



11th Generation Intel[®] Core[™] Processor

Datasheet, Volume 1 of 2

Supporting 11th Generation Intel[®] Core[™] Processor Families, Intel[®] Pentium[®] Processors, Intel[®] Celeron[®] Processors for UP3, UP4, H35 and H Platforms, formerly known as Tiger Lake

May 2023

Revision 012



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

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

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

All product plans and roadmaps are subject to change without notice.

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

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

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

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Contents

1	Introduction	11
1.1	Processor Volatility Statement	13
1.2	Package Support	13
1.3	Supported Technologies	14
1.3.1	API Support (Windows*)	15
1.4	Power Management Support	15
1.4.1	Processor Core Power Management	15
1.4.2	System Power Management	15
1.4.3	Memory Controller Power Management	16
1.4.4	Processor Graphics Power Management	16
1.4.4.1	Memory Power Savings Technologies	16
1.4.4.2	Display Power Savings Technologies	16
1.4.4.3	Graphics Core Power Savings Technologies	16
1.5	Thermal Management Support	16
1.6	Ball-out Information	17
1.7	Processor Testability	17
1.8	Operating Systems Support	17
1.9	Terminology and Special Marks	17
1.10	Related Documents	21
2	Technologies	22
2.1	Platform Environmental Control Interface (PECI)	22
2.1.1	PECI Bus Architecture	22
2.2	Intel® Virtualization Technology (Intel® VT)	24
2.2.1	Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)	24
2.2.2	Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)	27
2.2.3	Intel® APIC Virtualization Technology (Intel® APICv)	29
2.3	Security Technologies	30
2.3.1	Intel® Trusted Execution Technology (Intel® TXT)	30
2.3.2	Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)	31
2.3.3	Perform Carry-Less Multiplication Quad Word Instruction (PCLMULQDQ)	32
2.3.4	Intel® Secure Key	32
2.3.5	Execute Disable Bit	32
2.3.6	Boot Guard Technology	33
2.3.7	Intel® Supervisor Mode Execution Protection (SMEP)	33
2.3.8	Intel® Supervisor Mode Access Protection (SMAP)	33
2.3.9	Intel® Software Guard Extensions (Intel® SGX)	33
2.3.10	Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)	35
2.3.11	User Mode Instruction Prevention (UMIP)	35
2.3.12	Read Processor ID (RDPID)	35
2.3.13	Total Memory Encryption (Intel® TME)	36
2.3.14	Control-flow Enforcement Technology (Intel® CET)	36
2.3.14.1	Shadow Stack	36
2.3.14.2	Indirect Branch Tracking	37
2.3.15	KeyLocker Technology	37
2.3.16	Devil's Gate Rock (DGR)	37
2.4	Power and Performance Technologies	38
2.4.1	Intel® Smart Cache Technology	38
2.4.2	IA Core Level 1 and Level 2 Caches	38
2.4.3	Intel® Turbo Boost Max Technology 3.0	39

2.4.4	Power Aware Interrupt Routing (PAIR).....	39
2.4.5	Intel® Hyper-Threading Technology (Intel® HT Technology)	39
2.4.6	Intel® Turbo Boost Technology 2.0.....	40
2.4.6.1	Intel® Turbo Boost Technology 2.0 Power Monitoring	40
2.4.6.2	Intel® Turbo Boost Technology 2.0 Power Control.....	40
2.4.6.3	Intel® Turbo Boost Technology 2.0 Frequency	40
2.4.7	Enhanced Intel SpeedStep® Technology	41
2.4.8	Intel® Speed Shift Technology.....	41
2.4.9	Intel® Advanced Vector Extensions 2 (Intel® AVX2)	41
2.4.10	Advanced Vector Extensions 512 Bit (Intel® AVX-512).....	42
2.4.11	Intel® 64 Architecture x2APIC	43
2.4.12	Intel® Dynamic Tuning Technology (DTT)	44
2.4.13	Intel® GNA 2.0 (GMM and Neural Network Accelerator).....	45
2.4.14	Cache Line Write Back (CLWB)	45
2.4.15	Ring Interconnect	45
2.5	Intel® Image Processing Unit (Intel® IPU6).....	46
2.5.1	Platform Imaging Infrastructure.....	46
2.5.2	Intel® Image Processing Unit (Intel® IPU6).....	46
2.6	Debug Technologies	47
2.6.1	Intel® Processor Trace	47
2.6.2	Platform CrashLog	47
2.6.3	Telemetry Aggregator	47
2.7	Clock Topology	49
2.7.1	Integrated Reference Clock PLL	49
2.8	Intel Volume Management Device (VMD) Technology.....	49
2.8.1	Intel Volume Management Device Technology Objective.....	49
2.8.2	Intel Volume Management Device Technology	50
2.8.3	Key Features.....	51
2.9	Deprecated Technologies	51
3	Power Management	52
3.1	Advanced Configuration and Power Interface (ACPI) States Supported	55
3.2	Processor IA Core Power Management	56
3.2.1	OS/HW Controlled P-states	57
3.2.1.1	Enhanced Intel SpeedStep® Technology	57
3.2.1.2	Intel® Speed Shift Technology	57
3.2.2	Low-Power Idle States.....	57
3.2.3	Requesting the Low-Power Idle States	58
3.2.4	Processor IA Core C-State Rules	58
3.2.5	Package C-States	59
3.2.6	Package C-States and Display Resolutions.....	62
3.3	Processor Graphics Power Management	63
3.3.1	Memory Power Savings Technologies.....	63
3.3.1.1	Intel® Rapid Memory Power Management (Intel® RMPM).....	63
3.3.2	Display Power Savings Technologies.....	64
3.3.2.1	Intel® Seamless Display Refresh Rate Switching Technology (Intel® SDRRS Technology) with eDP* Port	64
3.3.2.2	Intel® Automatic Display Brightness	64
3.3.2.3	Smooth Brightness	64
3.3.2.4	Intel® Display Power Saving Technology (Intel® DPST) 6.3.....	64
3.3.2.5	Panel Self-Refresh 2 (PSR 2).....	65
3.3.2.6	Low-Power Single Display Pipe (LPSP).....	65
3.3.2.7	Intel® Smart 2D Display Technology (Intel® S2DDT)	65
3.3.3	Processor Graphics Core Power Savings Technologies	65
3.3.3.1	Intel® Graphics Dynamic Frequency.....	65
3.3.3.2	Intel® Graphics Render Standby Technology (Intel® GRST)	66

3.3.3.3	Dynamic FPS (DFPS)	66
3.4	System Agent Enhanced Intel SpeedStep® Technology	66
3.5	Voltage Optimization	66
3.6	ROP (Rest Of Platform) PMIC	66
3.7	PCI Express* Power Management	67
4	Thermal Management	68
4.1	Processor Thermal Management	68
4.1.1	Thermal Considerations	68
4.1.1.1	Package Power Control	69
4.1.1.2	Platform Power Control	70
4.1.1.3	Turbo Time Parameter (Tau)	71
4.1.2	Configurable TDP (cTDP) and Low-Power Mode	71
4.1.2.1	Configurable TDP	71
4.1.2.2	Low-Power Mode	72
4.1.3	Thermal Management Features	73
4.1.3.1	Adaptive Thermal Monitor	73
4.1.3.2	Digital Thermal Sensor	75
4.1.3.3	PROCHOT# Signal	76
4.1.3.4	PROCHOT Output Only	77
4.1.3.5	Bi-Directional PROCHOT#	77
4.1.3.6	PROCHOT Demotion Algorithm	77
4.1.3.7	Voltage Regulator Protection using PROCHOT#	78
4.1.3.8	Thermal Solution Design and PROCHOT# Behavior	78
4.1.3.9	Low-Power States and PROCHOT# Behavior	79
4.1.3.10	THRMTRIP# Signal	79
4.1.3.11	Critical Temperature Detection	79
4.1.3.12	On-Demand Mode	79
4.1.3.13	MSR Based On-Demand Mode	79
4.1.3.14	I/O Emulation-Based On-Demand Mode	80
4.1.4	Intel® Memory Thermal Management	80
4.2	Thermal and Power Specifications	81
5	Memory	85
5.1	System Memory Interface	85
5.1.1	Processor SKU Support Matrix	85
5.1.2	Supported Population	86
5.1.3	Supported Memory Modules and Devices	88
5.1.4	System Memory Timing Support	89
5.1.5	System Memory Timing Support	89
5.1.6	SAGV Points	90
5.1.7	Memory Controller (MC)	90
5.1.8	System Memory Controller Organization Mode (DDR4)	91
5.1.9	System Memory Frequency	92
5.1.10	Technology Enhancements of Intel® Fast Memory Access (Intel® FMA)	92
5.1.11	Data Scrambling	93
5.1.12	ECC DDR4 H-Matrix Syndrome Codes	93
5.1.13	Data Swapping	94
5.1.14	DDR I/O Interleaving	94
5.1.15	DRAM Clock Generation	95
5.1.16	DRAM Reference Voltage Generation	95
5.1.17	Data Swizzling	95
5.2	Integrated Memory Controller (IMC) Power Management	95
5.2.1	Disabling Unused System Memory Outputs	95
5.2.2	DRAM Power Management and Initialization	96
5.2.2.1	Initialization Role of CKE	97
5.2.2.2	Conditional Self-Refresh	97

	5.2.2.3	Dynamic Power-Down	97
	5.2.2.4	DRAM I/O Power Management	98
	5.2.3	DDR Electrical Power Gating	98
	5.2.4	Power Training	98
6		USB-C* Sub System	99
6.1		USB-C Sub-System General Capabilities.....	100
6.2		USB™ 4 Router	101
	6.2.1	USB 4 Host Router Implementation Capabilities	101
6.3		USB-C Sub-system xHCI/xDCI Controllers.....	102
	6.3.1	USB 3 Controllers	102
		6.3.1.1 Extensible Host Controller Interface (xHCI).....	102
		6.3.1.2 Extensible Device Controller Interface (xDCI).....	102
	6.3.2	USB-C Sub-System PCIe Interface	103
6.4		USB-C Sub-System Display Interface.....	103
7		PCIe* Interface	104
7.1		Processor PCI Express* Interface	104
	7.1.1	PCI Express* Support.....	104
	7.1.2	PCI Express* Lane Polarity Inversion.....	106
	7.1.3	PCI Express* Architecture.....	106
	7.1.4	PCI Express* Configuration Mechanism.....	107
	7.1.5	PCI Express* Equalization Methodology	107
	7.1.6	PCI Express* Hot-Plug.....	108
		7.1.6.1 Presence Detection.....	108
		7.1.6.2 SMI/SCI Generation.....	108
8		Direct Media Interface (DMI) and On Package Interface (OPI)	109
8.1		Direct Media Interface (DMI)	109
	8.1.1	DMI Lane Reversal and Polarity Inversion.....	109
	8.1.2	DMI Error Flow	110
	8.1.3	DMI Link Down	110
8.2		On Package Interface (OPI).....	110
	8.2.1	OPI Support:.....	110
	8.2.2	Functional description:	110
9		Graphics	111
9.1		Processor Graphics.....	111
	9.1.1	Media Support (Intel® QuickSync and Clear Video Technology HD)	111
		9.1.1.1 Hardware Accelerated Video Decode	111
		9.1.1.2 Hardware Accelerated Video Encode.....	112
		9.1.1.3 Hardware Accelerated Video Processing	113
		9.1.1.4 Hardware Accelerated Transcoding	113
9.2		Platform Graphics Hardware Feature	114
	9.2.1	Hybrid Graphics.....	114
10		Display.....	115
10.1		Display Technologies Support.....	115
10.2		Display Configuration	115
10.3		Display Features.....	117
	10.3.1	General Capabilities	117
	10.3.2	Multiple Display Configurations	117
	10.3.3	High-bandwidth Digital Content Protection (HDCP)	117
	10.3.4	DisplayPort*	118
		10.3.4.1 Multi-Stream Transport (MST).....	118
	10.3.5	High-Definition Multimedia Interface (HDMI*).....	120
	10.3.6	embedded DisplayPort* (eDP*)	121

10.3.7	MIPI* DSI	122
10.3.8	Integrated Audio	122
10.3.9	DisplayPort* Input (DP-IN)	123
11	Camera/MIPI	125
11.1	Camera Pipe Support	125
11.2	MIPI* CSI-2 Camera Interconnect	125
11.2.1	Camera Control Logic.....	125
11.2.2	Camera Modules	125
11.2.3	CSI-2 Lane Configuration	126
12	Signal Description	129
12.1	System Memory Interface	131
12.1.1	DDR4 Memory Interface	131
12.1.2	LPDDR4x Memory Interface	134
12.2	PCIe4 Gen4 Interface Signals.....	135
12.3	Direct Media Interface (DMI) Signals.....	135
12.4	PCIe16 Gen4 Interface Signals	136
12.5	Reset and Miscellaneous Signals.....	136
12.6	Display Interfaces	137
12.6.1	Embedded DisplayPort* (eDP*) Signals	137
12.6.2	MIPI DSI* Signals	138
12.6.3	Digital Display Interface (DDI) Signals	138
12.6.4	Digital Display Audio Signals	138
12.7	DP-IN Interface Signals.....	139
12.8	USB Type-C Signals.....	139
12.9	MIPI* CSI-2 Interface Signals	140
12.10	Processor Clocking Signals.....	141
12.11	Testability Signals	141
12.12	Error and Thermal Protection Signals	142
12.13	Processor Power Rails	142
12.14	Ground, Reserved and Non-Critical to Function (NCTF) Signals	143
12.15	Processor Internal Pull-Up / Pull-Down Terminations	144
13	Electrical Specifications	145
13.1	Processor Power Rails	145
13.1.1	Power and Ground Pins	145
13.1.2	Full Integrated Voltage Regulator (FIVR)	145
13.1.3	V _{CC} Voltage Identification (VID)	145
13.2	DC Specifications	146
13.2.1	Processor Power Rails DC Specifications	146
13.2.1.1	V _{CCIN} DC Specifications.....	146
13.2.1.2	V _{CCIN_AUX} DC Specifications.....	148
13.2.1.3	VDD2 DC Specifications	150
13.2.1.4	V _{CCST} DC Specifications.....	150
13.2.1.5	V _{CC1P8A} DC Specifications	152
13.2.1.6	DDR4 DC Specifications	152
13.2.1.7	LPDDR4x DC Specifications	154
13.2.1.8	PCI Express* Graphics (PEG) DC Specifications.....	155
13.2.1.9	Digital Display Interface (DDI) DC Specifications	155
13.2.1.10	embedded DisplayPort* (eDP*) DC Specification	156
13.2.1.11	MIPI* CSI-2 D-Phy Receiver DC Specifications	156
13.2.1.12	CMOS DC Specifications.....	157
13.2.1.13	GTL and OD DC Specification.....	157
13.2.1.14	PECI DC Characteristics	158
14	Package Mechanical Specifications	161
14.1	Package Mechanical Attributes	161



14.2	Package Loading and Die Pressure Specifications	161
14.2.1	Package Loading Specifications	161
14.2.2	Die Pressure Specifications.....	162
14.3	Package Storage Specifications.....	163
15	CPU And Device IDs	165
15.1	CPUID	165
15.2	PCI Configuration Header.....	165

Figures

1-1	11 th Generation Intel® Core™ UP3/UP4/H35/UP3-Refresh//H35-Refresh Processor Lines Platform Diagram	12
1-2	H Processor Line Platform Diagram	13
2-1	Example for PECI Host-Clients Connection.....	23
2-2	Example for PECI EC Connection.....	24
2-3	Device to Domain Mapping Structures	28
2-4	Processor Cache Hierarchy	38
2-5	Processor Camera System	46
2-6	Telemetry Aggregator	48
3-1	UP3 and UP4 Processor Lines Power States	53
3-2	H Processor Line Power States.....	54
3-3	Processor Package and IA Core C-States.....	55
3-4	Idle Power Management Breakdown of the Processor IA Cores.....	57
3-5	Package C-State Entry and Exit	60
4-1	Package Power Control.....	70
4-2	PROCHOT Demotion Signal Description.....	78
5-1	Intel® DDR4 Flex Memory Technology Operations	91
5-2	DDR4 Interleave (IL) and Non-Interleave (NIL) Modes Mapping	95
6-1	USB-C* Sub-system Block Diagram.....	99
7-1	PCI Express* Related Register Structures in the Processor	107
8-1	Example for DMI Lane Reversal Connection	109
10-1	Processor Display Architecture.....	116
10-2	DisplayPort* Overview	118
10-3	HDMI* Overview	120
10-4	MIPI* DSI Overview.....	122
10-5	DP-IN Block Diagram	124
13-1	Input Device Hysteresis.....	159

Tables

1-1	11 th Generation Intel® Core™ Processor Lines.....	12
1-2	Operating Systems Support	18
1-3	Terminology.....	18
1-4	Special marks	21
3-1	System States.....	56
3-2	Integrated Memory Controller (IMC) States	57
3-3	G, S, and C Interface State Combinations	57
3-4	Core C-states.....	59
3-5	Package C-States	61
3-6	Deepest Package C-State Available	64
3-7	TCSS Power State.....	64
3-8	Package C-States with PCIe* Link States Dependencies	68

4-1	Configurable TDP Modes.....	73
4-2	TDP Specifications.....	82
4-3	Package Turbo Specifications.....	84
4-4	Junction Temperature Specifications.....	85
5-1	DDR Support Matrix Table.....	86
5-2	Processor DDR Speed Support.....	86
5-4	LPDDR4x Channels Population Rules.....	87
5-5	DDR4 Memory Down Channels Population Rules.....	87
5-3	DDR Technology Support Matrix.....	87
5-6	H DDR4 SoDIMM Population Configuration.....	88
5-7	Supported DDR4 SoDIMM Module Configurations.....	89
5-8	Supported DDR4 Memory Down Device Configurations.....	89
5-9	Supported LPDDR4x x32 DRAMs Configurations.....	89
5-10	Supported LPDDR4x x64 DRAMs Configurations.....	90
5-11	DDR4 System Memory Timing Support.....	90
5-12	LPDDR4x System Memory Timing Support.....	91
5-13	System Agent Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies.....	91
5-14	ECC H-Matrix Syndrome Codes.....	94
5-15	Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping.....	95
6-1	USB-C* Port Configuration.....	101
6-2	USB-C* Lanes Configuration.....	102
6-3	USB-C* Non-Supported Lane Configuration.....	102
6-4	PCIe via USB4 Configuration.....	104
7-1	PCI Express* 4 -lane Bifurcation and Lane Reversal Mapping.....	105
7-2	PCI Express* 16-lane Bifurcation and Lane Reversal Mapping.....	105
7-3	PCI Express* Maximum Transfer Rates and Theoretical Bandwidth.....	107
9-1	Hardware Accelerated Video Decoding.....	113
9-2	Hardware Accelerated Video Encode.....	114
10-1	Display Ports Availability and Link Rate.....	117
10-2	Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations.....	121
10-3	DisplayPort Maximum Resolution.....	122
10-4	HDMI Maximum Resolution.....	123
10-5	Embedded DisplayPort Maximum Resolution.....	124
10-6	MIPI* DSI Maximum Resolution.....	125
10-7	Processor Supported Audio Formats over HDMI and DisplayPort*.....	125
11-1	CSI-2 Lane Configuration for UP3-Processor Line.....	130
11-2	CSI-2 Lane Configuration for UP4-Processor Line.....	130
11-3	CSI-2 Lane Configuration for H-Processor Line.....	131
12-1	Signal Tables Terminology.....	133
12-2	DDR4 Memory Interface.....	135
12-3	LPDDR4x Memory Interface.....	138
12-4	DMI Interface Signals.....	139
12-5	Processor Clocking Signals.....	145
12-6	Processor Power Rails Signals.....	146
12-7	Processor Pull-up Power Rails Signals.....	147
12-8	GND, RSVD, and NCTF Signals.....	148
13-1	Processor Vcc _{IN} Active and Idle Mode DC Voltage and Current Specifications.....	150
13-2	Vcc _{IN_AUX} Supply DC Voltage and Current Specifications.....	152
13-3	Memory Controller (VDD2) Supply DC Voltage and Current Specifications.....	154
13-4	Vcc Sustain (Vcc _{ST}) Supply DC Voltage and Current Specifications.....	154
13-5	Vcc Sustain Gated (Vcc _{STG}) Supply DC Voltage and Current Specifications.....	155
13-6	Vcc1P8A Supply DC Voltage and Current Specifications.....	156
13-7	DDR4 Signal Group DC Specifications.....	156
13-8	LPDDR4x Signal Group DC Specifications.....	158
13-10	DSI HS Transmitter DC Specifications.....	159



13-11DSI LP Transmitter DC Specifications	159
13-12Digital Display Interface Group DC Specifications (DP/HDMI)	160
13-13DP-IN Group DC Specifications	160
13-14CMOS Signal Group DC Specifications.....	161
13-15 GTL Signal Group and Open Drain Signal Group DC Specifications	161
13-16PECI DC Electrical Limits.....	162
13-17System Reference Clocks DC and AC Specifications	163
14-1 Package Mechanical Attributes.....	165
14-2 Package Loading Specifications.....	166
14-3 Package Loading Specifications.....	166
15-1 CPUID Format	169
15-2 PCI Configuration Header.....	170
15-3 Host Device ID (DID0)	170
15-4 Other Device ID UP3/UP4/H35/UP3-Refresh/H35-Refresh.....	170
15-5 Other Device IDs (H Processor Line).....	171
15-6 Graphics Device ID (DID2).....	172

Revision History

Revision Number	Description	Revision Date
001	Initial Release	September 2020
002	Added UP4 SKU information	October 2020
003	Chapter 4, "Thermal Management" <ul style="list-style-type: none"> Updated Table 4-2 	October 2020
004	Added H35 SKU information Added Chapter 8, "Direct Media Interface (DMI) and On Package Interface (OPI)"	January 2021
005	Added H SKU information Chapter 2, "Technologies" <ul style="list-style-type: none"> Updated Figure 2-5 Chapter 12, "Signal Description" <ul style="list-style-type: none"> Removed the Value column in Section 12.15 Chapter 13, "Electrical Specifications" <ul style="list-style-type: none"> Updated H VCCIN to 2.1V in Table 13-1 	May 2021
006	Added H35-Refresh and UP3-Refresh SKU information	June 2021
007	Chapter 4, "Thermal Management" <ul style="list-style-type: none"> Updated frequency data Chapter 13, "Electrical Specifications" <ul style="list-style-type: none"> Updated Table 13-14 and Table 13-16 Updated Chapter 14, "Package Mechanical Specifications" Chapter 15, "CPU And Device IDs" <ul style="list-style-type: none"> Updated Table 15-1 	August 2021
008	Chapter 10, "Display" <ul style="list-style-type: none"> Updated Section 10.3.9 	November 2021
009	Chapter 4, "Thermal Management" <ul style="list-style-type: none"> Updated Table 4-2. TDP Specifications 	December 2021
010	Chapter 1, "Introduction" <ul style="list-style-type: none"> Added note in Section 1.8, "Operating Systems Support" 	February 2022
011	Chapter 1, "Introduction" <ul style="list-style-type: none"> Updated additional support to include Open CL* 3.0 in Section 1.3.1 Updated information from "S3 is valid on H only" to "S3 is valid on both TGL UP3 and TGL H in Linux only." in Section 1.4.2 Added new note in Section 1.4.2 	May 2022
012	Chapter 10, "Display" <ul style="list-style-type: none"> Updated UP3 and UP4-Processor Line information for DDI A in Table 10-1 Updated UP3, UP4 and H-Processor Line information for DP* with DSC in Table 10-3 Added new note in Table 10-3 Chapter 12, "Signal Description" <ul style="list-style-type: none"> Updated Dir information for DRAM_RESET#/RESET# in Table 12-2 Updated Dir information for DRAM_RESET# in Table 12-3 	May 2023



1 Introduction

11th Generation Intel[®] Core[™] processor is a 64 bit, multi-core processor built on 10 nanometer process technology.

- The UP3-Processor Line and UP4-Processor Line are offered in a 1-Chip Platform that includes the Intel[®] 500 Series Chipset Family On-Package Platform Controller Hub die on the same package as the processor die.
- The UP3-Refresh is a derivative of the UP3-Processor line. In the rest of this document, all references to UP3 processor also covers the UP3-Refresh processor and will be called out separately only when changes exist.
- The H35-processor line and H35-Refresh are derivative of the UP3-Processor line. H35 processor and H35-Refresh have the same specs as the UP3 processor with the only change being the PL1 TDP value of 35W as seen in the table below. In the rest of this document, all references to UP3 processor also covers the H35 processor and H35-Refresh processor and will not be called out separately.
- The H-processor line is offered in a 2-Chip Platform and connected to a discrete PCH chipset on the motherboard.

The following table describes the different 11th Generation Intel[®] Core[™] processor lines:

Table 1-1. 11th Generation Intel[®] Core[™] Processor Lines

Processor Line ¹	Package	Base TDP ³	Processor IA Cores	Graphics Configuration	Platform Type
UP4-Processor Line	BGA1598	9W	4	Up to 96EU	1-Chip
			2		
UP3-Processor Line	BGA1449	28W	4	Up to 96EU	1-Chip
			2		
UP3-Refresh - Processor Line	BGA1449	28W	4	Up to 96EU	1-Chip
			2		
UP3- Pentium/Celeron Processor Line	BGA1449	15W	2	48EU	1-Chip
H35 Processor Line	BGA1449	35W	4	Up to 96EU	1-Chip
H35-Refresh Processor Line	BGA1449	35W	4	Up to 96EU	1-Chip
H- Processor Line	BGA1787	45W	8	Up to 32EU	2-Chip
			6		
<p>Note:</p> <ol style="list-style-type: none"> 1. Processor lines offering may change. 2. For additional TDP Configurations, refer to Chapter 4, "Thermal Management", for adjustment to the base TDP required to preserve base frequency associated with the sustained long-term thermal capability. 3. TDP workload does not reflect I/O connectivity cases such as Thunderbolt[™]. 					

Figure 1-1. 11th Generation Intel® Core™ UP3/UP4/H35/UP3-Refresh//H35-Refresh Processor Lines Platform Diagram

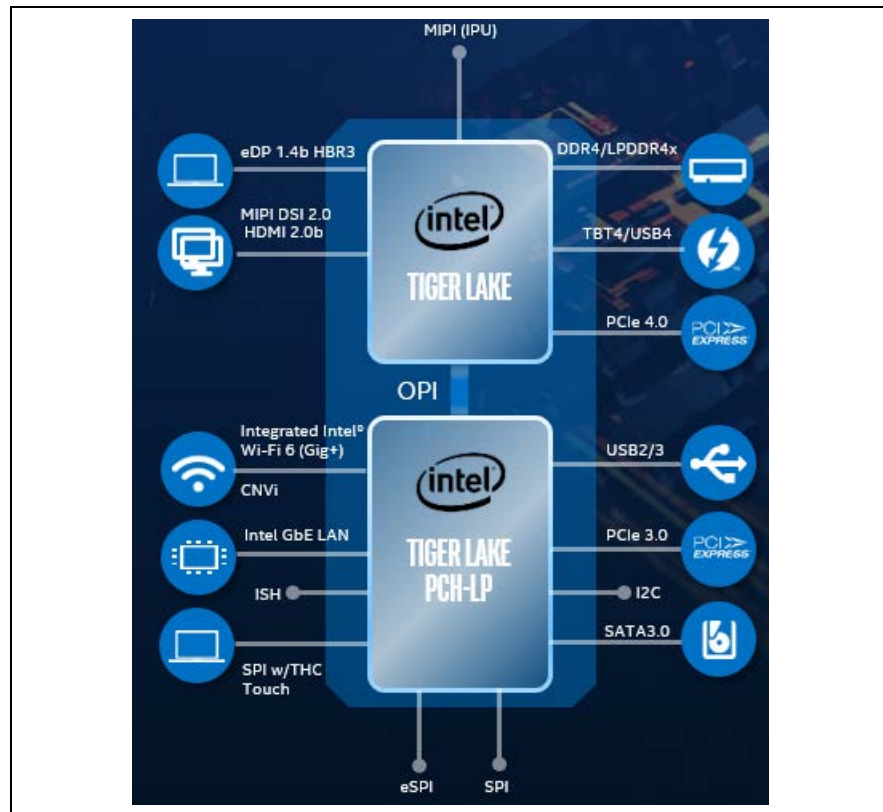
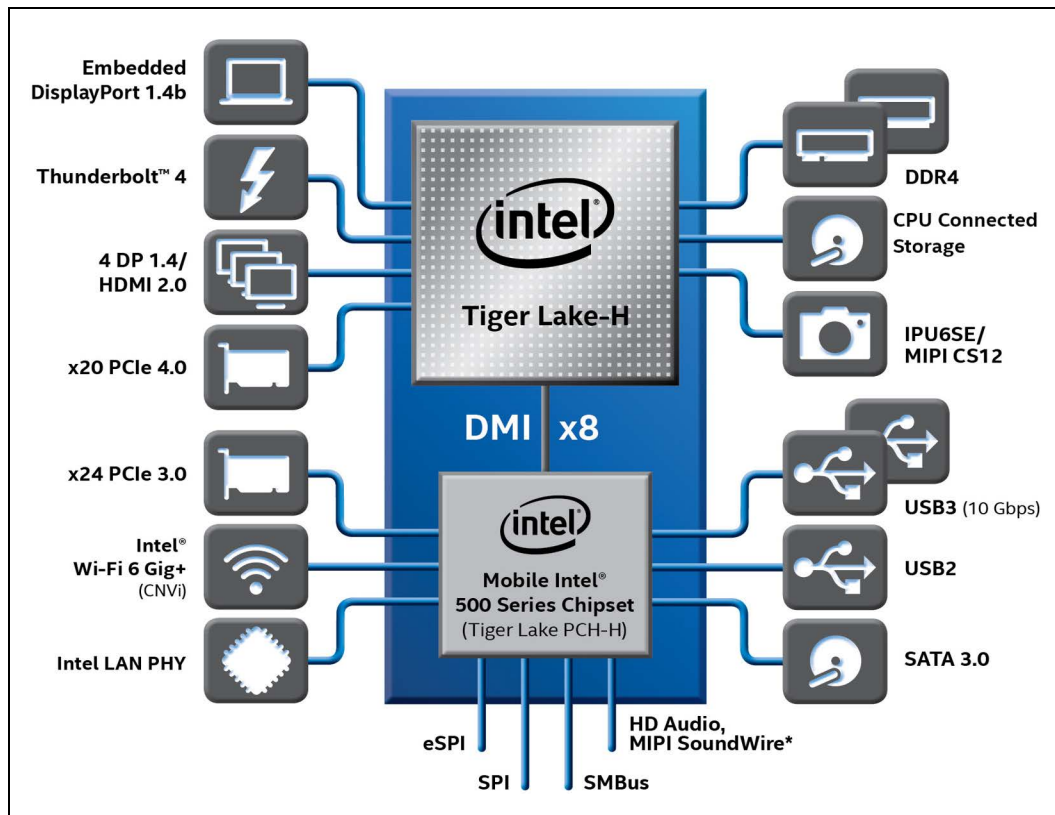


Figure 1-2. H Processor Line Platform Diagram



Not all processor interfaces and features are presented in all Processor Lines. The presence of various interfaces and features will be indicated within the relevant sections and tables.

Note: Throughout this document, the 11th Generation Intel® Core™ processor may be referred to as **processor** and the Intel® 500 Series Chipset Family On-Package Platform Controller Hub (LP) may be referred to as **PCH**.

1.1 Processor Volatility Statement

11th Generation Intel® Core™ Processor families do not retain any end-user data when powered down and/or when the processor is physically removed.

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

1.2 Package Support

The processor is available in the following packages:

- A 26.5 x 18.5 mm, Z= 0.963 +/-0.077 mm (the height from the bottom of the BGA to the top of the die), BGA package for 11th Generation Intel® Core™ UP4-Processor Line.

- A 45.5 x 25 mm, Z=1.185 +/-0.096 mm (the height from the bottom of the BGA to the top of the die), BGA package for 11th Generation Intel[®] Core™ UP3-Processor Line.
- A 50 x 26.5 mm, Z= 1.325 +/- 0.103 mm (the height from the bottom of the BGA to the top of the die), BGA package for H-Processor Line.

1.3 Supported Technologies

- PECCI – Platform Environmental Control Interface
- Intel[®] Virtualization Technology (Intel[®] VT)
- Intel[®] Trusted Execution Technology (Intel[®] TXT)
- Intel[®] Advanced Encryption Standard New Instructions (Intel[®] AES-NI)
- PCLMULQDQ (Perform Carry-Less Multiplication Quad word) Instruction
- Intel[®] Secure Key
- Execute Disable Bit
- Intel[®] Boot Guard
- SMEP – Supervisor Mode Execution Protection
- SMAP – Supervisor Mode Access Protection
- Intel[®] Software Guard Extensions (Intel[®] SGX)
 - May be available on Xeon Processor lines only
- SHA Extensions – Secure Hash Algorithm Extensions
- UMIP – User Mode Instruction Prevention
- RDPID – Read Processor ID
- Total Memory Encryption (Intel[®] TME)
 - Availability may vary between different processor lines.
- Control-flow Enforcement Technology (Intel[®] CET)
- KeyLocker Technology
- Intel[®] Smart Cache Technology
- IA Core Level 1 and Level 2 Caches
- Intel[®] Turbo Boost Technology 2.0
- Intel[®] Turbo Boost Max Technology 3.0
- Power Aware Interrupt Routing (PAIR)
- Intel[®] Hyper-Threading Technology (Intel[®] HT Technology)
- Intel SpeedStep[®] Technology
- Intel[®] Speed Shift Technology
- Intel[®] Advanced Vector Extensions 2 (Intel[®] AVX2)
- Intel[®] Advanced Vector Extensions 512 Bit (Intel[®] AVX-512)
- Intel[®] 64 Architecture x2APIC
- Intel[®] Transactional Synchronization Extensions (Intel[®] TSX-NI)
- Intel[®] Dynamic Tuning Technology (Intel[®] DTT)



- Intel® GNA 2.0 (Intel® GMM and Neural Network Accelerator)
- Cache Line Write Back (CLWB)
- Intel® Image Processing Unit (Intel® IPU)
- Intel® Processor Trace
- Platform CrashLog
- Integrated Reference Clock PLL

Note: The availability of the features above may vary between different processor SKUs. Refer [Chapter 2, “Technologies”](#) for more information.

1.3.1 API Support (Windows*)

- Direct3D* 2015, Direct3D* 12, Direct3D* 11.2, Direct3D* 11.1, Direct3D* 9, Direct3D* 10, Direct2D*
- OpenGL* 4.5
- Open CL* 2.1, Open CL* 2.0, Open CL* 1.2, Open CL* 3.0

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, C10
- Enhanced Intel SpeedStep® Technology
- Intel® Speed Shift Technology

Refer to [Section 3.2, “Processor IA Core Power Management”](#) for more information.

1.4.2 System Power Management

- S0/S0ix, S3, S4, S5
- S3 is valid on both Tiger Lake UP3 and Tiger Lake H in Linux* only.

Note: S3 is not supported in Windows* for both Tiger Lake UP3 and Tiger Lake H. Refer [Chapter 3, “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 [Section 5.2, “Integrated Memory Controller \(IMC\) Power Management”](#) 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 6.3)
- Panel Self-Refresh 2 (PSR 2)
- 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 [Chapter 9, “Graphics”](#) for more information.

1.5 Thermal Management Support

- Digital Thermal Sensor
- Intel® Adaptive Thermal Monitor
- THERMTRIP# and PROCHOT# support
- On-Demand Mode
- Memory Open and 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 - Intel® Dynamic Platform and Thermal Framework (DPTF).

Refer [Chapter 4, “Thermal Management”](#) for more information.

1.6 Ball-out Information

For the PCH ball information, refer to Intel® 500 Series Chipset Family Platform Controller Hub BGA Package Ballout.

1.7 Processor Testability

A DCI on-board connector should be placed, to enable full debug capabilities. For 11th Generation Intel® Core™ processor SKUs, a Direct Connect Interface Tool connector is highly recommended to enable lower C-state to debug. The processor includes boundary-scan for board and system level testability.

1.8 Operating Systems Support

Table 1-2. Operating Systems Support

Processor Line	Windows* 10 64 bit	OS X	Linux* OS	Chrome* OS
UP4-processor line	Yes	Yes	Yes	Yes
UP3-processor line	Yes	Yes	Yes	Yes
H35-processor line	Yes	Yes	Yes	No
UP3-Refresh processor line	Yes	Yes	Yes	Yes
H- processor line	Yes	Yes	Yes	No
Note: Refer to OS Vendor site for more information regarding latest OS revision support.				

