

Intel® Xeon® Processor Scalable Family

Specification Update

February 2018

Notice: The Intel® Xeon® Processor Scalable Family may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Reference Number: 336065-005



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

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

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

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

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

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

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

Intel, Intel Core, Intel® Core™ i7, Pentium®, Pentium® Pro, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2018, Intel Corporation. All Rights Reserved.



Contents

Revision History	4
Preface	5
Summary Tables of Changes	7
Identification Information	11
Errata	24
Specification Changes	40
Specification Clarifications	41
Documentation Changes	42



Revision History

Date	Revision	Description
July 2017	001	Initial Release (Intel Public).
July 2017	002	Added Errata 21 through 28.
August 2017	003	Added Errata 29 through 36.
November 2017	004	Added Errata 37 through 71.
February 2018	005	Updated Errata 23 and 24



Preface

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

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

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

Affected Documents

Document Title	Document Number/ Location
<i>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.



Nomenclature


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

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

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

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

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



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.



Errata (Sheet 1 of 3)

Number	Steppings			Status	Errata
	H-0	M-0	U-0		
1.	x	x	x	No Fix	A CAP Error While Entering Package C6 Might Cause DRAM to Fail to Enter Self-Refresh (Intel® Xeon® Processor Scalable Family)
2.	x	x	x	No Fix	PCIe* Lane Error Status Register Might Log False Correctable Error (Intel® Xeon® Processor SP)
3.	x	x	x	No Fix	In Memory Mirror Mode, DataErrorChunk Field Might be Incorrect (Intel® Xeon® Processor SP)
4.	x	x	x	No Fix	Intel® Resource Director Technology (Intel® RDT) MBM Does Not Accurately Track Write Bandwidth (Intel® Xeon® Processor SP)
5.	x	x	x	No Fix	Intel UPI Initialization Aborts Might be Logged (Intel® Xeon® Processor SP)
6.	x	x	x	No Fix	PCIe Port Might Incorrectly Log Malformed_TLP Error (Intel® Xeon® Processor SP)
7.	x	x	x	No Fix	CMCI Might Not be Signaled for Corrected Error (Intel® Xeon® Processor Scalable Family)
8.	x	x	x	No Fix	Intel® CAT/CDP Might Not Restrict Cacheline Allocation Under Certain Conditions (Intel® Xeon® Processor Scalable Family)
9.	x	x	x	No Fix	Credits Not Returned For PCIe Packets That Fail ECRC Check Problem (Intel® Xeon® Processor SP)
10.	x	x	x	No Fix	PCIe Link Might Fail to Train (Intel® Xeon® Processor SP)
11.	x	x	x	No Fix	Intel UPI CRC32 Rolling Mode is Not Functional (Intel® Xeon® Processor SP)
12.	x	x	x	No Fix	IODC Entry 0 Cannot be Masked (Intel® Xeon® Processor SP)
13.	x	x	x	No Fix	With eMCA2 Enabled a 3-Strike Might Cause an Unnecessary CATERR# Instead of Only MSMI (Intel® Xeon® Processor SP)
14.	x	x	x	No Fix	CMCI May Not be Signaled for Corrected Error (Intel® Xeon® Processor Scalable Family)
15.	x	x	x	No Fix	CSRs SVID and SDID Are Not Implemented For Some DDRIO and PCU Devices (Intel® Xeon® Processor SP)
16.	x	x	x	No Fix	Register Broadcast Read From DDRIO May Return a Zero Value (Intel® Xeon® Processor SP)
17.	x	x	x	No Fix	Intel® CMT Counters May Not Count Accurately (Intel® Xeon® Processor SP)
18.	x	x	x	No Fix	Intel® CAT Might Not Restrict Cacheline Allocation Under Certain Conditions (Intel® Xeon® Processor Scalable Family)
19.	x	x	x	No Fix	Intel® PCIe* Corrected Error Threshold Does Not Consider Overflow Count When Incrementing Error Counter (Intel® Xeon® Processor SP)
20.	x	x	x	No Fix	IIO RAS VPP Hangs During The Warm Reset Test (Intel® Xeon® Processor SP)
21.	x	x	x	No Fix	Intel UPI CRC Errors and PHY Init Aborts May Be Seen During UPI Slow Mode Training
22.	x	x	x	No Fix	A Core 3-Strike Event May Be Seen Under Certain Test Conditions
23.	x	x	x	No Fix	DDR4 Memory Bandwidth May Be Lower Than Expected at 2133 and 1866 Speeds
24.	x	x	x	No Fix	Lower Than Expected Performance May Be Seen Under Certain Intel AVX2 and Intel AVX-512 Workloads
25.	x	x	x	No Fix	A System Hang May Be Seen With Some 8S + XNC Type Platform Configurations
26.	x	x	x	No Fix	Sparing Per-Rank Error Masking Does Not Mask Correctable Errors



