



Intel[®] Core[™] Ultra 200S Series Processors

Datasheet, Volume 1 of 2

Rev. 001

October 2024



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.

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](https://www.intel.com).

Altering clock frequency, voltage, or memory interface speeds may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Intel has not validated processor running memory above Plan-Of-Record (POR) speed. DRAM/DIMM devices should support desired speed, check with DRAM/DIMM vendors for details. System manufacturers are responsible for all validation and assume the risk of any stability, security, performance, or other functional issues resulting from such alterations

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

Copyright © 2024, Intel Corporation. All rights reserved.

Contents

Revision History	11
1.0 Introduction	12
1.1 Processor Volatility Statement.....	13
1.2 Package Support.....	14
1.2.1 S Processor Package Support.....	14
1.3 Supported Technologies.....	14
1.3.1 API Support (Windows*).....	16
1.3.2 Firmware Resiliency.....	16
1.4 Power Management Support.....	17
1.4.1 Processor Core Power Management.....	17
1.4.2 System Power Management.....	17
1.4.3 Memory Controller Power Management.....	17
1.4.4 Processor Graphics Power Management.....	17
1.5 Thermal Management Support.....	18
1.6 Ballout Information.....	18
1.7 Processor Testability.....	18
1.8 Operating Systems Support.....	18
1.9 Terminology and Special Marks.....	19
1.10 Related Documents.....	21
2.0 Processor and Device IDs	22
2.1 CPUID.....	22
2.2 PCI Configuration Header.....	22
2.3 Device IDs.....	23
2.4 Revision IDs.....	25
3.0 Package Mechanical Specifications	26
3.1 Package Mechanical Attributes.....	26
3.2 Package Storage Specifications.....	26
4.0 Memory Mapping	28
4.1 Functional Description.....	28
4.1.1 PCI Devices and Functions.....	28
4.1.2 Fixed I/O Address Ranges.....	28
4.1.3 Variable I/O Decode Ranges.....	30
4.2 Memory Map.....	30
4.2.1 Boot Block Update Scheme.....	31
5.0 Security Technologies	33
5.1 Intel® Converged Boot Guard and Intel® TXT.....	33
5.2 Crypto Acceleration Instructions.....	34
5.2.1 Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI).....	34
5.2.2 Perform Carry-Less Multiplication Quad Word Instruction (PCLMULQDQ)	34
5.2.3 Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions).....	35
5.2.4 New Cryptographic Acceleration Instructions.....	35
5.3 Intel® Secure Key.....	35
5.4 Execute Disable Bit	36
5.5 Intel® Supervisor Mode Execution Prevention (Intel® SMEP).....	36

- 5.6 Intel® Supervisor Mode Access Prevention (Intel® SMAP)..... 36
- 5.7 User Mode Instruction Prevention (UMIP) 36
- 5.8 Read Processor ID (RDPID) 37
- 5.9 Intel® Total Memory Encryption - Multi-Key..... 37
- 5.10 Control-flow Enforcement Technology (Intel® CET)..... 37
 - 5.10.1 Shadow Stack..... 38
 - 5.10.2 Indirect Branch Tracking 38
- 5.11 KeyLocker Technology..... 38
- 5.12 Intel® System Resources Defense and Intel® System Security Report..... 39
- 5.13 BIOS Guard..... 39
- 5.14 Intel® Platform Trust Technology..... 39
- 5.15 Linear Address Space Separation (LASS)..... 39
- 5.16 Security Firmware Engines..... 39
 - 5.16.1 Intel® Silicon Security Engine..... 40
 - 5.16.2 Intel® Graphics System Controller (Intel® GSC)..... 40
- 6.0 Intel® Virtualization Technology (Intel® VT)..... 41**
 - 6.1 Intel® Virtualization Technology (Intel® VT) for Intel® 64 and Intel® Architecture (Intel® VT-x) 41
 - 6.2 Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d) 43
 - 6.3 Intel® APIC Virtualization Technology (Intel® APICv)..... 45
- 7.0 Instructions Set Enhancements..... 47**
 - 7.1 CMPccXADD..... 47
 - 7.2 LAM..... 47
- 8.0 Platform Environmental Control Interface (PECI)..... 48**
 - 8.1 PECI Bus Architecture..... 48
- 9.0 Intel GMM and Neural Network Accelerator (Intel GNA 3.5)..... 51**
- 10.0 Intel® Neural Processing Unit (Intel® NPU)..... 53**
 - 10.1 Functional Description..... 53
 - 10.1.1 Processor Subsystem..... 54
 - 10.1.2 NCE Subsystem..... 55
- 11.0 Power Management..... 57**
 - 11.1 System Power States, Advanced Configuration and Power Interface (ACPI) 57
 - 11.2 Processor IA Core Power Management..... 59
 - 11.2.1 OS/HW Controlled P-states..... 59
 - 11.3 Power and Performance Technologies..... 60
 - 11.3.1 Intel® Thread Director 60
 - 11.3.2 Intel® Smart Cache Technology..... 60
 - 11.3.3 P-core and E-core Level 1 and Level 2 Caches 60
 - 11.3.4 Intel® Adaptive Boost Technology..... 62
 - 11.3.5 **Intel® System Agent Enhanced SpeedStep® Technology** 62
 - 11.3.6 User Mode Wait Instructions 62
 - 11.4 Deprecated Technology..... 63
 - 11.5 Power and Internal Signals..... 63
 - 11.5.1 Signal Description..... 63
 - 11.5.2 Power Sequencing Signals..... 64
 - 11.5.3 I/O Signal Planes and States..... 64

- 12.0 Power Delivery.....65**
 - 12.1 Power and Ground Signals.....65
 - 12.2 Current Excursion Protection (CEP).....66
- 13.0 Electrical Specifications..... 67**
 - 13.1 Processor Power Rails..... 67
 - 13.1.1 Power and Ground Pins..... 67
 - 13.1.2 VCC Voltage Identification (VID).....67
 - 13.2 DC Specifications..... 67
 - 13.2.1 Processor Power Rails DC Specifications..... 68
 - 13.2.2 Processor Interfaces DC Specifications..... 76
- 14.0 Thermal Management..... 83**
 - 14.1 Processor Thermal Management.....83
 - 14.1.1 Thermal Considerations.....83
 - 14.1.2 Assured Power (cTDP).....86
 - 14.1.3 Thermal Management Features..... 88
 - 14.1.4 Intel® Memory Thermal Management94
 - 14.2 Processor Base Power Thermal and Power Specifications95
 - 14.3 Processor Line Thermal and Power Specifications.....97
 - 14.4 Error and Thermal Protection Signals..... 98
 - 14.5 Thermal Metrology 99
 - 14.6 Fan Speed Control Scheme with DTS.....100
 - 14.7 Thermal Sensor.....103
 - 14.7.1 Modes of Operation.....103
 - 14.7.2 Temperature Trip Point..... 104
 - 14.7.3 Thermal Sensor Accuracy (Taccuracy)..... 104
 - 14.7.4 Thermal Reporting to EC.....104
- 15.0 Clock Topology..... 105**
 - 15.1 Integrated Reference Clock PLL..... 105
 - 15.2 Processor Clocking Signals.....106
- 16.0 Memory..... 108**
 - 16.1 System Memory Interface.....108
 - 16.1.1 Processor SKU Support Matrix..... 108
 - 16.1.2 Supported Memory Modules and Devices..... 109
 - 16.1.3 System Memory Timing Support..... 110
 - 16.1.4 Memory Controller (MC)..... 111
 - 16.1.5 System Memory Controller Organization Mode..... 112
 - 16.1.6 System Memory Frequency..... 114
 - 16.1.7 Technology Enhancements of Intel® Fast Memory Access (Intel® FMA)..... 114
 - 16.1.8 Data Scrambling.....114
 - 16.1.9 ECC H-Matrix Syndrome Codes.....115
 - 16.1.10 Data Swapping 115
 - 16.1.11 DDR I/O Interleaving..... 115
 - 16.1.12 DRAM Clock Generation 116
 - 16.1.13 DRAM Reference Voltage Generation116
 - 16.1.14 Data Swizzling.....116
 - 16.1.15 Post Package Repair (PPR)..... 116
 - 16.2 Integrated Memory Controller (IMC) Power Management..... 116

16.2.1 Disabling Unused System Memory Outputs.....	116
16.2.2 DRAM Power Management and Initialization.....	116
16.2.3 DDR Electrical Power Gating.....	118
16.2.4 Power Training.....	118
16.3 Signal Description.....	119
17.0 USB Type-C* Sub System.....	120
17.1 General Capabilities.....	120
17.2 USB4* Router.....	122
17.2.1 USB4 Host Router Implementation Capabilities.....	122
17.3 xHCI/xDCI Controllers	123
17.3.1 USB 3 Controllers.....	123
17.4 Display Interface.....	124
17.5 USB Type-C Signals.....	124
17.6 LSx.....	124
17.6.1 LSx Signal Description.....	124
17.6.2 Integrated Pull-Ups and Pull-Downs.....	124
17.6.3 I/O Signal Planes and States.....	124
17.7 AUX BIAS Control.....	125
18.0 Intel® Volume Management Device (Intel® VMD) Technology	126
19.0 PCI Express* (PCIe*).....	129
19.1 Functional Description.....	129
19.1.1 PCI Express* Power Management.....	131
19.1.2 Port 80h Decode.....	131
19.1.3 Separate Reference Clock with Independent SSC (SRIS).....	131
19.1.4 Advanced Error Reporting.....	132
19.1.5 Single - Root I/O Virtualization (SR - IOV).....	132
19.1.6 PCI Express* Receiver Lane Polarity Inversion.....	132
19.1.7 Precision Time Measurement (PTM)	132
19.2 Signal Description.....	133
19.3 I/O Signal Planes and States.....	133
19.4 PCI Express* Root Port Support Feature Details.....	133
20.0 Graphics.....	138
20.1 Processor Graphics.....	138
20.1.1 Media Support (Intel® QuickSync and Clear Video Technology HD).....	138
20.1.2 Graphics Core Cache.....	141
20.2 Platform Graphics Hardware Feature	141
20.2.1 Hybrid Graphics.....	141
21.0 Display.....	142
21.1 Display Technologies Support.....	142
21.2 Display Interfaces	142
21.2.1 Digital Display Interface DDI Signals.....	142
21.2.2 Digital Display Interface TCP Signals.....	144
21.3 Display Features.....	144
21.3.1 General Capabilities.....	145
21.3.2 Multiple Display Configurations.....	146
21.3.3 High-bandwidth Digital Content Protection (HDCP).....	146
21.3.4 DisplayPort*.....	146

21.3.5 High-Definition Multimedia Interface (HDMI*).....	148
21.3.6 embedded DisplayPort* (eDP*).....	150
21.3.7 Integrated Audio.....	150
22.0 General Purpose Input and Output.....	152
22.1 Functional Description.....	152
22.1.1 Interrupt / IRQ via GPIO Requirement.....	152
22.1.2 Integrated Pull-ups and Pull-downs.....	152
22.1.3 SCI / SMI and NMI.....	152
22.1.4 Timed GPIO.....	153
22.2 Signal Description.....	153
23.0 Interrupt Timer Subsystem (ITSS).....	154
23.1 Feature Overview.....	154
23.2 Functional Description.....	154
23.2.1 8254 Timers.....	155
23.2.2 APIC Advanced Programmable Interrupt Controller.....	157
23.2.3 High Precision Event Timer (HPET).....	157
24.0 Direct Media Interface (DMI).....	162
24.1 DMI Lane Reversal and Polarity Inversion.....	162
24.2 DMI Error Flow.....	163
24.3 DMI Link Down.....	163
24.4 Signal Description.....	164
25.0 Direct Enhanced Serial Peripheral Interface (Direct eSPI).....	165
25.1 Functional Description.....	165
25.1.1 Processor-PCH eSPI Return Clock Support.....	165
25.1.2 PCH Multiple External eSPI Device Support.....	165
25.1.3 Processor Direct eSPI Channel Support.....	165
25.2 Signal Description.....	167
26.0 Testability and Monitoring.....	169
26.1 Signal Description.....	169
26.2 I/O Signal Planes and States.....	169
27.0 Miscellaneous Signals.....	171
27.1 Signal Description.....	171
27.2 Ground and Reserved Signals.....	171

Figures

1	S LGA Processor Line Platform Diagram.....	13
2	Device to Domain Mapping Structures	44
3	PECI Host-Clients Connection Example.....	49
4	PECI EC Connection Example.....	50
5	NPU IP Block Diagram	54
6	Power State Block Diagram.....	59
7	P-core and E-core Cache Hierarchy	61
8	Package Power Control.....	85
9	PROCHOT Demotion Description	92
10	Thermal Test Vehicle (TMTV) Case Temperature (T _{CASE}) Measurement Location	100
11	Digital Thermal Sensor (DTS) 1.1 Definition Points	101
12	Digital Thermal Sensor (DTS) 2.0 Definition Points.....	103
13	System Clock Block Diagram.....	105
14	Intel® DDR5 Flex Memory Technology Operations.....	113
15	GPIO - Virtual Wire Index Bit Mapping	125
16	Technology Description.....	127
17	Supported PCI Express* Link Configurations	134
18	PCIe Controller Bifurcation/Configuration Mode Strap Details	136
19	Processor Display Architecture.....	145
20	DisplayPort* Overview.....	147
21	HDMI* Overview	149
22	Example for DMI Lane Reversal Connection.....	163

