

Intel[®] Xeon[®] Processor Scalable Family

Specification Update

August 2020



Contents

Revision History	4
Preface	5
Summary Tables of Changes	7
Identification Information	13
Errata	24
Specification Changes	50
Specification Clarifications	51
Documentation Changes	52



Preface

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

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

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


Affected Documents

Document Title	Document Number/ Location
<i>Intel® Xeon® Processor Scalable Family Datasheet: Volume 1 - Electrical</i>	336062
<i>Intel® Xeon® Processor Scalable Family Datasheet: Volume 2 - Registers</i>	336063

Related Documents

Document Title	Document Number/ Location
<i>Intel® 64 and IA-32 Architecture Software Developer Manual, Volume 1: Basic Architecture</i>	253665 ¹
<i>Volume 2A: Instruction Set Reference, A-M</i>	253666 ¹
<i>Volume 2B: Instruction Set Reference, N-Z</i>	253667 ¹
<i>756B Volume 3A: System Programming Guide, Part 1</i>	253668 ¹
<i>Volume 3B: System Programming Guide, Part 2</i>	253669 ¹
<i>ACPI Specifications</i>	www.acpi.info ²

1. Document is available publicly at <http://developer.intel.com>.
2. Document available at www.acpi.info.



Summary Tables of Changes

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

Codes Used in Summary Tables

Stepping

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

Page

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

Status

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

Row



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




Specification Clarifications

No.	Specification Clarifications
1	None for this revision of this specification update.

Documentation Changes

No.	Documentation Changes
1	None for this revision of this specification update.



Identification Information

Component Identification via Programming Interface

The Intel® Xeon® Processor Scalable Family stepping can be identified by the following register contents:

Reserved	Extended Family ¹	Extended Model ²	Reserved	Processor Type ³	Family Code ⁴	Model Number ⁵	Stepping ID ⁶
31:28	27:20	19:16	15:13	12	11:8	7:4	3:0
	00000000b	0101b		0b	0110b	0101b	Varies per stepping

1. The Extended Family, bits [27:20] are used in conjunction with the Family Code, specified in bits [11:8], to indicate whether the processor belongs to the Intel386™, Intel486™, Pentium®, Pentium® Pro, Pentium® 4, Intel® Core™ processor family, or Intel® Core™ i7 family.
2. The Extended Model, bits [19:16] in conjunction with the Model Number, specified in bits [7:4], are used to identify the model of the processor within the processor’s family.
3. The Processor Type, specified in bit [12] indicates whether the processor is an original OEM processor, an Over Drive processor, or a dual processor (capable of being used in a dual processor system).
4. The Family Code corresponds to bits [11:8] of the EDX register after RESET, bits [11:8] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the generation field of the Device ID register accessible through Boundary Scan.
5. The Model Number corresponds to bits [7:4] of the EDX register after RESET, bits [7:4] of the EAX register after the CPUID instruction is executed with a 1 in the EAX register, and the model field of the Device ID register accessible through Boundary Scan.
6. The Stepping ID in bits [3:0] indicates the revision number of that model. See [Table 2, “FPGA Segment” on page 14](#) for the processor stepping ID number in the CPUID information.

When EAX is set to a value of one, the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number, and Stepping ID in the EAX register. Note that after reset, the EDX processor signature value equals 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.

Table 1. Server Segment (Sheet 1 of 2)

Physical Chop	Stepping	Segment Wayness	CPUID	CAPID0 (Segment)			CAPID0 (Wayness)		CAPID4 (Chop)		
				B:1, D:30, F:3, O:84						B:1, D:30 F:3, O:94	
				5	4	3	1	0	7	6	
XCC	B-0	Server, 2S	0x50652	1	1	1	0	1	1	1	
	B-0	Server, 4S	0x50652	1	1	1	1	0	1	1	
	B-0	Server, 8S	0x50652	1	1	1	1	1	1	1	
	H-0	Server, 2S	0x50654	1	1	1	0	1	1	1	
	H-0	Server, 4S	0x50654	1	1	1	1	0	1	1	
	H-0	Server, 8S	0x50654	1	1	1	1	1	1	1	



Table 1. Server Segment (Sheet 2 of 2)

Physical Chop	Stepping	Segment Wayness	CPUID	CAPID0 (Segment)			CAPID0 (Wayness)		CAPID4 (Chop)		
				B:1, D:30, F:3, O:84						B:1, D:30 F:3, O:94	
				5	4	3	1	0	7	6	
HCC	L-0	Server, 2S	0x50652	1	1	1	0	1	1	0	
	M-0	Server, 2S	0x50654	1	1	1	0	1	1	0	
LCC	U-0	Server, 2S	0x50654	1	1	1	0	1	0	0	

Table 2. FPGA Segment

- FPGA segment identified via bits 5:3=[011] of CAPID0

Physical Chop	Stepping	Segment Wayness	CPUID	CAPID0 (Segment)			CAPID0 (Wayness)		CAPID4 (Chop)		
				B:1, D:30, F:3, O:84						B:1, D:30 F:3, O:94	
				5	4	3	1	0	7	6	
XCC	B-0	FPGA, 2S	0x50652	0	1	1	0	1	1	1	
	H-0	FPGA, 2S	0x50654	0	1	1	0	1	1	1	

Table 3. Fabric Segment

- Fabric segment identified via bits 5:3=[001] of CAPID0

Physical Chop	Stepping	Segment Wayness	CPUID	CAPID0 (Segment)			CAPID0 (Wayness)		CAPID4 (Chop)		
				B:1, D:30, F:3, O:84						B:1, D:30 F:3, O:94	
				5	4	3	1	0	7	6	
XCC	B-0	Fabric, 2S	0x50652	0	0	1	0	1	1	1	
	H-0	Fabric, 2S	0x50654	0	0	1	0	1	1	1	



- The 8156 (not listed previously) has identical frequencies to 5122 but adds third Intel® UPI and 8-socket capability.
- The 8158 (not listed previously) has identical frequencies to 6136 but adds 8-socket capability.

Figure 3. Intel® Xeon® Processor Scalable Family Intel® AVX-512 Turbo Frequencies

81xx, 61xx, and 51xx processors for highest per-core performance.

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8180	28	38.50	205	1.7	3.5	3.5	3.3	3.3	3.2	3.2	3.2	3.2	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.3	2.3	2.3		
8168	24	33.00	205	1.9	3.5	3.5	3.3	3.3	3.2	3.2	3.2	3.2	3.2	3.2	3.2	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5						
6148	20	27.50	150	1.6	3.5	3.5	3.3	3.3	3.1	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.2	2.2	2.2										
6154	18	24.75	200	2.1	3.5	3.5	3.3	3.3	3.2	3.2	3.2	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.7	2.7												
6150	18	24.75	165	1.9	3.5	3.5	3.3	3.3	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.5	2.5												
6142	16	22.00	150	1.6	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.2	2.2	2.2														
6132	14	19.25	140	1.7	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.3	2.3															
6146	12	24.75	165	2.1	3.5	3.5	3.3	3.3	3.1	3.1	3.1	3.1	2.7	2.7	2.7	2.7																
6136	12	24.75	150	2.1	3.5	3.5	3.3	3.3	3.1	3.1	3.1	3.1	2.7	2.7	2.7	2.7																
6126	12	19.25	125	1.7	3.5	3.5	3.3	3.3	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3																
6144	8	24.75	150	2.2	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8																				
6134	8	24.75	130	2.1	3.5	3.5	3.3	3.3	2.7	2.7	2.7	2.7																				
6128	6	19.25	115	2.3	3.5	3.5	3.3	3.3	2.9	2.9																						
5122	4	16.50	105	2.7	3.5	3.5	3.3	3.3																								

- The 8180, 6142, and 6134 have 1.5 TB/socket memory capacity versions (8180M, 6142M, and 6134M – not listed previously) with identical frequencies.
- The 8156 (not listed previously) has identical frequencies to 5122 but adds third Intel® UPI and 8-socket capability.
- The 8158 (not listed previously) has identical frequencies to 6136 but adds 8-socket capability.

Figure 4. Intel® Xeon® Processor Scalable Family Non Intel® AVX Turbo Frequencies

81xx and 61xx processors

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8176	28	38.50	165	2.1	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9	2.8	2.8	2.8			
8170	26	35.75	165	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8	2.8	2.8	2.8			
8164	26	35.75	150	2.0	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.7	2.7	2.7	2.7			
8160	24	33.00	150	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8							
6152	22	30.25	140	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	2.9	2.9	2.9	2.9	2.8	2.8									
6138	20	27.50	125	2.0	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.2	3.2	3.2	2.9	2.9	2.9	2.7	2.7	2.7	2.7										
6140	18	24.75	140	2.3	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.0	3.0													
8153	16	22.00	125	2.0	2.8	2.8	2.6	2.6	2.5	2.5	2.5	2.5	2.5	2.5	2.5	2.3	2.3	2.3	2.3													
6130	16	22.00	125	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.1	3.1	3.1	2.8	2.8	2.8	2.8													

- The 8176, 8170, 8160, and 6140 have 1.5 TB/socket memory capacity versions (8180M, 8170M, 8160M, and 6140M – not listed previously) with identical frequencies.




Figure 5. Intel® Xeon® Processor Scalable Family Intel® AVX 2.0 Turbo Frequencies

81xx and 61xx, processors.

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																												
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	
8176	28	38.50	165	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.4	2.4	2.4	2.4
8170	26	35.75	165	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.2	3.2	3.2	3.2	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4	
8164	26	35.75	150	1.6	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.3	2.3	2.3	2.3	2.3	2.3	2.3		
8160	24	33.00	150	1.8	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5	2.5	2.5	2.5		
6152	22	30.25	140	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.4	2.4	2.4	2.4	2.4	2.4							
6138	20	27.50	125	1.6	3.6	3.6	3.4	3.4	3.2	3.2	3.2	3.2	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.3	2.3	2.3	2.3									
6140	18	24.75	140	1.9	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.6	2.6												
8153	16	22.00	125	1.6	2.7	2.7	2.5	2.5	2.4	2.4	2.4	2.4	2.2	2.2	2.2	2.2	2.0	2.0	2.0	2.0													
6130	16	22.00	125	1.7	3.6	3.6	3.4	3.4	3.1	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4													

- The 8176, 8170, 8160, and 6140 have 1.5 TB/socket memory capacity versions (8180M, 8170M, 8160M, and 6140M – not listed previously) with identical frequencies.