Errata (Sheet 2 of 3)

Number	Steppings			Status	Errata
	H-0	M-0	U-0		
27.	x	x	x	No Fix	PCIe* Root Port Electromechanical Interlock Control Register Can Be Written
28.	x	x	x	No Fix	Live Error Recovery Feature Being Disabled is not Getting Reflected in PXP2CAP Value
29.	x	x	x	No Fix	Performance Monitoring M2MEM Counters For Memory Controller Reads/Writes Are Not Counting Read/Write Retries
30.	x	x	x	No Fix	System Hangs May Occur When IPQ and IRQ Requests Happen at The Same Time
31.	x	x	x	No Fix	Two Intel® UPI Reads From XNC May Lead to a System Hang
32.	x	x	x	No Fix	IIO VPP May Hang During Warm Reset
33.	x	x	x	No Fix	Machine Check Events may be logged in banks 9, 10 and 11 that do not represent actual errors
34.	x	x	x	No Fix	Advanced RAS Dynamic Link Width Reduction may not resize the Intel UPI link
35.	x	x	x	No Fix	Lower than expected performance may be seen with some Intel AVX workloads due to incorrect uncore frequency scaling
36.	x	x	x	No Fix	Unexpected DDR ECC Errors May be Seen
37.	x	x	x	No Fix	Spurious Corrected Errors May be Reported
38.	x	x	x	No Fix	Dynamic Link Width Reduction May Not Resize the Intel UPI Link
39.	x	x	x	No Fix	Writing to LT_LOCK_MEMORY and LT_UNLOCK_MEMORY MSRs Simultaneously May Have Inconsistent Results
40.	x	x	x	No Fix	Masked Bytes in a Vector Masked Store Instructions May Cause Write Back of a Cache Line
41.	x	x	x	No Fix	ERROR_N[2:0] Pins May Not be Cleared After a Warm Reset
42.	x	x	x	No Fix	CRC Store Operation Corner Case May Result in Hang
43.	x	x	x	No Fix	Atomicity May Not be Preserved When Executing With RTM Enabled
44.	x	x	x	No Fix	Intel PCIe Slot Presence Detect and Presence Detect Changed Logic Not PCIe Specification Compliant
45.	x	x	x	No Fix	In Patrol Scrub System Address Mode, Address is Not Loaded from CSRs After Re-enable
46.	x	x	x	No Fix	Intel® Processor Trace (Intel® PT) TIP.PGD May Not Have Target IP Payload
47.	x	x	x	No Fix	The Corrected Error Count Overflow Bit in IA32_MCO_STATUS is Not Updated When The UC Bit is Set
48.	x	x	x	No Fix	SMRAM State-Save Area Above the 4 GB Boundary May Cause Unpredictable System Behavior
49.	x	x	x	No Fix	POPCNT Instruction May Take Longer to Execute Than Expected
50.	x	x	x	No Fix	Load Latency Performance Monitoring Facility May Stop Counting
51.	x	x	x	No Fix	Intel® PT PSB+ Packets May Contain Unexpected Packets
52.	x	x	x	No Fix	Performance Monitoring Counters May Undercount When Using CPL Filtering
53.	x	x	x	No Fix	Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters
54.	x	x	x	No Fix	Incorrect Branch Predicted Bit in BTS/BTM Branch Records
55.	x	x	x	No Fix	DR6.B0-B3 May Not Report All Breakpoints Matched When a MOV/POP SS is Followed by a Store or an MMX Instruction
56.	x	x	x	No Fix	Performance Monitoring Load Latency Events May Be Inaccurate For Gather Instructions