1.9 Terminology and Special Marks

Table 1-3. Terminology (Sheet 1 of 4)

Term	Description
4K	Ultra High Definition (UHD)
AES	Advanced Encryption Standard
AGC	Adaptive Gain Control
API	Application Programming Interface
AVC	Advanced Video Coding
BLT	Block Level Transfer
BPP	Bits per Pixel
CDR	Clock and Data Recovery

Table 1-3. Terminology (Sheet 2 of 4)

Term	Description
CTLE	Continuous Time Linear Equalizer
DDC	Digital Display Channel (Refer to PCH Datasheet (#631119) for more details)
DDI	Digital Display Interface for DP or HDMI/DVI
DSI	Display Serial Interface
DDR4	Fourth-Generation Double Data Rate SDRAM Memory Technology
DFE	Decision Feedback Equalizer
D0ix-states	USB controller power states ranging from D0i0 to D0i3, where D0i0 is fully powered on and D0i3 is primarily powered off. Controlled by SW.
DMA	Direct Memory Access
DPPM	Dynamic Power Performance Management
DPTF	Intel Dynamic Platform and Thermal Framework
DMI	Direct Media Interface
DP*	DisplayPort*
DSC	Display Stream Compression
DSI	Display Serial Interface
DTS	Digital Thermal Sensor
ECC	Error Correction Code - used to fix DDR transactions errors
eDP*	Embedded DisplayPort*
EU	Execution Unit in the Graphics Processor
FIVR	Fully Integrated Voltage Regulator
GSA	Graphics in System Agent
GNA	Gauss Newton Algorithm
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® TSX-NI	Intel® Transactional Synchronization Extensions
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.
ITH	Intel® Trace Hub
IOV	I/O Virtualization
IPU	Image Processing Unit
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]. For more information, refer to appropriate BIOS Specification.
LLC	Last Level Cache

Table 1-3. Terminology (Sheet 3 of 4)

Term	Description
LPDDR4x	Low Power Double Data Rate SDRAM memory technology /x- additional power save.
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.
LTR	The Latency Tolerance Reporting (LTR) mechanism enables Endpoints to report their service latency requirements for Memory Reads and Writes to the Root Complex, so that power management policies for central platform resources (such as main memory, RC internal interconnects, and snoop resources) can be implemented to consider Endpoint service requirements.
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]. For more information, refer to the appropriate BIOS specification.
MLC	Mid-Level Cache
MPEG	Motion Picture Expert Group, international standard body JTC1/SC29/WG11 under ISO/IEC that has defined audio and video compression standards such as MPEG-1, MPEG-2, and MPEG-4, etc.
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.
OPVR	On-Package Voltage Regulator
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 to as "chipset".
PECI	Platform Environment Control Interface
PEG	PCI Express* Graphics
PL1, PL2, PL3	Power Limit 1, Power Limit 2, Power Limit 3
PMIC	Power Management Integrated Circuit
Processor	The 64 bit multi-core component (package)
Processor Core	The term "processor core" refers to the 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
PSx	Power Save States (PS0, PS1, PS2, PS3, PS4)
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.
ROP	Rest of Platform
SCI	System Control Interrupt. SCI is used in the ACPI protocol.
SDP	Scenario Design Power
SGX	Software Guard Extension
SHA	Secure Hash Algorithm
SSC	Spread Spectrum Clock
SSIC	Super Speed Inter-Chip

Table 1-3. Terminology (Sheet 4 of 4)

Term	Description
Storage Conditions	Refer Section 14.3, "Package Storage Specifications"
STR	Suspend to RAM
S0ix-states	Processor residency idle standby power states.
TAC	Thermal Averaging Constant
TBT	Thunderbolt™ Interface
TCC	Thermal Control Circuit
TDP	Thermal Design Power
TTV TDP	Thermal Test Vehicle TDP
USB-R	The type of storage redirection used from AMT 11.0 onward. In contrast to IDE-R, which presents remote floppy or CD drives as though they were integrated into the host machine, USB-R presents remote drives as though they were connected via a USB port.
V _{CC}	Processor Core Power Supply
V _{CCGT}	Processor Graphics Power Supply
VCCIO_OUT	I/O Power Supply
V _{CCSA}	System Agent Power Supply
VLD	Variable Length Decoding
VMD	Volume Management Device
VPID	Virtual Processor ID
V _{SS}	Processor Ground

Table 1-4. Special marks

Mark	Definition
[]	Brackets ([]) sometimes follow a ball, pin, registers or a bit name. These brackets enclose a range of numbers, for example, TCP[2:0]_TXRX_P[1:0] may refer to four USB-C* pins or EAX[7:0] may indicate a range that is 8 bits length.
_N / # / B	A suffix of _N or # or B indicates an active low signal. For example, CATERR# Note: _N does not refer to a differential pair of signals such as CLK_P, CLK_N
0x000	Hexadecimal numbers are identified with an x in the number. All numbers are decimal (base 10) unless otherwise specified. Non-obvious binary numbers have the 'b' enclosed at the end of the number. For example, 0101b



1.10 Related Documents

Document	Document Number
Intel® 500 Series Chipset Family On-Package Platform Controller Hub Datasheet, Volume 1 of 2	631119
Intel® 500 Series Chipset Family On-Package Platform Controller Hub Datasheet, Volume 2 of 2	631120
11th Gen Intel® Core™ Processor Family for IoT Platforms - Datasheet Addendum	632133
11th Generation Intel® Core™ Processors Datasheet Volume 2 of 2	631122
11th Generation Intel® Core™ Processor Family Datasheet Volume 2b of 2	643524
11th Generation Intel® Core™ Processor Family Specification Update	631123
Intel® 500 Series Chipset Family On- Package Platform Controller Hub (PCH) Specification Update	630747

§ §

2 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/>

Note: The last section of this chapter is dedicated to deprecated technologies. These technologies are not supported in this processor but were supported in previous generations.

2.1 Platform Environmental Control Interface (PECI)

PECI is an Intel proprietary interface that provides a communication channel between Intel processors and external components such as 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.

Note: PECI interface can be implemented using a single BI_DIRECTIONAL I/O pin serial interface or using the eSPI parallel bus

2.1.1 PECI Bus Architecture

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

The idle state on the bus is '0' (logical low) and near zero (Logical voltage level).

Note: PECI supported frequency range is 3.2 KHz - 1 MHz.

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.**

Figure 2-1. Example for PECE Host-Clients Connection

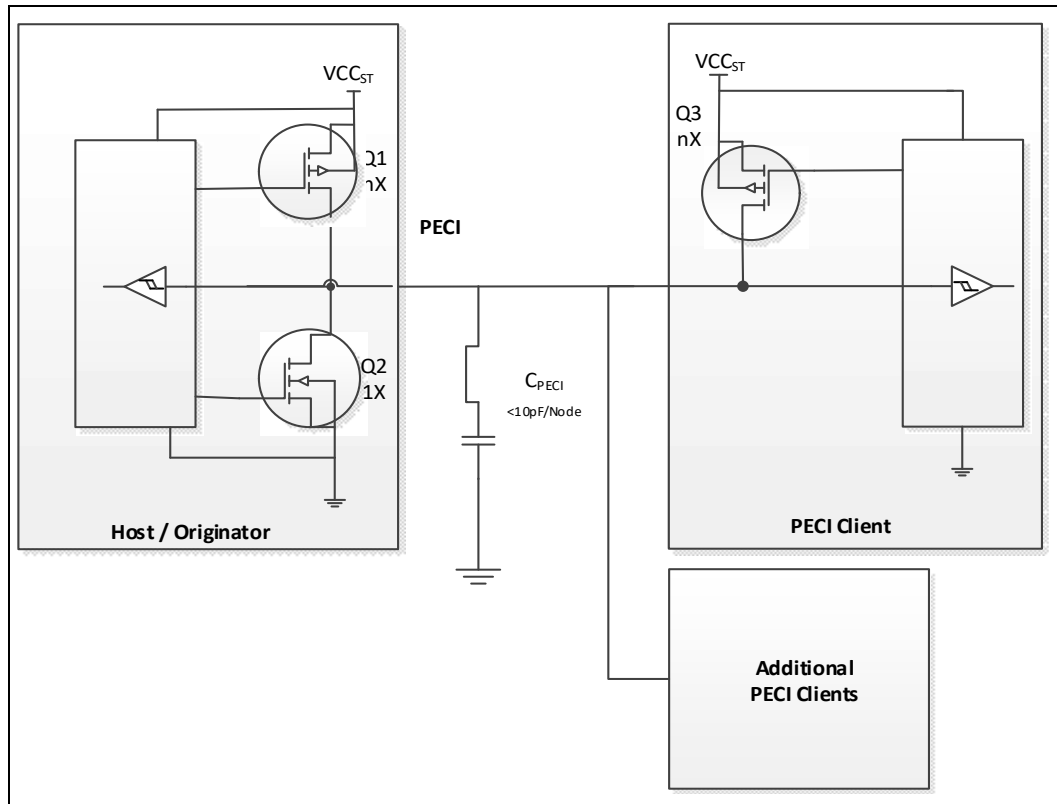
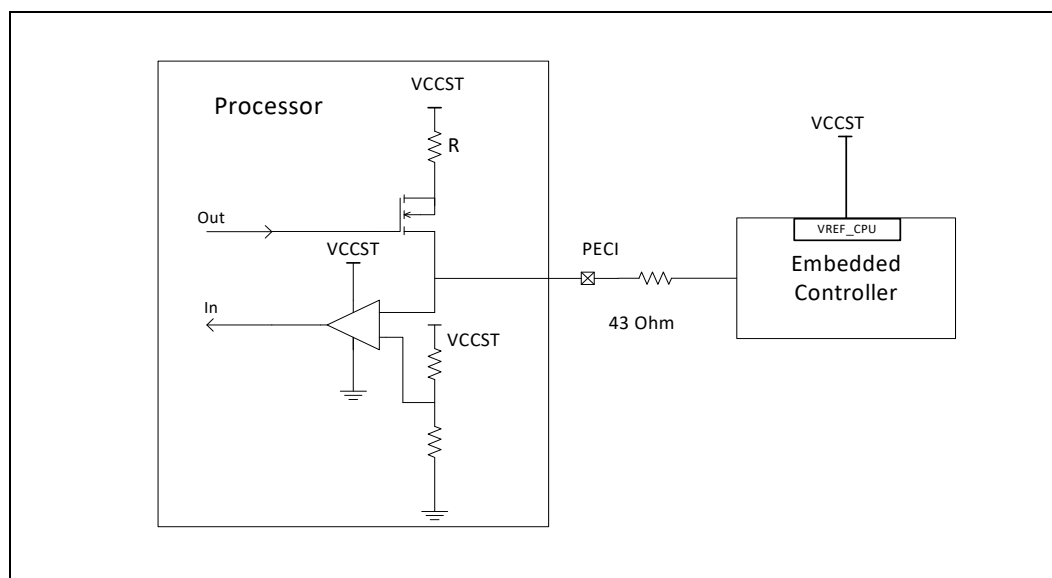


Figure 2-2. Example for PECI EC Connection



2.2 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>

2.2.1 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 Virtualization 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.
- **More 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.
- **More 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:

- **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 i.e. 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" VM-execution 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_CTL2 has bit to indicate the capability, RDMSR can be used to read and query whether the processor supports the capability or not.
- **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 defragmentation, 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. The 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:

- **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**
 - The 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 the 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.



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

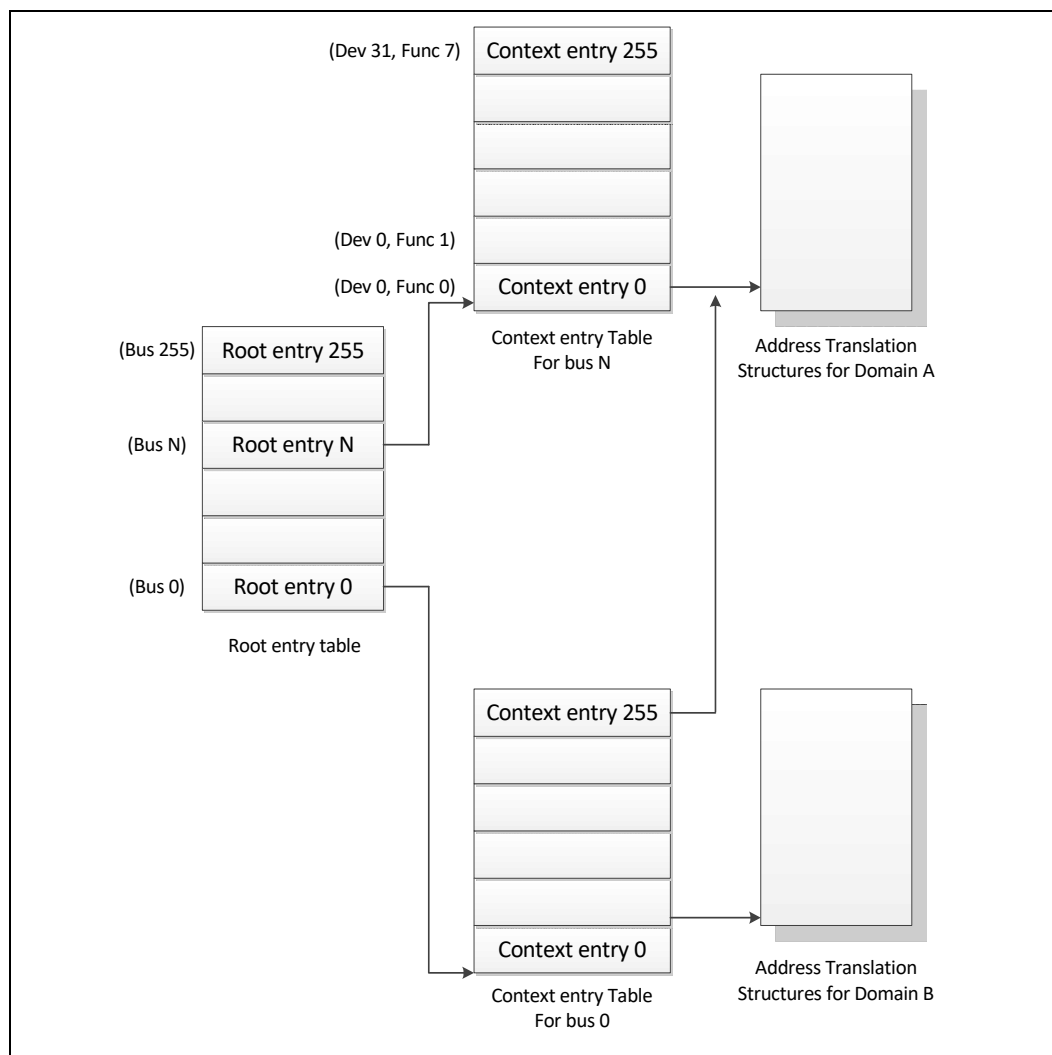
Intel® VT-d Objectives

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 virtualization 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 device assignment restrictions above and to include a multi-level translation table (VT-d Table) that contains Guest specific address translations.

Figure 2-3. 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 above 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 *Intel® Virtualization Technology for Directed I/O Architecture Specification* at <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 the default context
- 39 bit guest physical address and host physical address widths
- Support for 4 K 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 FEEEx_xxxxh) not translated Translation faults result in cycle forwarding to VBIOS region (byte enables masked for writes). Returned data may be bogus for internal agents, PEG/DMI interfaces return unsupported request status
- 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 super-page – support of Intel® VT-d super-page (2 MB, 1 GB) for default Intel® VT-d engine (that covers all devices except IGD)
IGD Intel® VT-d engine does not support super-page and BIOS should disable super-page in default Intel® VT-d engine when iGfx is enabled.

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

2.2.3 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 enable 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.

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

2.3 Security Technologies

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

For more information, refer 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.

2.3.2 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 is 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 industrial 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
- 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 <http://www.intel.com/products/processor/manuals>

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

2.3.3 Perform Carry-Less Multiplication Quad Word Instruction (PCLMULQDQ)

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

2.3.4 Intel® Secure Key

The processor supports Intel® Secure Key (formerly known as Digital Random Number Generator or 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 <http://www.intel.com/products/processor/manuals>

2.3.5 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.

2.3.6 Boot Guard Technology

Boot Guard technology is a part of boot integrity protection technology. Boot Guard can help protect the platform boot integrity by preventing the execution of unauthorized boot blocks. With 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, Boot Guard extends the trust boundary of the platform boot process down to the hardware level.

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 are that Boot Guard can help maintain platform integrity by preventing re-purposing of the manufacturer's hardware to run an unauthorized software stack.

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

2.3.7 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:

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

2.3.8 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:

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

2.3.9 Intel® Software Guard Extensions (Intel® SGX)

Software Guard Extensions (SGX) is a processor enhancement designed to help protect application integrity and confidentiality of secrets and withstands software and certain hardware attacks.

Software Guard Extensions (SGX) architecture provides the capability to create isolated execution environments named Enclaves that operate from a protected region of memory.

Enclave code can be accessed using new special ISA commands that jump into per Enclave predefined addresses. Data within an Enclave can only be accessed from that same Enclave code.

The latter security statements hold under all privilege levels including supervisor mode (ring-0), System Management Mode (SMM) and other Enclaves.

Intel® SGX features a memory encryption engine that both encrypts Enclave memory as well as protects it from corruption and replay attacks.

Intel® SGX benefits over alternative Trusted Execution Environments (TEEs) are:

- Enclaves are written using C/C++ using industry standard build tools.
- High processing power as they run on the processor.
- Large amount of memory is available as well as non-volatile storage (such as disk drives).
- Simple to maintain and debug using standard IDEs (Integrated Development Environment)
- Scalable to a larger number of applications and vendors running concurrently
- Dynamic memory allocation:
 - Heap and thread-pool management
 - On-demand stack growth
 - Dynamic module/library loading
 - Concurrency management in applications such as garbage collectors
 - Write-protection of EPC pages (Enclave Page Cache - Enclave protected memory) after initial relocation
 - On-demand creation of code pages (JIT, encrypted code modules)
- Allow Launch Enclaves other than the one currently provided by Intel
- Maximum protected memory size has increased to 256 MB.
 - Supports 64, 128 and 256 MB protected memory sizes.
- VMM Over-subscription. The VMM over-subscription mechanism allows a VMM to make more resources available to virtual machines than what is actually available on the platform. The initial Intel® SGX architecture was optimized for EPC partitioning/ballooning model for VMMs, where a VMM assigns a static EPC partition to each SGX guest OS without over-subscription and guests are free to manage (i.e. oversubscribe) their own EPC partitions. The Intel® SGX EPC Over subscription Extensions architecture provides a set of new instructions allowing VMMs to efficiently oversubscribe EPC memory for its guest operating systems.

For more information, refer Intel® SGX website at:

<https://software.intel.com/en-us/sgx>.

For more information, refer:

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

Note: Intel® SGX may be available in Xeon SKUs only.

2.3.10 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, but they may also 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/artTGLes/intel-sha-extensions>

2.3.11 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 instructions 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:

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

2.3.12 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:

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

2.3.13 Total Memory Encryption (Intel® TME)

This technology encrypts the platform's entire memory with a single key. TME, when enabled via BIOS configuration, ensures that all memory accessed from the Intel processor is encrypted.

TME encrypts memory accesses using the AES XTS algorithm with 128-bit keys. The encryption key used for memory encryption is generated using a hardened random number generator in the processor and is not exposed to software.

Data in-memory and on the external memory buses is encrypted and exists in plain text only inside the processor. This allows existing software to operate without any modification while protecting memory using TME. TME does not protect memory from modifications.

TME allows the BIOS to specify a physical address range to remain unencrypted. Software running on a TME enabled system has full visibility into all portions of memory that are configured to be unencrypted by reading a configuration register in the processor.

Note: Memory access to nonvolatile memory (Optane) is encrypted as well.

More information on Intel® TME can be found at:

<https://software.intel.com/sites/default/files/managed/a5/16/Multi-Key-Total-Memory-Encryption-Spec.pdf>

Note: Multi-Key Total Memory Encryption (MKTME) is not supported.

2.3.14 Control-flow Enforcement Technology (Intel® CET)

Return-oriented Programming (ROP), and similarly CALL/JMP-oriented programming (COP/JOP), have been the prevalent attack methodology for stealth exploit writers targeting vulnerabilities in programs.

CET provides the following components to defend against ROP/JOP style control-flow subversion attacks:

2.3.14.1 Shadow Stack

A shadow stack is a second stack for the program that is used exclusively for control transfer operations. This stack is separate from the data stack and can be enabled for operation individually in user mode or supervisor mode.

The shadow stack is protected from tamper through the page table protections such that regular store instructions cannot modify the contents of the shadow stack. To provide this protection the page table protections are extended to support an additional attribute for pages to mark them as "Shadow Stack" pages. When shadow stacks are enabled, control transfer instructions/flows such as near call, far call, call to interrupt/exception handlers, and so on. store their return addresses to the shadow stack. The RET instruction pops the return address from both stacks and compares them. If the

return addresses from the two stacks do not match, the processor signals a control protection exception (#CP). Stores from instructions such as MOV, XSAVE, and so on are not allowed to the shadow stack.

2.3.14.2 Indirect Branch Tracking

The ENDBR32 and ENDBR64 (collectively ENDBRANCH) are two new instructions that are used to mark valid indirect CALL/JMP target locations in the program. This instruction is a NOP on legacy processors for backward compatibility.

The processor implements a state machine that tracks indirect JMP and CALL instructions. When one of these instructions is seen, the state machine moves from IDLE to WAIT_FOR_ENDBRANCH state. In WAIT_FOR_ENDBRANCH state the next instruction in the program stream must be an ENDBRANCH. If an ENDBRANCH is not seen the processor causes a control protection fault (#CP), otherwise the state machine moves back to IDLE state.

More information on Intel® CET can be found at:

<https://software.intel.com/sites/default/files/managed/4d/2a/control-flow-enforcement-technology-preview.pdf>

2.3.15 KeyLocker Technology

A method to make long-term keys short-lived without exposing them. This protects against vulnerabilities when keys can be exploited and used to attack encrypted data such as disk drives.

An instruction (LOADIWKEY) allows the OS to load a random wrapping value (IWKey). The IWKey can be backed up and restored by the OS to/from the PCH in a secure manner.

The Software can wrap its own key via the ENCODEKEY instruction and receive a handle. The handle is used with the AES*KL instructions to handle encrypt and decrypt operations. Once a handle is obtained, the software can delete the original key from memory.

2.3.16 Devil's Gate Rock (DGR)

DGR is a BIOS hardening technology that splits SMI (System Management Interrupts) handlers into Ring 3 and Ring 0 portions.

Supervisor/user paging on the smaller Ring 0 portion will enforce access policy for all the ring 3 code with regard to the SMM state save, MSR registers, IO ports and other registers.

The Ring 0 portion can perform save/restore of register context to allow the Ring 3 section to make use of those registers without having access to the OS context or the ability to modify the OS context.

The Ring 0 portion is signed and provided by Intel. This portion is attested by the processor.

2.4 Power and Performance Technologies

2.4.1 Intel® Smart Cache Technology

The Intel® Smart Cache Technology is a shared Last Level Cache (LLC).

- The LLC is non-inclusive.
- The LLC may also be referred to as a 3rd level cache.
- The LLC is shared between all IA cores as well as the Processor Graphics.
- The 1st and 2nd 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 3 MB per physical core and is a 12-way associative cache.

2.4.2 IA Core Level 1 and Level 2 Caches

The 1st level cache is divided into a data cache (DCU) and an instruction cache (IFU). The processor 1st level cache size is 48 KB for data and 32 KB for instructions. The 1st level cache is an 8-way associative cache.

The 2nd level cache holds both data and instructions. It is also referred to as mid-level cache or MLC.

The processor 2nd level cache size is 1.25 MB and is a 20-way non-inclusive associative cache.

Figure 2-4. Processor Cache Hierarchy

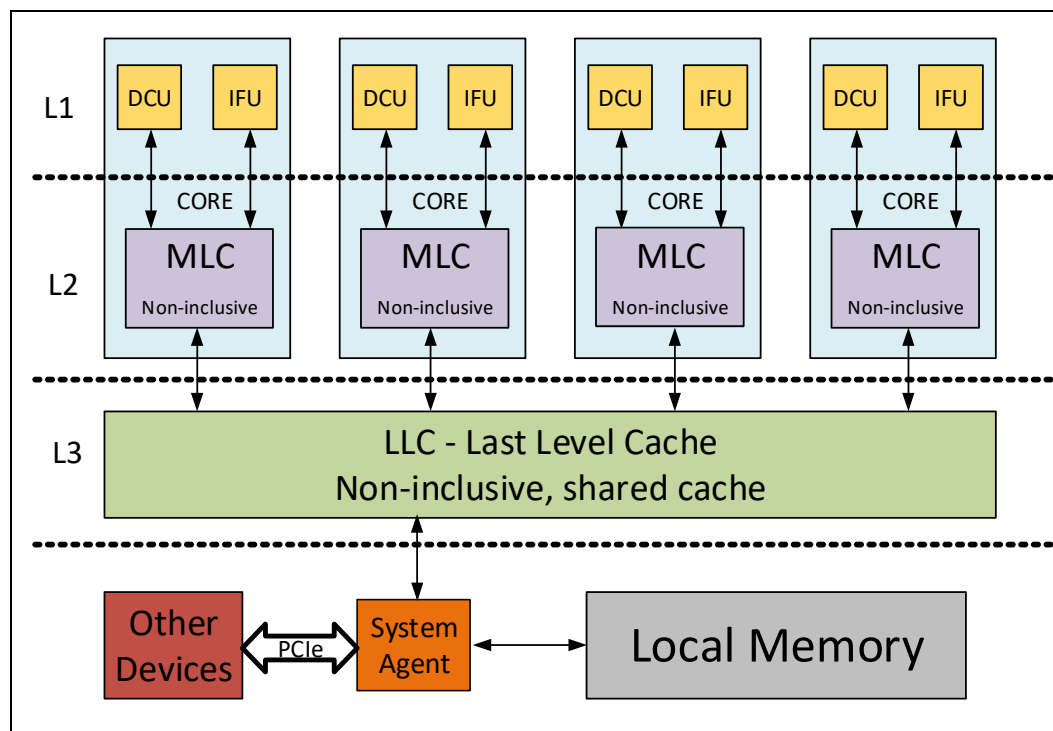


Figure 2-4. Processor Cache Hierarchy

- | |
|---|
| <p>Notes:</p> <ol style="list-style-type: none">1. L1 Data cache (DCU) - 48 KB (per core)2. L1 Instruction cache (IFU) - 32 KB (per core)3. MLC - Mid Level Cache - 1.25 MB (per core) |
|---|

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

Intel® Turbo Boost Max 3.0 Technology is only supported by H/H35 processor lines.

2.4.4 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 most beneficial for high-interrupt scenarios like Gigabit LAN, WLAN peripherals, and so on.

2.4.5 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.

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

2.4.6 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 the 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.

2.4.6.1 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 the 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.

2.4.6.2 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 systems thermal and power limitations. These controls allow for turbo optimizations within system constraints and are accessible using MSR, MMIO, and PECI interfaces.

2.4.6.3 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.

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 [Power Management](#).

2.4.7 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 frequencies 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 processors 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: Because there is low transition latency between P-states, a significant number of transitions per-second are possible.

2.4.8 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 requests the 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.

2.4.9 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 you should consult your 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.

2.4.10 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 prediction, 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.

For more information, refer:

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

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
- **AVX512VP2INTERSECT**: Compute Intersection Between DWORDS/QUADWORDS to a Pair of Mask Registers

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

2.4.11 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 the performance of interrupt delivery
- Reduces the 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 4 KByte page, identical to the xAPIC architecture.
- In the x2APIC mode, APIC registers are accessed through the 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 the x2APIC mode.
- The x2APIC architecture provides backward compatibility to the xAPIC architecture and forwards extensible for future Intel platform innovations.

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

For more information, refer:

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

2.4.12 Intel® Dynamic Tuning Technology (DTT)

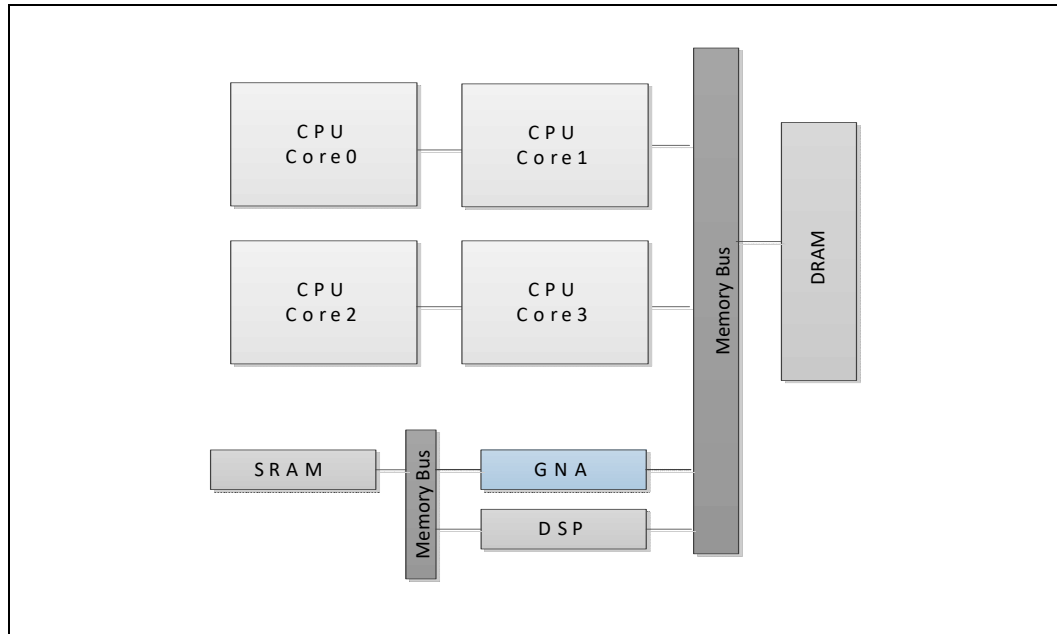
Intel Dynamic Tuning consists of a set of software drivers and applications that allow a system manufacturer to optimize system performance and usability by:

- 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.
- Refer <https://www.intel.com/content/www/us/en/architecture-and-technology/adaptix.html>

2.4.13 Intel® GNA 2.0 (GMM and Neural Network Accelerator)

GNA stands for Gaussian Mixture Model and Neural Network Accelerator.

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



2.4.14 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 the 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 a 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.

For more information, refer:

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

2.4.15 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.

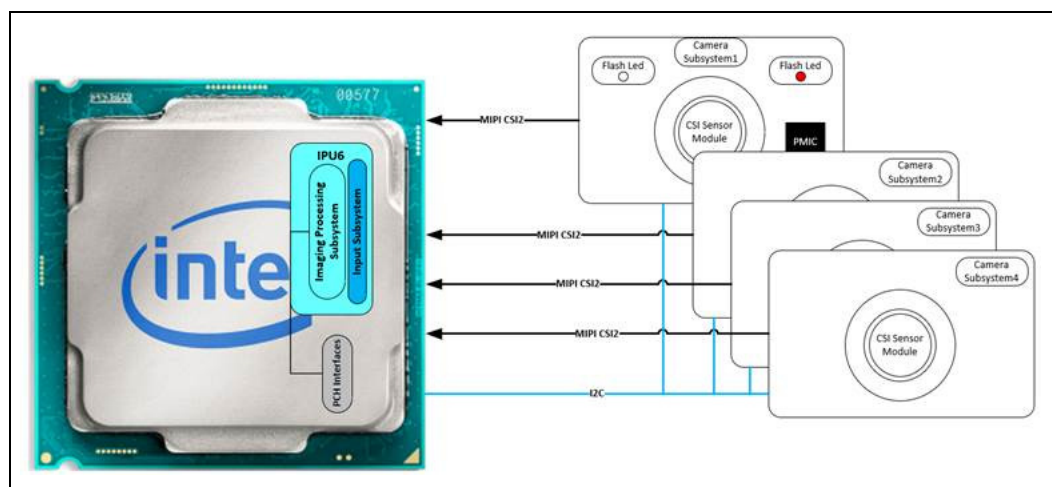
2.5 Intel® Image Processing Unit (Intel® IPU6)

2.5.1 Platform Imaging Infrastructure

The platform imaging infrastructure is based on the following hardware components:

- **Camera Subsystem:** Located in the lid of the system and contains CMOS sensor, flash, LED, I/O interface (MIPI* CSI-2 and I²C*), focus control and other components.
- **Camera I/O Controller:** The I/O controller is located in the processor and contains a MIPI-CSI2 host controller. The host controller is a PCI device (independent of the IPU device). The CSI-2 HCI brings imaging data from an external image into the system and provides a command and control channel for the image using I²C.
- **Intel® IPU (Image Processing Unit):** The IPU processes raw images captured by Bayer sensors. The result images are used by still photography and video capture applications (JPEG, H.264, and so on.).

Figure 2-5. Processor Camera System



2.5.2 Intel® Image Processing Unit (Intel® IPU6)

IPU6 is Intel's 6th generation solution for an Imaging Processing Unit, providing advanced imaging functionality for Intel® Core™ branded processors, as well as more specialized functionality for High Performance Mobile Phones, Automotive, Digital Surveillance Systems (DSS), and other market segments.

IPU6 is a continuing evolution of the architecture introduced in IPU4 and enhanced in IPU5. Additional image quality improvements are introduced, as well as hardware accelerated support for temporal de-noising and new sensor technologies such as Spatially Variant Exposure HDR and Dual Photo Diode, among others.

IPU6 provides a complete high quality hardware accelerated pipeline, and is therefore not dependent on algorithms running on the vector processors to provide the highest quality output.

UP4/UP3 processor has the most advanced IPU6, H has a lighter version of the IPU.

2.6 Debug Technologies

2.6.1 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 on software execution. This provides an 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 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.

2.6.2 Platform CrashLog

The CrashLog feature is intended for use by system builders (OEMs) as a means to triage and perform first level debug of failures.

Additionally, CrashLog enables the BIOS or the OS to collect data on failures with the intent to collect and classify the data as well as analyze failure trends.

CrashLog is a mechanism to collect debug information into a single location and then allow access to that data via multiple methods, including the BIOS and OS of the failing system.

CrashLog is initiated by a Crash Data Detector on observation of error conditions (TCO watchdog timeout, machine check exceptions, and so on).

Crash Data Detector notifies the Crash Data Requester of the error condition in order for the Crash Data Requester to collect Crash Data from several different IPs and/or Crash Nodes and stores the data to the Crash Data Storage (on-die SRAM) prior to the reset.

After the system has rebooted, the Crash Data Collector reads the Crash Data from the Crash Data Storage and makes the data available to either to software and/or back to a central server to track error frequency and trends.

2.6.3 Telemetry Aggregator

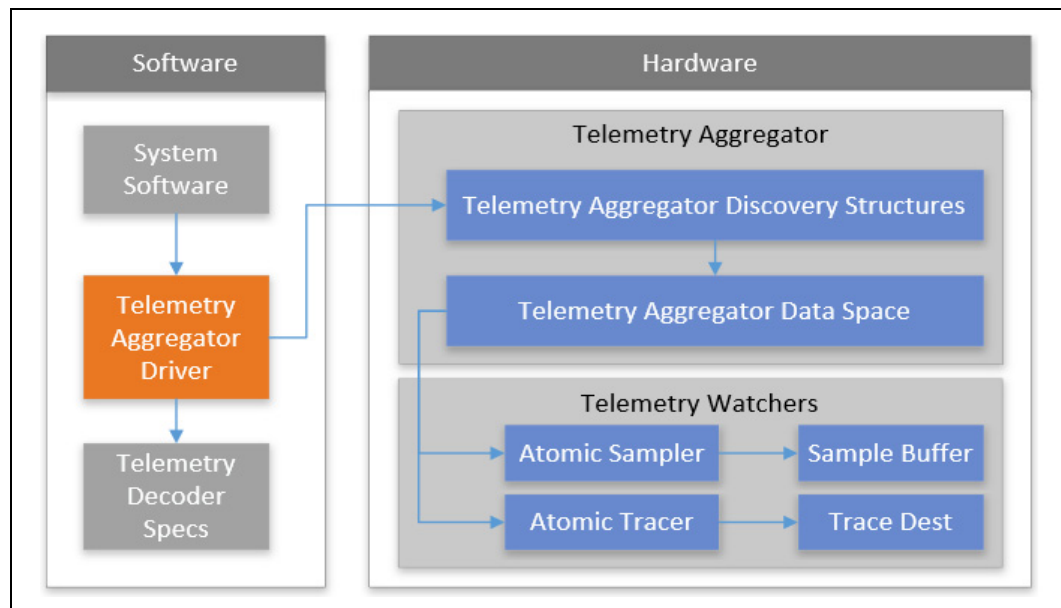
The Telemetry Aggregator serves as an architectural and discoverable interface to hardware telemetry:

- Standardized PCIe discovery solution that enables software to discover and manage telemetry across products
- Standardized definitions for telemetry decode, including data type definitions
- Exposure of commonly used telemetry for power and performance debug including:
 - P-State status, residency and counters
 - C-State status, residency and counters
 - Energy monitoring
 - Device state monitoring (for example, PCIe L1)
 - Interconnect/bus bandwidth counters
 - Thermal monitoring

Exposure of SoC state snapshot for atomic monitoring of package power states, uninterrupted by software that reads.