Tables

1	Processor Series	12
2	Terminology.....	19
3	Special Marks	21
4	CPUID Format.....	22
5	PCI Configuration Header.....	23
6	Host Device ID (DID0) and Processor Graphics Device ID (DID2).....	23
7	Other Device ID.....	24
8	ACPI Device ID for GPIO Controller.....	24
9	S LGA Package Mechanical Attributes.....	26
10	Fixed I/O Ranges Decoded by Processor.....	28
11	Variable I/O Decode Ranges	30
12	Processor Memory Decode Ranges (Processor Perspective).....	31
13	Boot Block Update Scheme.....	31
14	Acronyms.....	57
15	References.....	57
16	General System Power States	57
17	State Transition Rules for the Processor	58
18	Power Sequencing Signals	64
19	Power Rail Description.....	65
20	Power Rail Sense Signals.....	65
21	Processor VCC _{CORE} Active and Idle Mode DC Voltage and Current Specifications (S Processor Line).....	68
22	Processor Graphics (VccGT) Supply DC Voltage and Current Specifications (S Processor Line).....	70
23	VccSA Supply DC Voltage and Current Specifications (S Processor Line).....	71
24	Memory Controller (VDD2) Supply DC Voltage and Current Specifications	73
25	VCCPRIM_VNNAON Supply DC Voltage and Current Specifications.....	73
26	VCCPRIM_IO Supply DC Voltage and Current Specifications.....	74
27	VCCPRIM_1P8_PROC Supply DC Voltage and Current Specifications.....	74
28	VCCPRIM_1P8_PROC_SOC Supply DC Voltage and Current Specifications.....	75
29	VCCPRIM_1P8_PROC_DDR Supply DC Voltage and Current Specifications.....	75
30	VCCPRIM_1P8_PROC_FLTRA Supply DC Voltage and Current Specifications.....	76
31	DSI HS Transmitter DC Specifications.....	77
32	DSI LP Transmitter DC Specifications.....	78
33	Display Audio and Utility Pins DC Specification.....	78
34	CMOS Signal Group DC Specifications	81
35	GTL Signal Group DC Specifications.....	81
36	SVID Signal Group DC Specifications	82
37	Definitions/Acronyms.....	83
38	Assured Power (cTDP).....	87
39	General Notes.....	95
40	Processor Base Power Specifications (S Processor Line)	96
41	Package Turbo Specifications (S Processor Lines)	97
42	Operating Temperature Specifications (S Processor Line)	98
43	Low Power and TMTV Specifications (S Processor Line LGA).....	98
44	TCONTROL Offset Configuration (S Processor Line LGA - Client)	98
45	Error and Thermal Protection Signals.....	98
46	Digital Thermal Sensor (DTS) 1.1 Thermal Solution Performance Above T _{CONTROL}	102
47	Thermal Margin Slope.....	103
48	Signal Description.....	106
49	DDR Support Matrix Table.....	108
50	DDR Technology Support Matrix.....	108
51	Supported DDR5 Non-ECC SoDIMM/CSoDIMM Module Configurations (S -Series Processor).....	109

52	Supported DDR5 ECC SoDIMM/CSoDIMM Module Configurations (S-Series Processor)	109
53	Supported DDR5 Non-ECC UDIMM/CUDIMM Module Configurations (S-Series Processor)	109
54	Supported DDR5 ECC UDIMM/CUDIMM Module Configurations (S-Series Processor).....	110
55	DDR5 System Memory Timing Support.....	110
56	SA Speed Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies	111
57	DDR5 Memory Interface.....	119
58	USB Type-C* Port Configuration.....	121
59	USB Type-C* Lanes Configuration.....	121
60	USB Type-C* Non-Supported Lane Configuration.....	122
61	Acronym.....	129
62	Reference Table.....	129
63	Features Supported.....	129
64	Power Plane and States for PCI Express* Signals	133
65	PCI Express* Root Port Feature Details	133
66	Hardware Accelerated Video Decoding	139
67	Hardware Accelerated Video Encode	139
68	Display Ports Availability and Link Rate.....	142
69	Digital Display Interface DDI Signals.....	142
70	Digital Display Interface TCP Signals.....	144
71	Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations.....	147
72	DisplayPort Maximum Resolution.....	148
73	HDMI Maximum Resolution.....	149
74	Embedded DisplayPort Maximum Resolution.....	150
75	Processor Supported Audio Formats over HDMI* and DisplayPort*.....	150
76	Acronyms.....	152
77	Acronyms.....	154
78	References.....	154
79	Counter Operating Modes.....	156
80	Region Entries in the Descriptor.....	166
81	Acronyms.....	169
82	References.....	169
83	Testability Signals.....	169
84	Power Planes and States for Testability Signals.....	169
85	GND, RSVD, and NCTF Signals.....	172



Revision History

Document Number	Revision Number	Description	Revision Date
832586	001	Initial Release	October 2024

1.0 Introduction

This document is intended for Original Equipment Manufacturers (OEMs), Original Design Manufacturers (ODM) and BIOS vendors creating products based on the Intel® Core™ Ultra 200S Series Processors.

Intel® Core™ Ultra Processors includes the Intel® Performance Hybrid architecture, P-Cores for performance, and E-Cores for Efficiency. Refer to [Table 1](#) on page 12 for availability in Intel processor lines. For more details on P-Core and E-Core, refer to [Power Management](#) on page 57.

This document assumes a working knowledge of the vocabulary and principles of interfaces and architectures such as PCI Express* (PCIe*), Universal Serial Bus (USB), Advance Host Controller Interface (AHCI), eXtensible Host Controller Interface (xHCI), and so on.

This document abbreviates buses as Bn, devices as Dn and functions as Fn. For example, Device 31 Function 0 is abbreviated as D31:F0, Bus 1 Device 8 Function 0 is abbreviated as B1:D8:F0. Generally, the bus number will not be used, and can be considered to be Bus 0.

The S-Processor Line is offered in a 2-Chip Platform that includes the Processor Chip and Platform Controller Hub (PCH-S) Chip in LGA and BGA Package.

Naming Convention in this document:

- S-Processor refers to Intel® Core™ Ultra 200S Series processors.

The S-Processor is based on the disaggregated architecture.

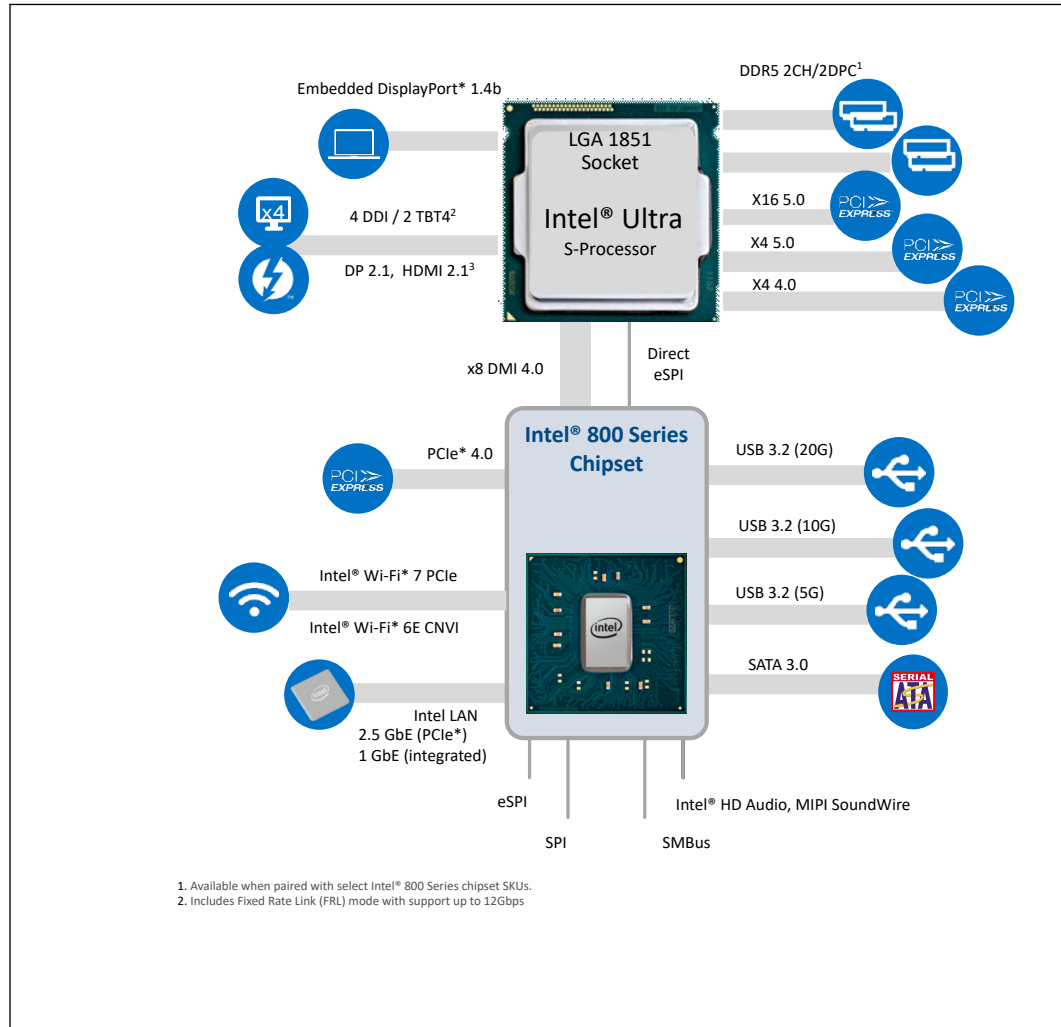
The following table describes the different processor series:

Table 1. Processor Series

Processor Line ¹	Package	Processor Base Power ^{3, 4}	Processor Max P-cores ²	Processor Max E-cores ²	Graphics Configuration Max Xe-cores ²	Platform Type
S-Processor	LGA1851	125 W	8	16	4	2-Chip

Notes: 1. Processor lines offering may change.
 2. Core count refers to native die with maximum physical cores. Some SKU may have reduced functional cores.
 3. For additional Processor Base Power Configurations, refer to [Processor Line Thermal and Power Specifications](#) on page 97, for adjustment to the Processor Base Power required to preserve base frequency associated with the sustained long-term thermal capability.
 4. Processor Base Power workload does not reflect I/O connectivity cases such as Thunderbolt, for power adders estimation for various I/O connectivity scenarios:

Figure 1. S LGA Processor Line Platform Diagram



NOTE

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.

1.1 Processor Volatility Statement

The 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

1.2.1 S Processor Package Support

The S-Processor line is available in the following package:

LGA1851

- A 45 x 37.5 mm
- Package Z-Height = 4.464 ± 0.133 mm

1.3 Supported Technologies

- PCI Express* (PCIe*)
- Xe Graphics Core Based Processor Graphics
- Display
 - DisplayPort* 2.1 (DP* 2.1)
 - Embedded DisplayPort* 1.4 (eDP* 1.4)
 - High-Definition Multimedia Interface* 2.1 (HDMI* 2.1)
- Intel® Gaussian & Neural Accelerator 3.5 (Intel® GNA 3.5)
- Intel® Neural Processing Unit (Intel® NPU)
- Memory
 - DDR5
 - Error Code Correction (ECC)
 - In Band Error Code Correction (IBECC)
- Platform Environmental Control Interface (PECI)
- Intel® Volume Management Device (Intel® VMD)
- Integrated Clock Controller (ICC)/Integrated Reference Clock PLL
- Real Time Clock Controller (RTCC)
- General Purpose Input Output (GPIO)
- Virtualization
 - Intel® Virtualization Technology (Intel® VT-x)
 - Intel® Virtualization Technology for Directed I/O (Intel® VT-d)
 - Intel® APIC Virtualization Technology (Intel® APICv)
- Security Technologies
 - Intel® Trusted Execution Technology (Intel® TXT)
 - Intel® Converged Boot Guard and Intel® Trusted Execution Technology (Intel® CBnT)
 - Intel® Secure Key
 - Execute Disable Bit
 - Intel® Supervisor Mode Execution Protection (Intel® SMEP)

- Intel® Supervisor Mode Access Protection (Intel® SMAP)
- User Mode Instruction Prevention (UMIP)
- Read Processor ID (RDPID)
- Intel® Total Memory Encryption (Intel® TME)
- Intel® Multi-Key Total Memory Encryption (Intel® MK-TME)
- Intel® Control-flow Enforcement Technology (Intel® CET)
- Linear Address Space Separation (LASS)
- KeyLocker Technology
- Security Technologies - Crypto Acceleration Instructions
 - Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
 - Perform Carry-Less Multiplication Quad Word Instruction (PCLMULQDQ)
 - Intel® Secure Hash Algorithm - 512 (Intel® SHA - 512)
 - Intel® Secure Hash Algorithm Extensions (Intel® SHA Extensions)
- Security Technologies - Intel® System Resources Defense and Intel® System Security Report
- Security Technologies - Security Firmware Engines
 - Intel® Converged Security and Management Engine (Intel® CSME)
 - Intel® Silicon Security Engine
 - Intel® Graphics System Controller (Intel® GSC)
 - Intel® Active Management Technology (Intel® AMT)
- Testability and Monitoring
 - JTAG Boundary Scan
 - Intel® Software Toolkit
 - Platform Crashlog
 - Platform Monitoring Technology (PMT)
 - Intel® Processor Trace
 - Intel® Trace Hub (Intel® TH)
 - Direct Connect Interface (DCI) for debug
 - Debug Island
 - Early Boot Debug
- Power Management Technologies
 - Advanced Configuration and Power Interface (ACPI) Power Management Logic Support
 - Intel® Smart Cache Technology
 - Power and Efficient Cores Level 1 and Level 2 Caches
 - Cache Line Write Back (CLWB)
 - Intel® Thread Director
 - Intel® Hybrid Technology
 - Intel® Turbo Boost Technology 2.0

- Intel® Turbo Boost Max Technology 3.0
- Intel® Adaptive Boost Technology
- Intel® Thermal Velocity Boost (Intel® TVB)
- Intel SpeedStep® Technology
- Intel® System Agent Enhanced SpeedStep Technology (Intel® SAGV)
- Intel® Speed Shift Technology
- Intel® Advanced Vector Extensions 2 (Intel® AVX2)
- Intel® AVX2 Vector Neural Network Instructions (Intel® AVX2 VNNI)
- Intel® Advanced Programmable Interrupt Controller (APIC)
- Intel® 64 Architecture x2APIC
- Intel® Dynamic Tuning technology (Intel® DTT)
- Power Delivery Technologies
 - Digital Linear Voltage Regulator (DLVR)
 - Fast V-Mode (FVM)
 - Current Excursion Protection (CEP)

NOTE

The availability of the features above may vary between different processor SKUs.