Errata (Sheet 3 of 3)

Number	Steppings			Status	Errata
	H-0	M-0	U-0		
57.	x	x	x	No Fix	VM Exit May Set IA32_EFER.NXE When IA32_MISC_ENABLE Bit 34 is Set to 1
58.	x	x	x	No Fix	x87 FPU Exception (#MF) May be Signaled Earlier Than Expected
59.	x	x	x	No Fix	CPUID TLB Associativity Information is Inaccurate
60.	x	x	x	No Fix	Vector Masked Store Instructions May Cause Write Back of Cache Line Where Bytes Are Masked
61.	x	x	x	No Fix	Incorrect FROM_IP Value For an RTM Abort in BTM or BTS May be Observed
62.	x	x	x	No Fix	MOVNTDQA From WC Memory May Pass Earlier Locked Instructions
63.	x	x	x	No Fix	#GP on Segment Selector Descriptor that Straddles Canonical Boundary May Not Provide Correct Exception Error Code
64.	x	x	x	No Fix	Intel® PT OVF Packet May be Lost if Immediately Preceding a TraceStop
65.	x	x	x	No Fix	The Intel PT CR3 Filter is Not Re-evaluated on VM Entry
66.	x	x	x	No Fix	BNDLDX and BNDSTX May Not Signal #GP on Non-Canonical Bound Directory Access
67.	x	x	x	No Fix	Performance Monitor Event For Outstanding Offcore Requests and Snoop Requests May be Incorrect
68.	x	x	x	No Fix	Branch Instructions May Initialize MPX Bound Registers Incorrectly
69.	x	x	x	No Fix	A Spurious APIC Timer Interrupt May Occur After Timed MWAIT
70.	x	x	x	No Fix	Writing a Non-Canonical Value to an LBR MSR Does Not Signal a #GP When Intel® PT is Enabled
71.	x	x	x	No Fix	VM Entry That Clears TraceEn May Generate a FUP

Specification Changes

Number	Specification Changes
1	None for this revision of this specification update.

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 12](#) 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	



Non-Intel® Advanced Vector Extensions (non-Intel® AVX), Intel® Advanced Vector Extensions (Intel® AVX), and Intel® Advanced Vector Extensions 512 (Intel® AVX-512) Turbo Frequencies

Figure 1. Intel® Xeon® Processor Scalable Family Non-Intel AVX Turbo Frequencies

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

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
8180	28	38.50	205	2.5	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.3	3.3	3.3	3.3	3.2	3.2	3.2	3.2		
8168	24	33.00	205	2.7	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4		
6148	20	27.50	150	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									
6154	18	24.75	200	3.0	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7		
6150	18	24.75	165	2.7	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.4		
6142	16	22.00	150	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3													
6132	14	19.25	140	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3													
6146	12	24.75	165	3.2	4.2	4.2	4.1	4.1	4.0	4.0	4.0	4.0	3.9	3.9	3.9																	
6136	12	24.75	150	3.0	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6													
6126	12	19.25	125	2.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3	3.3														
6144	8	24.75	150	3.5	4.2	4.2	4.1	4.1	4.1	4.1	4.1	4.1	4.1																			
6134	8	24.75	130	3.2	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7																			
6128	6	19.25	115	3.4	3.7	3.7	3.7	3.7	3.7	3.7																						
5122	4	16.50	105	3.6	3.7	3.7	3.7	3.7																								