The Telemetry Aggregator is also a companion to the CrashLog feature where data is captured about the SoC at the point of a crash. These counters can provide insights into the nature of the crash.

Figure 2-6. Telemetry Aggregator



2.7 Clock Topology

The processor has 3 reference clocks that drive the various components within the SoC:

- Processor reference clock or base clock (BCLK). 100 MHz with SSC.
- PCIe reference clock (PCTGLK). 100 MHz with SSC.
- Fixed clock. 38.4 MHz without SSC (crystal clock).

BCLK drives the following clock domains:

- Core
- Ring
- Graphics (GT)
- Memory Controller (MC)
- System Agent (SA)

PCTGLK drives the following clock domains:

- PCIe Controller(s)
- DMI/OPIO

Fixed clock drives the following clock domains:

- Display
- SVID controller
- Time Stamp Counters (TSC)
- Type C subsystem

2.7.1 Integrated Reference Clock PLL

The processor includes a phase lock loop (PLL) that generates the reference clock for the processor from a fixed crystal clock. The processor reference clock is also referred to as Base Clock or BCLK.

By integrating the BCLK PLL into the processor die, a cleaner clock is achieved at a lower power compared to the legacy PCH BCLK PLL solution.

The BCLK PLL has controls for RFI/EMI mitigations as well as Over-clocking capabilities.

2.8 Intel Volume Management Device (VMD) Technology

2.8.1 Intel Volume Management Device Technology Objective

Standard Operating Systems generally recognize individual PCIe Devices and load individual drivers. This is undesirable in some cases such as, for example, when there are several PCIe-based hard-drives connected to a platform where the user wishes to configure them as part of a RAID array. The Operating System current treats individual hard-drives as separate volumes and not part of a single volume.

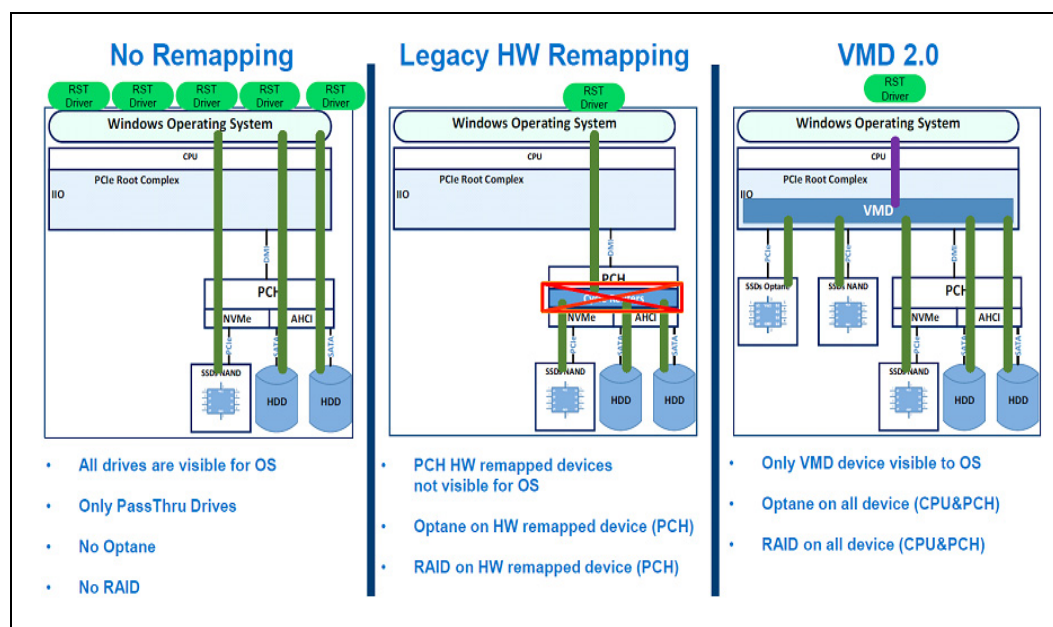
In other words, the Operating System requires multiple PCIe devices to have multiple driver instances, making volume management across multiple host bus adapters (HBAs) and driver instances difficult.

Intel Volume Management Device (VMD) technology provides a means to provide volume management across separate PCI Express HBAs and SSDs without requiring operating system support or communication between drivers. For example, the OS will see a single RAID volume instead of multiple storage volumes, when Volume Management Device is used.

2.8.2 Intel Volume Management Device Technology

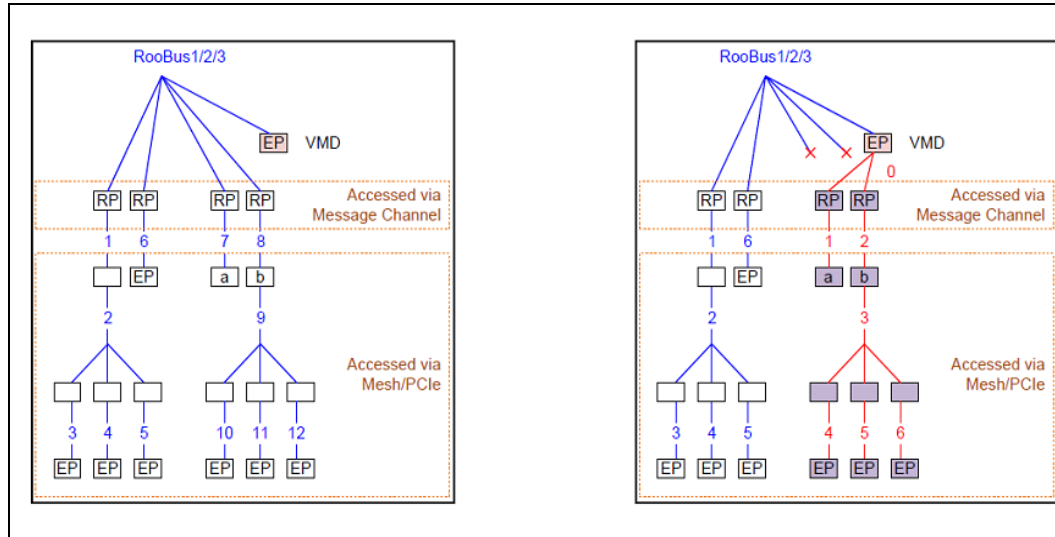
Intel Volume Management Device technology does this by obscuring each storage controller from the OS, while allowing a single driver to be loaded that would control each storage controller.

Intel Volume Management technology requires support in BIOS and driver, memory and configuration space management.



A Volume Management Device (VMD) exposes a single device to the operating system, which will load a single storage driver. The VMD resides in the processor's PCIe root complex and it appears to the OS as a root bus integrated endpoint. In the processor, the VMD is in a central location to manipulate access to storage devices which may be attached directly to the processor or indirectly through the PCH. Instead of allowing individual storage devices to be detected by the OS and therefore causing the OS to load a separate driver instance for each, VMD provides configuration settings to allow specific devices and root ports on the root bus to be invisible to the OS.

Access to these hidden target devices is provided by the VMD to the single, unified driver.



2.8.3 Key Features

Supports MMIO mapped Configuration Space (CFGBAR):

- Supports MMIO Low
- Supports MMIO High
- Supports Register Lock or Restricted Access
- Supports Device Assign
- Function Assign
- MSI Remapping Disable

2.9 Deprecated Technologies

The processor has deprecated the following technologies and they are not longer supported:

- Intel® Memory Protection Extensions (Intel® MPX)
- Branch Monitoring Counters
- Intel® Transactional Synchronization Extensions (Intel® TSX-NI)

§ §

3 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

Figure 3-1. UP3 and UP4 Processor Lines Power States

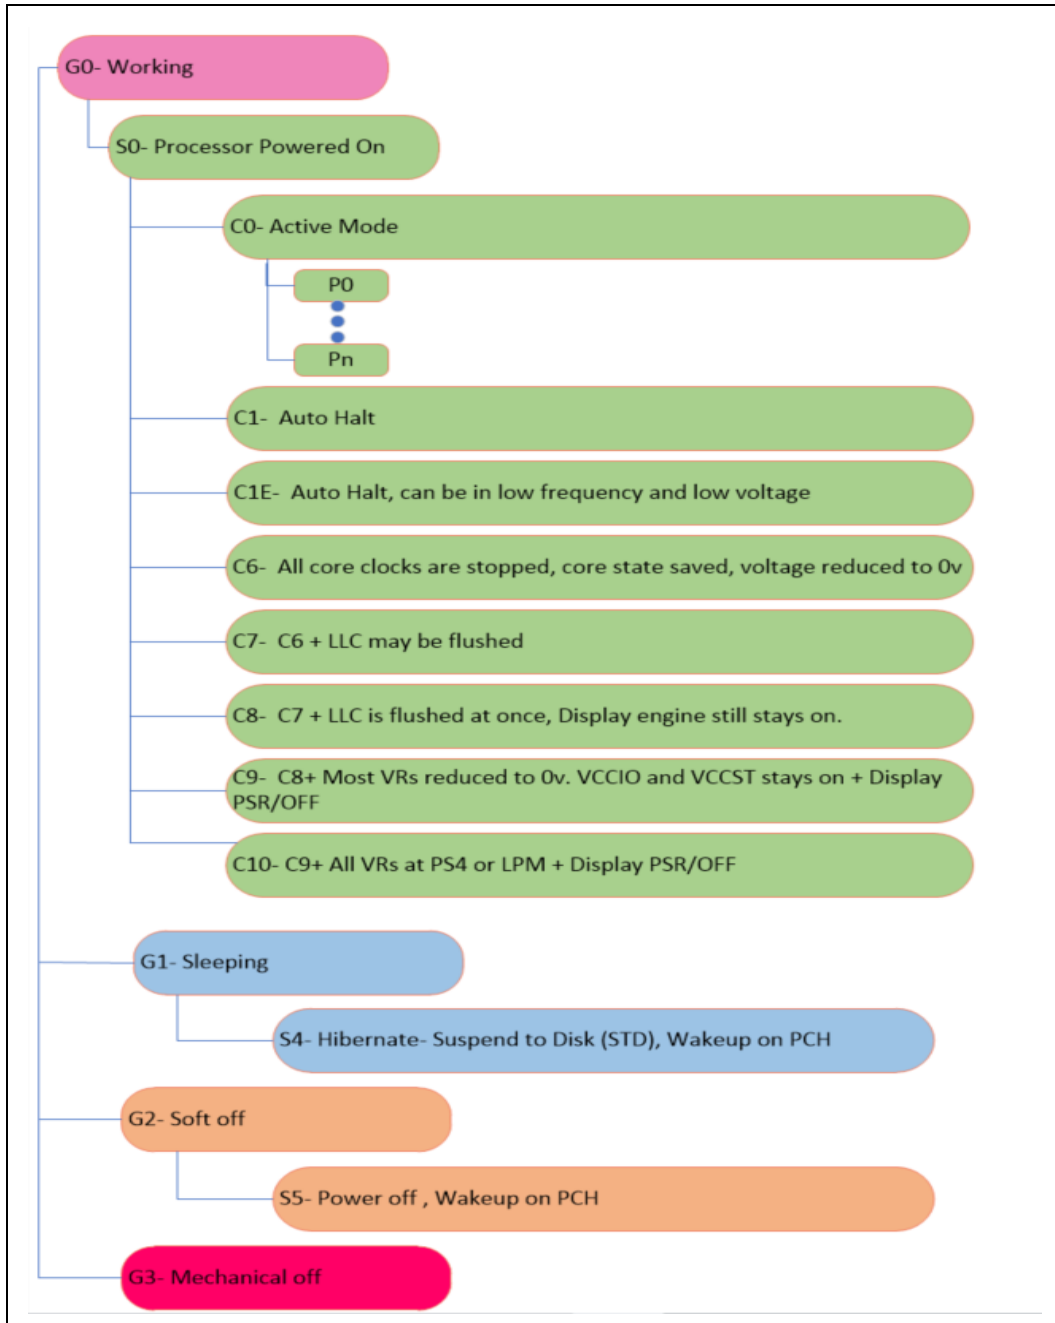


Figure 3-2. H Processor Line Power States

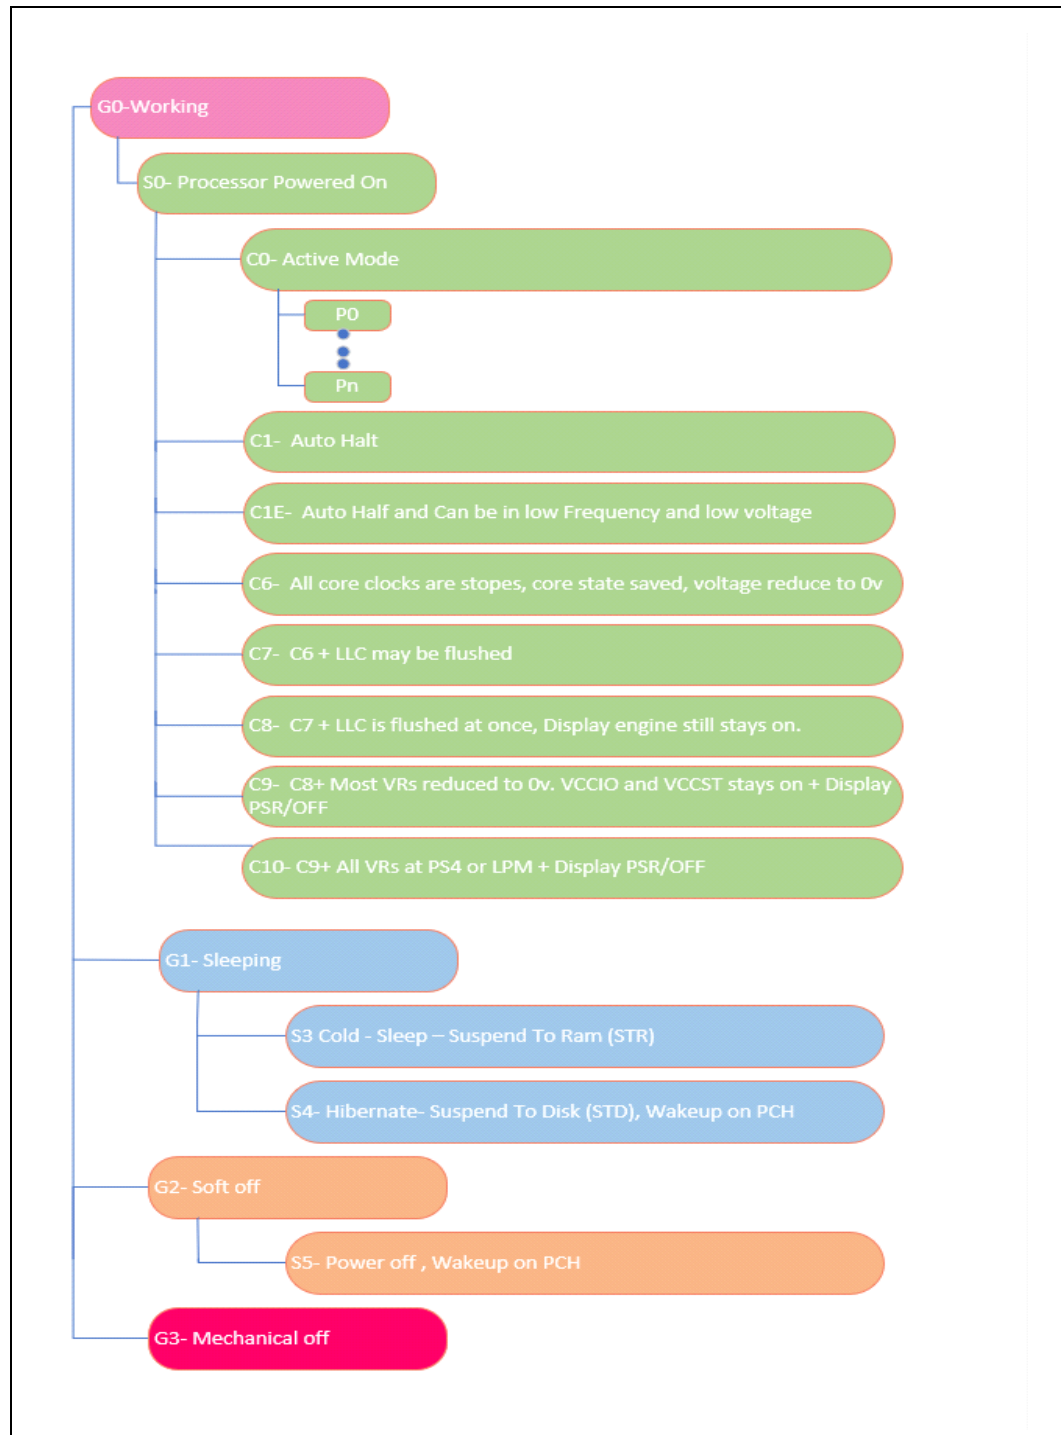
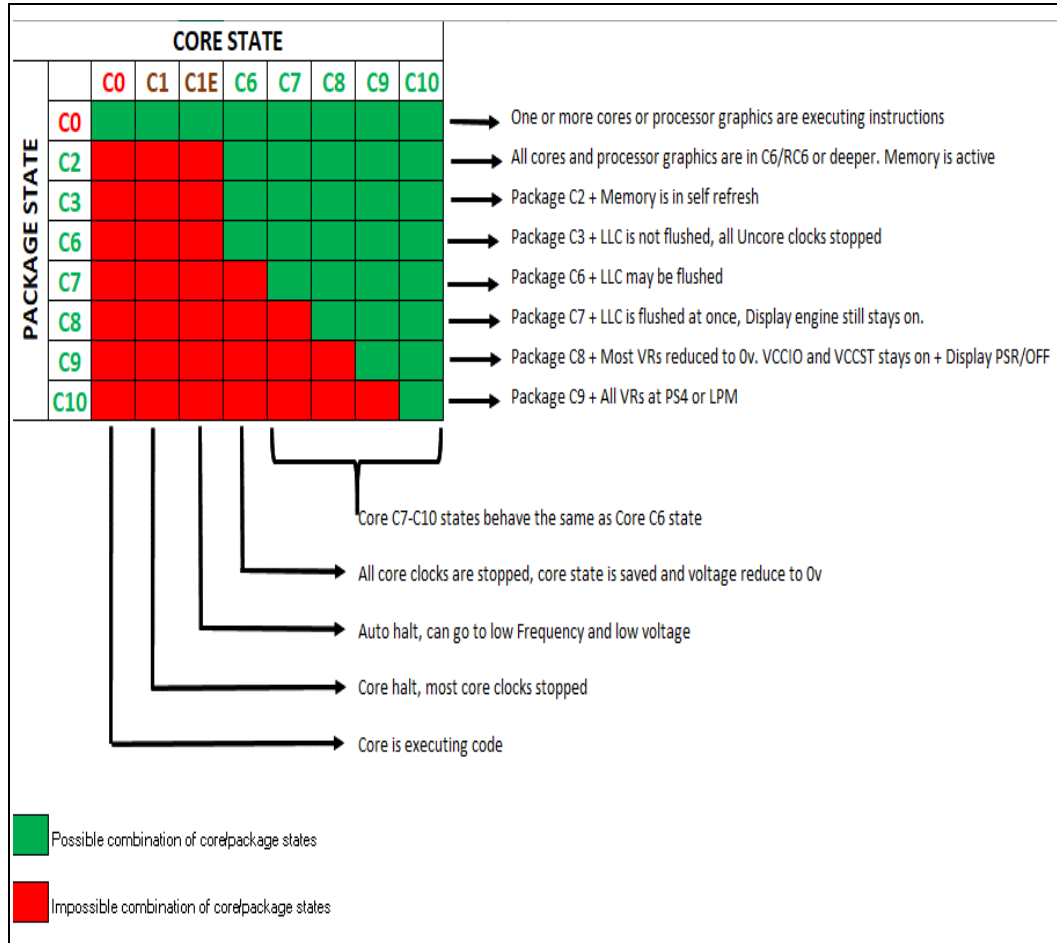


Figure 3-3. Processor Package and IA Core C-States



Notes:

1. **PkgC2/C3 are Non-architectural:** Software cannot request to enter these states explicitly. These states are intermediate states between PkgC0 and PkgC6.
2. There are constraints that prevent the system to go deeper.
3. The "core state" relates to the core which is in the HIGEST power state in the package (most active).

3.1 Advanced Configuration and Power Interface (ACPI) States Supported

This section describes the ACPI states supported by the processor.

Table 3-1. System States (Sheet 1 of 2)

State	Description
G0/S0/C0	Full On: CPU operating. Individual devices may be shut to save power. The different CPU operating levels are defined by Cx states.
G0/S0/Cx	Cx state: CPU manages C-states by itself and can be in low power state

Table 3-1. System States (Sheet 2 of 2)

State	Description
G1/S3	Suspend-To-RAM (STR): The system context is maintained in system DRAM, but power is shut to non-critical circuits. Memory is retained, and refreshes continue. All external clocks are shut off; RTC clock and internal ring oscillator clocks are still toggling. In S3 (H only), SLP_S3 signal stays asserted, SLP_S4 and SLP_S5 are inactive until a wake occurs.
G1/S4	Suspend-To-Disk (STD): The context of the system is maintained on the disk. All power is then shut to the system except to the logic required to resume. Externally appears same as S5 but may have different wake events. In S4, SLP_S3 and SLP_S4 both stay asserted and SLP_S5 is inactive until a wake occurs.
G2/S5	Soft Off: System context not maintained. All power is shut except for the logic required to restart. A full boot is required when waking. Here, SLP_S3, SLP_S4, and SLP_S5 are all active until a wake occurs.
G3	Mechanical OFF: System context not maintained. All power shut except for the RTC. No "Wake" events are possible because the system does not have any power. This state occurs if the user removes the batteries, turns off a mechanical switch, or if the system power supply is at a level that is insufficient to power the "waking" logic. When system power returns the transition will depend on the state just prior to the entry to G3.

Table 3-2. 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 3-3. G, S, and C Interface State Combinations

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	C2	Deep Sleep	On	Deep Sleep
G0	S0	C3	Deep Sleep	On	Deep Sleep
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. S3 valid for H only.
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

3.2 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.

3.2.1 OS/HW Controlled P-states

3.2.1.1 Enhanced Intel SpeedStep® Technology

Enhanced Intel SpeedStep® Technology enables OS to control and select P-state. For more information, refer [Section 2.4.7, “Enhanced Intel SpeedStep® Technology”](#).

3.2.1.2 Intel® Speed Shift Technology

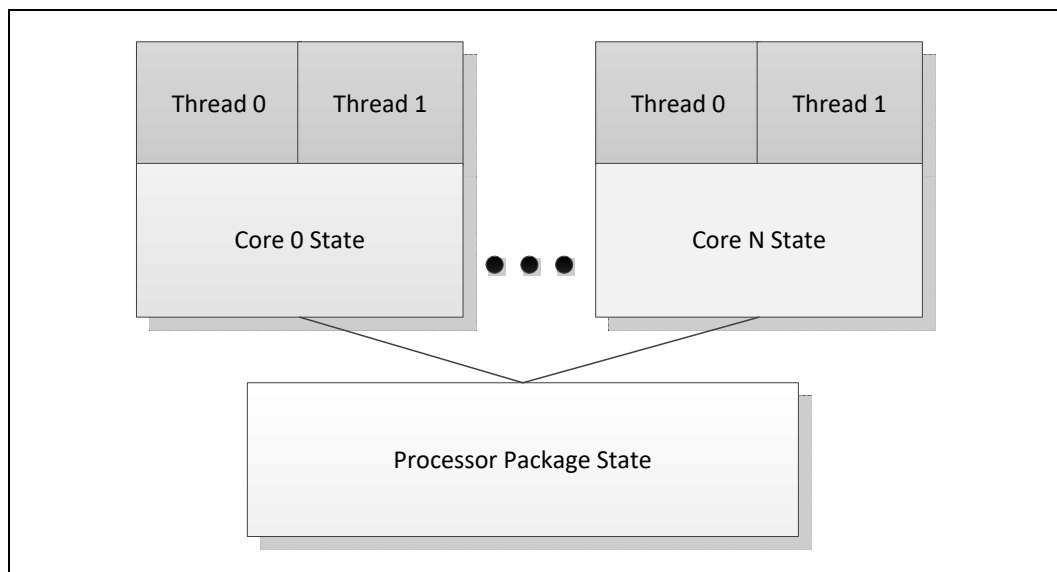
Intel® Speed Shift Technology is an energy efficient method of frequency control by the hardware rather than relying on OS control. For more details, refer [Section 2.4.8, “Intel® Speed Shift Technology”](#).

3.2.2 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 occurs 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 3-4. 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.

3.2.3 Requesting the 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, the 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 the 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.

3.2.4 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 C6 state, resulting in a processor IA core C1E state). Refer to 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.

Table 3-4. Core C-states (Sheet 1 of 2)

Core C-State	C-State Request Instruction	Description
C0	N/A	The normal operating state of a processor IA core where a code is being executed
C1	MWAIT(C1)	AutoHALT - core execution stopped, autonomous clock gating (package in C0 state)

Table 3-4. Core C-states (Sheet 2 of 2)

C1E	MWAIT(C1E)	Core C1 + lowest frequency and voltage operating point (package in C0 state)
C6-C10	MWAIT(C6/7/7s/ C8/9/10) or IO read=P_LVL3/4/5/ 6/7/8	Processor IA, flush their L1 instruction cache, the L1 data cache, and L2 cache to the LLC shared cache cores save their architectural state to an SRAM before reducing IA cores voltage, if possible may also be reduced to 0V. Core clocks are off. C7s is C7 with an additional PLL off.

Core 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.

C-State auto-demotion:

- C7/C6 to C1/C1E

The decision to demote a processor IA core from C6/C7 to C1/C1E 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 C1 until a sufficient amount of residency has been established. At that point, a processor IA core is allowed to go into C6 or C7. 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 C1 state.

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.

3.2.5 Package C-States

The processor supports C0, C2, C3, C6, 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 than 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 3-5. Package C-State Entry and Exit

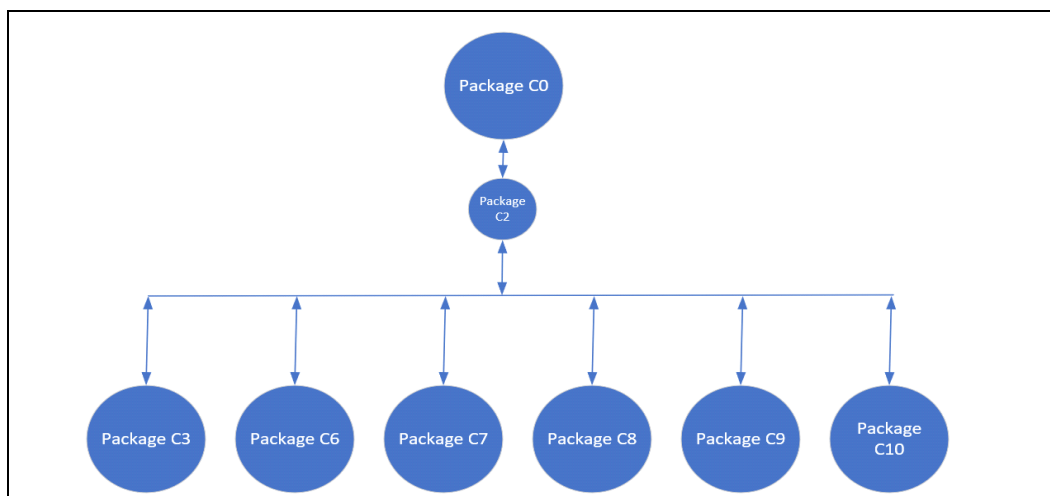


Table 3-5. Package C-States (Sheet 1 of 2)

Package C state	Description	Dependencies
PKG C0	Processor active state. At least one IA core in C0 or Processor Graphic in RC0 (Graphics active state).	-
PKG C2	Cannot be requested explicitly by the Software. All processor IA cores in C6 or deeper + Processor Graphic cores in RC6, memory path may be open. The processor will enter Package C2 when: <ul style="list-style-type: none"> • Transitioning from Package C0 to deep Package C state or from deep Package C state to Package C0. • All IA cores requested C6 or deeper + Processor Graphic cores in RC6 but there are constraints (LTR, programmed timer events in the near future and so forth) prevent entry to any state deeper than C2 state. • All IA cores requested C6 or deeper + Processor Graphic cores in RC6 but a device memory access request is received. Upon completion of all outstanding memory requests, the processor transitions back into a deeper package C-state. 	All processor IA cores in C6 or deeper. Processor Graphic cores in RC6.

Table 3-5. Package C-States (Sheet 2 of 2)

Package C state	Description	Dependencies
PKG C3	<p>All cores in C6 or deeper + Processor Graphics in RC6, LLC may be flushed and turned off, memory in self refresh, memory clock stopped.</p> <p>The processor will enter Package C3 when:</p> <ul style="list-style-type: none"> All IA cores in C6 or deeper + Processor Graphic cores in RC6. The platform components/devices allows proper LTR for entering Package C3. 	<p>All processor IA cores in C6 or deeper.</p> <p>Processor Graphics in RC6.</p> <p>memory in self refresh, memory clock stopped.</p> <p>LLC may be flushed and turned off.</p>
PKG C6	<p>Package C3 + BCLK is off + IMVP VRs voltage reduction/PSx state is possible.</p> <p>The processor will enter Package C6 when:</p> <ul style="list-style-type: none"> All IA cores in C6 or deeper + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C6. 	<p>Package C3.</p> <p>BCLK is off.</p> <p>IMVP VRs voltage reduction/PSx state is possible.</p>
PKG C7	<p>Package C6 + If all IA cores requested C7, LLC ways may be flushed until it is cleared. If the entire LLC is flushed, voltage will be removed from the LLC.</p> <p>The processor will enter Package C7 when:</p> <ul style="list-style-type: none"> All IA cores in C7 or deeper + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C7. 	<p>Package C6.</p> <p>If all IA cores requested C7.</p> <p>LLC ways may be flushed until it is cleared.</p> <p>If the entire LLC is flushed, voltage will be removed from the LLC.</p>
PKG C7S	<p>Package C6 + If all IA cores requested C7S, LLC is flushed in a single step, voltage will be removed from the LLC.</p> <p>The processor will enter Package C7 when:</p> <ul style="list-style-type: none"> All IA cores in C7S or deeper + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C7S. 	<p>Package C6</p> <p>If all IA cores requested C7S, LLC is flushed in a single step, voltage will be removed from the LLC.</p>
PKG C8	<p>Package C7 + LLC should be flushed at once.</p> <p>The processor will enter Package C8 when:</p> <ul style="list-style-type: none"> All IA cores in C8 or deeper + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C8. 	<p>Package C7 + LLC should be flushed at once.</p>
PKG C9	<p>Package C8 + display in PSR or powered off + most Uncore voltages at 0V. IA, GT and SA voltages are reduced to 0 V, while VCC_{IO_OUT} stays on.</p> <p>The processor will enter Package C9 when:</p> <ul style="list-style-type: none"> All IA cores in C9 or deeper + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C9. 	<p>Package C8.</p> <p>All IA cores in C9 or deeper.</p> <p>Display in PSR or powered off¹.</p> <p>VCC_{IO_OUT} stays on.</p>
PKG C10	<p>Package C9 + all VRs at PS4 or LPM + crystal clock off.</p> <p>The processor will enter Package C10 when:</p> <ul style="list-style-type: none"> All IA cores in C10 + Processor Graphic cores in RC6. The platform components/devices allow proper LTR for entering Package C10. 	<p>Package C9.</p> <p>All VRs at PS4 or LPM.</p> <p>Crystal clock off.</p>
<p>Notes:</p> <ul style="list-style-type: none"> Display In PSR is only on single embedded panel configuration and panel support PSR feature. TCSS may enter lowest power state (TC Cold) when no device attached to any of the TCSS ports. 		

Package C-State Auto-Demotion

The Processor may demote the Package C state to a shallower C state, for example instead of going into package C10, it will demote to package C8 (and so on as required). The processor decision to demote the package C state is based on the required C states latencies, entry/exit energy/power and devices LTR.

Modern Standby

Modern Standby 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. Modern Standby requires proper BIOS (refer BIOS specification in [Section 1.10, "Related Documents"](#)) and OS configuration.

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.

C6DRAM

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 savings. The SRAM modules operate on the sustained voltage rail (V_{ccST}).

The memory region used for C6DRAM 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.

Note: The availability of C6DRAM may vary between processor lines offers.

3.2.6 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.

Table 3-6. Deepest Package C-State Available

Resolution	Number of Displays	PSR Enabled ⁴	PSR Disabled ⁴
Up to 5120x3200 60 Hz ³	Single	PC10	PC8
Notes: 1. All Deep states are with Display on. 2. The deepest C-state has variance, dependent various parameters such as SW and Platform Devices. 3. Relevant to all Processor lines.			

Table 3-7. TCSS Power State

TCSS Power State	Allowed Package C States	Device Attached	Description
TC0	PC0-PC3	Yes	xHCI, xDCI, USB4 controllers may be active. USB4 DMA / PCIE may be active.
TC7	PC6-PC10	Yes	xHCI and xDCI are in D3. USB4 controller is in D3 or D0 idle. USB4 PCIe is inactive.
TC-Cold	PC3-PC10	No	xHCI/xDCI/TBT DMA/TBT PCIe are in D3. IOM is active.
TC10	PC6-PC10	No	Deepest power state. xHCI and xDCI are in D3. USB4 is in D3 or D0 idle. USB4 PCIe is inactive. IOM is inactive.
IOM - TCSS Input Output Manager: <ul style="list-style-type: none"> The IOM interacts with the SoC to perform power management, boot, reset, connect, and disconnect devices to Type-C sub-system. TCSS Devices (XHCI/XDCI/TBT Controller) - power state: <ul style="list-style-type: none"> D0 - Device at Active state. D3 - Device at lowest-powered state. 			

3.3 Processor Graphics Power Management

3.3.1 Memory Power Savings Technologies

3.3.1.1 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.

3.3.2 Display Power Savings Technologies

3.3.2.1 Intel® Seamless Display Refresh Rate Switching Technology (Intel® SDRRS Technology) 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 the user may be viewing web pages or social media sites while in battery mode.

3.3.2.2 Intel® Automatic Display Brightness

Intel® Automatic Display Brightness feature dynamically adjusts the back-light 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 luminance), the new back-light setting can be adjusted through BLC (Back Light Control). The converse applies for a brightly lit environment. Intel® Automatic Display Brightness increases the back-light setting.

3.3.2.3 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.

3.3.2.4 Intel® Display Power Saving Technology (Intel® DPST) 6.3

The Intel® DPST technique achieves back-light power savings while maintaining a good visual experience. This is accomplished by adaptively enhancing the displayed image while decreasing the back-light brightness simultaneously. The goal of this technique is to provide equivalent end-user-perceived image quality at a decreased back-light 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 back-light 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 back-light brightness is applied simultaneously to produce an image with similar user-perceived quality (such as brightness) as the original image.

Intel® DPST 6.3 has improved power savings without adversely affecting the performance.

3.3.2.5 Panel Self-Refresh 2 (PSR 2)

Panel Self-Refresh feature allows the Processor Graphics core to enter low-power state when the frame buffer content is not changing constantly. This feature is available on panels capable of supporting Panel Self-Refresh. Apart from being able to support, the eDP* panel should be eDP 1.4 compliant. PSR 2 adds partial frame updates and requires an eDP 1.4 compliant panel.

3.3.2.6 Low-Power Single Display Pipe (LPSP)

Low-power single display pipe is a power conservation feature that helps save power by keeping the inactive display 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 display 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.

3.3.2.7 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 can be enabled in single pipe mode only.

Intel® S2DDT is the 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, 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.

3.3.3 Processor Graphics Core Power Savings Technologies

3.3.3.1 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 thermal budget 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 thermal. 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.

