Intel® Xeon® E-2300 Processor Product Family

Datasheet, Volume 1 of 2

October 2021

Revision 001

Document Number: 662318

Contents

Figures

Tables

Revision History

Document Number	Revision Number	Description	Revision Date
662318	001	• Initial release	October 2021

1 Introduction

Intel® Xeon® E-2300 processor product family are 64-bit, multi-core processors built on 14-nanometer process technology.

The processor line is offered in a 2-Chip Platform that includes Intel® C250 Series Chipset Families Platform Controller Hub (PCH). Refer to Figure 1-1.

The following table describes the processor lines covered in this document.

Table 1-1. Processor Lines

Processor Line[1]	Package	Base TDP	Processor IA Cores	GT Cores	Platform Type
Intel® Xeon® E-2300 processor	LGA1200	95W	8	32	2-Chip[2]
		95W	6	32	
		80W	8	32	
		80W	6	32	
		80W	4	32	
		65W	8	N/A	
		65W	6	N/A	
		65W	4	32/0	

Notes:
1. Processor Lines offering may change.
2. The Intel® Xeon® E-2300 processor product family SKUs are paired with the Intel® C250 Series Platform Controller Hub (PCH).

Throughout this document, the Intel® Xeon® E-2300 processor product family may be referred to simply as "processor". The Intel® C250 Series Chipset Family Platform Controller Hub (PCH) may be referred to simply as "PCH".

Table 1-2. Intel® Xeon® E-2300 Processor Product Family SKUs

Processor Number	Cache Size	IA Cores	GT Cores	Base Core Freq. (GHz)	Turbo 1-Core Freq. (GHz)	Turbo 2 Core Freq. (GHz)	Turbo 3-Core Freq. (GHz)	Turbo 4-Core Freq. (GHz)	Turbo 5-Core Freq. (GHz)	Turbo 6-Core Freq. (GHz)	Turbo 7-Core Freq. (GHz)	Turbo 8-Core Freq. (GHz)	Thermal Design Power (W)
E-2388G	16 MB	8	32	3.2	5.1	5.0	5.0	4.9	4.9	4.8	4.7	4.6	95
E-2378G	16 MB	8	32	2.8	5.1	4.9	4.9	4.8	4.8	4.7	4.6	4.6	80
E-2378	16 MB	8	0	2.6	4.8	4.7	4.7	4.6	4.6	4.5	4.5	4.5	65
E-2386G	12 MB	6	32	3.5	5.1	5.0	4.9	4.9	4.8	4.7			95
E-2356G	12 MB	6	32	3.2	5.0	4.9	4.8	4.8	4.8	4.8			80
E-2336	12 MB	6	0	2.9	4.8	4.7	4.7	4.6	4.6	4.6			65
E-2374G	8 MB	4	32	3.7	5.0	4.9	4.9	4.9					80
E-2334	8 MB	4	0	3.4	4.8	4.7	4.6	4.6					65
E-2324G	8 MB	4	32	3.1	4.6	4.6	4.5	4.5					65
E-2314	8 MB	4	0	2.8	4.5	4.2	3.8	3.5					65

Figure 1-1. Intel® Xeon® E-2300 Processor Line Platforms Overview

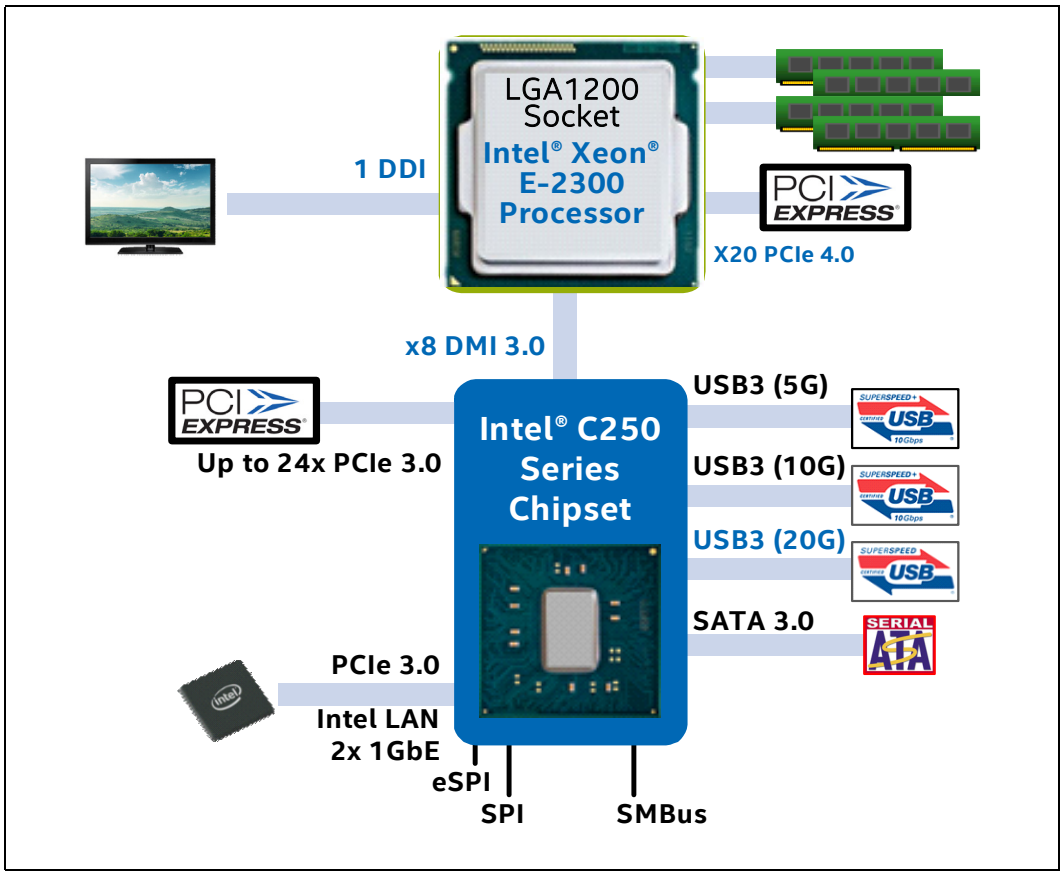

1.1 Processor Volatility Statement

The Processor do not retain any end user data when powered down and/or when the processor is physically removed.

Note: Power down refers to state which all processor power rails are off.

1.2 Supported Technologies

- PECI – Platform Environmental Control Interface
- SMEP – Supervisor Mode Execution Protection
- SMAP – Supervisor Mode Access Protection
- Intel® Virtualization Technology (Intel® VT)
  - Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)
  - Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)
  - Intel® APIC Virtualization Technology (Intel® APICv)
- Intel® Active Management Technology 11.0 (Intel® AMT 11.0)

- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2)
- Intel® Hyper-Threading Technology (Intel® HT Technology)
- Intel® 64 Architecture
- Execute Disable Bit
- Intel® Turbo Boost Technology 2.0
- Intel® Turbo Boost Max Technology 3.0
- Intel® Advanced Vector Extensions 2 (Intel® AVX2)
- Advanced Vector Extensions 512 Bit (Intel® AVX-512)
- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
- PCLMULQDQ (Perform Carry-Less Multiplication Quad word) Instruction
- Intel® Secure Key
- PAIR – Power Aware Interrupt Routing
- SMEP – Supervisor Mode Execution Protection
- Intel® Boot Guard
- Intel® Processor Trace
- High Definition Content Protection (HDCP)
- SHA Extensions – Secure Hash Algorithm Extensions
- UMIP – User Mode Instruction Prevention
- RDPID – Read Processor ID
- KeyLocker Technology
- Intel® Smart Cache Technology
- IA Core Level 1 and Level 2 Caches
- Intel® Hyper-Threading Technology (Intel® HT Technology)
- Intel® SpeedStep® Technology
- Intel® Speed Shift Technology
- Intel® 64 Architecture x2APIC
- Intel® Dynamic Tuning (Intel® Dynamic Platform and Thermal Framework (DTT (Previously DPTF))
- Intel® GNA 2.0 (GMM and Neural Network Accelerator)
- Cache Line Write Back (CLWB)

Note: The availability of the features may vary between processor SKUs.

Refer to Chapter 3, "Technologies" for more information.

1.3 API Support (Windows*)

- Direct3D* 2015, Direct3D 12, Direct3D 11.2, Direct3D 11.1, Direct3D 9, Direct3D10, Direct2D
- OpenGL* 4.5

- Open CL* 3.0, OpenCL 2.1, Open CL 2.0, Open CL 1.2
- DirectX* extensions: PixelSync, Instant Access, Conservative Rasterization, Render Target Reads, Floating-point De-norms, Shared a Virtual memory, Floating Point atomics, MSAA sample-indexing, Fast Sampling (Coarse LOD), Quilted Textures, GPU Enqueue Kernels, GPU Signals processing unit. Other enhancements include color compression.

Gen 12 architecture delivers hardware acceleration of Direct X* 12 Render pipeline comprising of the following stages: Vertex Fetch, Vertex Shader, Hull Shader, Tessellation, Domain Shader, Geometry Shader, Rasterizer, Pixel Shader, Pixel Output.

1.4 Power Management Support

1.4.1 Processor Core Power Management

- Full support of ACPI C-states as implemented by the following processor C-states:
  - C0, C1, C1E, C6, C7, C8, C9
- Enhanced Intel SpeedStep® Technology

Note: Refer to Section 4.2 for more information.

1.4.2 System Power Management

- S0/S0ix, S3, S4, S5

Refer to Chapter 4, "Power Management" for more information.

1.4.3 Memory Controller Power Management

- Disabling Unused System Memory Outputs
- DRAM Power Management and Initialization
- Initialization Role of CKE
- Conditional Self-Refresh
- Dynamic Power Down
- DRAM I/O Power Management
- DDR Electrical Power Gating (EPG)
- Power training

Refer to Section 4.3 for more information.

1.4.4 Processor Graphics Power Management

1.4.4.1 Memory Power Savings Technologies

- Intel® Rapid Memory Power Management (Intel® RMPM)
- Intel® Smart 2D Display Technology (Intel® S2DDT)

1.4.4.2 Display Power Savings Technologies

- Intel® (Seamless and Static) Display Refresh Rate Switching (DRRS) with eDP* port
- Intel® Automatic Display Brightness
- Smooth Brightness
- Intel® Display Power Saving Technology (Intel® DPST 6)
- Low Power Single Pipe (LPSP)

1.4.4.3 Graphics Core Power Savings Technologies

- Intel® Graphics Dynamic Frequency
- Intel® Graphics Render Standby Technology (Intel® GRST)
- Dynamic FPS (Intel® DFPS)

Refer to Section 4.4 for more information.

1.5 Thermal Management Support

- Digital Thermal Sensor
- Intel Adaptive Thermal Monitor
- THERMTRIP# and PROCHOT# support
- On-Demand Mode
- Memory closed Loop Throttling
- Memory Thermal Throttling
- External Thermal Sensor (TS-on-DIMM and TS-on-Board)
- Render Thermal Throttling
- Fan speed control with DTS
- Intel® Turbo Boost Technology 2.0 Power Control
- Intel® Dynamic Tuning - DTT (DPTF)

Refer to Chapter 5, "Thermal Management" for more information.

1.6 Processor Testability

Intel® DCI – DfX Trace Options Type A	Intel® DCI – DfX Trace Options Type C
DCI USB 3.0 Debug	DCI USB 3.0 Debug
DCI USB 2.0 Debug	DCI USB 2.0 Debug
DCI OOB (BSSB 4wire)	DCI OOB (BSSB 4wire)
2 wire DCI OOB	2 wire DCI OOB
Notes: 1. DCI USB2.0 Debug provides general support down to S0, Sx, S0iX 2. DCI OOB and 2-wire DCI OOB support equivalent early boot access 3. JTAG Pin Support needed to debug low power states and potential silicon issues	

Note: When separate XDP connectors will be used at C8–C10 states, the processor will need to be waked up using the PCH.

The processor includes boundary-scan for board and system level testability.

1.7 Operating System Support

Processor Line	Windows* 10 64-bit	Windows* Server 2016/2019/ 20H2	OS X	Linux* OS	Chrome* OS
Intel® Xeon® E-2300 (Server platform)	No	Yes	No	Yes	No

1.8 Terminology

Term	Description
4K	Ultra High Definition (UHD)
AES	Advanced Encryption Standard
AGC	Adaptive Gain Control
BLT	Block Level Transfer
BPP	Bits per Pixel
CDR	Clock and Data Recovery
CTLE	Continuous Time Linear Equalizer
DDI	Digital Display Interface for DP or HDMI*/DVI
DDR4	Dual Data Rate 4. Fourth-Generation Double Data Rate SDRAM Memory Technology RS - Reduced Standby Power
DFE	Decision Feedback Equalizer
DMA	Direct Memory Access
DMI	Direct Media Interface
Dynamic Tuning /DPTF	Dynamic Tuning
DP*	DisplayPort*
DTS	Digital Thermal Sensor
ECC	Error Correction Code - used to fix DDR transactions errors
eDP*	embedded DisplayPort*

Term	Description
EU	Execution Unit in the Processor Graphics
GSA	Graphics in System Agent
HDCP	High-bandwidth Digital Content Protection
HDMI*	High Definition Multimedia Interface
IMC	Integrated Memory Controller
Intel® 64 Technology	64-bit Memory Extensions to the IA-32 Architecture
Intel® DPST	Intel® Display Power Saving Technology
Intel® PTT	Intel® Platform Trust Technology
Intel® TXT	Intel® Trusted Execution Technology
Intel® VT	Intel® Virtualization Technology. Processor Virtualization, when used in conjunction with Virtual Machine Monitor software, enables multiple, robust independent software environments inside a single platform.
Intel® VT-d	Intel® Virtualization Technology (Intel® VT) for Directed I/O. Intel® VT-d is a hardware assist, under system software (Virtual Machine Manager or OS) control, for enabling I/O device Virtualization. Intel® VT-d also brings robust security by providing protection from errant DMAs by using DMA remapping, a key feature of Intel® VT-d.
IOV	I/O Virtualization
ISP	Image Signal Processor
LFM	Low Frequency Mode. corresponding to the Enhanced Intel SpeedStep® Technology's lowest voltage/frequency pair. It can be read at MSR CEh [47:40].
LLC	Last Level Cache
LPM	Low-Power Mode.The LPM Frequency is less than or equal to the LFM Frequency. The LPM TDP is lower than the LFM TDP as the LPM configuration limits the processor to single thread operation
LPSP	Low-Power Single Pipe
LSF	Lowest Supported Frequency.This frequency is the lowest frequency where manufacturing confirms logical functionality under the set of operating conditions.
MCP	Multi Chip Package - includes the processor and the PCH. In some SKUs it might have additional On-Package Cache.
MFM	Minimum Frequency Mode. MFM is the minimum ratio supported by the processor and can be read from MSR CEh [55:48].
MLC	Mid-Level Cache
NCTF	Non-Critical to Function. NCTF locations are typically redundant ground or non-critical reserved balls/lands, so the loss of the solder joint continuity at end of life conditions will not affect the overall product functionality.
PCH	Platform Controller Hub. The chipset with centralized platform capabilities including the main I/O interfaces along with display connectivity, audio features, power management, manageability, security, and storage features. The PCH may also be referred as "chipset".
PECI	Platform Environment Control Interface
PEG	PCI Express* (PCIe*) Graphics
PL1, PL2, PL3	Power Limit 1, Power Limit 2, Power Limit 3
Processor	The 64-bit multi-core component (package)
Processor Core	The term "processor core" refers to Si die itself, which can contain multiple execution cores. Each execution core has an instruction cache, data cache, and 256-KB L2 cache. All execution cores share the LLC.
Processor Graphics	Intel® Processor Graphics
PSR	Panel Self-Refresh

Term	Description
Rank	A unit of DRAM corresponding to four to eight devices in parallel, ignoring ECC. These devices are usually, but not always, mounted on a single side of a SODIMM.
SCI	System Control Interrupt. SCI is used in the ACPI protocol.
SDP	Scenario Design Power.
SHA	Secure Hash Algorithm
SSC	Spread Spectrum Clock
Storage Conditions	Refer Section 8.3, "Package Storage Specifications"
STR	Suspend to RAM
TAC	Thermal Averaging Constant
TCC	Thermal Control Circuit
TDP	Thermal Design Power
TOB	Tolerance Budget
TTV TDP	Thermal Test Vehicle TDP
V_{CC}	Processor Core Power Supply
V_{CCGT}	Processor Graphics Power Supply
V_{CCIO}	I/O Power Supply
V_{CCSA}	System Agent Power Supply
VCCIN_AUX	Vccin_aux Power Supply
V_{CCST}	Vcc Sustain Power Supply
V_{DDQ}	DDR Power Supply
VLD	Variable Length Decoding
VPID	Virtual Processor ID
V_{SS}	Processor Ground

1.9 Related Documents

Document	Document Number / Location
Intel® Xeon® E-2300 Processor Family Datasheet, Volume 2	662319
Intel® Xeon® E-2300 Processor Family Specification Update	679628
Advanced Configuration and Power Interface 3.0	http://www.acpi.info/
DDR4 Specification	http://www.jedec.org
High Definition Multimedia Interface specification revision 1.4	http://www.hdmi.org/manufacturer/specification.aspx
Embedded DisplayPort* Specification revision 1.4	http://www.vesa.org/vesa.standards/
DisplayPort* Specification revision 1.2	http://www.vesa.org/vesa.standards/
PCI Express* Base Specification Revision 4.0	http://www.pcisig.com/specifications
Intel® 64 and IA-32 Architectures Software Developer's Manuals	http://www.intel.com/products/processor/manuals/index.htm

2 Interfaces

2.1 System Memory Interface

- Two channels of DDR4 memory with a maximum of two DIMMs per channel. DDR technologies, number of DIMMs per channel, number of ranks per channel are SKU dependent.
- UDIMM support
- Single-channel and dual-channel memory organization modes
- Data burst length of eight for all memory organization modes
- DDR4 I/O Voltage of 1.2 V
- 64-bit wide channels
- ECC/Non-ECC UDIMM DDR4 support
- ECC is supported by Intel® Xeon® E-2300 server processor SKUs
- Theoretical maximum memory bandwidth of:
  - 50 GB/s in dual-channel mode assuming 3200 MT/s

Note: If the Intel® Xeon® E-2300 processor memory interface is configured to one DIMM per Channel, the processor can use either of the DIMMs, DIMM0 or DIMM1, signals CTRL[1:0] or CTRL[3:2].

2.1.1 System Memory Technology Supported

The Integrated Memory Controller (IMC) supports DDR4 protocols with two independent, 64-bit wide channels.

Table 2-1. Processor DDR Memory Speed Support

Processor Line	DDR4 1DPC[1] [MT/s]	DDR4 2DPC[1] [MT/s]
Intel® Xeon® E-2300 processor	3200[3]	3200[2,4]

Notes:
1. 1DPC refers to when only 1 DIMM slot per Channel is routed. 2DPC refers to when 2 DIMM slots per Channel are routed and are fully populated or partially populated with 1 DIMM only.
2. DDR4 2DPC is supported when channel is populated with the same DIMM part number.
3. DDR4 1DPC UDIMM:
   a. DDR4 1DPC 6 Layer board supports 3200 Gear2, 2933 Gear1 and 2666 Gear1 with Single Rank DIMMs and Dual Rank DIMMs.
4. DDR4 2DPC UDIMM:
   a. DDR4 2DPC 6 Layer board supports 3200 Gear2, 2933 Gear1 and 2666 Gear1 with 0R1R (single rank DIMM populated, other slot empty within the same channel), 1R1R (two single rank DIMM's populated in the same channel), 0R2R (Dual rank DIMM populated, other slot empty within the same channel).
   b. DDR4 2DPC 6 Layer board supports 2933 Gear1 and 2666 Gear1 with 2R2R (two dual rank DIMM's populated in the same channel),
5. **Gear 1**: CPU Memory Controller and Memory Speed are equal. **Gear 2**: CPU Memory Controller operates at half the Memory Speed (that is, CPU Memory Controller is at 1600 MHz while Memory Speed is at 3200 MHz when operating as Gear 2).

- DDR4 UDIMM Modules:
  - Standard 8 GB and 16 GB technologies and addressing are supported for x8 and x16 devices.

There is no support for memory modules with different technologies or capacities on opposite sides of the same memory module. If one side of a memory module is populated, the other side is either identical or empty.

2.1.1.1 DDR4 Supported Memory Modules and Devices

Table 2-2. Supported DDR4 Non-ECC UDIMM Module Configurations

Raw Card Version	DIMM Capacity	DRAM Device Technology	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/ Col Address Bits	# of Banks Inside DRAM	Page Size
A	8 GB	8 GB	1024M x 8	8	1	16/10	16	8K
A	16 GB	16 GB	2048M x 8	8	1	17/10	16	8K
C	4 GB	8 GB	512M x 16	4	1	16/10	8	8K
C	8 GB	16 GB	1024M x 16	4	1	17/10	8	8K
B	16 GB	8 GB	1024M x 8	16	2	16/10	16	8K
B	32 GB	16 GB	2048M x 8	16	2	17/10	16	8K

Table 2-3. Supported DDR4 ECC UDIMM Module Configurations

Raw Card Version	DIMM Capacity	DRAM Device Technology	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/ Col Address Bits	# of Banks Inside DRAM	Page Size
D	8 GB	8 GB	1024M x 8	8	1	16/10	16	8K
D	16 GB	16 GB	2048M x 8	8	1	17/10	16	8K
E	16 GB	8 GB	1024M x 8	16	2	16/10	16	8K
E	32 GB	16 GB	2048M x 8	16	2	17/10	16	8K

2.1.2 System Memory Timing Support

The IMC supports the following DDR Speed Bin, CAS Write Latency (CWL), and command signal mode timings on the main memory interface:

- tCL = CAS Latency
- tRCD = Activate Command to READ or WRITE Command delay
- tRP = PRECHARGE Command Period
- CWL = CAS Write Latency
- Command Signal modes:
  - 1N indicates a new DDR4 command may be issued every clock
  - 2N indicates a new DDR4 command may be issued every two clocks

Table 2-4. DRAM System Memory Timing Support (DDR4)

DRAM Device	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRP (ns)	CWL (tCK)	DPC (SoDIMM Only)	CMD Mode
DDR4	3200	22	13.75	13.75	9/10/11/ 12/14/16/ 18/20	1 or 2	2N

2.1.3 System Memory Organization Modes

The IMC supports two memory organization modes, single-channel and dual-channel. Depending upon how the DDR Schema and DIMM Modules are populated in each memory channel, a number of different configurations can exist.

Single-Channel Mode

In this mode, all memory cycles are directed to a single channel. Single-Channel mode is used when either the Channel A or Channel B DIMM connectors are populated in any order, but not both.

Dual-Channel Mode – Intel® Flex Memory Technology Mode

The IMC supports Intel® Flex Memory Technology Mode. Memory is divided into a symmetric and asymmetric zone. The symmetric zone starts at the lowest address in each channel and is contiguous until the asymmetric zone begins or until the top address of the channel with the smaller capacity is reached. In this mode, the system runs with one zone of dual-channel mode and one zone of single-channel mode, simultaneously, across the whole memory array.

Note: Channels A and B can be mapped for physical channel 0 and 1 respectively or vice versa. However, channel A size should be greater or equal to channel B size.

Figure 2-1. Intel® Flex Memory Technology Operations

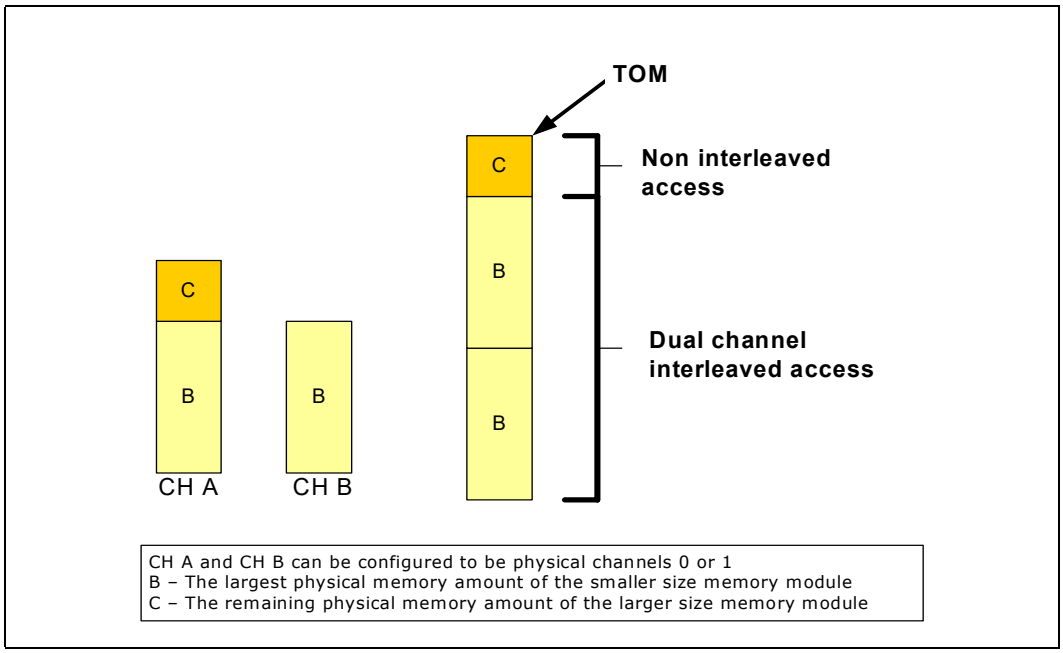

Dual-Channel Symmetric Mode (Interleaved Mode)

Dual-Channel Symmetric mode, also known as interleaved mode, provides maximum performance on real world applications. Addresses are ping-ponged between the channels after each cache line (64 byte boundary). If there are two requests, and the second request is to an address on the opposite channel from the first, that request can be sent before data from the first request has returned. If two consecutive cache lines are requested, both may be retrieved simultaneously, since they are ensured to be on opposite channels. Use Dual-Channel Symmetric mode when both Channel A and Channel B DIMM connectors are populated in any order, with the total amount of memory in each channel being the same.

When both channels are populated with the same memory capacity and the boundary between the dual channel zone and the single channel zone is the top of memory, IMC operates completely in Dual-Channel Symmetric mode.

Note: The DRAM device technology and width may vary from one channel to the other.

2.1.4 DIMM Slots Population Rule of Intel® Xeon® E-2300 server platform

The system memory controller supports up to two DIMM connectors per channel.

The frequency of system memory will be the lowest frequency of all the DIMMs placed in the system, as determined through the SPD registers on the DIMMs. DIMMs with a different parameter (different DIMM part number) cannot be installed on different slots either the same or cross memory channels, DDR4 2DPC is supported only when channel is populated with the same DIMM part number.

Note: Mixed DIMM part numbers on either the same or cross memory channels are not supported on Intel® Xeon® E-2300 server platform, and mixed DIMM configurations may reduce maximum achievable memory speed.

2.1.5 Technology Enhancements of Intel® Fast Memory Access (Intel® FMA)

The following sections describe the Just-in-Time Scheduling, Command Overlap, and Out-of-Order Scheduling Intel® FMA technology enhancements.

Just-in-Time Command Scheduling

The memory controller has an advanced command scheduler where all pending requests are examined simultaneously to determine the most efficient request to be issued next. The most efficient request is picked from all pending requests and issued to system memory Just-in-Time to make optimal use of Command Overlapping. Thus, instead of having all memory access requests go individually through an arbitration mechanism forcing requests to be executed one at a time, they can be started without interfering with the current request allowing for concurrent issuing of requests. This allows for optimized bandwidth and reduced latency while maintaining appropriate command spacing to meet system memory protocol.

Command Overlap

Command Overlap allows the insertion of the DRAM commands between the Activate, Pre-charge, and Read/Write commands normally used, as long as the inserted commands do not affect the currently executing command. Multiple commands can be issued in an overlapping manner, increasing the efficiency of system memory protocol.

Out-of-Order Scheduling

While leveraging the Just-in-Time Scheduling and Command Overlap enhancements, the IMC continuously monitors pending requests to system memory for the best use of bandwidth and reduction of latency. If there are multiple requests to the same open page, these requests would be launched in a back to back manner to make optimum use of the open memory page. This ability to reorder requests on the fly allows the IMC to further reduce latency and increase bandwidth efficiency.

2.1.6 Data Scrambling

The system memory controller incorporates a Data Scrambling feature to minimize the impact of excessive di/dt on the platform system memory VRs due to successive 1s and 0s on the data bus. Past experience has demonstrated that traffic on the data bus is not random and can have energy concentrated at specific spectral harmonics creating high di/dt which is generally limited by data patterns that excite resonance between the package inductance and on die capacitances. As a result, the system memory controller uses a data scrambling feature to create pseudo-random patterns on the system memory data bus to reduce the impact of any excessive di/dt.

2.1.7 ECC S-Matrix Syndrome Codes

Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit
0				No Error			
1	64	37	26	81	2	146	53
2	65	38	46	82	18	148	4
4	66	41	61	84	34	152	20

Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit	Syndrome Value	Flipped Bit
7	60	42	9	88	50	161	49
8	67	44	16	97	21	162	1
11	36	47	23	98	38	164	17
13	27	49	63	100	54	168	33
14	3	50	47	104	5	176	44
16	68	52	14	112	52	193	8
19	55	56	30	128	71	194	24
21	10	64	70	131	22	196	40
22	29	67	6	133	58	200	56
25	45	69	42	134	13	208	19
26	57	70	62	137	28	224	11
28	0	73	12	138	41	241	7
31	15	74	25	140	48	242	31
32	69	76	32	143	43	244	59
35	39	79	51	145	37	248	35

Notes:
1. All other syndrome values indicate unrecoverable error (more than one error).

2.1.8 DDR I/O Interleaving

The processor supports I/O interleaving, which has the ability to swap DDR bytes for routing considerations. BIOS configures the I/O interleaving mode before DDR training.

There are two supported modes:

- Interleave (IL)
- Non-Interleave (NIL)

The following table and figure describe the pin mapping between the IL and NIL modes.

Table 2-5. Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping (Sheet 1 of 2)

IL (DDR4)		NIL (DDR4)	
Channel	Byte	Channel	Byte
DDR0	Byte0	DDR0	Byte0
DDR0	Byte1	DDR0	Byte1
DDR0	Byte2	DDR0	Byte4
DDR0	Byte3	DDR0	Byte5
DDR0	Byte4	DDR1	Byte0
DDR0	Byte5	DDR1	Byte1
DDR0	Byte6	DDR1	Byte4
DDR0	Byte7	DDR1	Byte5
DDR1	Byte0	DDR0	Byte2
DDR1	Byte1	DDR0	Byte3

Table 2-5. Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping (Sheet 2 of 2)

IL (DDR4)		NIL (DDR4)	
Channel	**Byte**	**Channel**	**Byte**
DDR1	Byte2	DDR0	Byte6
DDR1	Byte3	DDR0	Byte7
DDR1	Byte4	DDR1	Byte2
DDR1	Byte5	DDR1	Byte3
DDR1	Byte6	DDR1	Byte6
DDR1	Byte7	DDR1	Byte7

Figure 2-2. Interleave (IL) and Non-Interleave (NIL) Modes Mapping

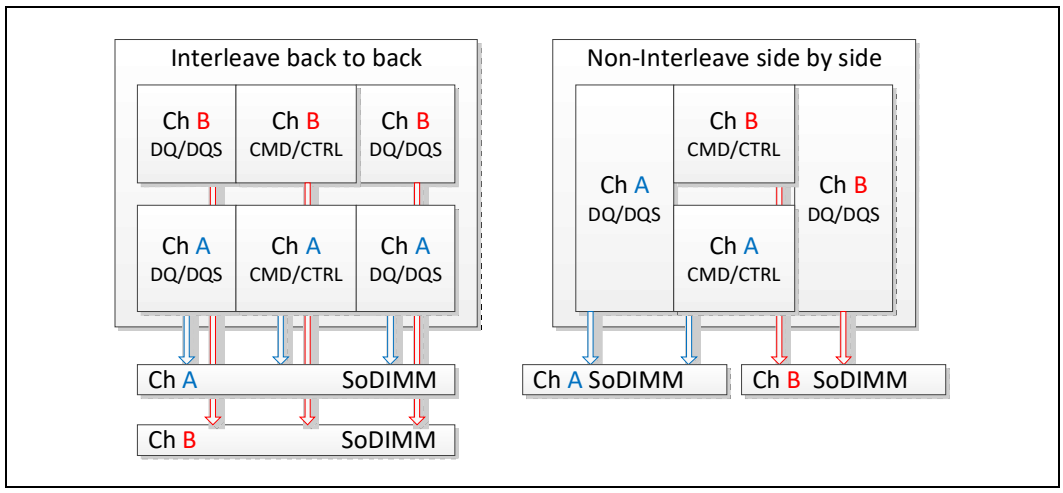

2.1.9 Data Swapping

By default, the processor supports on-board data swapping in two manners (for all segments and DRAM technologies):

- Byte (DQ+DQS) swapping between bytes in the same channel.
- Bit swapping within specific byte. ECC Byte swapping (with other Bytes) is not allowed, ECC bits swap is allowed.

Particular layouts may choose from the six options; each subchannel may pick a different option as the layout requires.