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

Figure 2. Intel® Xeon® Processor Scalable Family Intel AVX 2.0 Turbo Frequencies

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

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
8180	28	38.50	205	2.1	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	3.3	3.3	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9	2.8	2.8	2.8		
8168	24	33.00	205	2.3	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	3.3	3.3	3.2	3.2	3.2	3.2	3.0	3.0	3.0	3.0					
6148	20	27.50	150	1.9	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.6								
6154	18	24.75	200	2.6	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	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.3		
6150	18	24.75	165	2.3	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	3.1	3.1	3.1	3.1	3.0	3.0									
6142	16	22.00	150	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.2	3.2	3.2	3.2	2.9	2.9	2.9	2.9												
6132	14	19.25	140	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.9	2.9														
6146	12	24.75	165	2.6	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	3.3	3.3													
6136	12	24.75	150	2.6	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	3.3	3.3													
6126	12	19.25	125	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																
6144	8	24.75	150	2.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5																				
6134	8	24.75	130	2.7	3.6	3.6	3.4	3.4	3.4	3.4	3.4	3.4																				
6128	6	19.25	115	2.9	3.6	3.6	3.6	3.6	3.6	3.6																						
5122	4	16.50	105	3.3	3.6	3.6	3.6	3.6																								

- The 8180, 6142, and 6134 have 1.5 TB/socket memory capacity versions (8180M, 6142M, and 6134M – not listed above) with identical frequencies.



- The 8156 (not listed above) has identical frequencies to 5122 but adds third Intel UPI and 8-socket capability.
- The 8158 (not listed above) 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	2.8	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.3	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	2.9	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	2.2									
6154	18	24.75	200	2.1	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	2.8	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	3.2	2.9	2.9	2.9	2.9	2.6	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	2.4	2.2	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	2.4	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																					
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 above) with identical frequencies.
- The 8156 (not listed above) has identical frequencies to 5122 but adds third Intel UPI and 8-socket capability.
- The 8158 (not listed above) 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	3.4	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9	2.8	2.8	2.8	2.8	
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	3.3	3.0	3.0	3.0	3.0	2.8	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	3.4	3.2	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	3.4	3.2	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	3.4	3.1	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	3.4	3.2	3.2	3.2	3.2	2.9	2.9	2.9	2.9	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	3.4	3.1	3.1	3.1	3.1	3.0	3.0											
6153	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	3.4	3.1	3.1	3.1	3.1	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 above) 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	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	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	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	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	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	2.5	2.3	2.3	2.3	2.3	2.3	2.3	2.3	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	3.0	2.7	2.7	2.7	2.7	2.6	2.6	2.6	2.6	2.6	2.6	2.6	2.6	2.6	2.6	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	2.0	2.0	2.0	2.0	2.0	2.0	2.0	2.0	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	2.4	2.4	2.4	2.4	2.4	2.4	2.4	2.4	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 above) 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	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	2.1	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	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	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	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	1.9	1.9	1.9	1.9	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	2.1	2.1	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	1.6	1.6	1.6	1.6	1.6	1.6	1.6	1.6	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	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	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 above) 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	2.7	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	2.1	2.1	2.1	1.9	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	2.3	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	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									
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	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	2.6	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	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	2.2	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	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														