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

Intel® Graphics Render Standby Technology is a technique designed to optimize the average power of the graphics part. The Graphics Render engine will be put in a sleep state, or Render Standby (RS), during times of inactivity or basic video modes. While in Render Standby state, the graphics part will place the VR (Voltage Regulator) into a low voltage state. Hardware will save the render context to the allocated context buffer when entering RS state and restore the render context upon exiting RS state.

3.3.3.3 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.

3.4 System Agent Enhanced Intel SpeedStep® Technology

System Agent Enhanced Intel SpeedStep® Technology is a dynamic voltage frequency scaling of the System Agent clock based on memory utilization. Unlike processor core and package Enhanced Intel SpeedStep® Technology, System Agent Enhanced Intel SpeedStep® Technology has four valid operating points. When running light workload and SA Enhanced Intel SpeedStep® Technology is enabled, the DDR data rate may change as follows:

Before changing the DDR data rate, the processor sets DDR to self-refresh and changes the needed parameters. The DDR voltage remains stable and unchanged.

BIOS/MRC DDR training at maximum, mid and minimum frequencies sets I/O and timing parameters.

Refer [Table 5-13, "System Agent Enhanced Speed Steps \(SA-GV\) and Gear Mode Frequencies"](#).

3.5 Voltage Optimization

Voltage Optimization opportunistically provides reduction in power consumption, i.e., a boost in performance at a given PL1. Over time the benefit is reduced. There is no change to base frequency or turbo frequency. During system validation and tuning, this feature should be disabled to reflect processor power and performance that is expected over time.

This feature is available on selected SKUs.

3.6 ROP (Rest Of Platform) PMIC

In addition to discrete voltage regulators, Intel supports specific PMIC (Power Management Integrated Circuit) models to power the ROP rails. Intel supports ROP PMIC as part of UP3/UP4 - Processor Lines only.

3.7 PCI Express* Power Management

- Active power management support using L1 Substates (L1.1,L1.2)
- L0s power state is not supported on 11th Generation Intel® Core™ processor platform.
- All inputs and outputs disabled in L2/L3 Ready state.
- Processor PCIe* interface supports Hot-Plug.

Note: An increase in power consumption may be observed when PCI Express* ASPM capabilities are disabled.

Table 3-8. Package C-States with PCIe* Link States Dependencies

L-State	Description	Package C-State
L0s	L0s is not supported on 11th Generation Intel® Core™ Processor platform.	N/A
L0 or deeper	L0 is an active state.	C0-C3
L1.0 or deeper	L1- Higher latency, lower power state.	C3-C8
L1.2 or deeper	Lower power state.	C9-C10
L2	RTD3/Modern standby state.	C10

§ §

4 Thermal Management

4.1 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:

- Remains below the maximum junction temperature (T_{jMAX}) specification at the maximum thermal design power (TDP).
- 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.

4.1.1 Thermal Considerations

The processor TDP is the maximum sustained power that should be used for the design of the processor thermal solution. TDP is a power dissipation and junction 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 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.

The processor integrates multiple processing IA cores, graphics cores and for some SKUs a PCH on 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, 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/GTx) 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 is 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.

4.1.1.1 Package Power Control

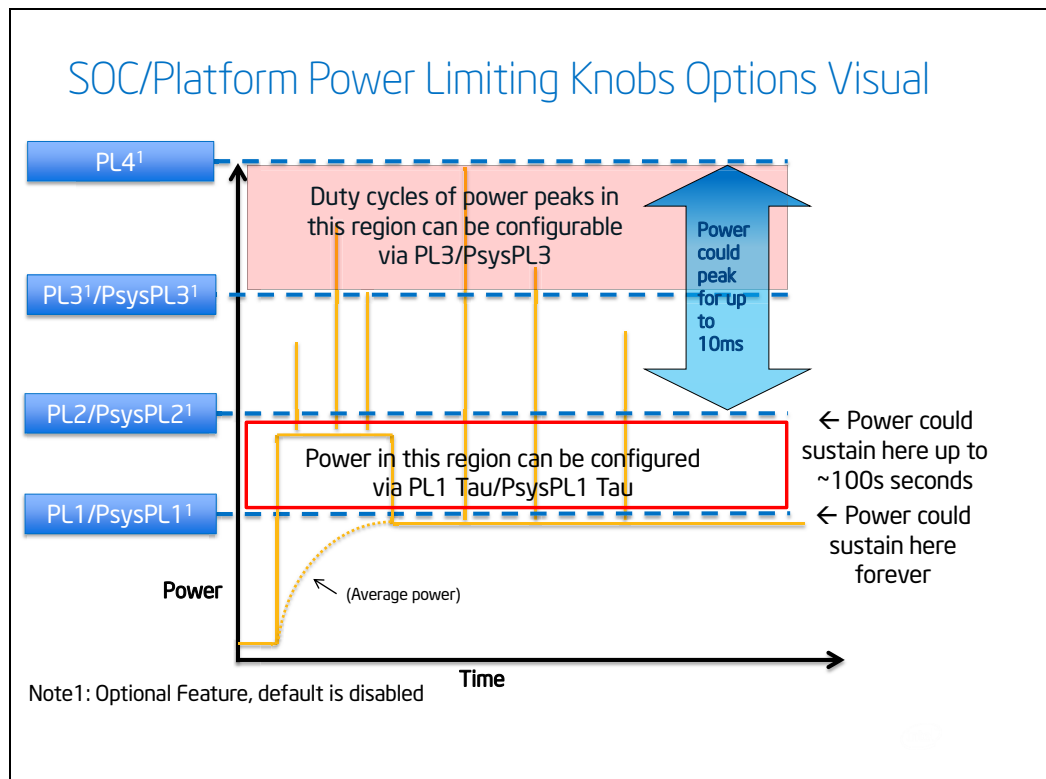
The package power control settings of PL1, PL2, PL3, PL4, and Tau allow the designer to configure Intel® Turbo Boost Technology 2.0 to match the platform power delivery and package thermal solution limitations.

- **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.

Notes:

1. Implementation of Intel® Turbo Boost Technology 2.0 only requires configuring PL1, PL1, Tau and PL2.
2. PL3 and PL4 are disabled by default.

Figure 4-1. Package Power Control



4.1.1.2 Platform Power Control

The processor introduces Psys (Platform Power) to enhance processor power management. The Psys signal needs to be sourced from a compatible charger circuit and routed to the IMVP9 (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 are analogous to the processor power limits described.

- **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.
- The Intel Dynamic Tuning (DTT/DPTF) is recommended for performance improvement in mobile platforms. Dynamic Tuning is configured by system manufacturers dynamically optimizing the processor power based on the currently platform thermal and power delivery conditions. Contact Intel Representatives for enabling details.

4.1.1.3 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.

4.1.2 Configurable TDP (cTDP) and Low-Power Mode

Configurable TDP (cTDP) and Low-Power Mode (LPM) form a design option where the processor's behavior and package TDP are dynamically adjusted to a desired system performance and power envelope. Configurable TDP and Low-Power Mode technologies offer opportunities to differentiate system design while running active workloads on select processor SKUs through scalability, configuration and adaptability. The scenarios or methods by which each technology is used are customizable but typically involve changes to PL1 and associated frequencies for the scenario with a resultant change in performance depending on system's usage. Either technology can be triggered by (but are not limited to) changes in OS power policies or hardware events such as docking a system, flipping a switch or pressing a button. cTDP and LPM are designed to be configured dynamically and do not require an operating system reboot.

Note: Configurable TDP and Low-Power Mode technologies are not battery life improvement technologies.

4.1.2.1 Configurable TDP

Note: Configurable TDP availability may vary between the different SKUs.

With cTDP (Configurable TDP), the processor is now capable of altering the maximum sustained power with an alternate processor IA core base frequency. Configurable TDP allows operation in situations where extra cooling is available or situations where a cooler and quieter mode of operation is desired. The requirements for developing a non-driver approach can be found by referencing the appropriate processor Configurable TDP and LPM Implementation Guide (refer).

cTDP consists of three modes as shown in the following table:

Table 4-1. Configurable TDP Modes

Mode	Description
Base	The average power dissipation and junction temperature operating condition limit, specified in Table 4-2, "TDP Specifications" and Table 4-4, "Junction Temperature Specifications" 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.
TDP-Up	The SKU-specific processor IA core frequency where manufacturing confirms logical functionality within the set of operating condition limits specified for the SKU segment and Configurable TDP-Up configuration in Table 4-2, "TDP Specifications" and Table 4-4, "Junction Temperature Specifications" . The Configurable TDP-Up Frequency and corresponding TDP is higher than the processor IA core Base Frequency and SKU Segment Base TDP.
TDP-Down	The processor IA core frequency where manufacturing confirms logical functionality within the set of operating condition limits specified for the SKU segment and Configurable TDP-Down configuration in Table 4-2, "TDP Specifications" and Table 4-4, "Junction Temperature Specifications" . The Configurable TDP-Down Frequency and corresponding TDP is lower than the processor IA core Base Frequency and SKU Segment Base TDP.

In each mode, the Intel® Turbo Boost Technology 2.0 power limits are reprogrammed along with a new OS controlled frequency range. The Intel Dynamic Tuning driver assists in TDP operation by adjusting processor PL1 dynamically. The cTDP mode does not change the maximum per-processor IA core turbo frequency.

4.1.2.2 Low-Power Mode

Low-Power Mode (LPM) can provide cooler and quieter system operation. By combining several active power limiting techniques, the processor can consume less power while running at equivalent low frequencies. Active power is defined as processor power consumed while a workload is running and does not refer to the power consumed during idle modes of operation. LPM is only available using the Intel Dynamic tuning (DTT/DPTF) driver.

Through the Dynamic tuning (DTT/DPTF) driver, LPM can be configured to use each of the following methods to reduce active power:

- Restricting package power control limits and Intel® Turbo Boost Technology availability
- Off-Lining processor IA core activity (Move processor traffic to a subset of cores)
- Placing a processor IA Core at LFM or LSF (Lowest Supported Frequency)
- Utilizing IA clock modulation
- LPM power as listed in the *TDP Specifications* table is defined at a point which processor IA core working at LSF, GT = Rpn and 1 IA core active

Off-lining processor IA core activity is the ability to dynamically scale a workload to a limited subset of cores in conjunction with a lower turbo power limit. It is one of the main vectors available to reduce active power. However, not all processor activity is ensured to be able to shift to a subset of cores. Shifting a workload to a limited subset of cores allows other processor IA cores to remain idle and save power. Therefore, when LPM is enabled, less power is consumed at equivalent frequencies.

Minimum Frequency Mode (MFM) of operation, which is the Lowest Supported Frequency (LSF) at the LFM voltage, has been made available for use under LPM for further reduction in active power beyond LFM capability to enable cooler and quieter modes of operation.

4.1.3 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.

4.1.3.1 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 T_{jMAX} .

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.

T_{jMAX} 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.

4.1.3.1.1 TCC Activation Offset

TCC Activation Offset can be set as an offset from T_{jMAX} to lower the onset of TCC and Adaptive Thermal Monitor. In addition, there is an optional time window (Tau) to manage processor performance at the TCC Activation offset value via an EWMA (Exponential Weighted Moving Average) of temperature.

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 T_{jMAX} value and used as a new maximum 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 T_j 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 T_{jMAX} thermal management. That is, whether the TCC activation offset is 0 or not, TCC Activation will occur at T_{jMAX} .

4.1.3.1.2 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 the 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.

4.1.3.1.3 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.

4.1.3.1.4 Thermal Throttling

As the processor approaches $T_{J_{Max}}$ a throttling mechanisms will engage to protect the processor from over-heating and provide control thermal budgets.

Achieving this is done by reducing IA and other subsystem agent's voltages and frequencies in a gradual and coordinated manner that varies depending on the dynamics of the situation. IA frequencies and voltages will be directed down as low as LFM (Lowest Frequency Mode). Further restricts are possible via Thermal Trolling point (TT1) under conditions where thermal budget cannot be re-gained fast enough with voltages and frequencies reduction alone. TT1 keeps the same processor voltage and clock frequencies the same yet skips clock edges to produce effectively slower clocking rates. This will effectively result in observed frequencies below LFM on the Windows PERF monitor.

4.1.3.2 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.

When the temperature is retrieved by the processor MSR, it is the instantaneous temperature of the given DTS. When the 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 (0x1B1) MSR and IA32_THERM_STATUS (0x19C) MSR.

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 ($T_{j_{MAX}}$), 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 (0x1A2) MSR. The temperature returned by the DTS is an implied negative integer indicating the relative offset from $T_{j_{MAX}}$. The DTS does not report temperatures greater than $T_{j_{MAX}}$. 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.

4.1.3.2.1 Digital Thermal Sensor Accuracy ($T_{accuracy}$)

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

4.1.3.2.2 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 $T_{j_{MAX}}$.

4.1.3.3 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

4.1.3.3.1 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 10Khz.

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 [Section 4.1.3.6, "PROCHOT Demotion Algorithm"](#)).

4.1.3.4 PROCHOT Output Only

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

4.1.3.5 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 de-asserts PROCHOT#. The processor can be configured to generate an interrupt upon assertion and de-assertion of the PROCHOT# signal.

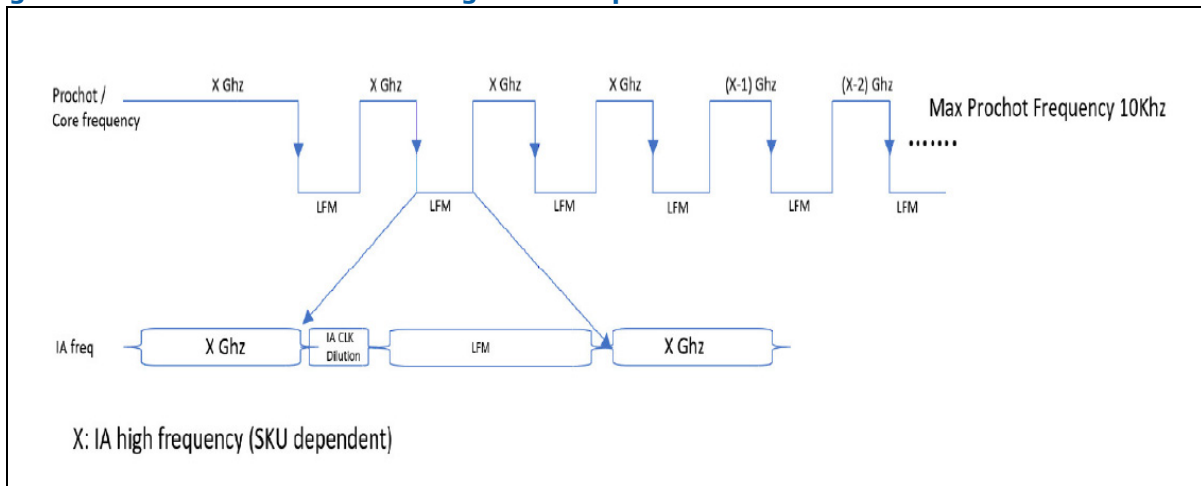
When PROCHOT# is configured as a bi-directional signal and PROCHOT# is asserted by the processor, it is impossible for the processor to detect a system assertion of PROCHOT#. The system assertion will have to wait until the processor de-asserts PROCHOT# before PROCHOT# action can occur due to the system 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.

4.1.3.6 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.

Figure 4-2. PROCHOT Demotion Signal Description



4.1.3.7 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.

4.1.3.8 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).

4.1.3.9 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 PECCI 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 PECCI.

PECCI can be implemented using either the single bit bidirectional I/O pin or using the eSPI interface

4.1.3.10 THRMTRIP# 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 THRMTRIP# signal will go active.

4.1.3.11 Critical Temperature Detection

Critical Temperature detection is performed by monitoring the package temperature. This feature is intended for graceful shutdown before the THRMTRIP# is activated. However, the processor execution is not guaranteed between critical temperature and THRMTRIP#. 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 (0x1B1) MSR and the condition also generates a thermal interrupt, if enabled.

4.1.3.12 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 the 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.

4.1.3.13 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.

4.1.3.14 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.

4.1.4 Intel® Memory Thermal Management

DRAM Thermal Aggregation

Punit firmware is responsible for aggregating DRAM temperature sources into a per-DIMM reading as well as an aggregated virtual 'max' sensor reading. At reset, MRC communicates to the MC the valid channels and ranks as well as DRAM type. At that time, Punit firmware sets up a valid channel and rank mask that is then used in the thermal aggregation algorithm to produce a single maximum temperature

DRAM Thermal Monitoring

- DRAM thermal sensing Periodic DDR thermal reads from DDR
- DRAM thermal calculation Punit reads of DDR thermal information direct from the memory controller (MR4 or MPR) Punit estimation of a virtual maximum DRAM temperature based on per-rank readings. Application of thermal filter to the virtual maximum temperature.
- Thermal Monitoring Collection of inputs from software for thermal interrupt management. Calculation of thermal threshold status and log bits as well as interrupt delivery.
- DRAM thermal reporting Writing out of filtered, virtual maximum DDR temperature to software visible registers.

DRAM Refresh Rate Control

- The new interface provides a direct override to refresh rate management, regardless of MR4 or MPR readings. From both in-band software and PECCI/ESPI.

DRAM Bandwidth Throttling

- Software and PECCI/ESPI control for throttling the memory controller and DRAM maximum bandwidth.

4.2 Thermal and Power Specifications

The following notes apply to the tables below, [Table 4-2, "TDP Specifications"](#), [Table 4-3, "Package Turbo Specifications"](#), and [Table 4-4, "Junction Temperature Specifications"](#)

Note	Note Definition
1	The TDP and Configurable TDP values are 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 4.1.1.2, "Platform Power Control" for further information.
5	The 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	The Processor will be controlled to a specified power limit as described in Section 2.4.6.1, "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.
8	For controllable turbo workloads, the PL2 limit may be exceeded for up to 10 ms.
9	LPM power level is an opportunistic power and is not a guaranteed value as usages and implementations may vary.
10	Power limits may vary depending on if the product supports the 'TDP-up' and/or 'TDP-down' modes. Default power limits can be found in the PKG_PWR_SKU MSR (614h).
11	The processor die do not reach maximum sustained power simultaneously since the sum of the 2 die's estimated power budget is controlled to be equal to or less than the package TDP (PL1) limit.
12	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.
13	May vary based on SKU.
14	<ul style="list-style-type: none"> The formula of $PL2=PL1*1.25$ is the hardware default. PL2- SoC opportunistic higher Average Power with limited duration controlled by Tau_PL1 setting, the larger the Tau, the longer the PL2 duration.
16	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=28 s.

Table 4-2. TDP Specifications (Sheet 1 of 3)

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 table above)
H-Processor Line	8- Core GT1 45W	Configurable TDP-Up	3.3GHz	1.45GHz	65 ²	1,9,10,15
		Base	2.6GHz		45 ¹	
		LFM	0.8GHz	0.35GHz	N/A	

Table 4-2. TDP Specifications (Sheet 2 of 3)

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 table above)
H-Processor Line	8- Core GT1 45W	Base	2.3GHz up to 2.6GHz	1.45GHz	45 ¹	1,9,10,15
		Configurable TDP-Down	1.9GHz up to 2.1GHz		35 ²	
		LFM	0.8GHz	0.35GHz	N/A	
H-Processor Line	6- Core GT1 45W	Base	2.6GHz up to 3.2GHz	1.4GHz up to 1.45GHz	45 ¹	1,9,10,15
		Configurable TDP-Down	2.1GHz up to 2.6GHz		35 ²	
		LFM	0.8GHz	0.35GHz	N/A	
UP3-Processor Line	4- Core GT2 28W	Base	2.4GHz to 3.0GHz	1.3GHz up to 1.35GHz	28	1,9,10,15
		Configurable TDP-Down 1	2.4GHz up to 3.0GHz		15	
		Configurable TDP-Down 2	0.9GHz up to 1.2GHz		12	
		LFM	0.4GHz		0.1GHz	
UP3-Processor Line	2- Core GT2 28W	Base	3.0GHz	1.25GHz	28	1,9,10,15
		Configurable TDP-Down 1	2.2GHz		15	
		Configurable TDP-Down 2	1.7GHz		12	
		LFM	0.4GHz		0.1GHz	
UP3-Pentium/Celeron Processor Line	2- Core GT2 15W	Base	1.8GHz-2.0GHz	1.25GHz	15	1,9,10,15
		LFM	0.4GHz		N/A	
UP4-Processor Line	4- Core GT2 9W	Base	1.1GHz to 1.2GHz	1.1GHz	9	1,9,10,11,15
		Configurable TDP-Up	1.5GHz up to 2.1GHz		15	
		Configurable TDP-Down	0.8GHz up to 0.9GHz		7	
		LFM	0.4GHz		0.1GHz	
UP4-Processor Line	2- Core GT2 9W	Base	1.8GHz	1.1GHz	9	1,9,10,11,15
		Configurable TDP-Up	2.5GHz		15	
		Configurable TDP-Down	1.5GHz		7	
		LFM	0.4GHz		0.1GHz	
H35-Processor Line	4- Core GT2 35W	Base	3.1GHz up to 3.3GHz	1.35GHz	35 ¹	1,9,10,15
		Configurable TDP-Down	2.6GHz up to 3.0GHz		28 ¹	
		LFM	0.4GHz	0.1GHz	N/A	



Table 4-2. TDP Specifications (Sheet 3 of 3)

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 table above)
H35-Refresh Processor Line	4- Core GT2 35W	Base	3.4GHz	1.4GHz	35	1,9,10,11,15
		Configurable TDP-Down	2.9GHz		28	
		LFM	0.4GHz	0.1GHz	N/A	
UP3-Refresh Processor Line	4- Core GT2 28W	Base	2.9 GHz	1.4GHz	28	1,9,10,15
		Configurable TDP-Down_1	1.8 GHz		15	
		Configurable TDP-Down_2	1.3 GHz		12	
		LFM	0.4GHz	0.1GHz	N/A	

Table 4-3. Package Turbo Specifications (Sheet 1 of 2)

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Parameter	Minimum	Hardware Default	MSR Max Value	Recommendation Value	Units	Notes
H-Processor Line	8- Core GT1 45W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	56	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	45	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	
H-Processor Line	6 - Core GT1 45W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	28	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	45	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	
UP3/ UP3-Refresh - Processor Line	4/2- Core GT2 28W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	28	S	3,4,5,6,7,8,14,16
		Power Limit 1 (PL1)	N/A	28	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	
UP3-Pentium/ Celeron Processor Line	2- Core GT2 15W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	28	S	3,4,5,6,7,8,14,16
		Power Limit 1 (PL1)	N/A	15	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	

Table 4-3. Package Turbo Specifications (Sheet 2 of 2)

Segment and Package	Processor IA Cores, Graphics Configuration and TDP	Parameter	Minimum	Hardware Default	MSR Max Value	Recommendation Value	Units	Notes
UP4-Processor Line	4/2- Core GT2 9W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	28	S	3,4,5,6,7,8,14,16
		Power Limit 1 (PL1)	N/A	9	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	
H35/ H35 Refresh - Processor Line	4- Core GT2 35W	Power Limit 1 Time (PL1 Tau)	0.01	1	448	28	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	35	N/A	N/A	W	
		Power Limit 2 (PL2)	N/A	PL1*1.25	N/A	N/A	W	

Notes:

- For the notes, refer to the 1st page of [Chapter 4, "Thermal Management"](#).
- 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=28 s.

Table 4-4. Junction Temperature Specifications

Segment	Symbol	Package Turbo Parameter	Temperature Range		TDP Specification Temperature Range		Units	Notes
			Minimum	Maximum	Minimum	Maximum		
H-Processor Line BGA	T_j	Junction temperature limit	0	100	0	100	°C	1, 2
UP3/UP3-Refresh// H35-H35 Refresh Processor Line BGA	T_j	Junction temperature limit	0	100	35	100	°C	1, 2
UP4-Processor Line BGA	T_j	Junction temperature limit	0	100	0	90	°C	1, 2, 3

Notes:

- The thermal solution needs to ensure that the processor temperature does not exceed the TDP Specification Temperature.
- The processor junction temperature is monitored by Digital Temperature Sensors (DTS). For DTS accuracy, refer [Section 4.1.3.2.1, "Digital Thermal Sensor Accuracy \(T_accuracy\)"](#).
- UP4 specification need to be compliance with the 90°C TDP specification temperature, TCC Offset = 10 and Tau value should be programed into MSR 1A2h. The recommended TCC_Offset averaging Tau value is 5 s.

§ §

5 Memory

5.1 System Memory Interface

5.1.1 Processor SKU Support Matrix

Table 5-1. DDR Support Matrix Table

Technology	DDR4 ⁵	LPDDR4x ⁴
Processor	UP3/H ^{9,10}	UP3/UP4
Maximum Frequency [MT/s]	3200	4266
VDDQ [V] ⁵	1.2	0.6
VDD2 [V] ^{5,7}	1.2	1.1
DPC ^{1,2}	1, (2 is for H only)	-
RPC	1,2	1,2
Die Density [Gb]	8, 16	8, 16
Ballmap Mode	IL ³ /NIL	NIL
ECC	H ⁸ only	N/A

Notes:

- 1DPC refer to when only 1DIMM 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. 2DPC supported in H segment only.
- An Interleave SoDIMM/MD placements like butterfly or back-to-back supported with a Non-Interleave Ballout mode at UP3.
- LPDDR4x technology should be implemented homogeneous means that all DRAM devices in the system should be from the same vendor and have the same part number. Implementing a mix of DRAM devices may cause serious signal integrity and functional issues.
- DDR4 supports asymmetric channel memory capacity. For best IA and GFx performance it is recommended to use symmetric channels capacity.
- 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.
- VDD2 is Processor and DRAM voltage, and VDDQ is DRAM voltage.
- H-processor supports ECC with Interleave Ballout only. ECC with non-interleave is not supported.
- H-processor DDR4 DIMM0 (Rank[1:0]) must be populated in case DIMM1 (Rank[3:2]) is populated. Processor will not boot in case DIMM1 (Rank[3:2]) is populated and DIMM0 (Rank[1:0]) is not populated.
- Some SKUs may configured to run up to 2933 MT/s

Table 5-2. Processor DDR Speed Support

Processor Line	DDR4 Memory Down [MT/s]	DDR4 1DPC [MT/s]	DDR4 2DPC [MT/s]	LPDDR4x [MT/s]
H-Processor Line ¹	3200	3200	3200	N/A
UP3-Processor Line	3200	3200	N/A	4266
UP4-Processor Line	N/A	N/A	N/A	4266
UP3-Refresh Processor Line	3200	3200	N/A	4266

Note:

1. H-Processor DDR4 2DPC is supported when the channel is populated with the same SoDIMM part number.

Table 5-3. DDR Technology Support Matrix

Technology	Form Factor	Ball count	Processor
DDR4	SoDIMM	260	UP3/H
DDR4	x8 SDP (1R)	78	UP3/H
DDR4	x16 SDP (1R)	96	UP3/H
DDR4	x16 DDP (1R)	96	UP3/H
LPDDR4x	x32 (1R, 2R)	200	UP4 ¹ /UP3
LPDDR4x	x64 (1R, 2R)	432	UP4/UP3
LPDDR4x	x64 (1R, 2R)	556	UP4

Notes:
1. Non POR configuration only

5.1.2 Supported Population

Table 5-4. LPDDR4x Channels Population Rules

Configuration			Memory Controller 0				Memory Controller 1			
			CH0	CH1	CH2	CH3	CH4	CH5	CH6	CH7
# of Chs	PKG Width	# of DQs	x16	x16	x16	x16	x16	x16	x16	x16
8	64	128	x64				x64			
8	32	128	x32		x32		x32		x32	
4	64	64	x64							
4	32	64	x32		x32					

Note: White blank means not populated



Table 5-6. H DDR4 SoDIMM Population Configuration

Configuration	Memory Controller 0		Memory Controller 1	
	Channel 0		Channel 1	
	DIMM 0	DIMM 1	DIMM 0	DIMM 1
	Rank[1:0]	Rank[3:2]	Rank[1:0]	Rank[3:2]
2 DIMM per channel	x64	x64	x64	x64
	X	X	X	X
	X	X	X	
	X		X	X
	X	X		
			X	X
	X		X	
	X			
1 DIMM per channel	X			
			X	
	X		X	

Notes:

1. X means SoDIMM populated, White blank means SoDIMM not populated, Gray blank means no slot
2. DDR4 DIMM0(Rank[1:0]) must be populated as a default DIMM, DDR4 DIMM1(Rank[3:2]) is optional. Populating Slot1(Ranks[3:2]) when Slot0(Ranks[1:0]) is not populate may cause system not to boot.
3. 2DPC required populated DIMM have same part number within each channel.
4. DDR4 Memory Down Channel must use Rank[0]

5.1.3 Supported Memory Modules and Devices

Table 5-7. Supported DDR4 SoDIMM 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	ECC
A	8 GB	8 Gb	1024M x 8	8	1	16/10	16	8 K	No
A	16 GB	16 Gb	2048M x 8	8	1	17/10	16	8 K	No
C	4 GB	8 Gb	512M x 16	4	1	16/10	8	8 K	No
C	8 GB	16 Gb	1024M x 16	4	1	17/10	8	8 K	No
E	16 GB	8 Gb	1024M x 8	16	2	16/10	16	8 K	No
E	32 GB	16 Gb	2048M x 8	16	2	17/10	16	8 K	No
D	8 GB	8 Gb	1024M x 8	9	1	16/10	16	8 K	Yes
D	16 GB	16 Gb	2048M x 8	9	1	17/10	16	8 K	Yes
G	16 GB	8 Gb	1024M x 8	18	2	16/10	16	8 K	Yes
G	32 GB	16 Gb	2048M x 8	18	2	17/10	16	8 K	Yes
H	16 GB	8 Gb	1024M x 8	18	2	16/10	16	8 K	Yes
H	32 GB	16 Gb	2048M x 8	18	2	17/10	16	8 K	Yes

Table 5-8. Supported DDR4 Memory Down Device Configurations

Maximum System Capacity ³	PKG Type (Die bits x Package bits)	DRAM Organization / Package Type	Package Density	Die Density	Dies Per Channel	Rank Per Channel	PKGs Per channel	Physical Device Rank	Banks Inside DRAM	Page Size
16 GB	SDP 8x8	1024M x 8	8 Gb	8 Gb	8	1	8	1	16	8 K
32 GB	SDP 8x8	2048M x 8	16 Gb	16 Gb	8	1	8	1	16	8 K
8 GB	SDP 16x16	512M x 16	8 Gb	8 Gb	4	1	4	1	8	8 K
16 GB ¹	SDP 16x16	1024M x 16	16 Gb	16 Gb	4	1	4	1	8	8 K
16 GB	DDP 8x16	1024M x 16	16 Gb	8 Gb	8	1	4	1	16	8 K
32 GB ^{2,3}	DDP 8x16	2048M x 16	32 Gb	16 Gb	8	1	4	1	16	8 K

Notes:

- For SDP: 1Rx16 using 16Gb die density - the maximum system capacity is 16 GB
- For DDP: 1Rx16 using 16Gb die density - the maximum system capacity is 32 GB.
- Pending DRAM samples availability.
- Maximum system capacity refer to system with 2 channels populated

Table 5-9. Supported LPDDR4x x32 DRAMs Configurations (Sheet 1 of 2)

Maximum System Capacity ⁴	PKG Type	(Die bits per Ch x PKG bits) ²	Die Density	PKG Density	Rank Per PKGs
8 GB	DDP	16x32	8 Gb	16 Gb	1
16 GB	QDP	16x32	8 Gb	32 Gb	2
32 GB	ODP	8x32(Byte mode)	8 Gb	64 Gb	2
16 GB	DDP	16x32	16 Gb	32 Gb	1
32 GB	QDP	16x32	16 Gb	64 Gb	2
64 GB ⁵	ODP	8x32(Byte mode)	16 Gb	128 Gb	2

Table 5-9. Supported LPDDR4x x32 DRAMs Configurations (Sheet 2 of 2)

Maximum System Capacity ⁴	PKG Type	(Die bits per Ch x PKG bits) ²	Die Density	PKG Density	Rank Per PKGs
Notes: 1. x32 BGA devices are 200 balls 2. QDP = Quad Die Package, ODP-Octal Die Package 3. Each LPDDR4x channel include two sub-channels 4. Maximum system capacity refers to system with all 8 sub-channels populated 5. Pending DRAM samples availability.					

Table 5-10. Supported LPDDR4x x64 DRAMs Configurations

Maximum System Capacity ⁴	PKG Type	(Die bits per Ch x PKG bits) ²	Die Density	Ball Count Per PKG	PKG Density	DRAM Channels Per PKGs	Processor Line	Rank Per PKGs
8 GB	QDP	16x64	8 Gb	432	32 Gb	4	UP3/UP4	1
16 GB	ODP	16x64	8 Gb	432	64 Gb	4	UP3/UP4	2
16 GB	QDP	16x64	16 Gb	432	64 Gb	4	UP3/UP4	1
8 GB	QDP	16x64	8 Gb	556	32 Gb	4	UP4	1
16 GB	ODP	16x64	8 Gb	556	64 Gb	4	UP4	2
16 GB ¹	QDP	16x64	16 Gb	556	64 Gb	4	UP4	1
Notes: 1. Maximum system capacity refers to system with all 8 sub-channels populated 2. QDP = Quad Die Package, ODP-Octal Die Package 3. Each LPDDR4x channel include two sub-channels								

5.1.4 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
- tRPb = per-bank PRECHARGE time
- tRPab = all-bank PRECHARGE time
- CWL = CAS Write Latency
- Command Signal modes:
 - 2N indicates a new DDR4/LPDDR4x command may be issued every 2 clocks
 - 1N indicates a new DDR4/LPDDR4x command may be issued every clock.

5.1.5 System Memory Timing Support

Table 5-11. DDR4 System Memory Timing Support

DRAM Device	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRP (ns)	CWL (tCK)	DPC	CMD Mode
DDR4	3200	22	13.75	13.75	9-12, 14,16,18,20	2	2N

Table 5-12. LPDDR4x System Memory Timing Support

DRAM Device	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRPpb (ns)	tRPab (ns)	WL (tCK) Set B
LPDDR4x	4266	36	18	18	21	34

5.1.6 SAGV Points

SAGV (System Agent Geyserville) is a way by which they SoC can dynamically scale the work point (V/F), by applying DVFS (Dynamic Voltage Frequency Scaling) based on memory bandwidth utilization and/or the latency requirement of the various workloads for better energy efficiency at System-Agent. Pcode heuristics are in charge of providing request for Qclock work points by periodically evaluating the utilization of the memory and IA stalls.

11th Generation Intel® Core™ processor adds support for a 4th point for SAGV. A 4th GV point allows Pcode to select a more optimal frequency so that SA and Qclk region are operating at a lower voltage/frequency but still providing the required BW.

Table 5-13. System Agent Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies

Processor Line	Technology	DDR Maximum Rate [MT/s]	SAGV-LowBW	SAGV-MedBW	SAGV-HighBW	SAGV- High Performance
H	DDR4	3200	2133, G2	2133, G2	3200, G2	2933, G1
UP3	DDR4	3200	2133, G2	2133, G2	3200, G2	2666, G1
UP3	LPDDR4x	4266	2133, G2	3200, G2	4266, G2	2666, G1
UP4	LPDDR4x	4266	2133, G2	3200, G2	4266, G2	4266, G2

Notes:

- 11th Generation Intel® Core™ Processor supports dynamic gearing technology where the Memory Controller can run at 1:1 (Gear-1, Legacy mode) or 1:2 (Gear-2 mode) ratio of DRAM speed. The gear ratio is the ratio of DRAM speed to Memory Controller Clock.
MC Channel Width equal to DDR Channel width multiply by Gear Ratio
- SA-GV operating points
 - LowBW**- Low frequency point, Minimum Power point. Characterized by low power, low BW, high latency. The system will stay at this point during low to moderate BW consumption.
 - MedBW** - this point is tuned for balance between power & performance (BW demand) Characterized by moderate power and latency, moderate BW. Only during IA performance workloads, the system will to switch to this point and only in case this point can provide enough BW
 - HighBW** Maximum Bandwidths Point, minimum memory latency point, Characterized by high power, low latency and high BW. This point intended for high GT and moderate-high IA BW
 - High Performance** - Lowest Latency point, low BW and highest power
- Refer to [Section 3.4, "System Agent Enhanced Intel® SpeedStep® Technology"](#) for more details

5.1.7 Memory Controller (MC)

The integrated memory controller is responsible of transferring data between the processor and the DRAM as well as the DRAM maintenance. In this, there are two instances of MC, one per memory slice. Each controller is capable of supporting up to four channels of LPDDR4x and one channel of DDR4.

The two controllers are independent and have no means of communicating with each other, so they need to be configured separately.

