

Second Generation Intel® Xeon® Scalable Processors

Specification Update

May 2019

Notice: The Second Generation Intel® Xeon® Scalable Processors 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: 338848-003US



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 © 2019, Intel Corporation. All Rights Reserved.



Contents

Revision History	4
Preface	5
Summary Tables of Changes	7
Identification Information	10
Errata	18
Specification Changes	25
Specification Clarifications	26
Documentation Changes	27



Revision History

Date	Revision	Description
April 2019	001	Initial Release (Intel Public).
April 2019	002	Added errata CLX18 and CLX19. Made clarifications to Turbo Frequency Tables.
May 2019	003	Added errata CLX20, CLX21, CLX22, CLX23, CLX24, CLX25 and CLX26



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>Second Generation Intel® Xeon® Scalable Processors Datasheet: Volume 1 - Electrical</i>	338845
<i>Second Generation Intel® Xeon® Scalable Processors Datasheet: Volume 2 - Registers</i>	338846

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

Number	Steppings			Status	Errata
	B-1	L-1	R-1		
CLX1.	x	x	x	No Fix	Intel® CAT/CDP Might Not Restrict Cacheline Allocation Under Certain Conditions (Intel® Xeon® Processor Scalable Family)
CLX2.	x	x	x	No Fix	Intel® PT PSB+ Packets May be Omitted on a C6 Transition
CLX3.	x	x	x	No Fix	IDI_MISC Performance Monitoring Events May be Inaccurate
CLX4.	x	x	x	No Fix	Intel® PT CYC Packets Can be Dropped When Immediately Preceding PSB
CLX5.	x	x	x	No Fix	Intel® PT VM-entry Indication Depends on The Incorrect VMCS Control Field
CLX6.	x	x	x	No Fix	Intel® MBA Read After MSR Write May Return Incorrect Value
CLX7.	x	x	x	No Fix	In eMCA2 Mode, When The Retirement Watchdog Timeout Occurs CATERR# May be Asserted
CLX8.	x	x	x	No Fix	VCVTPS2PH To Memory May Update MXCSR in The Case of a Fault on The Store
CLX9.	x	x	x	No Fix	Intel® PT May Drop All Packets After an Internal Buffer Overflow
CLX10.	x	x	x	No Fix	Non-Zero Values May Appear in ZMM Upper Bits After SSE Instructions
CLX11.	x	x	x	No Fix	ZMM/YMM Registers May Contain Incorrect Values
CLX12.	x	x	x	No Fix	When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions
CLX13.	x	x	x	No Fix	Intel® PT ToPA Tables Read From Non-Cacheable Memory During an Intel® TSX Transaction May Lead to Processor Hang
CLX14.	x	x	x	No Fix	Performing an XACQUIRE to an Intel® PT ToPA Table May Lead to Processor Hang
CLX15.	x	x	x	No Fix	Using Intel® TSX Instructions May Lead to Unpredictable System Behavior
CLX16.	x	x	x	No Fix	Reading Some C-state Residency MSRs May Result in Unpredictable System Behavior
CLX17.	x	x	x	No Fix	Performance in an 8sg System May Be Lower Than Expected
CLX18.	x	x	x	No Fix	Memory May Continue to Throttle after MEMHOT# De-assertion
CLX19.	x	x	x	No Fix	Unexpected Uncorrected Machine Check Errors May Be Reported
CLX20.	x	x	x	No Fix	CQM Counters May Decrement an Additional Time From During a FwdCode Flow
CLX21.	x	x	x	No Fix	Intel® MBM Counters May Double Count
CLX22.	x	x	x	No Fix	Intel® MBA May Incorrectly Throttle All Threads
CLX23.	x	x	x	No Fix	Setting Performance Monitoring IA32_PERF_GLOBAL_STATUS_SET MSR Bit 63 May Not #GP
CLX24.	x	x	x	No Fix	Branch Instruction Address May be Incorrectly Reported on TSX Abort When Using MPX
CLX25.	x	x	x	No Fix	x87 FDP Value May be Saved Incorrectly
CLX26.	x	x	x	No Fix	Intel® PT Trace May Drop Second Byte of CYC Packet

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 Second Generation Intel® Xeon® Scalable Processors 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 1, "Component Identification via registers"](#) on page 10 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. Component Identification via registers