Figure 6. Intel® Xeon® Processor Scalable Family Intel® AVX-512 Turbo Frequencies

81xx and 61xx processors.

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8176	28	38.50	165	1.3	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9
8170	26	35.75	165	1.3	3.5	3.5	3.3	3.3	2.9	2.9	2.9	2.9	2.5	2.5	2.5	2.5	2.2	2.2	2.2	2.2	2.1	2.1	2.1	2.1	1.9	1.9	1.9	1.9	1.9	1.9	1.9	
8164	26	35.75	150	1.2	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	1.9	1.9	1.9	1.9	1.8	1.8	1.8	1.8	1.8	1.8	1.8	
8160	24	33.00	150	1.4	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0				
6152	22	30.25	140	1.4	3.5	3.5	3.3	3.3	2.9	2.9	2.9	2.9	2.5	2.5	2.5	2.5	2.2	2.2	2.2	2.2	2.0	2.0	2.0	2.0	2.0	2.0						
6138	20	27.50	125	1.3	3.5	3.5	3.3	3.3	2.7	2.7	2.7	2.7	2.3	2.3	2.3	2.3	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9								
6140	18	24.75	140	1.5	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	2.1	2.1	2.1									
8153	16	22.00	125	1.2	2.6	2.6	2.4	2.4	2.0	2.0	2.0	2.0	1.7	1.7	1.7	1.7	1.6	1.6	1.6	1.6												
6130	16	22.00	125	1.3	3.5	3.5	3.1	3.1	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	1.9	1.9	1.9	1.9												

- The 8176, 8170, 8160, and 6140 have 1.5 TB/socket memory capacity versions (8180M, 8170M, 8160M, and 6140M – not listed previously) with identical frequencies.



Figure 7. Intel® Xeon® Processor Scalable Family Non Intel® AVX, Intel® AVX 2.0, and Intel® AVX-512 Turbo Frequencies

81xxT and 61xxT processors - 10-year use plus NEBS-friendly thermal specification

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8160T	24	33.00	150	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8							
6138T	20	27.50	125	2.0	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.7	2.7	2.7										
6130T	16	22.00	125	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.1	3.1	3.1	2.8	2.8	2.8														
6126T	12	19.25	125	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.3	3.3	3.3																	

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8160T	24	33.00	150	1.8	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5						
6138T	20	27.50	125	1.5	3.6	3.6	3.4	3.4	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.2	2.2	2.2										
6130T	16	22.00	125	1.7	3.6	3.6	3.4	3.4	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.4	2.4	2.4														
6126T	12	19.25	125	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	2.9	2.9	2.9																		

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																													
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28		
8160T	24	33.00	150	1.4	3.5	3.5	3.3	3.3	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0							
6138T	20	27.50	125	1.2	3.5	3.5	3.2	3.2	2.5	2.5	2.5	2.5	2.1	2.1	2.1	1.9	1.9	1.9	1.9	1.8	1.8	1.8												
6130T	16	22.00	125	1.3	3.5	3.5	3.1	3.1	2.4	2.4	2.4	2.1	2.1	2.1	2.1	1.9	1.9	1.9																
6126T	12	19.25	125	1.7	3.5	3.5	3.3	3.3	2.6	2.6	2.6	2.6	2.3	2.3	2.3																			

- The 6126T processor is optimized for highest per-core performance.




Figure 8. Intel® Xeon® Processor Scalable Family Non Intel® AVX and Intel® AVX 2.0 Turbo Frequencies

81xxF and 61xxF processors with integrated Intel® Omni-Path Fabric.

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8176F	28	38.50	173	2.1	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9	2.8	2.8	2.8	2.8
8160F	24	33.00	160	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.2	3.0	3.0	3.0	2.8	2.8	2.8						
6148F	20	27.50	160	2.4	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1									
6138F	20	27.50	135	2.0	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.7	2.7	2.7									
6142F	16	22.00	160	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3												
6130F	16	22.00	135	2.1	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.8												
6126F	12	19.25	135	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3																

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8176F	28	38.50	173	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.4	2.4	2.4	2.4
8160F	24	33.00	160	1.8	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5	2.5						
6148F	20	27.50	160	1.9	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.6									
6138F	20	27.50	135	1.6	3.6	3.6	3.4	3.4	3.2	3.2	3.2	3.2	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.3	2.3	2.3									
6142F	16	22.00	160	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.2	3.2	3.2	2.9	2.9	2.9	2.9													
6130F	16	22.00	135	1.7	3.6	3.6	3.4	3.4	3.1	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.4	2.4	2.4													
6126F	12	19.25	135	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	2.9	2.9	2.9	2.9																

The 6148F, 6142F, and 6126F processors are optimized for highest per-core performance.

Figure 9. Intel® Xeon® Processor Scalable Family Intel® AVX-512 Turbo Frequencies

81xxF and 61xxF processors with integrated Intel® Omni-Path Fabric (Intel® OP Fabric).

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
8176F	28	38.50	173	1.3	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9
8160F	24	33.00	160	1.4	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0					
6148F	20	27.50	160	1.6	3.5	3.5	3.3	3.3	3.1	3.1	3.1	3.1	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.2	2.2	2.2	2.2								
6138F	20	27.50	135	1.3	3.5	3.5	3.3	3.3	2.7	2.7	2.7	2.7	2.3	2.3	2.3	2.3	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9								
6142F	16	22.00	160	1.6	3.5	3.5	3.3	3.3	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.2	2.2	2.2	2.2												
6130F	16	22.00	135	1.3	3.5	3.5	3.1	3.1	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	1.9	1.9	1.9													
6126F	12	19.25	135	1.7	3.5	3.5	3.3	3.3	2.6	2.6	2.6	2.6	2.3	2.3	2.3																	

The 6148F, 6142F, and 6126F processors are optimized for highest per-core performance.




Figure 10. Intel® Xeon® Processor Scalable Family Non Intel® AVX Turbo Frequencies

51xx, 41xx, and 31xx processors.

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120	14	19.25	105	2.2	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.6	2.6														
5118	12	16.50	105	2.3	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.7	2.7	2.7																	
5115	10	13.75	85	2.4	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.8	2.8																		
4116	12	16.50	85	2.1	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.4	2.4	2.4																	
4114	10	13.75	85	2.2	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.5	2.5																		
4112	4	8.25	85	2.6	3.0	3.0	2.9	2.9																								
4110	8	11.00	85	2.1	3.0	3.0	2.8	2.8	2.4	2.4	2.4	2.4																				
4108	8	11.00	85	1.8	3.0	3.0	2.7	2.7	2.1	2.1	2.1	2.1																				
3106	8	11.00	85	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.7																				
3104	6	8.25	85	1.7	1.7	1.7	1.7	1.7	1.7																							

Figure 11. Intel® Xeon® Processor Scalable Family Intel® AVX 2.0 Turbo Frequencies

51xx, 41xx, and 31xx processors.

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120	14	19.25	105	1.8	3.1	3.1	2.9	2.9	2.7	2.7	2.7	2.7	2.3	2.3	2.3	2.3	2.2	2.2														
5118	12	16.50	105	1.9	3.1	3.1	2.9	2.9	2.6	2.6	2.6	2.6	2.3	2.3	2.3																	
5115	10	13.75	85	2.0	3.1	3.1	2.9	2.9	2.6	2.6	2.6	2.6	2.4	2.4																		
4116	12	16.50	85	1.7	2.9	2.9	2.7	2.7	2.4	2.4	2.4	2.4	2.1	2.1																		
4114	10	13.75	85	1.8	2.9	2.9	2.7	2.7	2.3	2.3	2.3	2.3	2.2	2.2																		
4112	4	8.25	85	2.2	2.9	2.9	2.6	2.6																								
4110	8	11.00	85	1.7	2.9	2.9	2.7	2.7	2.1	2.1	2.1	2.1																				
4108	8	11.00	85	1.4	2.9	2.9	2.3	2.3	1.8	1.8	1.8	1.8																				
3106	8	11.00	85	1.3	1.3	1.3	1.3	1.3	1.3	1.3	1.3	1.3																				
3104	6	8.25	85	1.3	1.3	1.3	1.3	1.3	1.3																							

Figure 12. Intel® Xeon® Processor Scalable Family Intel® AVX-512 Turbo Frequencies

51xx, 41xx, and 31xx processors.

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120	14	19.25	105	1.2	2.9	2.9	2.5	2.5	1.9	1.9	1.9	1.9	1.6	1.6	1.6	1.6	1.6															
5118	12	16.50	105	1.2	2.9	2.9	2.4	2.4	1.8	1.8	1.8	1.8	1.6	1.6	1.6																	
5115	10	13.75	85	1.2	2.9	2.9	2.2	2.2	1.7	1.7	1.7	1.7	1.6	1.6																		
4116	12	16.50	85	1.1	1.8	1.8	1.6	1.6	1.5	1.5	1.5	1.5	1.4	1.4	1.4																	
4114	10	13.75	85	1.1	1.8	1.8	1.6	1.6	1.5	1.5	1.5	1.5	1.4	1.4																		
4112	4	8.25	85	1.1	1.8	1.8	1.4	1.4																								
4110	8	11.00	85	1.0	1.8	1.8	1.6	1.6	1.3	1.3	1.3	1.3																				
4108	8	11.00	85	0.9	1.8	1.8	1.5	1.5	1.2	1.2	1.2	1.2																				
3106	8	11.00	85	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8																				
3104	6	8.25	85	0.8	0.8	0.8	0.8	0.8	0.8																							



Figure 13. Intel® Xeon® Processor Scalable Family Non Intel® AVX, Intel® AVX 2.0, and Intel® AVX-512 Turbo Frequencies

51xxT and 41xxT processors - 10-year use plus NEBS-friendly thermal specification.