In a symmetric memory population, each controller only view half of the total physical memory address space.

Both MC support only one technology in a system DDR4, LPDDR4x mix of technologies in one system is not allowed.

5.1.8 System Memory Controller Organization Mode (DDR4)

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

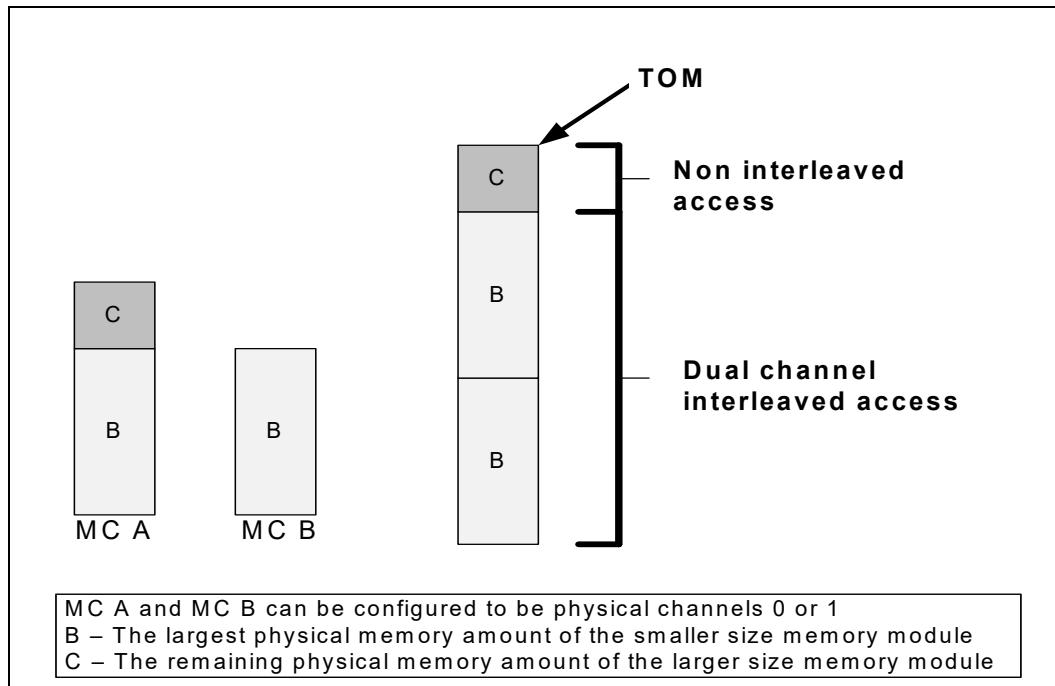
In this mode, all memory cycles are directed to a single channel. Single-Channel mode is used when either the Channel A or Channel B DIMM connectors are populated in any order, but not both.

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 5-1. Intel® DDR4 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.

Notes:

1. The DRAM device technology and width may vary from one channel to another. Different memory size between channels are relevant to DDR4 only.

5.1.9 System Memory Frequency

In all modes, the frequency of system memory is the lowest frequency of all memory modules placed in the system, as determined through the SPD registers on the memory modules. The system memory controller supports a single DIMM connector per channel. If DIMMs with different latency are populated across the channels, the BIOS will use the slower of the two latencies for both channels. For Dual-Channel modes, both channels should have a DIMM connector populated. For Single-Channel mode, only a single channel can have a DIMM connector populated.

5.1.10 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.

5.1.11 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 1 s and 0 s 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.

5.1.12 ECC DDR4 H-Matrix Syndrome Codes

Table 5-14. ECC H-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
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. This table is relevant only for H-Processor ECC supported SKUs.							

5.1.13 Data Swapping

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

- Bit swapping is allowed within each Byte for all DDR technologies.
- **LPDDR4x**: Byte swapping is allowed within each x16 sub channel.
- **LPDDR4x**: Upper/Lower four x16 sub channels to be connected to x64 DRAM or two x32 DRAMs. Swapping between four upper to four lower x16 sub channels is not allowed.
- **DDR4**: Byte swapping is allowed within each x64 channel.
- ECC bits swap is allowed within DDR4 ECC[7...0].

5.1.14 DDR I/O Interleaving

Note: 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. UP4/ UP3-Processor line packages are optimized only for Non-Interleaving mode (NIL).

There are two supported modes:

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

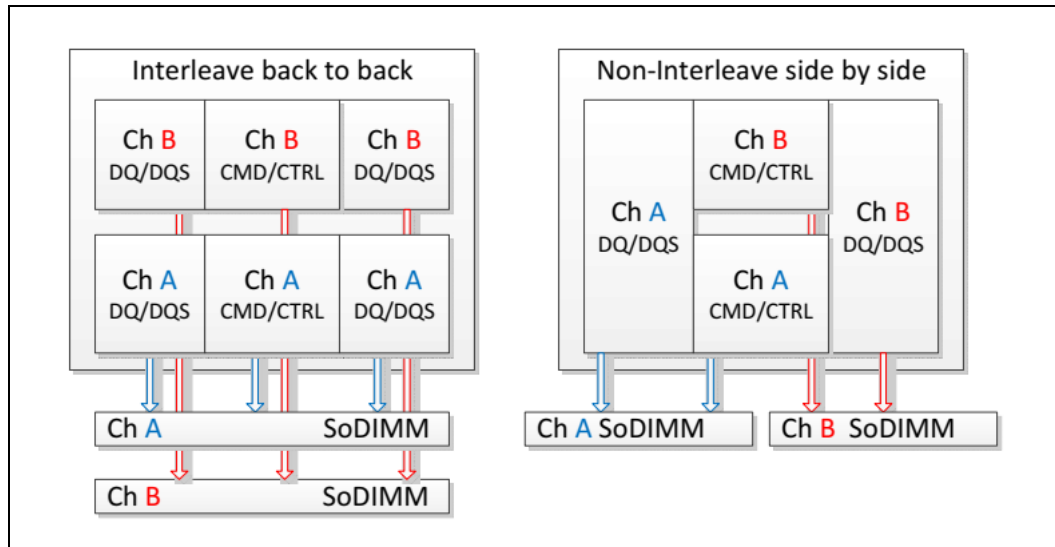
The following table and figure describe the pin mapping between the IL and NIL modes.

Table 5-15. Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping

IL (DDR4)		NIL (DDR4)		NIL (LPDDR4x)	
Channel	Byte	Channel	Byte	Sub Channel	Byte
DDR0	Byte0	DDR0	Byte0	DDR7	Byte1
DDR0	Byte1	DDR0	Byte2	DDR6	Byte0
DDR0	Byte2	DDR0	Byte4	DDR7	Byte1
DDR0	Byte3	DDR0	Byte6	DDR6	Byte0
DDR0	Byte4	DDR1	Byte0	DDR5	Byte1
DDR0	Byte5	DDR1	Byte2	DDR4	Byte0
DDR0	Byte6	DDR1	Byte4	DDR5	Byte1
DDR0	Byte7	DDR1	Byte6	DDR4	Byte0
DDR1	Byte0	DDR0	Byte1	DDR3	Byte1
DDR1	Byte1	DDR0	Byte3	DDR2	Byte0
DDR1	Byte2	DDR0	Byte5	DDR3	Byte1
DDR1	Byte3	DDR0	Byte7	DDR2	Byte0
DDR1	Byte4	DDR1	Byte1	DDR1	Byte1
DDR1	Byte5	DDR1	Byte3	DDR0	Byte0
DDR1	Byte6	DDR1	Byte5	DDR1	Byte1
DDR1	Byte7	DDR1	Byte7	DDR0	Byte0

Notes: UP4-processor line supports NIL only.

Figure 5-2. DDR4 Interleave (IL) and Non-Interleave (NIL) Modes Mapping



5.1.15 DRAM Clock Generation

Each support rank has a differential clock pair for DDR4. Each sub-channel has a differential clock pair for LPDDR4x.

5.1.16 DRAM Reference Voltage Generation

Read Vref is generated by the memory controller in all technologies. Write Vref is generated by the DRAM in all technologies. Command Vref is generated by the DRAM in LPDDR4x while the memory controller generates VrefCA per DIMM for DDR4. In all cases, it has small step sizes and is trained by MRC.

5.1.17 Data Swizzling

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

5.2 Integrated Memory Controller (IMC) Power Management

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

5.2.1 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 seen by the processor I/O buffer receivers caused by reflections from potentially unterminated 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.

5.2.2 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 three different types of power-down modes in package C0 state. The different power-down modes can be enabled through configuring PM PDWN config 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 – a small number of cycles.
- **Pre-charged Power-down (PPD):** This mode is entered if all banks in DDR are pre-charged when de-asserting CKE. Power-saving in this mode is intermediate – better than APD. 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 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 arrive 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 a 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
- 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 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.

5.2.2.1 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 microseconds after power and clocks to SDRAM devices are stable.

5.2.2.2 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 [Section 3.3.1.1, “Intel® Rapid Memory Power Management \(Intel® RMPM\)”](#) for more details on conditional self-refresh with Intel HD Graphics enabled.

When entering the 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.

5.2.2.3 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 pre-charge power-down (CKE de-assertion with all pages closed). Pre-charge 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 the refresh.

5.2.2.4 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. 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).

5.2.3 DDR Electrical Power Gating

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 and VDD2 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 VccSA for all non-critical state to reduce idle power.

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

5.2.4 Power Training

BIOS MRC performing Power Training steps to reduce DDR I/O power while keeping reasonable operational margins still guaranteeing 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 operating margins using advanced mathematical models.



6 USB-C* Sub System

USB-C* is a cable and connector specification defined by USB-IF.

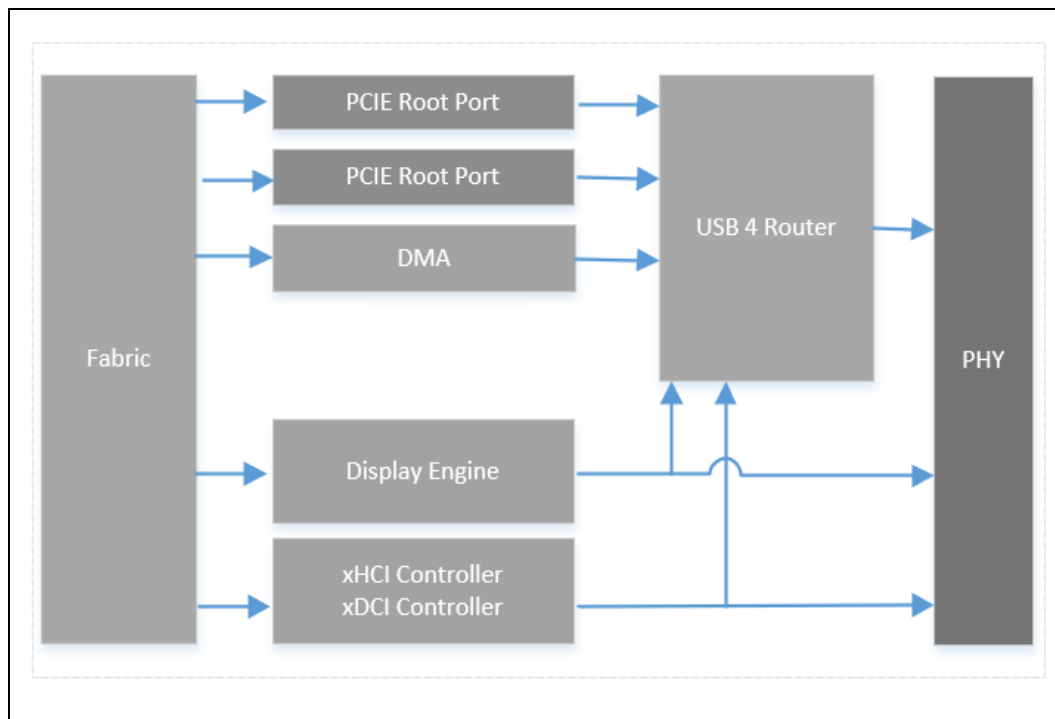
The USB-C sub-system supports all processor lines DPoC (DisplayPort over Type-C) protocols. The USB-C sub-system can also support be configured as native DisplayPort or HDMI v2.0b interfaces, for more information refer to [Chapter 10, "Display"](#).

Thunderbolt™ 4 is a USB-C solution brand which requires the following elements:

- USB2, USB3 (10 Gbps), USB3/DP implemented at the connector.
- In additional, it requires USB4 implemented up to 40 Gbps, including Thunderbolt™ 3 compatibility as defined by USB4/USB-PD specs and 15 W of bus power
- Thunderbolt™ 4 solutions use (and prioritize) the USB4 PD entry mode (while still supporting Thunderbolt™ 3 alt mode)
- This product has the ability to support these requirements

Note: If USB4 (20 Gbps) only solutions are implemented, Thunderbolt™ 3 compatibility as defined by USB4/USB-PD specs and 15 W of bus power are still recommended

Figure 6-1. USB-C* Sub-system Block Diagram



Note: USB-C sub-system support 2x USB 4 routers, each router can support up to two Type-C ports.

6.1 USB-C Sub-System General Capabilities

- xHCI (USB 3 host controller) and xDCI (USB 3 device controller) implemented in the processor in addition to the controllers in the PCH.
- No support for USB Type-A on the processor side.
- Intel® AMT/vPro over Thunderbolt™ docking.
- Support power saving when USB-C* disconnected.
- Support up to four simultaneous ports.
- DbC Enhancement for Low Power Debug until Pkg C6
- Host
 - Aggregate BW through the controller at least 3 GB/s, direct connection or over USB 4.
 - Wake capable on each host port from S0i3, Sx: Wake on Connects, Disconnects, Device Wake.
- Device
 - Aggregate BW through xHCI controller of at least 3 GB/s
 - D0i2 and D0i3 power gating
 - Wake capable on host initiated wakes when the system is in S0i3, Sx Available on all ports
- Port Routing Control for Dual Role Capability
 - Needs to support SW/FW and ID pin based control to detect host versus device attach
 - SW mode requires PD controller or other FW to control
- USB-R device to host controller connection is over UTMI+ links.

Table 6-1. USB-C* Port Configuration

Port Group	Port	UP4- Processor Line	UP3- Processor Line	H - Processor Line
Group A	TCP 0	USB 4 ⁴	USB 4 ⁴	USB 4 ⁴
	TCP 1	USB 3 ³	USB 3 ³	USB 3 ³
Group B		DisplayPort ¹	DisplayPort ¹	DisplayPort ¹
	TCP 2	HDMI ²	HDMI ²	HDMI ²
	TCP 3	N/A		

Notes:

1. Supported on Type-C or Native connector
2. HDMI v2.0b is supported only on Native connector.
3. USB 3 supported link rates:
 - a. USB 3 Gen 1x1 (5 Gbps)
 - b. USB 3 Gen 2x1 (10 Gbps)
4. USB4 operating link rates (including both rounded and non-rounded modes for Thunderbolt™ 3 compatibility):
 - a. USB 4 Gen 2x2 (20 Gbps)
 - b. USB 4 Gen 3x2 (40 Gbps)
 - c. 10.3125 Gbps, 20.625 Gbps - Compatible to Thunderbolt™ 3 non-rounded modes.
5. USB 2 interface supported over Type-C connector, sourced from PCH.
6. USB Type-A connector is not supported.
7. Port group is defined as two ports sharing the same USB4 router, each router supports up to two display interfaces.
8. Display interface can be connected directly to a DP/HDMI/Type-C port or thru USB 4 router (Tunneled) on a Type-C connector.
9. If two ports in the same group are configured to one as USB4 and the other as DP/HDMI fixed connection each port will support single display interface.

Table 6-2. USB-C* Lanes Configuration

Lane1	Lane2	Comments
USB 4	USB 4	Both lanes operate at Gen 2 (10G) or Gen 3 (20G) and also support non-rounded frequencies (10.3125G / 20.625G) for TBT3 compatibility.
USB3	No connect	Any combination of <ul style="list-style-type: none"> • USB3.2 Gen 1x1 (5 Gbps) • USB3.2 Gen 2x1 (10 Gbps)
No connect	USB3	
USB3	DPx2	Any of HBR3/HBR2/HBR1/RBR for DP and USB3.2 (10 Gbps)
DPx2	USB3	
DPx4		Both lanes at the same DP rate - no support for 2x DPx2 USB-C connector

Table 6-3. USB-C* Non-Supported Lane Configuration

Lane1	Lane2	Comments
#	PCIe* Gen3/2/1	No PCIe* native support
PCIe* Gen3/2/1	#	
#	USB4	No support for USB4 with any other protocol
USB4	#	
USB3	USB3	No support for Multi-lane USB3

6.2 USB™ 4 Router

USB4 is a Standard architecture (formerly known as CIO), but with the addition of USB3 (10G) tunneling, and rounded frequencies. USB4 adds a new USB4 PD entry mode, but fully documents mode entry, and negotiation elements of Thunderbolt™ 3.

USB4 architecture (formerly known as Thunderbolt™ 3 protocol) is a transformational high-speed, dual protocol I/O, and it provides flexibility and simplicity by encapsulating both data (PCIe* and USB3) and video

(DisplayPort*) on a single cable connection that can daisy-chain up to six devices. USB4/Thunderbolt™ controllers act as a point of entry or a point of exit in the USB4 domain. The USB4 domain is built as a daisy chain of USB4/Thunderbolt™ enabled products for the encapsulated protocols - PCIe, USB3 and DisplayPort. These protocols are encapsulated into the USB4 fabric and can be tunneled across the domain.

USB4 controllers can be implemented in various systems such as PCs, laptops and tablets, or devices such as storage, docks, displays, home entertainment, cameras, computer peripherals, high end video editing systems, and any other PCIe based device that can be used to extend system capabilities outside of the system's box.

The integrated connection maximum data rate is 20.625 Gbps per lane but supports also 20.0 Gbps, 10.3125 Gbps, and 10.0 Gbps and is compatible with older Thunderbolt™ device speeds.

6.2.1 USB 4 Host Router Implementation Capabilities

The integrated USB-C sub-system implements the following interfaces via USB 4:

- Two DisplayPort* sink interfaces each one capable of:

- DisplayPort 1.4 specification for tunneling
- 1.62 Gbps or 2.7 Gbps or 5.4 Gbps or 8.1 Gbps link rates
- x1, x2 or x4 lane operation
- Support for DSC compression
- Two PCI Express* Root Port interfaces each one capable of:
 - PCI Express* 3.0 x4 compliant @ 8.0 GT/s
- Two xHCI Port interfaces each one capable of:
 - USB 3.2 Gen2x1 (10 Gbps)
- USB 4 Host Interface:
 - PCI Express* 3.0 x4 compliant endpoint
 - Supports simultaneous transmit and receive on 12 paths
 - Raw mode and frame mode operation configurable on a per-path basis
 - MSI and MSI-X support
 - Interrupt moderation support
- USB 4 Time Management Unit (TMU):
- Two Interfaces to USB-C* connectors, each one supports:
 - USB4 PD entry mode, as well as TBT 3 compatibility mode, each supporting:
 - 20 paths per port
 - Each port support 20.625/20.0 Gbps or 10.3125/10.0 Gbps link rates.
 - 16 counters per port

6.3 USB-C Sub-system xHCI/xDCI Controllers

The processor supports xHCI/xDCI controllers. The native USB 3 path proceeds from the memory directly to PHY.

6.3.1 USB 3 Controllers

6.3.1.1 Extensible Host Controller Interface (xHCI)

Extensible Host Controller Interface (xHCI) is an interface specification that defines Host Controller for a universal Serial Bus (USB 3), which is capable of interfacing with USB 1.x, 2.0, and 3.x compatible devices.

In case that a device (example, USB3 mouse) was connected to the computer, the computer will work as Host and the xHCI will be activated inside the CPU.

The xHCI controller support link rate of up to USB 3.2 Gen 2x1 (10G).

6.3.1.2 Extensible Device Controller Interface (xDCI)

Extensible Device Controller Interface (xDCI) is an interface specification that defines Device Controller for a universal Serial Bus (USB 3), which is capable of interfacing with USB 1.x, 2.0, and 3.x compatible devices.



In case that the computer is connected as a device (example, tablet connected to desktop) to another computer then the xDCI controller will be activated inside the device and will talk to the Host at the other computer.

The xDCI controller support link rate of up to USB 3.2 Gen 1x1 (5G).

Notes: These controllers are instantiated in the processor die as a separate PCI function functionality for the USB-C* capable ports.

6.3.2 USB-C Sub-System PCIe Interface

Table 6-4. PCIe via USB4 Configuration

USB4 IPs	USB4_PCIe	UP4 USB-C* Ports	UP3 USB-C* Ports	H USB-C* Ports
USB4_DMA0	USB4_PCIE0	TCP0	TCP0	TCP0
	USB4_PCIE1	TCP1	TCP1	TCP1
USB4_DMA1	USB4_PCIE2	TCP2	TCP2	TCP2
	USB4_PCIE3	N/A	TCP3	TCP3

6.4 USB-C Sub-System Display Interface

Refer [Chapter 10, "Display"](#)

§ §

7 PCIe* Interface

7.1 Processor PCI Express* 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*.

7.1.1 PCI Express* Support

The UP4/UP3 processor PCI Express* interface a single 4-lane (x4) port. The interconnect between the UP4/UP3 processor and NVMe* storage provided through an M.2 connector (HotPlug is not supported for UP3/UP4 M.2 connector). In addition, it will support graphics PCIe Gen4 devices too.

The H processor PCI Express* has two interfaces:

- 4-lane (x4) port, traditionally used for an interconnect to NVMe* storage through an M.2 connector.
- 16-lane (x16) port that can also be configured as multiple ports at narrower widths. Traditionally used for an interconnect to graphics PCIe Gen4 devices through device down. dGPU support through the use of CEM form factor is supported on H-processor DT only. x16 bifurcated can also be used to interface NVMe* storage devices through an M.2 connector.

The UP4/UP3/H processor X4 port supports the configurations shown in the following table:

Table 7-1. 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
Note: PCIe 060 is a single x4 port without bifurcation capabilities, thus bifurcation pin straps are not applicable								

The H processor X16 port supports the configurations shown in the following table:

Table 7-2. 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								

Table 7-2. 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
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:

- For CFG bus details, refer to [Section 6.4](#).
- 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.
- 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 (PCIe 010 controller).
 - Connect lane 0 of 2nd device to lane 8 to (PCIe 011 controller).
 - Connect lane 0 of 3rd device to lane 12 to i(PCIe 012 controller).
 For example:
 - When using 1x8 + 2x4, the 8 lane device should use lanes 0:7.
 - When using 1x4 + 1x2, the 4 lane device should use lanes 0:3, and other 2 lanes device should use lanes 8:9.
 - 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.
- 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.

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.
- Dynamic width capability.

- Message Signaled Interrupt (MSI and MSI-X) messages.
- Lane reversal
- Advanced Error Reporting (AER)

TC/VC Mapping - Allows mapping Traffic Classes to different internal virtual channels. While default configuration may apply to most use cases, the use of certain Traffic Class may impact performance. This capability should be enabled using the BIOS.

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

Table 7-3. PCI Express* Maximum Transfer Rates and Theoretical Bandwidth

PCI Express* Generation	Encoding	Maximum Transfer Rate [GT/s]	Theoretical Bandwidth [GB/s]		
			All processor lines x4	H-processor line x8	H-processor line x16
Gen 1	8b/10b	2.5	1.0	2.0	4.0
Gen 2	8b/10b	5	2.0	4.0	8.0
Gen 3	128b/130b	8	3.9	7.9	15.8
Gen 4	128b/130b	16	7.9	15.8	31.5

Note: Theoretical BW is the Maximum BW during data streaming, without considering utilization factor and overheads.

7.1.2 PCI Express* Lane Polarity Inversion

The PCI Express* Base Specification requires polarity inversion to be supported independently by all receivers across a Link. Each differential pair within each Lane of a PCIe* Link handles its own polarity inversion. Polarity inversion is applied, as needed, during the initial training sequence of a Lane. In other words, a Lane will still function correctly even if a positive (Tx+) signal from a transmitter is connected to the negative (Rx-) signal of the receiver. Polarity inversion eliminates the need to untangle a trace route to reverse a signal polarity difference within a differential pair and no special configuration settings are necessary to enable it.

Note: The polarity inversion does not imply direction inversion or direction reversal; that is, the Tx differential pair from one device must still connect to the Rx differential pair on the receiving device, per the PCIe* Base Specification. Polarity Inversion is not the same as "PCI Express* Controller Lane Reversal".

7.1.3 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 at 16 GT/s uses a 128b/130b encoding.

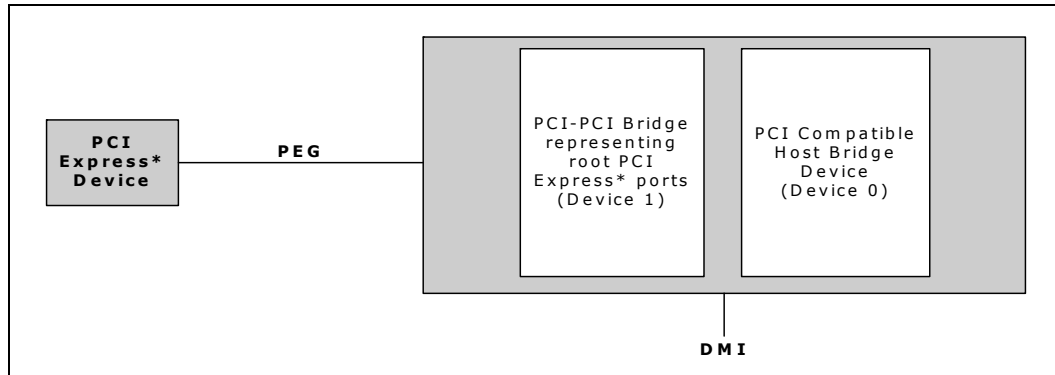
The four lanes port can operate at 2.5 GT/s, 5 GT/s, 8 GT/s or 16 GT/s.

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.

7.1.4 PCI Express* Configuration Mechanism

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

Figure 7-1. PCI Express* Related Register Structures in the 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 to the PCI Express Base Specification for details of both the PCI-compatible and PCI Express* Enhanced configuration mechanisms and transaction rules.

7.1.5 PCI Express* Equalization Methodology

Link equalization requires equalization for both TX and RX sides for the processor and for the Endpoint 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) values.
- 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* Gen 3 and Gen 4 specification.

7.1.6 PCI Express* Hot-Plug

All PCIe* Root Ports support Express Card 1.0 based hot - plug that performs the following:

- Presence Detect and Link Active Changed Support
- Interrupt Generation Support

7.1.6.1 Presence Detection

When a module is plugged in and power is supplied, the physical layer will detect the presence of the device, and the root port sets SLSTS.PDS and SLSTS.PDC. If SLCTL.PDE and SLCTL.HPE are both set, the root port will also generate an interrupt.

When a module is removed (using the physical layer detection), the root port clears SLSTS.PDS and sets SLSTS.PDC. If SLCTL.PDE and SLCTL.HPE are both set, the root port will also generate an interrupt.

7.1.6.2 SMI/SCI Generation

Interrupts for power - management events are not supported on legacy operating systems. To support power - management on non - PCI Express* aware operating systems, power management events can be routed to generate SCI. To generate SCI, MPC.HPCE must be set. When set, enabled hot - plug events will cause SMSCS.HPCS to be set

Additionally, BIOS workaround for hot - plug can be supported by setting MPC.HPME.

When this bit is set, hot - plug events can cause SMI status bits in SMSCS to be set.

Supported hot - plug events and their corresponding SMSCS bit are:

- Presence Detect Changed – SMSCS.HPPDM
- Link Active State Changed – SMSCS.HPLAS

When any of these bits are set, SMI# will be generated. These bits are set regardless of whether interrupts or SCI is enabled for hot - plug events. The SMI# may occur concurrently with an interrupt or SCI

Notes:

- SMI is referred to Serial management Interfaces
- SLSTS - Slot Status
- SLCTL – Slot Control



8 Direct Media Interface (DMI) and On Package Interface (OPI)

8.1 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)

8.1.1 DMI Lane Reversal and Polarity Inversion

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

Figure 8-1. Example for DMI Lane Reversal Connection

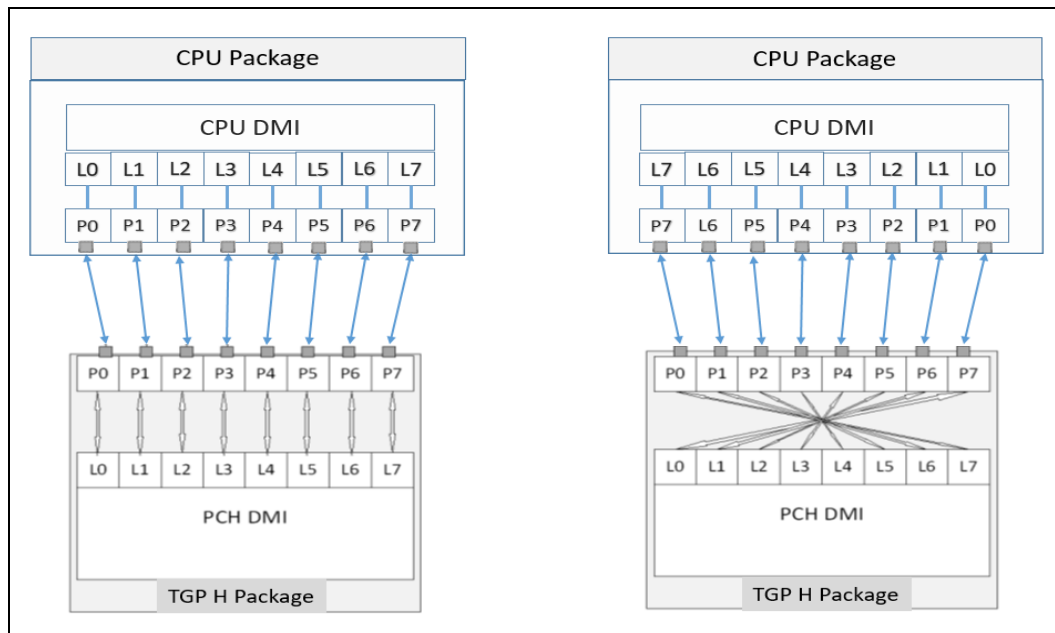


Figure 8-1. Example for DMI Lane Reversal Connection**Notes:**

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. For further data, refer to PCH Datasheet

8.1.2 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.

8.1.3 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.

8.2 On Package Interface (OPI)**8.2.1 OPI Support:**

The processor communicates with the PCIe using an internal interconnect BUS named OPI.

8.2.2 Functional description:

OPI supports either 4 GT/s or 2 GT/s data rates where 4GT/s is supported by UP3 and 2GT/s is supported by UP4. 4GT/s OPI operates at 0.95V and 2GT/s OPI operates at 0.85V.

§ §

9 Graphics

9.1 Processor Graphics

The processor graphics is based on X^e graphics core architecture that enables substantial gains in performance and lower-power consumption over prior generations. X^e architecture supports up to 96 Execution Units (EUs) depending on the processor SKU.

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.

9.1.1 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: HEVC and VP9 support additional 10 bpc, YCbCr 4:2:2 or 4:4:4 profiles. Refer additional detail support matrix.

9.1.1.1 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)
- Direct3D11 Video API
- Intel Media SDK
- MFT (Media Foundation Transform) filters.
- Intel VA API

X^e supports full HW accelerated video decoding for AVC/VC1/MPEG2/HEVC/VP9/JPEG/AV1.

Table 9-1. Hardware Accelerated Video Decoding (Sheet 1 of 2)

Codec	Profile	Level	Maximum Resolution
MPEG2	Main	Main High	1080p

Table 9-1. Hardware Accelerated Video Decoding (Sheet 2 of 2)

Codec	Profile	Level	Maximum Resolution
WMV9	Advanced Main Simple	L3 High Simple	3840x3840
AVC/H264	High Main	L5.2	4K
	4:2:0 8 bit		4K@60
JPEG/MJPEG	Baseline	Unified level	16Kx16K
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@60 8K@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 10 bit		5K@60
	4:2:0 12 bit		8K@60
AV1	0 (4:2:0 8 bit) 0 (4:2:0 10 bit)	L3	4Kx2K (video) 16Kx16K (still picture)

Expected performance:

- More than 16 simultaneous decode streams @ 1080p.

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

9.1.1.2 Hardware Accelerated Video Encode

Gen12 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/HEVC/VP9/JPEG.

Table 9-2. Hardware Accelerated Video Encode (Sheet 1 of 2)

Codec	Profile	Level	Maximum Resolution
AVC/H264	High Main	L5.1	2160p(4 K)

Table 9-2. Hardware Accelerated Video Encode (Sheet 2 of 2)

Codec	Profile	Level	Maximum Resolution
JPEG	Baseline	—	16 Kx16 K
HEVC/H265	Main Main10 Main 4:2:2 10 Main 4:4:4 Main 4:4:4 10	L5.1	4320p(8 K) 16 Kx4 K @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(8 K) 16 Kx4 K @higher freq

Note: Hardware encode for H264 SVC is not supported.

9.1.1.3 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.
- Intel VA API

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

9.1.1.4 Hardware Accelerated Transcoding

Transcoding is a combination of decode, video processing (optional) and encode. Using the above 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.

9.2 Platform Graphics Hardware Feature

9.2.1 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.