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* 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.

1.3.2 Firmware Resiliency

Intel's NVMe based recovery supports recovery of all firmware on Intel® Core™ Ultra 200S Series Processors from NVMe storage boot partition in a secure manner.

Firmware Resiliency and Recovery in-field is critical to keep PCs up and running while preventing the requirement of additional space on SPI flash to keep a backup firmware. Therefore, it decreases the Platform BOM cost.

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, C2, C3, C6, C8, and C10

Refer to [Processor IA Core Power Management](#) on page 59 for more information.

1.4.2 System Power Management

Supports the following power management system states:

- Modern Standby (S0ix)
- S3
- S4
- S5

Refer to [Power Management](#) on page 57 for more information.

1.4.3 Memory Controller Power Management

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

Refer to [Integrated Memory Controller \(IMC\) Power Management](#) on page 116 for more information.

1.4.4 Processor Graphics Power Management

Memory Power Savings Technologies

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

Display Power Savings Technologies

- Intel® (Seamless and Static) Display Refresh Rate Switching (DRRS) with eDP* port
- Intel® Display Power Saving Technology (Intel® DPST 8.0)
- Intel® OLED Power Saving Technology (Intel® OPST) 1.1
- Panel Self-Refresh 2 (PSR 2)
- Low-Power Single Pipe (LPSP)
- Low-Power Dual Pipe (LPDP)

- Intel® Smart 2D Display Technology (Intel® S2DDT)
- Intel® Low Refresh Rate (LRR)

Graphics Core Power Savings Technologies


- Graphics Dynamic Frequency
- Intel® Graphics Render Standby Technology (Intel® GRST)
- Intel Capped Frames Per Second (CFPS)

1.5 Thermal Management Support

- Intel® Adaptive Thermal Monitor
- Digital Thermal Sensor
- THERMTRIP# and PROCHOT# support
- Critical Temperature Detection
- Software Controlled Clock Modulation (On-Demand Mode)
- Memory Thermal Throttling
- Render Thermal Throttling
- Fan Speed Control with DTS

Refer to Thermal Management chapter for more information.

1.6 Ballout Information

For information on the Intel® Core™ Ultra 200S Series Processors ball information, download the pdf, click  on the navigation pane and refer the spreadsheet, **832586-001_Ballout.xlsx**.

1.7 Processor Testability

The processor includes boundary-scan for board and system level testability. Refer to the appropriate processor Testability Information - Boundary Scan Description Language (BSDL) file.

1.8 Operating Systems Support

Processor Line	Windows* 10 (21H2/22H2)	Windows* 11 (22H2/23H2/24H2)	Chrome* OS	Linux* OS
S LGA	Yes	Yes	No	Yes

NOTE

Refer to OS Vendor site for more information regarding latest OS revision support.

1.9 Terminology and Special Marks

Table 2. Terminology

Term	Description
4K	Ultra High Definition (UHD)
AES	Advanced Encryption Standard
AGC	Adaptive Gain Control
AI	Artificial Intelligence
API	Application Programming Interface
AVC	Advanced Video Coding
BLT	Block Level Transfer
BPP	Bits per Pixel
CDR	Clock and Data Recovery
CTLE	Continuous Time Linear Equalizer
DDC	Digital Display Channel
DDI	Digital Display Interface for DisplayPort or HDMI/DVI
DSI	Display Serial Interface
DDR5	Fifth-Generation Double Data Rate SDRAM Memory Technology
DFE	Decision Feedback Equalizer
DMA	Direct Memory Access
DPPM	Dynamic Power Performance Management
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
GSA	Graphics in System Agent
GNA	Gaussian & Neural-Network Accelerator
HDCP	High-Bandwidth Digital Content Protection
HDMI*	High Definition Multimedia Interface
IMC	Integrated Memory Controller
Intel® 64 Technology	64-bit memory extensions to the IA-32 architecture
Intel® DPST	Intel® Display Power Saving Technology
Intel® PTT	Intel® Platform Trust Technology
Intel® TXT	Intel® Trusted Execution Technology
<i>continued...</i>	

Term	Description
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.
Intel® TH	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].
LLC	Last Level Cache
LPDDR5/x	Low Power Double Data Rate SDRAM memory technology /x- additional power save.
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.
MFM	Minimum Frequency Mode. MFM is the minimum ratio supported by the processor and can be read from MSR CEh [55:48].
MLC	Mid-Level Cache
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.
NPU	Neural Processing Unit
PECI	Platform Environment Control Interface
PEG	PCI Express* Graphics
PCH	Platform Controller Hub
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 tile 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)
<i>continued...</i>	

Term	Description
Rank	A unit of DRAM corresponding to four to eight devices in parallel, ignoring ECC. These devices are usually, but not always, mounted on a single side of a DIMM.
SCI	System Control Interrupt. SCI is used in the ACPI protocol.
SDP	Scenario Design Power
SHA	Secure Hash Algorithm
SSC	Spread Spectrum Clock
STR	Suspend to RAM
TAC	Thermal Averaging Constant
TBT	Thunderbolt™ Interface
TCC	Thermal Control Circuit
TMTV Processor Base Power	Thermal Test Vehicle Processor Base Power
VLD	Variable Length Decoding
VPID	Virtual Processor ID
V _{SS}	Processor Ground
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.
S0ix-states	Processor and PCH residency idle standby power states.

Table 3. 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 / #	A suffix of _N or # indicates an active low signal. For example, CATERR# _N does not refer to a differential pair of signals such as CLK_P, CLK_N
h	Hexadecimal numbers are identified with 'h' appended at the end of the number. For example, CF9h. All numbers are decimal (base 10) unless otherwise specified. Non-obvious binary numbers have 'b' appended at the end of the number. For example, 0101b.

1.10 Related Documents

Document	Document Number
Intel® Core™ Ultra 200S Series Processors Datasheet Volume 2 of 2	834966
Intel® 800 Series Chipset Family Platform Controller Hub (PCH) Datasheet, Volume 1 of 2	833778

2.0 Processor and Device IDs

2.1 CPUID

Table 4. CPUID Format

Compute Tile	Stepping	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]
S 8P + 16E	B-Step	C0662h	Reserved	00000000b	1100b	Reserved	00b	0110b	0110b	0010b

- 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 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

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.

2.2 PCI Configuration Header

Every PCI-compatible function has a standard PCI configuration header, as shown in the table below. 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.

NOTE

For more information on other device IDs, refer to [Table 6](#) on page 23, and [Table 7](#) on page 24.

Table 5. PCI Configuration Header

Byte3	Byte2	Byte1	Byte0	Address
Device ID		Vendor ID (8086h)		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

2.3 Device IDs

This section specifies the device IDs of the processor.

Table 6. Host Device ID (DID0) and Processor Graphics Device ID (DID2)

Processor Line	Package	P-Cores	E-Cores	Graphics Configuration X ^e -Cores	Host Device ID (DID0)	Processor Graphics Device ID (DID2)
S- Processor	LGA1851	8	16	4	7D1Ah	7D67h
	LGA1851	8	12	4	7D1Bh	7D67h
	LGA1851	8	12	N/A	7D1Bh	N/A
	LGA1851	6	8	4	7D2Ah	7D67h
	LGA1851	6	8	N/A	7D2Ah	N/A

Table 7. Other Device ID

Device	Bus / Device / Function	S-Processor DID
PCI Express* Root Port #12	0 / 1 / 0	7ECCh
Dynamic Tuning Technology (DTT)	0 / 4 / 0	AD03h
PCI Express Root Port #13	0 / 6 / 0	AE4Dh
PCI Express Root Port #10	0 / 6 / 1	7ECAh
PCI Express Root Port #14	0 / 6 / 3	AE4Eh
PCI Express Root Port #15	0 / 6 / 4	AE4Fh
USB Type-C Subsystem PCIe Root Port #16	0 / 7 / 0	7EC4h
USB Type-C Subsystem PCIe Root Port #17	0 / 7 / 1	7EC5h
Gauss Newton Algorithm (GNA)	0 / 8 / 0	AE4Ch
Crash Log & Telemetry	0 / 10 / 0	AD0Dh
NPU	0 / 11 / 0	AD1Dh
USB xHCI	0 / 13 / 0	7EC0h
USB xDCI	0 / 13 / 1	7EC1h
Thunderbolt™ DMA0	0 / 13 / 2	7EC2h
Intel® Volume Management Device (VMD)	0 / 14 / 0	AD0Bh
P2SB (IOE)	0 / 19 / 0	7EC8h
PMC (IOE)	0 / 19 / 2	7ECEh
Shared SRAM (IOE)	0 / 19 / 3	7ECFh
Shared SRAM (SOC-S)	0 / 20 / 0	AE7Fh
Intel® CSME: HECI #1	0 / 22 / 0	AE70h
Intel® CSME: HECI #2	0 / 22 / 1	AE71h
Intel® CSME: HECI #3	0 / 22 / 4	AE74h
Direct eSPI Controller	0 / 31 / 0	AE00h - AE1Fh
P2SB (SOC-S)	0 / 31 / 1	AE20h
PMC (SOC-S)	0 / 31 / 2	AE21h
SMBus	0 / 31 / 4	AE22h
SPI (flash) Controller	0 / 31 / 5	AE23h
Intel® Trace Hub (Intel® TH)	0 / 31 / 7	AE24h

Table 8. ACPI Device ID for GPIO Controller

ACPI ID	Note
S LGA	INTC1082

2.4 Revision IDs

The Revision ID (RID) register is an 8-bit register located at offset 08h in the PCI header of every PCI/PCIe* function. The RID register is used by software to identify a particular component stepping when a driver change or patch unique to that stepping is needed.

3.0 Package Mechanical Specifications

3.1 Package Mechanical Attributes

The S LGA Processor Lines use a Flip Chip technology available in a Land Grid Array (LGA) package. The following table provides an overview of the package mechanical attributes.

Table 9. S LGA Package Mechanical Attributes

Package	Parameter	S Processor Line
Package Technology	Package Type	Flip Chip Land Grid Array
	Interconnect	Land Grid Array (LGA)
	Lead Free	Yes
	Halogenated Flame Retardant Free	Yes
Package Configuration	Solder Ball Composition	SAC405
	Ball/Pin Count	1851
	Grid Array Pattern	Grid Array
	Land Side Capacitors	Yes
	Tile Side Capacitors	Yes
	Tile Configuration	Foveros
Package Dimensions	Nominal Package Size	45 x 37.5 mm
	Maximum Package Z-Height	4.464 ± 0.133 mm
	Minimum Ball/Pin pitch	0.8 mm

3.2 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	
T _{SUSTAINED STORAGE}	The ambient storage temperature limit (in shipping media) for the sustained period of time	-5°C	40°C	
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		
<i>continued...</i>				

Parameter	Description	Minimum	Maximum	Notes
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	
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.			
<p><i>Notes:</i></p> <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. Consult your board manufacturer for storage specifications. 				

4.0 Memory Mapping

This chapter describes (from the processor perspective) the memory ranges that the Processor decodes.

4.1 Functional Description

4.1.1 PCI Devices and Functions

The Processor incorporates a variety of PCI devices and functions, as shown in the following table. If for some reason, the particular system platform does not want to support any one of the Device Functions, with the exception of D30:F0, they can individually be disabled. The integrated Gigabit Ethernet controller will be disabled if no Platform LAN Connect component is detected. When a function is disabled, it does not appear to the software. A disabled function will not respond to any register reads or writes, ensuring that these devices appear hidden to software.

4.1.2 Fixed I/O Address Ranges

The following table shows the Fixed I/O decode ranges from the processor perspective.

NOTE

For each I/O range, there may be separate behavior for reads and writes.

I/O cycles that go to target ranges that are marked as Reserved will be handled as follow : writes are ignored and reads will return all 1's. The P2SB will claim many of the fixed I/O accesses and forward those transactions over IOSF-SB to their functional target.

Address ranges that are not listed or marked Reserved are NOT positively decoded (unless assigned to one of the variable ranges) and will be internally terminated.

Table 10. Fixed I/O Ranges Decoded by Processor

I/O Address	Read Target	Write Target	Internal Unit (Unless[E]: External) ²	Separate Enable/Disable
20h – 21h	Interrupt Controller	Interrupt Controller	Interrupt	None
24h – 25h	Interrupt Controller	Interrupt Controller	Interrupt	None
28h – 29h	Interrupt Controller	Interrupt Controller	Interrupt	None
2Ch – 2Dh	Interrupt Controller	Interrupt Controller	Interrupt	None
30h – 31h	Interrupt Controller	Interrupt Controller	Interrupt	None
34h – 35h	Interrupt Controller	Interrupt Controller	Interrupt	None
<i>continued...</i>				

I/O Address	Read Target	Write Target	Internal Unit (Unless[E]: External) ²	Separate Enable/Disable
38h – 39h	Interrupt Controller	Interrupt Controller	Interrupt	None
3Ch – 3Dh	Interrupt Controller	Interrupt Controller	Interrupt	None
40h	Timer/Counter	Timer/Counter	8254 Timer	None
42h-43h	Timer/Counter	Timer/Counter	8254 Timer	None
50h	Timer/Counter	Timer/Counter	8254 Timer	None
52h-53h	Timer/Counter	Timer/Counter	8254 Timer	None
61h	NMI Controller	NMI Controller	CPU I/F	None
63h	NMI Controller ¹	NMI Controller ¹	CPU I/F	Yes, alias to 61h. GIC.P61AE
65h	NMI Controller ¹	NMI Controller ¹	CPU I/F	Yes, alias to 61h. GIC.P61AE
67h	NMI Controller ¹	NMI Controller ¹	CPU I/F	Yes, alias to 61h. GIC.P61AE
70h	RTC Controller	NMI and RTC Controller	RTC	None
80h ³	eSPI or PCIe	eSPI or PCIe	Read: [E] eSPI or PCIe Write: [E] eSPI or [E] PCIe	None. PCIe if GCS.RPR='1', else eSPI
84h - 86h	eSPI or PCIe	eSPI or PCIe	Read: [E] eSPI or PCIe Write: [E] eSPI or [E] PCIe	None. PCIe if GCS.RPR='1', else eSPI
88h	eSPI or PCIe	eSPI or PCIe	Read: [E] eSPI or PCIe Write: [E] eSPI or [E] PCIe	None. PCIe if GCS.RPR='1', else eSPI
8Ch - 8Eh	eSPI or PCIe	eSPI or PCIe	Read: [E] eSPI or PCIe Write: [E] eSPI or [E] PCIe	None. PCIe if GCS.RPR='1', else eSPI
92h	Reset Generator	Reset Generator	CPU I/F	None
A0h - A1h	Interrupt Controller	Interrupt Controller	Interrupt	None
A4h - A5h	Interrupt Controller	Interrupt Controller	Interrupt	None
A8h - A9h	Interrupt Controller	Interrupt Controller	Interrupt	None
ACh - ADh	Interrupt Controller	Interrupt Controller	Interrupt	None
B0h - B1h	Interrupt Controller	Interrupt Controller	Interrupt	None
B2h - B3h	Power Management	Power Management	Power Management	None

continued...

