



2nd Gen Intel[®] Xeon[®] Scalable Processors

Specification Update

July 2021



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.

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see <http://www.intel.com/technology/turboboost>.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

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, Pentium, Intel Optane and the Intel logo are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

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

Copyright © 2021, 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
July 2021	019	Added new errata CLX59.
June 2021	018	Added new errata CLX58.
May 2021	017	Added new errata CLX57.
April 2021	016	Added new errata CLX55 and CLX56.
March 2021	NA	Out of cycle.
February 2021	015	Added new erratum CLX54.
January 2021	014	Added new errata CLX47, CLX48, CLX49, CLX50, CLX51, CLX52, and CLX53.
December 2020	NA	Out of cycle.
November 2020	NA	Out of cycle.
October 2020	NA	Out of cycle.
September 2020	012/013	Added new erratum CLX46.
July 2020	011	Added new erratum CLX45.
June 2020	010	Added new errata CLX43 and CLX44. Out of cycle.
May 2020	009	Added new erratum CLX42. Made little fixes about title "Base" in column on figures 14 and 15.
April 2020	008	Added new errata CLX41. Added new section "Refresh Processors - Non-Intel® AVX, Intel® AVX, and Intel® AVX-512 Turbo Frequencies" on page 19 Added New Turbo Frequencies to Refresh Processors Figure 13, "2nd Gen Intel® Xeon® Scalable Processors Non-Intel® AVX Turbo Frequencies" on page 19 Added New Turbo Frequencies to Refresh Processors Figure 14, "2nd Gen Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies" on page 20 Added New Turbo Frequencies to Refresh Processors Figure 15, "2nd Gen Intel® Xeon® Scalable Processors Intel® AVX 512 Turbo Frequencies" on page 21
March 2020	007	Added new errata CLX37, CLX38, CLX39 and CLX40.
December 2019	006	Remove CLX15. Updated CLX11. Add a new erratum and numbered it as CLX15.
November 2019	005	Added errata CLX33, CLX34, CLX35 and CLX36.
September 2019	004	Added errata CLX27, CLX28, CLX29, CLX30, CLX31 and CLX32. Updated Turbo Frequency Tables.
May 2019	003	Added errata CLX20, CLX21, CLX22, CLX23, CLX24, CLX25 and CLX26
April 2019	002	Added errata CLX18 and CLX19. Made clarifications to Turbo Frequency Tables.
April 2019	001	Initial Release (Intel Public).

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>2nd Gen Intel® Xeon® Scalable Processors Datasheet: Volume 1 - Electrical</i>	338845
<i>2nd Gen 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>Intel® 64 and IA-32 Architecture Software Developer Manual, Volume 2A: Instruction Set Reference, A-M</i>	253666 ¹
<i>Intel® 64 and IA-32 Architecture Software Developer Manual, Volume 2B: Instruction Set Reference, N-Z</i>	253667 ¹
<i>Intel® 64 and IA-32 Architecture Software Developer Manual, Volume 3A: System Programming Guide, Part 1</i>	253668 ¹
<i>Intel® 64 and IA-32 Architecture Software Developer Manual, 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, and so on).

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 use the following notations:

Codes Used in Summary Tables

Stepping

X:	Errata exist 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 the left of table row indicates this erratum is either new or modified from the previous version of the document.



Errata (Sheet 1 of 2)

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	PCIe Root Port Does Not Increment REPLAY_NUM on Multiple NAKs of The Same TLP
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 Intel® 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
CLX27.	x	x	x	No fix	Intel® Speed Select Base Configuration P1 Frequency May Not be Selectable
CLX28.	x	x	x	No fix	IMC Patrol Scrubbing Engine May Hang
CLX29.	x	x	x	No fix	Intel® MBM Counters May Report System Memory Bandwidth Incorrectly
CLX30.	x	x	x	No fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

Errata (Sheet 2 of 2)