§ §

10 Display

10.1 Display Technologies Support

Technology	Standard
eDP* 1.4b	VESA* Embedded DisplayPort* Standard 1.4b
MIPI DSI	MIPI* Display Serial Interface (DSI) Specification Version 1.3
DisplayPort* 1.4a	VESA* DisplayPort* Standard 1.4a VESA* DisplayPort* PHY Compliance Test Specification 1.4a VESA* DisplayPort* Link Layer Compliance Test Specification 1.4a VESA* DisplayPort* Alt Mode on USB Type-C Standard Version 1.0b
HDMI* 2.0b	High-Definition Multimedia Interface Specification Version 2.0b

10.2 Display Configuration

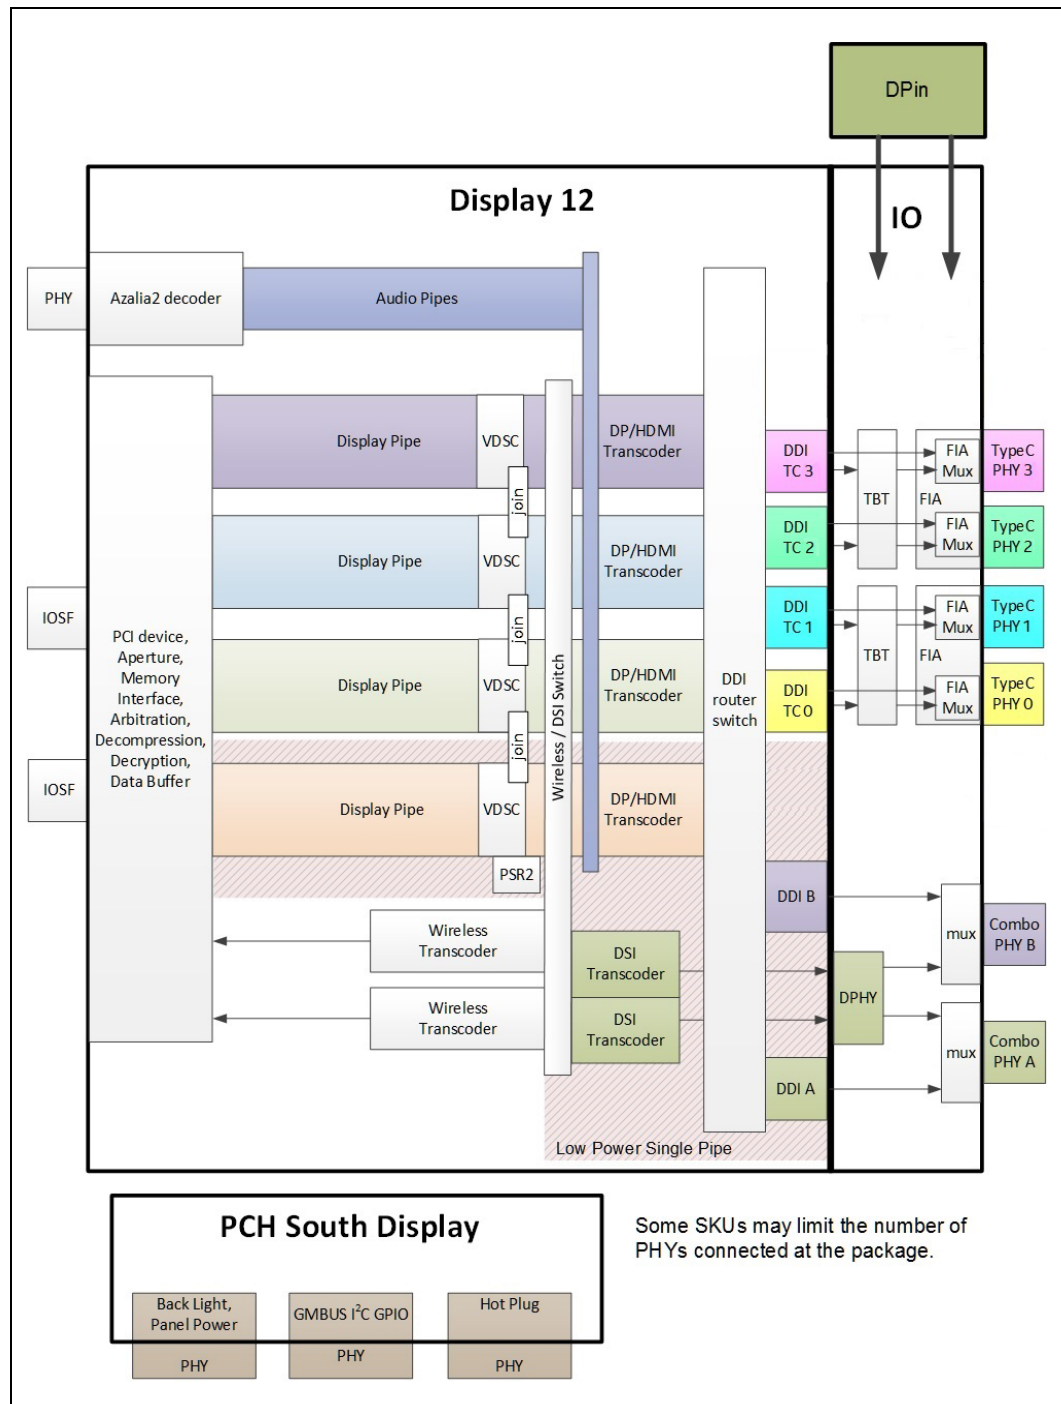
Table 10-1. Display Ports Availability and Link Rate

Port	UP3-Processor Line	UP4-Processor Line	H-Processor Line
DDI A ^{1,2}	DP* up to HBR2 HDMI* up to 5.94 Gbps eDP* up to HBR3 MIPI DSI up to 2.5 Gbps ⁵	DP* up to HBR2 HDMI* up to 5.94 Gbps eDP* up to HBR3 MIPI DSI up to 2.5 Gbps	eDP* up to HBR3
DDI B ²	eDP* up to HBR3 DP* up to HBR2 HDMI* up to 5.94 Gbps	eDP* up to HBR3 DP* up to HBR2 HDMI* up to 5.94 Gbps MIPI DSI up to 2.5 Gbps	eDP* up to HBR3 DP* up to HBR2 HDMI* up to 5.94 Gbps
DDI TCP0 ³	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
DDI TCP1 ³	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
DDI TCP2 ³	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps	DP* up to HBR3 HDMI* up to 5.94 Gbps
DDI TCP3 ³	DP* up to HBR3 HDMI* up to 5.94 Gbps	N/A	DP* up to HBR3 HDMI* up to 5.94 Gbps
DPIP0 ⁴	N/A	N/A	DP* up to HBR3
DPIP1 ⁴			DP* up to HBR3
DPIP2 ⁴			DP* up to HBR3
DPIP3 ⁴			DP* up to HBR3

Notes:

- Dual low power embedded panel are supported (each can be eDP and/or MIPI DSI).
 - PSR2 can be supported only on single low power display.
 - Highest Package C state supported for dual embedded display configuration is PC8.
- DDI - Digital Display Interface.
- Each of the four TCP ports can be implemented as HDMI, DP, or DPoC (DisplayPort over Type-C)
- DPIP_x are DisplayPort* Rx ports referred as DP-in port, for more information refer to DP-IN section.
- MIPI DSI supported on UP3 processor family but not fully validated.

Figure 10-1. Processor Display Architecture



Note: For port availability in each of the processor lines, refer [Table 10-1, "Display Ports Availability and Link Rate"](#).

10.3 Display Features

10.3.1 General Capabilities

- Up to four simultaneous displays.
 - Single 8K60Hz panel, supported by joining two pipes over single port.
 - Up to 4x4K60Hz display concurrent.
- Audio stream support on external ports.
- Up to four USB* Type-C for DisplayPort* Alt Mode, DisplayPort* over TBT, or DisplayPort*/HDMI* interfaces.
- Up to four DP-IN interfaces supporting discrete GPU DisplayPort* stream via USB-C* sub system
- HDR (High Dynamic Range) support.
- Four Display Pipes - Supporting blending, color adjustments, scaling and dithering.
- Transcoders - Containing the Timing generators supporting eDP*, DSI*, DP*, HDMI* interfaces.
- One Low Power optimized pipe.
 - LACE (Localized Adaptive Contrast Enhancement), supported up to 5 K resolutions.
 - 3D LUT - power efficient pixel modification function for color processing.
 - FBC (Frame Buffer Compression) - power saving feature.

Note: DP-IN supported only in H segment.

10.3.2 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 four 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 four 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.

10.3.3 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 both HDCP 2.3 and 1.4 content protection over wired displays (HDMI* 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.

10.3.4 DisplayPort*

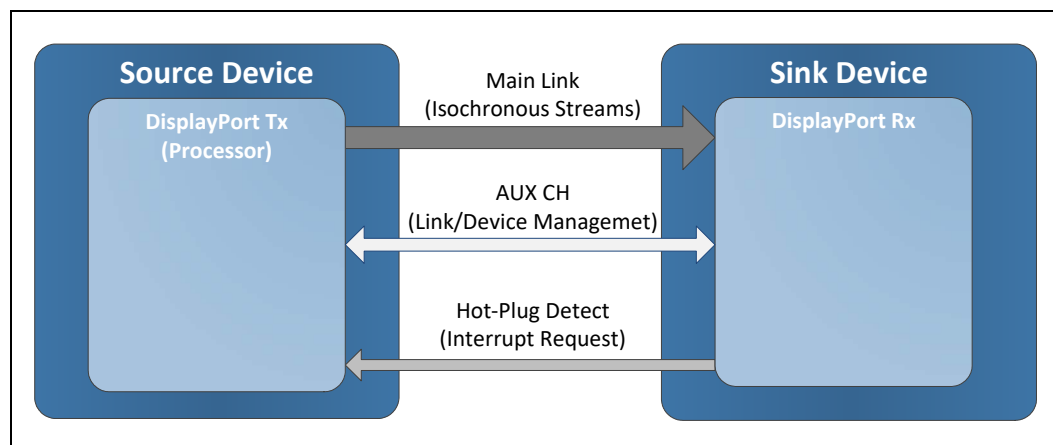
The DisplayPort* is a digital communication interface that uses differential signaling to achieve a high-bandwidth bus interface designed to support connections between PCs and monitors, projectors, and TV displays.

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 [Section 10.1, "Display Technologies Support"](#).

The DisplayPort* support DisplayPort* Alt mode over Type-C and DP tunneling via TBT. Refer to [Chapter 6, "USB-C* Sub System"](#) For DisplayPort* Alt mode support.

Figure 10-2. 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.

10.3.4.1 Multi-Stream Transport (MST)

- The processor supports Multi-Stream Transport (MST), enabling multiple monitors to be used via a single DisplayPort connector.
- Maximum MST DP supported resolution:

Table 10-2. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations

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
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 above 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
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@60 Hz = 16 + 4.62 = 20.62 Gbps [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 10-3. DisplayPort Maximum Resolution

Standard	UP3-Processor Line ¹	UP4-Processor Line ¹	H-Processor Line ¹
DP*	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp
DP* with DSC	5120x3200 120 Hz 30 bpp	5120x3200 120 Hz 30 bpp	5120x3200 120 Hz 30 bpp
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. 4. Due to the support of low power mode, no compressed pipe joiner feature is co-exist which leads to the maximum of 5K60 resolution.			

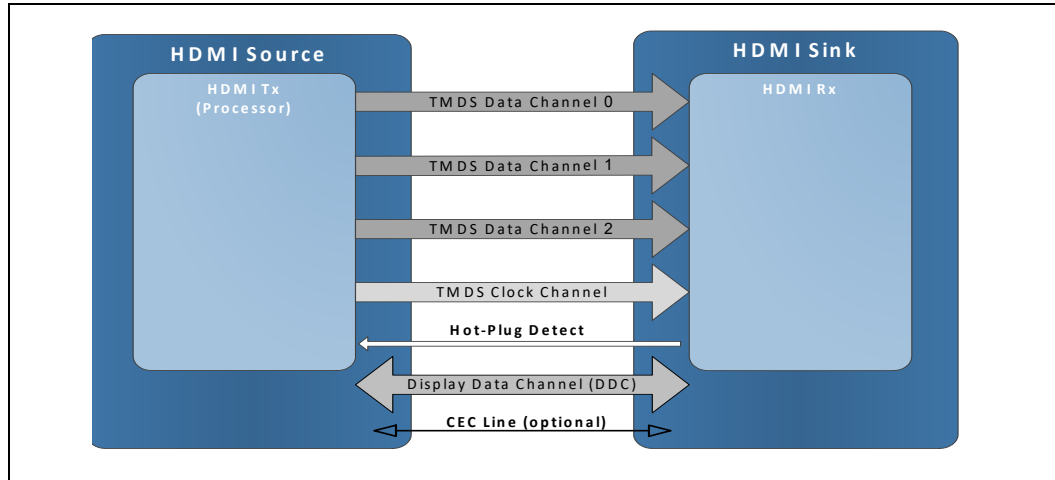
10.3.5 High-Definition Multimedia Interface (HDMI*)

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 10-3. 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 10-4. HDMI Maximum Resolution

Standard	UP3-Processor Line	UP4-Processor Line	H-Processor Line
HDMI 1.4	4 Kx2 K 24-30 Hz 24 bpp	4 Kx2 K 24-30 Hz 24 bpp	4 Kx2 K 24-30 Hz 24 bpp
HDMI 2.0b	4 Kx2 K 48-60 Hz 24 bpp (RGB/YUV444) 4 Kx2 K 48-60 Hz 12 bpc (YUV420)	4 Kx2 K 48-60 Hz 24 bpp (RGB/YUV444) 4 Kx2 K 48-60 Hz 12 bpc (YUV420)	4 Kx2 K 48-60 Hz 24bpp (RGB/YUV444) 4 Kx2 K 48-60 Hz 12bpc (YUV420)
Notes:			
1. bpp - bit per pixel.			
2. Resolution support is subject to memory BW availability.			
3. HDMI2.1 Can be supported using LSPCON (DP1.4 to HDMI2.1 protocol converter).			

10.3.6 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 HBR3 link rate.
- Support Backlight PWM control signal.
- Support VESA DSC (Data Stream Compression).
- Support SSC.
- Panel Self Refresh 1.
- Panel Self Refresh 2.
- MSO 2x2 (Multi Segment Operation).
- Dedicated Aux channel.

- Adaptive Sync.

Table 10-5. Embedded DisplayPort Maximum Resolution

Standard	UP3-Processor Line ¹	UP4-Processor Line ¹	H-Processor Line ¹
eDP*	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp	4096x2304 60 Hz 36 bpp 5120x3200 60 Hz 24 bpp
eDP* with DSC ⁵	5120x3200 120 Hz 30 bpp 7680x4320 60 Hz 30 bpp	5120x3200 120 Hz 30 bpp 7680x4320 60 Hz 24 bpp	5120x3200 120 Hz 30 bpp 7680x4320 60 Hz 30 bpp

Notes:

1. Maximum resolution is based on the implementation of 4 lanes at HBR3 link data rate.
2. PSR2 supported for up to 5 K resolutions.
3. bpp - bit per pixel.
4. Resolution support is subject to memory BW availability.
5. High resolution are supported, validation is depended on panel market availability.

10.3.7 MIPI* DSI

Display Serial Interface (DSI*) specifies the interface between a host processor and peripherals such as a display module. DSI is a high speed and high performance serial interface that offers efficient and low power connectivity between the processor and the display module.

- One link x8 data lanes or two links each with x4 lanes support.
- Supported on Low power optimized pipe.
- Support the Backlight control signal.
- Support VESA DSC (Data Stream Compression).

Figure 10-4. MIPI* DSI Overview

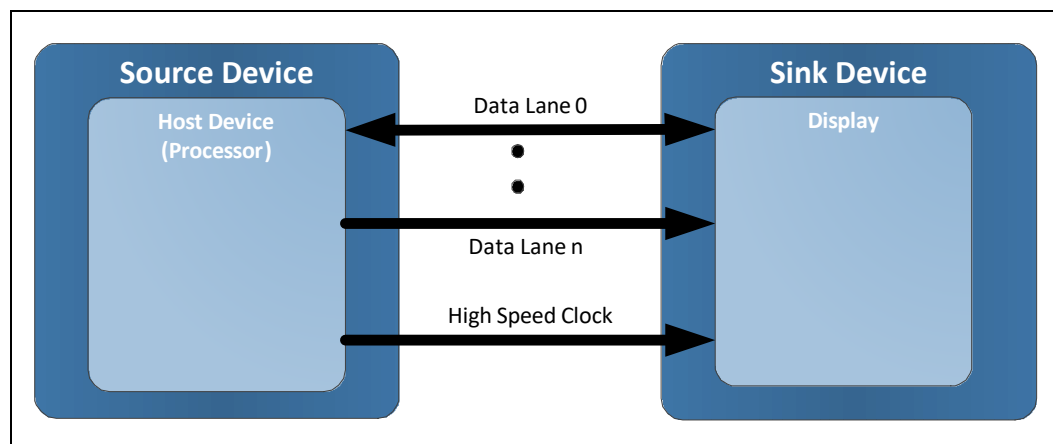


Table 10-6. MIPI* DSI Maximum Resolution

Standard	UP3-Processor Line	UP4-Processor Line	H-Processor Line
MIPI* DSI (Single Link)	3200x2000 @60 Hz 24 bpp	3200x2000 @60 Hz 24 bpp	N/A
MIPI* DSI (Single Link) with DSC	5120x3200 @60 Hz 24 bpp	5120x3200 @60 Hz 24bpp	
MIPI* DSI (Dual Link)	N/A	4096x2304 @60 Hz 24 bpp 3840x2160 @60 Hz 24 bpp	
MIPI* DSI (Dual Link) with DSC	N/A	5120x3200 @60 Hz 24 bpp	
Notes:			
1. bpp - bit per pixel.			
2. Resolution support is subject to memory BW availability.			

10.3.8 Integrated Audio

- HDMI* and DisplayPort interfaces can carry audio along with video.
- The processor supports three High Definition audio streams on four 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 10-7. 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.

10.3.9 DisplayPort* Input (DP-IN)

DP-IN interface supports platforms that are using discrete GPU with DisplayPort* output interfaces.

Each stream transmitted from the discrete GPU towards DP-IN Receiver interface can be internally routed to each of USB-C* sub system ports, as long as a Type-C solution have been implemented:

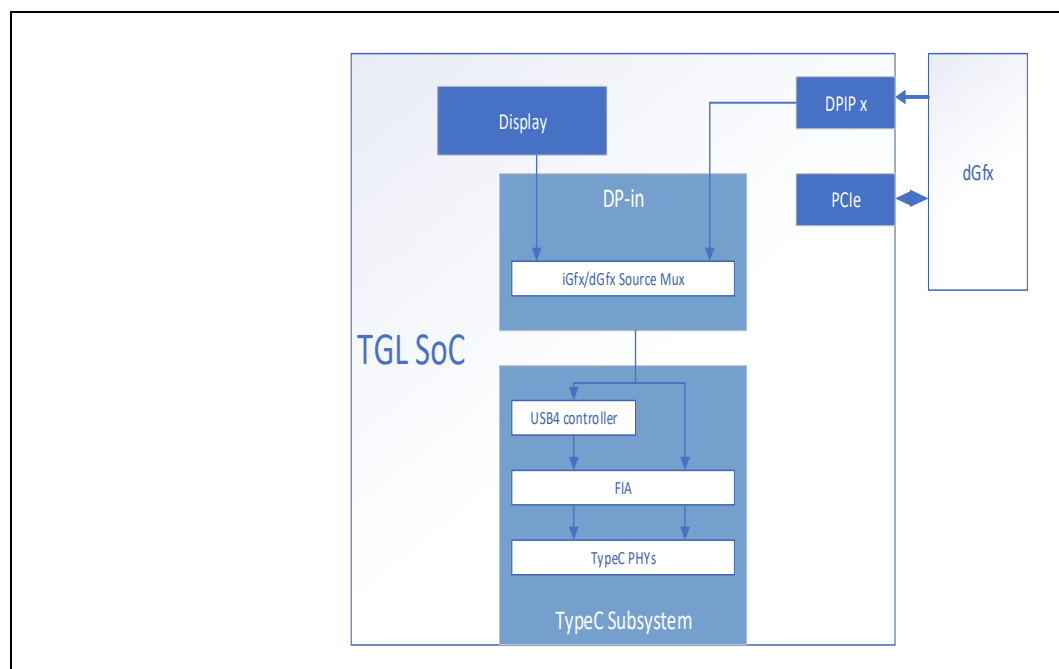
- DPoC port - Display Port Over Type-C.
- USB4 port - DisplayPort tunneled over USB4.

DP-IN interface support VESA* LTTPR (Link Training Tunable PHY Repeater).

Notes:

1. DP-IN is supported only in H processor line.
2. DP-IN requires an external display source that supports VESA* LTTPR (Link Training Tunable PHY Repeater). The following modes are supported
 - i. Non-transparent mode – recommended model.
 - ii. Transparent mode – this mode is limited only for TCP ports which is connected through a BBR re-timer.

Figure 10-5. DP-IN Block Diagram



Each DP-in port support:

- Hot Plug Detect, required on board Level shifter
- AUX channel
- Main Link supporting up to 4 lanes each with up to HBR3 link rate.
- VESA* Link Training Tunable PHY Repeater
- Supported for DisplayPort over Type-C (DPoC) and DisplayPort* tunneling via thunderbolt on each of USB-C* ports.

Note:

Supported for DisplayPort over Type-C (DPoC) and DisplayPort* tunneling via thunderbolt on each of USB-C* ports.

Not supported for Fixed DP, HDMI, or eDP interfaces.

11 Camera/MIPI

11.1 Camera Pipe Support

The IPU6se fixed function pipe supports the following functions:

- Black level correction;
- White balance;
- Color matching;
- Lens shading (vignette) correction;
- Color crosstalk (color shading) correction;
- Dynamic defect pixel replacement;
- Auto-focus-pixel (PDAF) hiding;
- High quality demosaic;
- Scaling and format conversion;
- Temporal noise reduction running on Intel graphics.

11.2 MIPI* CSI-2 Camera Interconnect

The Camera I/O Controller provides a native/integrated interconnect to camera sensors, compliant with MIPI* CSI-2 V2.0 protocol. Total of 12 data + 4 clock lanes (UP3), 15 data + 6 clock (UP4) and 4 data + 2 clock (H-processor) lanes are available for the camera interface supporting up to 4 sensors (UP3), 6 sensors (UP4) or 2 sensors (H-processor).

Data transmission interface (referred as CSI-2) is a unidirectional differential serial interface with data and clock signals; the physical layer of this interface is the MIPI* Alliance Specification for D-PHY.

The control interface (referred as CCI) is a bi-directional control interface compatible with I²C standard.

11.2.1 Camera Control Logic

The camera infrastructure supports several architectural options for camera control utilizing camera PMIC and/or discrete logic. IPU6 control options utilize I²C for bidirectional communication and PCH GPIOs to drive various control functions.

11.2.2 Camera Modules

Intel maintains an Intel User Facing Camera Approved Vendor List and Intel World-Facing Approved Vendor List to simplify system design. Additional services are available to support non-AVL options.



11.2.3 CSI-2 Lane Configuration

Table 11-1. CSI-2 Lane Configuration for UP3-Processor Line

Port Data/Clock	Configuration Option 1	Configuration Option 2
Port B Lane 0	x4	x4
Port B Lane 1		
Port B Clock		
Port B Lane 2		
Port B Lane 3		
Port C Clock	x4	x4
Port C Lane 0		
Port C Lane 1		
Port C Lane 2		
Port C Lane 3		
Port E Clock	Not Used	x2
Port E Lane 0	x4	
Port E Lane 1		
Port F Clock		
Port F Lane 0		x2
Port F Lane 1		

Table 11-2. CSI-2 Lane Configuration for UP4-Processor Line (Sheet 1 of 2)

Port Data/Clock	Configuration Option 1	Configuration Option 2
Port B Lane 0	x4	x4
Port B Lane 1		
Port B Clock		
Port B Lane 2		
Port B Lane 3		
Port C Clock	x4	x4
Port C Lane 0		
Port C Lane 1		
Port C Lane 2		
Port C Lane 3		
Port E Clock	Not Used	x2
Port E Lane 0	x4	
Port E Lane 1		
Port F Clock		
Port F Lane 0		x2
Port F Lane 1		
Port G Clock	x2	x2
Port G Lane 0		
Port G Lane 1		

Table 11-2. CSI-2 Lane Configuration for UP4-Processor Line (Sheet 2 of 2)

Port Data/Clock	Configuration Option 1	Configuration Option 2
Port H Lane 0	Not Used	x1
Port H Clock		
Notes: 1. Ports G and H available on UP4 only. 2. Port E,F selection of configuration 1 or 2 is orthogonal to Port G, H selection of configuration 1 or 2.		

Table 11-3. CSI-2 Lane Configuration for H-Processor Line

Port Data/Clock	Configuration Option 1	Configuration Option 2
Port A Clock	x2	x4
Port B Lane 0		
Port B Lane 1		
Port B Clock	x2	
Port B Lane 2 / Port A Lane 1		
Port B Lane 3 / Port A Lane 0		
Note: H-processor IPU6-Slim has 4 data and 2 clock lanes only.		

§ §

12 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 to the following table).

Table 12-1. Signal Tables Terminology (Sheet 1 of 2)

	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 (1.2 V tolerant)
	LPDDR4	LPDDR4 memory (1.1 V tolerant)
	LPDDR4x	LPDDR4x memory (1.1 V TX, 0.6 V RX tolerant)
	Analog	Analog reference or output. May be used as a threshold voltage or for buffer compensation.
	PCIE	PCI Express
	DMI	DMI
	GTL	Gunning Transceiver Logic signaling technology
	CMOS	CMOS
	AUDIO	AUDIO
	N/A	Not Applicable
	Async CMOS ¹	Async CMOS
	DP/HDMI	DP/HDMI
	DPHY	Used for DSI and CSI
	OD	Open Drain
	PECI Async	PECI Async
	Diff Amp Clock	Diff Amp Clock Input Buffer
	Ref	Voltage Reference signal
	Power	Power
	PWR_SENSE	Isolated, low impedance voltage sense pins.
	Ground	Ground
GND_SENSE	Isolated, low impedance reference ground sense pins.	
TCP	Type-C port	
Link Type	SE	Single Ended
	DIFF	Differential pair

Table 12-1. Signal Tables Terminology (Sheet 2 of 2)

	Short name	Functionality
<p>Notes:</p> <ol style="list-style-type: none"> 1. Qualifier for a buffer type¹ 2. CMOS - Complementary Metal Oxide Substrate 3. GTL - Gunning Transceiver Signaling Technology Logic 4. DP - DisplayPort 5. PECEI - 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 		

12.1 System Memory Interface

12.1.1 DDR4 Memory Interface

Table 12-2. DDR4 Memory Interface (Sheet 1 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
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] refers to DDR channel 0, Byte 2, Bit 5.	I/O	DDR4	SE	UP3/H Processor Lines
DDR0_DQ8[7:0] DDR1_DQ8[7:0]	ECC Data Buses: Data signals interface to the SDRAM data buses.	I/O	DDR4	SE	H Processor Line
DDR0_DQSP[7:0] DDR0_DQSN[7:0] DDR1_DQSP[7:0] DDR1_DQSN[7:0]	Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during reading and write transactions. Example: DDR0_DQSP0 refers to DQSP of DDR channel 0, Byte 0.	O	DDR4	Diff	UP3/H Processor Lines
DDR0_DQSP[8] DDR0_DQSN[8] DDR1_DQSP[8] DDR1_DQSN[8]	ECC Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during reading and write transactions.	O	DDR4	Diff	H Processor Line
DDR0_CLK_N[1:0] DDR0_CLK_P[1:0] DDR1_CLK_N[1:0] DDR1_CLK_P[1:0]	SDRAM Differential Clock: Differential clocks signal pairs, pair per rank. The crossing of the positive edge and the negative edge of their complement are used to sample the command and control signals on the SDRAM.	O	DDR4	Diff	UP3 Processor Line
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 and the negative edge of their complement are used to sample the command and control signals on the SDRAM.	O	DDR4	Diff	H Processor Line
DDR0_CKE[1:0] DDR1_CKE[1:0]	Clock Enable: (1 per rank). These signals are used to: <ul style="list-style-type: none"> 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	UP3 Processor Line
DDR0_CKE[3:0] DDR1_CKE[3:0]	Clock Enable: (1 per rank). These signals are used to: <ul style="list-style-type: none"> 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	H Processor Line
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	[1:0] applicable for All Processor Lines. [3:2] applicable only for H-Processor Line.

Table 12-2. DDR4 Memory Interface (Sheet 2 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_ODT[1:3] DDR1_ODT[1:3]	On Die Termination: (1 per rank). Active SDRAM Termination Control.	O	DDR4	SE	UP3/H Processor Lines
DDR0_MA[16:0] DDR1_MA[16:0]	<p>Address: These signals are used to provide the multiplexed row and column address to the SDRAM.</p> <ul style="list-style-type: none"> A[16:14] use also as command signals, see ACT_N 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) <p>DDR0_MA[16] uses as RAS# signal DDR0_MA[15] uses as CAS# signal DDR0_MA[14] uses as WE# signal DDR1_MA[16] uses as RAS# signal DDR1_MA[15] uses as CAS# signal DDR1_MA[14] uses as WE# signal</p>	O	DDR4	SE	UP3/H Processor Lines
DDR0_ACT# DDR1_ACT#	Activation Command: ACT# HIGH along with CS_N determines that the signals addresses below have command functionality.	O	DDR4	SE	UP3/H Processor Lines
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	UP3/H Processor Lines
DDR0_BG[1:0] DDR1_BG[1:0]	<p>Bank Group: BG[1:0] define to which bank group an Active, reading, Write or Precharge command is being applied.</p> <p>BG0 also determines which mode register is to be accessed during a MRS cycle.</p>	O	DDR4	SE	UP3/H Processor Lines
DDR0_BA[1:0] DDR1_BA[1:0]	Bank Address: BA[1:0] define to which bank an Active, reading, 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	UP3/H Processor Lines

Table 12-2. DDR4 Memory Interface (Sheet 3 of 3)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR1_CA[12:0]	Command Address: These signals are used to provide the multiplexed command and address to the SDRAM.	O	DDR4	SE	H Processor Line
DDR0_PAR DDR1_PAR	Command and Address Parity: These signals are used for parity check.	O	DDR4	SE	UP3/H Processor Lines
DDR0_VREF_CA[1:0] DDR1_VREF_CA[1:0]	Memory Reference Voltage for Command and Address	O	A	SE	UP3/H Processor Lines
DDR_RCOMP	System Memory Resistance Compensation	I	Analog	SE	UP3/H Processor Lines
DRAM_RESET#/ RESET#	Memory Reset	I	CMOS	SE	UP3/H Processor Lines
DDR_VTT_CTL	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 (H SKU only).	O	CMOS	SE	UP3/H Processor Lines

12.1.2 LPDDR4x Memory Interface

Table 12-3. LPDDR4x Memory Interface (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_DQ[1:0][7:0] DDR1_DQ[1:0][7:0] DDR2_DQ[1:0][7:0] DDR3_DQ[1:0][7:0] DDR4_DQ[1:0][7:0] DDR6_DQ[1:0][7:0] DDR7_DQ[1:0][7:0]	<p>Data Buses: Data signals interface to the SDRAM data buses.</p> <p>Example: DDR0_DQ1[5] refers to DDR channel 0, Byte 1, Bit 5.</p>	I/O	LPDDR4x	SE	UP3/UP4 Processor Lines
DDR0_DQSP[1:0] DDR1_DQSP[1:0] DDR2_DQSP[1:0] DDR3_DQSP[1:0] DDR4_DQSP[1:0] DDR6_DQSP[1:0] DDR7_DQSP[1:0] DDR0_DQSN[1:0] DDR1_DQSN[1:0] DDR2_DQSN[1:0] DDR3_DQSN[1:0] DDR4_DQSN[1:0] DDR6_DQSN[1:0] DDR7_DQSN[1:0]	<p>Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during reading and write transactions.</p>	O	LPDDR4x	Diff	UP3/UP4 Processor Lines
DDR0_CLK_N DDR0_CLK_P DDR1_CLK_N DDR1_CLK_P DDR2_CLK_N DDR2_CLK_P DDR3_CLK_N DDR3_CLK_P DDR4_CLK_N DDR4_CLK_P DDR6_CLK_N DDR6_CLK_P DDR7_CLK_N DDR7_CLK_P	<p>SDRAM Differential Clock: Differential clocks signal pairs, pair per channel and package. The crossing of the positive edge and the negative edge of their complement are used to sample the command and control signals on the SDRAM.</p>	O	LPDDR4x	Diff	UP3/UP4 Processor Lines
DDR0_CKE[1:0] DDR1_CKE[1:0] DDR2_CKE[1:0] DDR3_CKE[1:0] DDR4_CKE[1:0] DDR6_CKE[1:0] DDR7_CKE[1:0]	<p>Clock Enable: (1 per rank) These signals are used to:</p> <ul style="list-style-type: none"> Initialize the SDRAMs during power-up. Power-down SDRAM ranks. Place all SDRAM ranks into and out of self-refresh during STR. 	O	LPDDR4x	SE	UP3/UP4 Processor Lines

Table 12-3. LPDDR4x Memory Interface (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_CS[1:0] DDR1_CS[1:0] DDR2_CS[1:0] DDR3_CS[1:0] DDR4_CS[1:0] DDR6_CS[1:0] DDR7_CS[1:0]	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. The Chip select signal is Active High.	O	LPDDR4x	SE	UP3/UP4 Processor Lines
DDR0_CA[5:0] DDR1_CA[5:0] DDR2_CA[5:0] DDR3_CA[5:0] DDR4_CA[5:0] DDR6_CA[5:0] DDR7_CA[5:0]	Command Address: These signals are used to provide the multiplexed command and address to the SDRAM.	O	LPDDR4x	SE	UP3/UP4 Processor Lines
DDR_RCOMP	System Memory Resistance Compensation	A	Analog	SE	UP3/UP4 Processor Lines
DRAM_RESET#	Memory Reset	I	CMOS	SE	UP3/UP4 Processor Lines

12.2 PCIe4 Gen4 Interface Signals

Signal Name	Description	Dir	Buffer Type	Link Type	Availability
PCIE4_RCOMP_P PCIE4_RCOMP_N	Resistance Compensation for PEG channel	I	Analog	Diff	All Processor Lines
PCIE4_TX_P[3:0] PCIE4_TX_N[3:0]	PCIe Transmit Differential Pairs	O	PCIE	Diff	All Processor Lines
PCIE4_RX_P[3:0] PCIE4_RX_N[3:0]	PCIe Receive Differential Pairs	I	PCIE	Diff	All Processor Lines

12.3 Direct Media Interface (DMI) Signals

Table 12-4. DMI Interface Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DMI_RX_P[7:0] DMI_RX_N[7:0]	DMI Input from PCH: Direct Media Interface receive differential pairs.	I	DMI	Diff	H Processor Line
DMI_TX_P[7:0] DMI_TX_N[7:0]	DMI Output to PCH: Direct Media Interface transmit differential pairs.	O	DMI	Diff	
DMI_RCOMP_P DMI_RCOMP_N	Configuration Resistance Compensation.	I	DMI	Diff	

12.4 PCIe16 Gen4 Interface Signals

Signal Name	Description	Dir	Buffer Type	Link Type	Availability
PCI16_COM0_RCOMP_P PCI16_COM0_RCOMP_N	Resistance Compensation for PEG channel	I	Analog	Diff	H Processor Line
PCI16_TX_P[15:0] PCI16_TX_N[15:0]	PCIe Transmit Differential Pairs	O	PCIE	Diff	H Processor Line
PCI16_RX_P[15:0] PCI16_RX_N[15:0]	PCIe Receive Differential Pairs	I	PCIE	Diff	H Processor Line

12.5 Reset and Miscellaneous Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CFG[17:0]	<p>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.</p> <ul style="list-style-type: none"> • CFG[3], CFG[0]: Reserved configuration lane. • CFG[2]: UP4/UP3 Reserved • CFG[2]: H PCI Express* Static x16 Lanes Numbering Reversal. <ul style="list-style-type: none"> – 1 - (Default) Normal – 0 - Reversed • CFG[4]: eDP enable: <ul style="list-style-type: none"> – 1 = Disabled. – 0 = Enabled. • CFG[6:5]: UP4/UP3 Reserved • CFG[6:5]: H PCI Express* Bifurcation <ul style="list-style-type: none"> – 00 = 1 x8, 2 x4 PCI Express* – 01 = reserved – 10 = 2 x8 PCI Express* – 11 = 1 x16 PCI Express* • CFG[13:7]: Reserved configuration lanes. • CFG[14]: PEG60 (PCIe4) Lane Reversal: <ul style="list-style-type: none"> – 1 - (Default) Normal – 0 - Reversed • CFG[17:15]: Reserved configuration lanes. 	I	GTL	SE	All Processor Lines
CFG_RCOMP	Configuration Resistance Compensation	I	N/A	SE	All Processor Lines
EAR_N/ EAR_N_TEST_NCTF	<ul style="list-style-type: none"> • Stall reset sequence after PCU PLL lock until de-asserted: <ul style="list-style-type: none"> – 1 = (Default) Normal Operation; No stall. – 0 = Stall. 	I	CMOS	SE	All Processor Lines
PROC_POPIRCOMP	POPIO Resistance Compensation	I	N/A	SE	UP3/UP4 Processor Lines
CPU_ID	A PLATFORM indication signal, the signal defers between current Gen processor and the compatibility option, with the new Gen processor. Ensure to connect the signal pin according to usage of the platform. This Signal (PIN) is important for functional and booting.	I	CMOS	SE	H Processor Line

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
PROC_TRIGOUT	Debug pin	O	CMOS	SE	H Processor Line
PROC_TRIGIN	Debug pin	I	CMOS	SE	H Processor Line

12.6 Display Interfaces

12.6.1 Embedded DisplayPort* (eDP*) Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIA/B_TXP[3:0] DDIA/B_TXN[3:0]	embedded DisplayPort Transmit: differential pair	O	DP/HDMI	Diff	All Processor Lines
DDIA/B_AUX_P DDIA/B_AUX_N	embedded DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.	O	DP/HDMI	Diff	All Processor Lines
DISP_UTILS	embedded DisplayPort Utility: Output control signal used for brightness correction of embedded LCD displays with back light modulation. This pin will co-exist with functionality similar to existing BKLCTCL pin on PCH	O	Async CMOS	SE	All Processor Lines
DDI_RCOMP	DDI IO Compensation resistor, supporting DP*, eDP* and HDMI* channels.	I	Analog	SE	All Processor Lines
Note: eDP* implementation go along with additional sideband signals					

12.6.2 MIPI DSI* Signals

Signal Name	Description	Dir	Buffer Type	Link Type	Availability
DDIA/B_TXP[3:0] DDIA/B_TXN[3:0] DDIA/B_AUXP DDIA/B_AUXN	DPHY Transmit: differential pair DPHY Clock: differential pair	I/O	DPHY	Diff	UP4/UP3 Processor Lines UP3 processor lines support DDIA only.
DSI_DE_TE_1 DSI_DE_TE_2	Tearing Effect	I	Async CMOS	SE	UP4/UP3 Processor Lines UP3 processor lines support DSI_DE_TE_1 only.