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120T	14	19.25	105	2.2	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.6	2.6														
5119T	14	19.25	85	1.9	3.2	3.2	3.0	3.0	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.3	2.3														
4116T	12	16.50	85	2.1	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.4	2.4	2.4	2.4																
4114T	10	13.75	85	2.2	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.5	2.5																		
4109T	8	11.00	70	2.0	3.0	3.0	2.8	2.8	2.3	2.3	2.3	2.3																				

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120T	14	19.25	105	1.8	3.1	3.1	2.9	2.9	2.7	2.7	2.7	2.7	2.3	2.3	2.3	2.3	2.2	2.2														
5119T	14	19.25	85	1.5	3.1	3.1	2.9	2.9	2.3	2.3	2.3	2.3	2.0	2.0	2.0	2.0	1.9	1.9														
4116T	12	16.50	85	1.7	2.9	2.9	2.7	2.7	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1																
4114T	10	13.75	85	1.8	2.9	2.9	2.7	2.7	2.3	2.3	2.3	2.3	2.2	2.2																		
4109T	8	11.00	70	1.6	2.9	2.9	2.6	2.6	2.0	2.0	2.0	2.0																				

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5120T	14	19.25	105	1.2	2.9	2.9	2.5	2.5	1.9	1.9	1.9	1.9	1.6	1.6	1.6	1.6	1.6	1.6														
5119T	14	19.25	85	1.0	2.9	2.9	2.2	2.2	1.7	1.7	1.7	1.7	1.4	1.4	1.4	1.4	1.4	1.4														
4116T	12	16.50	85	1.1	1.8	1.8	1.6	1.6	1.5	1.5	1.5	1.5	1.4	1.4	1.4	1.4																
4114T	10	13.75	85	1.1	1.8	1.8	1.6	1.6	1.5	1.5	1.5	1.5	1.4	1.4																		
4109T	8	11.00	70	1.0	1.8	1.8	1.6	1.6	1.3	1.3	1.3	1.3																				

Figure 14. Intel® Xeon® Processor Scalable Family Non Intel® AVX, Intel® AVX 2.0, and Intel® AVX-512 Turbo Frequencies

51xx7 off-roadmap processor.

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5117	14	19.25	105	2.0	2.8	2.8	2.6	2.6	2.5	2.5	2.5	2.5	2.4	2.4	2.4	2.4	2.3	2.3														

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5117	14	19.25	105	1.3	2.8	2.8	2.5	2.5	1.9	1.9	1.9	1.9	1.6	1.6	1.6	1.6	1.6	1.6														

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5117	14	19.25	105	1.1	2.8	2.8	2.2	2.2	1.7	1.7	1.7	1.7	1.4	1.4	1.4	1.4	1.4	1.4														




Figure 15. Intel® Xeon® Processor Scalable Family Non Intel® AVX, Intel® AVX 2.0, and Intel® AVX-512 Turbo Frequencies

51xx7 off-roadmap processor with integrated Intel® Omni-Path Fabric.

					# of active cores / maximum core frequency in turbo mode (GHz)																												
SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	
5117F	14	19.25	113	2.0	2.8	2.8	2.6	2.6	2.5	2.5	2.5	2.5	2.4	2.4	2.4	2.4	2.3	2.3															

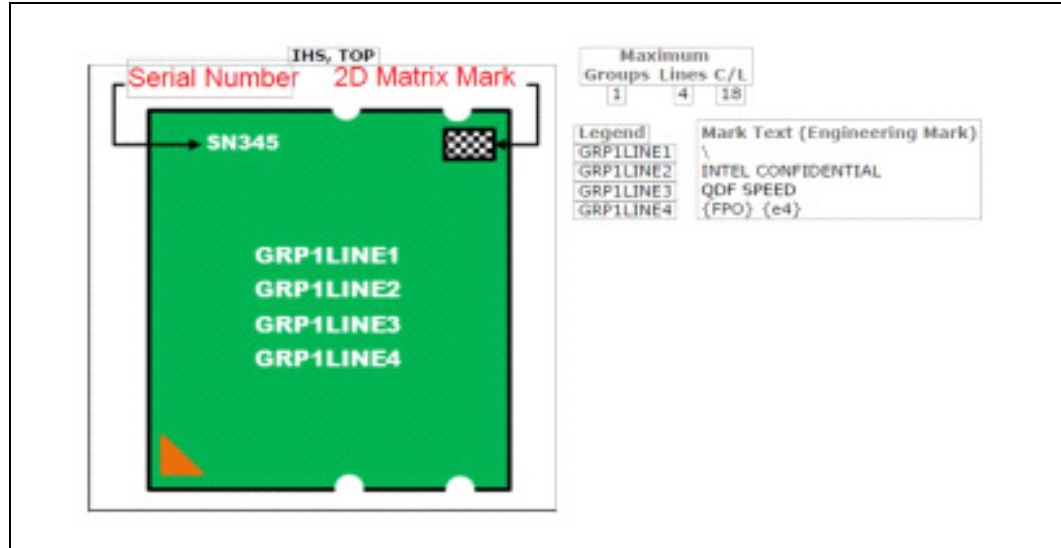
					# of active cores / maximum core frequency in turbo mode (GHz)																												
SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Frequency (GHz)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	
5117F	14	19.25	113	1.3	2.8	2.8	2.5	2.5	1.9	1.9	1.9	1.9	1.6	1.6	1.6	1.6	1.6	1.6															

					# of active cores / maximum core frequency in turbo mode (GHz)																											
SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
5117F	14	19.25	113	1.1	2.8	2.8	2.2	2.2	1.7	1.7	1.7	1.7	1.4	1.4	1.4	1.4	1.4	1.4														



Component Marking Information

Figure 16. Processor Preliminary Top Side Marking (Example)



For the Intel® Xeon® Processor Scalable Family SKUs, see <https://ark.intel.com/content/www/us/en/ark/products/series/125191/intel-xeon-scalable-processors.html>



Errata

SKX1. A CAP Error While Entering Package C6 Might Cause DRAM to Fail to Enter Self-Refresh (Intel® Xeon® Processor Scalable Family)

Problem: A Command/Address Parity (CAP) error that occurs on the command to direct DRAM to enter self-refresh might cause the DRAM to fail to enter self-refresh although the processor enters Package-C6

Implication: Due to this erratum, DRAM might fail to be refreshed, which might result in uncorrected errors being reported from the DRAM.

Workaround: None.

Status: No Fix.

SKX2. PCIe* Lane Error Status Register Might Log False Correctable Error (Intel® Xeon® Processor SP)

Problem: Due to this erratum, PCIe* LNERSTS (Device 0; Function 0; Offset 258h; bits [3:0]) might log false lane-based correctable errors.

Implication: Diagnostics cannot reliably use LNERSTS to report correctable errors.

Workaround: None.

Status: No Fix.

SKX3. In Memory Mirror Mode, DataErrorChunk Field Might be Incorrect (Intel® Xeon® Processor SP)

Problem: In Memory Mirror Mode, DataErrorChunk bits (IA32_MC7_MISC register MSR(41FH) bits [61:60]) might not correctly report the chunk containing an error.

Implication: Due to this erratum, this field is not accurate when Memory Mirror Mode is enabled.

Workaround: None.

Status: No Fix.

SKX4. Intel® Resource Director Technology (Intel® RDT) MBM Does Not Accurately Track Write Bandwidth (Intel® Xeon® Processor SP)

Problem: Intel® Resource Director Technology (RDT) Memory Bandwidth Monitoring (MBM) does not count cacheable write-back traffic to local memory. This results in the RDT MBM feature under counting total bandwidth consumed.

Implication: Applications using this feature might report incorrect memory bandwidth.

Workaround: None.

Status: No Fix.


SKX5. Intel® UPI Initialization Aborts Might be Logged (Intel® Xeon® Processor SP)

Problem: If Intel® Ultra Path Interconnect (Intel® UPI) is configured for slow mode operation, initialization aborts might occur.

Implication: Unexpected initialization aborts might be logged in the ktireut_ph_ctr1 register (Bus: 3; Device: 16-14; Function 1; Offset 12h; Bit 4).

Workaround: None.

Status: No Fix.



SKX6. PCIe* Port Might Incorrectly Log Malformed_TLP Error (Intel® Xeon® Processor SP)

Problem: If the PCIe* port receives a TLP that triggers both a Malformed_TLP error and an ECRC_TLP error, the processor should only log an ECRC_TLP error. However, the processor logs both errors.

Implication: Due to this erratum, the processor may incorrectly log Malformed_TLP errors.

Workaround: None.

Status: No Fix.

SKX7. CMCI Might Not be Signaled for Corrected Error (Intel® Xeon® Processor Scalable Family)

Problem: Machine check banks 9, 10, and 11 might not signal Corrected Machine Check Interrupt (CMCI) after the first corrected error is reported in the bank even if the MCI_STATUS register has been cleared.

Implication: After the first corrected error is reported in one of the affected machine check banks, subsequent errors will be logged but may not result in a CMCI.

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

Status: No Fix.

SKX8. Intel® CAT/CDP Might Not Restrict Cacheline Allocation Under Certain Conditions (Intel® Xeon® Processor Scalable Family)

Problem: Under certain microarchitectural conditions involving heavy memory traffic, cache lines might fill outside the allocated L3 capacity bitmask (CBM) associated with the current Class of Service (CLOS).

Implication: Cache Allocation Technology/Code and Data Prioritization (CAT/CDP) might see performance side effects and a reduction in the effectiveness of the CAT feature for certain classes of applications, including cache-sensitive workloads than seen on previous platforms.

Workaround: None identified. Contact your Intel representative for details of possible mitigations. None identified.

Status: No Fix.

SKX9. Credits Not Returned For PCIe* Packets That Fail ECRC Check Problem (Intel® Xeon® Processor SP)

Problem: The processor's IIO does not return credits back to the PCIe* link in case of end-to-end CRC (ECRC) errors.

Implication: Due to this erratum, the link might experience degraded performance or might eventually fail due to a loss of credits.

Workaround: For processors that support Live Error Recovery (LER) the link would be reset and credits would be restored. Processors that do not support LER should configure ECRC errors to be fatal.

Status: No Fix.

SKX10. PCIe* Link Might Fail to Train (Intel® Xeon® Processor SP)

Problem: When a pin on a PCIe* lane is not connected to the link partner, the PCIe* port's LTSSM might hang in the detect state.

Implication: When this erratum occurs, the PCIe* link fails to train and the corresponding link partner(s) are not enumerated.

Workaround: None.

Status: No Fix.

**SKX11. Intel® UPI CRC32 Rolling Mode is Not Functional (Intel® Xeon® Processor SP)**

Problem: With UPI CRC32 Rolling Mode enabled, UPI Rx CRC errors might be seen.