Number	Steppings			Status	Errata
	B-1	L-1	R-1		
CLX31.	x	x	x	No fix	Voltage/Frequency Curve Transitions May Result in Machine Check Errors or Unpredictable System Behavior
CLX32.	x	x	x	No fix	Processor May Behave Unpredictably on Complex Sequence of Conditions Which Involve Branches That Cross 64 Byte Boundaries
CLX33.	x	x	x	No fix	STIBP May Not Function as Intended
CLX34.	x	x	x	No fix	Intel® UPI, DMI and PCIe Interfaces May See Elevated Bit Error Rates
CLX35.	x	x	x	No fix	Unexpected Page Faults in Guest Virtualization Environment
CLX36.	x	x	x	No fix	Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation
CLX37.	x	x	x	No fix	Memory Controller May Hang While in Virtual Lockstep
CLX38.	x	x	x	No fix	MD_CLEAR Operations May Not Properly Overwrite All Buffers
CLX39.	x	x	x	No fix	ITD Algorithm May Not Select Correct Operating Voltage
CLX40.	x	x	x	No fix	Direct Branches With Partial Address Aliasing May Lead to Unpredictable System Behavior
CLX41.	x	x	x	No fix	Runtime Patch Load Enables Processor Capabilities That May Cause Performance Degradation
CLX42.	x	x	x	No fix	Performance Monitoring General Counter 2 May Have Invalid Value Written When TSX Is Enabled
CLX43.	x	x	x	No fix	Intel® QuickData Technology Engine May Hang With Any DMA Error if Completion Status is Improperly Set
CLX44.	x	x	x	No fix	Overflow Flag in MSR May be Incorrectly Set
CLX45.	x	x	x	No fix	A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes
CLX46.	x	x	x	No fix	A Fixed Interrupt May Be Lost When a Core Exits C6
CLX47.	x	x	x	No fix	A TOR TO may be seen with DDR4 16Gb 2DPC 3DS LRDIMMs
CLX48.	x	x	x	No fix	Certain Errors in Device 16 of a VLS Region Report Device 0 as the Failed Device
CLX49.	x	x	x	No fix	Memory errors in a VLS region on a certain device may not be properly corrected
CLX50.	x	x	x	No fix	Processor May Hang if Warm Reset Triggers During BIOS Initialization
CLX51.	x	x	x	No fix	When in CPGC Mode With Memory Refresh Disabled, DDR Scheduler May be Blocked From Issuing CPGC Commands
CLX52.	x	x	x	Fixed	High Levels of Posted Interrupt Traffic on the PCIe Port May Result in a Machine Check With a TOR Timeout
CLX53.	x	x	x	Fixed	Some Short Loops of Instructions May Cause a 3-Strike Machine Check Without a TOR Timeout
CLX54.	x	x	x	Fixed	Processor May Fail to Retry a Write to DDRT Memory
CLX55.	x	x	x	No Fix	Retried PECI PCI Config Local Register Accesses May Not Operate Correctly
CLX56.	x	x	x	Fixed	MD_CLEAR Operations May Not Properly Overwrite All Buffers
CLX57.	x	x	x	Fixed	An Incorrect Instruction Pointer May Be Reported for a REP MOVS Instruction
CLX58.	x	x	x	Fixed	Intel® Optane™ Persistent Memory Mode or Mixed Mode May Cause a Hang
CLX59.	x	x	x	No Fix	Systems Using Intel® Optane™ Persistent Memory in Mixed Mode May Experience a System Hang or Reset



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 2nd Gen 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](#) 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. 2nd Gen 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	3.7	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	3.6	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	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	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	3.6	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	3.6	3.5	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	3.6	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	3.4	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	3.4	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	2.8	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	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	2.6	2.6	2.6	2.6	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	2.7	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	3.4	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	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	3.0	2.7	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	3.1	2.8	2.8	2.8	2.8	2.7	2.7										

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



Figure 3. 2nd Gen 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	2.5	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	2.2	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	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	2.6	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	3.0	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	2.0	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	2.5	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	2.6	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								
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												
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	2.2	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	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	2.0	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	2.1	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	2.1	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	2.0	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	2.1	2.1	2.1	2.1	2.1	2.1													

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

Figure 4. 2nd Gen 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	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	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	1.9																							

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

Figure 5. 2nd Gen 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	2.7	2.5	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	1.5																							

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

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

Figure 8. 2nd Gen Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies

N and U Processors					# 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
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	3.2	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	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	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	3.3	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	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 previously shown are subject to change without notice.

Figure 9. 2nd Gen Intel® Xeon® Scalable Processors Intel® AVX-512 Turbo Frequencies

N and U Processors					# 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
6252N	24	35.75	150	1.4	3.4	3.4	3.2	3.2	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	3.0	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 previously shown are subject to change without notice.



Figure 10. Intel® Xeon® W-3200 Processors Non-Intel® AVX Turbo Frequencies

SKU	Cores	LLC (MB)	TDP (W)	Base non-AVX Core Frequency (GHz)	ITBM	# 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
W-3275	28	38.5	205	2.5	4.6	4.4	4.4	4.2	4.2	4.1	4.1	4.1	4.1	4.1	4.1	4.1	3.9	3.9	3.9	3.6	3.6	3.6	3.6	3.3	3.3	3.3	3.3	3.2	3.2	3.2			
W-3265	24	33	205	2.7	4.6	4.4	4.4	4.2	4.2	4.1	4.1	4.1	4.1	4.1	4.1	3.9	3.9	3.9	3.6	3.6	3.6	3.6	3.4	3.4	3.4								
W-3245	16	22	205	3.2	4.6	4.4	4.4	4.2	4.2	4.1	4.1	4.1	4.1	4.1	3.9	3.9	3.9																
W-3235	12	19.25	180	3.3	4.5	4.4	4.4	4.2	4.2	4.1	4.1	4.1	4.1	4.0	4.0	4.0																	
W-3225	8	16.5	160	3.7	4.4	4.3	4.3	4.2	4.2	4.2	4.2																						
W-3223	8	16.5	160	3.5	4.2	4.0	4.0	3.8	3.8	3.8	3.8																						