12.6.3 Digital Display Interface (DDI) Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDIA_TXP[3:0] DDIA_TXN[3:0] DDIB_TXP[3:0] DDIB_TXN[3:0]	Digital Display Interface Transmit: DisplayPort and HDMI Differential Pairs	O	DP*/HDMI	Diff	All Processor Lines
DDIA_AUX_P DDIA_AUX_N DDIB_AUX_P DDIB_AUX_N	Digital Display Interface Display Port Auxiliary: Half-duplex, bidirectional channel consist of one differential pair for each channel.	I/O	DP*	Diff	
Note: HDMI* implementation go along with additional sideband signals					

12.6.4 Digital Display Audio Signals

Signal Name	Description	Dir.	Link Type	Availability
AUDOUT	Serial Data Output for display audio interface	O	SE	H Processor Line
AUDIN	Serial Data Input for display audio interface	I	SE	
AUDCLK	Serial Data Clock	I	SE	

12.7 DP-IN Interface Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DPIPO_RXP/N[3:0] DPIP1_RXP/N[3:0] DPIP2_RXP/N[3:0] DPIP3_RXP/N[3:0]	DisplayPort* Receiver: DisplayPort Differential Pairs	I	DP*	Diff	H Processor Line
DPIPO_AUX_P/N DPIP1_AUX_P/N DPIP2_AUX_P/N DPIP3_AUX_P/N	DP-IN Display Port Auxiliary: Half-duplex, bidirectional channel consist of one differential pair for each channel.	I/O	DP*	Diff	
DPIPO_RCOMP DPIP1_RCOMP DPIP2_RCOMP DPIP3_RCOMP	IO Compensation resistor, supporting DP* channel.	N/A	Analog	SE	
DPIPO_HPD DPIP1_HPD DPIP2_HPD DPIP3_HPD	DisplayPort* Hot Plug Detect	O	DP*	SE	

Note: HPD signal require on-board level shifter

12.8 USB Type-C Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
TCP[2:0]_TX_P[1:0] TCP[2:0]_TX_N[1:0]	TX Data Lane.	O	TCP	Diff	All Processor Lines
TCP[3]_TX_P[1:0] TCP[3]_TX_N[1:0]	TX Data Lane.	O	TCP	Diff	UP3/H Processor Lines
TCP[2:0]_TXRX_P[1:0] TCP[2:0]_TXRX_N[1:0]	RX Data Lane, also serves as the secondary TX data lane.	I/O	TCP	Diff	All Processor Lines
TCP[3]_TXRX_P[1:0] TCP[3]_TXRX_N[1:0]	RX Data Lane, also serves as the secondary TX data lane.	I/O	TCP	Diff	UP3/H Processor Lines
TCP[2:0]_AUX_P TCP[2:0]_AUX_N	Common Lane AUX-PAD.	I/O	TCP	Diff	All Processor Lines
TCP[3]_AUX_P TCP[3]_AUX_N	Common Lane AUX-PAD.	I/O	TCP	Diff	UP3/H Processor Lines
TC_RCOMP_P TC_RCOMP_N	Type-C Resistance Compensation.	I	Analog	Diff	All Processor Lines
TCP0_MBIAS_RCOMP	Type-C Resistance Compensation.	I	Analog	Diff	All Processor Lines

12.9 MIPI* CSI-2 Interface Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
CSI_A_DP[1:0] CSI_A_DN[1:0]	CSI-2 Ports Data lane	I	DPHY	Diff	H Processor Line
CSI_B_DP[3:0] CSI_B_DN[3:0]	CSI-2 Ports Data lane	I	DPHY	Diff	All Processor Lines
CSI_C_DP[3:0] CSI_C_DN[3:0]					UP3/UP4 Processor Lines
CSI_E_DP[1:0] CSI_E_DN[1:0]					UP3/UP4 Processor Lines
CSI_F_DP[3:0] CSI_F_DN[3:0]					UP3/UP4 Processor Lines
CSI_G_DP[1:0] CSI_G_DN[1:0]	CSI-2 Ports Data lane	I	DPHY	Diff	UP4 Processor Line
CSI_H_DP[0] CSI_H_DN[0]					UP4 Processor Line
CSI_A_CLK_P CSI_A_CLK_N	CSI 2 Port A Clock lane	I	DPHY	Diff	H Processor Line
CSI_B_CLK_P CSI_B_CLK_N	CSI-2 Ports B-C Clock lane	I	DPHY	Diff	All Processor Lines
CSI_C_CLK_P CSI_C_CLK_N					UP3/UP4 Processor Lines
CSI_E_CLK_P CSI_E_CLK_N	CSI-2 Ports E-H Clock lane	I	DPHY	Diff	UP3/UP4 Processor Lines
CSI_F_CLK_P CSI_F_CLK_N					UP3/UP4 Processor Lines
CSI_G_CLK_P CSI_G_CLK_N	CSI-2 Ports E-H Clock lane	I	DPHY	Diff	UP4 Processor Line
CSI_H_CLK_P CSI_H_CLK_N					UP4 Processor Line
CSI_RCOMP	CSI Resistance Compensation	N/A	N/A	SE	All Processor Lines

12.10 Processor Clocking Signals

Table 12-5. Processor Clocking Signals

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
BCLK_P BCLK_N	100 MHz Differential bus clock input to the processor	I	CMOS	Diff	H Processor Line
CLK_XTAL_P CLK_XTAL_N	38.4 MHz Differential bus clock input to the processor	I	CMOS	Diff	
PCI_BCLKP PCI_BCLKN	100 MHz Clock for PCI Express* logic	I	CMOS	Diff	
RTC_CLK	Real time counter	I	CMOS	SE	

12.11 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	All Processor Lines
PROC_PRDY#	Probe Mode Ready: PROC_PRDY# is a processor output used by debug tools to determine processor debug readiness.	O	OD	SE	All Processor Lines
PROC_PREQ#	Probe Mode Request: PROC_PREQ# is used by debug tools to request debug operation of the processor.	I	GTL	SE	All Processor Lines
PROC_TCK	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	All Processor Lines
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	All Processor Lines
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	All Processor Lines
PROC_TMS	Test Mode Select: A JTAG specification support signal used by debug tools.	I	GTL	SE	All Processor Lines
PROC_TRST#	Test Reset: Resets the Test Access Port (TAP) logic. This signal should be driven low during power on Reset.	I	GTL	SE	All Processor Lines

Note: Refer to the electric spec in DC specification data for more details on the Buffer type power spec requirement. For the buffer type, refer to CMOS [Section 13.2.1.12, “CMOS DC Specifications”](#) data. For the buffer type, refer to TGTL in table, [Section 13.2.1.13, “GTL and OD DC Specification”](#) for electric DC specification data.

12.12 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	All Processor Lines
PECI	Platform Environment Control Interface: A serial sideband interface to the processor. It is used primarily for thermal, power, and error management. Details regarding the PEFI electrical specifications, protocols and functions can be found in the RS-Platform Environment Control Interface (PECI) Specification, Revision 3.0.	I/O	PECI Async	SE	All Processor Lines
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	GTL I OD O	SE	All Processor Lines
THERMTRIP#	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	All Processor Lines

12.13 Processor Power Rails

Table 12-6. Processor Power Rails Signals (Sheet 1 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCC _{IN}	Input FIVR, Processor IA Cores and Graphics Power Rail.	I	Power	-	All Processor Lines
VCC _{IN_AUX}	Input FIVR, SA and PCH components.	I	Power	-	All Processor Lines
VCC _{1P8A}	Power rail supporting PCIe PHY and DPIN PHY power supply	I	Power	-	H Processor Line
V _{DD2}	System Memory power rail	I	Power	-	All Processor Lines
VCC _{ST}	Sustain voltage for processor standby modes	I	Power	-	All Processor Lines
VCC _{STG}	Gated sustain voltage for processor standby modes	I	Power	-	All Processor Lines
VCC _{IN_SENSE} VCC _{IN_AUX_VCCSENSE}	Isolated, low impedance voltage sense pins. They can be used to sense or measure voltage near the silicon.	I	PWR SENSE	SE	All Processor Lines

Table 12-6. Processor Power Rails Signals (Sheet 2 of 2)

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
VCC _{IN_AUX} _VSSSENSE VSS _{IN} _SENSE	Isolated, low impedance reference ground sense pins. They can be used to sense or measure the reference ground to the adequate voltage rail near the silicon.	I	GND_SENSE	SE	All Processor Lines

Table 12-7. Processor Pull-up Power Rails Signals

Signal Name	Description	Dir.	Type	Availability
VCC _{STG_OUT_LGC}	Reference power rail for all Legacy Signals Pull-up on platform.	O	Reference Power	UP3 Processor Line
VCC _{ST_OUT}	Reference power rail for Legacy Signals Pull-up on platform	O	Reference Power	UP4 Processor Line
VCC _{STG_OUT}	Reference power rail for JTAG/ PROCHOT Signals Pull-up on platform, Supplier of the FPGM power rail	O	Reference Power	UP4/H Processor Lines
VCC _{STG_OUT}	VCC _{STG_OUT} Power rail	O	Power	UP3 Processor Line
VCC _{IO_OUT}	Reference power rail for all Debug/ Config Signals Pull-up on platform.	O	Reference Power	All Processor Lines

12.14 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
- **_NCTF**: These signals are non-critical to function and should not be connected.
- **RSVD_2.2K_PD**: This signal should be connected to GND via 2.2K ohm resistor.

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 [Section 12-8, “GND, RSVD, and NCTF Signals”](#).

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. Resistor values should be within $\pm 20\%$ of the impedance of the baseboard trace.

Table 12-8. GND, RSVD, and NCTF Signals

Signal Name	Description	Availability
Vss	Ground: Processor ground node	All Processor Lines
RSVD	Reserved: All signals that are RSVD should not be connected on the board.	All Processor Lines
RSVD_NCTF	Reserved Non-Critical To Function: RSVD_NCTF should not be connected on the board.	All Processor Lines
RSVD_TP	Test Point: Intel recommends to route each RSVD_TP to an accessible test point. Intel may require these test points for platform specific debug. Leaving these test points inaccessible could delay debug by Intel.	All Processor Lines
RSVD_2.2K_PD	RSVD_2.2K_PD - must connect to 2.2K Ohm resistor 1% to GND (pull down)	H Processor Line

12.15 Processor Internal Pull-Up / Pull-Down Terminations

Signal Name	Pull Up/Pull Down	Rail
BPM#[3:0]	Pull Up/Pull Down	VCC _{IO_OUT}
PROC_PREQ#	Pull Up	VCC _{STG}
PROC_TDI	Pull Up	VCC _{STG}
PROC_TMS	Pull Up	VCC _{STG}
PROC_TRST#	Pull Down	VCC _{STG}
PROC_TCK	Pull Down	VCC _{STG}
CFG[17:0]	Pull Up	VCC _{IO_OUT}

§ §

13 Electrical Specifications

13.1 Processor Power Rails

Power Rail	Description	UP3 Processor Line	UP4 Processor Line	H Processor Line
V _{CCIN}	Input FIVR ¹ , Processor IA Cores And Graphic Power Rail	SVID	SVID	SVID
V _{CCIN_AUX} ³	Input FIVR ¹ , SA And PCH components	PCH VID	PCH VID	PCH VID
V _{CC1P8A} ⁵	PCIe PHY and DPIN PHY power supply	N/A	N/A	Fixed
V _{CCST} ⁴	Sustain Power Rail	Fixed	Fixed	Fixed
V _{CCSTG} ⁴	Sustain Gated Power Rail	Fixed	Fixed	Fixed
V _{DD2}	Integrated Memory Controller Power Rail	Fixed (Memory technology dependent)	Fixed (Memory technology dependent)	Fixed (Memory technology dependent)

Notes:

1. FIVR = Fully Integrated Voltage Regulator, refer to 13.1.2, *Full Integrated Voltage Regulator (FIVR)*
2. For details regarding each rail's VR.
3. V_{CCIN_AUX} has a few discrete voltages defined by PCH VID.
4. V_{CCST} and V_{CCSTG} these rails are not connect to external voltage regulator moreover they are connected to the VCC1P05 power rail (from PCH) through a power gate.
5. Power rail exist only in H SKU.

13.1.1 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.

13.1.2 Full Integrated Voltage Regulator (FIVR)

Due to the integration of platform voltage regulators into the processor, the processor has one main voltage rail (V_{CCIN}), the PCH has one main voltage rail (V_{CCIN_AUX}) and a voltage rail for the memory interface (V_{DD2}).

The voltage rail V_{CCIN} will supply the integrated voltage regulators which in turn will regulate to the appropriate voltages for the Cores, cache, System Agent, TCSS and graphics. This integration allows the processor to better control on-die voltages to optimize between performance and power savings. The V_{CCIN} rail will remain a VID-based voltage with a loadline similar to the core voltage rail in previous processors.

13.1.3 V_{CC} Voltage Identification (VID)

Intel processors/chipsets 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/chipset 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.

For VID coding and further information, refer to the *IMVP9 PWM Specification* and *Serial VID (SVID) Protocol Specification*.

13.2 DC Specifications

The processor DC specifications in this section are defined at the processor signal pins, unless noted otherwise. For pin listing, refer to 11th Generation Intel[®] Core[™] Processor Line Package Ballout Mechanical Specification.

- The DC specifications for the LPDDR4x/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 1 MHz. Refer additional guidance for each rail.

13.2.1 Processor Power Rails DC Specifications

13.2.1.1 V_{CCIN} DC Specifications

Table 13-1. Processor V_{CCIN} Active and Idle Mode DC Voltage and Current Specifications (Sheet 1 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
Operating Voltage	Voltage Range for Processor Operating Mode	UP3/ UP3-Refresh/UP4/H35 Processor Lines	0	—	2.0	V	1,2,3, 7,12
Operating Voltage	Voltage Range for Processor Operating Mode	H Processor Line	0	—	2.05	V	1,2,3, 7,12
I _{ccMAX} (UP3 Processor)	Maximum Processor I _{cc}	UP3 Processor UP3-Refresh/Line (28W) 4-Core GT2	—	—	65	A	4,5,6,7,11
I _{ccMAX} (H35 Processor)	Maximum Processor I _{cc}	H35 Processor Line (35W) 4-Core GT2	—	—	65	A	4,5,6,7,11
I _{ccMAX} (UP3 Processor)	Maximum Processor I _{cc}	UP3 Processor Line (28W) 2-Core GT2	—	—	35	A	4,5,6,7,11

Table 13-1. Processor V_{CC1N} Active and Idle Mode DC Voltage and Current Specifications (Sheet 2 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
I _{ccMAX} (UP3 Processor)	Maximum Processor I _{cc}	UP3 Processor Line (15W) Celeron 2-Core GT2	—	—	35	A	4,5,6,7,11
I _{ccMAX} (UP4 Processor)	Maximum Processor I _{cc}	UP4 Processor Line (9W) 4-Core GT2	—	—	42	A	4,5,6,7,11
I _{ccMAX} (UP4 Processor)	Maximum Processor I _{cc}	UP4 Processor Line (9W) 2-Core GT2	—	—	25	A	4,5,6,7,11
I _{ccMAX} (H Processor)	Maximum Processor I _{cc}	H Processor Line (45W) 8-Core GT1	—	—	105	A	4,5,6,7,11
		H Processor Line (45W) 6-Core GT1	—	—	74	A	4,5,6,7,11
I _{ccTDC}	Thermal Design Current (TDC) for processor V _{CC1N} Rail	—	—	—	—	A	9
TOB _{VCC}	Voltage Tolerance	PS0, PS1	—	—	±20	mV	3, 6, 8
		PS2, PS3	—	—	±35		
Ripple	Ripple Tolerance	PS0, PS1	—	—	±15	mV	3, 6, 8
		PS2, PS3	—	—	±30		
DC_LL	Loadline slope within the VR regulation loop capability	UP3/UP3-Refresh/ H35 Processor Lines	0	—	2.0	mΩ	10,13,14
DC_LL	Loadline slope within the VR regulation loop capability	UP4 Processor Line	0	—	2.0	mΩ	10,13,14
DC_LL	Loadline slope within the VR regulation loop capability	H Processor Line (45W) 8-Core GT1	0	1.5	1.7	mΩ	10,13,14,15
DC_LL	Loadline slope within the VR regulation loop capability	H Processor Line (45W) 6 Core GT1	0	1.5	—	mΩ	10,13,14,15
AC_LL3	AC Loadline 3 (<1 MHz)	UP3/UP3-Refresh/ H35 4 Core Processor Lines	—	—	4.4	mΩ	10,13,14
AC_LL3	AC Loadline 3 (<1 MHz)	UP3 2 Core Processor Line	—	—	8.0	mΩ	10,13,14
AC_LL3	AC Loadline 3 (<1 MHz)	UP4 2 Core Processor Line	—	—	8.0	mΩ	10,13,14
AC_LL3	AC Loadline 3 (<1 MHz)	UP4 4 Cores Processor Line	—	—	4.7	mΩ	10,13,14
AC_LL3	AC Loadline 3 (<1 MHz)	H Processor Line (45W) 8-Core GT1	—	—	1.7	mΩ	10,13,14

Table 13-1. Processor Vcc_{IN} Active and Idle Mode DC Voltage and Current Specifications (Sheet 3 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
AC_LL3	AC Loadline 3 (<1 MHz)	H Processor Line (45W) 6 Core GT1	—	—	2.0	mΩ	10,13,14
T_OVS_TDP_MAX	Maximum Overshoot time TDP/virus mode	—	—	—	500	μs	
V_OVS TDP_MAX/ virus_MAX	Maximum Overshoot at TDP/virus mode	—	—	—	10	%	

Notes:

- 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.
- 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 Speed-step Technology, or low-power states).
- The voltage specification requirements are measured across Vcc_SENSE and Vss_SENSE as near as possible to the processor. The measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Ω minimum impedance. The maximum length of the ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
- Processor Vcc_{IN} VR to be designed to electrically support this current.
- Processor Vcc_{IN} VR to be designed to thermally support this current indefinitely.
- Long term reliability cannot be assured if tolerance, ripple, and core noise parameters are violated.
- Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
- PSx refers to the voltage regulator power state as set by the SVID protocol.
- LL measured at sense points.
- Typ column represents Icc_{MAX} for commercial application it is NOT a specification - it's a characterization of limited samples using limited set of benchmarks that can be exceeded.
- Operating voltage range in steady state.
- LL spec values should not be exceeded. If exceeded, power, performance and reliability penalty are expected.
- 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.
- H-processor DC LL spec value is 1.5 mOhm, DC LL can be lower or equal to AC LL=1.7 mOhm.

13.2.1.2 Vcc_{IN_AUX} DC Specifications

Table 13-2. Vcc_{IN_AUX} Supply DC Voltage and Current Specifications (Sheet 1 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
VCCIN _{AUX}	Voltage Range	UP3/UP3-Refresh/ H35 Processor Lines	—	1.8	—	V	1,2,3,7
VCCIN _{AUX}	Voltage Range	UP4 Processor Line	—	1.65	1.8	V	1,2,3,7
VCCIN _{AUX}	Voltage Range	H Processor Line	—	1.8	—	V	1,2,3,7
Icc _{MAX}	Maximum Vcc _{IN_AUX} Icc	UP3/UP3-Refresh Processor Line (28W) 4/2-Core GT2	0	—	27	A	1,2
Icc _{MAX}	Maximum Vcc _{IN_AUX} Icc	H35 Processor Line (35W) 4-Core GT2	0	—	27		

Table 13-2. VccIN_AUX Supply DC Voltage and Current Specifications (Sheet 2 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
I _{CCMAX}	Maximum VccIN_AUX I _{cc}	UP3 Processor Line (15W) Celeron 2-Core GT2	0	—	27	A	1,2
I _{CCMAX}	Maximum VccIN_AUX I _{cc}	UP4 Processor Line (9W) 4/2-Core GT2	0	—	25	A	1,2
I _{CCMAX}	Maximum VccIN_AUX I _{cc}	H Processor Line (45W) 8/6-Core GT1	0	—	42	A	1,2,8
TOB _{VCC}	Voltage Tolerance Budget	UP3/UP3-Refresh/ H35 Processor Lines	—	—	AC+DC: +5/-10	%	1,3,6
TOB _{VCC}	Voltage Tolerance Budget	UP4 Processor Line	—	—	DC Min: -4 AC+DC: ± 7.5	%	1,3,6
TOB _{VCC}	Voltage Tolerance Budget	H Processor Line	—	—	AC+DC: +5/-10	%	1,3,6
VOS	Maximum Overshoot	UP4 Processor Line	—	—	1.89	V	2,6
TVOS	Maximum Overshoot	UP4 Processor Line	—	—	5	us	2,6
VOS	Maximum Overshoot	UP3/UP3-Refresh/ H35 Processor	—	—	1.95	V	2,6
TVOS	Maximum Overshoot	UP3/H35 Processor Lines	—	—	5	us	2,6
VOS	Maximum Overshoot	H Processor Line	—	—	1.89	V	2,6
TVOS	Maximum Overshoot	H Processor Line	—	—	5	us	2,6
AC_LL	AC Loadline	UP4 Processor Line	—	—	• 10 KHz-1 Mhz: 5.8	mΩ	4,5
		UP3/UP3-Refresh/ H35 Processor Lines	—	—	• 10 KHz-1 Mhz: 6.9	mΩ	4,5
AC_LL	AC Loadline	H Processor Line	—	—	4.0	mΩ	4,5
DC_LL	DC Loadline	H Processor Line	0	—	2.1	mΩ	4,5
DC_LL	DC Loadline	UP3/UP3-Refresh/ H35 Processor Lines	0	—	3.3	mΩ	4,5
		UP4 Processor Line	0	—	2.1	mΩ	4,5

Table 13-2. VccIN_AUX Supply DC Voltage and Current Specifications (Sheet 3 of 3)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	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. Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.							
3. The voltage specification requirements are measured across Vcc_SENSE and Vss_SENSE as near as possible to the processor. The measurement needs to be performed with a 20 MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 Ω minimum impedance. The maximum length of the 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. LL measured at sense points. LL specification values should not be exceeded. If exceeded, power, performance, and reliability penalty are expected.							
5. The LL values are for reference. Must still need to meet the voltage tolerance specification.							
6. Voltage Tolerance budget values Include ripples							
7. VccIN_AUX is having few point of voltage define by PCH VID							
8. VCCIN_AUX ICCMAX includes both CPU and PCH, CPU will consume 27A and PCH 15A.							

13.2.1.3 V_{DD2} DC Specifications

Table 13-3. Memory Controller (V_{DD2}) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
V _{DD2} (LPDDR4x)	Processor I/O supply voltage for LPDDR4x	UP3/UP3-Refresh/UP4 Processor Lines	Typ-5%	1.115	Typ+5%	V	3,4,5
V _{DD2} (DDR4)	Processor I/O supply voltage for DDR4	UP3/UP3-Refresh/H Processor Lines	Typ-5%	1.2	Typ+5%	V	3,4,5
TOB _{VDD2}	VDD2 Tolerance	All	V _{DD2} MIN < AC+DC < V _{DD2} MAX			V	3,4,6
ICCMAX_VDD2 (LPDDR4x)	Maximum Current for V _{DD2} Rail (LPDDR4x)	UP3/UP3-Refresh/UP4 Processor Lines	—	—	1.5	A	2
ICCMAX_VDD2 (DDR4)	Maximum Current for V _{DD2} Rail (DDR4)	UP3/UP3-Refresh Processor Line	—	—	1.5	A	2
ICCMAX_VDD2 (DDR4)	Maximum Current for V _{DD2} Rail (DDR4)	H Processor Line	—	—	4.3		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 on package pins as near as possible to the processor with an oscilloscope set to 100 MHz bandwidth, 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.							
6. For Voltage less than 1 V TOB will be 50 mV.							

13.2.1.4 Vcc_{ST} DC Specifications

Table 13-4. Vcc Sustain (Vcc_{ST}) Supply DC Voltage and Current Specifications (Sheet 1 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes ^{1,2}
Vcc _{ST}	Processor Vcc Sustain supply voltage	UP3/UP3-Refresh/H35 Processor Lines	—	1.025	—	V	3
Vcc _{ST}	Processor Vcc Sustain supply voltage	UP4/H Processor Lines	—	1.065	—	V	3,6,7

Table 13-4. Vcc Sustain (Vcc_{ST}) Supply DC Voltage and Current Specifications (Sheet 2 of 2)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes ^{1,2}
TOB _{ST}	Vcc _{ST} Tolerance	All	AC+DC: ± 5			%	3,5,7
Icc _{MAX_ST}	Maximum Current for Vcc _{ST}	UP3/UP3-Refresh/H35 Processor Lines	—	—	500	mA	4
Icc _{MAX_ST}	Maximum Current for Vcc _{ST}	UP4 Processor Line	—	—	300	mA	4
Icc _{MAX_ST}	Maximum Current for Vcc _{ST}	H Processor Line	—	—	970	mA	4

Notes:

- 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.
- Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
- The voltage specification requirements are measured on package pins as near as possible to the processor with an oscilloscope set to 100 MHz bandwidth, 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.
- The maximum Icc_{MAX_ST} specification is preliminary and based on initial pre-silicon estimation and is subject to change.
- For Voltage less than 1 V TOB will be 50 mV.
- VCCST without PG will have typical of 1.025V, some collateral may indicate VCCST = 1.025 V which present the typical voltage without PG.
- VCCST momentarily can go to 1.15V during certain scenarios, there is no side effect.

Table 13-5. 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 gated supply voltage	UP3/UP3-Refresh/H35 Processor Lines	—	1.025	—	V	3,7
Vcc _{STG}	Processor Vcc Sustain gated supply voltage	UP4/H Processor Lines	—	1.065	—	V	3,6,7
TOB _{STG}	Vcc _{STG} Tolerance	All	AC+DC: ± 5%			%	3,5,7
Icc _{MAX_STG}	Maximum Current for Vcc _{STG}	UP3/UP3-Refresh/H35 Processor Lines	—	—	300	mA	4
Icc _{MAX_STG}	Maximum Current for Vcc _{STG}	UP4 Processor Line	—	—	250	mA	4
Icc _{MAX_STG}	Maximum Current for Vcc _{STG}	H Processor Line	—	—	340	mA	4

Notes:

- 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.
- Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
- The voltage specification requirements are measured on package pins as near as possible to the processor with an oscilloscope set to 100 MHz bandwidth, 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.
- The maximum Icc_{MAX_STG} specification is preliminary and based on initial pre-silicon estimation and is subject to change.
- For Voltage less than 1 V TOB will be 50 mV.
- VCCSTG without PG will have typical of 1.025 V, some collateral may indicate VCCSTG = 1.025 V which present the typical voltage without PG.
- VCCSTG momentarily can go to 1.15V during certain scenarios, there is no side effect.

13.2.1.5 Vcc_{1P8A} DC Specifications

Table 13-6. Vcc1P8A Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes 1,2
VCC1P8A	Processor PCIE/DPIN PHY supply voltage	H Processor Line	—	1.8	—	V	3
TOB1P8A	VCC1P8A Tolerance	H Processor Line	AC+DC: ± 5			%	3,5
ICC _{MAX_1P8A}	Maximum Current for VCC1P8A	H Processor Line	—	—	500	mA	4

Notes:

- 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.
- Long term reliability cannot be assured in conditions above or below Maximum/Minimum functional limits.
- The voltage specification requirements are measured on package pins as near as possible to the processor with an oscilloscope set to 100 MHz bandwidth, 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.
- The maximum ICC_{MAX_1P8A} specification is preliminary and based on initial pre-silicon estimation and is subject to change.
- For Voltage less than 1 V TOB will be 50 mV.

13.2.1.6 DDR4 DC Specifications

Table 13-7. DDR4 Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	UP3/H-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
V _{IL}	Input Low Voltage	—	0.75*V _{DD2}	0.68*V _{DD2}	V	2, 3, 4
V _{IH}	Input High Voltage	0.82*V _{DD2}	0.75*V _{DD2}	—	V	2, 3, 4
R _{ON_UP} (DQ)	Data Buffer pull-up Resistance	30	—	50	Ω	5,12
R _{ON_DN} (DQ)	Data Buffer pull-down Resistance	30	—	50	Ω	
R _{ODT} (DQ)	On-die termination equivalent resistance for data signals	40	—	200	Ω	6, 12
V _{ODT} (DC)	On-die termination DC working point (driver set to receive mode)	0.45*V _{DD2}	—	0.85*V _{DD2}	V	12
R _{ON_UP} (CK)	Clock Buffer pull-up Resistance	25	—	45	Ω	5, 12
R _{ON_DN} (CK)	Clock Buffer pull-down Resistance	25	—	45	Ω	5, 12
R _{ON_UP} (CMD)	Command Buffer pull-up Resistance	25	—	45	Ω	5, 12
R _{ON_DN} (CMD)	Command Buffer pull-down Resistance	25	—	45	Ω	5, 12
R _{ON_UP} (CTL)	Control Buffer pull-up Resistance	25	—	45	Ω	5, 12
R _{ON_DN} (CTL)	Control Buffer pull-down Resistance	25	—	45	Ω	5, 12
R _{ON_UP} (SM_PG_CNTL1)	System Memory Power Gate Control Buffer Pull-up Resistance	45	—	125	Ω	—
R _{ON_DN} (SM_PG_CNTL1)	System Memory Power Gate Control Buffer Pull-down Resistance	40	—	130	Ω	—
I _{LI}	Input Leakage Current (DQ, CK) 0 V 0.2* V _{DD2} 0.8* V _{DD2}	—	—	1.1	mA	—
DDR0_VREF_DQ DDR1_VREF_DQ	VREF output voltage	Trainable	V _{DD2} /2	Trainable	V	—
SM_RCOMP[0]	Command COMP Resistance	99	100	101	Ω	8
SM_RCOMP[1]	Data COMP Resistance	99	100	101	Ω	8
SM_RCOMP[2]	ODT COMP Resistance	99	100	101	Ω	8

Table 13-7. DDR4 Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	UP3/H-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
Notes:						
<ol style="list-style-type: none"> 1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. Timing specifications only depend on the operating frequency of the memory channel and not the maximum rated frequency 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_{OH} may experience excursions above V_{DD2}. However, input signal drivers should comply with the signal quality specifications. 5. Pull up/down resistance after compensation (assuming $\pm 5\%$ COMP inaccuracy). Note that BIOS power training may change these values significantly based on margin/power trade-off. 6. ODT values after COMP (assuming $\pm 5\%$ inaccuracy). BIOS MRC can reduce ODT strength towards 7. The minimum and maximum values for these signals are programmable by BIOS to one of the two sets. 8. SM_RCOMP[x] resistance should be provided on the system board with 1% resistors. SM_RCOMP[x] resistors are to VSS. Values are pre-silicon estimations and are subject to change. 9. SM_DRAMPWROK must have a maximum of 15 ns rise or fall time over $V_{DD2} * 0.30 \pm 100$ mV and the edge must be monotonic. 10. SM_VREF is defined as $V_{DD2}/2$ for DDR4 11. R_{ON} tolerance is preliminary and might be subject to change. 12. Maximum-minimum range is correct but center point is subject to change during MRC boot training. 13. Processor may be damaged if V_{IH} exceeds the maximum voltage for extended periods. 						

13.2.1.7 LPDDR4x DC Specifications

Table 13-8. LPDDR4x Signal Group DC Specifications

Symbol	Parameter	UP3/UP4-Processor Line			Units	Notes ¹
		Minimum	Typical	Maximum		
V _{IL}	Input Low Voltage	—	0.2*V _{DD2}	0.08*V _{DD2}	V	2, 3, 4
V _{IH}	0 = Input High Voltage	0.35*V _{DD2}	0.2*V _{DD2}	—	V	2, 3, 4
R _{ON_UP} (DQ)	Data Buffer pull-up Resistance	30	—	50	Ω	5,12
R _{ON_DN} (DQ)	Data Buffer pull-down Resistance	30	—	50	Ω	5,12
R _{ODT} (DQ)	On-die termination equivalent resistance for data signals	40	—	200	Ω	6, 12
V _{ODT} (DC)	On-die termination DC working point (driver set to receive mode)	0.1*V _{DD2}	—	0.3*V _{DD2}	V	10
R _{ON_UP} (CK)	Clock Buffer pull-up Resistance	30	—	45	Ω	5, 12
R _{ON_DN} (CK)	Clock Buffer pull-down Resistance	30	—	45	Ω	5, 12
R _{ON_UP} (CMD)	Command Buffer pull-up Resistance	30	—	45	Ω	5, 12
R _{ON_DN} (CMD)	Command Buffer pull-down Resistance	30	—	45	Ω	5, 12
R _{ON_UP} (CTL)	Control Buffer pull-up Resistance	30	—	45	Ω	5, 12
R _{ON_DN} (CTL)	Control Buffer pull-down Resistance	30	—	45	Ω	5, 12
R _{ON_UP} (SM_VTT_CTL1)	System Memory Power Gate Control Buffer Pull-up Resistance	N/A	—	N/A	Ω	N/A
R _{ON_DN} (SM_VTT_CTL1)	System Memory Power Gate Control Buffer Pull-down Resistance	N/A	—	N/A	Ω	N/A
I _{LI}	Input Leakage Current (DQ, CK) 0 V 0.2* V _{DD2} 0.8* V _{DD2}	—	—	1.1	mA	—
DDR0_VREF_DQ DDR1_VREF_DQ DDR_VREF_CA	VREF output voltage	Trainable			V	—
SM_RCOMP[0]	Command COMP Resistance	99	100	101	Ω	8
SM_RCOMP[1]	Data COMP Resistance	99	100	101	Ω	8
SM_RCOMP[2]	ODT COMP Resistance	99	100	101	Ω	8

Notes:

1. Unless otherwise noted, all specifications in this table apply to all processor frequencies. Timing specifications only depend on the operating frequency of the memory channel and not the maximum rated frequency
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_{OH} may experience excursions above V_{DD2}. However, input signal drivers should comply with the signal quality specifications.
5. Pull up/down resistance after compensation (assuming ±5% COMP inaccuracy). Note that BIOS power training may change these values significantly based on margin/power trade-off.
6. ODT values after COMP (assuming ±5% inaccuracy). BIOS MRC can reduce ODT strength towards
7. The minimum and maximum values for these signals are programmable by BIOS to one of the two sets.
8. SM_RCOMP[x] resistance should be provided on the system board with 1% resistors. SM_RCOMP[x] resistors are to VSS. Values are pre-silicon estimations and are subject to change.
9. SM_DRAMPWROK must have a maximum of 15 ns rise or fall time over V_{DD2} * 0.30 ±100 mV and the edge must be monotonic.
10. SM_VREF is defined as V_{DD2}/2 for DDR4/LPDDR4x
11. R_{ON} tolerance is preliminary and might be subject to change.
12. Maximum-minimum range is correct but center point is subject to change during MRC boot training.
13. Processor may be damaged if V_{IH} exceeds the maximum voltage for extended periods.

13.2.1.8 PCI Express* Graphics (PEG) DC Specifications

Symbol	Parameter	Min	Typ	Max	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 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 VCCIO_OUT. 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.

13.2.1.9 Digital Display Interface (DDI) DC Specifications

Table 13-10. DSI HS Transmitter DC Specifications

Parameter	Description	Minimum	Nom	Max	Units	Notes ¹
V _{CMTX}	HS transmit static common-mode voltage	150	200	250	mV	1
ΔV _{CMTX(1,0)}	V _{CMTX} mismatch when output is Differential-1 or Differential-0	—	—	5	mV	2
V _{OD}	HS transmit differential voltage	140	200	270	mV	1
ΔV _{OD}	V _{OD} mismatch when output is Differential-1 or Differential-0	—	—	14	mV	2
V _{OHS}	HS output high voltage	—	—	360	mV	1
Z _{OS}	Single ended output impedance	40	50	62.5	Ω	
ΔZ _{OS}	Single ended output impedance mismatch	—	—	10	%	

Notes:

1. Value when driving into load impedance anywhere in the ZID range.
2. A transmitter should minimize ΔV_{OD} and ΔV_{CMTX(1,0)} in order to minimize radiation, and optimize signal integrity

Table 13-11. DSI LP Transmitter DC Specifications (Sheet 1 of 2)

Parameter	Description	Minimum	Nominal	Maximum	Units	Notes ¹
V _{OH}	Thevenin output high level	1.1	1.05	1.3	V	1
		0.95	—	1.3	V	2
V _{OL}	Thevenin output low level	-50	—	50	mV	
Z _{OLP}	Output impedance of LP transmitter	110	—	—	Ω	3
V _{pin}	Pin signal voltage range	-50	—	1350	mV	
I _{LEAK}	Pin Leakage current	-10	—	10	uA	4
V _{GNDSH}	Ground shift	-50	—	50	mV	
V _{pin(ABSMAX)}	Transient pin voltage level	-0.15	—	1.45	V	6
T _{Vpin(ABSMAX)}	Maximum transient time above VPIN(MAX) or below VPIN(MIN)	—	—	20	ns	5

Table 13-11. DSI LP Transmitter DC Specifications (Sheet 2 of 2)

Parameter	Description	Minimum	Nominal	Maximum	Units	Notes ¹
Notes: 1. Applicable when the supported data rate \leq 1.5 Gbps. 2. Applicable when the supported data rate $>$ 1.5 Gbps. 3. Though no maximum value for ZOLP is specified, the LP transmitter output impedance shall ensure the TRLP/TFLP specification is met. 4. The voltage overshoot and undershoot beyond the VPIN is only allowed during a single 20ns window after any LP-0 to LP-1 transition or vice versa. For all other situations it must stay within the VPIN range. 5. This value includes ground shift.						