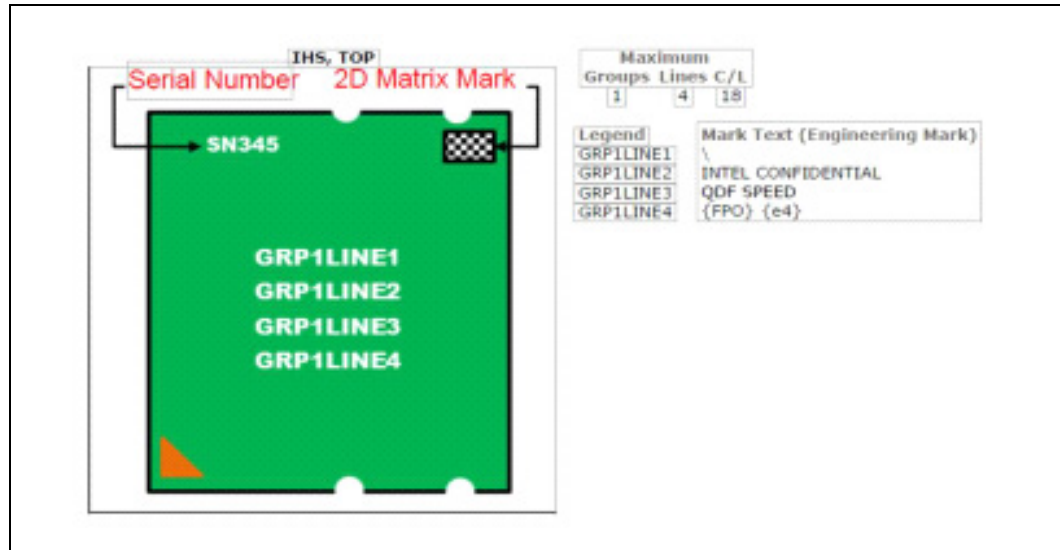
Figure 15. Intel® Xeon® Processor Scalable Family Non-Intel AVX, Intel AVX 2.0, and Intel AVX-512 Turbo Frequencies

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)



The Intel® Xeon® Processor Scalable Family stepping can be identified by the following register markings.

Table 4. Intel® Xeon® Processor Scalable Family Identification (Sheet 1 of 3)

S-Spec Number ¹	Die	Stepping	CPUID	Speed	DDR4 1DPC (MHz)	TDP (W)	# of Cores	LLC Cache Size (MB)	Maximum Supported Sockets/Intel® UPI Links
QM7W	XCC	H-0	0x50654	2.0	2666	145	24	33	Scalable 8 Socket/ 3
QM7X	XCC	H-0	0x50654	2.5	2666	205	28	38.5	Scalable 8 Socket/ 3
QM7Y	XCC	H-0	0x50654	2.5	2666	145	16	22	Scalable 4 Socket/ 3
QM7Z	XCC	H-0	0x50654	2.3	2666	145	20	27.5	Scalable 4 Socket/ 3
QM80	XCC	H-0	0x50654	2.1	2666	165	28	38.5	Scalable 8 Socket/ 3
QM86	XCC	H-0	0x50654	2.3	2666	135	18	24.75	Scalable 4 Socket/ 3
QM81	XCC	H-0	0x50654	3.0	2666	200	18	33	Scalable 4 Socket/ 3
QM8W	XCC	H-0	0x50654	2.0	2666	135	22	30.25	Scalable 4 Socket/ 3
QM8X	XCC	H-0	0x50654	2.0	2666	120	16	22	Scalable 4 Socket/ 3
QMAD	XCC	H-0	0x50654	2.1	2666	165	26	35.75	Scalable 8 Socket/ 3
QMAB	XCC	H-0	0x50654	2.7	2666	205	24	33	Scalable 8 Socket/ 3
QMA9	XCC	H-0	0x50654	2.7	2666	165	18	24.75	Scalable 4 Socket/ 3
QM7U	XCC	H-0	0x50654	2.5	2666	205	28	38.5	Scalable 8 Socket/ 3
QM7R	XCC	H-0	0x50654	2.1	2666	165	28	38.5	Scalable 8 Socket/ 3
QMAC	XCC	H-0	0x50654	2.1	2666	165	26	35.75	Scalable 8 Socket/ 3
QMQ7	XCC	H-0	0x50654	2.5	2666	205	28	38.5	Scalable 8 Socket/ 3
QMQ3	XCC	H-0	0x50654	2.1	2666	165	28	38.5	Scalable 8 Socket/ 3
QMQC	XCC	H-0	0x50654	2.1	2666	165	26	35.75	Scalable 8 Socket/ 3
QMQ8	XCC	H-0	0x50654	2.7	2666	205	24	33	Scalable 8 Socket/ 3