I/O Address	Read Target	Write Target	Internal Unit (Unless[E]: External) ²	Separate Enable/Disable
B4h - B5h	Interrupt Controller	Interrupt Controller	Interrupt	None
B8h - B9h	Interrupt Controller	Interrupt Controller	Interrupt	None
BCh - BDh	Interrupt Controller	Interrupt Controller	Interrupt	None
4D0h - 4D1h	Interrupt Controller	Interrupt Controller	Interrupt Controller	None
CF9h	Reset Generator	Reset Generator	Interrupt controller	None

Notes: 1. Only if the Port 61 Alias Enable bit (GIC.P61AE) bit is set. Otherwise, the cycle is internally terminated by the Processor.
 2. Destination of eSPI when eSPI Disabled pin strap is 0.
 3. This includes byte, word or double-word (DW) access at I/O address 80h.

4.1.3 Variable I/O Decode Ranges

The following Table shows the Variable I/O Decode Ranges. They are set using Base Address Registers (BARs) or other configuration bits in the various configuration spaces. The PnP software (PCI or ACPI) can use their configuration mechanisms to set and adjust these values.

WARNING

The Variable I/O Ranges should not be set to conflict with the Fixed I/O Ranges. There may be some unpredictable results if the configuration software allows conflicts to occur. The Processor does not perform any checks for conflicts.

Table 11. Variable I/O Decode Ranges

Range Name ¹	Mappable	Size (Bytes)	Target
ACPI	Anywhere in 64K I/O Space	256	Power Management
IO Trapping Ranges	Anywhere in 64K I/O Space	1 to 256 Bytes	Trap
DMI General Purpose I/O Ranges (1 to 3)	Anywhere in 64K I/O Space	4 to 256 Bytes	General Purpose

Note: All ranges are decoded directly from DMI.

4.2 Memory Map

The following table shows (from the processor perspective) the memory ranges that the processor will decode. Cycles that are not directed to any of the internal memory targets, will be host aborted.

PCIe cycles generated by external PCIe hosts will be positively decoded unless they fall in the PCI-PCI bridge memory forwarding ranges (those addresses are reserved for PCI peer-to-peer traffic). Software must not attempt locks to the processor’s memory-mapped I/O ranges.

NOTE

Total ports are different for the different SKUs.

Table 12. Processor Memory Decode Ranges (Processor Perspective)

Memory Range	Target	Dependency/Comments
FECX X000 - FECX X040	I/O(x)APIC inside SOC-S, or behind SOC-S's PCIe root-port	XX controlled via APIC Range Select (ASEL) field and APIC Enable (AEN) bit. For PCIe root port, I/OxApic Enable (PAE) bit
FFFC 0000 - FFFF FFFF	Intel® CSME	Always enabled. Refer to Section 6.3 on the Top-Block Swap
FED0 X000 - FED0 X3FF	HPET	BIOS determines "fixed" location which is one of four 1KB ranges where X (in the first column) is 0h, 1h, 2h, or 3h.
FED4 8000 - FED4 BFFF	LPC	Always enabled - LT address range NOTE: the LPC "bus" interface is removed, but the LPC HW is retained to support this register set.
FED4 C000 - FED4 FFFF	Internal (PSF Error Handler)	Always enabled
FED5 0000 - FED5 FFFF	ESE++	Always enabled
FED6 0000 - FED6 1FFF	xHCI	NOT positively decoded in PCH(DMI/PSF)
FED6 2000 - FED6 3FFF	xHCI (CPU)	Fixed range in CPU - never forwarded to PCH. First implemented in ICL
FED7 0000 - FED7 4FFF	Internal Device	Security feature related
DMI General Purpose Memory ranges (1 to 3)	General purpose	Enable via setting the Decode Enable bit of the respective Memory Range register
MMIO resources for VMD managed devices	Storage devices	Enable through the VMD device in SOC

4.2.1 Boot Block Update Scheme

The Processor supports a "Top-Block Swap" mode that has the Processor swap the top block in the FWH or SPI flash (the boot block) with another location. This allows for safe update of the Boot Block (even if a power failure occurs). When the "top-swap" enable bit is set, the Processor will invert A16 for cycles going to the upper two 64-KB blocks in the FWH or appropriate address lines as selected in Boot Block Size (BOOT_BLOCK_SIZE) soft strap for SPI.

For FWH when top swap is enabled, accesses to FFFF_0000h-FFFF_FFFFh are directed to FFFE_0000h-FFE_FFFFh and vice versa. When the Top Swap Enable bit is 0, the Processor will not invert A16.

For SPI when top swap is enabled, the behavior is as described below. When the Top Swap Enable bit is 0, the Processor will not invert any address bit.

Table 13. Boot Block Update Scheme

BOOT_BLOCK_SIZE Value	Accesses to	Being Directed to
000 (64KB)	FFFF_0000h - FFFF_FFFFh	FFFE_0000h - FFFE_FFFFh and vice versa
001 (128KB)	FFFE_0000h - FFFF_FFFFh	FFFC_0000h - FFFD_FFFFh and vice versa
010 (256KB)	FFFC_0000h - FFFF_FFFFh	FFF8_0000h - FFFB_FFFFh and vice versa
011 (512KB)	FFF8_0000h - FFFF_FFFFh	FFF0_0000h - FFF7_FFFFh and vice versa
100 (1MB)	FFF0_0000h - FFFF_FFFFh	FFE0_0000h - FFEF_FFFFh and vice versa
101 - 111	Reserved	Reserved
<i>Note:</i> This bit is automatically set to 0 by RTEST#, but not by PLTRST#.		

The scheme is based on the concept that the top block is reserved as the “boot” block, and the block immediately below the top block is reserved for doing boot-block updates.

The algorithm is:

1. Software copies the top block to the block immediately below the top
2. Software checks that the copied block is correct. This could be done by performing a checksum calculation.
3. Software sets the “Top-Block Swap” bit. This will invert the appropriate address bits for the cycles going to the FWH or the SPI.
4. Software erases the top block
5. Software writes the new top block
6. Software checks the new top block
7. Software clears the top-block swap bit
8. Software sets the Top_Swap Lock-Down bit

If a power failure occurs at any point after step 3, the system will be able to boot from the copy of the boot block that is stored in the block below the top. This is because the top-swap bit is backed in the RTC well.

There is one remaining unusual case that could occur if the RTC battery is not sufficiently high to maintain the RTC well. To avoid the potentially fatal case (where the Top-Swap bit is NOT set, but the top block is not valid), a pin strap will allow forcing the top-swap bit to be set. This would be a last resort to allow the user to get the system to boot (and avoid having to de-solder the system flash).

When the top-swap strap is used, the top-swap bit will be forced to 1 (cannot be cleared by software).

5.0 Security Technologies

5.1 Intel® Converged Boot Guard and Intel® TXT

Intel® Converged Boot Guard and Intel® TXT (Intel® CBnT) is a unification of Intel® Trusted Execution Technology (Intel® TXT) and Intel® Platform Protection Technology with Intel® Boot Guard. Intel® CBnT merges elements of Intel® TXT and Intel® Boot Guard to enhance platform boot security, while also simplifying the implementation. Although Intel® CBnT implements some architectural changes, it is not fundamentally a new technology, but rather a fusion of existing Intel® Boot Guard and Intel® TXT technologies.

Intel® CBnT has been designed to allow greater commonality between implementations for client platforms and server platforms. Previously, the architectural implementation of Intel® TXT was somewhat different between client and server platforms, which necessitated some differences in BIOS implementation depending on the platform. With Intel® CBnT, Intel has largely combined features across client and server providing greater alignment in design of the BIOS and ACMs.

Intel® Converged Boot Guard and Intel® TXT provides both a static root of trust for verifying the BIOS initial boot block and measuring the boot path, as well as a dynamic root of trust for measuring the OS or VMM.

The purpose of Intel® Boot Guard is to verify that the initial BIOS startup code is good, i.e., BIOS has not been maliciously nor inadvertently modified. Several different Boot Profiles are supported, which primarily differ in:

- **Enforcement Policy:** what actions are taken if BIOS cannot be verified.
- **Measurement Policy:** whether BIOS startup code is measured into the TPM for attestation.

The primary objective of Intel® TXT is to provide a dynamic root of trust for measuring the OS or VMM enabling platform boot into a secure measured launch environment (MLE). Intel® TXT relies on the static root of trust provided by Intel® Boot Guard to ensure validity of the MLE Trusted Compute Base (TCB), which is the BIOS code that is trusted to configure the platform. Intel® TXT provides the ability to allow only a known good OS/VMM to launch into a trusted environment via a Launch Control Policy (LCP). And once an OS/VMM is in a trusted environment, Intel® TXT protects memory secrets against surprise reset attacks.

With the modifications made to the Intel® TXT architecture in Intel® CBnT, it is now required that some of the verifications performed by Intel® Boot Guard be implemented for Intel® TXT support. Verifications of pre-boot objects such as FIT, key and policy manifests, and of Startup BIOS.

Still formally all four combinations of constituent technologies are supported at OEM choice:

- Intel® Boot Guard only enabled.
- Intel® TXT only enabled.

- Both Intel® Boot Guard and Intel® TXT enabled.

5.2 Crypto Acceleration Instructions

5.2.1 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 a great deal of flexibility.

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

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

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

NOTE

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

5.2.2 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 the *Intel® 64 Architectures Software Developer's Manual, Volume 2*. Available at:

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

5.2.3 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 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.

Information on Intel® SHA can be found at: <http://software.intel.com/en-us/artTGLes/intel-sha-extensions>

5.2.4 New Cryptographic Acceleration Instructions

The processor supports new extensions for acceleration of some common or emerging cryptographic algorithms:

1. AVX2 version of VPADD52 for acceleration of RSA signature verification
2. SHA2-512 (or 384)
3. Chinese crypto standards SM3 and SM4

5.3 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 and RDSEED instructions. 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 and RDSEED instructions include cryptographic key generation as used in a variety of applications, including communication, digital signatures, secure storage, etc.

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

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

5.4 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.

5.5 Intel® Supervisor Mode Execution Prevention (Intel® SMEP)

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

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

5.6 Intel® Supervisor Mode Access Prevention (Intel® SMAP)

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

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

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

5.7 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 the *Intel® 64 Architectures Software Developer's Manual, Volume 3*. Available at:

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

5.8 Read Processor ID (RDPID)

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

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

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

5.9 Intel® Total Memory Encryption - Multi-Key

This technology encrypts the platform's entire memory with multiple encryption keys. Intel® Total Memory Encryption (Intel® TME), when enabled via BIOS configuration, ensures that all memory accessed from the Intel processor is encrypted.

Intel TME encrypts memory accesses using the AES XTS algorithm with 256-bit keys. The global encryption key used for memory encryption is generated using a hardened random number generator in the processor and is not exposed to software.

Software (OS/VMM) manages the use of keys and can use each of the available keys for encrypting any page of the memory. Thus, Intel® Total Memory Encryption - Multi-key (Intel® TME-MK) allows page granular encryption of memory. By default Intel TME-MK uses the Intel TME encryption key unless explicitly specified by 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 Intel TME. Intel TME does not protect memory from modifications.

Intel TME allows the BIOS to specify a physical address range to remain unencrypted. Software running on Intel 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.

NOTES

- Memory access to nonvolatile memory (Intel® Optane™) is encrypted as well.
 - More information on Intel TME-MK can be found at:
<https://software.intel.com/sites/default/files/managed/a5/16/Total-Memory-Encryption-Multi-Key-Spec.pdf>
 - A cold boot is required when enable/ disable Intel TME feature on this platform.
-

5.10 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:

5.10.1 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, etc. 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, etc. are not allowed to the shadow stack.

5.10.2 Indirect Branch Tracking

The ENDBR32 and ENDBR64 (collectively ENDBRANCH) are two 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 exception (#CP), otherwise the state machine moves back to IDLE state.

More information on Intel® CET can be found at Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 1, Chapter 18:

<https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html>

5.11 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.

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 encrypt and decrypt operations. Once a handle is obtained, the software can delete the original key from memory.

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

NOTE

KeyLocker Technology may not be available on all SKUs.

5.12 Intel® System Resources Defense and Intel® System Security Report

Intel® System Resources Defense is the collection of techniques and code within the BIOS used to create and enforce HW access policy for the SMI handler. It consists of a collection of policy mechanisms that are configured by POST before the SMI handler is locked down. Once the SMI handler is locked, all accesses into the system must be compliant with the policy established during POST.