Table 13-12. Digital Display Interface Group DC Specifications (DP/HDMI)

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes ¹
V _{IL}	Aux Input Low Voltage	—	—	0.8	V	
V _{IH}	Aux Input High Voltage	2.25	—	3.6	V	
V _{OL}	DDIx_TX[3:0] Output Low Voltage TCPx_TX[3:0] Output Low Voltage	—	—	0.25*V _{CCIO_OUT}	V	1,2
V _{OH}	DDIx_TX[3:0] Output high Voltage TCPx_TX[3:0] Output high Voltage	0.75*V _{CCIO_OUT}	—	—	V	1,2
ZTX-DIFF-DC	DC Differential Tx Impedance	100	—	120	Ω	
Notes: 1. V _{CCIO_OUT} depends on segment. 2. V _{OL} and V _{OH} levels depends on the level chosen by the Platform.						

Table 13-13. DP-IN Group DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes ¹
V _{OL}	DPIPx HPD low level	0	—	-	V	1,2
V _{OH}	DPIPx HPD high level	—	—	V _{CCIO_OUT}	V	1,2
Notes: 1. V _{CCIO_OUT} depends on segment. 2. x is referring to ports 0-3.						

13.2.1.10 embedded DisplayPort* (eDP*) DC Specification

Symbol	Parameter	Minimum	Typical	Maximum	Units
V _{OL}	eDP_DISP_UTIL Output Low Voltage	—	—	0.1*V _{CCIO_OUT}	V
V _{OH}	eDP_DISP_UTIL Output High Voltage	0.9*V _{CCIO_OUT}	—	—	V
R _{UP}	eDP_DISP_UTIL Internal pull-up	45	—	—	Ω
R _{DOWN}	eDP_DISP_UTIL Internal pull-down	45	—	—	Ω
Notes: 1. COMP resistance is to V _{COMP_OUT} . 2. eDP_RCOMP resistor should be provided on the system board.					

13.2.1.11 MIPI* CSI-2 D-Phy Receiver DC Specifications

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes
V _{CMRX(DC)}	Common-mode voltage HS receive mode	70	—	330	mV	1,2

Symbol	Parameter	Minimum	Typical	Maximum	Units	Notes
V _{IDTH}	Differential input high threshold	—	—	70	mV	3
		—	—	40	mV	4
V _{IDTL}	Differential input low threshold	-70	—	—	mV	3
		-40	—	—	mV	4
V _{IHHS}	Single-ended input high voltage	—	—	460	mV	1
V _{ILHS}	Single-ended input low voltage	-40	—	—	mV	1
V _{TERM-EN}	Single-ended threshold for HS termination enable	—	—	450	mV	
Z _{ID}	Differential input impedance	80	100	125	Ω	

Notes:

- Excluding possible additional RF interference of 100 mV peak sine wave beyond 450 MHz
- This table value includes a ground difference of 50 mV between the transmitter and the receiver, the static common-mode level tolerance and variations below 450 MHz
- For devices supporting data rates < 1.5 Gbps.
- For devices supporting data rates > 1.5 Gbps.
- Associated Signals: MIPI* CSI2: Refer to MIPI® Alliance D-PHY Specification 1.2.

13.2.1.12 CMOS DC Specifications

Table 13-14. CMOS Signal Group DC Specifications

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
V _{IL}	Input Low Voltage	—	V _{CC} *0.3	V	2, 5
V _{IH}	Input High Voltage	V _{CC} *0.7	—	V	2, 4, 5
R _{ON}	Buffer on Resistance	20	70	Ω	-
I _{LI}	Input Leakage Current	—	±150	μA	3
V _{HYSTERESIS}	Hysteresis Voltage	0.15*V _{CC}	—	V	-

Notes:

- Unless otherwise noted, all specifications in this table apply to all processor frequencies.
- The V_{CC} referred to in these specifications refers to instantaneous V_{CCST/IO}.
- For V_{IH} between "0" V and V_{CCST}. Measured when the driver is tri-stated.
- V_{IH} may experience excursions above V_{CCST}. However, input signal drivers should comply with the signal quality specifications.
- V_{OH} and V_{OL} need to comply with V_{IL} and V_{IH} specs.

13.2.1.13 GTL and OD DC Specification

Table 13-15. GTL Signal Group and Open Drain Signal Group DC Specifications (Sheet 1 of 2)

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
V _{IL}	Input Low Voltage (TAP, except PROC_JTAG_TCK, PROC_JTAG_TRST#)	—	0.6*V _{CC}	V	2, 5
V _{IH}	Input High Voltage (TAP, except PROC_JTAG_TCK, PROC_JTAG_TRST#)	0.72*V _{CC}	—	V	2, 4, 5
V _{IL}	Input Low Voltage (PROC_JTAG_TCK, PROC_JTAG_TRST#)	—	0.3*V _{CC}	V	2, 5
V _{IH}	Input High Voltage (PROC_JTAG_TCK, PROC_JTAG_TRST#)	0.7*V _{CC}	—	V	2, 4, 5
V _{HYSTERESIS}	Hysteresis Voltage	0.2*V _{CC}	—	V	-
R _{ON}	Buffer on Resistance (TDO)	7	17	Ω	-

Table 13-15. GTL Signal Group and Open Drain Signal Group DC Specifications (Sheet 2 of 2)

Symbol	Parameter	Minimum	Maximum	Units	Notes ¹
V _{IL}	Input Low Voltage (other GTL)	—	0.6*V _{CC}	V	2, 5
V _{IH}	Input High Voltage (other GTL)	0.72*V _{CC}	—	V	2, 4, 5
R _{ON}	Buffer on Resistance (BPM)	12	28	Ω	-
R _{ON}	Buffer on Resistance (other GTL)	16	24	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 V_{CC} referred to in these specifications refers to instantaneous V_{CCST/IO}.
3. For V_{IN} between 0 V and V_{CC}. Measured when the driver is tri-stated.
4. V_{IH} and V_{OH} may experience excursions above V_{CC}. However, input signal drivers should comply with the signal quality specifications.

13.2.1.14 PECl DC Characteristics

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

V_{CCST} nominal levels will vary between processor families. All PECl devices will operate at the V_{CCST} level determined by the processor installed in the system.

Table 13-16. PECl DC Electrical Limits

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	V _{CCST} + 0.15	V	-
V _{hysteresis}	Hysteresis	0.1 * V _{CCST}	—	V	-
V _{IL}	Input Voltage Low- Edge Threshold Voltage	0.275 * V _{CCST}	0.525 * V _{CCST}	V	-
V _{IH}	Input Voltage High- Edge Threshold Voltage	0.550 * V _{CCST}	0.725 * V _{CCST}	V	-
C _{bus}	Bus Capacitance per Node	—	10	pF	-
C _{pad}	Pad Capacitance	0.7	1.8	pF	-
I _{leak000}	leakage current @ 0V	—	0.25	mA	-
I _{leak100}	leakage current @ V _{CCST}	—	0.15	mA	-


Notes:

1. V_{CCST} supplies the PECl interface. PECl behavior does not affect V_{CCST} minimum/maximum specifications.
2. The leakage specification applies to powered devices on the PECl bus.
3. The PECl buffer internal pull up resistance measured at 0.75* V_{CCST}.
4. I_{leak100} represents the worst case leakage at 100 DegC.

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 13-1. Input Device Hysteresis



§ §

14 Package Mechanical Specifications

14.1 Package Mechanical Attributes

The UP3/UP4 Processor Lines use a Flip Chip technology available in a Ball Grid Array (BGA) package. The following table provides an overview of the package mechanical attributes.

Table 14-1. Package Mechanical Attributes

Package	Parameter	UP4-Processor Line	UP3- Processor Line	H-Processor Line
		4/2 Core GT2	4/2 Core GT2	8/6 Core GT1
Package Technology	Package Type	Flip Chip Ball Grid Array	Flip Chip Ball Grid Array	Flip Chip Ball Grid Array
	Interconnect	Ball Grid Array (BGA)	Ball Grid Array (BGA)	Ball Grid Array (BGA)
	Lead Free	Yes	Yes	Yes
	Halogenated Flame Retardant Free	Yes	Yes	Yes
Package Configuration	Solder Ball Composition	SAC405	SAC405	SAC405
	Ball/Pin Count	1598	1449	1787
	Grid Array Pattern	Balls Anywhere	Balls Anywhere	Balls Anywhere
	Land Side Capacitors	Yes (250um maximum height)	Yes (250 um maximum height)	Yes (250 um maximum height)
	Die Side Capacitors	No	No	No
	Die Configuration	2 Dice Multi-Chip Package (MCP)	2 Dice Multi-Chip Package (MCP)	1 Die Single-Chip Package
Package Dimensions	Nominal Package Size	18.5x26.5 mm	25x45.5 mm	50x26.5
	Z	0.963 +/-0.077 mm	1.185 +/-0.096 mm	1.41 +/- 0.109 mm
	Minimum Ball/Pin pitch	0.4 mm	0.65 mm	0.65 mm

14.2 Package Loading and Die Pressure Specifications

Intel has defined the maximum total compressive load limits that can be applied to the package for the following processor lines. These values should not be exceeded by the system design.

14.2.1 Package Loading Specifications

- Considerations should be made to ensure steady state static loading on the packages that does not exceed the limits recommended. Excessive steady state static loading can induce solder ball cracks, especially over a period of time resulting in higher failure rate.
- This static compressive load is not to be exceeded, therefore the tolerance of the package and the tolerances of the thermal solution (including attach mechanism)

should be taken into account when calculating or measuring static load on the package.

- An ideal thermal solution design would apply a load as uniform as possible on all dies in order to optimize thermal performance and minimize mechanical risk.

Table 14-2. Package Loading Specifications

Package	Maximum Static Compressive Load [lbf]	Backing Plate Allowable	Adhesive	Minimum PCB Thickness Assumptions [mm/mils]
UP3-Processor Line	15	Yes ¹	No	0.8/32
	10	Yes ¹	No	0.7/28
UP4-Processor Line	10	No	No	0.7/28
	5	No	Yes ²	0.6/24
H-Processor Line	25	Yes ¹	No	0.8/32
Notes: 1. If using backing plate, recommended maximum back plate thickness is 0.5mm. 2. At a minimum, corner glue is required.				

14.2.2 Die Pressure Specifications

A more relevant metric for concentrated loading is chosen by Intel based on the physics of failure to evaluate die damage risk due to thermal solution enabling

Static Compressive pressure refers to the long term steady state pressure applied to the die from the thermal solution after system assembly is complete.

Transient Compressive pressure refers to the pressure on the dice at any moment during the thermal solution assembly/disassembly procedures. Other system procedures such as repair/rework can also cause high pressure loading to occur on the die and should be evaluated to ensure these limits are not exceeded.

Metric: This metric is pressure over a 2 mm x 2 mm area

Table 14-3. Package Loading Specifications

Package	Static Compressive Pressure ¹ [PSI]	Transient Compressive Pressure ¹ [PSI]
UP4-Processor Line	800	800
UP3-Processor Line	800	800
H-Processor Line	800	800
Note: This is the load and pressure that has been tested by Intel for a single assembly cycle. This metric is a pressure over 2 mm ² (2 mm x 2 mm) area.		

14.3 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.	-5 °C	40 °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% @ 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:				
<ol style="list-style-type: none"> 1. T_{ABSOLUTE 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 devices removed from the moisture barrier bag. 3. Post board attaches storage temperature limits are not specified for non-Intel branded boards. Contact your board manufacturer for storage specifications. 				



15 CPU And Device IDs

15.1 CPUID

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

Table 15-1. CPUID Format

SKU	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]
UP4/UP3/H35	806C1h	Reserved	0000000b	1000b	Reserved	00b	0110b	1100b	0001b
UP3-Refresh	806C2h	Reserved	0000000b	1000b	Reserved	00b	0110b	1100b	0010b
H35-Refresh	806C2h	Reserved	0000000b	1000b	Reserved	00b	0110b	1100b	0010b
H	806D1h	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 Pentium®, Celeron®, or 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.

15.2 PCI Configuration Header

Note: Every PCI-compatible function has a standard PCI configuration header, as shown in [Table 15-2, "PCI Configuration Header"](#). This includes mandatory registers (Bold) to

determine which driver to load for the device. Some of these registers define ID values for the PCI function, which are described in this chapter.

Table 15-2. 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 15-3. Host Device ID (DID0)

Platform	Device ID
UP4 4 Cores	9A12h
UP4 2 Cores	9A02h
UP3 4 Cores	9A14h
UP3 2 Cores	9A04h
H35 4 Cores	9A14h
UP3-Refresh/H35-Refresh 4 Cores	9A1Ah
H 8 Cores	9A36h
H 6 Cores	9A26h

Table 15-4. Other Device ID UP3/UP4/H35/UP3-Refresh/H35-Refresh (Sheet 1 of 2)

Device	Processor Line	Bus / Device / Function	DID
DTT	All	0 / 4 / 0	9A03h
IPU	UP4, UP3	0 / 5 / 0	9A19h
PEG60	All	0 / 6 / 0	9A09h
TBT_PCIe0	All	0 / 7 / 0	9A23h
TBT_PCIe1	All	0 / 7 / 1	9A25h
TBT_PCIe2	All	0 / 7 / 2	9A27h
TBT_PCIe3	All	0 / 7 / 3	9A29h
GNA	All	0 / 8 / 0	9A11h

Table 15-4. Other Device ID UP3/UP4/H35/UP3-Refresh/H35-Refresh (Sheet 2 of 2)

Device	Processor Line	Bus / Device / Function	DID
NPK	All	0 / 9 / 0	9A33h
Crash-log SRAM	All	0 / 10 / 0	9A0Dh
USB xHCI	UP4, UP3	0 / 13 / 0	9A13h
USB xDCI	UP4, UP3	0 / 13 / 1	9A15h
TBT DMA0	All	0 / 13 / 2	9A1Bh
TBT DMA1	All	0 / 13 / 3	9A1Dh
VMD	All	0 / 14 / 0	9A0Bh

Table 15-5. Other Device IDs (H Processor Line)

Device	Bus / Device / Function	DID
PCIe RC 010 (x16)	0 / 1 / 0	9A01h
PCIe RC 011 (x8)	0 / 1 / 1	9A05h
PCIe RC 012 (x4)	0 / 1 / 2	9A07h
DPTF/DPPM	0 / 4 / 0	9A03h
IPU	0 / 5 / 0	9A39h
PCIe RC 060 (x4)	0 / 6 / 0	9A0Fh
TBT_PCIe0	0 / 7 / 0	9a2Bh
TBT_PCIe1	0 / 7 / 1	9a2Dh
TBT_PCIe2	0 / 7 / 2	9a2fh
TBT_PCIe3	0 / 7 / 3	9a31h
GNA	0 / 8 / 0	9A11h
NPK	0 / 9 / 0	9A33h
Crash-log SRAM	0 / 10 / 0	9A0Dh
USB xHCI	0 / 13 / 0	9A17h
USB xDCI	0 / 13 / 1	9A15h
TBT DMA0	0 / 13 / 2	9a1fh
TBT DMA1	0 / 13 / 3	9a21h
VMD	0 / 14 / 0	9A0Bh

Table 15-6. Graphics Device ID (DID2)

SKU	#EUs	Processor Step	GT SKU	Device 2 ID	Device 2 Rev
UP4 4+2/2+2	96/80	B0	GT2	9A40h	1h
UP4 4+2/2+2	48	B0	GT2	9A78h	1h
UP3 4+2/2+2	96/80	B0	GT2	9A49h	1h
UP3 4+2/2+2	48	B0	GT2	9A78h	1h
H35 4+2	96/80	B0	GT2	9A49h	1h
UP3-Refresh 4+2	96/80	C0	GT2	9A49h	3h
H35-Refresh 4+2	96/80	C0	GT2	9A49h	3h
H 8+1/6+1	32	R0	GT1	9A60h	1h
H 8+1/6+1	16	R0	GT1	9A68h	1h

§ §

1-1	11 th Generation Intel® Core™ Processor Lines	12
1-2	Operating Systems Support	18
1-3	Terminology	18
1-4	Special marks	21
3-1	System States	56
3-2	Integrated Memory Controller (IMC) States	57
3-3	G, S, and C Interface State Combinations	57
3-4	Core C-states	59
3-5	Package C-States	61
3-6	Deepest Package C-State Available	64
3-7	TCSS Power State	64
3-8	Package C-States with PCIe* Link States Dependencies	68
4-1	Configurable TDP Modes	73
4-2	TDP Specifications	82
4-3	Package Turbo Specifications	84
4-4	Junction Temperature Specifications	85
5-1	DDR Support Matrix Table	86
5-2	Processor DDR Speed Support	86
5-4	LPDDR4x Channels Population Rules	87
5-5	DDR4 Memory Down Channels Population Rules	87
5-3	DDR Technology Support Matrix	87
5-6	H DDR4 SoDIMM Population Configuration	88
5-7	Supported DDR4 SoDIMM Module Configurations	89
5-8	Supported DDR4 Memory Down Device Configurations	89
5-9	Supported LPDDR4x x32 DRAMs Configurations	89
5-10	Supported LPDDR4x x64 DRAMs Configurations	90
5-11	DDR4 System Memory Timing Support	90
5-12	LPDDR4x System Memory Timing Support	91
5-13	System Agent Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies	91
5-14	ECC H-Matrix Syndrome Codes	94
5-15	Interleave (IL) and Non-Interleave (NIL) Modes Pin Mapping	95
6-1	USB-C* Port Configuration	101
6-2	USB-C* Lanes Configuration	102
6-3	USB-C* Non-Supported Lane Configuration	102
6-4	PCIe via USB4 Configuration	104
7-1	PCI Express* 4 -lane Bifurcation and Lane Reversal Mapping	105
7-2	PCI Express* 16-lane Bifurcation and Lane Reversal Mapping	105
7-3	PCI Express* Maximum Transfer Rates and Theoretical Bandwidth	107
9-1	Hardware Accelerated Video Decoding	112
9-2	Hardware Accelerated Video Encode	113
10-1	Display Ports Availability and Link Rate	116
10-2	Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations	120
10-3	DisplayPort Maximum Resolution	121
10-4	HDMI Maximum Resolution	122
10-5	Embedded DisplayPort Maximum Resolution	123
10-6	MIPI* DSI Maximum Resolution	124
10-7	Processor Supported Audio Formats over HDMI and DisplayPort*	124
11-1	CSI-2 Lane Configuration for UP3-Processor Line	128
11-2	CSI-2 Lane Configuration for UP4-Processor Line	128
11-3	CSI-2 Lane Configuration for H-Processor Line	129
12-1	Signal Tables Terminology	130
12-2	DDR4 Memory Interface	132
12-3	LPDDR4x Memory Interface	135
12-4	DMI Interface Signals	136
12-5	Processor Clocking Signals	142



12-6	Processor Power Rails Signals.....	143
12-7	Processor Pull-up Power Rails Signals	144
12-8	GND, RSVD, and NCTF Signals	145
13-1	Processor Vcc _{IN} Active and Idle Mode DC Voltage and Current Specifications	147
13-2	Vcc _{IN_AUX} Supply DC Voltage and Current Specifications	149
13-3	Memory Controller (VDD2) Supply DC Voltage and Current Specifications.....	151
13-4	Vcc Sustain (VccST) Supply DC Voltage and Current Specifications.....	151
13-5	Vcc Sustain Gated (VccSTG) Supply DC Voltage and Current Specifications.....	152
13-7	DDR4 Signal Group DC Specifications	153
13-6	Vcc1P8A Supply DC Voltage and Current Specifications.....	153
13-8	LPDDR4x Signal Group DC Specifications	155
13-9	155
13-10	DSI HS Transmitter DC Specifications.....	156
13-11	DSI LP Transmitter DC Specifications	156
13-12	Digital Display Interface Group DC Specifications (DP/HDMI)	157
13-13	DP-IN Group DC Specifications	157
13-14	CMOS Signal Group DC Specifications.....	158
13-15	GTL Signal Group and Open Drain Signal Group DC Specifications	158
13-16	PECI DC Electrical Limits.....	159
13-17	System Reference Clocks DC and AC Specifications	160
14-1	Package Mechanical Attributes.....	161
14-2	Package Loading Specifications.....	162
14-3	Package Loading Specifications.....	162
15-1	CPUID Format	164
15-2	PCI Configuration Header.....	165
15-3	Host Device ID (DID0)	165
15-4	Other Device ID UP3/UP4/H35/UP3-Refresh/H35-Refresh.....	165
15-5	Other Device IDs (H Processor Line).....	166
15-6	Graphics Device ID (DID2).....	167

1-1	11 th Generation Intel® Core™ UP3/UP4/H35/UP3-Refresh//H35-Refresh Processor Lines Platform Diagram	13
1-2	H Processor Line Platform Diagram	14
2-1	Example for PECI Host-Clients Connection	24
2-2	Example for PECI EC Connection	25
2-3	Device to Domain Mapping Structures	29
2-4	Processor Cache Hierarchy	39
2-5	Processor Camera System	47
2-6	Telemetry Aggregator	49
3-1	UP3 and UP4 Processor Lines Power States	54
3-2	H Processor Line Power States	55
3-3	Processor Package and IA Core C-States	56
3-4	Idle Power Management Breakdown of the Processor IA Cores	58
3-5	Package C-State Entry and Exit	61
4-1	Package Power Control	71
4-2	PROCHOT Demotion Signal Description	79
5-1	Intel® DDR4 Flex Memory Technology Operations	92
5-2	DDR4 Interleave (IL) and Non-Interleave (NIL) Modes Mapping	96
6-1	USB-C* Sub-system Block Diagram	100
7-1	PCI Express* Related Register Structures in the Processor	108
8-1	Example for DMI Lane Reversal Connection	110
10-1	Processor Display Architecture	117
10-2	DisplayPort* Overview	119
10-3	HDMI* Overview	122
10-4	MIPI* DSI Overview	123
10-5	DP-IN Block Diagram	125
13-1	Input Device Hysteresis	160

1	Introduction	12
1.1	Processor Volatility Statement	14
1.2	Package Support	14
1.3	Supported Technologies	15
1.3.1	API Support (Windows*)	16
1.4	Power Management Support	16
1.4.1	Processor Core Power Management	16
1.4.2	System Power Management	16
1.4.3	Memory Controller Power Management	17
1.4.4	Processor Graphics Power Management	17
1.4.4.1	Memory Power Savings Technologies	17
1.4.4.2	Display Power Savings Technologies	17
1.4.4.3	Graphics Core Power Savings Technologies	17
1.5	Thermal Management Support	17
1.6	Ball-out Information	18
1.7	Processor Testability	18
1.8	Operating Systems Support	18
1.9	Terminology and Special Marks	18
1.10	Related Documents	22
2	Technologies	23
2.1	Platform Environmental Control Interface (PECI)	23
2.1.1	PECI Bus Architecture	23
2.2	Intel® Virtualization Technology (Intel® VT)	25
2.2.1	Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-X)	25
2.2.2	Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)	28
2.2.3	Intel® APIC Virtualization Technology (Intel® APICv)	30
2.3	Security Technologies	31
2.3.1	Intel® Trusted Execution Technology (Intel® TXT)	31
2.3.2	Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)	32
2.3.3	Perform Carry-Less Multiplication Quad Word Instruction (PCLMULQDQ)	33
2.3.4	Intel® Secure Key	33
2.3.5	Execute Disable Bit	33
2.3.6	Boot Guard Technology	34
2.3.7	Intel® Supervisor Mode Execution Protection (SMEP)	34
2.3.8	Intel® Supervisor Mode Access Protection (SMAP)	34
2.3.9	Intel® Software Guard Extensions (Intel® SGX)	34
2.3.10	Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)	36
2.3.11	User Mode Instruction Prevention (UMIP)	36
2.3.12	Read Processor ID (RDPID)	36
2.3.13	Total Memory Encryption (Intel® TME)	37
2.3.14	Control-flow Enforcement Technology (Intel® CET)	37
2.3.14.1	Shadow Stack	37
2.3.14.2	Indirect Branch Tracking	38
2.3.15	KeyLocker Technology	38
2.3.16	Devil's Gate Rock (DGR)	38
2.4	Power and Performance Technologies	39
2.4.1	Intel® Smart Cache Technology	39
2.4.2	IA Core Level 1 and Level 2 Caches	39
2.4.3	Intel® Turbo Boost Max Technology 3.0	40
2.4.4	Power Aware Interrupt Routing (PAIR)	40
2.4.5	Intel® Hyper-Threading Technology (Intel® HT Technology)	40
2.4.6	Intel® Turbo Boost Technology 2.0	41
2.4.6.1	Intel® Turbo Boost Technology 2.0 Power Monitoring	41

2.4.6.2	Intel® Turbo Boost Technology 2.0 Power Control	41
2.4.6.3	Intel® Turbo Boost Technology 2.0 Frequency	41
2.4.7	Enhanced Intel SpeedStep® Technology	42
2.4.8	Intel® Speed Shift Technology	42
2.4.9	Intel® Advanced Vector Extensions 2 (Intel® AVX2)	42
2.4.10	Advanced Vector Extensions 512 Bit (Intel® AVX-512)	43
2.4.11	Intel® 64 Architecture x2APIC	44
2.4.12	Intel® Dynamic Tuning Technology (DTT)	45
2.4.13	Intel® GNA 2.0 (GMM and Neural Network Accelerator)	46
2.4.14	Cache Line Write Back (CLWB)	46
2.4.15	Ring Interconnect	46
2.5	Intel® Image Processing Unit (Intel® IPU6)	47
2.5.1	Platform Imaging Infrastructure	47
2.5.2	Intel® Image Processing Unit (Intel® IPU6)	47
2.6	Debug Technologies	48
2.6.1	Intel® Processor Trace	48
2.6.2	Platform CrashLog	48
2.6.3	Telemetry Aggregator	48
2.7	Clock Topology	50
2.7.1	Integrated Reference Clock PLL	50
2.8	Intel Volume Management Device (VMD) Technology	50
2.8.1	Intel Volume Management Device Technology Objective	50
2.8.2	Intel Volume Management Device Technology	51
2.8.3	Key Features	52
2.9	Deprecated Technologies	52
3	Power Management	53
3.1	Advanced Configuration and Power Interface (ACPI) States Supported	56
3.2	Processor IA Core Power Management	57
3.2.1	OS/HW Controlled P-states	58
3.2.1.1	Enhanced Intel SpeedStep® Technology	58
3.2.1.2	Intel® Speed Shift Technology	58
3.2.2	Low-Power Idle States	58
3.2.3	Requesting the Low-Power Idle States	59
3.2.4	Processor IA Core C-State Rules	59
3.2.5	Package C-States	60
3.2.6	Package C-States and Display Resolutions	63
3.3	Processor Graphics Power Management	64
3.3.1	Memory Power Savings Technologies	64
3.3.1.1	Intel® Rapid Memory Power Management (Intel® RMPM)	64
3.3.2	Display Power Savings Technologies	65
3.3.2.1	Intel® Seamless Display Refresh Rate Switching Technology (Intel® SDRRS Technology) with eDP* Port	65
3.3.2.2	Intel® Automatic Display Brightness	65
3.3.2.3	Smooth Brightness	65
3.3.2.4	Intel® Display Power Saving Technology (Intel® DPST) 6.3	65
3.3.2.5	Panel Self-Refresh 2 (PSR 2)	66
3.3.2.6	Low-Power Single Display Pipe (LPSP)	66
3.3.2.7	Intel® Smart 2D Display Technology (Intel® S2DDT)	66
3.3.3	Processor Graphics Core Power Savings Technologies	66
3.3.3.1	Intel® Graphics Dynamic Frequency	66
3.3.3.2	Intel® Graphics Render Standby Technology (Intel® GRST)	67
3.3.3.3	Dynamic FPS (DFPS)	67
3.4	System Agent Enhanced Intel SpeedStep® Technology	67
3.5	Voltage Optimization	67
3.6	ROP (Rest Of Platform) PMIC	67

3.7	PCI Express* Power Management	68
4	Thermal Management	69
4.1	Processor Thermal Management	69
4.1.1	Thermal Considerations	69
4.1.1.1	Package Power Control	70
4.1.1.2	Platform Power Control	71
4.1.1.3	Turbo Time Parameter (Tau)	72
4.1.2	Configurable TDP (cTDP) and Low-Power Mode	72
4.1.2.1	Configurable TDP	72
4.1.2.2	Low-Power Mode	73
4.1.3	Thermal Management Features	74
4.1.3.1	Adaptive Thermal Monitor	74
4.1.3.2	Digital Thermal Sensor	76
4.1.3.3	PROCHOT# Signal	77
4.1.3.4	PROCHOT Output Only	78
4.1.3.5	Bi-Directional PROCHOT#	78
4.1.3.6	PROCHOT Demotion Algorithm	78
4.1.3.7	Voltage Regulator Protection using PROCHOT#	79
4.1.3.8	Thermal Solution Design and PROCHOT# Behavior	79
4.1.3.9	Low-Power States and PROCHOT# Behavior	80
4.1.3.10	THRMTRIP# Signal	80
4.1.3.11	Critical Temperature Detection	80
4.1.3.12	On-Demand Mode	80
4.1.3.13	MSR Based On-Demand Mode	80
4.1.3.14	I/O Emulation-Based On-Demand Mode	81
4.1.4	Intel® Memory Thermal Management	81
4.2	Thermal and Power Specifications	82
5	Memory	86
5.1	System Memory Interface	86
5.1.1	Processor SKU Support Matrix	86
5.1.2	Supported Population	87
5.1.3	Supported Memory Modules and Devices	89
5.1.4	System Memory Timing Support	90
5.1.5	System Memory Timing Support	90
5.1.6	SAGV Points	91
5.1.7	Memory Controller (MC)	91
5.1.8	System Memory Controller Organization Mode (DDR4)	92
5.1.9	System Memory Frequency	93
5.1.10	Technology Enhancements of Intel® Fast Memory Access (Intel® FMA)	93
5.1.11	Data Scrambling	94
5.1.12	ECC DDR4 H-Matrix Syndrome Codes	94
5.1.13	Data Swapping	95
5.1.14	DDR I/O Interleaving	95
5.1.15	DRAM Clock Generation	96
5.1.16	DRAM Reference Voltage Generation	96
5.1.17	Data Swizzling	96
5.2	Integrated Memory Controller (IMC) Power Management	96
5.2.1	Disabling Unused System Memory Outputs	96
5.2.2	DRAM Power Management and Initialization	97
5.2.2.1	Initialization Role of CKE	98
5.2.2.2	Conditional Self-Refresh	98
5.2.2.3	Dynamic Power-Down	98
5.2.2.4	DRAM I/O Power Management	99
5.2.3	DDR Electrical Power Gating	99
5.2.4	Power Training	99

6	USB-C* Sub System	100
6.1	USB-C Sub-System General Capabilities.....	101
6.2	USB™ 4 Router	102
6.2.1	USB 4 Host Router Implementation Capabilities	102
6.3	USB-C Sub-system xHCI/xDCI Controllers.....	103
6.3.1	USB 3 Controllers	103
6.3.1.1	Extensible Host Controller Interface (xHCI).....	103
6.3.1.2	Extensible Device Controller Interface (xDCI).....	103
6.3.2	USB-C Sub-System PCIe Interface	104
6.4	USB-C Sub-System Display Interface.....	104
7	PCIe* Interface	105
7.1	Processor PCI Express* Interface	105
7.1.1	PCI Express* Support.....	105
7.1.2	PCI Express* Lane Polarity Inversion.....	107
7.1.3	PCI Express* Architecture.....	107
7.1.4	PCI Express* Configuration Mechanism.....	108
7.1.5	PCI Express* Equalization Methodology	108
7.1.6	PCI Express* Hot-Plug.....	109
7.1.6.1	Presence Detection.....	109
7.1.6.2	SMI/SCI Generation.....	109
8	Direct Media Interface (DMI) and On Package Interface (OPI)	110
8.1	Direct Media Interface (DMI)	110
8.1.1	DMI Lane Reversal and Polarity Inversion.....	110
8.1.2	DMI Error Flow	111
8.1.3	DMI Link Down	111
8.2	On Package Interface (OPI).....	111
8.2.1	OPI Support:.....	111
8.2.2	Functional description:	111
9	Graphics	112
9.1	Processor Graphics.....	112
9.1.1	Media Support (Intel® QuickSync and Clear Video Technology HD)	112
9.1.1.1	Hardware Accelerated Video Decode	112
9.1.1.2	Hardware Accelerated Video Encode.....	113
9.1.1.3	Hardware Accelerated Video Processing	114
9.1.1.4	Hardware Accelerated Transcoding	114
9.2	Platform Graphics Hardware Feature	115
9.2.1	Hybrid Graphics.....	115
10	Display	116
10.1	Display Technologies Support.....	116
10.2	Display Configuration	116
10.3	Display Features.....	118
10.3.1	General Capabilities	118
10.3.2	Multiple Display Configurations	118
10.3.3	High-bandwidth Digital Content Protection (HDCP)	118
10.3.4	DisplayPort*	119
10.3.4.1	Multi-Stream Transport (MST).....	119
10.3.5	High-Definition Multimedia Interface (HDMI*).....	121
10.3.6	embedded DisplayPort* (eDP*)	122
10.3.7	MIPI* DSI.....	123
10.3.8	Integrated Audio	124
10.3.9	DisplayPort* Input (DP-IN)	124
11	Camera/MIPI	127

11.1	Camera Pipe Support	127
11.2	MIPI* CSI-2 Camera Interconnect	127
11.2.1	Camera Control Logic.....	127
11.2.2	Camera Modules	127
11.2.3	CSI-2 Lane Configuration	128
12	Signal Description	130
12.1	System Memory Interface	132
12.1.1	DDR4 Memory Interface.....	132
12.1.2	LPDDR4x Memory Interface.....	135
12.2	PCIe4 Gen4 Interface Signals.....	136
12.3	Direct Media Interface (DMI) Signals.....	136
12.4	PCIe16 Gen4 Interface Signals.....	137
12.5	Reset and Miscellaneous Signals.....	137
12.6	Display Interfaces	138
12.6.1	Embedded DisplayPort* (eDP*) Signals	138
12.6.2	MIPI DSI* Signals	139
12.6.3	Digital Display Interface (DDI) Signals	139
12.6.4	Digital Display Audio Signals.....	139
12.7	DP-IN Interface Signals.....	140
12.8	USB Type-C Signals.....	140
12.9	MIPI* CSI-2 Interface Signals	141
12.10	Processor Clocking Signals.....	142
12.11	Testability Signals	142
12.12	Error and Thermal Protection Signals	143
12.13	Processor Power Rails	143
12.14	Ground, Reserved and Non-Critical to Function (NCTF) Signals	144
12.15	Processor Internal Pull-Up / Pull-Down Terminations	145
13	Electrical Specifications	146
13.1	Processor Power Rails	146
13.1.1	Power and Ground Pins	146
13.1.2	Full Integrated Voltage Regulator (FIVR)	146
13.1.3	V _{CC} Voltage Identification (VID)	146
13.2	DC Specifications	147
13.2.1	Processor Power Rails DC Specifications	147
13.2.1.1	V _{CCIN} DC Specifications.....	147
13.2.1.2	V _{CCIN_AUX} DC Specifications.....	149
13.2.1.3	VDD2 DC Specifications	151
13.2.1.4	V _{CCST} DC Specifications.....	151
13.2.1.5	V _{CC1P8A} DC Specifications	152
13.2.1.6	DDR4 DC Specifications	153
13.2.1.7	LPDDR4x DC Specifications	155
13.2.1.8	PCI Express* Graphics (PEG) DC Specifications.....	156
13.2.1.9	Digital Display Interface (DDI) DC Specifications.....	156
13.2.1.10	embedded DisplayPort* (eDP*) DC Specification	157
13.2.1.11	MIPI* CSI-2 D-Phy Receiver DC Specifications.....	157
13.2.1.12	CMOS DC Specifications.....	158
13.2.1.13	GTL and OD DC Specification.....	158
13.2.1.14	PECI DC Characteristics	159
14	Package Mechanical Specifications	161
14.1	Package Mechanical Attributes	161
14.2	Package Loading and Die Pressure Specifications	161
14.2.1	Package Loading Specifications	161
14.2.2	Die Pressure Specifications	162
14.3	Package Storage Specifications	163



15	CPU And Device IDs	164
15.1	CPUID	164
15.2	PCI Configuration Header	164