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-1	Server, 2S	0x50657	1	1	1	0	1	1	1	
	B-1	Server, 4S	0x50657	1	1	1	1	0	1	1	
	B-1	Server, 8S	0x50657	1	1	1	1	1	1	1	
HCC	L-1	Server, 2S	0x50657	1	1	1	0	1	1	0	
	L-1	Server, 4S	0x50657	1	1	1	1	0	1	0	
LCC	R-1	Server, 2S	0x50657	1	1	1	0	1	0	0	



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. Second Generation Intel® Xeon® Scalable Processors Non Intel® AVX Turbo Frequencies

82xx, 62xx, and 52xx Processors

SKU	Cores	LLC (MB)	TDP (W)	# of active cores / maximum core frequency in turbo mode (GHz)																													
				Base																													
				non-AVX Core Freq. (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	
8280	28	38.5	205	2.7	4.0	4.0	3.8	3.8	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	3.5	3.5	3.5	3.5	3.3	3.3	3.3	3.3			
8276	28	38.5	165	2.2	4.0	4.0	3.8	3.8	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1	3.0	3.0	3.0		
8270	26	35.75	205	2.7	4.0	4.0	3.8	3.8	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	3.5	3.5	3.5	3.5	3.4	3.4					
8268	24	35.75	205	2.9	3.9	3.9	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	3.6	3.5	3.5	3.5	3.5							
8260	24	35.75	165	2.4	3.9	3.9	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	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1						
8256	4	16.5	105	3.8	3.9	3.9	3.9	3.9																									
8253	16	22	125	2.2	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.7	2.7	2.7	2.5	2.5	2.5															
6254	18	24.75	200	3.1	4.0	4.0	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9	3.9												
6252	24	35.75	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						
6248	20	27.5	150	2.5	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.2	3.2	3.2											
6246	12	24.75	165	3.3	4.2	4.2	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1																		
6244	8	24.75	150	3.6	4.4	4.4	4.3	4.3	4.3	4.3	4.3	4.3																					
6242	16	22	150	2.8	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.5	3.5	3.5															
6240	18	24.75	150	2.6	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.3	3.3												
6238	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								
6234	8	24.75	130	3.3	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0																					
6230	20	27.5	125	2.1	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8										
6226	12	19.25	125	2.7	3.7	3.7	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5																		
5222	4	16.5	105	3.8	3.9	3.9	3.9	3.9																									
5220	18	24.75	125	2.2	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.7	2.7												
5218	16	22	125	2.3	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	2.8	2.8	2.8	2.8														

SKU	Cores	LLC (MB)	TDP (W)	# of active cores / maximum core frequency in turbo mode (GHz)																												
				Base																												
				non-AVX Core Freq. (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
6262V	24	33	135	1.9	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	2.5					
6222V	20	27.5	115	1.8	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.4	2.4	2.4	2.4									
6238T	22	30.25	125	1.9	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8	2.7	2.7							
6230T	20	27.5	125	2.1	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8									
5220T	18	24.75	105	1.9	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.7	2.7											
5218T	16	22	105	2.1	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.0	3.0	3.0	2.7	2.7	2.7														
4209T	8	11	70	2.2	3.2	3.2	3.0	3.0	2.5	2.5	2.5	2.5																				
5220S	18	24.75	125	2.7	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.7	2.7											

- 8280, 8276, 8260, 6240 and 6138 have 2TB/socket and 4.5TB/socket memory capacity versions (8280M, 8280L, 8276M, 8276L, 8260M, 8260L, 6240M, 6240L, 6138M and 6138L) with identical frequencies.
- All details shown above are subject to change without notice.



Figure 2. Second Generation Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies

82xx, 62xx, and 52xx Processors

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Freq. (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
8280	28	38.5	205	2.2	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	3.3	3.3	3.3	3.0	3.0	3.0	2.9	2.9	2.9			
8276	28	38.5	165	1.7	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.6	2.6	2.6				
8270	26	35.75	205	2.2	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.9	2.9	2.9	2.9				
8268	24	35.75	205	2.4	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	3.2	3.2	3.2	3.0	3.0	3.0	3.0								
8260	24	35.75	165	1.9	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.0	3.0	3.0	2.7	2.7	2.7	2.6	2.6	2.6	2.6							
8256	4	16.5	105	3.3	3.7	3.7	3.7	3.7																								
8253	16	22	125	1.7	2.7	2.7	2.5	2.5	2.4	2.4	2.4	2.2	2.2	2.2	2.0	2.0	2.0	2.0														
6254	18	24.75	200	2.7	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	3.4	3.4													
6252	24	35.75	150	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.1	3.1	3.1	2.8	2.8	2.8	2.5	2.5	2.5	2.4	2.4	2.4	2.4								
6248	20	27.5	150	1.9	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.4	3.4	3.4	3.0	3.0	3.0	2.8	2.8	2.8	2.8											
6246	12	24.75	165	2.9	4.0	4.0	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8																	
6244	8	24.75	150	3.0	4.0	4.0	3.9	3.9	3.9	3.9	3.9																					
6242	16	22	150	2.3	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.4	3.4	3.4	3.1	3.1	3.1	3.1														
6240	18	24.75	150	2.0	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.2	3.2	3.2	2.9	2.9	2.9	2.8	2.8													
6238	22	30.25	140	1.7	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.1	3.1	3.1	2.8	2.8	2.8	2.5	2.5	2.5	2.5	2.5	2.5									
6234	8	24.75	130	2.8	3.9	3.9	3.7	3.7	3.7	3.7	3.7																					
6230	20	27.5	125	1.6	3.8	3.8	3.6	3.6	3.4	3.4	3.4	2.9	2.9	2.9	2.6	2.6	2.6	2.4	2.4	2.4												
6226	12	19.25	125	2.3	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.1	3.1	3.1	3.1																	
5222	4	16.5	105	3.3	3.8	3.8	3.8	3.8																								
5220	18	24.75	125	1.8	3.8	3.8	3.6	3.6	3.4	3.4	3.4	2.9	2.9	2.9	2.6	2.6	2.6	2.5	2.5													
5218	16	22	125	1.8	2.9	2.9	2.7	2.7	2.6	2.6	2.6	2.5	2.5	2.5	2.3	2.3	2.3	2.3														

SKU	Cores	LLC (MB)	TDP (W)	Base AVX 2.0 Core Freq. (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
6262V	24	33	135	1.6	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.0	3.0	3.0	2.8	2.8	2.8											
6222V	20	27.5	115	1.6	3.3	3.3	3.1	3.1	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.3	2.3	2.2	2.2	2.2												
6238T	22	30.25	125	1.5	3.6	3.6	3.4	3.4	3.2	3.2	3.2	2.7	2.7	2.7	2.4	2.4	2.4	2.2	2.2	2.2	2.2	2.2										
6230T	20	27.5	125	1.6	3.8	3.8	3.6	3.6	3.4	3.4	3.4	2.9	2.9	2.9	2.6	2.6	2.6	2.4	2.4	2.4												
5220T	18	24.75	105	1.5	3.8	3.8	3.6	3.6	3.4	3.4	3.4	2.9	2.9	2.9	2.6	2.6	2.6	2.5	2.5													
5218T	16	22	105	1.7	2.8	2.8	2.6	2.6	2.5	2.5	2.5	2.4	2.4	2.4	2.2	2.2	2.2															
4209T	8	11	70	2.1	3.0	3.0	2.7	2.7	2.1	2.1	2.1	2.1																				
5220S	18	24.75	125	1.8	3.8	3.8	3.6	3.6	3.4	3.4	3.4	2.9	2.9	2.9	2.6	2.6	2.6	2.5	2.5													

- 8280, 8276, 8260, 6240 and 6138 have 2TB/socket and 4.5TB/socket memory capacity versions (8280M, 8280L, 8276M, 8276L, 8260M, 8260L, 6240M, 6240L, 6138M and 6138L) with identical frequencies.
- All details shown above are subject to change without notice.



Figure 3. Second Generation Intel® Xeon® Scalable Processors Intel® AVX-512 Turbo Frequencies