Implication: Due to this erratum, when UPI CRC32 Rolling Mode is enabled, UPI Rx CRC errors might be seen.

Workaround: None. Do not enable UPI CRC32 setting in BIOS.

Status: No Fix

SKX12. IODC Entry 0 Cannot be Masked (Intel® Xeon® Processor SP)

Problem: The individual I/O Directory Cache (IODC) Entry 0 cannot be masked using HA_COH_CFG_1, (Bus 1; Devices 11-8; Functions 7-0, Offset 0x11C, bit 0) therefore Entry 0 is always allocated.

Implication: No functional implications.

Workaround: None.

Status: No Fix.

SKX13. With eMCA2 Enabled a 3-Strike Might Cause an Unnecessary CATERR# Instead of Only MSMI (Intel® Xeon® Processor SP)

Problem: When eMCA2 is enabled to cause an MSMI due to a 3-strike event, a pulsed CATERR# and MSMI# event might both be observed on the pins.

Implication: When this erratum occurs, an unnecessary CATERR# pulse might be observed.

Workaround: None.

Status: No Fix.

SKX14. CMCI May Not be Signaled for Corrected Error (Intel® Xeon® Processor Scalable Family)

Problem: Machine check banks 9, 10, and 11 might not signal CMCI after the first corrected error is reported in the bank even if the MCI_STATUS register has been cleared.

Implication: After the first corrected error is reported in one of the affected machine check banks, subsequent errors are logged but might not result in a CMCI.

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

Status: No Fix.

SKX15. CSRs SVID and SDID Are Not Implemented For Some DDRIO and PCU Devices (Intel® Xeon® Processor SP)

Problem: The DDRIO (Bus: 3; Device {19,22}; Function {6,7} and Bus: 0; Device: {20,23}; Function: {4,5,6,7};) and PCU (Bus: 3; Device 31; Functions {0,2}) do not implement the SVID (Offset 0x2C) and SDID (Offset 0x2E) CSRs. Read accesses to these register locations return all zeros.


Implication: Software relying on DDRIO and PCU SVID and SDID CSR support might not function correctly.

Workaround: None identified. Do not use SVID and SDID for these devices and functions.

Status: No Fix.

SKX16. Register Broadcast Read From DDRIO May Return a Zero Value (Intel® Xeon® Processor Scalable Family)

Problem: When performing a BIOS broadcast register read to DDRIO a value of zero is always returned.



Implication: When this erratum occurs, BIOS might not be able to proceed due to always reading a value of zero.

Workaround: None. Use unicast register read for each instance instead of broadcast register read for all instances at once.

Status: No Fix.

SKX17. Intel® CMT Counters May Not Count Accurately (Intel® Xeon® Processor SP)

Problem: Under complex micro-architectural conditions, the Cache Monitoring Technology (CMT) counters might over-count.

Implication: Software relying on CMT registers to enable resource allocation might not operate correctly. This can lead to reporting of more cachelines used than the cache supports or the counter wrapping and returning a too small value. WBINVD might not result in the CMT counters being zeroed. Intel has not observed this erratum in commercially available software.

Workaround: None.

Status: No Fix.

SKX18. Intel® CAT Might Not Restrict Cacheline Allocation Under Certain Conditions (Intel® Xeon® Processor Scalable Family)

Problem: Under certain micro-architectural conditions involving heavy memory traffic, cachelines might fill outside the allocated L3 capacity bit-mask (CBM) associated with the current Class of Service (CLOS).

Implication: CAT might appear less effective at protecting certain classes of applications, including cache-sensitive workloads than on previous platforms.

Workaround: None identified. Contact your Intel representative for details of possible mitigations.

Status: No Fix.

SKX19. Intel® PCIe* Corrected Error Threshold Does Not Consider Overflow Count When Incrementing Error Counter (Intel® Xeon® Processor SP)

Problem: The PCIe* corrected error counter feature does not take the overflow bit in the count (bit 15 of XPCORERRCOUNTER (Bus; RootBus Device; 0 Function; 0 Offset; 4D0h)) into account when comparing the count to the threshold in XPCORERRTHRESHOLD.ERROR_THRESHOLD. Therefore, users end up with another interrupt once the counter has rolled over and hit the threshold + 0x8000.

Implication: Due to this erratum, the PCIe* corrected error signaling might occur even after the error count has exceeded the corrected error count threshold, not just a single time when reaching the threshold. Intel has not observed this erratum with any commercially available system.

Workaround: None identified.

Status: No Fix.

SKX20. IIO RAS VPP Hangs During The Warm Reset Test (Intel® Xeon® Processor SP)

Problem: When VPPCL bit 0 of VPP_reset_Mode (Bus 1; Device 30; Function 5; Offset 0xF0) bit is set to 0, and the CPU is undergoing reset flow while PCIe* hot-plug operation is in process, the Virtual Pin Port (VPP) hot-plug commands might stop responding.

Implication: Due to this erratum, during CPU reset hot-plug commands might not complete.

Workaround: None. Do not set VPP reset mode to zero.

Status: No Fix.

**SKX21. Intel® UPI CRC Errors and PHY Init Aborts May Be Seen During UPI Slow Mode Training**

Problem: During a normal cold boot or cold reset, UPI CRC errors and PHY init aborts may be seen due to a random miscalculation of UPI lane skewing during training

Implication: Intel® UPI CRC errors and PHY init aborts may be seen during boot or reset

Workaround: PLR3 contains a workaround for this issue. Details can be found in the BIOS release notes.

Status: No Fix

SKX22. A Core 3-Strike Event May Be Seen Under Certain Test Conditions

Problem: When running some stress tests and/or related applications, a core 3-strike event may be seen. This similar 3-strike event may also occur when system is at idle.

Implication: A core 3-strike event may be seen resulting in a system hang and/or a shutdown.

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

Status: No fix.

SKX23. DDR4 Memory Bandwidth May Be Lower Than Expected at 2133 and 1866 Speeds

Problem: In some DDR4 memory configurations running 2133 or 1866, lower than expected memory bandwidth may be seen. When running at these speeds, there may also be a possibility of seeing socket-to-socket variation in performance as well.

Implication: DDR4 Memory Bandwidth may be lower than expected at 2133 and 1866 speeds.

Workaround: Intel® Xeon® processor scalable family-based platform BIOS 132R08 contains a workaround for this issue.

Status: No Fix.

SKX24. Lower Than Expected Performance May Be Seen With Certain Intel® AVX-512 Workloads

Problem: Under certain Intel® AVX-512 workloads, the Uncore frequency may not scale with Core frequency as expected.

Implication: Lower than expected performance may be seen under with some Intel® AVX-512 workloads.

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

Status: No fix.

SKX25. A System Hang May Be Seen With Some 8S + XNC Type Platform Configurations

Problem: A KTI write back credit starvation event may occur in some 8S + XNC platform configurations leading to a CHA deadlock. This may eventually cause a system hang.

Implication: A system hang may occur in some 8S + XNC platform configurations.

Workaround: A future BIOS workaround is in development.


Status: No Fix.

SKX26. Sparing Per-Rank Error Masking Does Not Mask Correctable Errors

Problem: The IMC (Integrated Memory Controller) Sparing PREM (Per-Rank Error Masking) capability does not mask off correctable error logging and signaling as expected.

Implication: Due to this erratum, errors will continue to be logged and signaled despite per-rank error masking. Per-rank error counters are still masked.

Workaround: None Identified



Status: No Fix

SKX27. PCIe* Root Port Electromechanical Interlock Control Register Can Be Written

Problem: Electromechanical Interlock Control (bit 11) in the Slot Control register (B: Root Port; D: 0-3; F: 0 bits offset 0x18) in the PCIe* Capability table should be read-only and always return 0. Due to this erratum, this register can be written.

Implication: Writes to this bit can cause later reads to return the written value. However, this has no other effect on functionality.

Workaround: None Identified.

Status: No Fix

SKX28. Live Error Recovery Feature Being Disabled is not Getting Reflected in PXP2CAP Value

Problem: When Live Error Recovery (LER) feature is disabled, the LER capability register still remains in the PCIe* extended header space and is linked to pxp2cap. This register will indicate that LER feature is available when it is not.

Implication: Due to this erratum, Intel® Xeon® (SP) 4100 series and 3100 series CPU SKUs with standard RAS features that have LER disabled may not correctly indicate the status of this feature to software which may indicate the LER capability still exists. Software may incorrectly assume that uncorrectable errors will be downgraded to correctable errors.

Workaround: None Identified.

Status: No Fix

SKX29. Performance Monitoring M2MEM Counters For Memory Controller Reads/Writes Are Not Counting Read/Write Retries

Problem: PMON M2MEM counters for read and write events do not account for scrub reads and scrub writes during the error flow.

Implication: Due to this erratum, a mismatch in the counters for Read/Write retries in M2MEM and iMC (integrated memory controller) may be observed.

Workaround: When doing error injection testing, counting reads and writes in the presence of ECC errors will only be precise using the iMC counter, not the M2MEM counter

Status: No Fix

SKX30. System Hangs May Occur When IPQ and IRQ Requests Happen at The Same Time

Problem: When IPQ and IRQ requests happen at the same time, and the IPQ request is starved due to PAMatch/NotAllowSnoop on a TORID (Table of Request ID) then the IRQ request that is waiting for the TORID's SF/LLC may become invalid.

Implication: Due to this erratum, if IPQ and IRQ requests do not need to snoop any cores, then IPQ requests may block IRQ requests resulting in a system hang. Intel® has only observed this erratum in a synthetic test environment.

Workaround: None identified.

Status: No Fix

SKX31. Two Intel® UPI Reads From XNC May Lead to a System Hang

Problem: If Intel® UPI non-snoop reads are targeted to the prefetchable memory region, then two outstanding reads to the same system address can merge into the same prefetch request.

Implication: Due to this erratum, an eXternal Node Controller (XNC) issuing non-snoop reads to the prefetchable memory region may result in one of the read's completions being dropped leading to a system hang.



Workaround: XNCs should not target the prefetchable memory region with UPI non-snoop reads.