2.1.10 DRAM Clock Generation

Every supported rank has a differential clock pair. There are a total of four clock pairs driven directly by the processor to DRAM.

2.1.11 DRAM Reference Voltage Generation

The memory controller has the capability of generating the DDR4 Reference Voltage (VREF) internally for both read and write operations. The generated VREF can be changed in small steps, and an optimum VREF value is determined for both during a cold boot through advanced training procedures in order to provide the best voltage to achieve the best signal margins.

2.1.12 Data Swizzling

All Processor Lines do not have die-to-package DDR swizzling.

2.2 PCI Express* (PCIe*) Interface

This section describes the PCI Express* interface capabilities of the processor. Refer to *PCI Express Base* Specification 4.0* for details on PCI Express*.

The Intel® Xeon® E-2300 processor PCI Express* has two interfaces, 16-lane (x16) port and a 4-lane (x4) port.

The interconnect between the processor and the devices can be on board (soldered down) or be provided through either M.2 or AIC connectors.

2.2.1 PCI Express* Support

2.2.1.1 Bifurcation and Lane Reversal

The Intel® Xeon® E-2300 processor PCI Express* has two interfaces:

• 4-lane (x4) port (refer Table 2-6)

• 16-lane (x16) port that can also be configured as multiple ports at narrower widths (refer Table 2-7).

Table 2-6. PCI Express* 4 -lane Bifurcation and Lane Reversal Mapping

Bifurcation	Link Width	CFG Signals	Lanes			
	0:6:0	CFG [14]	0	1	2	3
PCIe Controller			PCIe 060			
1x4	x4	1	0	1	2	3
1x4 **Reversed**	x4	0	3	2	1	0
PCIe 060 is a single x4 port without bifurcation capabilities thus bifurcation pin straps are not applicable						

Table 2-7. PCI Express* 16-lane Bifurcation and Lane Reversal Mapping

Bifurcation	Link Width			CFG Signals			Lanes															
	0:1:0	0:1:1	0:1:2	CFG [6]	CFG [5]	CFG [2]	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
PCIe Controller							PCIe 010															
1x16	x16	N/A	N/A	1	1	1	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1x16 **Reversed**	x16	N/A	N/A	1	1	0	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
PCIe Controller							PCIe 010								PCIe 011							
2x8	x8	x8	N/A	1	0	1	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7
PCIe Controller							PCIe 011								PCIe 010							
2x8 **Reversed**	x8	x8	N/A	1	0	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
PCIe Controller							PCIe 010								PCIe 011				PCIe 012			
1x8+2x4	x8	x4	x4	0	0	1	0	1	2	3	4	5	6	7	0	1	2	3	0	1	2	3
PCIe Controller							PCIe 012				PCIe 011				PCIe 010							
1x8+2x4 **Reversed**	x8	x4	x4	0	0	0	3	2	1	0	3	2	1	0	7	6	5	4	3	2	1	0

Notes:
1. Support is also provided for narrow width and use devices with lower number of lanes (that is, usage on x4 configuration), however further bifurcation is not supported.
2. In case that more than one device is connected, the device with the highest lane count, should always be connected to the lower lanes, as follows:
   - Connect lane 0 of 1st device to lane 0.
   - Connect lane 0 of 2nd device to lane 8.
   - Connect lane 0 of 3rd device to lane 12.

   For example:
   a. When using 1x8 + 2x4, the 8 lane device should use lanes 0:7.
   b. When using 1x4 + 1x2, the 4 lane device should use lanes 0:3, and other 2 lanes device should use lanes 8:9.
   c. When using 1x4 + 1x2 + 1x1, 4 lane device should use lanes 0:3, two lane device should use lanes 8:9, one lane device should use lane 12.
3. For reversal lanes, for example: When using 1x8, the 8 lane device should use lanes 8:15, so lane 15 will be connected to lane 0 of the Device.
4. It is recommended to use PCIe 060 (x4 PCIe lanes) for SSD storage and PCIe 010 (x16 PCIe lanes) for GFx.

2.2.2 Supported Features

The processor supports the following:

- Hierarchical PCI-compliant configuration mechanism for downstream devices.
- Traditional PCI style traffic (asynchronous snooped, PCI ordering).
- PCI Express* extended configuration space. The first 256 bytes of configuration space aliases directly to the PCI Compatibility configuration space. The remaining portion of the fixed 4-KB block of memory-mapped space above that (starting at 100h) is known as extended configuration space.
- PCI Express* Enhanced Access Mechanism. Accessing the device configuration space in a flat memory-mapped fashion.
- Automatic discovery, negotiation, and training of link out of reset.
- Multiple Virtual Channel.
- 64-bit downstream address format, but the processor never generates an address above 512 GB (Bits 63:39 will always be zeros).

- 64-bit upstream address format, but the processor responds to upstream read transactions to addresses above 512 GB (addresses where any of Bits 63:39 are nonzero) with an Unsupported Request response. Upstream write transactions to addresses above 512 GB will be dropped.
- Re-issues Configuration cycles that have been previously completed with the Configuration Retry status.
- PCI Express* reference clock is a 100 MHz differential clock.
- Power Management Event (PME) functions.
- Modern Standby
- Dynamic width capability.
- Message Signaled Interrupt (MSI and MSI-X) messages.
- Lane reversal
- Advanced Error Reporting (AER)

2.2.3 Transfer Rates and Bandwidth

The following table summarizes the transfer rates and theoretical bandwidth of PCI Express* link.

Table 2-8. PCI Express* Maximum Transfer Rates and Theoretical Bandwidth

PCI Express* Generation	Encoding	Maximum Transfer Rate [GT/s]	Theoretical Bandwidth [GB/s]
			x16
1.0	8b/10b	2.5	4.0
2.0	8b/10b	5	8.0
3.0	128b/130b	8	15.8
4.0	128b/130b	16	31.5

2.2.4 PCI Express* Architecture

Compatibility with the PCI addressing model is maintained to ensure that all existing applications and drivers operate unchanged.

The PCI Express* configuration uses standard mechanisms as defined in the PCI Plug-and-Play specification. The processor PCI Express* port supports Gen 4.0 at 16 GT/s uses a 128b/130b encoding.

The four lanes port can operate at 2.5 GT/s, 5 GT/s, 8 GT/s or 16 GT/s.

The PCI Express* architecture is specified in three layers – Transaction Layer, Data Link Layer, and Physical Layer. Refer to the *PCI Express Base Specification 4.0* for details of PCI Express* architecture.

2.2.5 PCI Express* Configuration Mechanism

The PCI Express* (external graphics) link is mapped through a PCI-to-PCI bridge structure.

Figure 2-3. PCI Express* Related Register Structures in Processor

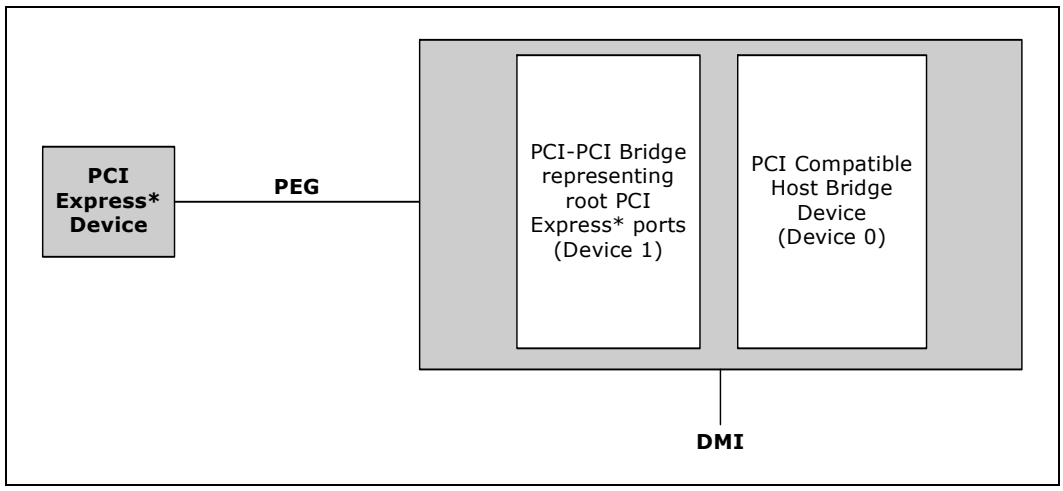

PCI Express* extends the configuration space to 4096 bytes per-device/function, as compared to 256 bytes allowed by the conventional PCI specification. PCI Express* configuration space is divided into a PCI-compatible region (that consists of the first 256 bytes of a logical device's configuration space) and an extended PCI Express* region (that consists of the remaining configuration space). The PCI-compatible region can be accessed using either the mechanisms defined in the PCI specification or using the enhanced PCI Express* configuration access mechanism described in the PCI Express* Enhanced Configuration Mechanism section.

The PCI Express* Host Bridge is required to translate the memory-mapped PCI Express* configuration space accesses from the host processor to PCI Express* configuration cycles. To maintain compatibility with PCI configuration addressing mechanisms, it is recommended that system software access the enhanced configuration space using 32-bit operations (32-bit aligned) only. Refer the PCI Express Base Specification for details of both the PCI-compatible and PCI Express* Enhanced configuration mechanisms and transaction rules.

2.2.6 PCI Express* Equalization Methodology

Link equalization requires equalization for both TX and RX sides for the processor and for the End point device.

Adjusting transmitter and receiver of the lanes is done to improve signal reception quality and for improving link robustness and electrical margin.

The link timing margins and voltage margins are strongly dependent on equalization of the link.

The processor supports the following:

- Full TX Equalization: Three Taps Linear Equalization (Pre, Current and Post cursors), with FS/LF (Full Swing /Low Frequency) 24/8 values respectively.

- Full RX Equalization and acquisition for: AGC (Adaptive Gain Control), CDR (Clock and Data Recovery), adaptive DFE (decision feedback equalizer) and adaptive CTLE peaking (continuous time linear equalizer).
- Full adaptive phase 3 EQ compliant with PCI Express* 3.0 specification.

Refer the *PCI Express* Base Specification 3.0* for details on PCI Express* equalization.

2.3 Direct Media Interface (DMI)

Note: The DMI interface is only present in 2-Chip platform processors.

Direct Media Interface (DMI) connects the processor and the PCH.

Main characteristics:

- 8 lanes Gen 3 DMI support
- Reduced 4 lane DMI support
- DC coupling - no capacitors between the processor and the PCH
- PCH end-to-end lane reversal across the link
- Half-Swing support (low-power/low-voltage)

2.3.1 DMI Lane Reversal and Polarity Inversion

Note: Polarity Inversion and Lane Reversal on DMI Link are not allowed in processor side. Lane reversal can only be allowed on the PCH side.

Figure 2-4. Example for DMI Lane Reversal Connection

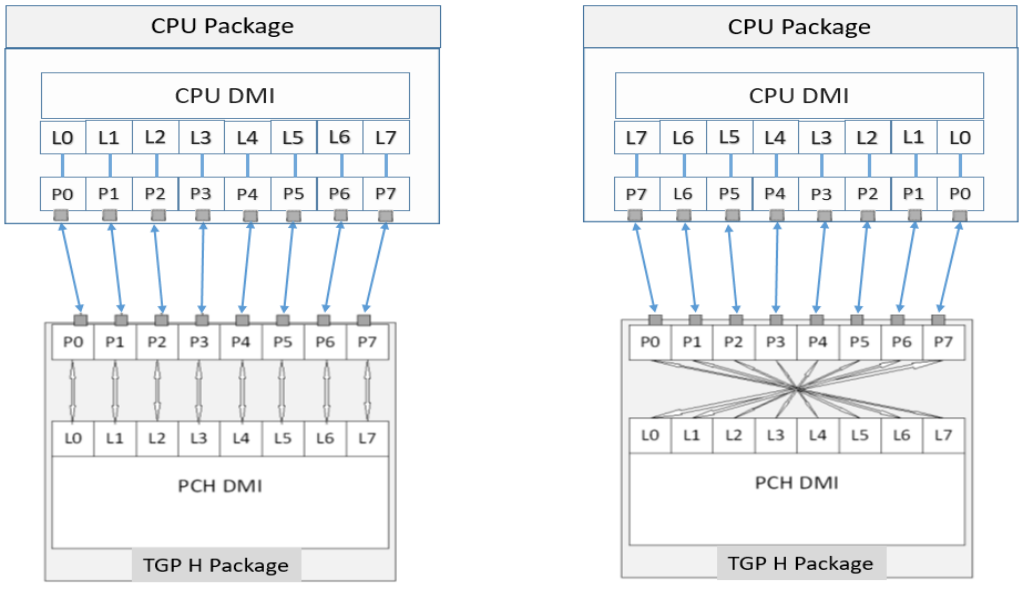

Notes:
1. DMI Lane Reversal is supported only on PCH-H and not on the Processor.
2. L[7:0] - Processor and PCH DMI Controller Logical Lane Numbers.
3. P[7:0] - Processor and PCH DMI Package Pin Lane Numbers.
4. For reduced 4 lane reversal connection, CPU L0 - CPU P0 <-> PCH P3 - PCH L0 and CPU L3 - CPU P3 <-> PCH P0 - PCH L3.

2.3.2 DMI Error Flow

DMI can only generate SERR in response to errors; never SCI, SMI, MSI, PCI INT, or GPE. Any DMI related SERR activity is associated with Device 0.

2.3.3 DMI Link Down

The DMI link going down is a fatal, unrecoverable error. If the DMI data link goes to data link down, after the link was up, then the DMI link hangs the system by not allowing the link to retrain to prevent data corruption. This link behavior is controlled by the PCH.

Downstream transactions that had been successfully transmitted across the link prior to the link going down may be processed as normal. No completions from downstream, non-posted transactions are returned upstream over the DMI link after a link down event.

2.4 Processor Graphics

The processor graphics architecture delivers high dynamic range of scaling to address segments spanning low power to high power, increased performance per watt, support for next generation of APIs. X^e scalable architecture is partitioned by usage domains along Render/Geometry, Media, and Display. The architecture also delivers very low-power video playback and next generation analytics and filters for imaging related applications. The new Graphics Architecture includes 3D compute elements, Multi-format HW assisted decode/encode pipeline, and Mid-Level Cache (MLC) for superior high definition playback, video quality, and improved 3D performance and media.

2.4.1 Media Support (Intel® QuickSync and Clear Video Technology HD)

X^e implements multiple media video codecs in hardware as well as a rich set of image processing algorithms.

Note: Supports 12 bit HEVC and VP9 decode and additional 10 bit decode in YCbCr 4:2:2 or 4:4:4 profiles. Refer additional detail support matrix.

2.4.1.1 Hardware Accelerated Video Decode

X^e implements a high-performance and low-power HW acceleration for video decoding operations for multiple video codecs.

The HW decode is exposed by the graphics driver using the following APIs:

- Direct3D* 9 Video API (DXVA2)
- Direct3D12 Video API
- Intel® Media SDK
- MFT (Media Foundation Transform) filters.

X^e supports full HW accelerated video decoding for AVC/VC1/MPEG2/HEVC/VP8/JPEG.

Note: HEVC – 10 bit support.

Table 2-9. Hardware Accelerated Video Decoding

Codec	Profile	Level	Maximum Resolution
MPEG2	Main	Main High	1080p
VC1/WMV9	Advanced Main Simple	L3 High Simple	3840x3840
AVC/H264	High Main	L5.2	4K
	4:2:0 8bit		4K at 60
JPEG/MJPEG	Baseline	Unified level	16K x16K
HEVC/H265	Main 12 Main 422 10 Main 422 12 Main 444 Main 444 10 Main 444 12 SCC main SCC main 10 SCC main 444 SCC main 444 10	L6.2	5K at 60 8K at 60
VP9	0 (4:2:0 Chroma 8-bit) 1 (4:4:4 8 bit) 2 (4:2:0 Chroma 10/12 bit)	Unified level	4320p(8K) 16Kx4K
	4:4:4 10bit		5K at 60
	4:2:0 12bit		8K at 60
AV1	0 (4:2:0 8-bit) 0 (4:2:0 10-bit)	L3	4K x 2K (video) 16K x 16K (still picture)

Expected performance:

- More than 16 simultaneous decode streams at 1080p.

Note: Actual performance depends on the processor SKU, content bit rate, and memory frequency. Hardware decode for H264 SVC is not supported

Intel® Xeon® E-2300 processor supports up to 5K at 60

2.4.1.2 Hardware Accelerated Video Encode

X^e implements a low-power low-latency fixed function encoder and a high-quality customizable encoder with hardware assisted motion estimation engine which supports AVC, MPEG-2, HEVC, and VP9.

The HW encode is exposed by the graphics driver using the following APIs:

- Intel® Media SDK
- MFT (Media Foundation Transform) filters

X^e supports full HW accelerated video encoding for AVC/MPEG2/HEVC/VP9/JPEG.

Table 2-10. Hardware Accelerated Video Encode

Codec	Profile	Level	Maximum Resolution
MPEG2	Main10	High	8K
AVC/H264	High Main	L5.1	2160p(4K)
JPEG	Baseline	—	16Kx16K
HEVC/H265	Main Main10 Main 4:2:2 10 Main 4:4:4 Main 4:4:4 10	L5.1	4320p(8K) 16Kx4K at higher freq
VP9	0 (4:2:0 Chroma 8 bit) 1 (partial: 4:4:4 8 bit) 2 (partial: 4:2:0 10 bit) 3 (partial: 4:4:4 10 bit)	—	4320p(8K) 16Kx4K at higher freq

Note: Hardware encode for H264 SVC is not supported.

2.4.1.3 Hardware Accelerated Video Processing

There is hardware support for image processing functions such as De-interlacing, Film cadence detection, Advanced Video Scaler (AVS), detail enhancement, image stabilization, gamut compression, HD adaptive contrast enhancement, skin tone enhancement, total color control, Chroma de-noise, SFC (Scalar and Format Conversion), memory compression, Localized Adaptive

Contrast Enhancement (LACE), spatial de-noise, Out-Of-Loop De-blocking (from AVC decoder), 16 bpc support for de-noise/de-mosaic.

The HW video processing is exposed by the graphics driver using the following APIs:

- Direct3D* 9 Video API (DXVA2).
- Direct3D* 11 Video API.
- Intel® Media SDK.
- MFT (Media Foundation Transform) filters.
- Intel® CUI SDK.

Note: Not all features are supported by all the previous APIs. Refer to the relevant documentation for more details.

2.4.1.4 Hardware Accelerated Transcoding

Transcoding is a combination of decode, video processing (optional) and encode. Using the previous hardware capabilities can accomplish a high-performance transcode pipeline. There is not a dedicated API for transcoding.

The processor graphics supports the following transcoding features:

- High performance high quality flexible encoder for video editing, video archiving.
- Low-power low latency encoder for video conferencing, wireless display, and game streaming.
- Lossless memory compression for media engine to reduce media power.

- High-quality Advanced Video Scaler (AVS)
- Low power Scaler and Format Converter.

2.5 Platform Graphics Hardware Feature

2.5.1 Hybrid Graphics

Microsoft* Windows* 10 operating system enables the Windows*10 Hybrid graphics framework wherein the GPUs and their drivers can be simultaneously utilized to provide users with the benefits of both performance capability of discrete GPU (dGPU) and low power display capability of the processor GPU (iGPU). For instance, when there is a high-end 3D gaming workload in progress, the dGPU will process and render the game frames using its graphics performance, while iGPU continues to perform the display operations by compositing the frames rendered by dGPU. We recommend that OEMS should seek further guidance from Microsoft* to confirm that the design fits all the latest criteria defined by Microsoft* to support HG.

Microsoft* Hybrid Graphics definition includes the following:

1. The system contains a single integrated GPU and a single discrete GPU.
2. It is a design assumption that the discrete GPU has a significantly higher performance than the integrated GPU.
3. Both GPUs shall be physically enclosed as part of the system.
   — Microsoft* Hybrid DOES NOT support hot-plugging of GPUs
   — OEMS should seek further guidance from Microsoft* before designing systems with the concept of hot-plugging
4. Starting with Windows*10 Th1 (WDDM 2.0), a previous restriction that the discrete GPU is a render-only device, with no displays connected to it, has been removed. A render-only configuration with NO outputs is still allowed, just NOT required.

It must be noted that systems that have outputs available off of the discrete GPU will NOT support previous versions of the OS (Windows* 8.1 and Older).

Note: Intel® Quick Sync Video (QSV) is supported even when a discrete graphics card is plugged in.

2.6 Display Interfaces

2.6.1 Display Technologies Support

Technology	Standard
eDP* 1.4b	VESA* Embedded DisplayPort* Standard 1.4b
DisplayPort* 1.4a	VESA* DisplayPort* Standard 1.4a VESA* DisplayPort* PHY Compliance Test Specification 1.4a VESA* DisplayPort* Link Layer Compliance Test Specification 1.4
HDMI* 2.0b	High-Definition Multimedia Interface Specification Version 2.0b

Table 2-11. Display Ports Availability and Link Rate

SKU	S81-Processor Line
DDI A / EDP	eDP* up to HBR2 DP* up to HBR3 HDMI* up to 5.94 Gbps
DDI B / DDI 1	DP* up to HBR3 HDMI* up to 5.94 Gbps
DDI C / DDI TC1 / DDI 2	
DDI D / DDI TC2 / DDI 3	
Notes: 1. HBR3 - 8.1 Gbps lane rate, supported using on board re-timer for DP* interface. 2. HBR2 - 5.4 Gbps lane rate	

Figure 2-5. Processor Display Architecture

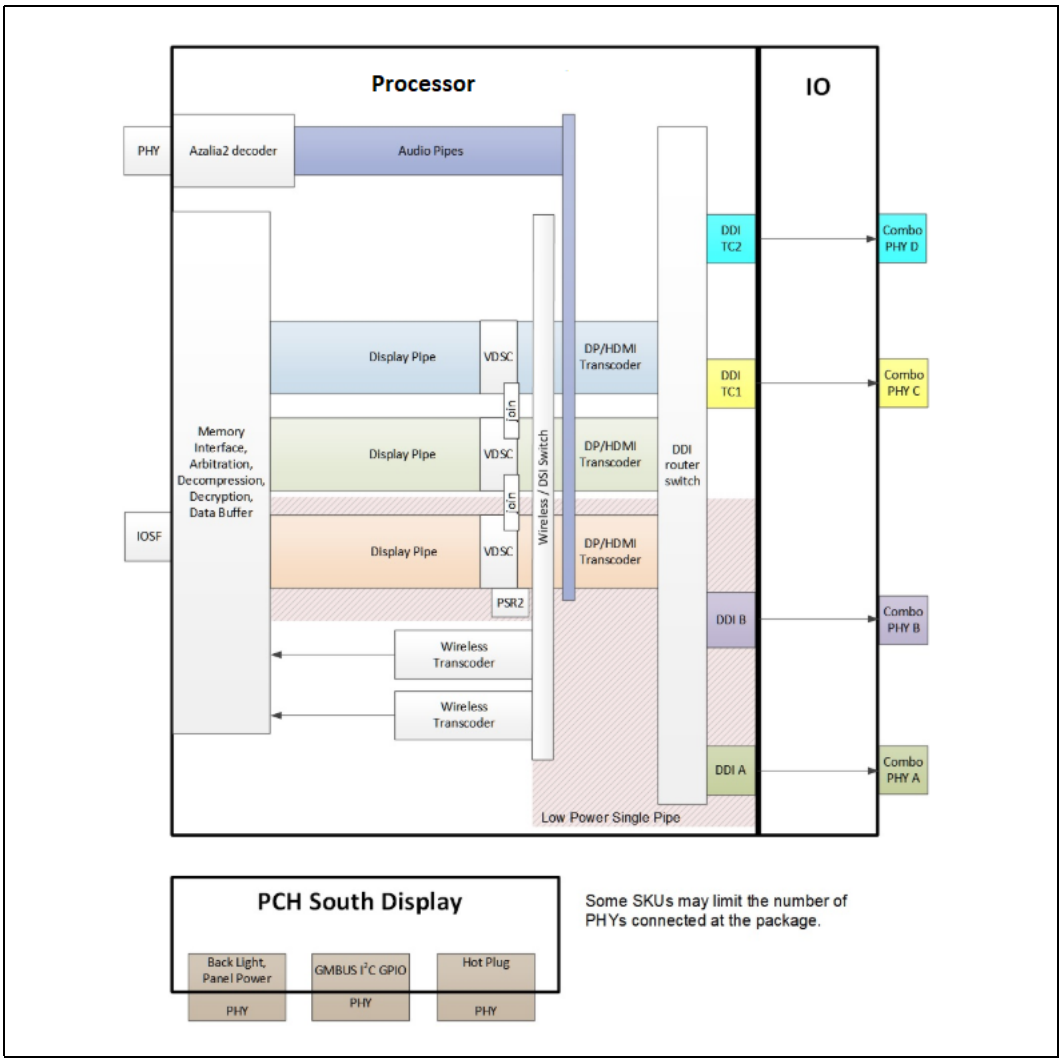

Note: For port availability in each of the processor lines, refer to Table 2-11, "Display Ports Availability and Link Rate".

2.6.2 Display Features

2.6.2.1 General Capabilities

- Up to three simultaneous displays.
  - — Up to three 4K 60 Hz displays.
  - — Up to two 5K 60 Hz displays.
- Audio stream support on external ports.
- Up to four display interfaces that can be configured as eDP*, DisplayPort*, HDMI* refer to Table 2-11, "Display Ports Availability and Link Rate".
- HDR (High Dynamic Range) support.

- Three Display Pipes - Supporting blending, color adjustments, scaling and dithering.
- Transcoders - Containing the Timing generators supporting eDP*, DP*, HDMI* interfaces.
- DPST - Display Power Saving Technology.
- Low Power optimized pipe.
  - LACE (Localized Adaptive Contrast Enhancement).
  - 3D LUT - power efficient pixel modification function for color processing.
  - FBC (Frame Buffer Compression) - power saving feature.

2.6.3 Multiple Display Configurations

The following multiple display configuration modes are supported (with appropriate driver software):

- Single Display is a mode with one display port activated to display the output to one display device.
- Display Clone is a mode with up to three display ports activated to drive the display content of same color depth setting but potentially different refresh rate and resolution settings to all the active display devices connected.
- Extended Desktop is a mode with up to three display ports activated to drive the content with potentially different color depth, refresh rate, and resolution settings on each of the active display devices connected.

2.6.4 High-bandwidth Digital Content Protection (HDCP)

HDCP is the technology for protecting high-definition content against unauthorized copy or unreceptive between a source (computer, digital set top boxes, and so on) and the sink (panels, monitor, and TVs). The processor supports HDCP 1.4, 2.2, 2.3 content protection over wired displays (HDMI*, DVI, and DisplayPort*).

The HDCP 1.4/2.2/2.3 keys are integrated into the processor and customers are not required to physically configure or handle the keys.

2.6.5 DisplayPort*

The DisplayPort* is a digital communication interface that uses differential signaling to achieve a high-bandwidth bus interface designed to support connections between PCs and monitors, projectors, and TV displays.

A DisplayPort* consists of a Main Link (four lanes), Auxiliary channel, and a Hot-Plug Detect signal. The Main Link is a unidirectional, high-bandwidth, and low-latency channel used for transport of isochronous data streams such as uncompressed video and audio. The Auxiliary Channel (AUX CH) is a half-duplex bi-directional channel used for link management and device control. The Hot-Plug Detect (HPD) signal serves as an interrupt request from the sink device to the source device.

The processor is designed in accordance with VESA* DisplayPort* specification. Refer to Section 2.6.1, “Display Technologies Support”.

Figure 2-6. DisplayPort* Overview

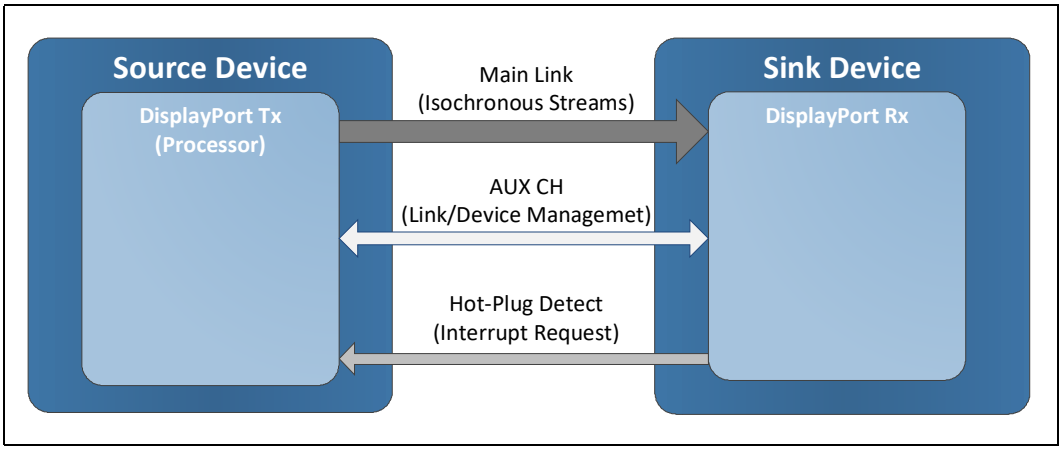

- Support main link of 1, 2, or 4 data lanes.
- Aux channel for Link/Device management.
- Support up to 36 BPP (Bit Per Pixel).
- Support SSC.
- Support YCbCR 4:4:4, YCbCR 4:2:0, and RGB color format.
- Support MST (Multi-Stream Transport).
- Support VESA DSC 1.1.
- Adaptive Sync.

2.6.5.1 Multi-Stream Transport (MST)

- The processor supports Multi-Stream Transport (MST), enabling multiple monitors to be used via a single DisplayPort connector.
- Maximum Multi Stream Transport DP supported resolution are listed in the table below.

Table 2-12. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations (Sheet 1 of 2)

Pixels per Line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
640	480	60	25.2	0.76
800	600	60	40	1.20
1024	768	60	65	1.95
1280	720	60	74.25	2.23
1280	768	60	68.25	2.05
1360	768	60	85.5	2.57
1280	1024	60	108	3.24
1400	1050	60	101	3.03
1680	1050	60	119	3.57
1920	1080	60	148.5	4.46

Table 2-12. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations (Sheet 2 of 2)

Pixels per Line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
1920	1200	60	154	4.62
2048	1152	60	156.75	4.70
2048	1280	60	174.25	5.23
2048	1536	60	209.25	6.28
2304	1440	60	218.75	6.56
2560	1440	60	241.5	7.25
3840	2160	30	262.75	7.88
2560	1600	60	268.5	8.06
2880	1800	60	337.5	10.13
3200	2400	60	497.75	14.93
3840	2160	60	533.25	16.00
4096	2160	60	556.75	16.70
4096	2304	60	605	18.15
5120	3200	60	1042.5	31.28

Notes:
1. All the previous is related to bit depth of 24.
2. The data rate for a given video mode can be calculated as- Data Rate = Pixel Frequency * Bit Depth.
3. The bandwidth requirements for a given video mode can be calculated as: Bandwidth = Data Rate * 1.25 (for 8B/10B coding overhead).
4. The link bandwidth depends if the standards is reduced blanking or not.
   If the standard is not reduced blanking - the expected bandwidth may be higher.
   For more details refer to VESA and Industry Standards and Guidelines for Computer Display Monitor Timing (DMT). Version 1.0, Rev. 13 February 8, 2013
5. To calculate what are the resolutions that can be supported in MST configurations, follow the below guidelines:
   a. Identify what is the link bandwidth column according to the requested display resolution.
   b. Summarize the bandwidth for two of three displays accordingly, and make sure the final result is below 21.6 Gbps. (for example: 4 lanes HBR2 bit rate)

   For example:
   a. Docking two displays: 3840x2160@60 Hz + 1920x1200@60hz = 16 + 4.62 = 20.62Gbps [Supported]
   b. Docking three displays: 3840x2160@30 Hz + 3840x2160@30 Hz + 1920x1080@60 Hz = 7.88 + 7.88 + 4.16 = 19.92 Gbps [Supported]

Table 2-13. DisplayPort Maximum Resolution

Standard	S81-Processor Line[1]
DP*	4096x2304 60Hz 36bpp 5120x3200 60Hz 24bpp
DP* with DSC	5120x3200 60Hz 30bpp

Notes:
1. Maximum resolution is based on the implementation of 4 lanes at HBR3 link data rate.
2. bpp - bit per pixel.
3. Resolution support is subject to memory BW availability.

2.6.6 High-Definition Multimedia Interface (HDMI*)

The High-Definition Multimedia Interface (HDMI*) is provided for transmitting uncompressed digital audio and video signals from DVD players, set-top boxes, and other audio-visual sources to television sets, projectors, and other video displays. It can carry high-quality multi-channel audio data and all standard and high-definition

consumer electronics video formats. The HDMI display interface connecting the processor and display devices uses transition minimized differential signaling (TMDS) to carry audiovisual information through the same HDMI cable.

HDMI* includes three separate communications channels: TMDS, DDC, and the optional CEC (consumer electronics control). CEC is not supported on the processor. As shown in the following figure, the HDMI* cable carries four differential pairs that make up the TMDS data and clock channels. These channels are used to carry video, audio, and auxiliary data. In addition, HDMI carries a VESA DDC. The DDC is used by an HDMI* Source to determine the capabilities and characteristics of the Sink.