82xx, 62xx, and 52xx Processors

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Freq. (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
8280	28	38.5	205	1.8	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	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	
8276	28	38.5	165	1.3	3.7	3.7	3.5	3.5	3.3	3.3	3.3	3.3	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.2	2.2	2.2	2.2	2.1	2.1	2.1	2.1	
8270	26	35.75	205	1.8	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.2	3.2	3.2	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.4	2.4	2.4			
8268	24	35.75	205	1.9	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	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					
8260	24	35.75	165	1.5	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.3	2.3	2.3	2.3					
8256	4	16.5	105	2.7	3.7	3.7	3.5	3.5																								
8253	16	22	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	1.6	1.6	1.6	1.6													
6254	18	24.75	200	2.2	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.9	2.9											
6252	24	35.75	150	1.3	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0					
6248	20	27.5	150	1.6	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5									
6246	12	24.75	165	2.4	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4																
6244	8	24.75	150	2.6	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5																				
6242	16	22	150	1.9	3.7	3.7	3.5	3.5	3.2	3.2	3.2	3.2	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5												
6240	18	24.75	150	1.6	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5										
6238	22	30.25	140	1.3	3.6	3.6	3.4	3.4	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1	2.1	2.1	2.1						
6234	8	24.75	130	2.3	3.7	3.7	3.5	3.5	3.1	3.1	3.1	3.1																				
6230	20	27.5	125	1.1	3.7	3.7	3.5	3.5	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0									
6226	12	19.25	125	1.9	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6																
5222	4	16.5	105	2.7	3.7	3.7	3.5	3.5																								
5220	18	24.75	125	1.4	3.7	3.7	3.5	3.5	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.1	2.1	2.1	2.1	2.1	2.1	2.1										
5218	16	22	125	1.5	2.9	2.9	2.7	2.7	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1												

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Freq. (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
6262V	24	33	135	1.1	3.2	3.2	3.0	3.0	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.2	2.2	2.2	2.2	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9				
6222V	20	27.5	115	1.1	3.0	3.0	2.8	2.8	2.5	2.5	2.5	2.5	2.1	2.1	2.1	2.1	1.9	1.9	1.9	1.9	1.8	1.8	1.8	1.8								
6238T	22	30.25	125	1.1	3.5	3.5	3.3	3.3	2.6	2.6	2.6	2.6	2.2	2.2	2.2	2.2	2.0	2.0	2.0	2.0	1.8	1.8	1.8	1.8	1.8	1.8						
6230T	20	27.5	125	1.1	3.7	3.7	3.5	3.5	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0								
5220T	18	24.75	105	1.1	3.7	3.7	3.5	3.5	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										
5218T	16	22	105	1.3	2.8	2.8	2.6	2.6	2.5	2.5	2.5	2.5	2.2	2.2	2.2	2.2	2.0	2.0	2.0	2.0												
4209T	8	11	70	1.2	2.0	2.0	1.8	1.8	1.5	1.5	1.5	1.5																				
5220S	18	24.75	125	1.4	3.7	3.7	3.5	3.5	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										

- 8280, 8276, 8260, 6240 and 6138 have 2TB/socket and 4.5TB/socket memory capacity versions (8280M, 8280L, 8276M, 8276L, 8260M, 8260L, 6240M, 6240L, 6138M and 6138L) with identical frequencies.
- All details shown above are subject to change without notice.




Figure 4. Second Generation Intel® Xeon® Scalable Processors Non Intel® AVX Turbo Frequencies

52xx, 42xx, and 32xx 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
5220	18	24.75	125	2.2	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.7	2.7										
5218	16	22	125	2.3	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.8												
5217	8	11	115	3	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4																				
5215	10	13.75	85	2.5	3.4	3.4	3.2	3.2	3.1	3.1	3.1	3.1	3.0	3.0																		
4216	16	22	100	2.1	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7													
4215	8	11	85	2.5	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0																				
4214	12	16.5	85	2.2	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.7	2.7	2.7																	
4210	10	13.75	85	2.2	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.7	2.7																		
4208	8	11	85	2.1	3.2	3.2	3.0	3.0	2.5	2.5	2.5	2.5																				
3204	6	8.25	85	1.9	1.9	1.9	1.9	1.9																								

- 5215 has 2TB/socket and 4.5TB/socket memory capacity versions (5215M and 5215L) with identical frequencies.
- 4214 has an Intel® Speed Select Technology option (4214Y) with identical frequencies.
- All details shown above are subject to change without notice.

Figure 5. Second Generation Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies

52xx, 42xx, and 32xx 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
5220	18	24.75	125	1.8	3.8	3.8	3.6	3.6	3.4	3.4	3.4	3.4	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.5	2.5										
5218	16	22	125	1.8	2.9	2.9	2.7	2.7	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5	2.3	2.3	2.3	2.3												
5217	8	11	115	2.5	3.5	3.5	3.3	3.3	3.0	3.0	3.0	3.0																				
5215	10	13.75	85	2	3.1	3.1	2.9	2.9	2.8	2.8	2.8	2.8	2.6	2.6																		
4216	16	22	100	1.4	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.5	2.5	2.5	2.3	2.3	2.3	2.3														
4215	8	11	85	2	3.3	3.3	3.1	3.1	2.6	2.6	2.6	2.6																				
4214	12	16.5	85	1.8	3.1	3.1	2.9	2.9	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4																
4210	10	13.75	85	1.9	3.0	3.0	2.8	2.8	2.5	2.5	2.5	2.5	2.3	2.3																		
4208	8	11	85	1.6	3.0	3.0	2.6	2.6	2.0	2.0	2.0	2.0																				
3204	6	8.25	85	1.5	1.5	1.5	1.5	1.5																								

- 5215 has 2TB/socket and 4.5TB/socket memory capacity versions (5215M and 5215L) with identical frequencies.
- 4214 has an Intel® Speed Select Technology option (4214Y) with identical frequencies.