Table 4. Intel® Xeon® Processor Scalable Family Identification (Sheet 2 of 3)

S-Spec Number ¹	Die	Stepping	CPUID	Speed	DDR4 1DPC (MHz)	TDP (W)	# of Cores	LLC Cache Size (MB)	Maximum Supported Sockets/Intel® UPI Links
QMQ9	XCC	H-0	0x50654	3.0	2666	200	18	33	Scalable 4 Socket/ 3
QMQ6	XCC	H-0	0x50654	2.7	2666	165	18	24.75	Scalable 4 Socket/ 3
QMJ7	XCC	H-0	0x50654	2.6	2666	155	16	22	2 / 2
QMHX	XCC	H-0	0x50654	2.1	2666	130	16	22	2 / 2
QMHS	XCC	H-0	0x50654	2.6	2666	130	12	19.25	2 / 2
QMS8	XCC	H-0	0x50654	2.0	2666	150	26	35.75	Scalable 8 Socket/ 3
QMRV	XCC	H-0	0x50654	2.1	2666	150	24	33	Scalable 8 Socket/ 3
QMRG ²	XCC	H-0	0x50654	2.1	2666	150	24	33	Scalable 8 Socket/ 3
QMS3 ³	XCC	H-0	0x50654	2.1	2666	150	24	33	Scalable 8 Socket/ 3
QMRZ	XCC	H-0	0x50654	2.1	2666	140	22	30.25	Scalable 4 Socket/ 3
QMS1	XCC	H-0	0x50654	2.4	2666	150	20	27.5	Scalable 4 Socket/ 3
QMS0	XCC	H-0	0x50654	2.0	2666	125	20	27.5	Scalable 4 Socket/ 3
QMS9 ²	XCC	H-0	0x50654	2.0	2666	125	20	27.5	Scalable 4 Socket/ 3
QMRS	XCC	H-0	0x50654	2.3	2666	140	18	24.75	Scalable 4 Socket/ 3
QMRU ³	XCC	H-0	0x50654	2.3	2666	140	18	24.75	Scalable 4 Socket/ 3
QMRT	XCC	H-0	0x50654	2.6	2666	150	16	22	Scalable 4 Socket/ 3
QMRW ³	XCC	H-0	0x50654	2.6	2666	150	16	22	Scalable 4 Socket/ 3
QMS6	XCC	H-0	0x50654	2.1	2666	125	16	22	Scalable 4 Socket/ 3
QMRH ²	XCC	H-0	0x50654	2.1	2666	125	16	22	Scalable 4 Socket/ 3
QMRX	XCC	H-0	0x50654	3.0	2666	150	12	24.75	Scalable 4 Socket/ 3
QMRY	XCC	H-0	0x50654	2.6	2666	125	12	19.25	Scalable 4 Socket/ 3
QMRJ ²	XCC	H-0	0x50654	2.6	2666	125	12	19.25	Scalable 4 Socket/ 3
QMRM ³	XCC	H-0	0x50654	3.2	2666	130	8	24.75	Scalable 4 Socket/ 3
QMRL	XCC	H-0	0x50654	3.2	2666	130	8	24.75	Scalable 4 Socket/ 3
QMRN	XCC	H-	0x50654	3.6	2666	105	4	16.5	Scalable 4 Socket/ 3
QMRQ	XCC	H-0	0x50654	3.6	2666	105	4	16.5	Scalable 8 Socket/ 3
QN33	XCC	H-0	0x50654	2.6	2666	140	14	19.25	Scalable 4 Socket/ 3
QN34	XCC	H-0	0x50654	3.4	2666	115	6	19.25	Scalable 4 Socket/ 3
QN35	XCC	H-0	0x50654	3.0	2666	200	18	24.75	Scalable 4 Socket/ 3
QN36	XCC	H-0	0x50654	2.1	2666	150	24	33	Scalable 8 Socket/ 3
QN37	XCC	H-0	0x50654	2.0	2666	125	20	27.5	Scalable 4 Socket/ 3
QN38	XCC	H-0	0x50654	2.1	2666	125	16	22	Scalable 4 Socket/ 3
QN39	XCC	H-0	0x50654	2.6	2666	125	12	19.25	Scalable 4 Socket/ 3
QN3B	XCC	H-0	0x50654	2.1	2666	135	16	22	2 / 2
QN3C	XCC	H-0	0x50654	2.6	2666	135	12	19.25	2 / 2
QN3D	XCC	H-0	0x50654	2.1	2666	165	28	38.5	2 / 2
QN3E	XCC	H-0	0x50654	2.1	2666	160	24	33	2 / 2
QN3F	XCC	H-0	0x50654	2.6	2666	160	16	22	2 / 2
QN3G	XCC	H-0	0x50654	2.4	2666	160	20	27.5	2 / 2