Audio, video, and auxiliary (control/status) data is transmitted across the three TMDS data channels. The video pixel clock is transmitted on the TMDS clock channel and is used by the receiver for data recovery on the three data channels. The digital display data signals driven natively through the PCH are AC coupled and needs level shifting to convert the AC coupled signals to the HDMI* compliant digital signals. The processor HDMI* interface is designed in accordance with the High-Definition Multimedia Interface.

Figure 2-7. HDMI* Overview

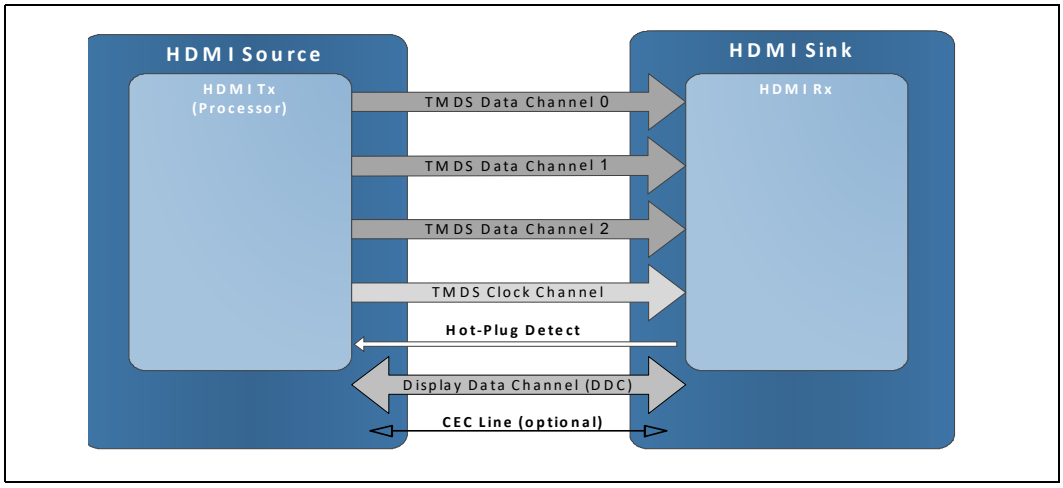

- DDC (Display Data Channel) channel.
- Support YCbCR 4:4:4, YCbCR 4:2:0, and RGB color format.
- Support up to 36 BPP (Bit Per Pixel).

Table 2-14. HDMI Supported Resolutions

Standard	S81-Processor Line
HDMI 1.4	4Kx2K 24-30Hz 24bpp
HDMI 2.0b	4Kx2K 48-60Hz 24bpp (RGB/YUV444) 4Kx2K 48-60Hz 12bpc (YUV420)
Notes: 1. bpp - bit per pixel. 2. Resolution support is subject to memory BW availability.	

2.6.7 embedded DisplayPort* (eDP*)

The embedded DisplayPort* (eDP*) is an embedded version of the DisplayPort standard oriented towards applications such as notebook and All-In-One PCs. Like DisplayPort, embedded DisplayPort* also consists of the Main Link, Auxiliary channel, and an optional Hot-Plug Detect signal.

- Supported on Low power optimized pipe.
- Support up to HBR2 link rate.
- Support Backlight PWM control signal.
- Support SSC
- Panel Self Refresh 1.
- MSO 2x2 (Multi Segment Operation).
- Dedicated Aux channel.

Table 2-15. Embedded DisplayPort Maximum Resolution

Standard	S81-Processor Line [1]
eDP*	4096x2304 60Hz 30bpp
eDP* with DSC	4096x2304 60Hz 30bpp

Notes:
1. Maximum resolution is based on the implementation of 4 lanes at HBR2 link data rate.
2. bpp - bit per pixel.
3. Resolution support is subject to memory BW availability.

2.6.8 Integrated Audio

- HDMI* and DisplayPort interfaces carry audio along with the video.
- The processor supports three High Definition audio streams on three digital ports simultaneously (the DMA controllers are in PCH).
- The integrated audio processing (DSP) is performed by the PCH and delivered to the processor using the AUDIO_SDI and AUDIO_CLK inputs pins.
- The AUDIO_SDO output pin is used to carry responses back to the PCH.
- Supports only the internal HDMI and DP CODECs.

Table 2-16. Processor Supported Audio Formats over HDMI and DisplayPort*

Audio Formats	HDMI*	DisplayPort*
AC-3 Dolby* Digital	Yes	Yes
Dolby* Digital Plus	Yes	Yes
DTS-HD*	Yes	Yes
LPCM, 192 kHz/24 bit, 6 Channel	Yes	Yes
Dolby* TrueHD, DTS-HD Master Audio* (Lossless Blu-Ray Disc* Audio Format)	Yes	Yes

The processor will continue to support Silent stream. A Silent stream is an integrated audio feature that enables short audio streams, such as system events to be heard over the HDMI* and DisplayPort* monitors. The processor supports silent streams over the HDMI and DisplayPort interfaces at 44.1 kHz, 48 kHz, 88.2 kHz, 96 kHz, 176.4 kHz, and 192 kHz sampling rates and silent multi-stream support.

2.7 Platform Environmental Control Interface (PECI)

PECI is an Intel proprietary interface that provides a communication channel between Intel processors and external components like Super IO (SIO) and Embedded Controllers (EC) to provide processor temperature, Turbo, Configurable TDP, and memory throttling control mechanisms and many other services. PECI is used for platform thermal management and real time control and configuration of processor features and performance.

2.7.1 PECI Bus Architecture

The PECI architecture is based on a wired OR bus that the clients (as processor PECI) can pull up (with strong drive).

The idle state on the bus is near zero.

The following figures demonstrate PECI design and connectivity:

- PECI Host-Clients Connection: While the host/originator can be third party PECI host and one of the PECI client is a processor PECI device.
- PECI EC Connection.
- PECI over eSPI is supported.

Figure 2-8. Example for PECI Host-Clients Connection

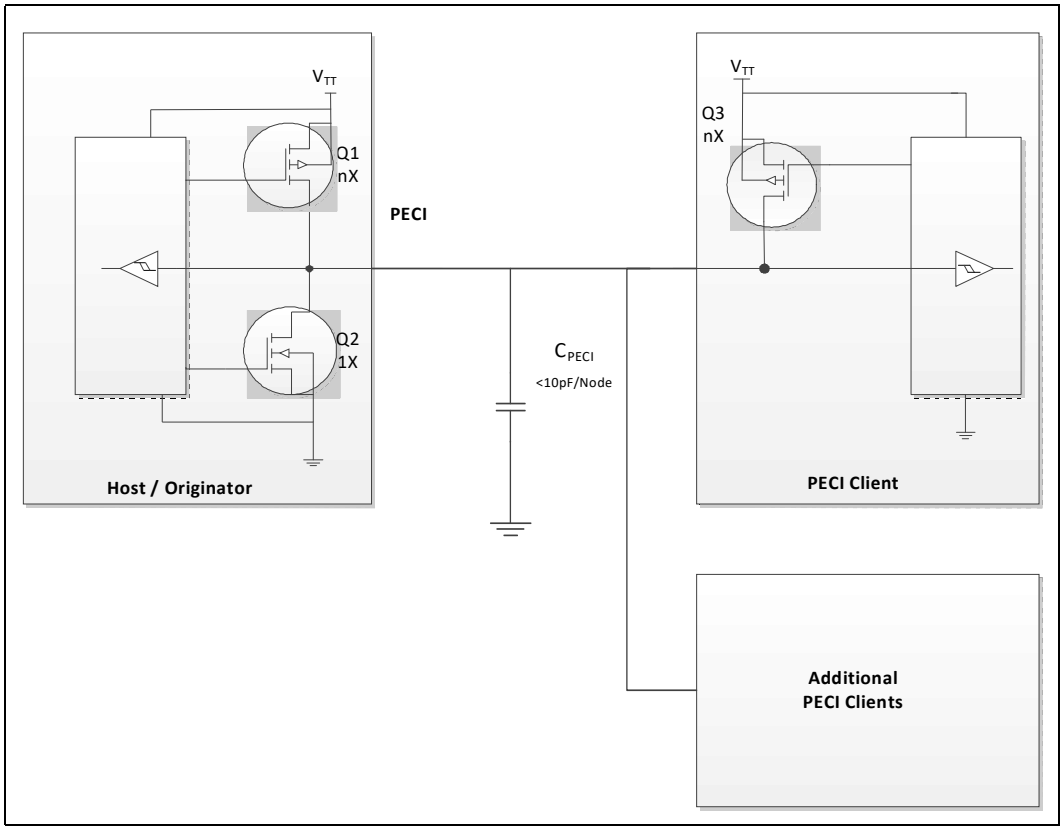

Figure 2-9. Example for PECI EC Connection

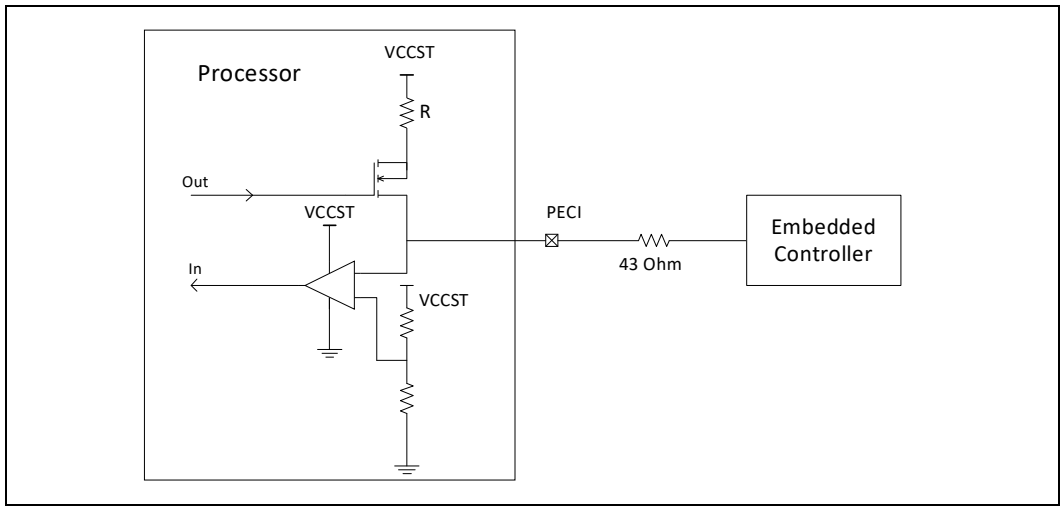

3 Technologies

This chapter provides a high-level description of Intel technologies implemented in the processor.

The implementation of the features may vary between the processor SKUs.

Details on the different technologies of Intel processors and other relevant external notes are located at the Intel technology web site: http://www.intel.com/technology/

3.1 Intel® Virtualization Technology (Intel® VT)

Intel® Virtualization Technology (Intel® VT) makes a single system appear as multiple independent systems to software. This allows multiple, independent operating systems to run simultaneously on a single system. Intel® VT comprises technology components to support virtualization of platforms based on Intel architecture microprocessors and chipsets.

Intel® Virtualization Technology (Intel® VT), Intel® 64, and Intel® Architecture (Intel® VT-x) added hardware support in the processor to improve the virtualization performance and robustness. Intel® Virtualization Technology for Directed I/O (Intel® VT-d) extends Intel® VT-x by adding hardware assisted support to improve I/O device virtualization performance.

Intel® VT-x specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 3*. Available at:

http://www.intel.com/products/processor/manuals

The Intel® VT-d specification and other VT documents can be referenced at:

http://www.intel.com/content/www/us/en/virtualization/virtualization-technology/intel-virtualization-technology.html

3.1.1 Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)

Intel® VT-x Objectives

Intel® VT-x provides hardware acceleration for virtualization of IA platforms. Virtual Machine Monitor (VMM) can use Intel® VT-x features to provide an improved reliable virtualized platform. By using Intel® VT-x, a VMM is:

- **Robust:** VMMs no longer need to use para-virtualization or binary translation. This means that VMMs will be able to run off-the-shelf operating systems and applications without any special steps.
- **Enhanced:** Intel® VT enables VMMs to run 64-bit guest operating systems on IA x86 processors.
- **Reliable:** Due to the hardware support, VMMs can now be smaller, less complex, and more efficient. This improves reliability and availability and reduces the potential for software conflicts.

- **Secure:** The use of hardware transitions in the VMM strengthens the isolation of VMs and further prevents corruption of one VM from affecting others on the same system.

Intel® VT-x Key Features

The processor supports the following added new Intel® VT-x features:

- Extended Page Table (EPT) Accessed and Dirty Bits
  - — EPT A/D bits enabled VMMs to efficiently implement memory management and page classification algorithms to optimize VM memory operations, such as de-fragmentation, paging, live migration, and check-pointing. Without hardware support for EPT A/D bits, VMMs may need to emulate A/D bits by marking EPT paging-structures as not-present or read-only, and incur the overhead of EPT page-fault VM exits and associated software processing.
- EPTP (EPT Pointer) Switching
  - — EPTP switching is a specific VM function. EPTP switching allows guest software (in VMX non-root operation, supported by EPT) to request a different EPT paging-structure hierarchy. This is a feature by which software in VMX non-root operation can request a change of EPTP without a VM exit. Software will be able to choose among a set of potential EPTP values determined in advance by software in VMX root operation.
- Pause Loop Exiting
  - — Support VMM schedulers seeking to determine when a virtual processor of a multiprocessor virtual machine is not performing useful work. This situation may occur when not all virtual processors of the virtual machine are currently scheduled and when the virtual processor in question is in a loop involving the PAUSE instruction. The new feature allows detection of such loops and is thus called PAUSE-loop exiting.

The processor IA core supports the following Intel® VT-x features:

- **Mode-based Execute Control for EPT (MBEC)**
  - — A mode of EPT operation which enables different controls for executability of Guest Physical Address (GPA) based on Guest specified mode (User/ Supervisor) of linear address translating to the GPA. When the mode is enabled, the executability of a GPA is defined by two bits in EPT entry. One bit for accesses to user pages and other one for accesses to supervisor pages.
  - — This mode requires changes in VMCS and EPT entries. VMCS includes a bit "Mode-based execute control for EPT" which is used to enable/disable the mode. An additional bit in EPT entry is defined as "execute access for user-mode linear addresses"; the original EPT execute access bit is considered as "execute access for supervisor-mode linear addresses". If the "mode-based execute control for EPT" VM-execution control is disabled the additional bit is ignored and the system work with one bit that is, the original bit, for execute control for both user and supervisor pages.
  - — Behavioral changes - Behavioral changes are across three areas:
- **Access to GPA -** If the "Mode-based execute control for EPT" VMexecution control is 1, treatment of guest-physical accesses by instruction fetches depends on the linear address from which an instruction is being fetched.

  a. If the translation of the linear address specifies user mode (the U/S bit was set in every paging structure entry used to translate the linear address), the resulting guest-physical address is executable under EPT only if the XU bit (at position 10) is set in every EPT paging-structure entry used to translate the guest-physical address.

b. If the translation of the linear address specifies supervisor mode (the U/S bit was clear in at least one of the paging-structure entries used to translate the linear address), the resulting guest-physical address is executable under EPT only if the XS bit is set in every EPT paging-structure entry used to translate the guest-physical address.

  - The XU and XS bits are used only when translating linear addresses for guest code fetches. They do not apply to guest page walks, data accesses, or A/D-bit updates.

- **VMEntry -** If the **"**activate secondary controls" and "Mode-based execute control for EPT" VM-execution controls are both 1, VM entries ensure that the "enable EPT" VM-execution control is 1. VM entry fails if this check fails. When such a failure occurs, control is passed to the next instruction,

- **VMExit** - The exit qualification due to EPT violation reports clearly whether the violation was due to User mode access or supervisor mode access.
  - Capability Querying: IA32_VMX_PROCBASED_CTLS2 has bit to indicate the capability, RDMSR can be used to read and query whether the processor supports the capability or not.
- Extended Page Tables (EPT)
  - EPT is hardware assisted page table virtualization
  - It eliminates VM exits from guest OS to the VMM for shadow page-table maintenance
- Virtual Processor IDs (VPID)
  - Ability to assign a VM ID to tag processor IA core hardware structures (such as TLBs)
  - This avoids flushes on VM transitions to give a lower-cost VM transition time and an overall reduction in virtualization overhead.
- Guest Preemption Timer
  - Mechanism for a VMM to preempt the execution of a guest OS after an amount of time specified by the VMM. The VMM sets a timer value before entering a guest
  - The feature aids VMM developers in flexibility and Quality of Service (QoS) guarantees
- Descriptor-Table Exiting
  - Descriptor-table exiting allows a VMM to protect a guest OS from internal (malicious software based) attack by preventing relocation of key system data structures like IDT (interrupt descriptor table), GDT (global descriptor table), LDT (local descriptor table), and TSS (task segment selector).
  - A VMM using this feature can intercept (by a VM exit) attempts to relocate these data structures and prevent them from being tampered by malicious software.

3.1.2 Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)

Intel® VT-d Objectives

The key Intel® VT-d objectives are domain-based isolation and hardware-based virtualization. A domain can be abstractly defined as an isolated environment in a platform to which a subset of host physical memory is allocated. Intel® VT-d provides accelerated I/O performance for a virtualized platform and provides software with the following capabilities:

- **I/O Device Assignment and Security**: for flexibly assigning I/O devices to VMs and extending the protection and isolation properties of VMs for I/O operations.
- **DMA Remapping**: for supporting independent address translations for Direct Memory Accesses (DMA) from devices.
- **Interrupt Remapping**: for supporting isolation and routing of interrupts from devices and external interrupt controllers to appropriate VMs.
- **Reliability**: for recording and reporting to system software DMA and interrupt errors that may otherwise corrupt memory or impact VM isolation.

Intel® VT-d accomplishes address translation by associating transaction from a given I/O device to a translation table associated with the Guest to which the device is assigned. It does this by means of the data structure in the following illustration. This table creates an association between the device's PCI Express* Bus/Device/Function (B/D/F) number and the base address of a translation table. This data structure is populated by a VMM to map devices to translation tables in accordance with the previous device assignment restrictions, and to include a multi-level translation table (Intel® VT-d Table) that contains Guest specific address translations.

Figure 3-1. Device to Domain Mapping Structures

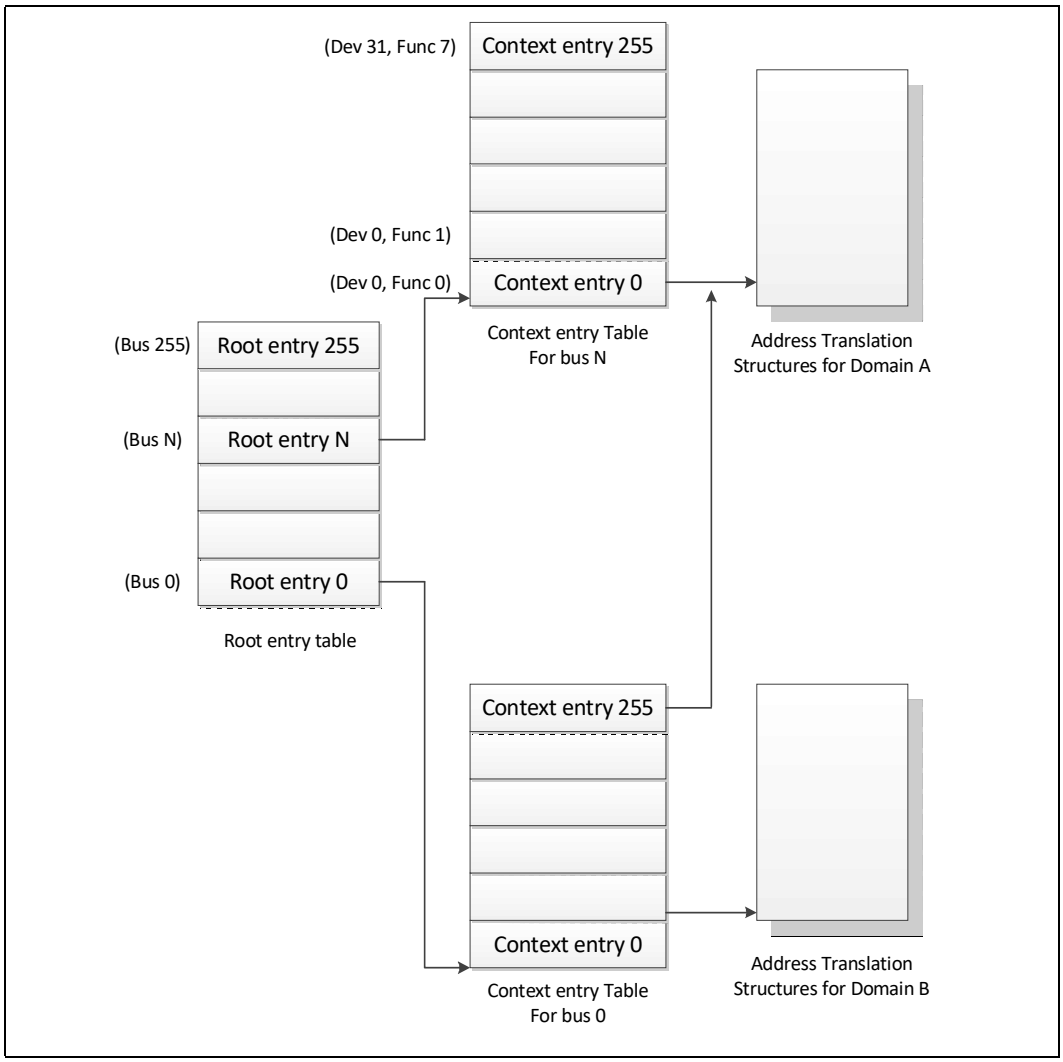

Intel® VT-d functionality, often referred to as an Intel® VT-d Engine, has typically been implemented at or near a PCI Express* host bridge component of a computer system. This might be in a chipset component or in the PCI Express* functionality of a processor with integrated I/O. When one such VT-d engine receives a PCI Express* transaction from a PCI Express* bus, it uses the B/D/F number associated with the transaction to search for an Intel® VT-d translation table. In doing so, it uses the B/D/F number to traverse the data structure shown in the previous figure. If it finds a valid Intel® VT-d table in this data structure, it uses that table to translate the address provided on the PCI Express bus. If it does not find a valid translation table for a given translation, this results in an Intel® VT-d fault. If Intel® VT-d translation is required, the Intel® VT-d engine performs an N-level table walk.

For more information, refer to Intel® Virtualization Technology for Directed I/O Architecture Specification http://www.intel.com/content/dam/www/public/us/en/documents/product-specifications/vt-directed-io-spec.pdf

Intel® VT-d Key Features

The processor supports the following Intel® VT-d features:

- Memory controller and processor graphics comply with the Intel® VT-d 2.1 Specification.
- Two Intel® VT-d DMA remap engines.
  - — iGFX DMA remap engine
  - — Default DMA remap engine (covers all devices except iGFX)
- Support for root entry, context entry, and default context
- 39-bit guest physical address and host physical address widths
- Support for 4K page sizes only
- Support for register-based fault recording only (for single entry only) and support for MSI interrupts for faults
- Support for both leaf and non-leaf caching
- Support for boot protection of default page table
- Support for non-caching of invalid page table entries
- Support for hardware based flushing of translated but pending writes and pending reads, on IOTLB invalidation
- Support for Global, Domain specific and Page specific IOTLB invalidation
- MSI cycles (MemWr to address FEEx_xxxxh) not translated
- Interrupt Remapping is supported
- Queued invalidation is supported
- Intel® VT-d translation bypass address range is supported (Pass Through)

The processor supports the following added new Intel® VT-d features:

- 4-level Intel® VT-d Page walk – both default Intel® VT-d engine as well as the Processor Graphics VT-d engine are upgraded to support 4-level Intel® VT-d tables (adjusted guest address width of 48 bits)
- Intel® VT-d superpage – support of Intel® VT-d superpage (2 MB, 1 GB) for default Intel® VT-d engine (that covers all devices except IGD)
  IGD Intel® VT-d engine does not support superpage and BIOS should disable superpage in default Intel® VT-d engine when iGfx is enabled.

Note: Intel® VT-d Technology may not be available on all SKUs.

3.1.3 Intel® APIC Virtualization Technology (Intel® APICv)

APIC virtualization is a collection of features that can be used to support the virtualization of interrupts and the Advanced Programmable Interrupt Controller (APIC).

When APIC virtualization is enabled, the processor emulates many accesses to the APIC, tracks the state of the virtual APIC, and delivers virtual interrupts — all in VMX non-root operation without a VM exit.

The following are the VM-execution controls relevant to APIC virtualization and virtual interrupts

- **Virtual-interrupt Delivery.** This controls enables the evaluation and delivery of pending virtual interrupts. It also enables the emulation of writes (memory-mapped or MSR-based, as enabled) to the APIC registers that control interrupt prioritization.
- **Use TPR Shadow.** This control enables emulation of accesses to the APIC's task-priority register (TPR) via CR8 and, if enabled, via the memory-mapped or MSR-based interfaces.
- **Virtualize APIC Accesses.** This control enables virtualization of memory-mapped accesses to the APIC by causing VM exits on accesses to a VMM-specified APIC-access page. Some of the other controls, if set, may cause some of these accesses to be emulated rather than causing VM exits.
- **Virtualize x2APIC Mode.** This control enables virtualization of MSR-based accesses to the APIC.
- **APIC-register Virtualization.** This control allows memory-mapped and MSR-based reads of most APIC registers (as enabled) by satisfying them from the virtual-APIC page. It directs memory-mapped writes to the APIC-access page to the virtual-APIC page, following them by VM exits for VMM emulation.
- **Process Posted Interrupts.** This control allows software to post virtual interrupts in a data structure and send a notification to another logical processor; upon receipt of the notification, the target processor will process the posted interrupts by copying them into the virtual-APIC page.

Note: Intel® APIC Virtualization Technology may not be available on all SKUs.

Intel® APIC Virtualization specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 3* available at:

http://www.intel.com/products/processor/manuals

3.2 Security Technologies

3.2.1 Intel® Trusted Execution Technology (Intel® TXT)

Intel® Trusted Execution Technology (Intel® TXT) defines platform-level enhancements that provide the building blocks for creating trusted platforms.

The Intel® TXT platform helps to provide the authenticity of the controlling environment such that those wishing to rely on the platform can make an appropriate trust decision. The Intel® TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

Another aspect of the trust decision is the ability of the platform to resist attempts to change the controlling environment. The Intel® TXT platform will resist attempts by software processes to change the controlling environment or bypass the bounds set by the controlling environment.

Intel® TXT is a set of extensions designed to provide a measured and controlled launch of system software that will then establish a protected environment for itself and any additional software that it may execute.

These extensions enhance two areas:

- The launching of the Measured Launched Environment (MLE).

- The protection of the MLE from potential corruption.

The enhanced platform provides these launch and control interfaces using Safer Mode Extensions (SMX).

The SMX interface includes the following functions:

- Measured/Verified launch of the MLE.
- Mechanisms to ensure the previous measurement is protected and stored in a secure location.
- Protection mechanisms that allow the MLE to control attempts to modify itself.

The processor also offers additional enhancements to System Management Mode (SMM) architecture for enhanced security and performance. The processor provides new MSRs to:

- Enable a second SMM range
- Enable SMM code execution range checking
- Select whether SMM Save State is to be written to legacy SMRAM or to MSRs
- Determine if a thread is going to be delayed entering SMM
- Determine if a thread is blocked from entering SMM
- Targeted SMI, enable/disable threads from responding to SMIs, both VLWs and IPI

For the previous features, BIOS should test the associated capability bit before attempting to access any of the previous registers.

For more information, refer to the Intel® Trusted Execution Technology Measured Launched Environment Programming Guide at:

http://www.intel.com/content/www/us/en/software-developers/intel-txt-software-development-guide.html

Note: Intel® TXT Technology may not be available on all SKUs.

3.2.2 Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)

The processor supports Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) that are a set of Single Instruction Multiple Data (SIMD) instructions that enable fast and secure data encryption and decryption based on the Advanced Encryption Standard (AES). Intel® AES-NI are valuable for a wide range of cryptographic applications, such as applications that perform bulk encryption/ decryption, authentication, random number generation, and authenticated encryption. AES is broadly accepted as the standard for both government and industry applications, and is widely deployed in various protocols.

Intel® AES-NI consists of six Intel® SSE instructions. Four instructions, AESENC, AESENCLAST, AESDEC, and AESDELAST facilitate high performance AES encryption and decryption. The other two, AESIMC and AESKEYGENASSIST, support the AES key expansion procedure. Together, these instructions provide full hardware for supporting AES; offering security, high performance, and a great deal of flexibility.

This generation of the processor has increased the performance of the Intel® AES-NI significantly compared to previous products.

The Intel® AES-NI specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2* available at:

http://www.intel.com/products/processor/manuals

Note: Intel® AES-NI Technology may not be available on all SKUs.

3.2.3 PCLMULQDQ (Perform Carry-Less Multiplication Quad word) Instruction

The processor supports the carry-less multiplication instruction, PCLMULQDQ. PCLMULQDQ is a Single Instruction Multiple Data (SIMD) instruction that computes the 128-bit carry-less multiplication of two 64-bit operands without generating and propagating carries. Carry-less multiplication is an essential processing component of several cryptographic systems and standards. Hence, accelerating carry-less multiplication can significantly contribute to achieving high speed secure computing and communication.

PCLMULQDQ specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2*. Available at:

http://www.intel.com/products/processor/manuals

3.2.4 Intel® Secure Key

The processor supports Intel® Secure Key (formerly known as Digital Random Number Generator (DRNG), a software visible random number generation mechanism supported by a high quality entropy source. This capability is available to programmers through the RDRAND instruction. The resultant random number generation capability is designed to comply with existing industry standards in this regard (ANSI X9.82 and NIST SP 800-90).

Some possible usages of the RDRAND instruction include cryptographic key generation as used in a variety of applications, including communication, digital signatures, secure storage, and so on.

RDRAND specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2* available at:

http://www.intel.com/products/processor/manuals

3.2.5 Execute Disable Bit

The Execute Disable Bit allows memory to be marked as non-executable when combined with a supporting operating system. If code attempts to run in non-executable memory, the processor raises an error to the operating system. This feature can prevent some classes of viruses or worms that exploit buffer overrun vulnerabilities and can, thus, help improve the overall security of the system.

3.2.6 Intel® Boot Guard Technology

Intel® Boot Guard technology is a part of boot integrity protection technology. Boot Guard can help protect the platform boot integrity by preventing execution of unauthorized boot blocks. With Intel® Boot Guard, platform manufacturers can create boot policies such that invocation of an unauthorized (or untrusted) boot block will trigger the platform protection per the manufacturer's defined policy.

With verification based in the hardware, Intel® Boot Guard extends the trust boundary of the platform boot process down to the hardware level.

Intel® Boot Guard accomplishes this by:

- Providing of hardware-based Static Root of Trust for Measurement (S-RTM) and the Root of Trust for Verification (RTV) using Intel architectural components.
- Providing of architectural definition for platform manufacturer Boot Policy.
- Enforcing of manufacture provided Boot Policy using Intel architectural components.

Benefits of this protection is that Intel® Boot Guard can help maintain platform integrity by preventing re-purposing of the manufacturer’s hardware to run an unauthorized software stack.

Note: Intel® Boot Guard availability may vary between the different SKUs.

3.2.7 Intel® Supervisor Mode Execution Protection (SMEP)

Intel® Supervisor Mode Execution Protection (SMEP) is a mechanism that provides the next level of system protection by blocking malicious software attacks from user mode code when the system is running in the highest privilege level. This technology helps to protect from virus attacks and unwanted code from harming the system. For more information, refer to *Intel® 64 Architectures Software Developer’s Manual, Volume 3* at:
http://www.intel.com/products/processor/manuals

3.2.8 Intel® Supervisor Mode Access Protection (SMAP)

Intel® Supervisor Mode Access Protection (SMAP) is a mechanism that provides next level of system protection by blocking a malicious user from tricking the operating system into branching off user data. This technology shuts down very popular attack vectors against operating systems.

For more information, refer to the *Intel® 64 Architectures Software Developer’s Manual, Volume 3 at*:

http://www.intel.com/products/processor/manuals

3.2.9 Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)

The Secure Hash Algorithm (SHA) is one of the most commonly employed cryptographic algorithms. Primary usages of SHA include data integrity, message authentication, digital signatures, and data de-duplication. As the pervasive use of

security solutions continues to grow, SHA can be seen in more applications now than ever. The Intel® SHA Extensions are designed to improve the performance of these compute-intensive algorithms on Intel® architecture-based processors.

The Intel® SHA Extensions are a family of seven instructions based on the Intel® Streaming SIMD Extensions (Intel® SSE) that are used together to accelerate the performance of processing SHA-1 and SHA-256 on Intel® architecture-based processors. Given the growing importance of SHA in our everyday computing devices, the new instructions are designed to provide a needed boost of performance to hashing a single buffer of data. The performance benefits will not only help improve responsiveness and lower power consumption for a given application, they may enable developers to adopt SHA in new applications to protect data while delivering to their user experience goals. The instructions are defined in a way that simplifies their mapping into the algorithm processing flow of most software libraries, thus enabling easier development.