- The W-3275, W-3265 and W-3245 have 2 TB/socket memory capacity versions (W-3275M, W-3265M and W-3245M) with identical frequencies
- ITBM = Intel® Turbo Boost Max Technology 3.0

Figure 11. Intel® Xeon® W-3200 Processors Intel® AVX 2.0 Turbo Frequencies

SKU	Cores	LLC (MB)	TDP (W)	Base AVX2.0 Core Frequency (GHz)	ITBM	# 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
W-3275	28	38.5	205	2.1	N/A	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.3	3.3	3.3	3.3	3.0	3.0	3.0	2.8	2.8	2.8	2.7	2.7	2.7				
W-3265	24	33	205	2.2	N/A	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.3	3.3	3.3	3.3	3.1	3.1	3.1	2.9	2.9	2.9								
W-3245	16	22	205	2.8	N/A	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.2	3.2	3.2																
W-3235	12	19.25	180	3	N/A	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5																			
W-3225	8	16.5	160	3.3	N/A	3.8	3.8	3.8	3.8	3.8	3.8																						
W-3223	8	16.5	160	3	N/A	3.8	3.8	3.6	3.6	3.5	3.5	3.5																					

- The W-3275, W-3265 and W-3245 have 2 TB/socket memory capacity versions (W-3275M, W-3265M and W-3245M) with identical frequencies
- ITBM = Intel® Turbo Boost Max Technology 3.0

Figure 12. Intel® Xeon® W-3200 Processors Intel® AVX-512 Turbo Frequencies

SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	ITBM	# 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
W-3275	28	38.5	205	1.6	N/A	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.0	3.0	3.0	2.7	2.7	2.7	2.5	2.5	2.5	2.3	2.3	2.3	2.2	2.2	2.2	2.2				
W-3265	24	33	205	1.8	N/A	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.1	3.1	3.1	2.8	2.8	2.8	2.5	2.5	2.5	2.4	2.4	2.4									
W-3245	16	22	205	2.3	N/A	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.1	3.1	3.1	2.8	2.8	2.8															
W-3235	12	19.25	180	2.5	N/A	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.0	3.0	3.0																		
W-3225	8	16.5	160	2.8	N/A	3.6	3.6	3.5	3.5	3.5	3.5																						
W-3223	8	16.5	160	2.5	N/A	3.3	3.3	3.1	3.1	3.0	3.0	3.0																					

- The W-3275, W-3265 and W-3245 have 2 TB/socket memory capacity versions (W-3275M, W-3265M and W-3245M) with identical frequencies
- ITBM = Intel® Turbo Boost Max Technology 3.0

Refresh Processors - Non-Intel® AVX, Intel® AVX, and Intel® AVX-512 Turbo Frequencies

Figure 13. 2nd Gen Intel® Xeon® Scalable Processors Non-Intel® AVX Turbo Frequencies

62xx, 52xx, 42xx and 32xx Processors					# of active cores / maximum core frequency in turbo mode (GHz)																													
SKU	Cores	LLC (MB)	TDP (W)	Base	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		
				non-AVX Core Freq. (GHz)																														
6256	12	33	205	3.6	4.5	4.5	4.3	4.3	4.3	4.3	4.3	4.3	4.3	4.3	4.3																			
6250	8	35.75	185	3.9	4.5	4.5	4.5	4.5	4.5	4.5	4.5	4.5																						
6208U					# of active cores / maximum core frequency in turbo mode (GHz)																													
SKU	Cores	LLC (MB)	TDP (W)	Base	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		
				non-AVX Core Freq. (GHz)																														
6208U	16	22	150	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													
4210T	10	13.75	95	2.3	3.2	3.2	3.0	3.0	2.7	2.7	2.7	2.7	2.5	2.5																				
6258R	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	3.7	3.7	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.4	3.4	
6248R	24	35.75	205	3.0	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	3.7	3.6	3.6	3.6	3.6							
6246R	16	35.75	205	3.4	4.1	4.1	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0	4.0													
6242R	20	35.75	205	3.1	4.1	4.1	3.9	3.9	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8									
6240R	24	35.75	165	2.4	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	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.2					
6238R	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	3.0	3.0	
6230R	26	35.75	150	2.1	4.0	4.0	3.8	3.8	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	3.0	3.0				
6226R	16	22	150	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														
5220R	24	35.75	150	2.2	4.0	4.0	3.8	3.8	3.7	3.7	3.7	3.7	3.7	3.7	3.7	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9							
5218R	20	27.5	125	2.1	4.0	4.0	3.8	3.8	3.7	3.7	3.7	3.7	3.5	3.5	3.5	3.5	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9										
4215R	8	11	130	3.2	4.0	4.0	3.8	3.8	3.6	3.6	3.6	3.6																						
4214R	12	16.5	100	2.4	3.5	3.5	3.3	3.3	3.2	3.2	3.2	3.2	3.0	3.0	3.0	3.0																		
4210R	10	13.75	100	2.4	3.2	3.2	3.0	3.0	2.9	2.9	2.9	2.9	2.9	2.9																				
3206R	8	11	85	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9																						