Intel® Runtime BIOS Resilience is a subset of Intel® System Resources Defense covering SMM memory policy only. Intel® Runtime BIOS Resilience Protection hardens the SMI handler via hardware enforced BIOS policy regarding SMI handler access to memory using an enhanced paging policy. This paging policy covers SMI handler access to both BIOS and MLE resources. Intel® Runtime BIOS Resilience Protection is extended using a technology codenamed Intel® System Security Report.

The Platform Properties Assessment Module (PPAM) is the primary component of Intel® System Security Report. It collects and reports information about platform SMM implementation and configuration, in order to provide trustworthy attestation of the resulting SMI memory policy regarding SMM secure configuration and access to MLE owned memory. Intel® System Security Report is used to create a trustworthy report describing the SMM policy. PPAM is a major/core component of Intel® System Security Report 1.0/1.1 technology

5.13 BIOS Guard

The platform must implement hardware controls to provide the platform manufacturer a robust mechanism to prevent unauthorized flash updates, while still allowing platform manufacturer approved updates. Intel® Platform Protection Technology with BIOS Guard accomplishes this by providing a very robust environment from which signed update images can be cryptographically verified and host flash writes can be done. Furthermore, a BIOS Guard enabled system does not allow host flash writes from any other environment.

5.14 Intel® Platform Trust Technology

Intel® Platform Trust Technology (Intel® PTT) offers the capabilities of discrete TPM 2.0. Intel PTT is a platform functionality for credential storage and key management used by Windows* 11. Intel PTT supports BitLocker* for hard drive encryption and supports all Microsoft* requirements for Trusted Platform Module (TPM) 2.0.

5.15 Linear Address Space Separation (LASS)

Linear Address Space Separation (LASS) can harden an OS kernel against specific classes of side channel exploit techniques.

5.16 Security Firmware Engines

5.16.1 Intel® Silicon Security Engine

A Security engine which is HW IP is based on CSxE HW IP and new FW IP design to be silicon Root of Trust providing secure FW loading, measurements and on-tile certification authority.

The firmware is based on a new design which focus on security, simplicity of architecture and isolated environment.

5.16.2 Intel® Graphics System Controller (Intel® GSC)

Intel® Graphics System Controller (Intel® GSC) is a HW IP block embedded within the media IP block of the graphics component to support content and display protection services such as DRM and HDCP.

NOTE

All graphics security functionalities are handled by GSC which was previously implemented by CSxE.

6.0 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.

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/>.

6.1 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 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.

- **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 nonroot 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 feature allows detection of such loops and is thus called PAUSE-loop exiting.

- **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.

- **Virtualization Exceptions**

A virtualization exception is a new processor exception. It uses vector 20 and is abbreviated #VE. A virtualization exception can occur only in VMX non-root operation. Virtualization exceptions occur only with certain settings of certain VM-execution controls. Generally, these settings imply that certain conditions that would normally cause VM exits instead cause virtualization exceptions

- **Translation of Guest-Physical Addresses Used by Intel Processor Trace**

With the "Intel PT uses guest physical addresses" feature, the addresses used by Intel PT can be treated as guest-physical addresses and translated using EPT. These addresses include the addresses of the output regions as well as the addresses of the ToPA entries that contain the output-region addresses.

6.2 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. 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 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 to *Intel® Virtualization Technology for Directed I/O Architecture Specification* <http://www.intel.com/content/dam/www/public/us/en/documents/product-specifications/vt-directed-io-spec.pdf>

Intel® VT-d Key Features

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

- Memory controller and processor graphics comply with the Intel® VT-d 2.1 Specification.
- Two Intel® VT-d DMA remap engines.
 - iGFX DMA remap engine
 - Default DMA remap engine (covers all devices except iGFX)
- 46-bit guest physical address and host physical address widths
- 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 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
- Interrupt Remapping is supported
- Queued invalidation is supported
- 4-level Intel®VT-d Page walk - all VTd engines support 4-level tables only (adjusted guest address width of 48 bits)
- Intel®VT-d super-page - all VTd engines support super-page (2 MB, 1 GB)
- Scalable Mode - all VTd engines support Scalable mode operation (using RID_PASID only)
- Nested - default Intel® VT-d engine support Nested translation

NOTE

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

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

NOTE

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

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

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

7.0 Instructions Set Enhancements

7.1 CMPccXADD

CMPccXADD is a new set of instructions that can be used to optimize certain highly contended synchronization scenarios that today use CMPXCHG (semaphores, shared queues, etc).

7.2 LAM

LAM repurposes the upper (untranslated) linear address bits to make them available for software use (for example, as metadata). This is accomplished by removing exiting canonical address checks when LAM is enabled.

In 64-bit mode, linear address have 64 bits and are translated either with 4-level paging, which translates the low 48 bits of each linear address, or with 5-level paging, which translates 57 bits. The upper linear-address bits are reserved through the concept of canonicity. A linear address is 48-bit canonical if bits 63:47 of the address are identical; it is 57-bit canonical if bits 63:56 are identical. (Clearly, any linear address that is 48-bit canonical is also 57-bit canonical.)

When 4-level paging is active, the processor requires all linear addresses used to access memory to be 48-bit canonical; similarly, 5-level paging ensures that all linear addresses are 57-bit canonical.

Software usages that associate metadata with a pointer might benefit from being able to place metadata in the upper (untranslated) bits of the pointer itself. However, the canonicity enforcement mentioned earlier implies that software would have to mask the metadata bits in a pointer (making it canonical) before using it as a linear address to access memory or alternatively create Paging translation tables with redundancies.

LAM allows software to use pointers with metadata without having to mask the metadata bits. A LAM enabled processor internally ignores the metadata bits in a pointer before using it as a linear address to access memory.

NOTE

LAM is supported only in 64-bit mode and applies only to addresses used for data accesses. LAM does not apply to addresses used for instruction fetches or to those that specify the targets of jump and call instructions.

8.0 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, Assured Power (cTDP), and Memory Throttling Control mechanisms and many other services. PEFI is used for platform thermal management and real-time control and configuration of processor features and performance.

NOTE

- PEFI over eSPI is supported.
-

8.1 PEFI Bus Architecture

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

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

NOTE

PEFI supported frequency range is 100 Khz-1 MHz.

The following figures demonstrate PEFI design and connectivity:

- PEFI Host-Clients Connection: While the host/originator can be third party PEFI host and one of the PEFI clients is a processor PEFI device.
- PEFI EC Connection.

Figure 3. PECI Host-Clients Connection Example





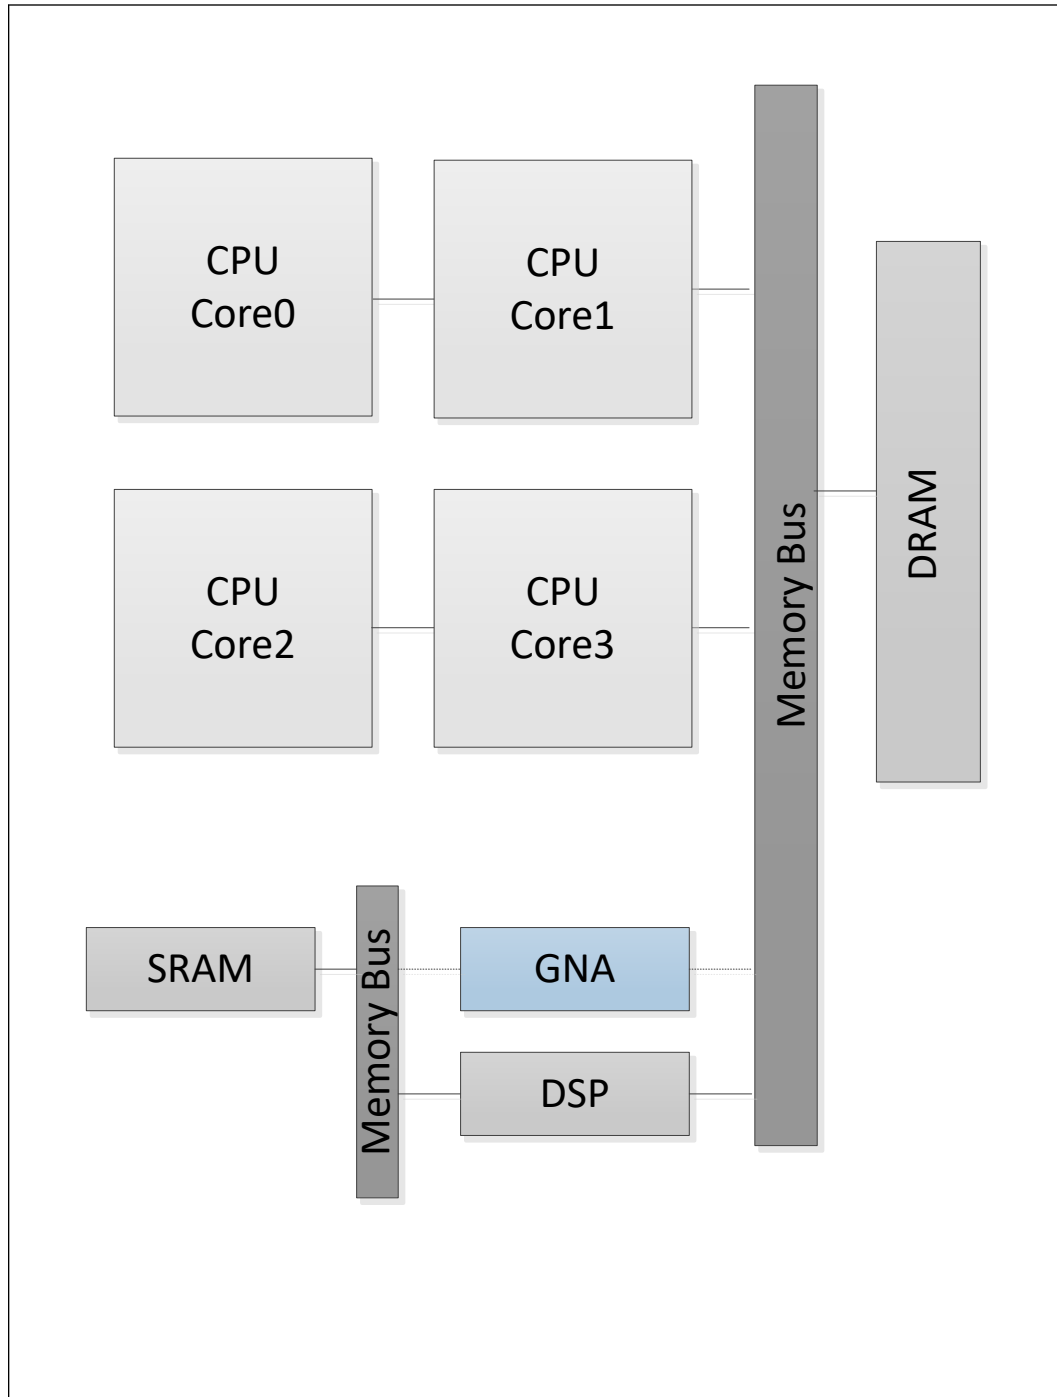
Figure 4. PEFI EC Connection Example



9.0 Intel GMM and Neural Network Accelerator (Intel GNA 3.5)

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.



10.0 Intel® Neural Processing Unit (Intel® NPU)

The NPU IP in the Intel® Core™ Ultra 200S Series Processors configuration is a Deep Learning accelerator enumerated to a host processor as an integrated PCIe device. It delivers the cutting-edge processing throughput required to satisfy the demands of Deep Learning applications. The NPU technology is applicable to personal computing devices such as tablets and PCs as a way to encourage AI based applications and services on power and performance sensitive platforms.

The functionality of the Intel® NPU is exposed to a Host system (enumerated as a PCIe device) via a base set of registers. These registers provide access to control and data path interfaces and reside in the Host and Processor subsystems of the Intel® NPU. All host communications are consumed by the scheduler of the Intel® NPU, a 32-bit LeonRT micro-controller. The LeonRT manages the command and response queues as well as the runtime management of the IP itself.

The NPU IP Deep Learning capability is provided by two Neural Compute Engine (NCE) Tiles. Both NCE Tiles are managed by the NPU Scheduler. Each Tile includes a configurable number of Multiply Accumulate (MAC) engines, purpose built for Deep Learning workloads, and two Intel® Movidius SHAVE DSP processors for optimal processing of custom deep learning operations.

The Intel® NPU of Intel® Core™ Ultra 200S Series Processors is configured with 2k MACs per tile totaling 4k MACs across both tiles and 4 MB of associated **near compute** memory.

The NPU plugin supports the following data types as inference precision of internal primitives: **INT8(I8/U8), FP16**.

10.1 Functional Description


NPU IP comprises 3 subsystems, as follows:

- Processor subsystem
- Host subsystem
- NCE subsystem

Apart from the subsystems, it has a Host interface for data exchange with the system memory. Details of these blocks are provided in the next sections.

Below is the block diagram of NPU IP:

Figure 5. NPU IP Block Diagram



10.1.1 Processor Subsystem

This subsystem provides the SW services through which all functions of the NPU are accessed. Those services are provided by firmware executing on the LeonRT Processor. The LeonRT is the first core out of reset in the NPU (before both the LeonNN and SHAVE cores).

10.1.2 NCE Subsystem

The Neural Compute Subsystem is a hardware accelerator for Deep Neural Network (DNN) workloads. It features a highly configurable pipeline to support DNN (Deep Neural Network) operations such as CNN (Convolutional Neural Networks), LSTM (Long Short-Term memory) and LRN (Local Response Norm). It also leverages activation and weight sparsity optimal performance.

Neural Compute Subsystem is built from up to 2 NCE Tiles (fixed) where each Tile is a primary unit of compute. Each Tile can support 2K Multiply Accumulate circuits (MACs) and two Activation SHAVE Engines (ACTShave). Tiles can be deployed to operate independently across multiple networks (threads) or be aggregated to form a multi cluster engine processing a single network (thread). Refer to the diagram below showing the 4K4M configuration.

NCE Subsystem supports two DMA engines. Each engine supports in-line weight decompression and write data broadcast capability into the local Connection Matrix (CMX) memory (dedicated SRAM).