More information on Intel® SHA can be found at:

http://software.intel.com/en-us/articles/intel-sha-extensions

3.2.10 User Mode Instruction Prevention (UMIP)

User Mode Instruction Prevention (UMIP) provides additional hardening capability to the OS kernel by allowing certain instructions to execute only in supervisor mode (Ring 0).

If the OS opt-in to use UMIP, the following instruction are enforced to run in supervisor mode:

- **SGDT**: Store the GDTR register value
- **SIDT**: Store the IDTR register value
- **SLDT**: Store the LDTR register value
- **SMSW**: Store Machine Status Word
- **STR**: Store the TR register value

An attempt at such execution in user mode causes a general protection exception (#GP).

UMIP specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 3* available at:

http://www.intel.com/products/processor/manuals

3.2.11 Read Processor ID (RDPID)

A companion instruction that returns the current logical processor's ID and provides a faster alternative to using the RDTSCP instruction.

RDPID specifications and functional descriptions are included in the *Intel® 64 Architectures Software Developer's Manual, Volume 2* available at:

http://www.intel.com/products/processor/manuals

3.3 Power and Performance Technologies

3.3.1 Intel® Smart Cache Technology

- The Intel® Smart Cache Technology is a shared Last Level Cache (LLC).
- The LLC may also be referred to as a third level cache.
- The LLC is shared between all IA cores as well as the Processor Graphics.
- The first and second level caches are not shared between physical cores and each physical core has a separate set of caches.
- The size of the LLC is SKU specific with a maximum of 2 MB per physical core and is a 16 way associative cache.

3.3.2 IA Core Level 1 and Level 2 Caches

The first-level cache is divided into a data cache and an instruction cache. The processor first level cache size is 48 KB for data and 32 KB for instructions. The first level cache is an eight way associative cache.

The second level cache holds both data and instructions. It is also referred to as mid-level cache or MLC. The processor second level cache size is 512 KB and is an eight way associative cache.

Figure 3-2. Processor Cache Hierarchy

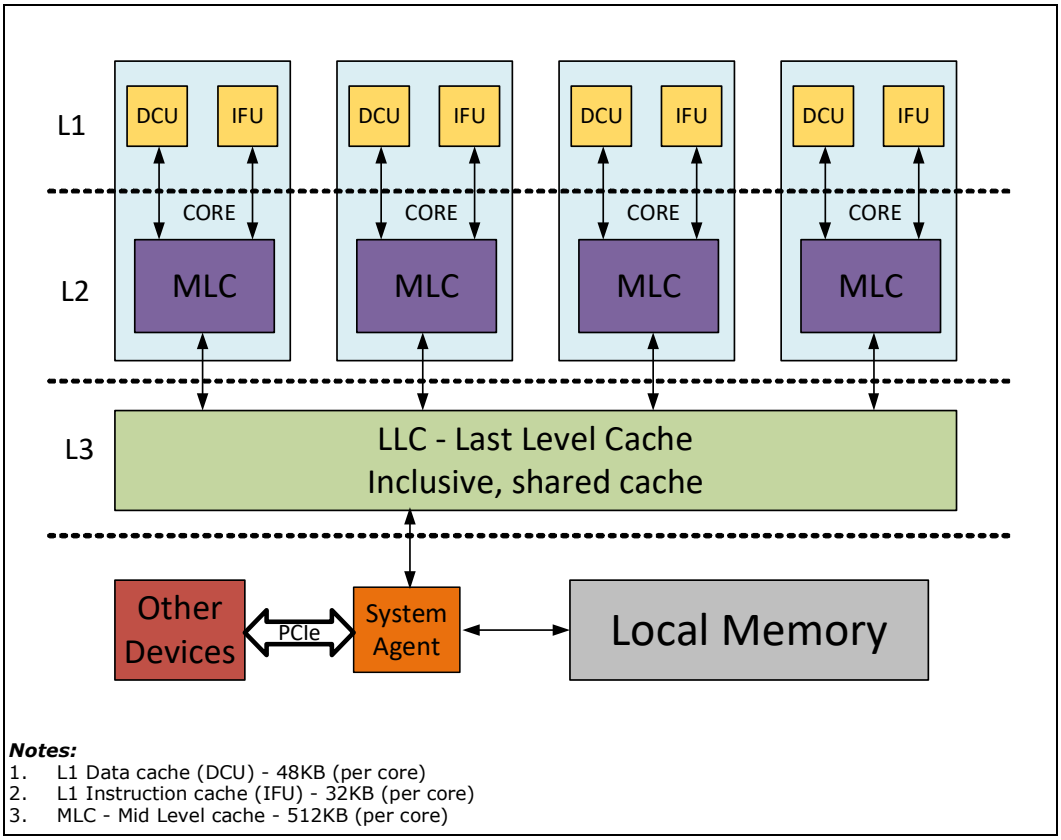

Notes:
1. L1 Data cache (DCU) - 48KB (per core)
2. L1 Instruction cache (IFU) - 32KB (per core)
3. MLC - Mid Level cache - 512KB (per core)

3.3.3 Intel® Turbo Boost Technology 2.0

The Intel® Turbo Boost Technology 2.0 allows the processor IA core / processor graphics core to opportunistically and automatically run faster than the processor IA core base frequency / processor graphics base frequency if it is operating below power, temperature, and current limits. The Intel® Turbo Boost Technology 2.0 feature is designed to increase performance of both multi-threaded and single-threaded workloads.

Compared with previous generation products, Intel® Turbo Boost Technology 2.0 will increase the ratio of application power towards TDP and also allows to increase power above TDP as high as PL2 for short periods of time. Thus, thermal solutions and platform cooling that are designed to less than thermal design guidance might experience thermal and performance issues since more applications will tend to run at the maximum power limit for significant periods of time.

Note: Intel® Turbo Boost Technology 2.0 may not be available on all SKUs.

3.3.3.1 Intel® Turbo Boost Technology 2.0 Power Monitoring

When operating in turbo mode, the processor monitors its own power and adjusts the processor and graphics frequencies to maintain the average power within limits over a thermally significant time period. The processor estimates the package power for all components on package. In the event that a workload causes the temperature to exceed program temperature limits, the processor will protect itself using the Adaptive Thermal Monitor.

3.3.3.2 Intel® Turbo Boost Technology 2.0 Power Control

Illustration of Intel® Turbo Boost Technology 2.0 power control is shown in the following sections and figures. Multiple controls operate simultaneously allowing customization for multiple system thermal and power limitations. These controls allow for turbo optimizations within system constraints and are accessible using MSR, MMIO, and PECI interfaces.

3.3.3.3 Intel® Turbo Boost Technology 2.0 Frequency

To determine the highest performance frequency amongst active processor IA cores, the processor takes the following into consideration:

- The number of processor IA cores operating in the C0 state.
- The estimated processor IA core current consumption and I_{CCMax} settings.
- The estimated package prior and present power consumption and turbo power limits.
- The package temperature and thermal limits.
- Required operating voltage and voltage limits

Any of these factors can affect the maximum frequency for a given workload. If the power, current, or thermal limit is reached, the processor will automatically reduce the frequency to stay within its TDP limit. Turbo processor frequencies are only active if the operating system is requesting the P0 state. For more information on P-states and C-states, refer Chapter 4, “Power Management”.

3.3.4 Intel® Turbo Boost Max Technology 3.0

The Intel® Turbo Boost Max Technology 3.0 (ITBMT 3.0) grants a different maximum Turbo frequency for individual processor cores.

To enable ITBMT 3.0 the processor exposes individual core capabilities; including diverse maximum turbo frequencies.

An operating system that allows for varied per core frequency capability can then maximize power savings and performance usage by assigning tasks to the faster cores, especially on low core count workloads.

Processors enabled with these capabilities can also allow software (most commonly a driver) to override the maximum per-core Turbo frequency limit and notify the operating system via an interrupt mechanism.

For more information on the Intel® Turbo Boost Max 3.0 Technology, refer

http://www.intel.com/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-max-technology.html

Note: Intel® Turbo Boost Max 3.0 Technology may not be available on all SKUs.

3.3.5 Power Aware Interrupt Routing (PAIR)

The processor includes enhanced power-performance technology that routes interrupts to threads or processor IA cores based on their sleep states. As an example, for energy savings, it routes the interrupt to the active processor IA cores without waking the deep idle processor IA cores. For performance, it routes the interrupt to the idle (C1) processor IA cores without interrupting the already heavily loaded processor IA cores. This enhancement is mostly beneficial for high-interrupt scenarios like Gigabit LAN, WLAN peripherals, and so on.

3.3.6 Intel® Hyper-Threading Technology (Intel® HT Technology)

The processor supports Intel® Hyper-Threading Technology (Intel® HT Technology) that allows an execution processor IA core to function as two logical processors. While some execution resources such as caches, execution units, and buses are shared, each logical processor has its own architectural state with its own set of general-purpose registers and control registers. This feature should be enabled using the BIOS and requires operating system support.

Intel recommends enabling Intel® Hyper-Threading Technology with Microsoft* Windows* 7 or newer and disabling Intel® Hyper-Threading Technology using the BIOS for all previous versions of Windows* operating systems. For more information on Intel® Hyper-Threading Technology, refer http://www.intel.com/technology/platform-technology/hyper-threading/

Note: Intel® HT Technology may not be available on all SKUs.

3.3.7 Intel® Thermal Velocity Boost (Intel® TVB)

Intel® Thermal Velocity Boost (Intel® TVB) allows the processor IA core to opportunistically and automatically increase the Intel® Turbo Boost Technology 2.0 frequency by up to two speed bins whenever processor temperature allows. The Intel® Thermal Velocity Boost feature is designed to increase performance of both multi-threaded and single-threaded workloads.

Note: Intel® Thermal Velocity Boost (TVB) may not be available on all SKUs.

3.3.8 Enhanced Intel SpeedStep® Technology

Enhanced Intel SpeedStep® Technology enables OS to control and select P-state. The following are the key features of Enhanced Intel SpeedStep® Technology:

- Multiple frequency and voltage points for optimal performance and power efficiency. These operating points are known as P-states.
- Frequency selection is software controlled by writing to processor MSRs. The voltage is optimized based on the selected frequency and the number of active processor IA cores.
  - — Once the voltage is established, the PLL locks on to the target frequency.
  - — All active processor IA cores share the same frequency and voltage. In a multi-core processor, the highest frequency P-state requested among all active IA cores is selected.
  - — Software-requested transitions are accepted at any time. If a previous transition is in progress, the new transition is deferred until the previous transition is completed.
- The processor controls voltage ramp rates internally to ensure glitch-free transitions.

Note: Due to a low transition latency between P-states, a significant number of transitions per-second is possible.

3.3.9 Intel® Speed Shift Technology

Intel® Speed Shift Technology is an energy efficient method of frequency control by the hardware rather than relying on OS control. OS is aware of available hardware P-states and request a desired P-state or it can let the hardware determine the P-state. The OS request is based on its workload requirements and awareness of processor capabilities. Processor decision is based on the different system constraints for example: Workload demand, thermal limits while taking into consideration the minimum and maximum levels and activity window of performance requested by the Operating System.For more details refer to the *Intel® 64 Architectures Software Developer's Manual, Volume 3* available at:

http://www.intel.com/products/processor/manuals

3.3.10 Intel® Advanced Vector Extensions 2 (Intel® AVX2)

Intel® Advanced Vector Extensions 2.0 (Intel® AVX2) is the latest expansion of the Intel instruction set. Intel® AVX2 extends the Intel® Advanced Vector Extensions (Intel® AVX) with 256-bit integer instructions, floating-point fused multiply add (FMA) instructions, and gather operations. The 256-bit integer vectors benefit math, codec,

image, and digital signal processing software. FMA improves performance in face detection, professional imaging, and high performance computing. Gather operations increase vectorization opportunities for many applications. In addition to the vector extensions, this generation of Intel processors adds new bit manipulation instructions useful in compression, encryption, and general purpose software.
For more information on Intel® AVX, refer http://www.intel.com/software/avx

Intel® Advanced Vector Extensions (Intel® AVX) are designed to achieve higher throughput to certain integer and floating point operation. Due to varying processor power characteristics, utilizing AVX instructions may cause a) parts to operate below the base frequency b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software and system configuration and user should consult the system manufacturer for more information.

Intel® Advanced Vector Extensions refers to Intel® AVX, Intel® AVX2 or Intel® AVX-512.

For more information on Intel® AVX, refer https://software.intel.com/en-us/isa-extensions/intel-avx.

Note: Intel® AVX and AVX2 Technologies may not be available on all SKUs.

3.3.11 Intel® Advanced Vector Extensions 512 Bit (Intel® AVX-512)

Intel® AVX support is widened to 512 bit SIMD operations. Programs can pack eight double precision and sixteen single precision floating numbers within the 512-bit vectors, as well as eight 64-bit and sixteen 32-bit integers. This enables processing of twice the number of data elements that Intel® AVX/AVX2 can process with a single instruction and four times the capabilities of Intel® SSE.

Intel® AVX-512 instructions are important because they open up higher performance capabilities for the most demanding computational tasks. Intel® AVX-512 instructions offer the highest degree of compiler support by including an unprecedented level of richness in the design of the instruction capabilities.

Intel® AVX-512 features include 32 vector registers each 512-bit wide and eight dedicated mask registers. Intel® AVX-512 is a flexible instruction set that includes support for broadcast, embedded masking to enable predication, embedded floating point rounding control, embedded floating-point fault suppression, scatter instructions, high speed math instructions, and compact representation of large displacement values.

Intel® AVX-512 offers a level of compatibility with Intel® AVX which is stronger than prior transitions to new widths for SIMD operations. Unlike Intel® SSE and Intel® AVX which cannot be mixed without performance penalties, the mixing of Intel® AVX and Intel® AVX-512 instructions is supported without penalty. Intel® AVX registers YMM0-YMM15 map into Intel® AVX-512 registers ZMM0-ZMM15 (in x86-64 mode), very much like Intel® SSE registers map into Intel® AVX registers. Therefore, in processors with Intel® AVX-512 support, Intel® AVX and Intel® AVX2 instructions operate on the lower 128 or 256 bits of the first 16 ZMM registers.

The Intel® AVX-512 instructions are documented in the Intel® Architecture Instruction Set Extensions Programming Reference (future architectures):

https://software.intel.com/sites/default/files/managed/b4/3a/319433-024.pdf

Intel® AVX-512 has multiple extensions that CPUID has been enhanced to expose.

- **AVX512F (Foundation)**: expands most 32-bit and 64-bit based AVX instructions with EVEX coding scheme to support 512-bit registers, operation masks, parameter broadcasting, and embedded rounding and exception control
- **AVX512CD (Conflict Detection)**: efficient conflict detection to allow more loops to be vectorized
- **AVX512BW (Byte and Word)**: extends AVX-512 to cover 8-bit and 16-bit integer operations
- **AVX512DQ (Doubleword and Quadword)**: extends AVX-512 to cover 32-bit and 64-bit integer operations
- **AVX512VL (Vector Length)**: extends most AVX-512 operations to also operate on XMM (128-bit) and YMM (256-bit) registers
- **AVX512IFMA (Integer Fused Multiply-Add)**: fused multiply-add of integers using 52-bit precision
- **AVX512VBMI (Vector Byte Manipulation Instructions)**: adds vector byte permutation instructions which were not present in AVX-512BW
- **AVX512VBMI2 (Vector Byte Manipulation Instructions 2)**: adds byte/word load, store and concatenation with shift
- **VPOPCNTDQ**: count of bits set to 1
- **VPCLMULQDQ**: carry-less multiplication of quadwords
- **AVX-512VNNI (Vector Neural Network Instructions)**: vector instructions for deep learning
- **AVX512GFNI (Galois Field New Instructions)**: vector instructions for calculating Galois Fields
- **AVX512VAES (Vector AES instructions)**: vector instructions for AES coding
- **AVX512BITALG (Bit Algorithms)**: byte/word bit manipulation instructions expanding VPOPCNTDQ

Note: Intel® AVX-512 may not be available on all SKUs.

3.3.12 Intel® 64 Architecture x2APIC

The x2APIC architecture extends the xAPIC architecture that provides key mechanisms for interrupt delivery. This extension is primarily intended to increase processor addressability.

Specifically, x2APIC:

- Retains all key elements of compatibility to the xAPIC architecture:
  - — Delivery modes
  - — Interrupt and processor priorities
  - — Interrupt sources
  - — Interrupt destination types

- Provides extensions to scale processor addressability for both the logical and physical destination modes
- Adds new features to enhance performance of interrupt delivery
- Reduces complexity of logical destination mode interrupt delivery on link based architectures

The key enhancements provided by the x2APIC architecture over xAPIC are the following:

- Support for two modes of operation to provide backward compatibility and extensibility for future platform innovations:
  - In xAPIC compatibility mode, APIC registers are accessed through memory mapped interface to a 4K-Byte page, identical to the xAPIC architecture.
  - In x2APIC mode, APIC registers are accessed through Model Specific Register (MSR) interfaces. In this mode, the x2APIC architecture provides significantly increased processor addressability and some enhancements on interrupt delivery.
- Increased range of processor addressability in x2APIC mode:
  - Physical xAPIC ID field increases from 8 bits to 32 bits, allowing for interrupt processor addressability up to 4G-1 processors in physical destination mode. A processor implementation of x2APIC architecture can support fewer than 32-bits in a software transparent fashion.
  - Logical xAPIC ID field increases from 8 bits to 32 bits. The 32-bit logical x2APIC ID is partitioned into two sub-fields – a 16-bit cluster ID and a 16-bit logical ID within the cluster. Consequently, ((2^20) - 16) processors can be addressed in logical destination mode. Processor implementations can support fewer than 16 bits in the cluster ID sub-field and logical ID sub-field in a software agnostic fashion.
- More efficient MSR interface to access APIC registers:
  - To enhance inter-processor and self-directed interrupt delivery as well as the ability to virtualize the local APIC, the APIC register set can be accessed only through MSR-based interfaces in x2APIC mode. The Memory Mapped IO (MMIO) interface used by xAPIC is not supported in x2APIC mode.
- The semantics for accessing APIC registers have been revised to simplify the programming of frequently-used APIC registers by system software. Specifically, the software semantics for using the Interrupt Command Register (ICR) and End Of Interrupt (EOI) registers have been modified to allow for more efficient delivery and dispatching of interrupts.
- The x2APIC extensions are made available to system software by enabling the local x2APIC unit in the "x2APIC" mode. To benefit from x2APIC capabilities, a new operating system and a new BIOS are both needed, with special support for x2APIC mode.
- The x2APIC architecture provides backward compatibility to the xAPIC architecture and forward extensible for future Intel platform innovations.

Note: Intel® x2APIC Technology may not be available on all SKUs.

For more information, refer the Intel® *64 Architecture x2APIC Specification* at http://www.intel.com/products/processor/manuals/.

3.3.13 Intel® GNA 2.0 (GMM and Neural Network Accelerator) GNA stands for Gaussian Mixture Model and Neural Network Accelerator

The Intel® GNA is used to process speech recognition without user training sequence. The Intel® GNA is designed to unload the processor cores and the system memory with complex speech recognition tasks and improve the speech recognition accuracy. The Intel® GNA is designed to compute millions of Gaussian probability density functions per second without loading the processor cores while maintaining low power consumption.

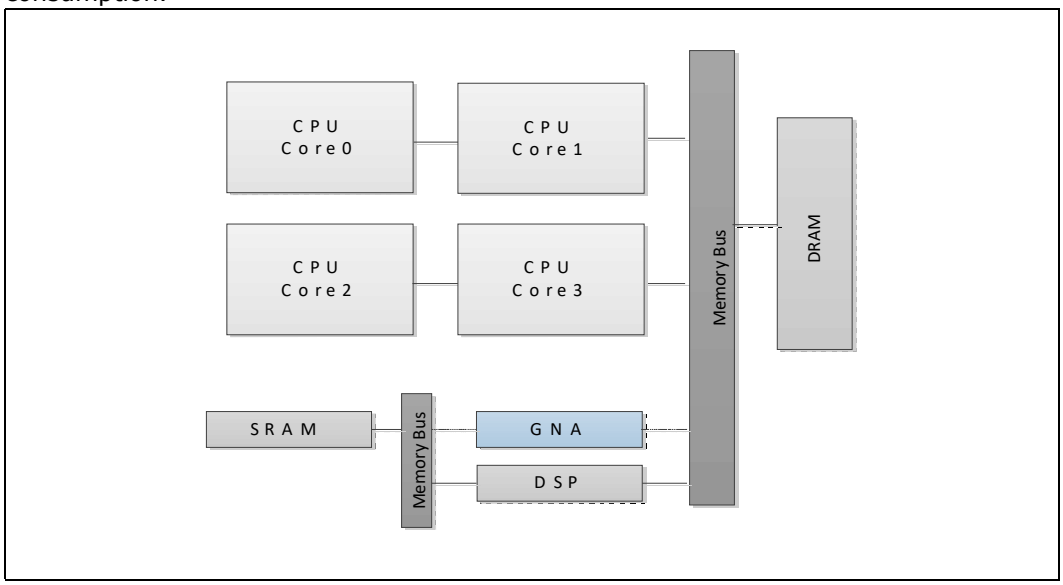

3.3.14 Cache Line Write Back (CLWB)

Writes back to memory the cache line (if dirty) that contains the linear address specified with the memory operand from any level of the cache hierarchy in the cache coherence domain. The line may be retained in the cache hierarchy in non-modified state. Retaining the line in the cache hierarchy is a performance optimization (treated as a hint by hardware) to reduce the possibility of cache miss on a subsequent access. Hardware may choose to retain the line at any of the levels in the cache hierarchy, and in some cases, may invalidate the line from the cache hierarchy. The source operand is a byte memory location.

The CLWB instruction is documented in the Intel® Architecture Instruction Set Extensions Programming Reference (future architectures):

https://software.intel.com/sites/default/files/managed/b4/3a/319433-024.pdf

3.3.15 Ring Interconnect

The Ring is a high speed, wide interconnect that links the processor cores, processor graphics and the System Agent.The Ring shares frequency and voltage with the Last Level Cache (LLC).The Ring's frequency dynamically changes. Its frequency is relative to both processor cores and processor graphics frequencies.

3.3.16 Intel® Dynamic Tuning Technology

Intel® Dynamic Tuning Technology consists of a set of software drivers and applications that allow a system manufacturer to optimize system performance and usability by:

- Dynamically optimize turbo settings of IA processors, power and thermal states of the platform for optimal performance
- Dynamically adjust the processor's peak power based on the current power delivery capability for optimal system usability
- Dynamically mitigate radio frequency interference for better RF throughput.

Note: This technology was previously referred to as Intel® Dynamic Platform and Thermal Framework (DPTF)

3.4 Debug Technologies

3.4.1 Intel® Processor Trace

Intel® Processor Trace (Intel® PT) is a tracing capability added to Intel® Architecture, for use in software debug and profiling. Intel® PT provides the capability for more precise software control flow and timing information, with limited impact to software execution. This provides enhanced ability to debug software crashes, hangs, or other anomalies, as well as responsiveness and short-duration performance issues.

Intel® VTune™ Amplifier for Systems and the Intel® System Debugger are part of Intel® System Studio 2015 (and newer) product, which includes updates for the new debug and trace features, including Intel® PT and Intel® Trace Hub.

Intel® System Studio 2015 is available for download at https://software.intel.com/en-us/system-studio.

An update to the Linux* performance utility, with support for Intel® PT, is available for download at https://github.com/virtuoso/linux-perf/tree/intel_pt. It requires rebuilding the kernel* and the perf utility.

3.5 Deprecated Technologies

- Intel® Memory Protection Extensions (Intel® MPX)
- Branch Monitoring Counters
- Hardware Lock Elision (HLE), part of Intel® TSX-NI
- Intel® Software Guard Extensions (Intel® SGX)

4 Power Management

This chapter provides information on the following power management topics:

- Advanced Configuration and Power Interface (ACPI) States
- Processor IA Core Power Management
- Integrated Memory Controller (IMC) Power Management
- PCI Express* Power Management
- Direct Media Interface (DMI) Power Management
- Processor Graphics Power Management

Notes:

- Package C-State C6 supported when Intel® Xeon® E-2300 processor is paired with an Intel® C250 Series Chipset Family Platform Controller Hub.

Figure 4-1. Processor Power States

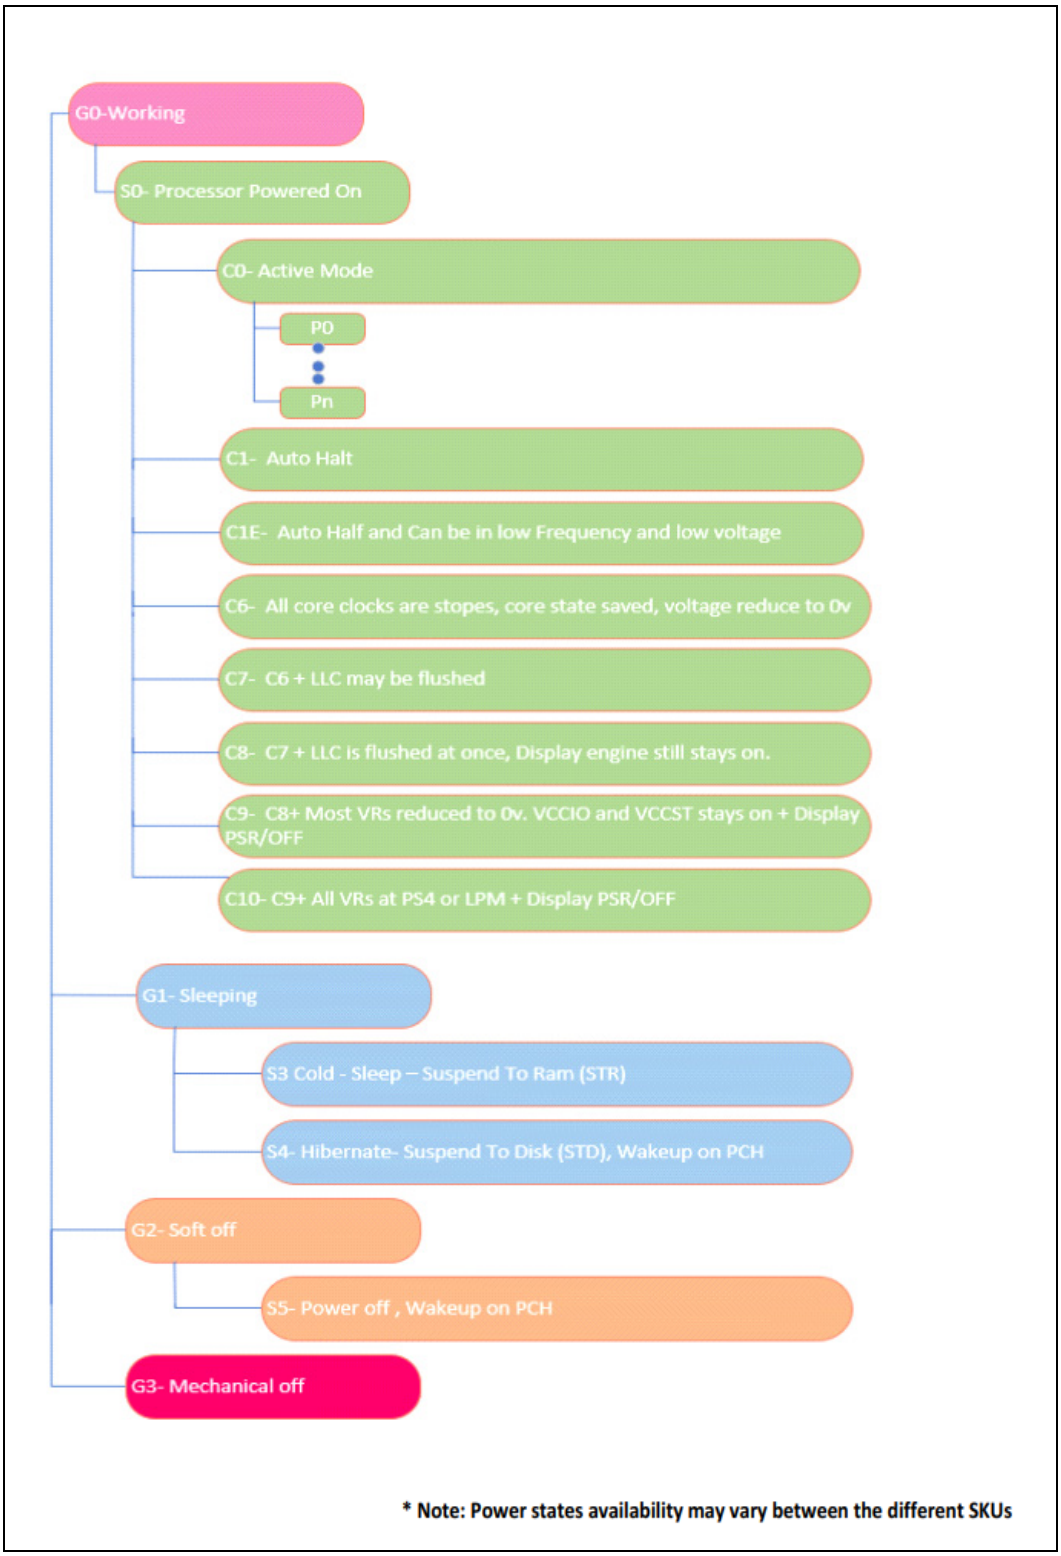

Figure 4-2. Processor Package and IA Core C-States

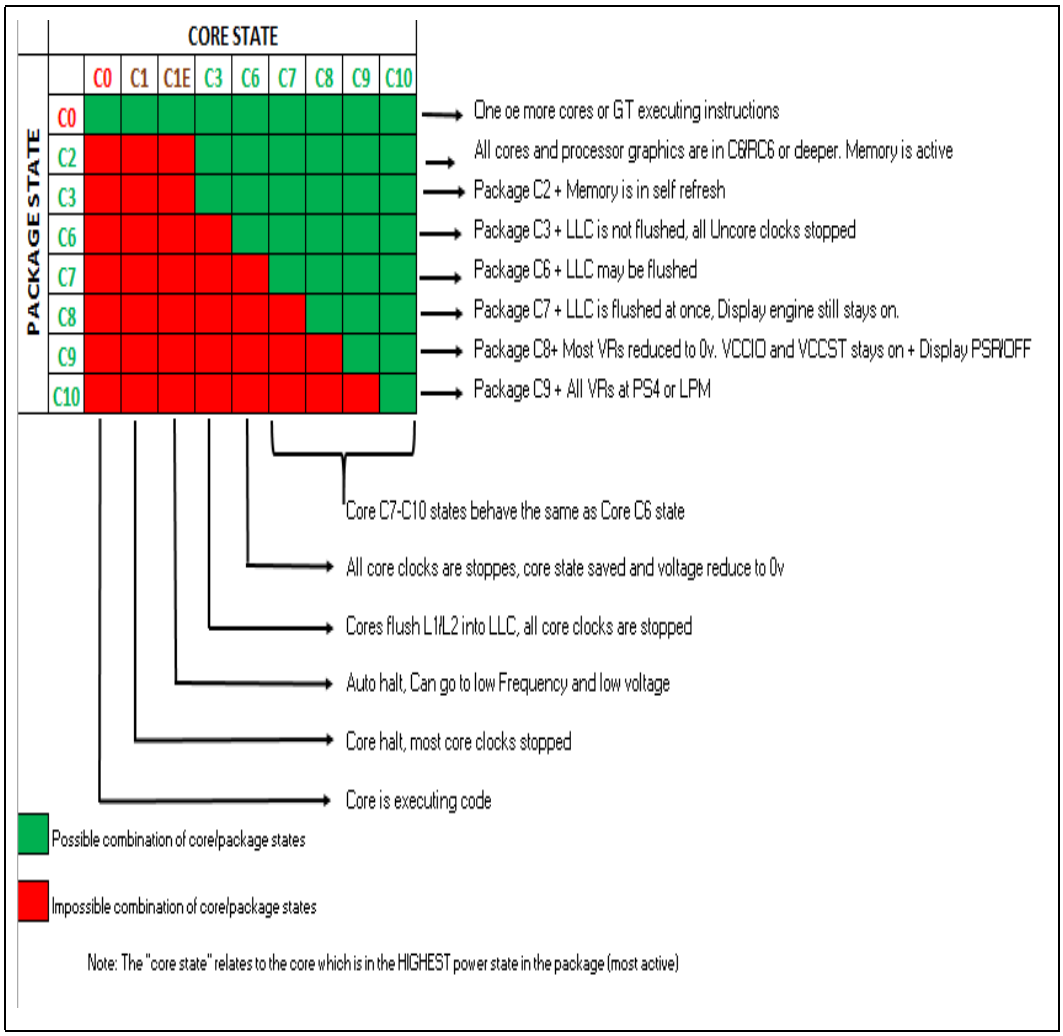

Notes:
1. There are constraints that prevent the system to go deeper.
2. The "core state" relates to the core which is in the HIGEST power state in the package (most active).

4.1 Advanced Configuration and Power Interface (ACPI) States Supported

This section describes the ACPI states supported by the processor.