Note: 6250 has a large memory (4.5 TB/Socket) capacity version (6250L) with identical frequencies

Figure 14. 2nd Gen Intel® Xeon® Scalable Processors Intel® AVX 2.0 Turbo Frequencies

62xx, 52xx, 42xx and 32xx Processors					# of active cores / maximum core frequency in turbo mode (GHz)																												
SKU	Cores	LLC (MB)	TDP (W)	Base																													
				AVX2.0 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	
6256	12	33	205	3.2	4.0	4.0	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8																	
6250	8	35.75	185	3.6	4.0	4.0	4.0	4.0	4.0	4.0	4.0																						

					# of active cores / maximum core frequency in turbo mode (GHz)																																		
SKU	Cores	LLC (MB)	TDP (W)	Base																																			
				AVX2.0 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							
6258R	28	38.5	205	2.3	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	3.5	3.5	3.3	3.3	3.3	3.3	3.3	3.3	3.0	3.0	3.0	3.0	2.9	2.9	2.9	2.9	
6248R	24	35.75	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	3.5	3.3	3.3	3.3	3.3	3.1	3.1	3.1	3.1								
6246R	16	35.75	205	3.0	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																			
6242R	20	35.75	205	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	3.5	3.5	3.5	3.5	3.3	3.3	3.3	3.3														
6240R	24	35.75	165	1.9	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	3.2	3.2	3.2	3.2	3.0	3.0	3.0	3.0											
6238R	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	3.5	3.5	3.5	3.5	3.5	3.2	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7	2.6	2.6	2.6	2.6			
6230R	26	35.75	150	1.7	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	3.5	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						
6226R	16	22	150	2.3	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1																		
6208U	16	22	150	2.3	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1																		
5220R	24	35.75	150	1.8	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	3.2	3.2	3.2	3.2	3.0	3.0	3.0	3.0	2.8	2.8	2.8	2.8								
5218R	20	27.5	125	1.7	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.5	3.2	3.2	3.2	2.9	2.9	2.9	2.9	2.9	2.7	2.7	2.7	2.7														
4215R	8	11	130	2.0	3.6	3.6	3.3	3.3	2.6	2.6	2.6	2.6																											
4214R	12	16.5	100	2.1	3.1	3.1	2.9	2.9	2.8	2.8	2.8	2.8	2.7	2.7	2.7	2.7																							
4210R	10	13.75	100	2.0	3.0	3.0	2.8	2.8	2.7	2.7	2.7	2.7	2.5	2.5																									
4210T	10	13.75	95	1.9	3.0	3.0	2.8	2.8	2.6	2.6	2.6	2.6	2.4	2.4																									
3206R	8	11	85	1.8	1.8	1.8	1.8	1.8	1.8	1.8	1.8	1.8																											

Note: 6250 has a large memory (4.5 TB/Socket) capacity version (6250L) with identical frequencies

Figure 15. 2nd Gen Intel® Xeon® Scalable Processors Intel® AVX 512 Turbo Frequencies