Status: No Fix

SKX32. IIO VPP May Hang During Warm Reset

Problem: When VPP_Reset_Mode bit 0 of VPPCTL (Bus 1; Device 30; Function 5; Offset 0xF0) is set to 0, and there is a PCIe* hot-plug event in progress, if the processor performs a warm reset, the Virtual Pin Port hot-plug flow may hang.

Implication: Due to this erratum, the Virtual Pin Port may hang.

Workaround: Do not set VPP_Reset_Mode to 0.

Status: No Fix.

SKX33. Machine Check Events may be logged in banks 9, 10 and 11 that do not represent actual errors

Problem: In some previous CPU Microcode + BIOS code combinations MCEs in banks 9, 10 and 11 may be seen. These do not represent actual errors and normally are processed out by early BIOS execution.

Implication: MCEs may be seen on banks 9, 10 and 11 that represent incorrect error data. These MCEs have the potential to be forwarded to the OS and may be end-user visible while not representing actual errors.

Workaround: Contact your Intel representative for additional information regarding this issue.

Status: No Fix.

SKX34. Advanced RAS Dynamic Link Width Reduction may not resize the Intel® UPI link

Problem: The Advanced RAS Dynamic Link Width Reduction feature may not be properly detected and enabled prior to UPI initialization.

Implication: Due to this erratum, if there is a hard failure of an Intel® UPI lane at boot time, the Advanced RAS Dynamic Link Width Reduction feature may not function.

Workaround: None identified.


Status: No Fix.

SKX35. I

Problem: This erratum may cause a machine-check error (IA32_MCi_STATUS.MCACOD=005H with IA32_MCi_STATUS.MSCOD=00FH or IA32_MCi_STATUS.MCACOD=0150H with IA32_MCi_STATUS.MSCOD=00FH) on the fetch of an instruction. It applies only if (1) instruction bytes are fetched from a linear address translated using a 4-Kbyte page and cached in the processor; (2) the paging structures are later modified so that these bytes are translated using a large page (2-Mbyte, 4-Mbyte or 1-GByte) with a different physical address (PA), memory type (PWT, PCD and PAT bits), or User/Supervisor (U/S) bit; and (3) the same instruction is fetched 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 with MSCOD 00FH may occur, possibly resulting in a shutdown. This erratum could also lead to unexpected correctable machine check (IA32_MCi_STATUS.UC=0) with error code 005H with MSCOD 00FH.

Workaround: Software should not write to a paging-structure entry in a way that would change the page size and either the physical address, memory type or User/Supervisor bit. It can instead use one of the following algorithms: 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. An alternative algorithm: first change the physical page attributes (combination of physical address, memory type and User/Supervisor bit) in all 4K pages in the affected linear addresses; then invalidate any translations for the



affected linear addresses; and then modify the relevant paging-structure entry to establish the new page size.

Status: No Fix.

SKX36. Unexpected DDR ECC Errors May be Seen

Problem: The processor may incorrectly an incorrectly configured DDR VCCP value, which may lead to unexpected DDR ECC errors.

Implication: Due to this erratum, unexpected DDR4 ECC errors may occur

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

Status: No Fix.

SKX37. Spurious Corrected Errors May be Reported

Problem: Due to this erratum, spurious corrected errors may be logged in the IA32_MC0_STATUS MSR (401H) register with the valid field (bit 63) set, the uncorrected error field bit (bit 61) not set, a Model Specific Error Code (bits [31:16]) of 0x0001, and an MCA Error Code (bits [15:0]) of 0x0005. If CMCI is enabled, these spurious corrected errors also signal interrupts.

Implication: When this erratum occurs, software may see an unusually high rate of reported corrected errors. As it is not possible to distinguish between spurious and non-spurious errors, this erratum may interfere with reporting non-spurious corrected errors.

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

Status: No Fix.

SKX38. Dynamic Link Width Reduction May Not Resize the Intel® UPI Link

Problem: The Advanced RAS Dynamic Link Width Reduction feature may not be properly detected and enabled prior to UPI initialization.

Implication: If there is a hard failure of a UPI lane at boot time, then due to this erratum, the Advanced RAS Dynamic Link Width Reduction feature may not function, allowing the system to hang.

Workaround: None identified.

Status: No Fix.

SKX39. Writing to LT_LOCK_MEMORY and LT_UNLOCK_MEMORY MSRs Simultaneously May Have Inconsistent Results

Problem: Writing to LT_LOCK_MEMORY MSR (2e7H) and to LT_UNLOCK_MEMORY MSR (2e6H) simultaneously from different physical cores may have inconsistent results. Some of the memory ranges may get locked as requested by the write to LT_LOCK_MEMORY MSR while some may get unlocked as requested by the write to LT_UNLOCK_MEMORY MSR.

Implication: Writing to LT_LOCK_MEMORY MSR and to LT_UNLOCK_MEMORY MSRs may not operate as expected if they are done on different cores simultaneously. Intel has not observed this erratum in any commercially available system.

Workaround: None identified. Software (BIOS) should write to these MSRs only on the BSP (boot strap processor).

Status: No Fix.

SKX40. Masked Bytes in a Vector Masked Store Instructions May Cause Write Back of a Cache Line

Problem: Vector masked store instructions to WB (write-back) memory-type that cross cache lines may lead to CPU writing back cached data even for cache lines where all of the bytes are masked.



Implication: The processor may generate writes of un-modified data. This can affect Memory Mapped I/O (MMIO) or non-coherent agents in the following ways:

1. For MMIO range that is mapped as WB memory type, this erratum may lead to Machine Check Exception (MCE) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.
2. If the CPU cached data is stale, for example in the case of memory written directly by a non-coherent agent (agent that uses non-coherent writes), this erratum may lead to writing back stale cached data even if these bytes are masked.

Workaround: Platforms should not map MMIO memory space or non-coherent device memory space as WB memory. If WB is used for MMIO range, software or VMM should not map such MMIO page adjacent to a regular WB page (adjacent on the linear address space, before or after the I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

Status: No Fix.

SKX41. ERROR_N[2:0] Pins May Not be Cleared After a Warm Reset

Problem: The processor's ERROR_N[2:0] pins may not be cleared after a warm reset.

Implication: Due to this erratum, the ERROR_N[2:0] pins may incorrectly indicate a pending error after a warm reset.

Workaround: The BIOS can contain code changes to work around this erratum.

Status: No Fix.

SKX42. CRC Store Operation Corner Case May Result in Hang

Problem: Intel® QuickData Technology Local and Remote CRC Store operations may result in a DMA channel hang when the CRC Store transfer size is less than 32 bytes and the destination offset is not DWORD-aligned.

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

Workaround: Software must configure Intel® QuickData Technology Local and Remote CRC Store operations to have descriptor destination offset addresses DWORD-aligned.

Status: No Fix.

SKX43. Atomicity May Not be Preserved When Executing With RTM Enabled

Problem: In multi-socket platforms, in very rare situations, when a thread is executing an Restricted Transactional Memory (RTM) transaction, the processor may allow a different socket's thread to write to an address used by the RTM transaction, without causing the first thread to abort its transaction. This prevents the first thread's transaction from completing atomically.

Implication: Loss of atomicity may occur when using RTM.


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

Status: No Fix.

SKX44. Intel PCIe* Slot Presence Detect and Presence Detect Changed Logic Not PCIe* Specification Compliant

Problem: When Hot-Plug Surprise is set in the Slot Capabilities register (Bus: RootBus, Dev: 1-3, Function: 0, Offset: A4h, Bit: 5), the Presence Detect State and Presence Detect Change in the Slot Status register (Bus: RootBus, Dev: 1-3, Function: 0, Offset: A2h), incorrectly ignores the out-of-band presence detect mechanism and only reflects the Physical Layer in-band presence detect mechanism.

Implication: Due to this erratum, if the Hot-Plug Surprise bit is set in the Slot Capabilities register, software will not be able to detect the presence of an adapter inserted while a slot is



powered down. Therefore, Hot-Plug Surprise must only be set in configurations where the slot power is always enabled.

Workaround: None Identified.

Status: No Fix.

SKX45. In Patrol Scrub System Address Mode, Address is Not Loaded from CSRs After Re-enable

Problem: The patrol scrub starting address registers [scrubaddressshi (Bus 2; Devices 12, 10; Function 0; Offset 910) and scrubaddresslo Bus 2; Devices 12, 10; Function 0; Offset 90c] should indicate when the first memory address from which patrol logic should start scrubs [when scrubctl.startscrub (Bus 2; Devices 12, 10; Function 0; Offset 914; Bit 24) is set]. Due to this erratum, after patrol is disabled, if the patrol scrub engine is re-enabled in System Address Mode with scrubctl.startscrub set, the patrol scrubbing engine may ignore the starting address registers. Re-enabling patrol after S3 exit or other warm reset event is not impacted by this.

Implication: Due to this erratum, when configured in system address mode, Patrol scrubs will not start from the address specified in the starting address registers. This may cause certain memory lines to be scrubbed more or less frequently than expected. Intel has not seen this erratum to affect the operation of any commercially available software.

Workaround: None identified. Contact your Intel representative for details of possible mitigations.

Status: No Fix.

SKX46. Intel® Processor Trace (Intel® PT) TIP.PGD May Not Have Target IP Payload

Problem: When Intel® PT 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: No Fix.

SKX47. 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: No Fix.

SKX48. SMRAM State-Save Area Above the 4 GB 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 System-Management Mode (SMM) 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 4 GB address boundary.

Status: No Fix.

SKX49. 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: No Fix.

SKX50. Load Latency Performance Monitoring Facility May Stop Counting

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 (PEBS extension). However due to this erratum, load latency facility may stop counting load instructions when Intel® Hyper-Threading Technology (Intel® HT Technology) is enabled.

Implication: Counters programmed with the affected events stop incrementing and do not generate PEBS records.

Workaround: None identified.

Status: No Fix.

SKX51. Intel® PT PSB+ Packets May be Omitted on a C6 Transition

Problem: An Intel® PT (Processor Trace) PSB+ (Packet Stream Boundary+) set of packets may not be generated as expected when IA32_RTIT_STATUS.PacketByteCnt[48:32] (MSR 0x571) reaches the PSB threshold and a logical processor C6 entry occurs within the following one KByte of trace output.

Implication: After a logical processor enters C6, Intel® PT output may be missing PSB+ sets of packets.

Workaround: None identified.

Status: No Fix.