Table 4-1. System States

State	Description
G0/S0	Full On
G1/S3-Cold	Suspend-to-RAM (STR). Context saved to memory (S3-Hot is not supported by the processor).
G1/S4	Suspend-to-Disk (STD). All power lost (except wake-up on PCH).
G2/S5	Soft off. All power lost (except wake-up on PCH). Total reboot.
G3	Mechanical off. All power removed from system.

Table 4-2. Processor IA Core / Package State Support

State	Description
C0	Active mode, processor executing code.
C1	AutoHALT processor IA core state (package C0 state).
C1E	AutoHALT processor IA core state with lowest frequency and voltage operating point (package C0 state).
C3	Processor IA execution cores in C3 or deeper, flush their L1 instruction cache, L1 data cache, and L2 cache to the LLC shared cache. LLC may be flushed. Clocks are shut off to each core.
C6	Processor IA execution cores in this state save their architectural state before removing core voltage. BCLK is off.
C6R	Processor IA saves its state into DRAM in and power gates the SRAM control module.
C7	Processor IA execution cores in this state behave similarly to the C6 state. If all execution cores request C7, LLC ways may be flushed until it is cleared. If the entire LLC is flushed, voltage will be removed from the LLC.
C8	C7 plus LLC should be flushed.
C9	C8 plus most Uncore voltages at 0 V. IA, GT and SA reduced to 0 V, while Vcc_{IO} stays on.
C10	C9 plus all VRs at PS4 or LPM. 24 MHz clock off

Table 4-3. Integrated Memory Controller (IMC) States

State	Description
Power up	CKE asserted. Active mode.
Pre-charge Power down	CKE de-asserted (not self-refresh) with all banks closed.
Active Power down	CKE de-asserted (not self-refresh) with minimum one bank active.
Self-Refresh	CKE de-asserted using device self-refresh.

Table 4-4. PCI Express* Link States

State	Description
L0	Full on – Active transfer state.
L1	Lowest Active Power Management – Longer exit latency (L1 Substates included)
L2	Deep energy saving state – Longer exit latency
L3	Lowest power state (power-off) – Longest exit latency

Table 4-5. Direct Media Interface (DMI) States

State	Description
L0	Full on – Active transfer state
L1	Lowest Active Power Management – Longer exit latency
L3	Lowest power state (power-off) – Longest exit latency

Table 4-6. G, S, and C Interface State Combinations (Sheet 1 of 2)

Global (G) State	Sleep (S) State	Processor Package (C) State	Processor State	System Clocks	Description
G0	S0	C0	Full On	On	Full On
G0	S0	C1/C1E	Auto-Halt	On	Auto-Halt
G0	S0	C3	Deep Sleep	On	Deep Sleep

Table 4-6. G, S, and C Interface State Combinations (Sheet 2 of 2)

Global (G) State	Sleep (S) State	Processor Package (C) State	Processor State	System Clocks	Description
G0	S0	C6/C7	Deep Power Down	On	Deep Power Down
G0	S0	C8/C9/C10	Off	On	Deeper Power Down
G1	S3	Power off	Off	Off, except RTC	Suspend to RAM
G1	S4	Power off	Off	Off, except RTC	Suspend to Disk
G2	S5	Power off	Off	Off, except RTC	Soft Off
G3	N/A	Power off	Off	Power off	Hard off

4.2 Processor IA Core Power Management

While executing code, Enhanced Intel SpeedStep® Technology and Intel® Speed Shift Technology optimizes the processor's IA core frequency and voltage based on workload. Each frequency and voltage operating point is defined by ACPI as a P-state. When the processor is not executing code, it is idle. A low-power idle state is defined by ACPI as a C-state. In general, deeper power C-states have longer entry and exit latencies.

4.2.1 OS/HW Controlled P-states

4.2.1.1 Enhanced Intel SpeedStep® Technology

Enhanced Intel SpeedStep® Technology enables OS to control and select P-state. The following are the key features of Enhanced Intel SpeedStep® Technology:

- Multiple frequency and voltage points for optimal performance and power efficiency. These operating points are known as P-states.
- Frequency selection is software controlled by writing to processor MSRs. The voltage is optimized based on the selected frequency and the number of active processor IA cores.
  - Once the voltage is established, the PLL locks on to the target frequency.
  - All active processor IA cores share the same frequency and voltage. In a multi-core processor, the highest frequency P-state requested among all active IA cores is selected.
  - Software-requested transitions are accepted at any time. If a previous transition is in progress, the new transition is deferred until the previous transition is completed.
- The processor controls voltage ramp rates internally to ensure glitch-free transitions.
- Because there is low transition latency between P-states, a significant number of transitions per-second are possible.

4.2.1.2 Intel® Speed Shift Technology

Intel® Speed Shift Technology is an energy efficient method of frequency control by the hardware rather than relying on OS control. OS is aware of available hardware P-states and request a desired P-state or it can let Hardware determine the P-state. The OS request is based on its workload requirements and awareness of processor capabilities.

Processor decision is based on the different system constraints for example: Workload demand, thermal limits while taking into consideration the minimum and maximum levels and activity window of performance requested by the operating system.

For more details, refer to the following document:

- Intel® 64 and IA-32 Architectures Software Developer's Manual (SDM), Volume 3B

4.2.2 Low-Power Idle States

When the processor is idle, low-power idle states (C-states) are used to save power. More power savings actions are taken for numerically higher C-states. However, deeper C-states have longer exit and entry latencies. Resolution of C-states occur at the thread, processor IA core, and processor package level. Thread-level C-states are available if Intel® Hyper-Threading Technology is enabled.

Caution: Long term reliability cannot be assured unless all the Low-Power Idle States are enabled.

Figure 4-3. Idle Power Management Breakdown of the Processor IA Cores

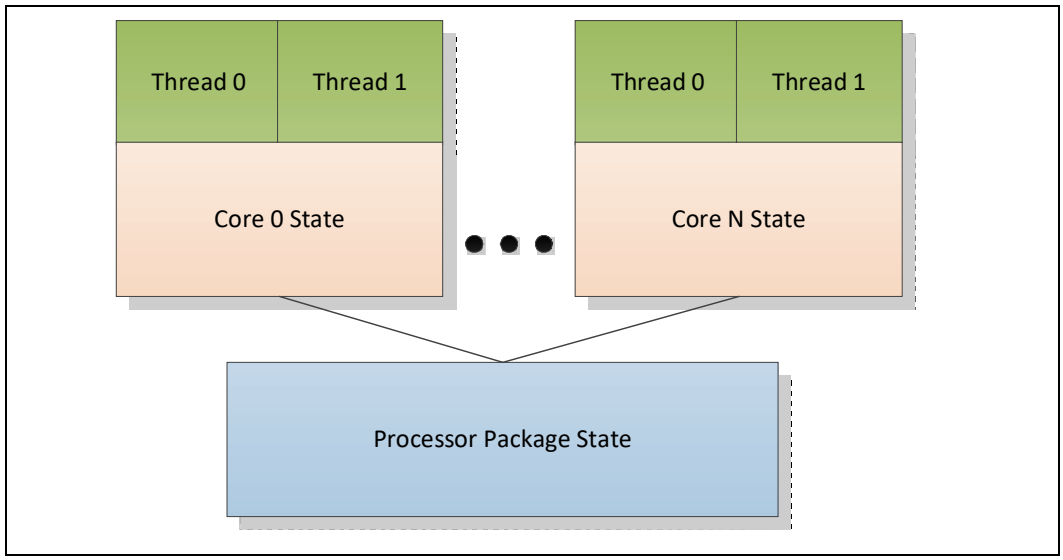

While individual threads can request low-power C-states, power saving actions only take place once the processor IA core C-state is resolved. processor IA core C-states are automatically resolved by the processor. For thread and processor IA core C-states, a transition to and from C0 state is required before entering any other C-state.

4.2.3 Requesting Low-Power Idle States

The primary software interfaces for requesting low-power idle states are through the MWAIT instruction with sub-state hints and the HLT instruction (for C1 and C1E). However, software may make C-state requests using the legacy method of I/O reads from the ACPI-defined processor clock control registers, referred to as P_LVLx. This method of requesting C-states provides legacy support for operating systems that initiate C-state transitions using I/O reads.

For legacy operating systems, P_LVLx I/O reads are converted within the processor to the equivalent MWAIT C-state request. Therefore, P_LVLx reads do not directly result in I/O reads to the system. The feature, known as I/O MWAIT redirection, should be enabled in the BIOS.

The BIOS can write to the C-state range field of the PMG_IO_CAPTURE MSR to restrict the range of I/O addresses that are trapped and emulate MWAIT like functionality. Any P_LVLx reads outside of this range do not cause an I/O redirection to MWAIT(Cx) like request. They fall through like a normal I/O instruction.

When P_LVLx I/O instructions are used, MWAIT sub-states cannot be defined. The MWAIT sub-state is always zero if I/O MWAIT redirection is used. By default, P_LVLx I/O redirections enable the MWAIT 'break on EFLAGS.IF' feature that triggers a wake up on an interrupt, even if interrupts are masked by EFLAGS.IF.

4.2.4 Processor IA Core C-State Rules

The following are general rules for all processor IA core C-states, unless specified otherwise:

- A processor IA core C-State is determined by the lowest numerical thread state (such as Thread 0 requests C1E while Thread 1 requests C3 state, resulting in a processor IA core C1E state). Refer the *G, S, and C Interface State Combinations* table.
- A processor IA core transitions to C0 state when:
  - An interrupt occurs.
  - There is an access to the monitored address if the state was entered using an MWAIT/Timed MWAIT instruction.
  - The deadline corresponding to the Timed MWAIT instruction expires.
- An interrupt directed toward a single thread wakes up only that thread.
- If any thread in a processor IA core is active (in C0 state), the core's C-state will resolve to C0.
- Any interrupt coming into the processor package may wake any processor IA core.
- A system reset re-initializes all processor IA cores.

Processor IA core C0 State

The normal operating state of a processor IA core where code is being executed.

At least one IA core is in C0. Processor Graphics is RC0 (Graphics active state) or RC6 (Graphics core power down state)

Processor IA core C1/C1E State

C1/C1E is a low-power state entered when all threads within a processor IA core execute a HLT or MWAIT(C1/C1E) instruction.

A System Management Interrupt (SMI) handler returns execution to either Normal state or the C1/C1E state. Refer the *Intel® 64 and IA-32 Architectures Software Developer's Manual* for more information.

While a processor IA core is in C1/C1E state, it processes bus snoops and snoops from other threads. For more information on C1E, Refer Section 4.2.5, "Package C-States".

Processor IA core C3 State

Individual threads of a processor IA core can enter the C3 state by initiating a P_LVL2 I/O read to the P_BLK or an MWAIT(C3) instruction. A processor IA core in C3 state flushes the contents of its L1 instruction cache, L1 data cache, and L2 cache to the shared LLC, while maintaining its architectural state. All processor IA core clocks are stopped at this point. Because the processor IA core's caches are flushed, the processor does not wake any processor IA core that is in the C3 state when either a snoop is detected or when another processor IA core accesses cacheable memory.

Processor IA core C6 State

Individual threads of a processor IA core can enter the C6 state by initiating a P_LVL3 I/O read or an MWAIT(C6) instruction. Before entering processor IA core C6 state, the processor IA core will save its architectural state to a dedicated SRAM. Once complete, a processor IA core will have its voltage reduced to zero volts. During exit, the processor IA core is powered on and its architectural state is restored.

Processor IA core C7-C10 States

Individual threads of a processor IA core can enter the C7, C8, C9, or C10 state by initiating a P_LVL4, P_LVL5, P_LVL6, P_LVL7 I/O read (respectively) to the P_BLK or by an MWAIT(C7/C8/C9/C10) instruction. The processor IA core C7-C10 state exhibits the same behavior as the processor IA core C6 state.

C-State Auto-Demotion

In general, deeper C-states, such as C6 or C7, have long latencies and have higher energy entry/exit costs. The resulting performance and energy penalties become significant when the entry/exit frequency of a deeper C-state is high. Therefore, incorrect or inefficient usage of deeper C-states have a negative impact on battery life and idle power. To increase residency and improve battery life and idle power in deeper C-states, the processor supports C-state auto-demotion.

There are two C-State auto-demotion options:

- C7/C6 to C3
- C7/C6/C3 To C1

The decision to demote a processor IA core from C6/C7 to C3 or C3/C6/C7 to C1 is based on each processor IA core's immediate residency history. Upon each processor IA core C6/C7 request, the processor IA core C-state is demoted to C3 or C1 until a sufficient amount of residency has been established. At that point, a processor IA core is allowed to go into C3/C6 or C7. Each option can be run concurrently or individually. If the interrupt rate experienced on a processor IA core is high and the processor IA core is rarely in a deep C-state between such interrupts, the processor IA core can be demoted to a C3 or C1 state. A higher interrupt pattern is required to demote a processor IA core to C1 as compared to C3.

This feature is disabled by default. BIOS should enable it in the PMG_CST_CONFIG_CONTROL register. The auto-demotion policy is also configured by this register.

4.2.5 Package C-States

The processor supports C0, C3, C6, C6DRAM, C7, C8, C9, and C10 package states. The following is a summary of the general rules for package C-state entry. These apply to all package C-states, unless specified otherwise:

- A package C-state request is determined by the lowest numerical processor IA core C-state amongst all processor IA cores.
- A package C-state is automatically resolved by the processor depending on the processor IA core idle power states and the status of the platform components.
  - Each processor IA core can be at a lower idle power state than the package if the platform does not grant the processor permission to enter a requested package C-state.
  - The platform may allow additional power savings to be realized in the processor.
  - For package C-states, the processor is not required to enter C0 before entering any other C-state.
  - Entry into a package C-state may be subject to auto-demotion – that is, the processor may keep the package in a deeper package C-state then requested by the operating system if the processor determines, using heuristics, that the deeper C-state results in better power/performance.

The processor exits a package C-state when a break event is detected. Depending on the type of break event, the processor does the following:

- If a processor IA core break event is received, the target processor IA core is activated and the break event message is forwarded to the target processor IA core.
  - If the break event is not masked, the target processor IA core enters the processor IA core C0 state and the processor enters package C0.
  - If the break event is masked, the processor attempts to re-enter its previous package state.
- If the break event was due to a memory access or snoop request,
  - But the platform did not request to keep the processor in a higher package C-state, the package returns to its previous C-state.
  - And the platform requests a higher power C-state, the memory access or snoop request is serviced and the package remains in the higher power C-state.

Figure 4-4. Package C-State Entry and Exit

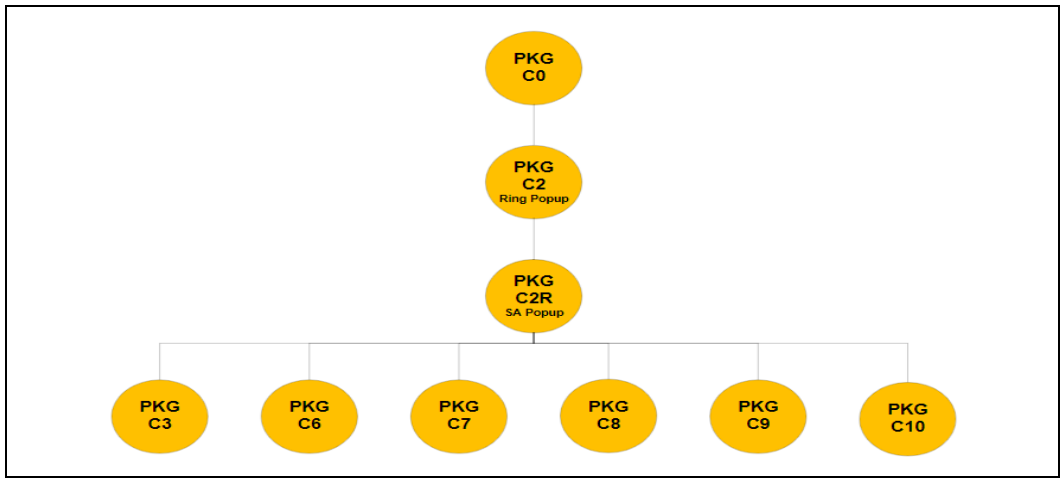

Package C0

This is the normal operating state for the processor. The processor remains in the normal state when at least one of its processor IA cores is in the C0 or C1 state or when the platform has not granted permission to the processor to go into a low-power state. Individual processor IA cores may be in deeper power idle states while the package is in C0 state.

Package C2 State

Package C2 state is an internal processor state that cannot be explicitly requested by software. A processor enters Package C2 state when either:

- All processor IA cores have requested a C3 or deeper power state and all graphics processor IA cores requested are in RC6, but constraints (LTR, programmed timer events in the near future, and so forth) prevent entry to any state deeper than C2 state.
- Or, all processor IA cores have requested a C3 or deeper power state and all graphics processor IA cores requested are in RC6 and a memory access request is received. Upon completion of all outstanding memory requests, the processor transitions back into a deeper package C-state.

Package C3 State

A processor enters the package C3 low-power state when:

- At least one processor IA core is in the C3 state.
- The other processor IA cores are in a C3 or deeper power state, and the processor has been granted permission by the platform.
- The platform has not granted a request to a package C6/C7 state or deeper state but has allowed a package C3 state.

In package C3-state, the LLC shared cache is valid.

Package C6 State

A processor enters the package C6 low-power state when:

- At least one processor IA core is in the C6 state.
- The other processor IA cores are in a C6 or deeper power state, and the processor has been granted permission by the platform.
- The platform has not granted a package C7 or deeper request but has allowed a C6 package state.

In package C6 state, all processor IA cores have saved their architectural state and have had their voltages reduced to zero volts. It is possible the LLC shared cache is flushed and turned off in package C6 state.

Package C6DRAM State

The C6DRAM feature saves the processor internal state at package C6 and deeper to DRAM instead of on-die SRAM. When the processor state has been saved to DRAM, the dedicated save/restore SRAM modules are power gated enabling IDLE power saving. The SRAM modules operate on the sustained voltage rail (VccST).

The Memory region used for C6SRAM resides in the Processor Reserved Memory region (PRMRR) which is encrypted and replay protected. The processor issues a Machine Check exception (#MC) if the processor state has been corrupted.

In package C6 state, all processor IA cores have saved their architectural state and

Package C7 State

The processor enters the package C7 low-power state when all processor IA cores are in the C7 or deeper state and the operating system may request that the LLC will be flushed.

processor IA core break events are handled the same way as in package C3 or C6.

Upon exit of the package C7 state, the LLC will be partially enabled once a processor IA core wakes up if it was fully flushed, and will be fully enabled once the processor has stayed out of C7 for a preset amount of time. Power is saved since this prevents the LLC from being re-populated only to be immediately flushed again. Some VRs are reduce to 0 V.

Package C8 State

The processor enters C8 states when the processor IA cores lower numerical state is C8.

The C8 state is similar to C7 state, but in addition, the LLC is flushed in a single step, Vcc and Vcc_{GT} are reduced to 0V. The display engine stays on.

Package C9 State

The processor enters C9 states when the processor IA cores lower numerical state is C9.

Package C9 state is similar to C8 state; the VRs are off, Vcc, Vcc_{GT} and at 0 V, Vcc_{IO} and Vcc_{ST} stays on.

Package C10 State

The processor enters C10 states when the processor IA cores lower numerical state is C10.

Package C10 state is similar to the package C9 state, but in addition the IMVP8 VR is in PS4 low-power state, which is near to shut off of the IMVP8 VR. The Vcc_{IO} is in low-power mode as well.

InstantGo

InstantGo is a platform state. On display time out the OS requests the processor to enter package C10 and platform devices at RTD3 (or disabled) in order to attain low power in idle.

Dynamic LLC Sizing

When all processor IA cores request C7 or deeper C-state, internal heuristics dynamically flushes the LLC. Once the processor IA cores enter a deep C-state, depending on their MWAIT sub-state request, the LLC is either gradually flushed N-ways at a time or flushed all at once. Upon the processor IA cores exiting to C0 state, the LLC is gradually expanded based on internal heuristics.

4.2.6 Package C-States and Display Resolutions

The integrated graphics engine has the frame buffer located in system memory. When the display is updated, the graphics engine fetches display data from system memory. Different screen resolutions and refresh rates have different memory latency requirements. These requirements may limit the deepest Package C-state the processor can enter. Other elements that may affect the deepest Package C-state available are the following:

- Display is on or off
- Single or multiple displays
- Native or non-native resolution
- Panel Self Refresh (PSR) technology

Note: Display resolution is not the only factor influencing the deepest Package C-state the processor can get into. Device latencies, interrupt response latencies, and core C-states are among other factors that influence the final package C-state the processor can enter.

The following table lists display resolutions and deepest available package C-State.The display resolutions are examples using common values for blanking and pixel rate. Actual results will vary. The table shows the deepest possible Package C-state.System workload, system idle, and AC or DC power also affect the deepest possible Package C-state.

4.2.7 DE FREQ and DE States

When pkgc is woken up from pkgc9/10, DE should be moved to DC5 from DC6. DE should not be filled first (since fill will require DC3 state).

CD Clock Frequency	DC6 -> DC5 (Delay) - Ref Clock - 24MHz	DC5 -> DC3 (Delay)
168 MHz	15	~100 µs(ref clock) + ~301 µs(cdclk) + 400 µs(PLL lock worst case) = 801 µs
337 MHz	15	~100 µs(ref clock) + ~150 µs(cdclk) + 400 µs(PLL lock worst case) = 650 µs
450 MHz	15	~100 µs(ref clock) + ~113 µs(cdclk) + 400 µs(PLL lock worst case) = 613 µs
600 MHz	15	~100 µs(ref clock) + ~85 µs(cdclk) + 400 µs(PLL lock worst case) = 585 µs

4.2.7.1 Deepest Package C-State Available

Processor Line	S Processor Line[1,2,3]	
PSR	PSR Enabled[4]	PSR Disabled
Deepest Pkg C State	PC10	PC8

Notes:
1. All Deep states are with Display ON.
2. The deepest package C-state dependents on various factors, including Platform devices, HW configuration and peripheral software.
3. Intel® Xeon® E-2300 processor supporting up to PC6 when paired with Intel® C250 Series Chipset Families Platform Controller Hub.
4. All are referring to 800x600, 1024x768, 1280x1024, 1920x1080, 1920x1200, 1920x1440, 2048x1536, 2560x1600, 2560x1920, 2880x1620, 2880x1800, 3200x1800, 3200x2000, 3840x2160 and 4096x2160 resolutions, up to 60 Hz.

4.3 Integrated Memory Controller (IMC) Power Management

The main memory is power managed during normal operation and in low-power ACPI C-states.

4.3.1 Disabling Unused System Memory Outputs

Any system memory (SM) interface signal that goes to a memory in which it is not connected to any actual memory devices (such as SODIMM connector is unpopulated, or is single-sided) is tri-stated. The benefits of disabling unused SM signals are:

- Reduced power consumption.
- Reduced possible overshoot/undershoot signal quality issues Refer by the processor I/O buffer receivers caused by reflections from potentially un-terminated transmission lines.

When a given rank is not populated, the corresponding control signals (CLK_P/CLK_N/ CKE/ODT/CS) are not driven.

At reset, all rows should be assumed to be populated, until it can be proven that they are not populated. This is due to the fact that when CKE is tri-stated with a DRAMs present, the DRAMs are not ensured to maintain data integrity. CKE tri-state should be enabled by BIOS where appropriate, since at reset all rows should be assumed to be populated.

4.3.2 DRAM Power Management and Initialization

The processor implements extensive support for power management on the memory interface. Each channel drives 4 CKE pins, one per rank.

The CKE is one of the power-saving means. When CKE is off, the internal DDR clock is disabled and the DDR power is reduced. The power-saving differs according to the selected mode and the DDR type used. For more information, refer to the IDD table in the DDR specification.

The processor supports four different types of power-down modes in package C0 state. The different power-down modes can be enabled through configuring PM PDWN configuration register. The type of CKE power-down can be configured through PDWN_mode (bits 15:12) and the idle timer can be configured through PDWN_idle_counter (bits 11:0). The different power-down modes supported are:

- **No power-down** (CKE disable)
- **Active power-down (APD):** This mode is entered if there are open pages when de-asserting CKE. In this mode the open pages are retained. Power-saving in this mode is the lowest. Power consumption of DDR is defined by IDD3P. Exiting this mode is fined by tXP – small number of cycles. For this mode, DRAM DLL should be on.
- **PPD/DLL-off:** In this mode the data-in DLLs on DDR are off. Power-saving in this mode is the best among all power modes. Power consumption is defined by IDD2P. Exiting this mode is defined by tXP, but also tXPDLL (10–20 according to DDR type) cycles until first data transfer is allowed. For this mode, DRAM DLL should be off.
- **Precharged power-down (PPD):** This mode is entered if all banks in DDR are precharged when de-asserting CKE. Power-saving in this mode is intermediate – better than APD, but less than DLL-off. Power consumption is defined by IDD2P. Exiting this mode is defined by tXP. The difference from APD mode is that when waking-up, all page-buffers are empty.) The LPDDR does not have a DLL. As a result, the power savings are as good as PPD/DDL-off but will have lower exit latency and higher performance.

The CKE is determined per rank, whenever it is inactive. Each rank has an idle counter. The idle-counter starts counting as soon as the rank has no accesses, and if it expires, the rank may enter power-down while no new transactions to the rank arrives to queues. The idle-counter begins counting at the last incoming transaction arrival.

It is important to understand that since the power-down decision is per rank, the IMC can find many opportunities to power down ranks, even while running memory intensive applications; the savings are significant (may be few Watts, according to DDR specification). This is significant when each channel is populated with more ranks.

Selection of power modes should be according to power-performance or thermal trade-off of a given system:

- When trying to achieve maximum performance and power or thermal consideration is not an issue: use no power-down.

- In a system which tries to minimize power-consumption, try using the deepest power-down mode possible – PPD/DLL-off with a low idle timer value.
- In high-performance systems with dense packaging (that is, tricky thermal design) the power-down mode should be considered in order to reduce the heating and avoid DDR throttling caused by the heating.

The default value that BIOS configures in PM PDWN configuration register is 6080 – that is, PPD/DLL-off mode with idle timer of 0x80, or 128 DCLKs. This is a balanced setting with deep power-down mode and moderate idle timer value.

The idle timer expiration count defines the # of DCLKs that a rank is idle that causes entry to the selected power mode. As this timer is set to a shorter time the IMC will have more opportunities to put the DDR in power-down. There is no BIOS hook to set this register. Customers choosing to change the value of this register can do it by changing it in the BIOS. For experiments, this register can be modified in real time if BIOS does not lock the IMC registers.

4.3.2.1 Initialization Role of CKE

During power-up, CKE is the only input to the SDRAM that has its level recognized (other than the reset pin) once power is applied. It should be driven LOW by the DDR controller to make sure the SDRAM components float DQ and DQS during power-up. CKE signals remain LOW (while any reset is active) until the BIOS writes to a configuration register. Using this method, CKE is ensured to remain inactive for much longer than the specified 200 micro-seconds after power and clocks to SDRAM devices are stable.

4.3.2.2 Conditional Self-Refresh

During S0 idle state, system memory may be conditionally placed into self-refresh state when the processor is in package C3 or deeper power state. Refer to Section 4.6.1.1 for more details on conditional self-refresh with Intel® HD Graphics enabled.

When entering the S3 – Suspend-to-RAM (STR) state or S0 conditional self-refresh, the processor IA core flushes pending cycles and then enters SDRAM ranks that are not used by the processor graphics into self-refresh. The CKE signals remain LOW so the SDRAM devices perform self-refresh.

The target behavior is to enter self-refresh for package C3 or deeper power states as long as there are no memory requests to service.

Table 4-7. Targeted Memory State Conditions

State	Memory State with Processor Graphics	Memory State with External Graphics
C0, C1, C1E	Dynamic memory rank power-down based on idle conditions.	Dynamic memory rank power-down based on idle conditions.
C3, C6, C7 or deeper	If the processor graphics engine is idle and there are no pending display requests, then enter self-refresh. Otherwise use dynamic memory rank power-down based on idle conditions.	If there are no memory requests, then enter self-refresh. Otherwise use dynamic memory rank power-down based on idle conditions.
S3	Self-Refresh Mode	Self-Refresh Mode
S4	Memory power-down (contents lost)	Memory power-down (contents lost)

4.3.2.3 Dynamic Power-Down

Dynamic power-down of memory is employed during normal operation. Based on idle conditions, a given memory rank may be powered down. The IMC implements aggressive CKE control to dynamically put the DRAM devices in a power-down state. The processor IA core controller can be configured to put the devices in active power-down (CKE de-assertion with open pages) or precharge power-down (CKE de-assertion with all pages closed). Precharge power-down provides greater power savings but has a bigger performance impact, since all pages will first be closed before putting the devices in power-down mode.

If dynamic power-down is enabled, all ranks are powered up before doing a refresh cycle and all ranks are powered down at the end of refresh.

4.3.2.4 DRAM I/O Power Management

Unused signals should be disabled to save power and reduce electromagnetic interference. This includes all signals associated with an unused memory channel. Clocks, CKE, ODT and CS signals are controlled per DIMM rank and will be powered down for unused ranks.

The I/O buffer for an unused signal should be tri-stated (output driver disabled), the input receiver (differential sense-amp) should be disabled, and any DLL circuitry related ONLY to unused signals should be disabled. The input path should be gated to prevent spurious results due to noise on the unused signals (typically handled automatically when input receiver is disabled).

4.3.3 DDR Electrical Power Gating (EPG)

The DDR I/O of the processor supports Electrical Power Gating (DDR-EPG) while the processor is at C3 or deeper power state.

In C3 or deeper power state, the processor internally gates VDDQ for the majority of the logic to reduce idle power while keeping all critical DDR pins such as CKE and VREF in the appropriate state.

In C7 or deeper power state, the processor internally gates V_{CCIO} for all non-critical state to reduce idle power.

In S3 or C-state transitions, the DDR does not go through training mode and will restore the previous training information.

4.3.4 Power Training

BIOS MRC performing Power Training steps to reduce DDR I/O power while keeping reasonable operational margins, still ensuring platform operation. The algorithms attempt to weaken ODT, driver strength and the related buffers parameters both on the MC and the DRAM side and find the best possible trade-off between the total I/O power and the operational margins using advanced mathematical models.

4.4 PCI Express* Power Management

- Active power management support using L1 Substates (L1.2, L1.2).
- All inputs and outputs disabled in L2/L3 Ready state.

Notes:

1. Processor PEG-PCIe interface does not support Hot-Plug.
2. Hot Plug like^ is only supported at Processor PEG-PCIe using Thunderbolt™ Device.
3. ^Turning Thunderbolt™ Power On and Off electrically RTD3 Like.
4. The PCI Express* and DMI interfaces are present only in 2-Chip platform processors.
5. An increase in power consumption may be observed when PCI Express* ASPM capabilities are disabled.

4.4.0.1 Package C-States with PCIe* Link States Dependencies

PEG/DMI	L-State	Description	Package C-State
DMI	L1	Higher latency, lower power "standby" state	PC6-PC10
PEG	L1, L1.1, L1.2 L2, Disabled, NDA (no device attached)	L1- Higher latency, lower power "standby" state L2 – Auxiliary-powered Link, deep-energy-saving state. Disabled - The intent of the Disabled state is to allow a configured Link to be disabled until directed or Electrical Idle is exited (that is, due to a hot removal and insertion) after entering Disabled. NDA- no physical device is attached on PEG port	PC6-PC7
PEG	L2, Disabled, NDA (no device attached)	L2 – Auxiliary-powered Link, deep-energy-saving state. Disabled - The intent of the Disabled state is to allow a configured Link to be disabled until directed or Electrical Idle is exited (that is, due to a hot removal and insertion) after entering Disabled. NDA- no physical device is attached on PEG port	PC8-PC10

4.5 Direct Media Interface (DMI) Power Management

Active power management support using L1 state.

Note: The PCI Express* and DMI interfaces are present only in 2-Chip platform processors.

4.6 Processor Graphics Power Management

4.6.1 Memory Power Savings Technologies

4.6.1.1 Intel® Rapid Memory Power Management (Intel® RMPM)

Intel® Rapid Memory Power Management (Intel® RMPM) conditionally places memory into self-refresh when the processor is in package C3 or deeper power state to allow the system to remain in the deeper power states longer for memory not reserved for graphics memory. Intel® RMPM functionality depends on graphics/display state (relevant only when processor graphics is being used), as well as memory traffic patterns generated by other connected I/O devices.

4.6.1.2 Intel® Smart 2D Display Technology (Intel® S2DDT)

Intel® S2DDT reduces display refresh memory traffic by reducing memory reads required for display refresh. Power consumption is reduced by less accesses to the IMC. Intel® S2DDT is only enabled in single pipe mode. Intel® S2DDT is most effective with:

- Display images well suited to compression, such as text windows, slide shows, and so on. Poor examples are 3D games.
- Static screens such as screens with significant portions of the background showing 2D applications, processor benchmarks, and so on, or conditions when the processor is idle. Poor examples are full-screen 3D games and benchmarks that flip the display image at or near display refresh rates.