62xx, 52xx, 42xx and 32xx Processors

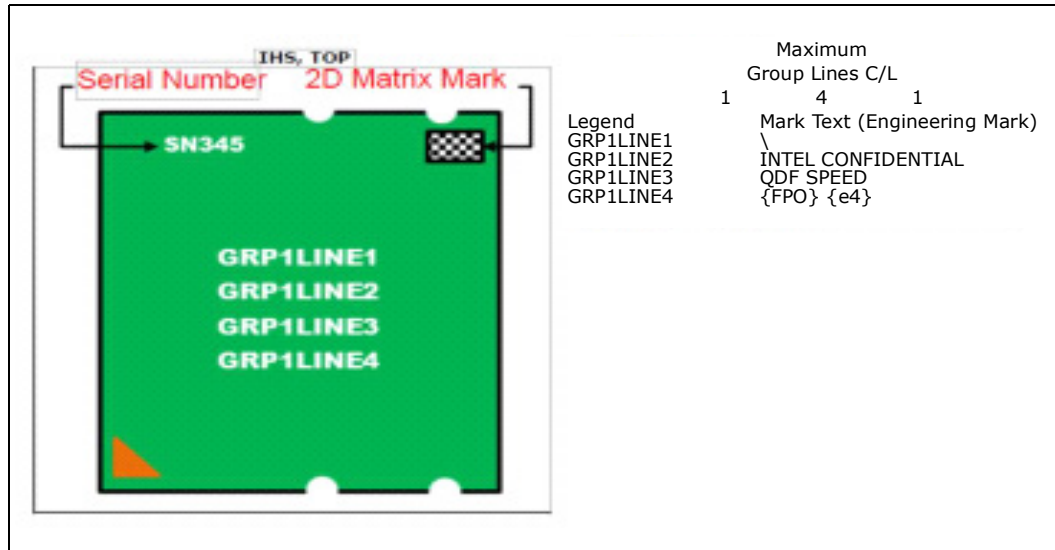
SKU	Cores	LLC (MB)	TDP (W)	Base AVX-512 Core Frequency (GHz)	# of active cores / maximum core frequency in turbo mode (GHz)																											
					1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
6256	12	33	205	2.7	3.8	3.8	3.6	3.6	3.5	3.5	3.5	3.5	3.3	3.3	3.3																	
6250	8	35.75	185	3.1	3.8	3.8	3.8	3.8	3.8	3.8	3.8	3.8																				

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
6258R	28	38.5	205	1.8	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	2.8	2.8	2.8	2.8	2.6	2.6	2.6	2.6	2.5	2.5	2.5	2.5	
6248R	24	35.75	205	2	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					
6246R	16	35.75	205	2.5	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1														
6242R	20	35.75	205	2.2	3.7	3.7	3.5	3.5	3.4	3.4	3.4	3.4	3.4	3.4	3.1	3.1	3.1	3.1	2.9	2.9	2.9	2.9										
6240R	24	35.75	165	1.5	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					
6238R	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	2.1	2.1	2.1	2.1	2.1	
6230R	26	35.75	150	1.3	3.6	3.6	3.4	3.4	3.3	3.3	3.3	3.3	3.0	3.0	3.0	2.7	2.7	2.7	2.7	2.5	2.5	2.5	2.5	2.3	2.3	2.3	2.3	2.3				
6226R	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												
6208U	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												
5220R	24	35.75	150	1.4	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	2.5	2.5	2.5	2.5	2.4	2.4	2.4	2.4				
5218R	20	27.5	125	1.1	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	2.3	2.3	2.3	2.3									
4215R	8	11	130	1.5	3.3	3.3	2.6	2.6	2.0	2.0	2.0	2.0																				
4214R	12	16.5	100	1.6	2.3	2.3	2.1	2.1	2.0	2.0	2.0	2.0	1.9	1.9	1.9	1.9																
4210R	10	13.75	100	1.4	2.2	2.2	2.0	2.0	1.8	1.8	1.8	1.8	1.7	1.7																		
4210T	10	13.75	95	1.2	3.0	3.0	2.8	2.8	2.3	2.3	2.3	2.3	2.2	2.2																		
3206R	8	11	85	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5	1.5																				

Note: 6250 has a large memory (4.5 TB/Socket) capacity version (6250L) with identical frequencies

Component Marking Information

Figure 16. Processor Preliminary Top Side Marking (Example)



For the 2nd Gen 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 that can be 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® Processor Trace (Intel® PT) Packet Stream Boundary+ (PSB+) 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 Packet Stream Boundary (PSB) packet can cause a single Cycle Count (CYC) packet, possibly along with an associated Mini Time Counter (MTC) 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 Paging Information Packet (PIP), which includes indication of entry into non-root operation, will be generated on VM-entry as long as the "Conceal VMX in Intel® PT" field (bit 19) in Secondary Execution Control register (IA32_VMX_PROCBASED_CTLDS2, MSR 048BH) is clear. This diverges from expected behavior, since this PIP should instead be generated only with a zero value of the "Conceal VMX entries from Intel® PT" field (Bit 17) in the Entry Control register (IA32_VMX_ENTRY_CTLDS MSR 0484H).

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

Workaround: A Virtual Machine Monitor (VMM) 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 Memory Bandwidth Allocation (MBA) 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 (for example, #PF) or VM exit. The value written to the MXCSR exceptions flag 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 Table of Physical Addresses (ToPA) entry transition can cause an internal buffer overflow that may result in all trace packets, including the Overflow (OVF) packet, being dropped.

Implication: When this erratum occurs, all trace data will be lost until either Intel® 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 four 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 Intel® AVX512 instruction writes to ZMM0-15. This erratum will not occur in software that uses VZEROUPPER between Intel® 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 Intel® 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: None identified.

Status: No fix.