For hardware assisted task synchronization, the NCE Subsystem provides barriers and workload FIFOs. Barriers remove as much software overhead as possible through ISR loops and programming sequences to keep the computing and data-movement pipelines full.

10.1.2.1 Some NCE Subsystem Features

- Dedicated real-time scheduler for job dispatching to DPU and Activation SHAVE engines. This is a LEON core (LeonNN) executing to two levels of cache.
- Two NCE Tiles with 2K MACs per tile.
- Activation SHAVE processors to support custom activation functions. These are vectorized processing units with a 128 bit data bus.
- 2MB of dedicated SRAM memory per tile

10.1.2.2 NCE Tile

The NCE Tile is the building block of the NCE Subsystem. The NCE subsystem supports a fixed two tile configuration. Each NCE tile supports the following:

- Single Data Processing Units (DPU) that supports 2048 MACs built from 512 MAC Processing Engines (MPE) with 4MACs in each MPE.
- An NCE Tile is capable of delivering:
 - 4 TOPS (8-bit) or 2 TFLOPS(FP16) @ 1 GHz¹ DPU Clock Frequency for a single DPU configuration

NOTE

1. 1 GHz is not the maximum frequency of DPU.
-

- Two ACT-SHAVE DSP with shared data and instruction L2 Cache used for flexible tensor compute operation.

10.1.2.3 ACT-SHAVE

ACT-SHAVE is DSP Processor which supports 128 bit vector operations. Two of ACT-SHAVE DSPs are placed in each NCE Tile and are used for custom layer and standard layers that do not map well to the DPU. All ACT-SHAVE DSP functions shall be included in the graph-file and barriers shall be used for HW Synchronization of the DSP operation and the rest of the schedule.

11.0 Power Management

The Power Management Controller (PMC) is the PCH unit that handles all PCH power management related activities. For more information, refer to Intel® 800 Series Chipset Family Platform Controller Hub (PCH) Datasheet, Volume 1 of 2 (#833778)

NOTE

In this chapter, Sx refers to S4/S5 states.

Table 14. Acronyms

Acronyms	Description
VR	Voltage Regulator

Table 15. References

Specification	Location
Advanced Configuration and Power Interface (ACPI)	https://uefi.org/sites/default/files/resources/ACPI_Spec_6_5_Aug29.pdf

11.1 System Power States, Advanced Configuration and Power Interface (ACPI)

This section describes System Power States and ACPI states supported by the processor.

For more information, refer to Intel® 800 Series Chipset Family Platform Controller Hub (PCH) Datasheet, Volume 1 of 2 (#833778).

Table 16. General System Power States

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
G0/S0ix/Cx	S0ix: The south supports an S0ix state which also requires the CPU be in a Cx state. Additional south power actions such as voltage reduction, chip-wide voltage rail removal may occur in this state.
<i>continued...</i>	

State	Description
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.
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.
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.

The table below shows the transitions rules among the various states.

NOTE

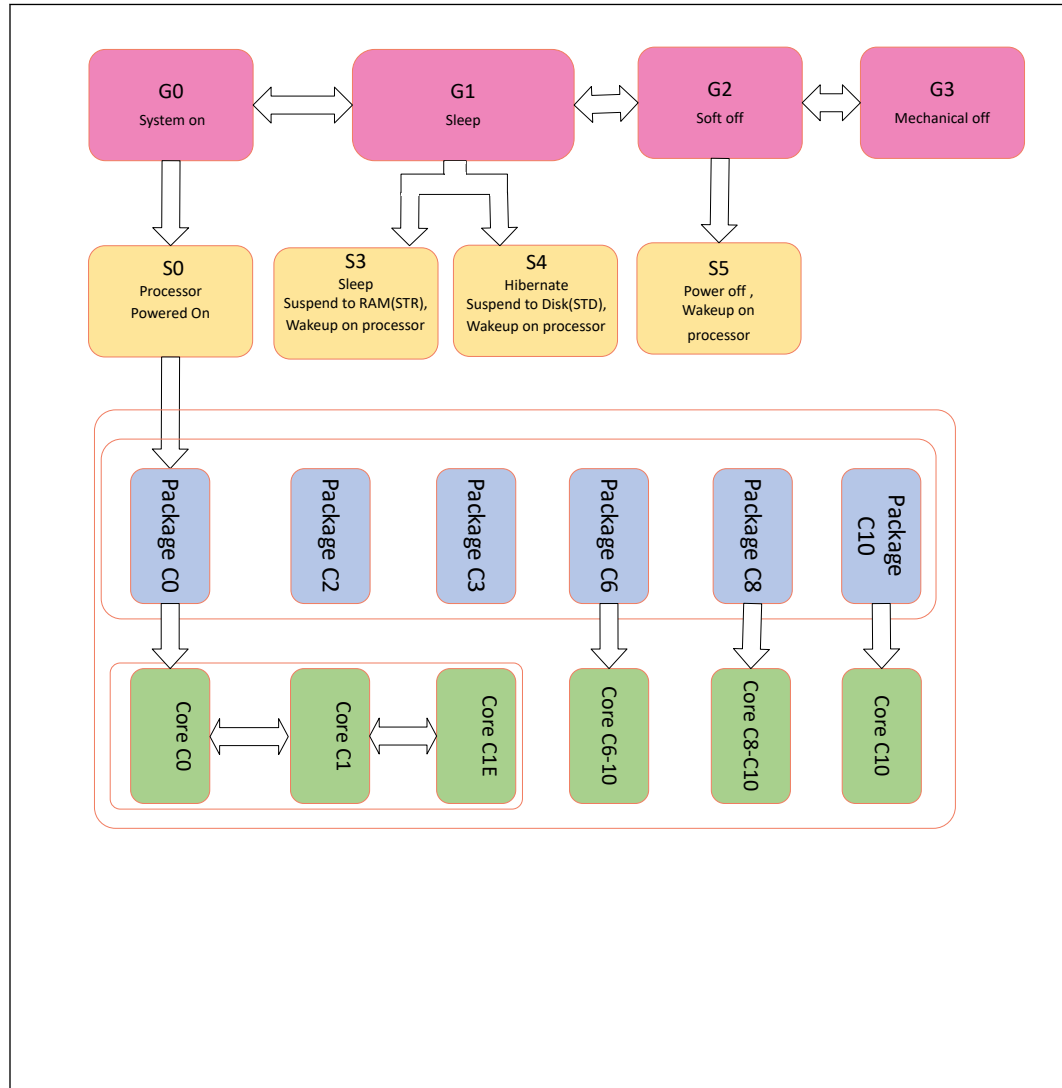
Transitions among the various states may appear to temporarily transition through intermediate states. For example, in going from S0 to S5, it may appear to pass through the G1/S4 state. These intermediate transitions and states are not listed in the table below.

Table 17. State Transition Rules for the Processor

Present State	Transition Trigger	Next State
G0/S0/C0	<ul style="list-style-type: none"> • SLP_EN bit set • Power Button Override³ • Mechanical Off/Power Failure 	<ul style="list-style-type: none"> • G0/S0/Cx • G1/S4, or G2/S5 state • G2/S5 • G3
G0/S0/Cx	<ul style="list-style-type: none"> • Power Button Override³ • Mechanical Off/Power Failure 	<ul style="list-style-type: none"> • G0/S0/C0 • S5 • G3
G0/S0ix/Cx	<ul style="list-style-type: none"> • Any south action which is blocked from occurring while in S0ix • CPU or south IP request for CPU C-state exit 	<ul style="list-style-type: none"> • G0/S0ix/Cx • G0/S0/Cx • G0/S0/C2(or C0)
G1/S3, G1/S4	<ul style="list-style-type: none"> • Any Enabled Wake Event • Power Button Override³ • Mechanical Off/Power Failure 	<ul style="list-style-type: none"> • G0/S0/C0² • G2/S5 • G3
G2/S5	<ul style="list-style-type: none"> • Any Enabled Wake Event • Mechanical Off/Power Failure 	<ul style="list-style-type: none"> • G0/S0/C0² • G3
G3	<ul style="list-style-type: none"> • Power Returns 	<ul style="list-style-type: none"> • S0/C0 (reboot) or G2/S5⁴ (stay off until power button pressed or other wake event)^{1,2}

Notes: 1. Some wake events can be preserved through power failure.
2. Transitions from the S4-S5 states to the S0 state are deferred until BATLOW# is inactive.
3. Includes all other applicable types of events that force the host into and stay in G2/S5.
4. If the system was in G1/S4 before G3 entry, then the system will go to S0/C0 or G1/S4.
5. On G3 exit, prior to the first transition to S0, S5 power may be higher than S5 power after the first S0 to S5 transition.
Some processor settings required to achieve minimum S5 power are loaded during first boot to S0 after a G3 exit. Consequently, entry into S5 from S0 will result in a more power-optimized S5 state than entry into S5 from G3 without an S5-S0-S5 transition. The difference is expected to be in the few mW range

Figure 6. Power State Block Diagram



11.2 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.

11.2.1 OS/HW Controlled P-states

11.2.1.1 Enhanced Intel SpeedStep® Technology

Enhanced Intel SpeedStep® Technology enables OS to control and select P-state.

11.2.1.2 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.

11.3 Power and Performance Technologies

11.3.1 Intel® Thread Director

Intel® Thread Director helps monitor and analyze performance data in real time to seamlessly place the right application thread on the right core and optimize performance per watt.

Built directly into the hardware, Intel® Thread Director uses machine learning to schedule tasks on the right core at the right time (as opposed to relying on static rules). This helps to ensure that Performance-cores and Efficient-cores work in concert, background tasks do not slow you down, and you can have more applications open simultaneously.

- Monitors the runtime instruction mix of each thread and the state of each core with nanosecond precision.
- Provides runtime feedback to the OS to make the optimal decision for any workload.
- Dynamically adapts its guidance according to the Assured Based Power (ABP) of the system, operating conditions, and power settings.

11.3.2 Intel® Smart Cache Technology

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

- The LLC is shared between all Compute tile cores (of any type). The maximal size of LLC is 3MB (12 ways, set associative) per P-core or E-core module (bundle of 4 E-Cores).
- The LLC is non-inclusive.
- The LLC may also be referred to as a 3rd level cache.

11.3.3 P-core and E-core Level 1 and Level 2 Caches

The 1st level cache is not shared between physical cores and each physical core has a separate set of caches.

The P-Core 1st level cache hierarchy is divided into:

- A Data Cache (DL0, DL1)
- An Instruction Cache (IL1)

On the data side, it is built as two-level cache, with L0 of 48KB and L1 of 192KB, both of which are 12-way set-associative.

On the instruction side, there is a single L1 cache of 64KB, which is 16-way set associative.

The E-Core 1st level cache hierarchy is divided into:

- A Data Cache (DL1)
- An Instruction Cache (IL1)

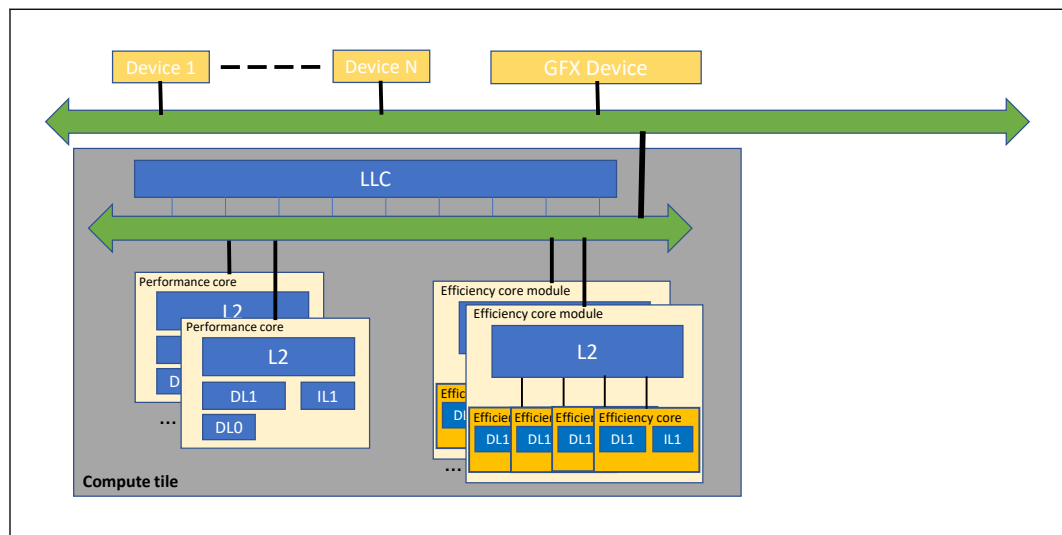
On the data side, it is built as one-level cache, with L1 of 32KB, 8-way set-associative.

On the instruction side, there is a single L1 cache of 64KB, which is 16-way set associative.

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

- The P-core 2nd level caches are not shared between physical cores and each physical core has a separate set of caches. Its size is 3MB and it is a 12-way associative non-inclusive cache.
- The E-core 2nd level cache is shared between E-Cores within a module of 4 E-Cores in the Compute tile. Its size is 4MB and it is a 16-way associative non-inclusive cache.

Figure 7. P-core and E-core Cache Hierarchy



NOTE

The above figure does not represent the exact number of cores.

Cache	P-core	E-core
L0 DL0	48KB 12-way set-associative per core	None
L1 DL1	192KB 12-way set-associative per core	32KB 8-way set-associative per core
<i>continued...</i>		

Cache	P-core	E-core
L1 IL1	64KB 16-way set-associative per core	64KB 16-way set-associative per core
L2	3MB 12-way set-associative per core	4MB 16-way set-associative within a module of 4 Compute tile E-cores
L3	Maximum of 3 MB per P-core / module of 4 E-cores shared across Compute tile	

11.3.4 Intel® Adaptive Boost Technology

Intel® Adaptive Boost Technology (Intel® ABT) opportunistically increases the multicore turbo frequency while operating within IccMAX and temperature spec limitations.