4.6.2 Display Power Savings Technologies

4.6.2.1 Display Power States

State	State Definition	LTR	Usage
DC0	All wells on	0	Boot or no driver installed
DC1	All wells on and LTR allowed	>0[1]	Driver loaded and using external display
DC2	External wells disabled	>0[1]	Only internal display
DC3	IO off	Max	In PSR[2] or all pipes disabled
DC4	Transition	Max	Display idle, transitioning to DC5
DC5	Internal wells disabled	Max	Display idle
DC6	Power lost, except sustain	Max	Display and SA idle

Notes:
1. LTR depends on resolution, FBC compression, number of planes, bits per pixel, and so on.
2. Refer PSR tables for the flavors of PSR that can reach DC3 and those that can proceed to DC5 and DC6.

4.7 GT / IA CORE Power States Relation

GT	IA core C-state**	Additional condition	Pkg-C Target
RC0	X		PC0
X	X	Deepest* state is limited to PC0	PC0
X	CC0|CC1		PC0
RC6	>=CC6	Deepest* state is limited to PC2	PC2
RC6	>=CC6	Deepest state is limited to PC3	PC3
RC6	=CC6	Deepest state allows PC6 or deeper	PC6
RC6	>=CC6	Deepest state limits to PC6	PC6
RC6	=CC7	Deepest state allows PC7 or deeper	PC7
RC6	>=CC7	Deepest state limits to PC7	PC7
RC6	=CC7	Deepest state allows PC7 or deeper **And** LLC is fully flushed	PC7S
RC6	>=CC7	Deepest state limits to PC7 **And** LLC is fully flushed	PC7S
RC6	=CC8	Deepest state allows PC8 or deeper	PC8
RC6	>=CC8	Deepest state limits to PC8	PC8
RC6	=CC9	Deepest state allows PC9 or deeper	PC9

GT	IA core C-state**	Additional condition	Pkg-C Target
RC6	>=CC9	Deepest state limits to PC9	PC9
RC6	=CC10	Deepest state allows PC9 or deeper **and** At least for one **enabled** core PCU.core_status<wbr>[C6_SRAM_flushed]=0	PC9
RC6	=CC10	Deepest state allows PC10 or deeper **and** For all **enabled** cores PCU.core_status<wbr>[C6_SRAM_flushed]=1	PC10

4.7.0.1 Intel® (Seamless and Static) Display Refresh Rate Switching (DRRS) with eDP* Port

Intel® DRRS provides a mechanism where the monitor is placed in a slower refresh rate (the rate at which the display is updated). The system is smart enough to know that the user is not displaying either 3D or media like a movie where specific refresh rates are required. The technology is very useful in an environment such as a plane where the user is in battery mode doing E-mail, or other standard office applications. It is also useful where user may be viewing web pages or social media sites while battery mode.

4.7.0.2 Intel® Automatic Display Brightness

Intel® Automatic Display Brightness feature dynamically adjusts the backlight brightness based upon the current ambient light environment. This feature requires an additional sensor to be on the panel front. The sensor receives the changing ambient light conditions and sends the interrupts to the Intel® Graphics driver. As per the change in Lux, (current ambient light illuminance), the new backlight setting can be adjusted through BLC. The converse applies for a brightly lit environment. Intel® Automatic Display Brightness increases the backlight setting.

4.7.0.3 Smooth Brightness

The Smooth Brightness feature is the ability to make fine grained changes to the screen brightness. All Windows* 10 system that support brightness control are required to support Smooth Brightness control and it should be supporting 101 levels of brightness control. Apart from the Graphics driver changes, there may be few System BIOS changes required to make this feature functional.

4.7.0.4 Intel® Display Power Saving Technology (Intel® DPST) 6.0

The Intel® DPST technique achieves backlight power savings while maintaining a good visual experience. This is accomplished by adaptively enhancing the displayed image while decreasing the backlight brightness simultaneously. The goal of this technique is to provide equivalent end-user-perceived image quality at a decreased backlight power level.

1. The original (input) image produced by the operating system or application is analyzed by the Intel® DPST subsystem. An interrupt to Intel® DPST software is generated whenever a meaningful change in the image attributes is detected. (A meaningful change is when the Intel® DPST software algorithm determines that enough brightness, contrast, or color change has occurred to the displaying images that the image enhancement and backlight control needs to be altered).
2. Intel® DPST subsystem applies an image-specific enhancement to increase image contrast, brightness, and other attributes.
3. A corresponding decrease to the backlight brightness is applied simultaneously to produce an image with similar user-perceived quality (such as brightness) as the original image.

Intel® DPST 6.0 has improved the software algorithms and has minor hardware changes to better handle backlight phase-in and ensures the documented and validated method to interrupt hardware phase-in.

4.7.0.5 Low-Power Single Pipe (LPSP)

Low-power single pipe is a power conservation feature that helps save power by keeping the inactive pipes powered OFF. This feature is enabled only in a single display configuration without any scaling functionalities. This feature is supported from 4th Generation Intel® Core™ processor family onwards. LPSP is achieved by keeping a single pipe enabled during eDP* only with minimal display pipeline support. This feature is panel independent and works with any eDP* panel (port A) in single display mode.

4.7.1 Processor Graphics Core Power Savings Technologies

4.7.1.1 Intel® Graphics Dynamic Frequency

Intel® Turbo Boost Technology 2.0 is the ability of the processor IA cores and graphics (Graphics Dynamic Frequency) cores to opportunistically increase frequency and/or voltage above the guaranteed processor and graphics frequency for the given part. Intel® Graphics Dynamic Frequency is a performance feature that makes use of unused package power and thermals to increase application performance. The increase in frequency is determined by how much power and thermal budget is available in the package, and the application demand for additional processor or graphics performance. The processor IA core control is maintained by an embedded controller. The graphics driver dynamically adjusts between P-States to maintain optimal performance, power, and thermals. The graphics driver will always place the graphics engine in its lowest possible P-State. Intel® Graphics Dynamic Frequency requires BIOS support. Additional power and thermal budget should be available.

4.7.1.2 Intel® Graphics Render Standby Technology (Intel® GRST)

The final power savings technology from Intel happens while the system is asleep. This is another technology where the voltage is adjusted down. For RC6 the voltage is adjusted very low, or very close to zero, what may reduced power by over 1000.

4.7.1.3 Dynamic FPS (DFPS)

Dynamic FPS (DFPS) or dynamic frame-rate control is a runtime feature for improving power-efficiency for 3D workloads. Its purpose is to limit the frame-rate of full screen 3D applications without compromising on user experience. By limiting the frame rate, the load on the graphics engine is reduced, giving an opportunity to run the Processor Graphics at lower speeds, resulting in power savings. This feature works in both AC/DC modes.

5 Thermal Management

5.1 Processor Thermal Management

The thermal solution provides both component-level and system-level thermal management. To allow optimal operation and long-term reliability of Intel processor-based systems, the system/processor thermal solution should be designed so that the processor:

- **Bare Die Parts**: Remains below the maximum junction temperature (Tj_{MAX}) specification at the maximum thermal design power (TDP).
- **Lidded Parts**: Remains below the maximum case temperature (Tcmax) specification at the maximum thermal design power.
- Conforms to system constraints, such as system acoustics, system skin-temperatures, and exhaust-temperature requirements.

Caution: Thermal specifications given in this chapter are on the component and package level and apply specifically to the processor. Operating the processor outside the specified limits may result in permanent damage to the processor and potentially other components in the system.

5.1.1 Thermal Considerations

The processor TDP is the maximum sustained power that should be used for design of the processor thermal solution. TDP is a power dissipation and component temperature operating condition limit, specified in this document, that is validated during manufacturing for the base configuration when executing a near worst case commercially available workload without AVX as specified by Intel for the SKU segment. TDP may be exceeded for short periods of time or if running a very high power workload.

To allow the optimal operation and long-term reliability of Intel processor-based systems, the processor must remain within the minimum and maximum component temperature specifications. For lidded parts, the appropriate case temperature (TCASE) specifications is defined by the applicable thermal profile. For bare die parts the component temperature specification is the applicable Tj_max.

Thermal solutions not designed to provide this level of thermal capability may affect the long-term reliability of the processor and system.

The processor integrates multiple processing IA cores, graphics cores in a single package.This may result in power distribution differences across the package and should be considered when designing the thermal solution.

Intel® Turbo Boost Technology 2.0 allows processor IA cores to run faster than the base frequency. It is invoked opportunistically and automatically as long as the processor is conforming to its temperature, voltage, power delivery and current control limits. When Intel® Turbo Boost Technology 2.0 is enabled:

- Applications are expected to run closer to TDP more often as the processor will attempt to maximize performance by taking advantage of estimated available energy budget in the processor package.

- The processor may exceed the TDP for short durations to utilize any available thermal capacitance within the thermal solution. The duration and time of such operation can be limited by platform runtime configurable registers within the processor.
- Graphics peak frequency operation is based on the assumption of only one of the graphics domains (GT) being active. This definition is similar to the IA core Turbo concept, where peak turbo frequency can be achieved when only one IA core is active. Depending on the workload being applied and the distribution across the graphics domains the user may not observe peak graphics frequency for a given workload or benchmark.
- Thermal solutions and platform cooling that are designed to less than thermal design guidance may experience thermal and performance issues.

Note: Intel® Turbo Boost Technology 2.0 availability may vary between the different SKUs.

5.1.2 Intel® Turbo Boost Technology 2.0 Power Monitoring

When operating in turbo mode, the processor monitors its own power and adjusts the processor and graphics frequencies to maintain the average power within limits over a thermally significant time period. The processor estimates the package power for all components on package. In the event that a workload causes the temperature to exceed program temperature limits, the processor will protect itself using the Adaptive Thermal Monitor.

5.1.3 Intel® Turbo Boost Technology 2.0 Power Control

Illustration of Intel® Turbo Boost Technology 2.0 power control is shown in the following sections and figures. Multiple controls operate simultaneously allowing customization for multiple system thermal and power limitations. These controls allow for turbo optimizations within system constraints and are accessible using MSR, MMIO, or PECI interfaces.

5.1.3.1 Package Power Control

The package power control settings of PL1, PL2, PL3, PL4 and Tau allow the designer to configure Intel® Turbo Boost Technology 2.0 to match the platform power delivery and package thermal solution limitations.

- **Power Limit 1 (PL1)**: A threshold for average power that will not exceed - recommend to set to equal TDP power. PL1 should not be set higher than thermal solution cooling limits.
- **Power Limit 2 (PL2)**: A threshold that if exceeded, the PL2 rapid power limiting algorithms will attempt to limit the spike above PL2.
- **Power Limit 3 (PL3)**: A threshold that if exceeded, the PL3 rapid power limiting algorithms will attempt to limit the duty cycle of spikes above PL3 by reactively limiting frequency. This is an optional setting
- **Power Limit 4 (PL4)**: A limit that will not be exceeded, the PL4 power limiting algorithms will preemptively limit frequency to prevent spikes above PL4.
- **Turbo Time Parameter (Tau)**: An averaging constant used for PL1 exponential weighted moving average (EWMA) power calculation.

Note:

1. Implementation of Intel® Turbo Boost Technology 2.0 only requires configuring PL1, PL1 Tau, PL2 and PL4.dt
2. PL3 is disabled by default.
3. PL4 hardware default is disabled, refer to additional documentation for recommended settings.

Figure 5-1. Package Power Control

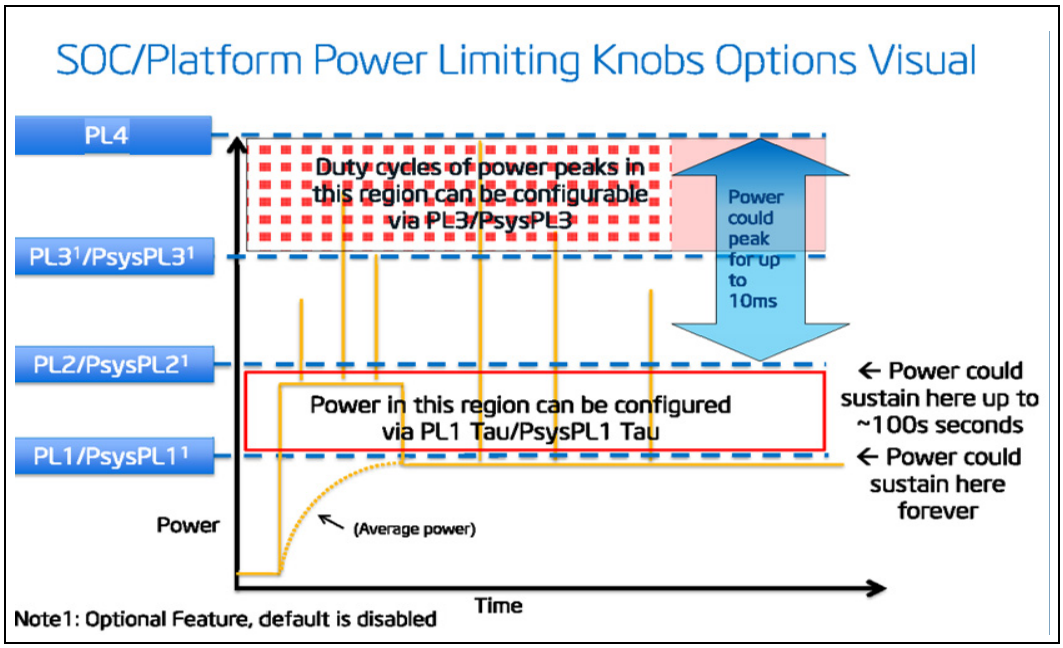

5.1.3.2 Platform Power Control

The processor supports Psys (Platform Power) to enhance processor power management. The Psys signal needs to be sourced from a compatible charger circuit and routed to the IMVP8 (voltage regulator). This signal will provide the total thermally relevant platform power consumption (processor and rest of platform) via SVID to the processor.

When the Psys signal is properly implemented, the system designer can utilize the package power control settings of PsysPL1/Tau, PsysPL2 and PsysPL3 for additional manageability to match the platform power delivery and platform thermal solution limitations for Intel® Turbo Boost Technology 2.0. The operation of the PsysPL1/tau, PsysPL2 and PsysPL3 is analogous to the processor power limits described in Section 5.1.3.1, “Package Power Control”.

- **Platform Power Limit 1 (PsysPL1)**: A threshold for average platform power that will not be exceeded - recommend to set to equal platform thermal capability.
- **Platform Power Limit 2 (PsysPL2)**: A threshold that if exceeded, the PsysPL2 rapid power limiting algorithms will attempt to limit the spikes above PsysPL2.
- **Platform Power Limit 3 (PsysPL3)**: A threshold that if exceeded, the PsysPL3 rapid power limiting algorithms will attempt to limit the duty cycle of spikes above PsysPL3 by reactively limiting frequency.

- **PsysPL1 Tau**: An averaging constant used for PsysPL1 exponential weighted moving average (EWMA) power calculation.
- The Psys signal and associated power limits / Tau are optional for the system designer and disabled by default.
- The Psys data will not include power consumption for charging.

5.1.3.3 Turbo Time Parameter (Tau)

Turbo Time Parameter (Tau) is a mathematical parameter (units of seconds) that controls the Intel® Turbo Boost Technology 2.0 algorithm. During a maximum power turbo event, the processor could sustain PL2 for a duration longer than the Turbo Time Parameter. If the power value and/or Turbo Time Parameter is changed during runtime, it may take some time based on the new Turbo Time Parameter level for the algorithm to settle at the new control limits. The time varies depending on the magnitude of the change, power limits, and other factors. There is an individual Turbo Time Parameter associated with Package Power Control and Platform Power Control.

5.1.4 Thermal Management Features

Occasionally the processor may operate in conditions that are near to its maximum operating temperature. This can be due to internal overheating or overheating within the platform. In order to protect the processor and the platform from thermal failure, several thermal management features exist to reduce package power consumption and thereby temperature in order to remain within normal operating limits. Furthermore, the processor supports several methods to reduce memory power.

5.1.4.1 Adaptive Thermal Monitor

The purpose of the Adaptive Thermal Monitor is to reduce processor IA core power consumption and temperature until it operates below its maximum operating temperature. Processor IA core power reduction is achieved by:

- Adjusting the operating frequency (using the processor IA core ratio multiplier) and voltage.
- Modulating (starting and stopping) the internal processor IA core clocks (duty cycle).

The Adaptive Thermal Monitor can be activated when the package temperature, monitored by any digital thermal sensor (DTS), meets its maximum operating temperature. The maximum operating temperature implies maximum junction temperature Tj_{MAX}.

Reaching the maximum operating temperature activates the Thermal Control Circuit (TCC). When activated the TCC causes both the processor IA core and graphics core to reduce frequency and voltage adaptively. The Adaptive Thermal Monitor will remain active as long as the package temperature remains at its specified limit. Therefore, the Adaptive Thermal Monitor will continue to reduce the package frequency and voltage until the TCC is de-activated.

Tj_{MAX} is factory calibrated and is not user configurable. The default value is software visible in the TEMPERATURE_TARGET (0x1A2) MSR, bits [23:16].

The Adaptive Thermal Monitor does not require any additional hardware, software drivers, or interrupt handling routines. It is not intended as a mechanism to maintain processor thermal control to PL1 = TDP. The system design should provide a thermal solution that can maintain normal operation when PL1 = TDP within the intended usage range.

Adaptive Thermal Monitor protection is always enabled.

5.1.4.1.1 TCC Activation Offset

TCC Activation Offset can be set as an offset from the maximum allowed component temperature to lower the onset of TCC and Adaptive Thermal Monitor. In addition, the processor has added an optional time window (Tau) to manage processor performance at the TCC Activation offset value via an EWMA (Exponential Weighted Moving Average) of temperature. For more information on TCC Activation offset.

TCC Activation Offset with Tau=0

An offset (degrees Celsius) can be written to the TEMPERATURE_TARGET (0x1A2) MSR, bits [29:24], the offset value will be subtracted from the value found in bits [23:16]. When the time window (Tau) is set to zero, there will be no averaging, the offset, will be subtracted from the Tj_{MAX} value and used as a new max temperature set point for Adaptive Thermal Monitoring. This will have the same behavior as in prior products to have TCC activation and Adaptive Thermal Monitor to occur at this lower target silicon temperature.

If enabled, the offset should be set lower than any other passive protection such as ACPI _PSV trip points.

TCC Activation Offset with Tau

To manage the processor with the EWMA (Exponential Weighted Moving Average) of temperature, an offset (degrees Celsius) is written to the TEMPERATURE_TARGET (0x1A2) MSR, bits [29:24], and the time window (Tau) is written to the TEMPERATURE_TARGET (0x1A2) MSR [6:0]. The Offset value will be subtracted from the value found in bits [23:16] and be the temperature.

The processor will manage to this average temperature by adjusting the frequency of the various domains. The instantaneous Tj can briefly exceed the average temperature. The magnitude and duration of the overshoot is managed by the time window value (Tau).

This averaged temperature thermal management mechanism is in addition, and not instead of Tj_{MAX} thermal management. That is, whether the TCC activation offset is 0 or not, TCC Activation will occur at Tj_{MAX}.

5.1.4.1.2 Frequency / Voltage Control

Upon Adaptive Thermal Monitor activation, the processor attempts to dynamically reduce processor temperature by lowering the frequency and voltage operating point. The operating points are automatically calculated by the processor IA core itself and do not require the BIOS to program them as with previous generations of Intel processors. The processor IA core will scale the operating points such that:

- The voltage will be optimized according to the temperature, the processor IA core bus ratio and number of processor IA cores in deep C-states.

- The processor IA core power and temperature are reduced while minimizing performance degradation.

Once the temperature has dropped below the trigger temperature, the operating frequency and voltage will transition back to the normal system operating point.

Once a target frequency/bus ratio is resolved, the processor IA core will transition to the new target automatically.

- On an upward operating point transition, the voltage transition precedes the frequency transition.
- On a downward transition, the frequency transition precedes the voltage transition.
- The processor continues to execute instructions. However, the processor will halt instruction execution for frequency transitions.

If a processor load-based Enhanced Intel SpeedStep® Technology/P-state transition (through MSR write) is initiated while the Adaptive Thermal Monitor is active, there are two possible outcomes:

- If the P-state target frequency is higher than the processor IA core optimized target frequency, the P-state transition will be deferred until the thermal event has been completed.
- If the P-state target frequency is lower than the processor IA core optimized target frequency, the processor will transition to the P-state operating point.

5.1.4.1.3 Clock Modulation

If the frequency/voltage changes are unable to end an Adaptive Thermal Monitor event, the Adaptive Thermal Monitor will utilize clock modulation. Clock modulation is done by alternately turning the clocks off and on at a duty cycle (ratio between clock "on" time and total time) specific to the processor. The duty cycle is factory configured to 25% on and 75% off and cannot be modified. The period of the duty cycle is configured to 32 microseconds when the Adaptive Thermal Monitor is active. Cycle times are independent of processor frequency. A small amount of hysteresis has been included to prevent excessive clock modulation when the processor temperature is near its maximum operating temperature. Once the temperature has dropped below the maximum operating temperature, and the hysteresis timer has expired, the Adaptive Thermal Monitor goes inactive and clock modulation ceases. Clock modulation is automatically engaged as part of the Adaptive Thermal Monitor activation when the frequency/voltage targets are at their minimum settings. Processor performance will be decreased when clock modulation is active. Snooping and interrupt processing are performed in the normal manner while the Adaptive Thermal Monitor is active.

Clock modulation will not be activated by the Package average temperature control mechanism.

5.1.4.2 Digital Thermal Sensor

Each processor has multiple on-die Digital Thermal Sensor (DTS) that detects the processor IA, GT and other areas of interest instantaneous temperature.

Temperature values from the DTS can be retrieved through:

- A software interface using processor Model Specific Register (MSR).
- A processor hardware interface as described in Platform Environmental Control Interface (PECI)

When temperature is retrieved by the processor MSR, it is the instantaneous temperature of the given DTS. When temperature is retrieved using PECI, it is the average of the highest DTS temperature in the package over a 256 ms time window. Intel recommends using the PECI reported temperature for platform thermal control that benefits from averaging, such as fan speed control. The average DTS temperature may not be a good indicator of package Adaptive Thermal Monitor activation or rapid increases in temperature that triggers the Out of Specification status bit within the PACKAGE_THERM_STATUS MSR 1B1h and IA32_THERM_STATUS MSR 19Ch.

Code execution is halted in C1 or deeper C- states. Package temperature can still be monitored through PECI in lower C-states.

Unlike traditional thermal devices, the DTS outputs a temperature relative to the maximum supported operating temperature of the processor (Tj_{MAX}), regardless of TCC activation offset. It is the responsibility of software to convert the relative temperature to an absolute temperature. The absolute reference temperature is readable in the TEMPERATURE_TARGET MSR 1A2h. The temperature returned by the DTS is an implied negative integer indicating the relative offset from Tj_{MAX}. The DTS does not report temperatures greater than Tj_{MAX}. The DTS-relative temperature readout directly impacts the Adaptive Thermal Monitor trigger point. When a package DTS indicates that it has reached the TCC activation (a reading of 0x0, except when the TCC activation offset is changed), the TCC will activate and indicate an Adaptive Thermal Monitor event. A TCC activation will lower both processor IA core and graphics core frequency, voltage, or both. Changes to the temperature can be detected using two programmable thresholds located in the processor thermal MSRs. These thresholds have the capability of generating interrupts using the processor IA core's local APIC.

5.1.4.2.1 Digital Thermal Sensor Accuracy (Taccuracy)

The error associated with DTS measurements will not exceed ±5 °C within the entire operating range.

5.1.4.2.2 Fan Speed Control with Digital Thermal Sensor

Digital Thermal Sensor based fan speed control (T_{FAN}) is a recommended feature to achieve optimal thermal performance. At the T_{FAN} temperature, Intel recommends full cooling capability before the DTS reading reaches Tj_{MAX}.

5.1.4.3 PROCHOT# Signal

The PROCHOT# (processor hot) signal is asserted by the processor when the TCC is active. Only a single PROCHOT# pin exists at a package level. When any DTS temperature reaches the TCC activation temperature, the PROCHOT# signal will be asserted. PROCHOT# assertion policies are independent of Adaptive Thermal Monitor enabling.

The PROCHOT# signal can be configured to the following modes:

- **Input only**: PROCHOT is driven by an external device.
- **Output only**: PROCHOT is driven by processor.
- **Bi-Directional**: Both Processor and external device can drive PROCHOT signal

5.1.4.4 PROCHOT Input Only

The PROCHOT# signal should be set to input only by default. In this state, the processor will only monitor PROCHOT# assertions and respond by setting the maximum frequency to 10 kHz.

The following two features are enabled when PROCHOT is set to Input only:

- **Fast PROCHOT**: Respond to PROCHOT# within 10 uS of PROCHOT# pin assertion, reducing the processor frequency by 50%.
- **PROCHOT Demotion Algorithm**: Designed to improve system performance during multiple PROCHOT assertions (refer to Section 5.1.4.7, "PROCHOT Demotion Algorithm").

5.1.4.5 PROCHOT Output Only

Legacy state, PROCHOT is driven by the processor to external device.

5.1.4.6 Bi-Directional PROCHOT#

By default, the PROCHOT# signal is set to input only. When configured as an input or bi-directional signal, PROCHOT# can be used for thermally protecting other platform components should they overheat as well. When PROCHOT# is driven by an external device:

- The package will immediately transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores. This is contrary to the internally-generated Adaptive Thermal Monitor response.
- Clock modulation is not activated.

The processor package will remain at the lowest supported P-state until the system deasserts PROCHOT#. The processor can be configured to generate an interrupt upon assertion and de-assertion of the PROCHOT# signal.

When PROCHOT# is configured as a bi-directional signal and PROCHOT# is asserted by the processor, it is impossible for the processor to detect a system assertion of PROCHOT#. The system assertion will have to wait until the processor de-asserts PROCHOT# before PROCHOT# action can occur due to the system assertion. While the processor is hot and asserting PROCHOT#, the power is reduced but the reduction rate is slower than the system PROCHOT# response of < 100 us. The processor thermal control is staged in smaller increments over many milliseconds. This may cause several milliseconds of delay to a system assertion of PROCHOT# while the output function is asserted.

5.1.4.7 PROCHOT Demotion Algorithm

PROCHOT demotion algorithm designed to improve system performance following multiple EC PROCHOT consecutive assertions. During each PROCHOT assertion processor will immediately transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores (LFM). When detecting several PROCHOT consecutive assertions the processor will reduce the max frequency in order to reduce the PROCHOT assertions events. The processor will keep reducing the frequency until no consecutive assertions detected. The processor will raise the frequency if no consecutive PROCHOT assertion events will occur. PROCHOT demotion algorithm enabled only when the PROCHOT is configured as input.

5.1.4.8 PROCHOT Demotion Signal Description

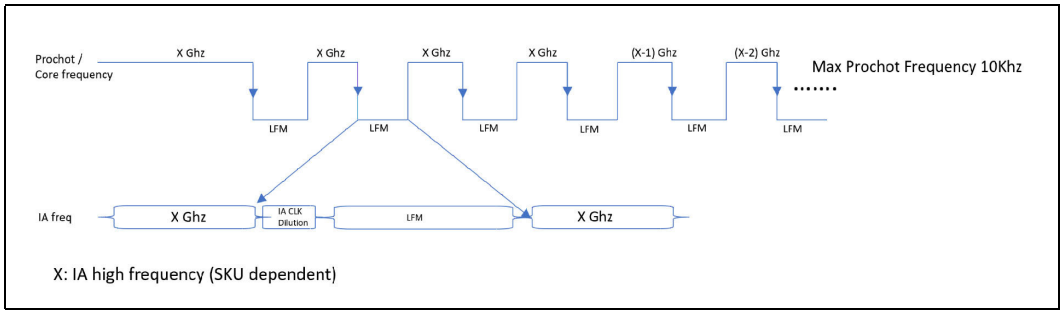

5.1.4.9 Voltage Regulator Protection Using PROCHOT#

PROCHOT# may be used for thermal protection of voltage regulators (VR). System designers can create a circuit to monitor the VR temperature and assert PROCHOT# and, if enabled, activate the TCC when the temperature limit of the VR is reached. When PROCHOT# is configured as a bi-directional or input only signal, if the system assertion of PROCHOT# is recognized by the processor, it will result in an immediate transition to the lowest P-State (Pn) supported by the processor IA cores and graphics cores. Systems should still provide proper cooling for the VR and rely on bi-directional PROCHOT# only as a backup in case of system cooling failure. Overall, the system thermal design should allow the power delivery circuitry to operate within its temperature specification even while the processor is operating at its TDP.

5.1.4.10 Thermal Solution Design and PROCHOT# Behavior

With a properly designed and characterized thermal solution, it is anticipated that PROCHOT# will only be asserted for very short periods of time when running the most power intensive applications. The processor performance impact due to these brief periods of TCC activation is expected to be so minor that it would be immeasurable. However, an under-designed thermal solution that is not able to prevent excessive assertion of PROCHOT# in the anticipated ambient environment may:

- Cause a noticeable performance loss.
- Result in prolonged operation at or above the specified maximum junction temperature and affect the long-term reliability of the processor.
- May be incapable of cooling the processor even when the TCC is active continuously (in extreme situations).

5.1.4.11 Low-Power States and PROCHOT# Behavior

Depending on package power levels during package C-states, outbound PROCHOT# may de-assert while the processor is idle as power is removed from the signal. Upon wake up, if the processor is still hot, the PROCHOT# will re-assert. Although, typically package idle state residency should resolve any thermal issues. The PECI interface is fully operational during all C-states and it is expected that the platform continues to manage processor IA core and package thermals even during idle states by regularly polling for thermal data over PECI.

5.1.4.12 THERMTRIP# Signal

Regardless of enabling the automatic or on-demand modes, in the event of a catastrophic cooling failure, the package will automatically shut down when the silicon has reached an elevated temperature that risks physical damage to the product. At this point, the THERMTRIP# signal will go active.

5.1.4.13 Critical Temperature Detection

Critical Temperature detection is performed by monitoring the package temperature. This feature is intended for graceful shutdown before the THERMTRIP# is activated. However, the processor execution is not guaranteed between critical temperature and THERMTRIP#. If the Adaptive Thermal Monitor is triggered and the temperature remains high, a critical temperature status and sticky bit are latched in the PACKAGE_THERM_STATUS MSR 1B1h and the condition also generates a thermal interrupt, if enabled.

5.1.4.14 On-Demand Mode

The processor provides an auxiliary mechanism that allows system software to force the processor to reduce its power consumption using clock modulation. This mechanism is referred to as "On-Demand" mode and is distinct from Adaptive Thermal Monitor and bi-directional PROCHOT#. The processor platforms should not rely on software usage of this mechanism to limit the processor temperature. On-Demand Mode can be accomplished using processor MSR or chipset I/O emulation. On-Demand Mode may be used in conjunction with the Adaptive Thermal Monitor. However, if the system software tries to enable On-Demand mode at the same time the TCC is engaged, the factory configured duty cycle of the TCC will override the duty cycle selected by the On-Demand mode. If the I/O based and MSR-based On-Demand modes are in conflict, the duty cycle selected by the I/O emulation-based On-Demand mode will take precedence over the MSR-based On-Demand Mode.

5.1.4.15 MSR Based On-Demand Mode

If Bit 4 of the IA32_CLOCK_MODULATION MSR is set to 1, the processor will immediately reduce its power consumption using modulation of the internal processor IA core clock, independent of the processor temperature. The duty cycle of the clock modulation is programmable using bits [3:1] of the same IA32_CLOCK_MODULATION MSR. In this mode, the duty cycle can be programmed in either 12.5% or 6.25% increments (discoverable using CPUID). Thermal throttling using this method will modulate each processor IA core's clock independently.

5.1.4.16 I/O Emulation-Based On-Demand Mode

I/O emulation-based clock modulation provides legacy support for operating system software that initiates clock modulation through I/O writes to ACPI defined processor clock control registers on the chipset (PROC_CNT). Thermal throttling using this method will modulate all processor IA cores simultaneously.

5.1.5 Intel® Memory Thermal Management

The processor provides thermal protection for system memory by throttling memory traffic when using either DIMM modules or a memory down implementation. Two levels of throttling are supported by the processor, either a warm threshold or hot threshold that is customizable through memory mapped I/O registers. Throttling based on the

warm threshold should be an intermediate level of throttling. Throttling based on the hot threshold should be the most severe. The amount of throttling is dynamically controlled by the processor.

The on Die Thermal Sensor (ODTS) uses a physical thermal sensor on DRAM dies. ODTS is available for DDR4. It is used to set refresh rate according to DRAM temperature. The memory controller reads DDR4 MR3 and configures the DDR refresh rate accordingly.

Memory temperature may be acquired through an on-board thermal sensor (TS-on-Board), retrieved by an embedded controller and reported to the processor through the PECI 3.1 interface. This methodology is known as PECI injected temperature. This is a method of Closed Loop Thermal Management (CLTM).