CLX12. When Virtualization Exceptions are Enabled, EPT Violations May Generate Erroneous Virtualization Exceptions

Problem: An access to a Guest-Physical Address (GPA) may cause an Extended Page Table (EPT)-violation VM exit. When the "EPT-violation #VE" VM-execution control is 1, an EPT violation may cause a virtualization exception (#VE) 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 an 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 ToPA table is placed in Uncacheable (UC) or Uncacheable Speculative Write Combining (USWC) memory, and a ToPA output region is filled during an Intel® Transaction Synchronization (Intel® TSX) 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 ToPA table, and that table is later read by the CPU during the Hardware Lock Elision (HLE) 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. PCIe Root Port Does Not Increment REPLAY_NUM on Multiple NAKs of The Same TLP

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

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

Workaround: None identified.

Status: No fix.

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 8 sg (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, 8 sg 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 microarchitectural 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 Cache Quality Monitoring (CQM) 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 Non-Temporal (NT) writes are used or in remote socket cases. The performance counters in the Integrated Memory Controller (IMC) 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 Memory Bandwidth Allocation (MBA) feature may select an incorrect MBA throttling value to apply to the core. A disabled logical processor may behave as though the 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 Intel® MPX

Problem: When using Intel® 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 Intel® 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® PT. Due to this erratum, the FROM_IP field in LBR/BTS/BTM, as well as the FUP source IP address that corresponds to the Intel® 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 Intel® TSX may operate incorrectly when Intel® 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.

CLX27. Intel® Speed Select Base Configuration P1 Frequency May Not be Selectable

Problem: To configure Intel® Speed Select (ISS), BIOS may program FLEX_RATIO MSR to select the target ratio for ISS Configuration 1 or Configuration 2. Programming FLEX_RATIO[15:8] for ISS precludes the ability to retrieve the Base Configuration frequency information.

Implication: If Intel® ISS Configuration 1 or Configuration 2 is selected, BIOS will not be able to discover the base frequency P1 for Base Configuration from the processor.

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

Status: No fix.

CLX28. IMC Patrol Scrubbing Engine May Hang

Problem: Under rare microarchitectural conditions, the processor's IMC Patrol Scrubbing Engine may hang.

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

Workaround: None identified.

Status: No fix.

CLX29. Intel® MBM Counters May Report System Memory Bandwidth Incorrectly

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

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

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

Status: No fix.

CLX30. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

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

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

Workaround: None identified.

Status: No fix.

CLX31. Voltage/Frequency Curve Transitions May Result in Machine Check Errors or Unpredictable System Behavior

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

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



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

Status: No fix.

CLX32. Processor May Behave Unpredictably on Complex Sequence of Conditions Which Involve Branches That Cross 64 Byte Boundaries

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

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

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

Status: No fix.

CLX33. STIBP May Not Function as Intended

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

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

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

Status: No fix.

CLX34. Intel® UPI, DMI and PCIe Interfaces May See Elevated Bit Error Rates

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

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

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

Status: No fix.

CLX35. Unexpected Page Faults in Guest Virtualization Environment

Problem: Under complex microarchitectural conditions, a virtualized guest could observe unpredictable system behavior.

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

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

Status: No fix.

CLX36. Instruction Fetch May Cause Machine Check if Page Size Was Changed Without Invalidation

Problem: This erratum may cause a machine-check error (IA32_MCi_STATUS.MCACOD=005H with IA32_MCi_STATUS.MSCOD=00FH or IA32_MCi_STATUS.MCACOD=0150H with IA32_MCi_STATUS.MSCOD=00FH) on the fetch of an instruction. It applies only if (1)

instruction bytes are fetched from a linear address translated using a 4-Kbyte page and cached in the processor; (2) the paging structures are later modified so that these bytes are translated using a large page (2-Mbyte, 4-Mbyte or 1-Gbyte) with a different Physical Address (PA), memory type (PWT, PCD and PAT bits), or User/Supervisor (U/S) bit; and (3) the same instruction is fetched after the paging structure modification but before software invalidates any TLB entries for the linear region.

Implication: Due to this erratum an unexpected machine check with error code 0150H with MSCOD 00FH may occur, possibly resulting in a shutdown. This erratum could also lead to unexpected correctable machine check (IA32_MCI_STATUS.UC=0) with error code 005H with MSCOD 00FH.

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

Status: No fix.

CLX37. Memory Controller May Hang While in Virtual Lockstep

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

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

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

Status: No fix.

CLX38. MD_CLEAR Operations May Not Properly Overwrite All Buffers

Problem: On processors that enumerate the IA32_ARCH_CAPABILITIES.TSX_CTRL MSR bit and are affected by TSX Asynchronous Abort (TAA), the VERW mem instruction should overwrite affected buffers with constant data. On processors also affected by this erratum, VERW may not overwrite upper store buffer data at byte offsets 32-63 of each entry, and may not overwrite upper load port data at byte offsets 32-63 of each port. This behavior may also occur on other MD_CLEAR operations, which overwrite microarchitectural structures: specifically the L1D_FLUSH command, and RSM.