SKX52. Performance Monitoring Counters May Undercount When Using CPL Filtering

Problem: Performance Monitoring counters configured to count only OS or only USR events (by setting only one of bits 16 or 17 in IA32_PERFEVTSELx) may undercount for a short cycle period of typically less than 100 processor clock cycles after the processor transitions to a new CPL. Events affected may include those counting CPL transitions (by additionally setting the edge-detect bit 18 in IA32_PERFEVTSELx).


Implication: Due to this erratum, Performance Monitoring counters may report counts lower than expected.

Workaround: None identified.

Status: No Fix.

SKX53. Incorrect Branch Predicted Bit in BTS/BTM Branch Records

Problem: Branch Trace Store (BTS) and Branch Trace Message (BTM) 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: No Fix.

SKX54. 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: No Fix.

SKX55. 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: No Fix.

SKX56. 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: No Fix.

SKX57. 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: No Fix.

SKX58. CPUID TLB Associativity Information is Inaccurate

Problem: CPUID leaf 2 (EAX=02H) TLB information inaccurately reports that the shared second-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 second-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 second-Level TLB associativity information reported by CPUID for the affected processors.

Status: No Fix.

SKX59. Vector Masked Store Instructions May Cause Write Back of Cache Line Where Bytes Are Masked

Problem: Vector masked store instructions to write-back (WB) memory-type that cross cache lines may lead to CPU writing back cached data even for cache lines where all of the bytes are masked.

This can affect Memory Mapped I/O (MMIO) or non-coherent agents in the following ways:

1. For MMIO range that is mapped as WB memory type, this erratum may lead to Machine Check Exception (MCE) due to writing back data into the MMIO space. This applies only to cross page vector masked stores where one of the pages is in MMIO range.
2. If the CPU cached data is stale, for example in the case of memory written directly by a non-coherent agent (agent that uses non-coherent writes), this erratum may lead to writing back stale cached data even if these bytes are masked.

Implication: CPU may generate writes into MMIO space which lead to MCE, or may write stale data into memory also written by non-coherent agents.

Workaround: It is recommended not to map MMIO range as WB. If WB is used for MMIO range, OS or VMM should not map such MMIO page adjacent to a regular WB page (adjacent on the linear address space, before or after the I/O page). Memory that may be written by non-coherent agents should be separated by at least 64 bytes from regular memory used for other purposes (on the linear address space).

Status: No Fix.


SKX60. Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed

Problem: During Restricted Transactional Memory (RTM) operation when branch tracing is enabled using Branch Trace Message (BTM) or Branch Trace Store (BTS), 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: No Fix.

**SKX61. MOVNTDQA From WC Memory May Pass Earlier Locked Instructions**

Problem: An execution of (V)MOVNTDQA (streaming load instruction) that loads from Write Combining (WC) memory may appear to pass an earlier locked instruction to a different cache line.

Implication: Software that expects a lock to fence subsequent (V)MOVNTDQA instructions may not operate properly.

Workaround: Software should not rely on a locked instruction to fence subsequent executions of MOVNTDQA. Software should insert an MFENCE instruction if it needs to preserve order between streaming loads and other memory operations.

Status: No Fix.

SKX62. #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: No Fix.

SKX63. Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop

Problem: If an Intel® PT 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: No Fix.

SKX64. 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 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: No Fix.

SKX65. 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: No Fix.

SKX66. Performance Monitor Event For Outstanding Offcore 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: No Fix.

SKX67. Branch Instructions May Initialize MPX Bound Registers Incorrectly

Problem: Depending on the current Intel® Memory Protection Extensions (Intel® MPX) 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: No Fix.

SKX68. A Spurious APIC Timer Interrupt May Occur After Timed MWAIT

Problem: Due to this erratum, a Timed MWAIT that completes for a reason other than the Timestamp Counter reaching the target value may be followed by a spurious APIC timer interrupt. This erratum can occur only if the APIC timer is in TSC-deadline mode and only if the mask bit is clear in the LVT Timer Register.

Implication: Spurious APIC timer interrupts may occur when the APIC timer is in TSC-deadline mode.


Workaround: TSC-deadline timer interrupt service routines should detect and deal with spurious interrupts.

Status: No Fix.

SKX69. 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: No Fix.

SKX70. VM Entry That Clears TraceEn May Generate a FUP

Problem: If VM entry clears Intel® PT IA32_RTIT_CTL.TraceEn (MSR 570H, bit 0) while PacketEn is 1 then a Flow Update Packet (FUP) 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: No Fix.

SKX71. Reading Some C-state Residency MSRs May Result in Unpredictable System Behavior

Problem: Under complex microarchitectural conditions, an MSR read of MSR_CORE_C3_RESIDENCY MSR (3FCh), MSR_CORE_C6_RESIDENCY MSR (3FDh), or MSR_CORE_C7_RESIDENCY MSR (3FEh) may result in unpredictable system behavior.

Implication: Unexpected exceptions or other unpredictable system behavior may occur.

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

Status: No Fix

SKX72. Processor May Hang When Executing Code In an HLE Transaction Region

Problem: Under certain conditions, if the processor acquires an HLE (Hardware Lock Elision) lock via the XACQUIRE instruction in the Host Physical Address range between 40000000H and 403FFFFFFH, it may hang with an internal timeout error (MCACOD 0400H) logged into IA32_MCI_STATUS.

Implication: Due to this erratum, the processor may hang after acquiring a lock via XACQUIRE.

Workaround: BIOS can reserve the host physical address ranges of 40000000H and 403FFFFFFH (e.g. map it as UC/MMIO). Alternatively, the VMM (Virtual Machine Monitor) can reserve that address range so no guest can use it. In non-virtualized systems, the OS can reserve that memory space.

Status: No fix.

SKX73. IDI_MISC Performance Monitoring Events May be Inaccurate

Problem: The IDI_MISC.WB_UPGRADE and IDI_MISC.WB_DOWNGRADE performance monitoring events (Event FEH; UMask 02H and 04H) counts cache lines evicted from the L2 cache. Due to this erratum, the per logical processor count may be incorrect when both logical processors on the same physical core are active. The aggregate count of both logical processors is not affected by this erratum.

Implication: IDI_MISC performance monitoring events may be inaccurate.

Workaround: None identified.

Status: No fix.

**SKX74. Intel® PT CYC Packets Can be Dropped When Immediately Preceding PSB**

Problem: Due to a rare microarchitectural condition, generation of an Intel® PT (Processor Trace) PSB (Packet Stream Boundary) packet can cause a single CYC (Cycle Count) packet, possibly along with an associated MTC (Mini Time Counter) packet, to be dropped.

Implication: An Intel® PT decoder that is using CYCs to track time or frequency will get an improper value due to the lost CYC packet.

Workaround: If an Intel® PT decoder is using CYCs and MTCs to track frequency, and either the first MTC following a PSB shows that an MTC was dropped, or the CYC value appears to be 4095 cycles short of what is expected, the CYC value associated with that MTC should not be used. The decoder should wait for the next MTC before measuring frequency again.

Status: No fix.

SKX75. Intel® PT VM-entry Indication Depends on The Incorrect VMCS Control Field

Problem: An Intel® Processor Trace PIP (Paging Information Packet), which includes indication of entry into non-root operation, will be generated on VM-entry as long as the "Conceal VMX in Intel® PT" field (bit 19) in Secondary Execution Control register (IA32_VMX_PROCBASED_CTLDS2, MSR 048BH) is clear. This diverges from expected behavior, since this PIP should instead be generated only with a zero value of the "Conceal VMX entries from Intel® PT" field (Bit 17) in the Entry Control register (IA32_VMX_ENTRY_CTLDS MSR 0484H).

Implication: An Intel® PT trace may incorrectly expose entry to non-root operation.

Workaround: A VMM (virtual machine monitor) should always set both the "Conceal VMX entries from Intel® PT" field in the Entry Control register and the "Conceal VMX in Intel® PT" in the Secondary Execution Control register to the same value.

Status: No fix.

SKX76. System May Hang Due to Lock Prefixes on Instructions That Access IIO's MMCFG

Problem: If a core uses a lock prefix on an access to an IIO's MMCFG space, then it might lead to a hang if that same IIO has a pending level-triggered interrupt.

Implication: The system may hang and cause a log a machine check timeout if issuing lock prefixes on MMCFG accesses.

Workaround: Do not use lock prefixes on accesses to MMCFG lines.

Status: No fix.


SKX77. Intel® MBA Read After MSR Write May Return Incorrect Value

Problem: The MBA (Memory Bandwidth Allocation) feature defines a series of MSRs (0xD50-0xD57) to specify MBA Delay Values per Class of Service (CLOS), in the IA32_L2_QoS_Ext_BW_Thrtl_n MSR range. Certain values when written then read back may return an incorrect value in the MSR. Specifically, values greater than or equal to 10 (decimal) and less than 39 (decimal) written to the MBA Delay Value (Bits [15:0]) may be read back as 10%.

Implication: The values written to the registers will be applied; however, software should be aware that an incorrect value may be returned.

Workaround: None identified.

Status: No fix.



SKX78. In eMCA2 Mode, When The Retirement Watchdog Timeout Occurs CATERR# May be Asserted

Problem: A Retirement Watchdog Timeout (MCACOD = 0x0400) in Enhanced MCA2 (eMCA2) mode will cause the CATERR# pin to be pulsed in addition to an MSMI# pin assertion. In addition, a Machine Check Abort (#MC) will be pended in the cores along with the MSMI.

Implication: Due to this erratum, systems that expect to only see MSMI# will also see CATERR# pulse when a Retirement Watchdog Timeout occurs. The CATERR# pulse can be safely ignored.

Workaround: None identified.

Status: No fix.

SKX79. Intel® MBA May Incorrectly Throttle All Threads

Problem: When one logical processor is disabled, the MBA (Memory Bandwidth Allocation) feature may select an incorrect MBA throttling value to apply to the core. A disabled logical processor may behave as though the Class of Service (CLOS) field in its associated IA32_PQR_ASSOC MSR (0xC8F) is set to zero (appearing to be set to CLOS[0]). When this occurs, the MBA throttling value associated with CLOS[0] may be incorrectly applied to both threads on the core.