- All details shown above are subject to change without notice.

Figure 6. Second Generation Intel® Xeon® Scalable Processors Intel® AVX-512 Turbo Frequencies

52xx, 42xx, and 32xx 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
5220	18	24.75	125	1.4	3.7	3.7	3.5	3.5	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										
5218	16	22	125	1.5	2.9	2.9	2.7	2.7	2.6	2.6	2.6	2.6	2.3	2.3	2.3	2.3	2.1	2.1	2.1	2.1												
5217	8	11	115	2.0	2.9	2.9	2.7	2.7	2.4	2.4	2.4	2.4																				
5215	10	13.75	85	1.4	2.9	2.9	2.5	2.5	1.9	1.9	1.9	1.9	1.8	1.8																		
4216	16	22	100	1.1	2.0	2.0	1.8	1.8	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.7	1.6	1.6	1.6	1.6												
4215	8	11	85	1.5	2.3	2.3	2.1	2.1	2.0	2.0	2.0	2.0																				
4214	12	16.5	85	1.3	2.0	2.0	1.8	1.8	1.7	1.7	1.7	1.7	1.6	1.6	1.6	1.6																
4210	10	13.75	85	1.2	2.0	2.0	1.8	1.8	1.6	1.6	1.6	1.6	1.5	1.5																		
4208	8	11	85	1.1	2.0	2.0	1.8	1.8	1.4	1.4	1.4	1.4																				
3204	6	8.25	85	1.0	1.0	1.0	1.0	1.0	1.0	1.0																						

- 5215 has 2TB/socket and 4.5TB/socket memory capacity versions (5215M and 5215L) with identical frequencies.
- 4214 has an Intel® Speed Select Technology option (4214Y) with identical frequencies.
- All details shown above are subject to change without notice.

Figure 7. Second Generation Intel® Xeon® Scalable Processors Non Intel® AVX Turbo Frequencies

N and U 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
6252N	24	35.75	150	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	3.3	3.2	3.2	3.2	3.2	3.0	3.0	3.0	3.0				
6230N	20	27.5	125	2.3	3.5	3.5	3.3	3.3	3.2	3.2	3.2	3.2	3.2	3.2	3.2	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9									
5218N	16	22	105	2.3	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0												
6212U	24	35.75	165	2.4	3.9	3.9	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	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1					
6210U	20	27.5	150	2.5	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.2								
6209U	20	27.5	125	2.1	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8								



- All details shown above are subject to change without notice.

Figure 8. Second Generation Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies

N and U 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
6252N	24	35.75	150	1.8	3.5	3.5	3.3	3.3	3.2	3.2	3.2	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	2.7	2.7					
6230N	20	27.5	125	1.6	3.4	3.4	3.2	3.2	3.1	3.1	3.1	3.1	3.1	3.1	3.1	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.6									
5218N	16	22	105	1.6	2.9	2.9	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8	2.8													
6212U	24	35.75	165	1.9	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	3.0	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.6	2.6	2.6	2.6			
6210U	20	27.5	150	1.9	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8								
6209U	20	27.5	125	1.6	3.8	3.8	3.6	3.6	3.4	3.4	3.4	3.4	3.4	2.9	2.9	2.9	2.9	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4							


- All details shown above are subject to change without notice.