5.2 All-Processor Line Thermal and Power Specifications

Note	Definition
1	The TDP values is the average power dissipation in junction temperature operating condition limit, for the SKU Segment and Configuration, for which the processor is validated during manufacturing when executing an associated Intel-specified high-complexity workload at the processor IA core frequency corresponding to the configuration and SKU.
2	TDP workload may consist of a combination of processor IA core intensive and graphics core intensive applications.
3	Can be modified at runtime by MSR writes, with MMIO and with PECI commands.
4	'Turbo Time Parameter' is a mathematical parameter (units of seconds) that controls the processor turbo algorithm using a moving average of energy usage. Do not set the Turbo Time Parameter to a value less than 0.1 seconds. refer to Section 5.1.3.2, "Platform Power Control" for further information.
5	Shown limit is a time averaged power, based upon the Turbo Time Parameter. Absolute product power may exceed the set limits for short durations or under virus or uncharacterized workloads.
6	Processor will be controlled to specified power limit as described in Section 5.1.2, "Intel® Turbo Boost Technology 2.0 Power Monitoring". If the power value and/or 'Turbo Time Parameter' is changed during runtime, it may take a short period of time (approximately 3 to 5 times the 'Turbo Time Parameter') for the algorithm to settle at the new control limits.
7	This is a hardware default setting and not a behavioral characteristic of the part. The reference BIOS code may override the hardware default power limit values to optimize performance
8	For controllable turbo workloads, the PL2 limit may be exceeded for up to 10 ms.
9	n/a
10	n/a
11	Default power limits can be found in the PKG_PWR_SKU MSR (614h).
12	The processor die and OPCM die do not reach maximum sustained power simultaneously since the sum of the two dies estimated power budget is controlled to be equal to or less than the package TDP (PL1) limit. For additional information, refer to the appropriate Mobile TMDG for more information (Refer Related Documents).N/A
13	cTDP down power is based on GT2 equivalent graphics configuration. cTDP down does not decrease the number of active Processor Graphics EUs, but relies on Power Budget Management (PL1) to achieve the specified power level.
14	May vary based on SKU, Not all SKUs have cTDP up/down, each SKU has a different base Frequency and cTDP frequency respective.
15	Sustained residencies at high voltages and temperatures may temporarily limit turbo frequency.
16	PL2 - SoC opportunistic higher Average Power with limited duration controlled by Tau_PL1 setting, The larger the Tau, the longer the PL2 duration.
17	PL1 Tau max recommendation value is the default value in the BIOS/BKC and this value is the one that has been tested.

5.3 Thermal and Power Specifications

Table 5-1. TDP Specifications

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Configuration	Processor IA Core Frequency	Graphics Core Frequency	Thermal Design Power (TDP) [w]	Notes (from section 5.2)
Intel® Xeon® E-2300 processor Family	8 Core / pGFX 95W	Base	3.2GHz	1.2GHz	95W	1,11,12,15
		LFM	0.8GHz	0.35GHz		
	6 Core / pGFX 95W	Base	3.5GHz	1.2GHz	95W	
		LFM	0.8GHz	0.35GHz		
	8 Core / pGFX 80W	Base	2.8GHz	1.2GHz	80W	
		LFM	0.8GHz	0.35GHz		
	6 Core / pGFX 80W	Base	3.2GHz	1.2GHz	80W	
		LFM	0.8GHz	0.35GHz		
	4 Core / pGFX 80W	Base	3.7GHz	1.2GHz	80W	
		LFM	0.8GHz	0.35GHz		
	8 Core / no pGFX 65W	Base	2.6GHz	N/A	65W	
		LFM	0.8GHz	N/A		
	6 Core / no pGFX 65W	Base	2.9GHz	N/A	65W	
		LFM	0.8GHz	N/A		
	4 Core / no pGFX 65W	Base	3.4GHz	N/A	65W	
		LFM	0.8GHz	N/A		
	4 Core / pGFX 65W	Base	3.1GHz	1.2GHz	65W	
		LFM	0.8GHz	0.35GHz		
	4 Core / no pGFX 65W	Base	2.8GHz	N/A	65W	
		LFM	0.8GHz	N/A		

Note: The ~ sign stands for approximation.

Table 5-2. Package Turbo Specifications

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Parameter	Minimum	Hardware Default	MSR MAX value	Units
Intel® Xeon® E-2300 processor Family	8/6 Core 95W	Power Limit 1 Time (PL1 Tau)	0.01	56	448	S
		Power Limit 1 (PL1)	N/A	95	N/A	W
		Power Limit 2 (PL2)	N/A	251 (8 Cores) 210 (6 Cores)	N/A	W
	8/6/4 Core 80W	Power Limit 1 Time (PL1 Tau)	0.01	28	448	S
		Power Limit 1 (PL1)	N/A	80	N/A	W
		Power Limit 2 (PL2)	N/A	235 (8 Cores) 205 (6 Cores) 170 (4 Cores)	N/A	W
	8/6/4 Core 65W	Power Limit 1 Time (PL1 Tau)	0.01	28	448	S
		Power Limit 1 (PL1)	N/A	65	N/A	W
		Power Limit 2 (PL2)	N/A	224 (8 Cores) 200 (6 Cores) 160 (4 Cores)	N/A	W

Notes:

- No Specifications for Min/Max PL1/PL2 values.
- Hardware default of PL1 Tau=1 s, By including the benefits available from power and thermal management features the recommended is to use PL1 Tau=56 s for DT 125W and 28 s for other SKU.

Table 5-3. T_{CASE} Specification

Processor IA Cores, Graphics Configuration and TDP	TDP	Tla, Psi_ca @Arizola TTV	Min Tcase (°C)	Max Tcase (°C)
8 Core / pGFX	95W	40 °C, 0.253 °C/W	0	62.1
6 Core / pGFX			0	62.7
8 Core / pGFX	80W	40 °C, 0.378 °C/W	0	68.5
6 Core / pGFX			0	69
4 Core / pGFX			0	70.3
8 Core / no pGFX	65W	40 °C, 0.483 °C/W	0	69.9
6 Core / no pGFX			0	70.3
4 Core / pGFX			0	71.4
4 Core / no pGFX			0	71.4

Table 5-4. $T_{CONTROL}$ Offset Configuration

Processor IA Cores, Graphics Configuration and TDP	TDP	TEMP_TARGET (Tcontrol) [°C]
8 Core / pGFX	95W	10
6 Core / pGFX		10
8 Core / pGFX	80W	10
6 Core / pGFX		10
4 Core / pGFX		10

Table 5-4. $T_{CONTROL}$ Offset Configuration

Processor IA Cores, Graphics Configuration and TDP	TDP	TEMP_TARGET (Tcontrol) [°C]
8 Core / no pGFX	65W	10
6 Core / no pGFX		10
4 Core / pGFX		10
4 Core / no pGFX		10

Notes:
1. 1. Digital Thermal Sensor (DTS) based fan speed control is recommended to achieve optimal thermal performance.
2. Intel recommends full cooling capability at approximately the DTS value of -1, to minimize TCC activation risk.
3. For example, if TCONTROL = 10 ºC, Fan acceleration operation will start at 90 ºC (100 ºC - 10 ºC).

5.3.1 Thermal Metrology

The maximum TTV case temperatures ($T_{CASE-MAX}$) can be derived from the data in the appropriate TTV thermal profile earlier in this chapter. The TTV T_{CASE} is measured at the geometric top center of the TTV integrated heat spreader (IHS). Next figure illustrates the location where T_{CASE} temperature measurements should be made.

Figure 5-2. Thermal Test Vehicle (TTV) Case Temperature (T_{CASE}) Measurement Location

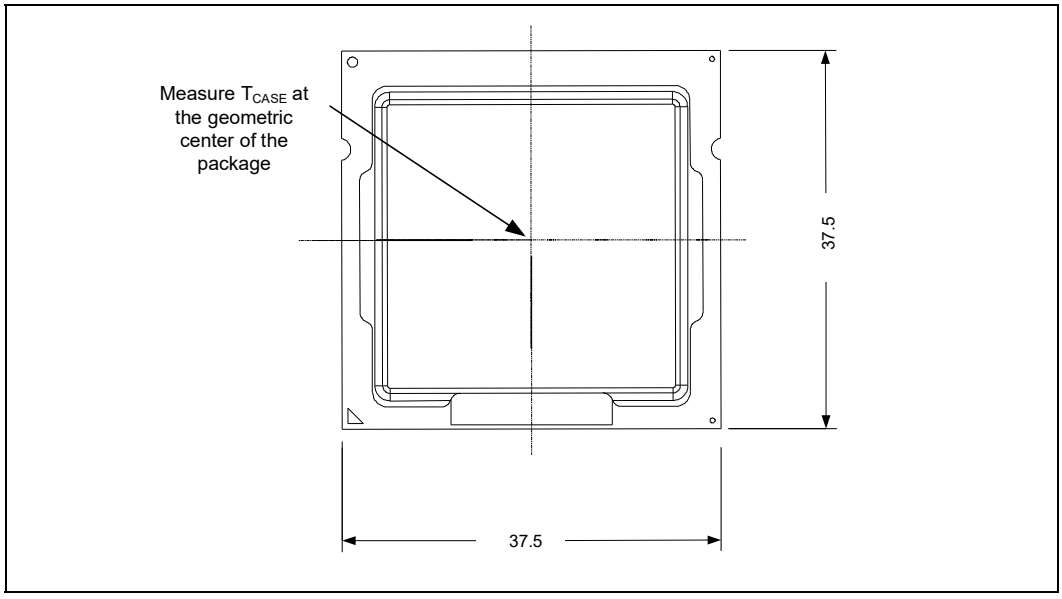

5.3.2 Fan Speed Control Scheme with Digital Thermal Sensor (DTS) 2.0

To simplify processor thermal specification compliance, the processor calculates the DTS Thermal Profile from $T_{CONTROL}$ Offset, TCC Activation Temperature, TDP, and the Thermal Margin Slope provided in the following table.

Note: TCC Activation Offset is 0 for the processors.

Using the DTS Thermal Profile, the processor can calculate and report the Thermal Margin, where a value less than 0 indicates that the processor needs additional cooling, and a value greater than 0 indicates that the processor is sufficiently cooled.

Figure 5-3. Digital Thermal Sensor (DTS) 2.0 Definition Points

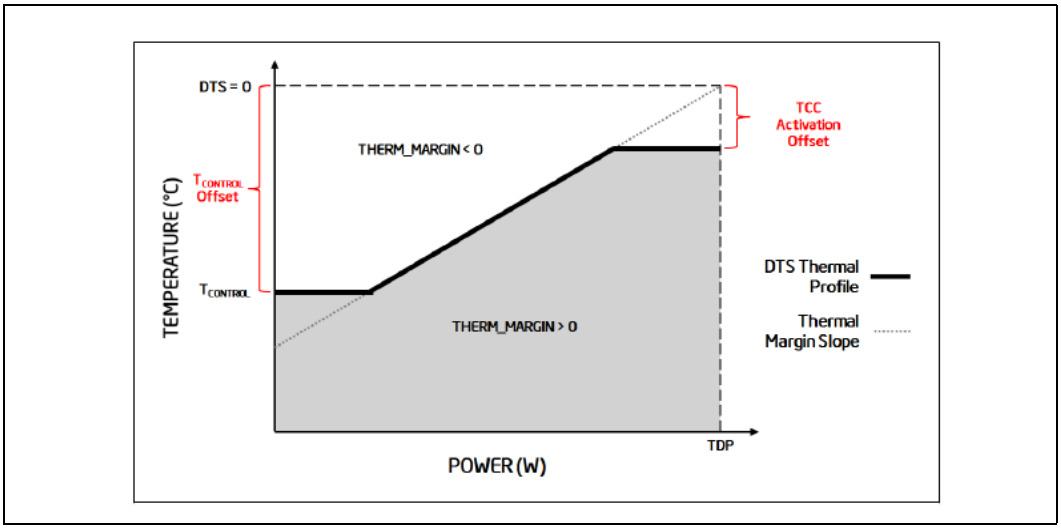

Table 5-5. T_{CASE} and DTS Thermal Profile

Processor IA Cores, Graphics Configuration and TDP	TDP [W]	TCC Activation [°C]	$T_{CONTROL}$	T_{CASE} Thermal Profile	T_{CASE_MAX} @ TDP	DTS Thermal Profile
8 Core / pGFX	95	100	10	0.21 x TDP + 42.2	62.1	0.608 x TDP + 42.2
6 Core / pGFX	95	100	10	0.22 x TDP + 41.8	62.7	0.613 x TDP + 41.8
8 Core / pGFX	80	100	10	0.32 x TDP + 42.9	68.5	0.714 x TDP + 42.9
6 Core / pGFX	80	100	10	0.33 x TDP + 42.6	69	0.717 x TDP + 42.6
4 Core / pGFX	80	100	10	0.34 x TDP + 43.1	70.3	0.711 x TDP + 43.1
8 Core / no pGFX	65	100	10	0.41 x TDP + 43.3	69.9	0.872 x TDP + 43.3
6 Core / no pGFX	65	100	10	0.42 x TDP + 43.0	70.3	0.877 x TDP + 43.0
4 Core / pGFX	65	100	10	0.43 x TDP + 43.4	71.4	0.87 x TDP + 43.4
4 Core / no pGFX	65	100	10	0.43 x TDP + 43.4	71.4	0.87 x TDP + 43.4

6 Signal Description

This chapter describes the processor signals. They are arranged in functional groups according to their associated interface or category. The notations in the following table are used to describe the signal type.

The signal description also includes the type of buffer used for the particular signal (Refer the following table).

Table 6-1. Signal Tables Terminology

	Short name	Functionality
Direction	I/O	Input or Output
	O	Output only
	I	Input only
	N/A	Not applicable (Mostly for power rails and RSVD signals)
Buffer Type	DDR4	DDR4 memory (1.2v tolerant)
	LPDDR4	LPDDR4 memory (1.1v tolerant)
	LPDDR4x	LPDDR4x memory (1.1v TX, 0.6v RX tolerant)
	LPDDR5	LPDDR5 memory (1.05v)
	Analog	Analog reference or output. May be used as a threshold voltage or for buffer compensation.
	PCI Express	PCI Express
	DMI	DMI
	GTL	Gunning Transceiver Logic signaling technology
	CMOS	CMOS (1.05v tolerant)
	AUDIO	AUDIO
	N/A	Not Applicable
	Async CMOS[1]	Async CMOS
	DP/HDMI	DP/HDMI
	OD	Open Drain
	PECI Async	PECI Async
	Diff Amp Clock	Diff Amp Clock Input Buffer
	Power	Power
	Ground	Ground
Link Type	SE	Single Ended
	DIFF	Differential pair

Notes:
1. Qualifier for a buffer type
2. CMOS - Complementary Metal Oxide Substrate
3. GTL - Gunning Transceiver Signaling Technology Logic
4. DP - Display Port
5. PECI - Platform Environment Control Interface
6. Async - Signal is not related to any clock in the system
7. DDR - Double Data Rate Synchronous Dynamic Random Access Memory
8. LPDDR - Low Power DDR
9. On some case I/O may be split into: I=GTL, O=OD

6.1 System Memory Interface

Table 6-2. DDR4 Memory Interface (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_ECC[7:0] DDR1_ECC[7:0]	**ECC Data Buses:** Data buses for ECC Check Byte.	I/O	DDR4	SE	ECC UDIMM/SODIM Modules with Intel® Xeon® E-2300 processor product family.
DDR0_DQ[7:0][[7:0]] DDR1_DQ[7:0][[7:0]]	**Data Buses:** Data signals interface to the SDRAM data buses. **Example:** DDR0_DQ2[5]	I/O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_DQSP[8:0] DDR0_DQSN[8:0] DDR1_DQSP[8:0] DDR1_DQSN[8:0]	**Data Strobes:** Differential data strobe pairs. The data is captured at the crossing point of DQS during read and write transactions.	I/O	DDR4	Diff	Intel® Xeon® E-2300 processor product family
DDR0_CLK_N[3:0] DDR0_CLK_P[3:0] DDR1_CLK_N[3:0] DDR1_CLK_P[3:0]	**SDRAM Differential Clock:** Differential clocks signal pairs, pair per rank. The crossing of the positive edge of DDR0_CKP/DDR1_CKP and the negative edge of their complement DDR0_CKN / DDR1_CKN are used to sample the command and control signals on the SDRAM.	O	DDR4	Diff	Intel® Xeon® E-2300 processor product family
DDR0_CKE[3:0] DDR1_CKE[3:0]	**Clock Enable:** (1 per rank). These signals are used to: • Initialize the SDRAMs during power-up. • Power-down SDRAM ranks. • Place all SDRAM ranks into and out of self-refresh during STR (Suspend to RAM).	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_CS#[3:0] DDR1_CS#[3:0]	**Chip Select:** (1 per rank). These signals are used to select particular SDRAM components during the active state. There is one Chip Select for each SDRAM rank.	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_ODT[3:0] DDR1_ODT[3:0]	**On Die Termination:** (1 per rank). Active SDRAM Termination Control.	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_MA[16:0] DDR1_MA[16:0]	**Address:** These signals are used to provide the multiplexed row and column address to the SDRAM. • A[16:14] use also as command signals, Refer ACT# signal description. • A10 is sampled during Read/Write commands to determine whether Autoprecharge should be performed to the accessed bank after the Read/Write operation. HIGH: Autoprecharge; LOW: no Autoprecharge). • A10 is sampled during a Precharge command to determine whether the Precharge applies to one bank (A10 LOW) or all banks (A10 HIGH). If only one bank is to be precharged, the bank is selected by bank addresses. • A12 is sampled during Read and Write commands to determine if burst chop (on-the-fly) will be performed. HIGH, no burst chop; LOW: burst chopped).	O	DDR4	SE	Intel® Xeon® E-2300 processor product family

Table 6-2. DDR4 Memory Interface (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_ACT# DDR1_ACT#	**Activation Command:** ACT# HIGH along with CS# determines that the signals addresses below have command functionality. A16 use as RAS# signal A15 use as CAS# signal A14 use as WE# signal	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_BA[1:0] DDR1_BA[1:0]	**Bank Address:** BA[1:0] define to which bank an Active, Read, Write or Precharge command is being applied. Bank address also determines which mode register is to be accessed during a MRS cycle.	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_ALERT# DDR1_ALERT#	**Alert:** This signal is used at command training only. It is getting the Command and Address Parity error flag during training. CRC feature is not supported.	I	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_PAR DDR1_PAR	**Command and Address Parity:** These signals are used for parity check.	O	DDR4	SE	Intel® Xeon® E-2300 processor product family
DDR0_VREF_CA DDR1_VREF_CA	**Memory Reference Voltage for DQ:**	O	Analog	SE	Intel® Xeon® E-2300 processor product family
DDR_VREF_CA[3:0]					

Table 6-3. System Memory Reference and Compensation Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR_VTT_CNTL	**System Memory Power Gate Control:** When signal is high – platform memory VTT regulator is enable, output high. When signal is low - Disables the platform memory VTT regulator in C8 and deeper and S3.	O	DDR4	SE	Intel® Xeon® E-2300 processor product family

6.2 PCI Express* Graphics (PEG) Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PCIE_RXP[15:0] PCIE_RXN[15:0]	PCI Express* Receive Differential Pairs.	I	PCI Express	Diff	Intel® Xeon® E-2300 processor product family
PCIE_TXP[15:0] PCIE_TXN[15:0]	PCI Express* Transmit Differential Pairs.	O	PCI Express	Diff	

6.3 Direct Media Interface (DMI) Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DMI_RXP[7:0] DMI_RXN[7:0]	**DMI Input from PCH:** Direct Media Interface receive differential pairs.	I	DMI	Diff	Intel® Xeon® E-2300 processor product family
DMI_TXP[7:0] DMI_TXN[7:0]	**DMI Output from PCH:** Direct Media Interface transmit differential pairs.	O	DMI	Diff	

6.4 Reset and Miscellaneous Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CFG[15:0]	**Configuration Signals:** The CFG signals have a default value of '1' if not terminated on the board. Intel recommends placing test points on the board for CFG pins. • **CFG[0]**: Stall reset sequence after PCU PLL lock until de-asserted: — 1 = (Default) Normal Operation; No stall. — 0 = Stall. • **CFG[1]**: Reserved configuration lane. • **CFG[2]**: PCI Express* Static x16 Lane Numbering Reversal. — 1 = Normal operation — 0 = Lane numbers reversed. • **CFG[3]:** Reserved configuration lane. • **CFG[4]:** eDP enable: — 1 = Disabled. — 0 = Enabled. • **CFG[6:5]:** PCI Express* Bifurcation — 00 = 1 x8, 2 x4 PCI Express* — 01 = reserved — 10 = 2 x8 PCI Express* — 11 = 1 x16 PCI Express* • **CFG[15:7]:** Reserved configuration lanes.	I	GTL	SE	Intel® Xeon® E-2300 processor product family
CFG[16P,17P] CFG[16N,17N]	• Reserved configuration lanes. • Reserved configuration lanes.	I	GTL	SW	
SYS_RESET#	Platform Reset pin driven by the PCH.	I	CMOS	SE	Intel® Xeon® E-2300 processor product family
IST_TRIG	**Impedance Spectrum Tool Trigger:** trigger point to support debug of possible power issues.	O	GTL	SE	Intel® Xeon® E-2300 processor product family
PROC_TRIGIN	Debug pin.	I	CMOS	SE	Intel® Xeon® E-2300 processor product family
PROC_TRIGOUT	Debug pin.	O	CMOS	SE	Intel® Xeon® E-2300 processor product family
PROC_AUDIO_SDI	**Processor Audio Serial Data Input:** This signal is an input to the processor from the PCH.	I	AUDIO	SE	Intel® Xeon® E-2300 processor product family
PROC_AUDIO_SDO	**Processor Audio Serial Data Output:** This Line signal is an output from the processor to the PCH.	O	AUDIO	SE	Intel® Xeon® E-2300 processor product family
PROC_AUDIO_CLK	**Processor Audio Clock**	I	AUDIO	SE	Intel® Xeon® E-2300 processor product family

6.5 Display Interface Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIB/C/D_TXP[3:0] DDIB/C/D_TXN[3:0]	**Digital Display Interface Transmit:** Differential Pairs	O	DP/HDMI	Diff	Intel® Xeon® E-2300 processor product family

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIB/C/D_AUXP_TXP DDIB/C/D_AUXP_TXN	**Digital Display Interface Display Port Auxiliary:** Half-duplex, bidirectional channel consist of one differential pair for each channel.	O	DP/HDMI	Diff	Intel® Xeon® E-2300 processor product family

6.6 Digital Display Audio Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
AUDOUT	**Serial Data Output for display audio interface**	O	AUDIO	SE	Intel® Xeon® E-2300 processor product family
AUDIN	**Serial Data Input for display audio interface**	I	AUDIO	SE	
AUDCLK	**Serial Data Clock**	I	AUDIO	SE	

6.7 Processor Clocking Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
BCLKP BCLKN	100 MHz Differential bus clock input to the processor.	I	Diff Amp Clock	Diff	Intel® Xeon® E-2300 processor product family
CLK24P CLK24N	24 MHz Differential bus clock input to the processor.	I	Diff Amp Clock	Diff	
PCI_BCLKP PCI_BCLKN	100 MHz Clock for PCI Express* logic	I	Diff Amp Clock	Diff	

6.8 Testability Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
BPM#[3:0]	**Breakpoint and Performance Monitor Signals:** Outputs from the processor that indicate the status of breakpoints and programmable counters used for monitoring processor performance.	I/O	GTL	SE	Intel® Xeon® E-2300 processor product family
PROC_PRDY#	**Probe Mode Ready:** PROC_PRDY# is a processor output used by debug tools to determine processor debug readiness.	O	OD	SE	Intel® Xeon® E-2300 processor product family
PROC_PREQ#	**Probe Mode Request:** PROC_PREQ# is used by debug tools to request debug operation of the processor.	I	GTL	SE	Intel® Xeon® E-2300 processor product family
PROC_TCLK	**Test Clock:** This signal provides the clock input for the processor Test Bus (also known as the Test Access Port). This signal should be driven low or Allowed to float during power on Reset.	I	GTL	SE	Intel® Xeon® E-2300 processor product family
PROC_TDI	**Test Data In:** This signal transfers serial test data into the processor. This signal provides the serial input needed for JTAG specification support.	I	GTL	SE	Intel® Xeon® E-2300 processor product family
PROC_TDO	**Test Data Out:** This signal transfers serial test data out of the processor. This signal provides the serial output needed for JTAG specification support.	O	OD	SE	Intel® Xeon® E-2300 processor product family
PROC_TMS	**Test Mode Select:** A JTAG specification support signal used by debug tools.	I	GTL	SE	Intel® Xeon® E-2300 processor product family

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PROC_TRST#	**Test Reset:** Resets the Test Access Port (TAP) logic. This signal should be driven low during power on Reset.	I	GTL	SE	Intel® Xeon® E-2300 processor product family

6.9 Error and Thermal Protection Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CATERR#	**Catastrophic Error:** This signal indicates that the system has experienced a catastrophic error and cannot continue to operate. The processor will set this signal for non-recoverable machine check errors or other unrecoverable internal errors. CATERR# is used for signaling the following types of errors: Legacy MCERRs, CATERR# is asserted for 16 BCLKs. Legacy IERRs, CATERR# remains asserted until warm or cold reset.	O	OD	SE	Intel® Xeon® E-2300 processor product family
PECI	**Platform Environment Control Interface:** A serial sideband interface to the processor. It is used primarily for thermal, power, and error management. PECI over eSPI is supported on S-Processor, for more details, refer to Intel 400 Series Chipset Family.	I/O	PECI Async	SE	Intel® Xeon® E-2300 processor product family
PROCHOT#	**Processor Hot:** PROCHOT# goes active when the processor temperature monitoring sensor(s) detects that the processor has reached its maximum safe operating temperature. This indicates that the processor Thermal Control Circuit (TCC) has been activated, if enabled. This signal can also be driven to the processor to activate the TCC.	I/O	I:GTL O:OD	SE	Intel® Xeon® E-2300 processor product family
THRMTRIP#	**Thermal Trip:** The processor protects itself from catastrophic overheating by use of an internal thermal sensor. This sensor is set well above the normal operating temperature to ensure that there are no false trips. The processor will stop all executions when the junction temperature exceeds approximately 130 °C. This is signaled to the system by the THERMTRIP# pin.	O	OD	SE	Intel® Xeon® E-2300 processor product family

6.10 Power Sequencing Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PROCPWRGD	**Processor Power Good:** The processor requires this input signal to be a clean indication that the V_{CC} and V_{DDQ} power supplies are stable and within specifications. This requirement applies regardless of the S-state of the processor. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until they come within specification. The signal should then transition monotonically to a high state.	I	CMOS	SE	Intel® Xeon® E-2300 processor product family

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCCST_PWRGD	**VCCST Power Good:** The processor requires this input signal to be a clean indication that the VCCST and VDDQ power supplies are stable and within specifications. This signal should have a valid level during both S0 and S3 power states. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until they come within specification. The signal should then transition monotonically to a high state.	I	CMOS	SE	Intel® Xeon® E-2300 processor product family
SKTOCC#	**Processor Detect / Socket Occupied:** Pulled down directly (0 Ohms) on the processor package to the ground. There is no connection to the processor silicon for this signal. System board designers may use this signal to determine if the processor is present.	N/A	N/A	SE	Intel® Xeon® E-2300 processor product family
VIDSOUT VIDSCK VIDALERT#	**VIDSOUT**, **VIDSCK, VIDALERT#:** These signals comprise a three-signal serial synchronous interface used to transfer power management information between the processor and the voltage regulator controllers.	I/O O I	I:GTL/O:OD OD CMOS	SE	Intel® Xeon® E-2300 processor product family
PM_SYNC	**Power Management Sync:** A sideband signal to communicate power management status from the PCH to the processor. PCH report EXTTS#/EVENT# status to the processor.	I	CMOS	SE	Intel® Xeon® E-2300 processor product family
PM_DOWN	**Power Management Down:** Sideband to PCH. Indicates processor wake up event EXTTS# on PCH. The processor combines the pin status into the OLTM/CLTM.	O	CMOS	SE	Intel® Xeon® E-2300 processor product family

6.11 Processor Power Rails

Table 6-4. Processor Power Rails Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
Vcccore	Processor IA cores power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{GT}	Processor Graphics power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
V_{DDQ}	System Memory power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{SA}	Processor System Agent power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{IO-0}	Processor PCIe I/O power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{IO-1-2}	Processor PCIe / DDR I/O power rail	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{ST}	Sustain voltage for processor standby modes	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{STG}	Gated sustain voltage for processor standby modes	I	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{PLL_OC}	Processor PLLs power rails	I	Power	N/A	Intel® Xeon® E-2300 processor product family

Table 6-4. Processor Power Rails Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
Vcc_SENSE	Isolated, low impedance voltage sense pins. They can be used to sense or measure voltage near the silicon.	N/A	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{GT}_SENSE		N/A	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{IO}_SENSE		N/A	Power	N/A	Intel® Xeon® E-2300 processor product family
Vcc_{SA}_SENSE		N/A	Power	N/A	Intel® Xeon® E-2300 processor product family

Table 6-5. Processor Ground Rails Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
Vss_SENSE	Isolated, low impedance Ground sense pins. They can be used for the reference ground near the silicon.	N/A	Ground	N/A	Intel® Xeon® E-2300 processor product family
Vss_{GT}_SENSE		N/A	Ground	N/A	Intel® Xeon® E-2300 processor product family
Vss_{SA}_Vss_{IO}_SENSE		N/A	Ground	N/A	Intel® Xeon® E-2300 processor product family
VCCIN_AUX_VSSSENSE		N/A	Ground	N/A	Intel® Xeon® E-2300 processor product family
Note:	VssSA_VssIO_SENSE is designated to have the GND reference for sense both VCCSA and VCCIO power rails.				

6.12 Ground, Reserved and Non-Critical to Function (NCTF) Signals

The following are the general types of reserved (RSVD) signals and connection guidelines:

- **RSVD**: these signals should not be connected
- **RSVD_TP**: these signals should be routed to a test point
- **RSVD_NCTF**: these signals are non-critical to function and may be left un-connected

Arbitrary connection of these signals to VCC, VDDQ, VSS, or to any other signal (including each other) may result in component malfunction or incompatibility with future processors. Refer below table. For reliable operation, always connect unused inputs or bi-directional signals to an appropriate signal level. Unused active high inputs should be connected through a resistor to ground (V_{SS}). Unused outputs may be left unconnected however, this may interfere with some Test Access Port (TAP) functions, complicate debug probing and prevent boundary scan testing. A resistor should be used when tying bi-directional signals to power or ground. When tying any signal to power or ground, the resistor can also be used for system testability.

Table 6-6. GND, RSVD, EDGECAP and NCTF Signals (Sheet 1 of 2)

Signal Name	Description
Vss	Processor ground node
RSVD	**Reserved:** All signals that are RSVD should not be connected on the board.
RSVD_NCTF	**Reserved Non-critical To Function:** RSVD_NCTF should not be connected on the board.

Table 6-6. GND, RSVD, EDGECAP and NCTF Signals (Sheet 2 of 2)

Signal Name	Description
RSVD_TP	**Test Point:** Intel recommends to route each RSVD_TP to an accessible test point. Intel may required these test point for platform specific debug. Leaving these test point inaccessible could delay debug bu Intel.
VDDQ_EDGE[1] VDDQ_EDGE[2]	Internal power pin, this pin should be connected to a decoupling capacitor and ground.
VCCIO_EdgeCap	Internal power pin, this pin should be connected to a decoupling capacitor and ground.

6.13 Processor Internal Pull-Up / Pull-Down Terminations

Signal Name	Pull Up/Pull Down	Rail	Value
BPM[3:0]	Pull Down	Vcc_{IO}	16-60 ohms
PROC_PREQ#	Pull Up	Vcc_{ST}	3 kohms
PROC_TDI	Pull Up	Vcc_{ST}	3 kohms
PROC_TMS	Pull Up	Vcc_{ST}	3 kohms
CFG[19:0]	Pull Up	Vcc_{IO}	3 kohms

7 Electrical Specifications

7.1 Processor Power Rails

Power Rail	Description	Control	Availability
V_{CC}	Processor IA Cores Power Rail	SVID	Intel® Xeon® E-2300 processor product family
Vcc_{GT}	Processor Graphics Power Rails	SVID	Intel® Xeon® E-2300 processor product family
Vcc_{SA}	System Agent Power Rail	SVID	Intel® Xeon® E-2300 processor product family
Vcc_{IO_0}	IO Power Rail VCCIO_0	2 bit VID	Intel® Xeon® E-2300 processor product family
$Vcc_{IO\_1\_2}$	IO Power Rail	Fixed	Intel® Xeon® E-2300 processor product family
Vcc_{ST}	Sustain Power Rail	Fixed	Intel® Xeon® E-2300 processor product family
Vcc_{PLL_OC}[1]	Processor PLLs OC power Rail	Fixed	Intel® Xeon® E-2300 processor product family
V_{DDQ}	Integrated Memory Controller Power Rail	Fixed (Memory technology dependent)	Intel® Xeon® E-2300 processor product family

Notes:
1. Vcc_{PLL_OC} power rail should be sourced from the VDDQ VR for DDR4 only. The connection can be direct or through a load switch, depending desired power optimization. In case of direct connection (Vcc_{PLL_OC} is shorted to V_{DDQ}, no load switch), platform should ensure that Vcc_{ST} is ON (high) while Vcc_{PLL_OC} is ON (high).
2. Vcc_{STG} power rail should be sourced from the VR as V_{CCST}. The connection can be direct or through a load switch, depending desired power optimization.

7.1.1 Power and Ground Pins

All power pins should be connected to their respective processor power planes, while all VSS pins should be connected to the system ground plane. Use of multiple power and ground planes is recommended to reduce I*R drop.