Implication: When Intel® Hyper-Threading technology is disabled or one logical thread on the core is disabled, the disabled thread is interpreted to have CLOS=0 set in its IA32_PQR_ASSOC MSR by hardware, which affects the calculation for the actual throttling value applied to the core. When this erratum occurs, the MBA throttling value associated with a given core may be incorrect.

Workaround: To work around this erratum, CLOS[0] should not be used if any logical cores are disabled. Alternately, software may leave all threads enabled.

Status: No fix.

SKX80. VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store

Problem: Execution of the VCVTPS2PH instruction with a memory destination may update the MXCSR exceptions flags (bits [5:0]) if the store to memory causes a fault (e.g., #PF) or VM exit. The value written to the MXCSR exceptions flags is what would have been written if there were no fault.

Implication: Software may see exceptions flags set in MXCSR, although the instruction has not successfully completed due to a fault on the memory operation. Intel has not observed this erratum to affect any commercially available software.

Workaround: None identified.

Status: No fix.

SKX81. Intel® PT May Drop All Packets After an Internal Buffer Overflow

Problem: Due to a rare microarchitectural condition, an Intel® PT (Processor Trace) ToPA (Table of Physical Addresses) entry transition can cause an internal buffer overflow that may result in all trace packets, including the OVF (Overflow) packet, being dropped.

Implication: When this erratum occurs, all trace data will be lost until either PT is disabled and re-enabled via IA32_RTIT_CTL.TraceEn [bit 0] (MSR 0570H) or the processor enters and exits a C6 or deeper C state.

Workaround: None identified.

Status: No fix.

**SKX82. Non-Zero Values May Appear in ZMM Upper Bits After SSE Instructions**

Problem: Under complex microarchitectural conditions, a VGATHER instruction with ZMM16-31 destination register followed by an SSE instruction in the next 4 instructions, may cause the ZMM register that is aliased to the SSE destination register to have non-zero values in bits 256-511. This may happen only when ZMM0-15 bits 256-511 are all zero, and there are no other instructions that write to ZMM0-15 in between the VGATHER and the SSE instruction. Subsequent SSE instructions that write to the same register will reset the affected upper ZMM bits and XSAVE will not expose these ZMM values as long as no other AVX512 instruction writes to ZMM0-15. This erratum will not occur in software that uses VZEROUPPER between AVX instructions and SSE instructions as recommended in the SDM.

Implication: Due to this erratum, an unexpected value may appear in a ZMM register aliased to an SSE destination. Software may observe this value only if the ZMM register aliased to the SSE instruction destination is used and VZEROUPPER is not used between AVX and SSE instructions. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: No fix.

SKX83. ZMM/YMM Registers May Contain Incorrect Values

Problem: Under complex microarchitectural conditions values stored in ZMM and YMM registers may be incorrect.

Implication: Due to this erratum, YMM and ZMM registers may contain an incorrect value. Intel® has not observed this erratum with any commercially available software.

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

Status: No fix.

SKX84. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions

Problem: An access to a GPA (guest-physical address) may cause an EPT-violation VM exit. When the "EPT-violation #VE" VM-execution control is 1, an EPT violation may cause a #VE (virtualization exception) instead of a VM exit. Due to this erratum, an EPT violation may erroneously cause a #VE when the "suppress #VE" bit is set in the EPT paging-structure entry used to map the GPA being accessed. This erratum does not apply when the "EPT-violation #VE" VM-execution control is 0 or when delivering an event through the IDT. This erratum applies only when the GPA in CR3 is used to access the root of the guest paging-structure hierarchy (or, with PAE paging, when the GPA in a PDPT is used to access a page directory).


Implication: When using PAE paging mode, an EPT violation that should cause a VMexit in the VMM may instead cause a VE# in the guest. In other paging modes, in addition to delivery of the erroneous #VE, the #VE may itself cause an EPT violation, but this EPT violation will be correctly delivered to the VMM.

Workaround: A VMM may support an interface that guest software can invoke with the VMCALL instruction when it detects an erroneous #VE.

Status: No fix.

SKX85. Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang

Problem: If an Intel® PT (Processor Trace) ToPA (Table of Physical Addresses) table is placed in UC (Uncacheable) or USWC (Uncacheable Speculative Write Combining) memory, and a ToPA output region is filled during an Intel® TSX (Transaction Synchronization) transaction, the resulting ToPA table read may cause a processor hang.



Implication: Placing Intel® PT ToPA tables in non-cacheable memory when Intel® TSX is in use may lead to a processor hang.

Workaround: None identified. Intel® PT ToPA tables should be located in WB memory if Intel® TSX is in use.

Status: No fix.

SKX86. Performing an XACQUIRE to an Intel® PT ToPA Table May Lead to Processor Hang

Problem: If an XACQUIRE lock is performed to the address of an Intel® PT (Processor Trace) ToPA (Table of Physical Addresses) table, and that table is later read by the CPU during the HLE (Hardware Lock Elision) transaction, the processor may hang.

Implication: Accessing ToPA tables with XACQUIRE may result in a processor hang.

Workaround: None identified. Software should not access ToPA tables using XACQUIRE. An OS or hypervisor may wish to ensure all application or guest writes to ToPA tables to take page faults or EPT violations.

Status: No fix.

SKX87. Use of VMX TSC scaling or TSC offsetting will result in corrupted Intel PT packets

Problem: When Intel® Processor Trace (Intel® PT) is enabled within a Virtual Machine Extensions (VMW) guest, and TSC (Time Stamp Counter) offsetting or TSC scaling is enabled for that guest, by setting primary processor-based execution control bit 3 or secondary processor-based execution control bit 25, respectively, in the VMCS (Virtual Machine Control Structure) for that guest, any TMA (TSC/MTC Alignment) packet generated will have corrupted values in the CTC (Core Timer Copy) and FastCounter fields. Additionally, the corrupted TMA packet will be followed by a bogus data byte.

Implication: An Intel® PT decoder will be confused when using the TMA packet to align cycle time with wall-clock time. The byte that follows the TMA will likely cause a decoder error for an unexpected or unrecognized packet.

Workaround: None identified. If a TMA packet with any reserved payload bits set is encountered by an Intel® PT decoder it should be ignored, along with the byte that immediately follows it. Alternatively, Intel PT users may opt to disable MTC and TMA packets by clearing IA32_RTIT_CTL.MTCEn[bit 9].

Status: No fix.

SKX88. Using Intel® TSX Instructions May Lead to Unpredictable System Behavior

Problem: Under complex microarchitectural conditions, software using Intel® Transactional Synchronization Extensions (Intel® TSX) may result in unpredictable system behavior. Intel has only seen this under synthetic testing conditions. Intel is not aware of any commercially available software exhibiting this behavior.

Implication: Due to this erratum, unpredictable system behavior may occur.

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

Status: No fix.

SKX89. Viral Mode of Error Containment (R_CPU07) may not properly handle data corruption containment as intended

Problem: Viral error notifications may not properly propagate to all other CPU agents (Intel UPI, CHA, M2MEM, IIO, and so forth).

Implication: Due to this erratum, data corruption containment may not be guaranteed.

Workaround: None identified.

Status: No fix.

**SKX90. Performance Monitoring General Purpose Counter 3 May Contain Unexpected Values**

Problem: When Restricted Transactional Memory (RTM) is supported (CPUID.07H.EBX.RTM [bit 11] = 1) and when TSX_FORCE_ABORT=0, Performance Monitor Unit (PMU) general purpose counter 3 (IA32_PMC3, MSR C4H and IA32_A_PMC3, MSR 4C4H) may contain unexpected values. Further, IA32_PREFEVTSEL3 (MSR 189H) may also contain unexpected configuration values.

Implication: Due to this erratum, software that uses PMU general purposes counter 3 may read an unexpected count and configuration.

Workaround: Software can avoid this erratum by writing 1 to bit 0 of TSX_FORCE_ABORT (MSR 10FH) which will cause all Restricted Transactional Memory (RTM) transactions to abort with EAX code 0. TSX_FORCE_ABORT MSR is available when CPUID.07H.EDX[bit 13]=1.

Status: No fix.

SKX91. Performance in an 8sg System May Be Lower Than Expected

Problem: In 8sg (8-socket glueless) systems, certain workloads may generate a significant stream of accesses to remote nodes, leading to unexpected congestion in the processor's snoop responses.

Implication: Due to this erratum, 8sg system performance may be lower than expected.

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

Status: No fix.

SKX92. Memory May Continue to Throttle after MEMHOT# De-assertion

Problem: When MEMHOT# is asserted by an external agent, the CPU may continue to throttle memory after MEMHOT# de-assertion.

Implication: When this erratum occurs, memory throttling occurs even after de-assertion of MEMHOT#.

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

Status: No fix.

SKX93. Unexpected Uncorrected Machine Check Errors May Be Reported

Problem: In rare micro-architectural conditions, the processor may report unexpected machine check errors. When this erratum occurs, IA32_MC0_STATUS (MSR 401H) will have the valid bit set (bit 63), the uncorrected error bit set (bit 61), a model specific error code of 03H (bits [31:16]) and an MCA error code of 05H (bits [15:0]).

Implication: Due to this erratum, software may observe unexpected machine check exceptions.

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

Workaround:

Status: No fix.

SKX94. Intel® PT Trace May Silently Drop Second Byte of CYC Packet

Problem: Due to a rare microarchitectural condition, the second byte of a 2-byte CYC (Cycle Count) packet may be dropped without an OVF (Overflow) packet.

Implication: A trace decoder may signal a decode error due to the lost trace byte.

Workaround: None identified. A mitigation is available for this erratum. If a decoder encounters a multi-byte CYC packet where the second byte has bit 0 (Ext) set to 1, it should assume that 4095 cycles have passed since the prior CYC packet, and it should ignore the first byte of the CYC and treat the second byte as the start of a new packet.

Status: No fix.

**SKX95. Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP**

Problem: Bit 63 of IA32_PERF_GLOBAL_STATUS_SET MSR (391H) is reserved. Due to this erratum, setting the bit will not result in General Protection Fault (#GP).

Implication: Software that attempts to set bit 63 of IA32_PERF_GLOBAL_STATUS_SET MSR does not generate #GP. There are no other system implications to this behavior.

Workaround: None identified.

Status: No fix.