Intel® ABT opportunistically delivers in-spec performance gains that are incremental to existing Turbo technologies. In systems equipped with performance spec power delivery, Intel® ABT allows additional multi-core turbo frequency while still operating within specified current and temperature limits.

11.3.5 Intel® System Agent Enhanced SpeedStep® Technology

Intel® System Agent Enhanced SpeedStep® Technology

Intel® System Agent Enhanced SpeedStep® Technology also known as SAGV (System Agent Geyserville) 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 SpeedStep® Technology has three valid operating points. When running light workload and Intel System Agent Enhanced SpeedStep® Technology is enabled, the DDR data rate may change as follows:

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

In order to achieve the optimal levels of performance and power, the memory initialization and training process performed during first system boot or after CMOS clear or after a BIOS update will take a longer time than a typical boot. During this initialization and training process, end users may see a blank screen. More information on the memory initialization process can be found in the industry standard JEDEC Specifications found on www.JEDEC.org.

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.

11.3.6 User Mode Wait Instructions

The *UMONITOR* and *UMWAIT* are user mode (Ring 3) instructions similar to the supervisor mode (Ring 0) *MONITOR/MWAIT* instructions without the C-state management capability.

TPAUSE is an enhanced *PAUSE* instruction.

The mnemonics for the three new instructions are:

- **UMONITOR:** operates just like *MONITOR* but allowed in all rings.
- **UMWAIT:** allowed in all rings, and no specification of target C-state.

- **TPAUSE:** similar to *PAUSE* but with a software-specified delay. Commonly used in spin loops.

11.4 Deprecated Technology

The processor has deprecated the following technology and it is no longer supported:

- DDR Running Average Power Limit (DDR RAPL)

11.5 Power and Internal Signals

11.5.1 Signal Description

For more information, refer to Intel® 800 Series Chipset Family Platform Controller Hub (PCH) Datasheet, Volume 1 of 2 (#833778)

Signal Name	Type	Description
PLT_PWROK	I	Platform Power OK: When asserted, PLT_PWROK is an indication to the processor that all of its core power rails have been stable. The platform may drive asynchronously. When PLT_PWROK is de-asserted, the processor asserts PLTRST#. <i>Notes:</i> <ul style="list-style-type: none"> • PLT_PWROK must not glitch, even if RSMRST# is low • An external pull-down resistor is required.
RSMRST#	I	Primary Well Reset: This signal is used for resetting the primary power plane logic. This signal must be asserted for at least 10ms before de-asserting. <i>Note:</i> An external pull down resistor is required.
SLP_S0#	O	S0 Sleep Control: When the processor is in C10 state, this pin will assert to indicate VR controller can go into a light load mode. This signal can also be connected to EC for other power management related optimizations. <i>Note:</i> An external pull-up resistor is required.
GPP_SD00/ TIME_SYNC0 GPP_SD01/ TIME_SYNC1	I	Time Synchronization: Used for synchronization both input (latch time when pin asserted) and output (toggle pin when programmed time is hit).
PS_ON#	O	Used to indicate to PSU when to turn off its main rails
RESET_SYNC#	I/O	Bidirectional signal used to synchronize reset events between the processor and PCH. Sync reset exits with the PCH in following two steps: <ol style="list-style-type: none"> 1. The processor floats the RESET_SYNC# pin to allow it to be pulled high by the platform 2. The processor waits until it sees the RESET_SYNC# pin go high before proceeding to reset exit. The PCH will stop driving the pin low when the eSPI controller in the PCH is out of reset and ready to receive traffic over the eSPI link. It is required that the PCH eSPI controller is ready before the processor eSPI controller is released from reset.
PROC_C10_GATE#	O	When asserted, PROC_C10_GATE# is the indication to the system that the processor is entering C10 and can handle the voltages on the VCCPRIM_IO and VCCPRIM_VNNAON rails being lowered to 0 V. When de-asserted must ramp back up to their operational voltage levels. The power good indicators for these rails must still be asserted high when these rails are lowered to 0 V during PROC_C10_GATE# assertion and while these rails ramp back up to their operational levels after PROC_C10_GATE# de-assertion.

11.5.2 Power Sequencing Signals

Table 18. Power Sequencing Signals

Signal Name	Description	Dir.	Link Type	Availability
VIDSOUT	VIDSOUT, VIDSCK, VIDALERT#: These signals comprise a three-signal serial synchronous interface used to transfer power management information between the processor and the voltage regulator controllers.	I/O	SE	S Processor Lines
VIDSCK		O		
VIDALERT#		I		

11.5.3 I/O Signal Planes and States

For more information, refer to Intel® 800 Series Chipset Family Platform Controller Hub (PCH) Datasheet, Volume 1 of 2 (#833778)

Signal Name	Power Plane	During Reset	Immediately after Reset	S4/S5
PROC_C10_GATE#	Primary	Driven High	Driven High	Driven High
PLT_PWROK	RTC	Undriven	Undriven	Undriven
RSMRST#	RTC	Undriven	Undriven	Undriven
PS_ON#	Primary	Indeterministic, no deglitch PD.	Depending on the PMC HW default value	Depending on PMC HW default value
SLP_S0# ¹	Primary	Driven High	Driven High	Driven High
RESET_SYNC#	Primary	Hi-Z	Driven High	Driven High

- Notes:
1. Driven High during S0 and driven Low during S0i3 when all criteria for assertion are met.
 2. SLP_S4# is driven low in S4/S5.
 3. SLP_S5# is driven high in S4, driven low in S5.
 4. .PRIMPWRDNACK is always '0' while in M0 or M3, but can be driven to '0' or '1' while in M0ff state. PRIMPWRDNACK is the default mode of operation.
 5. The pad should only be pulled low momentarily when the corresponding buffer power supply is not stable.
 6. Based on wake event and Intel CSME state.
 7. Internal weak pull-down resistor is enabled during power sequencing.
 8. Pin state is a function of whether the platform is configured to have Intel CSME on or off in Sx.
 9. Output High-Z, not glitch free.
 10. Output High-Z, glitch free with ~1 kohm Pull-down during respective power sequencing.
 11. Output High-Z, glitch free with ~20 kohm Pull-down during respective power sequencing.
 12. Output High-Z, glitch free with ~20 kohm Pull-up during respective power sequencing.

12.0 Power Delivery

12.1 Power and Ground Signals

This section describes the power rails.

Table 19. Power Rail Description

Signal Name	Description
VCCPRIM_1P8_PROC	Fixed 1.8 V for the CPU
VCCPRIM_1P8_PROC_SOC	Fixed 1.8 V for SOC-S
VCCPRIM_1P8_PROC_DDR	Fixed 1.8 V for DDR
VCCPRIM_1P8_PROC_FLTRA	VCCPRIM_1P8 with filter requirements.
VCCPRIM_IO	Fixed 1.25 V for IO blocks.
VCCPRIM_IO_OUT_PCH	VCCPRIM_IO voltage source for PCH VCCPRIM_JTAGPROC
VCCPRIM_IO_OUT_SVID	VCCPRIM_IO voltage source for IMVP SVID buses pull up voltage
VCCPRIM_VNNAON	Fixed 0.77 V for digital core blocks.
VCCCORE	Dynamic SVID power rail for IA cores.
VCCGT	Dynamic SVID power rail for graphics.
VCCSA	Dynamic SVID power rail for system agent.
VDD2	Fixed 1.05/1.10 V power rail for memory host controller.
VSS	Ground

Table 20. Power Rail Sense Signals

Signal Name	Description
VCCCORE_SENSE	VCCCORE sense pin.
VCCGT_SENSE	VCCGT sense pin.
VCCSA_SENSE	VCCSA sense pin.
VCCPRIM_IO_SENSE	VCCPRIM_IO sense pin.
VCCPRIM_VNNAON_SENSE	VCCPRIM_VNNAON sense pin.
VCCCORE_VSS_SENSE	VCCCORE VSS sense pin.
VCCGT_VSS_SENSE	VCCGT VSS sense pin.
VCCSA_VSS_SENSE	VCCSA VSS sense pin.

continued...

Signal Name	Description
VCCPRIM_1P8_PROC_SENSE	VCCPRIM_1P8_PROC sense pin.
VCCPRIM_VNNAON_VSS_SENSE	VCCPRIM_VNNAON_VSS sense pin.
VDD2_SENSE	VDD2 sense pin.

12.2 Current Excursion Protection (CEP)

This power management is a Processor integrated detector which senses when the Processor load current exceeds a preset threshold by monitoring for a Processor power domain voltage droop at the Processor power domain IMVPVR sense point. The Processor compares the IMVPVR output voltage with a preset threshold voltage (VTRIP) and when the IMVPVR output voltage is equal to or less than VTRIP , the Processor internally throttles itself to reduce the Processor load current and the power.

IMVP9.2 VRs enhance the CEP detector by adding a cycle by cycle current limiting feature where the IMVPVR quickly enters cycle by cycle current limit (becomes a current source) with the VR output current limited to a preset value (ITRIP) as set in the ICC_limit register.

13.0 Electrical Specifications

13.1 Processor Power Rails

Power Rail	Description	S Processor Line Controls
VCC _{CORE}	Processor IA Cores Power Rail	SVID
VCC _{GT}	Graphics Power Rail	SVID
VCC _{SA}	Processor System Agent Power Rail	SVID
VCC _{PRIM_1P8_PROC}	PCIE IO PHY Power 1.8V Rail	Fixed
VCC _{PRIM_IO}	Support IO	Fixed
VCC _{PRIM_VNNAON}	Support internal rails, TCSS, Display, PCIE and other internal Blocks	Fixed
V _{DD2}	Integrated Memory Controller Power Rail	Fixed (Memory technology dependent)

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 VCC 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.

13.2 DC Specifications

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

- The DC specifications for the DDR5 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 operating temperature, clock frequency, and input voltages. Read all notes associated with each parameter.
- IccMAX is the maximum current processor can draw, typically seen running a virus application (stress applications specifically designed to push the processor to maximum Power).
- With Fast V-Mode enabled, the output decoupling would see this IccMAX current and would need to be able to take transient load step up to IccMAX whereas the power stage (FET/Inductor) would only see ITRIP_MAX current.
- IccMAX.App is less than IccMax and is the electrical current drawn by the processor (per power rail) while running a typical user realistic application(s) scenario at P0nmax and Maximum Operating Temperature.
The processor VR must be able to sustain this current for at least 10ms.
- AC tolerances for all rails include voltage transients and voltage regulator voltage ripple up to 1 MHz. Refer to additional guidance for each rail.

13.2.1 Processor Power Rails DC Specifications

Altering clock frequency, power delivery, or voltage may void any product warranties and reduce stability, security, performance, and life of the processor and other components.

13.2.1.1 VCCCORE DC Specifications

Table 21. Processor VCC_{CORE} Active and Idle Mode DC Voltage and Current Specifications (S Processor Line)

Segment	Power Delivery Configuration	Symbol	Parameter	Minimum	Typical	Maximum	Unit	Note ¹
All S-Processor Line	Performance	Operating Voltage	Active Voltage Range for Vcc CORE	0	-	1.77	V	1,2,3,7,12,15,17
8P+16E Core 125W	Performance	IccMAX (S Processor)	Max. Current for Processor Rail	-	-	347	A	4,5,6,7,11
8P+16E Core 125W	Performance	IccMAX.App (S Processor)	Max. Application Current for Processor Rail	-	-	245	A	4,5,6,7,11
8P+12E Core 125W	Performance	IccMAX (S Processor)	Max. Current for Processor Rail	-	-	347	A	4,5,6,7,11
8P+12E Core 125W	Performance	IccMAX.App (S Processor)	Max. Application Current for Processor Rail	-	-	245	A	4,5,6,7,11
6P+8E Core 125W	Performance	IccMAX (S Processor)	Max. Current for Processor Rail	-	-	242	A	4,5,6,7,11

continued...

Segment	Power Delivery Configuration	Symbol	Parameter	Minimum	Typical	Maximum	Unit	Note ¹
6P+8E Core 125W	Performance	IccMAX.App (S Processor)	Max. Application Current for Processor Rail	-	-	174	A	4,5,6,7,11
All	Performance	Icc TDC	Thermal Design Current for Vcc CORE	-	-	Refer to VR_TDC in the Platform Power Map	A	9
PS0, PS1, PS2, PS3	Performance	TOB VCC	DC Voltage Tolerance	-	-	±20	mV	3, 6, 8
PS0, PS1, PS2, PS3	Performance	TOB VCC +Ripple	DC + Ripple Voltage Tolerance	-	-	-35 /+50	mV	3, 6, 8,16
8P+16E Core 125W (Performance and Extreme configurations) 8P+12E Core 125W 6P+8E Core 125W	Performance	DC_LL (S Processor)	DC Loadline	0	-	1.2	mΩ	10,13,14,18
8P+16E Core 125W 8P+12E Core 125W 6P+8E Core 125W	Performance	AC_LL (S Processor)	AC Loadline	0	-	<ul style="list-style-type: none"> Below 200kHz: 1.2mOhms. 200kHz -1MHz: linear decrease with log(frequency) from 1.2mohms to 1.1mohms. Above 1MHz: 1.1mOhms 	mΩ	10,13,14,18
All	Performance	V_OVS_MAX	Max Overshoot Allowance from IccMAX	-	-	200	mV	
All	Performance	T_OVS_MAX	Max Overshoot Time from IccMAX	-	-	500	us	

continued...