Table 4. Intel® Xeon® Processor Scalable Family Identification (Sheet 3 of 3)

S-Spec Number ¹	Die	Stepping	CPUID	Speed	DDR4 1DPC (MHz)	TDP (W)	# of Cores	LLC Cache Size (MB)	Maximum Supported Sockets/Intel® UPI Links
QN3H	XCC	H-0	0x50654	2.0	2666	135	20	27.5	2 / 2
QN98	XCC	H-0	0x50654	2.0	2400	113	14	19.25	2 / 2
QM8S	XCC	H-0	0x50654	2.0	2400	105	14	19.25	4 / 2
QN7D	XCC	H-0	0x50654	1.8	2666	150	8	24.75	Scalable 4 Socket/ 3
QN7C	XCC	H-0	0x50654	1.8	2666	165	12	24.75	Scalable 4 Socket/ 3
QMXG	HCC	M-0	0x50654	2.4	2400	85	10	13.75	Scalable 4 Socket/ 2
QMXK	HCC	M-0	0x50654	2.1	2400	85	12	16.5	2 / 2
QMXH	HCC	M-0	0x50654	2.3	2400	105	12	16.5	Scalable 4 Socket/ 2
QMXJ	HCC	M-0	0x50654	2.2	2400	105	14	19.25	Scalable 4 Socket/ 2
QMXL	HCC	M-0	0x50654	2.2	2400	105	14	19.25	Scalable 4 Socket/ 2
QN1J	HCC	M-0	0x50654	2.1	2400	85	12	16.5	2 / 2
QN09	LCC	U-0	0x50654	2.1	2400	85	8	11	2 / 2
QN0A	LCC	U-0	0x50654	1.8	2400	85	8	11	2 / 2
QN0B	LCC	U-0	0x50654	2.2	2400	85	10	13.75	2 / 2
QN0C	LCC	U-0	0x50654	1.7	2133	85	8	11	2 / 2
QN0D	LCC	U-0	0x50654	1.7	2133	85	6	8.25	2 / 2
QN0E	LCC	U-0	0x50654	2.6	2400	85	4	8.25	2 / 2
QN0F	LCC	U-0	0x50654	2	2400	70	8	11	2 / 2

1. These are QS samples of the Intel® Xeon® Processor Scalable Family.
2. These are QS samples of the Intel® Xeon® Processor Scalable Family high T_{CASE} + extended reliability.
3. These are QS samples of the Intel® Xeon® Processor Scalable Family with high memory (1.5 TB).



Errata

1. **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.

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

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

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

Workaround: None.

Status: No Fix.

3. **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.

4. **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.


5. **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.



6. **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.

7. **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.

8. **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.

9. **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.

10. **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.



11. 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

12. 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.

13. 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.

14. 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.

15. 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.



16. Register Broadcast Read From DDRIO May Return a Zero Value (Intel® Xeon® Processor SP)

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.

17. 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.

18. 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.

19. 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.



20. 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.

21. 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

22. 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.