Figure 9. Second Generation Intel® Xeon® Scalable Processors Intel® AVX-512 Turbo Frequencies

N and U Processors

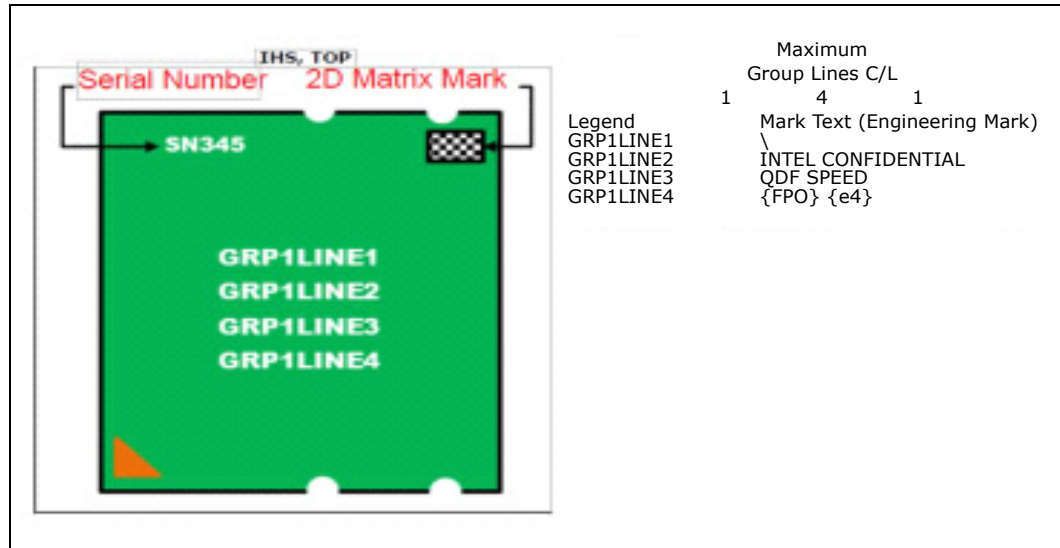
SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
6252N	24	35.75	150	1.4	3.4	3.4	3.2	3.2	3.1	3.1	3.1	3.1	3.1	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.3	2.3	2.3	2.3			
6230N	20	27.5	125	1.2	3.4	3.4	3.2	3.2	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								
5218N	16	22	105	1.2	2.9	2.9	2.7	2.7	2.6	2.6	2.6	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5													
6212U	24	35.75	165	1.5	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.0	3.0	3.0	3.0	2.6	2.6	2.6	2.6	2.4	2.4	2.4	2.4	2.3	2.3	2.3	2.3			
6210U	20	27.5	150	1.6	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.0	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5								
6209U	20	27.5	125	1.1	3.7	3.7	3.5	3.5	2.8	2.8	2.8	2.8	2.4	2.4	2.4	2.4	2.1	2.1	2.1	2.1	2.0	2.0	2.0	2.0								

- All details shown above are subject to change without notice.



Component Marking Information

Figure 10. Processor Preliminary Top Side Marking (Example)



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



Errata

CLX1. 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.

Status: No Fix.

CLX2. 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.

CLX3. 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.

CLX4. 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.



CLX5. 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_CTLSS2, 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_CTLSS 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.

CLX6. 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.

CLX7. 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.

CLX8. 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.

**CLX9. 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.

CLX10. 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.

CLX11. 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.

CLX12. 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 PDPTE 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.

CLX13. 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.

CLX14. 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.

CLX15. 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.

CLX16. 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.

CLX17. 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 workaround for this erratum

Status: No fix.

CLX18. 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.

CLX19. 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.

Status: No fix.

CLX20. CQM Counters May Decrement an Additional Time From During a FwdCode Flow

Problem: It is possible during a FwdCode flow that the CQM (Cache Quality Monitoring) counter may be decremented an additional time. This scenario would not result in a less than 0 counter.

Implication: Due to this erratum, CQM counters may be lower than expected.

Workaround: None identified.

Status: No fix.

CLX21. Intel® MBM Counters May Double Count

Problem: The MBM (Memory Bandwidth Monitoring) counters (accessible via the IA32_QM_EVTSEL / IA32_QM_CTR MSR pair) may double count when NT (Non-Temporal) writes are used or in remote socket cases. The performance counters in the IMC (integrated memory controller) are not affected and can report the read and write memory bandwidths.

Implication: For workloads utilizing NT operations the MBM accuracy may be reduced, which can affect performance monitoring or bandwidth-aware scheduling software.

Workaround: None identified. This erratum can be mitigated by using the IMC performance monitoring counters or per-core performance monitoring counters to derive a read/write ratio or per-core statistics that can be used to adjust the MBM counters.

Status: No fix.

CLX22. 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.

CLX23. 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.

CLX24. 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 (Intel® 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 Intel® Processor Trace (Intel® 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.

CLX25. 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.



CLX26. Intel® PT Trace May 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.



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.

§