7.1.2 V_{CC} Voltage Identification (VID)

Intel processors are individually calibrated in the factory to operate on a specific voltage/frequency and operating-condition curve specified for that individual processor. In normal operation, the processor autonomously issues voltage control requests according to this calibrated curve using the serial voltage-identifier (SVID) interface. Altering the voltage applied at the processor causing operation outside of this calibrated curve is considered out-of-specification operation.

The SVID bus consists of three open-drain signals: clock, data, and alert# to both set voltage-levels and gather telemetry data from the voltage regulators. Voltages are controlled per an 8-bit integer value, called a VID, that maps to an analog voltage level. An offset field also exists that allows altering the VID table. Alert can be used to inform the processor that a voltage-change request has been completed or to interrupt the processor with a fault notification.

7.2 DC Specifications

The processor DC specifications in this section are defined at the processor signal pins, unless noted otherwise.

- The DC specifications for the DDR4 signals are listed in the *Voltage and Current Specifications* section.
- The *Voltage and Current Specifications* section lists the DC specifications for the processor and are valid only while meeting specifications for junction temperature, clock frequency, and input voltages. Read all notes associated with each parameter.
- AC tolerances for all rails include voltage transients and voltage regulator voltage ripple up to 1MHz. Refer to additional guidance for each rail.

7.2.1 Processor Power Rails DC Specifications

7.2.1.1 Vcc DC Specifications

Table 7-1. Processor IA core (Vcc) Active and Idle Mode DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Min	Typ	Max			Unit	Note[1]
Operating Voltage	Voltage Range for Processor Operating Modes	Intel® Xeon® E-2300 processor product family	0	-	1.52 + Offset voltage = 1.72V			V	1,2,3, 7,16
Icc_{MAX}	Maximum Processor IA Core cc_{MAX}	Intel® Xeon® E-2300 processor product family	—	—	245			A	4, 6, 7,16
Icc_{TDC}	Thermal Design Current (TDC) for processor IA Cores Rail	—	—	—				A	9
TOB_{VCC}	Voltage Tolerance	PS0, PS1	—	—	±20			mV	3, 6, 8
		PS2, PS3	—	—	±20				
Ripple	Ripple Tolerance				$I_L \leq 0.5$	$0.5<I_L<$	$Icc_{TDC}<I_L<$	mV	3, 6, 8
		PS0	—	—	+30/ -10	+/-10	+/-15		
		PS1	—	—	+30/ -10	+/-15	+/-15		
		PS2	—	—	+30/ -10	+30/- 10	+30/-10		
		PS3	—	—	+30/ -10	+30/- 10	+30/-10		
DC_LL	Loadline slope within the VR regulation loop capability	Intel® Xeon® E-2300 processor product family	10, 13, 14	—	1.1			mΩ	10, 13, 14
AC_LL (S Processors)	AC Loadline	Intel® Xeon® E-2300 processor product family	—	—	Up to 20MHz: DC LL value			mΩ	10, 13, 14
T_OVS_TDP_MAX	Max Overshoot time **TDP/virus mode**	—	—	—	10/30			µs	
V_OVS TDP_MAX/ virus_MAX	Max Overshoot at **TDP/virus mode**	—	—	—	70/200			mV	

Table 7-1. Processor IA core (Vcc) Active and Idle Mode DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Min	Typ	Max	Unit	Note[1]

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Each processor is programmed with a maximum valid voltage identification value (VID) that is set at manufacturing and cannot be altered. Individual maximum VID values are calibrated during manufacturing such that two processors at the same frequency may have different settings within the VID range. Note that this differs from the VID employed by the processor during a power management event (Adaptive Thermal Monitor, Enhanced Intel® SpeedStep Technology, or low-power states).
3. The voltage specification requirements are measured across Vcc_SENSE and Vss_SENSE as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. Processor IA core VR to be designed to electrically support this current.
5. Processor IA core VR to be designed to thermally support this current indefinitely.
6. Long term reliability cannot be assured if tolerance, ripple, and core noise parameters are violated.
7. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.
8. PSx refers to the voltage regulator power state as set by the SVID protocol.
9. N/A
10. LL measured at sense points.
11. Typ column represents Icc_{MAX} for commercial application it is NOT a specification - it is a characterization of limited samples using limited set of benchmarks that can be exceeded.
12. Operating voltage range in steady state.
13. LL specification values should not be exceeded. If exceeded, power, performance and reliability penalty are expected.
14. Load Line (AC/DC) should be measured by the VRTT tool and programmed accordingly via the BIOS Load Line override setup options. AC/DC Load Line BIOS programming directly affects operating voltages (AC) and power measurements (DC). A superior board design with a shallower AC Load Line can improve on power, performance, and thermals compared to boards designed for POR impedance.

7.2.1.2 Vcc_{GT} DC Specifications

Table 7-2. Processor Graphics (Vcc_{GT}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note[1]
Operating voltage	Active voltage Range for Vcc_{GT}	Intel® Xeon® E-2300 processor product family	0	—	1.52	V	2, 3, 6, 8
Icc_{MAX_GT}	Max. Current for Processor Graphics Rail	Intel® Xeon® E-2300 processor product family	—	—	55	A	6
Icc_{TDC_GT}	Thermal Design Current (TDC) for Processor Graphics Rail	—	—	—	—	A	6
TOB_{GT}	Vcc_{GT} Tolerance	PS0, PS1	—	—	+/-20	mV	3, 4
		PS2, PS3	—	—	+/-20	mV	3, 4

Table 7-2. Processor Graphics (Vcc_{GT}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum			Unit	Note[1]
Ripple	Ripple Tolerance	—			$I_L \leq 0.5$	$0.5<I_L<Icc_{TDC}$	$Icc_{TDC}<I_L<Icc_{MAX}$	mV	3, 4
		PS0	—	—	+30/-10	+/-10	+/-15		
		PS1	—	—	+30/-10	+/-15	+/-15		
		PS2	—	—	+30/-10	+30/-10	+30/-10		
		PS3	—	—	+30/-10	+30/-10	+30/-10		
Icc_{MAX_GT}	Max. Current for Processor Graphics Rail	Intel® Xeon® E-2300 processor product family	—	—	55			A	6
DC_LL	AC Loadline	Intel® Xeon® E-2300 processor product family	—	—	4.0			mΩ	7, 9
AC_LL	AC Loadline	Intel® Xeon® E-2300 processor product family	—	—	AC LL same as DC LL			mΩ	7, 9
T_OVS_MAX	Max Overshoot time	—	—	—	10			µs	
V_OVS_MAX	Max Overshoot	—	—	—	70			mV	

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Each processor is programmed with a maximum valid voltage identification value (VID), which is set at manufacturing and cannot be altered. Individual maximum VID values are calibrated during manufacturing such that two processors at the same frequency may have different settings within the VID range. This differs from the VID employed by the processor during a power or thermal management event (Intel Adaptive Thermal Monitor, Enhanced Intel® SpeedStep Technology, or low-power states).
3. The voltage specification requirements are measured across Vcc_{GT_SENSE} and Vss_{GT_SENSE} as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. PSx refers to the voltage regulator power state as set by the SVID protocol.
5. Each processor is programmed with a maximum valid voltage identification value (VID), which is set at manufacturing and cannot be altered. Individual maximum VID values are calibrated during manufacturing such that two processors at the same frequency may have different settings within the VID range. This differs from the VID employed by the processor during a power or thermal management event (Intel Adaptive Thermal Monitor, Enhanced Intel® SpeedStep Technology, or low-power states).
6. LL measured at sense points.
7. Operating voltage range in steady state.
8. LL specification values should not be exceeded. If exceeded, power, performance and reliability penalty are expected.
9. Load Line (AC/DC) should be measured by the VRTT tool and programmed accordingly via the BIOS Load Line override setup options. AC/DC Load Line BIOS programming directly affects operating voltages (AC) and power measurements (DC). A superior board design with a shallower AC Load Line can improve on power, performance, and thermals compared to boards designed for POR impedance.

Table 7-3. Memory Controller (V_{DDQ}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note[1]
$V_{DDQ\ (DDR4)}$	Processor I/O supply voltage for DDR4	All	Typ-5%	1.2	Typ+5%	V	3, 4, 5
TOB_{VDDQ}	VDDQ Tolerance	All	VDDQmin<AC+DC<VDDQmax			V	3, 4

Table 7-3. Memory Controller (V_{DDQ}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note[1]
Icc_{MAX_VDDQ} (DDR4)	Max Current for V_{DDQ} Rail (DDR4)	All	—	—	3.5	A	2

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. The current supplied to the DIMM modules is not included in this specification.
3. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins.
4. No requirement on the breakdown of AC versus DC noise.
5. The voltage specification requirements are measured as near as possible to the processor. measurement needs to be performed with a 100MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.

7.2.1.3 Vcc_{SA} DC Specifications

Table 7-4. System Agent (Vcc_{SA}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum			Unit	Note[1,2]
Vcc_{SA}	Voltage for the System Agent	Intel® Xeon® E-2300 processor product family	0	—	1.52			V	3,5
TOB_{VCCSA}	Vcc_{SA} Tolerance	Intel® Xeon® E-2300 processor product family			±20			mV	3,8
I_{CCMAX_VCCSA}	Max Current for V_{CCSA} Rail	Intel® Xeon® E-2300 processor product family	—	—	22.1			A	1,2
Ripple	Ripple Tolerance	—			$I_L \leq 0.5$	0.5<I_L<Icc_{TDC}	I_{CCTDC}<I_L<Icc_{MAX}	mV	3, 4
		PS0	—	—	+30/-10	+/-10	+/-15		
		PS1	—	—	+30/-10	+/-15	+/-15		
		PS2	—	—	+30/-10	+30/-10	+30/-10		
		PS3	—	—	+30/-10	+30/-10	+30/-10		
T_OVS_MAX	Maximum Overshoot time	—	—	—	10			µs	
V_OVS_MAX	Maximum Overshoot	—	—	—	70			mV	

Table 7-4. System Agent (Vcc_{SA}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note[1,2]

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
3. The voltage specification requirements are measured across Vcc_{SA_SENSE} and Vss_{SA_SENSE} as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. PSx refers to the voltage regulator power state as set by the SVID protocol.
5. LL measured at sense points.
6. LL specification values should not be exceeded. If exceeded, power, performance and reliability penalty are expected.
7. Load Line (AC/DC) should be measured by the VRTT tool and programmed accordingly via the BIOS Load Line override setup options. AC/DC Load Line BIOS programming directly affects operating voltages (AC) and power measurements (DC). A superior board design with a shallower AC Load Line can improve on power, performance, and thermals compared to boards designed for POR impedance.
8. For Voltage less than 1V, TOB will be 50 mV.

7.2.1.4 Vcc_{IO} DC Specifications

Table 7-5. Processor I/O (Vcc_{IO}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note[1,2]
Vcc_{IO_0}	Voltage for the memory controller and shared cache	Intel® Xeon® E-2300 processor product family	Typ-5%	1.05	Typ+5%	V	3,6
$Vcc_{IO\_1\_2}$	Voltage for the memory and PCIe subsystem	Intel® Xeon® E-2300 processor product family	Typ-5%	1.0	Typ+5%	V	3
TOB_{VCCIO}	Vcc_{IO} Tolerance	All	+/-5 (AC + DC + Ripple) Up to 1 MHz			%	3,5
Icc_{MAX_VCCIO}	Maximum Current for V_{CCIO_0} Rail	Intel® Xeon® E-2300 processor product family	—	—	7.50	A	
	Maximum Current for $V_{CCIO\_1\_2}$ Rail				9.60		
T_OVS_MAX	Maximum Overshoot time	All	—	—	150	µS	4
V_OVS_MAX	Maximum Overshoot at TDP	All	—	—	30	mV	4

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
3. The voltage specification requirements are measured across Vcc_{IO_SENSE} and Vss_{IO_SENSE} as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. OS occurs during power on only, **not** during normal operation
5. For Voltage less than 1v, TOB will be +/-50mV (AC + DC + Ripple) up to 1 MHz.
6. The voltage is a two bit VID and this is controlled from VCCIO CPU strap.

7.2.2 Vcc_{ST} DC Specifications

Table 7-6. Vcc Sustain (Vcc_{ST}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes [1,2]
Vcc_{ST}	Processor Vcc Sustain supply voltage	Intel® Xeon® E-2300 processor product family	—	1.05	—	V	3
TOB_{ST}	Vcc_{ST} Tolerance	All	+/-5			%	3,4
Icc_{MAX_ST}	Max Current for Vcc_{ST}	Intel® Xeon® E-2300 processor product family	—	—	2300	mA	

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
3. The voltage specification requirements are measured on package pins as near as possible to the processor. Measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. For Voltage less than 1V, TOB will be 50 mV.

Table 7-7. Vcc Sustain Gated (Vcc_{STG}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes [1,2]
Vcc_{STG}	Processor Vcc Sustain supply voltage	Intel® Xeon® E-2300 processor product family	—	1.05	—	V	3
TOB_{STG}	Vcc_{STG} Tolerance	All	+/-5			%	3,4
Icc_{MAX_STG}	Max Current for Vcc_{STG}	Intel® Xeon® E-2300 processor product family	—	—	870	mA	

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.
3. The voltage specification requirements are measured on package pins as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. For Voltage less than 1V, TOB will be 50 mV.

Table 7-8. Processor PLL_OC (Vcc_{PLL_OC}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes[1,2]
Vcc_{PLL_OC}	PLL_OC supply voltage (DC + AC specification)	Intel® Xeon® E-2300 processor product family	—	1.2	—	V	3
TOB_{CCPLL_OC}	Vcc_{PLL_OC} Tolerance	Intel® Xeon® E-2300 processor product family	AC+DC:± 5			%	3,4
$Icc_{MAX\_VCCPLL\_OC}$	Max Current for Vcc_{PLL_OC} Rail	Intel® Xeon® E-2300 processor product family	—	—	251	mA	

Table 7-8. Processor PLL_OC (Vcc_{PLL_OC}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes[1,2]

Notes:
1. Unless otherwise noted, all specifications in this table are based on estimates and simulations or empirical data. These specifications will be updated with characterized data from silicon measurements at a later date.
2. Long term reliability cannot be assured in conditions above or below Max/Min functional limits.
3. The voltage specification requirements are measured on package pins as near as possible to the processor. measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. For Voltage less than 1V, TOB will be 50 mV.

7.2.3 Processor Interfaces DC Specifications

7.2.3.1 DDR4 DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes[1]
V_{IL}	Input Low Voltage	—	—	VREF(INT) - 0.07*VDDQ	V	2, 4, 8, 9, 13
V_{IH}	Input High Voltage	VREF(INT) + 0.07*VDDQ	—	—	V	3, 4, 8, 9, 13
$R_{ON\_UP/DN(DQ)}$	DDR4 Data Buffer pull-up/ down Resistance	Trainable			Ω	11
$R_{ODT(DQ)}$	DDR4 On-die termination equivalent resistance for data signals	Trainable			Ω	11
$V_{ODT(DC)}$	DDR4 On-die termination DC working point (driver set to receive mode)	$0.45*V_{DDQ}$	$0.5*V_{DDQ}$	$0.55*V_{DDQ}$	V	9
$R_{ON\_UP/DN(CK)}$	DDR4 Clock Buffer pull-up/ down Resistance	0.8*Typ	26	1.2*Typ	Ω	5, 11
$R_{ON\_UP/DN(CMD)}$	DDR4 Command Buffer pull-up/ down Resistance	0.8*Typ	20	1.2*Typ	Ω	11
$R_{ON\_UP/DN(CTL)}$	DDR4 Control Buffer pull-up/ down Resistance	0.8*Typ	20	1.2*Typ	Ω	5, 11
$R_{ON\_UP/DN}$ (DDR_VTT_CNTL)	System Memory Power Gate Control Buffer Pull-Up/ down Resistance	40	—	140	Ω	-
I_{LI}	Input Leakage Current (DQ, CK) 0 V $0.2*V_{DDQ}$ $0.8*V_{DDQ}$	—	—	1	mA	-
DDR0_VREF_DQ DDR1_VREF_DQ DDR_VREF_CA	VREF output voltage	$V_{DDQ}/2$-0.06	$V_{DDQ}/2$	$V_{DDQ}/2$+0.06	V	12,14, 15
DDR_RCOMP[0]	ODT resistance compensation	RCOMP values are memory topology dependent.			Ω	6
DDR_RCOMP[1]	Data resistance compensation				Ω	6
DDR_RCOMP[2]	Command resistance compensation				Ω	6

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes[1]

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. V_{IL} is defined as the maximum voltage level at a receiving agent that will be interpreted as a logical low value.
3. V_{IH} is defined as the minimum voltage level at a receiving agent that will be interpreted as a logical high value.
4. V_{IH} and V_{IL} may experience excursions above V_{DDQ}. However, input signal drivers should comply with the signal quality specifications.
5. This is the pull up/down driver resistance after compensation. Note that BIOS power training may change these values significantly based on margin/power trade-off. Refer processor I/O Buffer Models for I/V characteristics.
6. DDR_RCOMP resistance should be provided on the system board with ±1% resistors (except for S-Processor Line, resistors installed on package). DDR_RCOMP resistors are to V_{SS}. DDR_RCOMP resistors are installed on the package.
7. DDR_VREF is defined as $V_{DDQ}/2$ for DDR4
8. R_{ON} tolerance is preliminary and might be subject to change.
9. The value will be set during the MRC boot training within the specified range.
10. Processor may be damaged if V_{IH} exceeds the maximum voltage for extended periods.
11. Final value determined by BIOS power training, values might vary between bytes and/or units.
12. VREF values determined by BIOS training, values might vary between units.
13. VREF(INT) is a trainable parameter whose value is determined by BIOS for margin optimization.
14. DDR1_Vref_DQ connected to Channel 1 VREF_CA.
15. DDR_Vref_CA connected to Channel 0 VREF_CA.

7.2.3.2 PCI Express* Graphics (PEG) DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes[1]
$Z_{TX-DIFF-DC}$	DC Differential Tx Impedance	80	100	120	Ω	1, 5
Z_{RX-DC}	DC Common Mode Rx Impedance	40	50	60	Ω	1, 4
$Z_{RX-DIFF-DC}$	DC Differential Rx Impedance	80	—	120	Ω	1
PEG_RCOMP	resistance compensation	24.75	25	25.25	Ω	2, 3

Notes:
1. Refer to the PCI Express Base Specification for more details.
2. Low impedance defined during signaling. Parameter is captured for 5.0 GHz by RLTX-DIFF.
3. PEG_RCOMP resistance should be provided on the system board with 1% resistors. COMP resistors are to V_{CCIO}. PEG_RCOMP - Intel allows using 24.9 Ω 1% resistors.
4. DC impedance limits are needed to ensure Receiver detect.
5. The Rx DC Common Mode Impedance should be present when the Receiver terminations are first enabled to ensure that the Receiver Detect occurs properly.Compensation of this impedance can start immediately and the 15 Rx Common Mode Impedance (constrained by RLRX-CM to 50 Ω ±20%) should be within the specified range by the time Detect is entered.

7.2.3.3 CMOS DC Specifications (Clock and Power MGMT/Thermal Domains)

Symbol	Parameter	Minimum	Maximum	Units	Notes[1]
V_{IL}	Input Low Voltage	—	VccST * 0.3	V	2, 5, 6
V_{IH}	Input High Voltage	VccST * 0.7	—	V	2, 4, 5, 6
R_{ON}	Buffer on Resistance	23	73	Ω	6
I_{LI}	Input Leakage Current	—	±150	µA	3

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. The Vcc referred to in these specifications refers to instantaneous Vcc levels.
3. For VIN between "0" V and Vcc Measured when the driver is tri-stated.
4. V_{IH} and V_{OH} may experience excursions above Vcc. However, input signal drivers should comply with the signal quality specifications.
5. N/A
6. Voh and Vol need to comply with Vil and Vih specs. This is done by properly calculating pull-up and pull-down resistors.

7.2.3.4 CMOS DC Specifications (Debug and Display Domains)

Symbol	Parameter	Minimum	Maximum	Units	Notes[1]
V_{IL}	Input Low Voltage	—	VccIO * 0.3	V	2, 5, 6

Symbol	Parameter	Minimum	Maximum	Units	Notes[1]
V_{IH}	Input High Voltage	VccIO * 0.7	—	V	2, 4, 5, 6
R_{ON}	Buffer on Resistance	23	73	Ω	-
I_{LI}	Input Leakage Current	—	±150	µA	3

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. The Vcc referred to in these specifications refers to instantaneous Vcc levels.
3. For VIN between "0" V and Vcc Measured when the driver is tri-stated.
4. V_{IH} and V_{OH} may experience excursions above Vcc. However, input signal drivers should comply with the signal quality specifications.
5. N/A
6. Voh and Vol need to comply with Vil and Vih specs. This is done by properly calculating pull-up and pull-down resistors.

7.2.3.5 GTL and Open Drain DC Specifications

Symbol	Parameter	Minimum	Maximum	Units	Notes[1]
V_{IL}	Input Low Voltage (TAP, except PROC_TCK, PROC_TRST#)	-	VccIO * 0.6	V	2, 5, 6
V_{IH}	Input High Voltage (TAP, except PROC_TCK, PROC_TRST#)	VccIO * 0.72	-	V	2, 4, 5, 6
V_{IL}	Input Low Voltage (PROC_TCK,PROC_TRST#)	—	Vcc * 0.3	V	2, 5, 6
V_{IH}	Input High Voltage (PROC_TCK,PROC_TRST#)	Vcc * 0.3	—	V	2, 4, 5, 6
$V_{HYSTERESIS}$	Hysteresis Voltage	Vcc * 0.2	—	V	-
R_{ON}	Buffer on Resistance (TDO)	7	17	Ω	-
V_{IL}	Input Low Voltage (other GTL)	—	Vcc * 0.6	V	2, 5, 6
V_{IH}	Input High Voltage (other GTL)	Vcc * 0.72	—	V	2, 4, 5, 6
R_{ON}	Buffer on Resistance (CFG/BPM)	16	24	Ω	-
R_{ON}	Buffer on Resistance (other GTL)	12	28	Ω	-
I_{LI}	Input Leakage Current	—	±150	µA	3

Notes:
1. Unless otherwise noted, all specifications in this table apply to all processor frequencies.
2. The Vcc_{ST} referred to in these specifications refers to instantaneous $Vcc_{ST/IO}$.
3. For V_{IN} between 0 V and Vcc_{ST}. Measured when the driver is tri-stated.
4. V_{IH} and V_{OH} may experience excursions above Vcc_{ST}. However, input signal drivers should comply with the signal quality specifications.
5. N/A
6. Those V_{IL}/V_{IH} values are based on ODT disabled (ODT Pull-up not exist).

7.2.3.6 PECI DC Characteristics

The PECI interface operates at a nominal voltage set by Vcc_{ST}. The set of DC electrical specifications shown in the following table is used with devices normally operating from a Vcc_{ST} interface supply.

Vcc_{ST} nominal levels will vary between processor families. All PECI devices will operate at the Vcc_{ST} level determined by the processor installed in the system.

Table 7-9. PECI DC Electrical Limits (Sheet 1 of 2)

Symbol	Definition and Conditions	Minimum	Maximum	Units	Notes[1]
R_{up}	Internal pull up resistance	15	45	Ω	3
V_{IN}	Input Voltage Range	-0.15	Vcc_{ST} + 0.15	V	-

Table 7-9. PECI DC Electrical Limits (Sheet 2 of 2)

Symbol	Definition and Conditions	Minimum	Maximum	Units	Notes[1]
$V_{Hysteresis}$	Hysteresis	0.15 * Vcc_{ST}	—	V	-
V_{IL}	Input Voltage Low- Edge Threshold Voltage	—	0.3 * Vcc_{ST}	V	-
V_{IH}	Input Voltage High-Edge Threshold Voltage	0.7 * Vcc_{ST}	—	V	-
C_{bus}	Bus Capacitance per Node	N/A	10	pF	-
C_{pad}	Pad Capacitance	0.7	1.8	pF	-
Ileak000	leakage current @ 0V	—	0.6	mA	-
Ileak025	leakage current @ 0.25* Vcc_{ST}	—	0.4	mA	-
Ileak050	leakage current @ 0.50* Vcc_{ST}	—	0.2	mA	-
Ileak075	leakage current @ 0.75* Vcc_{ST}	—	0.13	mA	-
Ileak100	leakage current @ Vcc_{ST}	—	0.10	mA	-

Notes:
1. Vcc_{ST} supplies the PECI interface. PECI behavior does not affect Vcc_{ST} min/max specifications.
2. The leakage specification applies to powered devices on the PECI bus.
3. The PECI buffer internal pull up resistance measured at 0.75* Vcc_{ST}.

Input Device Hysteresis

The input buffers in both client and host models should use a Schmitt-triggered input design for improved noise immunity. Use the following figure as a guide for input buffer design.

Figure 7-1. Input Device Hysteresis

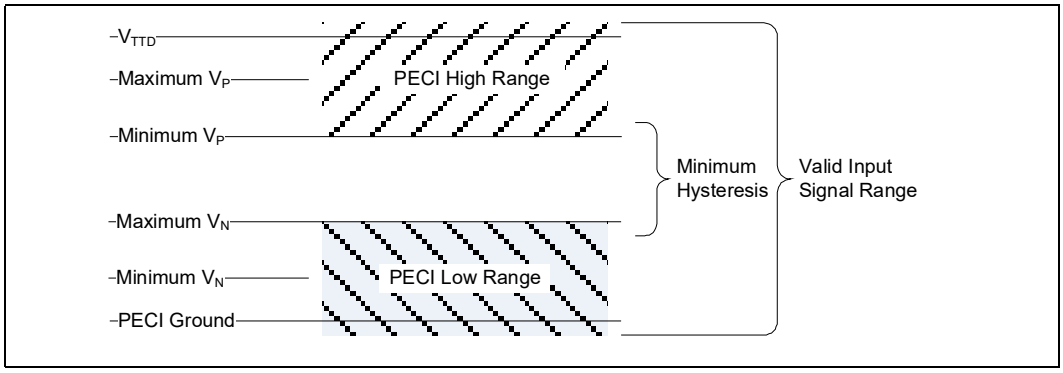

8 Package Mechanical Specifications

8.1 Package Mechanical Attributes

The Intel® Xeon® E-2300 processor uses a Flip Chip technology available in Land Grid Array (LGA). The following table provides an overview of the mechanical attributes of the package.

Package	Parameter	S-Processor Line
		GT1
Package Technology	Package Type	Flip Chip Land Grid Array
	Interconnect	Land Grid Array (LGA)
	Lead Free	N/A
	Halogenated Flame Retardant Free	Yes
Package Configuration	Solder Ball Composition	N/A
	Ball/Pin Count	1200
	Grid Array Pattern	LGA Grid Array
	Land Side Components	Yes
	Die Side Components	Yes
	Die Configuration	1 Die Single-Chip Package with IHS
Package Dimensions	Nominal Package Size	37.5x37.5 mm
	Minimum Ball/Pin pitch	0.9144 mm

Parameter	Minimum	Maximum
Static Compressive per Contact	0.098 N [10gf]	0.254 N [25gf]
Static Pre-Load Compressive	400 N [80 lbf; End of life]	845 N [190 lbf; Beginning of life]
Static Total Compressive	534 N [120 lbf, Beginning of Life] 400 N [80 lbf; End of life]	1068 N [240 lbf; Beginning of life]
Dynamic Compressive	N/A	489.5 N [110 lbf]
Board Transient Bend Strain	N/A	600ue
Maximum Heatsink Mask	N/A	550 g
PnP cover vertical removal for SMT	0.5 lb	Not recommended for system assy

8.2 Package Loading Specifications

Package	Limit	Minimum PCB Thickness Assumptions	Notes
S-Processor Line	67 N (15 lbf)	1.0 mm	1, 2, 3
	111 N (25 lbf)	1.0 mm	1, 2, 3, 4

Notes:
1. The thermal solution attach mechanism should not induce continuous to the package. It may only apply a uniform load to the die to maintain a thermal interface.
2. This specification applies to the uniform compressive load in the direction perpendicular to the dies' top surface. Load should be centred on processor die center.
3. This specification is based on limited testing for design characterization.
4. This load limit assumes the use of a backing plate.

8.3 Package Storage Specifications

Parameter	Description	Minimum	Maximum	Notes
$T_{ABSOLUTE\ STORAGE}$	The non-operating device storage temperature. Damage (latent or otherwise) may occur when subjected to this temperature for any length of time in Intel Original sealed moisture barrier bag and / or box.	-25 °C	125 °C	1, 2, 3
$T_{SUSTAINED\ STORAGE}$	The ambient storage temperature limit (in shipping media) for the sustained period of time.	-40°C	80 °C	1, 2, 3
$RH_{SUSTAINED\ STORAGE}$	The maximum device storage relative humidity for the sustained period of time as specified below in Intel Original sealed moisture barrier bag and / or box.	60% at 24 °C		1, 2, 3
$TIME_{SUSTAINED\ STORAGE}$	Maximum time: associated with customer shelf life in Intel Original sealed moisture barrier bag and / or box.	NA	Moisture Sensitive Devices: 60 months from bag seal date; Non-moisture sensitive devices: 60 months from lot date	1, 2, 3
Storage Conditions	Processors in a non-operational state may be installed in a platform, in a tray, boxed, or loose and may be sealed in airtight package or exposed to free air. Under these conditions, processor landings should not be connected to any supply voltages, have any I/Os biased, or receive any clocks. Upon exposure to "free air" (that is, unsealed packaging or a device removed from packaging material), the processor should be handled in accordance with moisture sensitivity labeling (MSL) as indicated on the packaging material. Boxed Land Grid Array packaged (LGA) processors are MSL 1 ('unlimited' or unaffected) as they are not heated in order to be inserted in the socket.			

Notes:
1. TABSOLUTE STORAGE applies to the un-assembled component only and does not apply to the shipping media, moisture barrier bags or desiccant. Refers to a component device that is not assembled in a board or socket that is not to be electrically connected to a voltage reference or I/O signals.
2. Specified temperatures are based on data collected. Exceptions for surface mount re-flow are specified by applicable JEDEC J-STD-020 and MAS documents. The JEDEC, J-STD-020 moisture level rating and associated handling practices apply to all moisture sensitive de-vices removed from moisture barrier bag.
3. Post board attach storage temperature limits are not specified for non-Intel branded boards. Consult the board manufacturer for storage specifications.

9 CPU And Device IDs

9.1 CPUID

The processor ID and stepping can be identified by the following register contents:

Table 9-1. CPUID Format

CPUID	Reserved [31:28]	Extended Family [27:20]	Extended Model [19:16]	Reserved [15:14]	Processor Type [13:12]	Family Code [11:8]	Model Number [7:4]	Stepping ID [3:0]
A0671h	Reserved	0000000b	1000b	Reserved	00b	0110b	1101b	0001b

- 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 Intel® Core™ processor family.
- 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.
- 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.
- 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.
- The Stepping ID in Bits [3:0] indicates the revision number of that model.
- When EAX is initialized to a value of '1', the CPUID instruction returns the Extended Family, Extended Model, Processor Type, Family Code, Model Number and Stepping ID value in the EAX register. Note that the EDX processor signature value after reset is equivalent to the processor signature output value in the EAX register.

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

9.2 PCI Configuration Header

Table 9-2. Host Device ID (DID0)

Platform	Device ID
Intel® Xeon® E-2300 processor 8 Core	4C43h
Intel® Xeon® E-2300 processor 6 Core	4C53h
Intel® Xeon® E-2300 processor 4 Core	4C63h

Table 9-3. PCI Configuration Header

Byte3	Byte2	Byte1	Byte0	Address
Device ID		**Vendor ID (0x8086)**		00h
Status		Command		04h
Class Code			**Revision ID**	08h
BIST	Header Type	Latency Timer	Cache Line Size	0Ch
Base Address Register0 (BAR0)				10h
Base Address Register1 (BAR1)				14h
Base Address Register2 (BAR2)				18h
Base Address Register3 (BAR3)				1Ch
Base Address Register4 (BAR4)				20h
Base Address Register5 (BAR5)				24h
Card-bus CIS Pointer				28h
Subsystem ID		**Subsystem Vendor ID**		2Ch
Expansion ROM Base Address				30h
Reserved			Capabilities Pointer	34h
Reserved				38h
Maximum Latency	Minimum Grant	Interrupt Pin	Interrupt Line	3Ch

Table 9-4. Graphics Device ID (DID2)

SKU	Processor Step	number of EU's	Device 2 ID	Device 2 Rev
Intel® Xeon® E-2300 processor /pGFX	B0	up to 32	4C9Ah	4h

Table 9-5. Other Device ID (Sheet 1 of 2)

Device	Bus / Device / Function	Device ID
PEG10 G3	0 / 1 / 0	4C01h
PEG11 G3	0 / 1 / 1	4C05h
PEG12 G3	0 / 1 / 2	4C07h
DPTF/DPPM	0 / 4 / 0	4C03h

Table 9-5. Other Device ID (Sheet 2 of 2)

Device	Bus / Device / Function	Device ID
PEG10 G3	0 / 1 / 0	4C01h
PEG11 G3	0 / 1 / 1	4C05h
PEG12 G3	0 / 1 / 2	4C07h
PEG60 x4 PCIe	0 / 6 / 0	4C09h
GNA	0 / 8 / 0	4C11h
NPK	0 / 9 / 0	4C19h