23. 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.

24. Lower Than Expected Performance May Be Seen Under Certain Intel AVX2 and Intel AVX-512 Workloads

Problem: A workload balancing mechanism in Intel Xeon processor scalable family CPU may have the potential to incorrectly apply a larger than expected weighting value to some Intel AVX2 and Intel AVX-512 workloads. This may cause a lower than expected frequency drop to occur, potentially affecting CPU performance under these workloads.

Implication: Lower than expected performance may be seen under certain Intel AVX2 and Intel AVX-512 workloads.

Workaround: None identified.

Status: No fix.

25. 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.

26. 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

27. 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

28. 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, 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

29. 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

30. 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

31. 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

32. 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.

33. 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 & may be end-user visible while not representing actual errors.

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

Status: No Fix.

34. 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.

35. Lower than expected performance may be seen with some Intel AVX workloads due to incorrect uncore frequency scaling

Problem: Due to a problem with UFS (Uncore Frequency Scaling), lower than expected performance may be seen with some Intel AVX workloads. The CPU may not ramp uncore frequency when running some Intel AVX workloads depending on the number of active cores.



Implication: Lower than expected performance may be seen with some Intel AVX workloads due to incorrect uncore frequency scaling.

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

Status: No Fix.

36. Unexpected DDR ECC Errors May be Seen

Problem: The processor may incorrectly configure the processor's Vccp rail voltage.

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.

37. 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.

38. 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.

39. 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.

40. 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.

41. 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.

42. 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.

43. 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.

44. 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.

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

Problem: The patrol scrub starting address registers [scrubaddresshi (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.

46. 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.

47. 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.

48. 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.

49. 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.

50. 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.

51. Intel® PT PSB+ Packets May Contain Unexpected Packets

Problem: Some Intel Processor Trace packets should be issued only between TIP.PGE (Target IP Packet.Packet Generation Enable) and TIP.PGD (Target IP Packet.Packet Generation Disable) packets. Due to this erratum, when a TIP.PGE packet is generated it may be preceded by a PSB+ (Packet Stream Boundary) that incorrectly includes Flow Update Packet (FUP) and MODE.Exec packets.

Implication: Due to this erratum, FUP and MODE.Exec may be generated unexpectedly.

Workaround: Decoders should ignore FUP and MODE.Exec packets that are not between TIP.PGE and TIP.PGD packets.

Status: No Fix.

52. 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.

53. Intel® PT ToPA PMI Does Not Freeze Performance Monitoring Counters

Problem: Due to this erratum, if IA32_DEBUGCTL.FREEZE_PERFMON_ON_PMI (MSR 1D9H, bit 12) is set to 1 when Intel PT triggers a Table of Physical Addresses (ToPA) PerfMon Interrupt (PMI), performance monitoring counters are not frozen as expected.



Implication: Performance monitoring counters will continue to count for events that occur during PMI handler execution.

Workaround: PMI handler software can programmatically stop performance monitoring counters upon entry.

Status: No Fix.

54. 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.

55. 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.

56. 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.

57. 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.

58. 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.

59. 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.

60. 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.



61. 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.

62. 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.

63. #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.

64. 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.

65. 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.



66. **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.

67. **Performance Monitor Event For Outstanding Offcore Requests and Snoop Requests May be Incorrect**

Problem: The performance monitor event OFFCORE_REQUESTS_OUTSTANDING (Event 60H, any Umask Value) should count the number of offcore outstanding transactions each cycle. Due to this erratum, the counts may be higher or lower than expected.

Implication: The performance monitor event OFFCORE_REQUESTS_OUTSTANDING may reflect an incorrect count.

Workaround: None identified.

Status: No Fix.

68. **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.

69. **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.



70. 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.

71. 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.



Specification Changes

There are no Specification Changes in this Specification Update revision.



Specification Clarifications

There are no Specification Clarifications in this Specification Update revision.



Documentation Changes

There are no Documentation Changes in this Specification Update revision.

§