCI command RdPkgConfig (Index 0) with Data Max Thread ID always returns a value of 0x12F regardless of the number of tiles enabled.

Implication: This PECCI command does not report an accurate Max Thread ID.

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

Status: No Fix


KNL19. BIST Results Always Indicate BIST Failure

Problem: BIST results in BIST_RESULTS_CFG_2 (Bus 0, Device 8, Function1, Offset 0xB8, bits [5:0]) returns 0 regardless of actual BIST status.

Implication: BIST results in BIST_RESULTS_CFG_2 are incorrect.

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

Status: No Fix



KNL20. Incorrect Linear-Address Translation May Be Used for Instruction Fetch

Problem: Under complex micro-architectural conditions when MT is active, operations which should invalidate instruction TLB entries may fail to do so. This may lead a later instruction fetch using a stale linear address translation.

Implication: When this erratum occurs, the processor may use incorrect translations, This may result in unexpected faults or other unpredictable system behavior.

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

Status: No Fix

KNL21. Accesses Between TOHM and $2^{45} - 1$ May Lead to Unpredictable System Behavior

Problem: A CPU access to addresses at or above TOHM (Top of High Memory) as configured in the TOHM CSR (Bus: 0; Device: 5; Function: 0; Offset: D8h) and below 2^{45} may lead to unpredictable system behavior when the L1 data prefetcher is enabled.

Implication: Unpredictable system behavior may occur.

Workaround: This address range inclusively between TOHM and $2^{45} - 1$ should be marked as not present in page or EPT tables and not used. Alternatively, MSR_MISC_FEATURE_CONTROL.L1_DATA_PREFETCH_DISABLE (MSR 01A4H, bit 0) can be set to 1 to disable the L1 data prefetcher.

Status: No Fix

KNL22. System May Hang When Loading a Second Microcode Update

Problem: System may hang during warm reset with a PCU_MC_STATUS.MCCOD = 0x0402 and PCU_MC_STATUS.MSEC_FW = 0x9C or 0x9E.

Implication: When this erratum occurs, the system may hang during a warm reset.

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

Status: No Fix

KNL23. Programmatic and PECI SMBus Access May Not Work as Intended

Problem: When BIOS locks access to SMBCNTL_0 (Bus 1; Device 30; Function 0; Offset 108H) and SMBCNTL_1 (Bus 1; Device 30; Function; Offset 108H), programmatic access outside SMM and PECI access may not be able to select their intended SMBus devices.

Implication: SMBus transactions will reference the previously addressed devices.

Workaround: It is possible for BIOS to contain a partial mitigation for this erratum that enables PECI access to change the unlocked portions of SMBCNTL_0 and SMBCNTL_1.

Status: No Fix

KNL24. System May Hang During Warm Reset

Problem: System may hang during warm reset with a PCU_MC_STATUS.MCCOD = 0x0402 and PCU_MC_STATUS.MSEC_FW = 0x9C or 0x9E.

Implication: When this erratum occurs, the system may hang during a warm reset.

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

Status: No Fix



KNL25. Operating With DDR4-2400 Memory May Cause Unpredictable System Behavior

Problem: Operating the processor with DDR4 memory configured to operate at 2400 MT/s may cause unpredictable system behavior.

Implication: When the erratum occurs, the system will exhibit unpredictable system behavior.

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

Status: No Fix

KNL26. Enabling DDR Opportunistic Self Refresh May Lead to Memory Errors

Problem: Correctable and uncorrectable memory errors may occur when DDR opportunistic self-refresh is enabled.

Implication: The system may experience a higher rate of memory errors when DDR opportunistic self-refresh is enabled, potentially leading to a system crash.

Workaround: It is possible for BIOS to contain code changes to work around this erratum.

Status: No Fix

KNL27. Certain Memory Controller Uncorrectable Errors Do Not Signal a Machine Check

Problem: Uncorrectable errors logged into the IA32_MCi_STATUS registers for bank 7 to bank 16 do not signal a machine check exception if the error cause was CA Parity, Data Parity, or Byte En Parity (MSCOD value of 0001H, 0002H, or 0004H, respectively).

Implication: The system may continue execution after encountering an uncorrectable error instead of responding to a machine check exception. Uncorrectable errors may lead to unpredictable system behavior.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. Implementing this workaround will cause signaling of IERR# instead of MCERR# when this uncorrectable error occurs. There may be an increased delay in error reporting.

Status: No Fix


KNL28. PCC is Not Set For Certain Memory Controller Uncorrectable Errors When Poison is Enabled

Problem: PCC field (bit[57]) of IA32_MCi_STATUS registers for bank 7 to bank 16 is not set as expected for a CA Parity error (MSCOD value of 0001H) when Poison is enabled.

Implication: The machine check handlers incorrectly log the error as an Uncorrectable Error with No Action, instead of logging an Uncorrectable Error and signaling a machine check exception.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. Implementing this workaround will cause signaling of IERR# when this uncorrectable error type occurs, in addition to the machine check exception. There may be an increased delay in error reporting.

Status: No Fix



KNL29. Memory Controller Machine Check Errors May be Incorrectly Logged

Problem: UC (Uncorrectable) errors logged into the IA32_MCi_STATUS registers for bank 7 to bank 16 will not indicate PCC=1 (Processor Context Corrupt, bit 57) and OVER=1 (bit 62) if the error is overwriting a previous UCNA (Uncorrected No Action Required) error.

Implication: The system may continue execution after encountering an uncorrectable error instead of responding to a machine check exception. Uncorrectable errors may lead to unpredictable system behavior.

Workaround: It is possible for the BIOS to contain a workaround for this erratum. Implementing this workaround will cause signaling of IERR# instead of MCERR# when this uncorrectable error occurs. There may be an increased delay in error reporting.

Status: No Fix

KNL30. Complex Set of Conditions May Result in Unpredictable System Behavior

Problem: A complex set of micro-architectural conditions may result in unpredictable system behavior.

Implication: When this erratum occurs, the system may exhibit unpredictable system behavior.

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

Status: No Fix

KNL31. Processor May Hang and Machine Check

Problem: Under complex micro-architectural conditions, the processor may hang, resulting in an Internal Timer Error Machine Check (IA32_MCi_STATUS.MCACOD=400H; bits [15:0]).

Implication: When this erratum occurs, the system may hang.

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

Status: No Fix

KNL32. Unpredictable System Behavior May Occur With MCDRAM Scrubbing Enabled

Problem: When the MCDRAM controller is configured in a Cache or Hybrid mode and error scrubbing is enabled, a complex set of micro-architectural conditions may lead to unpredictable system behavior.

Implication: When this erratum occurs, unpredictable system behavior may occur.

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

Status: No Fix

§



Specification Changes

None for this revision of the Specification update

Specification Clarifications

None for this revision of the Specification update

Documentation Changes

None for this revision of the Specification update

§