Implication: Software using MD_CLEAR operations to prevent TAA side channel methods from revealing previous accessed data may not prevent those side channel methods from inferring the value of the upper bytes of preceding vector loads or stores.

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

Status: No fix.

CLX39. ITD Algorithm May Not Select Correct Operating Voltage

Problem: Implementation of Inverse Temperature Dependency (ITD) compensation may exhibit incorrect voltage compensation under specific voltage and temperature conditions.

Implication: Due to this erratum, unpredictable system behavior may occur. This erratum has only been observed in a synthetic testing environment at Intel. Intel has not observed this erratum in any commercially available system.

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

Status: No fix.

CLX40. Direct Branches With Partial Address Aliasing May Lead to Unpredictable System Behavior

Problem: Under complex microarchitectural conditions involving direct branch instructions with partial address aliasing, unpredictable system behavior may occur. Intel has only seen this under synthetic testing conditions. Intel has not observed this under any commercially available software.

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

Workaround: None identified.

Status: No fix.

CLX41. Runtime Patch Load Enables Processor Capabilities That May Cause Performance Degradation

Problem: When loading certain microcode updates, some processor capabilities may be inadvertently enabled as part of the patch load procedure. Enabling these capabilities may cause a performance degradation on certain workloads.

Implication: When this erratum occurs, the process may exhibit unexpected performance degradation. There are no functional implications to this erratum.

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

Status: No fix.

CLX42. Performance Monitoring General Counter 2 May Have Invalid Value Written When TSX Is Enabled

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

Implication: Software may read invalid value from IA32_PMC2.

Workaround: None identified.

Status: No fix.

CLX43. Intel® QuickData Technology Engine May Hang With Any DMA Error if Completion Status is Improperly Set

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

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

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

Status: No fix.

CLX44. Overflow Flag in MSR May be Incorrectly Set

Problem: Under internal processor conditions, a single internal parity error seen in IA32_MC0_STATUS MSR (401h) with MCACOD (bits 15:0) value of 5h and MSCOD (bits31:16) value of 7h, may set the overflow flag (bit 62) in the same MSR.

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

Workaround: None identified.

Status: No fix.

CLX45. A Pending Fixed Interrupt May Be Dispatched Before an Interrupt of The Same Priority Completes

Problem: Resuming from C6 Sleep-State, with Fixed Interrupts of the same priority queued the processor may dispatch the second interrupt before the first interrupt has completed and written to the EOI register, causing the first interrupt to never complete.

Implication: Due to this erratum, Software may behave unexpectedly.

Workaround: None identified.

Status: No fix.

CLX46. A Fixed Interrupt May Be Lost When a Core Exits C6

Problem: Under complex microarchitectural conditions, when performance throttling happens during a core C6 exit, a fixed Interrupt may be lost.

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

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

Status: No fix.

CLX47. A TOR TO may be seen with DDR4 16Gb 2DPC 3DS LRDIMMs

Problem: The values for thrt_crit and thrt_hi limits may not be updated for DDR4 16 Gb 3 DS LRDIMM 2 DPC. This may lead to an incorrect interpretation of these values, which can lead to a subsequent TOR TO after Memory Reference Code (MRC) is done.

Implication: Due to this erratum, A TOR TO may be seen.

Workaround: None identified.

Status: No fix. For 2nd Gen Intel® Xeon® Processor Scalable Family was disabled 2X refresh for both 1DPC and 2 DPC 16 Gb 3DSL RDIMM and 16 Gb 3 DS RDIMM configurations.

CLX48. Certain Errors in Device 16 of a VLS Region Report Device 0 as the Failed Device

Problem: Under complex microarchitectural conditions, when Adaptive Data Correction (ADC) or Adaptive Double Device Data Correction (ADDDC) is enabled; and is in Virtual Lockstep (VLS) mode, then if a limited subset of multi bit errors is detected on primary device 16 in the VLS region, then specific memory controller register logs 0 instead of the actual failing device 16.

Implication: System Software that takes action based on register log may implicate the incorrect device.

Workaround: A workaround for this erratum is available in BIOS. Alternatively, software may detect primary DRAM Device corrected errors.

Status: No fix.

CLX49. Memory errors in a VLS region on a certain device may not be properly corrected

Problem: Under complex microarchitectural conditions, when ADC or ADDDC is enabled; and the system has spared out a DRAM device 0, 1, 3, 4, 5, 8, 13, 15 or 16; and is in VLS mode, then if a limited subset of multi bit errors is detected on primary DRAM device 16 in the VLS region, the error may not be properly corrected.