SKX96. Branch Instruction Address May be Incorrectly Reported on TSX Abort When Using MPX

Problem: When using Intel® Memory Protection Extensions (MPX), an Intel® Transactional Synchronization Extensions (TSX) transaction abort will occur in case of legacy branch (that causes bounds registers INIT) when at least one MPX bounds register was in a NON-INIT state. On such an abort, the branch Instruction address should be reported in the FROM_IP field in the Last Branch Records (LBR), Branch Trace Store (BTS) and Branch Trace Message (BTM) as well as in the Flow Update Packets (FUP) source IP address for Processor Trace (PT). Due to this erratum, the FROM_IP field in LBR/BTS/BTM, as well as the Flow Update Packets (FUP) source IP address that correspond to the TSX abort, may point to the preceding instruction.

Implication: Software that relies on the accuracy of the FROM_IP field / FUP source IP address and uses TSX may operate incorrectly when MPX is used.

Workaround: None identified.

Status: No fix.

SKX97. 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. Intel has not observed this erratum in any commercially available software.

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: No fix.

SKX98. IMC Patrol Scrubbing Engine May Hang

Problem: Under rare microarchitectural conditions, the processor's Integrated Memory Controller (IMC) Patrol Scrubbing Engine may hang.

Implication: When this erratum occurs, IMC Patrol Scrubbing will cease. Intel has only observed this erratum in a synthetic test environment when testing with high rates of ECC errors.

Workaround: None identified.

Status: No fix.

SKX99. Intel® MBM Counters May Report System Memory Bandwidth Incorrectly

Problem: Intel® Memory Bandwidth Monitoring (MBM) counters track metrics according to the assigned Resource Monitor ID (RMID) for that logical core. The IA32_QM_CTR register (MSR 0xC8E), used to report these metrics, may report incorrect system bandwidth for certain RMID values.



Implication: Due to this erratum, system memory bandwidth may not match what is reported.

Workaround: It is possible for software to contain code changes to work around this erratum. Please see the white paper titled Intel® Resource Director Technology (Intel® RDT) Reference Manual found at <https://software.intel.com/en-us/intel-resource-director-technology-rdt-reference-manual> for more information.

Status: No fix.

SKX100. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

Problem: Resuming from C6 Sleep-State, with Fixed Interrupts of the same priority queued (in the corresponding bits of the IRR and ISR APIC registers), the processor may dispatch the second interrupt (from the IRR bit) before the first interrupt has completed and written to the EOI register, causing the first interrupt to never complete.

Implication: Due to this erratum, Software may behave unexpectedly when an earlier call to an Interrupt Handler routine is overridden with another call (to the same Interrupt Handler) instead of completing its execution.

Workaround: None identified.

Status: No fix.

SKX101. Voltage/Frequency Curve Transitions May Result in Machine Check Errors or Unpredictable System Behavior

Problem: Under complex microarchitecture conditions, during voltage/frequency curve transitions, 3-strike machine check errors or other unpredictable system behavior may occur due to an issue in the FIVR logic.

Implication: When this erratum occurs, the system may cause a 3 strike machine check error or other unpredictable system behavior.

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

Status: No fix.

SKX102. Processor May Behave Unpredictably on Complex Sequence of Conditions Which Involve Branches That Cross 64 Byte Boundaries

Problem: Under complex micro-architectural conditions involving branch instructions bytes that span multiple 64 byte boundaries (cross cache line), unpredictable system behavior may occur.

Implication: When this erratum occurs, the system may behave unpredictably.

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

Status: No fix.

SKX103. Executing Some Instructions May Cause Unpredictable Behavior

Problem: Under complex micro-architectural conditions, executing an X87, AVX, or integer divide instruction may result in unpredictable system behavior.


Implication: When this erratum occurs, the system may behave unpredictably. Intel has not observed this erratum with any commercially available software.

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

Status: No fix.

SKX104. STIBP May Not Function as Intended

Problem: The Single Thread Indirect Branch Predictors bit (IA32_SPEC_CTL[STIBP] (MSR 48H, bit 1)) prevents the predicted targets of indirect branches on any logical processor of that core from being controlled by software that executes (or executed previously) on another logical processor of the same core. Under specific microarchitectural conditions one logical processor may be able to control the predicted targets of indirect branches



on the other logical processor even when one of the logical processors has set the STIBP bit.

Implication: Software relying on STIBP to mitigate against cross-thread speculative branch target injection may allow an attacker running on one logical processor to induce another logical processor on the same core to speculatively execute a disclosure gadget that could reveal confidential data through a side-channel method called Branch Target Injection. This erratum does not affect processors with Hyper-Threading disabled or enabling the cross thread protections of Indirect Branch Restricted Speculation bit (IA32_SPEC_CTL[IBRS] (MSR 48H, bit 0)).

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

Status: No fix.

SKX105. Intel® UPI, DMI and PCIe Interfaces May See Elevated Bit Error Rates

Problem: The Intel® Ultra Path Interconnect (Intel® UPI), Direct Media Interface (DMI) or Peripheral Component Interconnect Express (PCIe) interfaces may be subject to a high bit error rate.

Implication: Due to this erratum, an elevated rate of packet CRC errors may be observed on these interfaces which may lead to a Machine Check Error and/or may hang the system.

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

Status: No fix.

SKX106. Unexpected Page Faults in Guest Virtualization Environment

Problem: Under complex micro-architectural conditions, a virtualized guest could observe unpredictable system behavior.

Implication: When this erratum occurs, systems operating in a virtualization environment may exhibit unexpected page faults (double faults) leading to guest OS shutdown.

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

Status: No fix.

SKX107. PCIe Root Port Does Not Increment REPLAY_NUM on Multiple NAKs of The Same TLP

Problem: PCIe Root Port does not increment REPLAY_NUM on a replay initiated by a duplicate NAK for the same TLP (Transaction Layer Packet) and does not retain the Link.

Implication: If a non-compliant Endpoint NAKs the same TLP repeatedly, the lack of forward progress can lead to (PCIe Completion, TOR, Internal Timer MCE) timeout.

Workaround: None identified.

Status: No fix.

SKX108. Memory Controller May Hang While in Virtual Lockstep

Problem: Under complex microarchitectural conditions, a memory controller that is in Virtual Lockstep (VLS) may hang on a partial write transaction.

Implication: The memory controller hangs with a mesh-to-mem timeout Machine Check Exception (MSCOD=20h, MCACOD=400h). The memory controller hang may lead to other machine check timeouts that can lead to an unexpected system shutdown.

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

Status: No fix.

SKX109. Direct Branches With Partial Address Aliasing May Lead to Unpredictable System Behavior

Problem: Under complex micro-architectural conditions involving direct branch instructions with partial address aliasing, unpredictable system behavior may occur. Intel has only seen



this under synthetic testing conditions. Intel has not observed this under any commercially available software.

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

Workaround: None identified.

Status: No Fix.

SKX110. PCIe Function Level Reset May Generate an #NMI Exception

Problem: Under either of the following conditions, the processor will log a parity error in OTC_IRP_DAT_PAR register bit (RootBus, Device 5, Function 2, Offset 0X288, bits 11):

1. If CRS Software Visibility Enable bit (RootBus, Device [0-3], Function 0, Offset ACh, bit[4]) is not set,
2. Or if the first transaction sent to the endpoint is not a CfgRd (Configuration Read) transaction following a PCIe Function Level Reset or Secondary Bus Reset event affecting the root port.

Implication: When this erratum occurs, the processor will generate an unexpected #NMI exception, which may lead to a system hang or shutdown.

Workaround: Software should set CRS Software Visibility Enable bit to '1. Alternatively, software must ensure that the initial request targeting the end point is a CfgRd request.

Status: No Fix.

SKX111. Performance Monitoring General Counter 2 May Have Invalid Value Written When TSX Is Enabled

Problem: When Transactional Synchronization Extensions (TSX) is enabled, and there are aborts (HLE or RTM) overlapping with access or manipulation of the IA32_PMC2 general-purpose performance counter (Offset: C3h) it may return invalid value.

Implication: Software may read invalid value from IA32_PMC2.

Workaround: None identified.

Status: No Fix.

SKX112. Intel® QuickData Technology Engine May Hang With Any DMA Error if Completion Status is Improperly Set

Problem: If the Intel® QuickData Technology Engine (CBDMA) Error Completion Enable register (CHANCTRL.ERR_CMP_EN; CB_BAR Offset 80h; bit 2) is set, but the DMA descriptor's Generate completion status update is not enabled, the CBDMA engine may hang on any DMA error.

Implication: When this erratum occurs, software using the Intel® QuickData Technology Engine may not behave as expected due to a DMA error.

Workaround: Always enable the Generate completion status update in the DMA descriptor when setting CHANCTRL.ERR_CMP_EN.

Status: No fix.


SKX113. Overflow Flag in IA32_MC0_STATUS MSR May be Incorrectly Set

Problem: Under complex micro-architectural conditions, a single internal parity error seen in IA32_MC0_STATUS MSR (401h) with MCACOD (bits 15:0) value of 5h and MSCOD (bits 31:16) value of 7h, may set the overflow flag (bit 62) in the same MSR.

Implication: Due to this erratum, the IA32_MC0_STATUS overflow flag may be set after a single parity error. Intel has not observed this erratum with any commercially available software.

Workaround: None identified.

Status: No fix.



SKX114. A Fixed Interrupt May Be Lost When a Core Exits C6

Problem: Under complex micro-architectural conditions, when performance throttling happens during a core C6 exit, a fixed interrupt may be lost.

Implication: Due to this erratum, a fixed interrupt may be lost when internal throttling happens during a core C6 exit. Intel has only observed this erratum in synthetic test conditions.

Workaround: None identified.

Status: No fix.



Specification Changes

SKX1C. SMM Handler Code Access Control May Not Be Available

Problem: The SMM Handler Code Access Control is not enabled on some processors.

Implication: Due to this, SMM cannot use the SMM Handler Code Access Control. The lack of support for this feature is properly enumerated through MSR_SMM_MCA_CAP (MSR 17DH) SMM_Code_Access_Chk bit (bit 58) being set to 0 indicating that the feature is not available.

Workaround: It is possible for BIOS to contain a workaround that adds support for the SMM Handler Code Access Control.



Specification Clarifications

There are no Specification Clarifications in this Specification Update revision.



Documentation Changes

There are no Documentation Changes in this Specification Update revision.

§