Implication: The system may experience unpredictable system behavior. Intel has only observed this under synthetic testing conditions.

Workaround: None identified.

CLX50. Processor May Hang if Warm Reset Triggers During BIOS Initialization

Problem: Under complex microarchitectural conditions, when the processor receives a warm reset during BIOS initialization, the processor may hang with a machine check error reported in IA32_MCi_STATUS register.

Implication: Due to this erratum, the processor may hang. Intel has only observed this erratum in a synthetic test environment.

Workaround: None identified.

CLX51. When in CPGC Mode With Memory Refresh Disabled, DDR Scheduler May be Blocked From Issuing CPGC Commands

Problem: When memory refresh is disabled during Converged Pattern Generation and Checking (CPGC) mode, the Integrated Memory Controller (IMC) scheduler may become blocked from issuing CPGC read and write commands.

Implication: Due to this erratum, a system hang or continuous restart may occur.

Workaround: None identified.

CLX52. High Levels of Posted Interrupt Traffic on the PCIe Port May Result in a Machine Check With a TOR Timeout

Problem: High levels of posted interrupt traffic on the PCIe* port may lead to a TOR Timeout Machine Check Exception (MSCOD=000Ch, MCACOD = "Cache Hierarchy Errors") in bank IA32_MC9_STATUS (MSR 425h), IA32_MC10_STATUS (MSR 429h), or IA32_MC11_STATUS (MSR 42Dh).

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

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

Status: Fixed.

CLX53. Some Short Loops of Instructions May Cause a 3-Strike Machine Check Without a TOR Timeout

Problem: Under complex microarchitectural conditions, some short loops of instructions may cause a three-strike machine check [(MSCOD=80h, MCACOD=0400h) logged into IA32_MC3_STATUS (MSR 40Dh)] without a TOR timeout. This can only happen when both logical processors on the same physical processor are active.

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

Workaround: It is possible for the BIOS to contain a fix for this erratum (WW18 available).

Status: Fixed.

CLX54. Processor May Fail to Retry a Write to DDRT Memory

Problem: If a system DDR4 memory module signifies a transaction error, such as a Command Address Parity error or Write CRC error, the processor may fail to retry a write to DDRT memory. This erratum can only occur in memory configurations that have both DDR4 and DDRT memory in the same channel.

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

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

Status: Fixed.

CLX55. Retried PECI PCI Config Local Register Accesses May Not Operate Correctly

Problem: When the processor requests a PECI PCI Config Local Read or Write command to be retried, and the PECI host immediately retries the command (within 150 us), the processor may fail to correctly process the retried PECI command.

Implication: Due to this erratum, the PECE PCICongigLocal Read command may return incorrect data and the PECE PCICongigLocal Write command may incorrectly update the target.

Workaround: None Identified.

Status: No Fix.

CLX56. MD_CLEAR Operations May Not Properly Overwrite All Buffers

Problem: On processors that enumerate the MD_CLEAR CPUID, L1D_FLUSH, RSM, and VERW memory instructions should overwrite affected buffers with constant data. Under complex micro-architectural conditions, these instructions may not overwrite all affected buffers.

Implication: Due to this erratum, the use of MD_CLEAR operations to prevent MDS (Microarchitectural Data Sampling) or TAA (Intel® Transactional Synchronization Extensions Asynchronous Abort) side-channel methods from revealing previously accessed data may not be fully effective.

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

Status: Fixed.

CLX57. An Incorrect Instruction Pointer May Be Reported for a REP MOVS Instruction

Problem: When a REP MOVS instruction reports a Software Recoverable Action Required (SRAR) error for memory that software did not intend to access, the instruction pointer is reported incorrectly on the thread.

Implication: Due to this erratum, the processor may report SRAR errors with an incorrect instruction pointer.

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

Status: Fixed.

CLX58. Intel® Optane™ Persistent Memory Mode or Mixed Mode May Cause a Hang

Problem: When the processor is utilizing Intel® Optane™ Persistent Memory Mode or Mixed Mode, a Core MMIO read request may incorrectly block a write request from the IO subsystem, leading to a system hang.

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

Workaround: It is possible for the BIOS to contain a fix for this erratum. A Purley Refresh BIOS (610.D03).

Status: Fixed.

CLX59. Systems Using Intel® Optane™ Persistent Memory in Mixed Mode May Experience a System Hang or Reset

Problem: When the processor is utilizing Intel® Optane™ persistent memory in Mixed Mode with both App Direct and Memory Mode Snoopy Modes enabled, a system hang or reset may occur when running a workload.

Implication: Due to this erratum, the system may reset or hang, logging an Internal Timer Error or logging a Table of Requests (TOR) timeout.

Workaround: None identified. Systems can avoid this erratum by not using Mixed Mode with both App Direct and Memory Mode Snoopy Modes enabled.

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.

§