Segment	Power Delivery Configuration	Symbol	Parameter	Minimum	Typical	Maximum	Unit	Note ¹
All	Performance	V_OVS_MAX_APP	Max Overshoot Allowance from IccMAX.App	-	-	100	mV	
All	Performance	T_OVS_MAX_APP	Max Overshoot Time from IccMAX.App	-	-	50	us	

Notes:

- 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.
- Maximum operating voltage given for motherboard design purposes. 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). Failure of product operation, including potential irreversible damage to the part, can occur from operating the part at voltages above the individual VID that is programmed.
- 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 20MHz bandwidth limit on the oscilloscope, 1.5pF maximum probe capacitance, and 1Ω minimum impedance. The maximum length of the ground wire on the probe should be less than 5mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
- Processor VccCORE VR to be designed to electrically support this current.
- Processor VccCORE 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. Refer to the IMVP9.2 Specification for more information.
- Refer to the Platform Power Map for Thermal Design Current (TDC) guidance.
- LL measured at sense points.
- Typ column represents IccMAX for commercial application. It is NOT a specification but rather a characterization of limited samples using a 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 a 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.
- An IMVP9.2 controller to support VccCORE needs to have offset voltage capability to potentially support voltages (VID+Offset) higher than 1.5V.
- Ripple can be higher if DC TOB is below 20mV, as long as Total TOB is within TOB_{VCC}+Ripple spec.
- 160 mV leakage may be observed when rail is powered off. There are no known functional or power implications due to this leakage.
- 1.2mOhm is the design limit for -S, achieving a lower design value will improve performance and power efficiency.

13.2.1.2 VccGT DC Specifications

Table 22. Processor Graphics (VccGT) Supply DC Voltage and Current Specifications (S Processor Line)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
Vcc GT	Active Voltage Range for Vcc GT	All	0	-	1.52	V	2, 3, 5, 7
Icc MAX_GT (S Processor)	Max. Current for Processor Graphics Rail	S-Processor Line 8P +16E Core 125W	-	-	40	A	5
Icc MAX_GT (S Processor)	Max. Current for Processor Graphics Rail	S-Processor Line 8P +12E Core 125W	-	-	40	A	5
Icc MAX_GT (S Processor)	Max. Current for Processor Graphics Rail	S-Processor Line 6P +8E Core 125W	-	-	40	A	5
Icc TDC_GT	Thermal Design Current (TDC) for Processor Graphics Rail	All	-	-		A	5
TOB VCC GT	DC Voltage Tolerance	S -Processor Line: PS0, PS1 ,PS2, PS3	-	-	±20	mV	3,4

continued...

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note ¹
TOBVCC GT+Ripple	DC + Ripple Voltage Tolerance	S Processor Lines: PS0, PS1, PS2, PS3	-	-	-35 /+50	mV	3, 4
DC_LL (S Processor)	DC Loadline	S -Processor Lines	-	-	4.3	mΩ	6,8
AC_LL (S Processor)	AC Loadline	S -Processor Lines	-	-	<ul style="list-style-type: none"> Below 700kHz: 4.3mOhms. 700kHz -800kHz: linear decrease with log(frequency) from 4.3mOhms to 4.1mOhms. Above 800kHz: 4.1mOhms. 	mΩ	6,8,9
V_OVS_MAX	Max Overshoot Allowance from IccMAX	All	-	-	70	mV	
T_OVS_MAX	Max Overshoot Time from IccMAX	All	-	-	10	μs	

- Notes: 1. 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. Maximum operating voltage given for motherboard design purposes. 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). Failure of product operation, including potential irreversible damage to the part, can occur from operating the part at voltages above the individual VID that is programmed.
3. The voltage specification requirements are measured across VccGT_SENSE and VssGT_SENSE as near as possible to the processor. The measurement needs to be performed with a 20MHz bandwidth limit on the oscilloscope, 1.5 pF maximum probe capacitance, and 1 MΩ minimum impedance. The maximum length of ground wire on the probe should be less than 5 mm. Ensure external noise from the system is not coupled into the oscilloscope probe.
4. PSx refers to the voltage regulator power state as set by the SVID protocol. Refer to the IMVP9.2 Specification for more information.
5. Refer to the Platform Power Map for Thermal Design Current (TDC) guidance.
6. Operating voltage range in steady state.
7. LL spec values should not be exceeded. If exceeded, power, performance and a reliability penalty are expected.
8. Load Line measured at the sense point.
9. Ripple can be higher if DC TOB is below 20mV, as long as Total TOB is within TOBVCCGT+Ripple spec.

13.2.1.3 VccSA DC Specification

Table 23. VccSA Supply DC Voltage and Current Specifications (S Processor Line)

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes
VCC _{SA}	Active Voltage Range for Vcc SA	All S-Processor Line	0	-	1.52	V	1,2,3,7
Icc _{MAX_SA}	Max. Current for Processor System Agent Rail	S-Processor Line 8P+16E Core 125W	0	—	39	A	1,2
Icc _{MAX_SA}	Max. Current for Processor System Agent Rail	S-Processor Line 8P+12E Core 125W	0	—	39	A	1,2
Icc _{MAX_SA}	Max. Current for Processor System Agent Rail	S-Processor Line 6P+8E Core 125W	0	—	39	A	1,2

continued...

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Notes
ICC _{TDC_SA}	Thermal Design Current (TDC) for Processor System Agent Rail	All	-	-	Refer to VccSA IPL2 in the Platform Power Map	A	5
TOB _{VCCSA}	DC Voltage Tolerance	S - Processor Line PS0, PS1, PS2, PS3	—	—	±20	mV	1,3,6
TOB _{VCCSA} +Ripple	DC + Ripple Voltage Tolerance	S Processor Lines: PS0, PS1, PS2, PS3	—	—	-35 /+50	mV	3, 6, 8,16
DC_LL	DC Loadline	S -Processor Line	—	—	5.9	mΩ	4,5
AC_LL	AC Loadline	S Processor Line	—	—	<ul style="list-style-type: none"> Below 400kHz: 5.9mOhms 400kHz -1500kHz: linear decrease with log(frequency) from 5.9mOhms to 4mOhms. Above 1500kHz: 4mOhms 	mΩ	4,5
V_OVS_MAX	Max Overshoot Allowance from IccMAX	All	-	-	70	mV	
T_OVS_MAX	Max Overshoot Time from IccMAX	All	-	-	10	μs	

Notes:

- 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.
- Maximum operating voltage given for motherboard design purposes. 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). Failure of product operation, including potential irreversible damage to the part, can occur from operating the part at voltages above the individual VID that is programmed.
- 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.
- LL measured at sense points. LL specification values should not be exceeded. If exceeded, power, performance, and reliability penalty are expected.
- The LL values are for reference. Must still need to meet the voltage tolerance specification.**
- Voltage Tolerance budget values Include ripples
- Vcc_{SA} is having few point of voltage define by CPU VID

13.2.1.4 VDD2 DC Specifications

Table 24. Memory Controller (VDD2) Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note1
VDD2 (DDR5)	Processor I/O Supply Voltage for DDR5	All	-	1.1	-	V	3,4,5
TOB _{VDD2}	Voltage Tolerance	All	-	-	± 5	%	3,4
IccMAX_VDD2	Maximum Current for VDD2 Rail	S-Processor Line	-	-	4.22	A	2

Notes: 1. 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.

13.2.1.5 VCCPRIM_VNNAON DC Specifications

Table 25. VCCPRIM_VNNAON Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note
VCCPRIM_VNNAON	VCCPRIM_VNNAON Power Rail Voltage	All	-	0.77	-	V	2,3,4
TOB _{VCCPRIM_VNNAON}	Voltage Tolerance	All	-50	-	+50	mV	2,3
IccMAX_VNNAON	Maximum Current	S-Processor Line	-	-	29	A	3

Notes: 1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins. Keep 1st harmonic ripple ≤ 1MHz.
 3. No requirement on the breakdown of AC versus DC noise.
 4. 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

13.2.1.6 VCCPRIM_IO DC Specifications

Table 26. VCCPRIM_IO Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Units	Notes 1,2,5
VCCPRIM_IO	Processor I/O Power Rail Voltage	All	-	1.25	-	V	2,3,4
TOB VCCPRIM_IO	Voltage Tolerance	All	-	-	± 5	%	2,3
IccMAX_VCCPRIM_IO	Maximum Current	S-Processor Line	-	-	7.16	A	3

Notes: 1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins. Keep 1st harmonic ripple ≤ 1MHz.
 3. No requirement on the breakdown of AC versus DC noise.
 4. 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

13.2.1.7 VCCPRIM_1P8 DC Specifications

Table 27. VCCPRIM_1P8_PROC Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note
VCCPRIM_1P8_PROC	VCCPRIM_1P8_PROC Power Rail Voltage	All	-	1.8	-	V	2,3,4
TOB VCCPRIM_1P8_PROC	Voltage Tolerance	All	-	-	±5	%	2,3
IccMAX_VCCPRIM_1P8_PROC	Maximum Current	S-Processor Line	-	-	0.44	A	3

Notes: 1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins.
 3. No requirement on the breakdown of AC versus DC noise.
 4. 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

Table 28. VCCPRIM_1P8_PROC_SOC Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note
VCCPRIM_1P8_PROC_SOC	VCCPRIM_1P8_PROC_SOC Power Rail Voltage	All	-	1.8	-	V	2,3,4
TOB VCCPRIM_1P8_PROC_SOC	Voltage Tolerance	All	-	-	±5	%	2,3
IccMAX_VCCPRIM_1P8_PROC_SOC	Maximum Current	S-Processor Line	-	-	1.21	A	3

Notes: 1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins.
 3. No requirement on the breakdown of AC versus DC noise.
 4. 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

Table 29. VCCPRIM_1P8_PROC_DDR Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note
VCCPRIM_1P8_PROC_DDR	VCCPRIM_1P8_PROC_DDR Power Rail Voltage	All	-	1.8	-	V	2,3,4
TOB VCCPRIM_1P8_PROC_DDR	Voltage Tolerance	All	-	-	±5	%	2,3
IccMAX_VCCPRIM_1P8_PROC_DDR	Maximum Current	S-Processor Line	-	-	0.15	A	3

Notes: 1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins.
 3. No requirement on the breakdown of AC versus DC noise.
 4. 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

Table 30. VCCPRIM_1P8_PROC_FLTRA Supply DC Voltage and Current Specifications

Symbol	Parameter	Segment	Minimum	Typical	Maximum	Unit	Note
VCCPRIM_1P8_PROC_FLTRA	VCCPRIM_1P8_PROC_FLTRA Power Rail Voltage	All	-	1.8	-	V	2,3,4
TOBVCCPRIM_1P8_PROC_FLTRA	Voltage Tolerance	All	-	-	±5	%	2,3
IccMAX_VCCPRIM_1P8_PROC_FLTRA	Maximum Current	S-Processor Line	-	-	168	mA	3

Notes:

1. 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. Includes AC and DC error, where the AC noise is bandwidth limited to under 1 MHz, measured on package pins.
3. No requirement on the breakdown of AC versus DC noise.
4. 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

13.2.2 Processor Interfaces DC Specifications

13.2.2.1 DDR5 DC Specifications

Symbol	Minimum	Units
V _{IL}	-	V
V _{IH}	0.85*V _{dd2}	V
R _{ON_UP(DQ)}	30	Ω
R _{ON_DN(DQ)}	30	
R _{ODT(DQ)}	30	Ω
V _{ODT(DC)}	0.4*V _{dd2}	V
R _{ON_UP(CK)}	30	Ω
R _{ON_DN(CK)}	30	Ω
R _{ON_UP(CMD)}	30	Ω
R _{ON_DN(CMD)}	30	Ω
R _{ON_UP(CTL)}	30	Ω
R _{ON_DN(CTL)}	30	Ω
R _{ON_UP} (SM_PG_CNTL1)		Ω
R _{ON_DN} (SM_PG_CNTL1)		Ω
I _{LI}		mA
SM_RCOMP[0]	99	Ω

continued...

Symbol	Minimum	Units
SM_RCOMP[1]	99	Ω
SM_RCOMP[2]	99	Ω
<p><i>Notes:</i></p> <ol style="list-style-type: none"> 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. VIL is defined as the maximum voltage level at a receiving agent that will be interpreted as a logical low value. VIH is defined as the minimum voltage level at a receiving agent that will be interpreted as a logical high value. VIH and VOH may experience excursions above VDD2. However, input signal drivers should comply with the signal quality specifications. 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. Refer to processor I/O Buffer Models for I/V characteristics. ODT values after COMP (assuming ±5% inaccuracy). BIOS MRC can reduce ODT strength towards The minimum and maximum values for these signals are programmable by BIOS to one of the two sets. 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. SM_DRAMPWROK must have a maximum of 15 ns rise or fall time over VDD2 * 0.30 ±100 mV and the edge must be monotonic. RON tolerance is preliminary and might be subject to change. Maximum-minimum range is correct but center point is subject to change during MRC boot training. Processor may be damaged if VIH exceeds the maximum voltage for extended periods. 		

13.2.2.2 PCIe* DC and Timing Specifications

The PCIe Controller(s) and Transmit/Receive Physical Layer PHYs are compliant with the **PCI Express* Base Specification Revision 5.0 Version 1.0, 22 May 2019**.

For PCIe* electrical specifications, refer to the PCI Express* Base Specification Revision 5.0 Version 1.0, 22 May 2019; which is available at <https://pcisig.com/>.

13.2.2.3 Digital Display Interface (DDI) DC Specifications

Table 31. 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 _{OHHs}	HS output high voltage			360	mV	1
<i>continued...</i>						

Parameter	Description	Minimum	Nom	Max	Units	Notes 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 32. DSI LP Transmitter DC Specifications

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
TV _{pin(ABSMAX)}	Maximum transient time above VPIN(MAX) or below VPIN(MIN)			20	ns	5

Notes: 1. Applicable when the supported data rate <= 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 20 ns 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 33. Display Audio and Utility Pins DC Specification

Symbol	Parameter	Minimum	Typical	Maximum	Units
V _{OL}	Output Low Voltage	—	—	VCCIO *0.1	V
V _{OH}	Output High Voltage	VCCIO * 0.9	—	—	V
Output Impedance	Output Impedance	—	50	—	Ω
V _{IL}	Input Low Voltage	—	—	VCCIO *0.25	V
V _{IH}	Input Low Voltage	VCCIO * 0.75	—	—	V

1. DC specification for Disp_Utills_1 and Disp_Utills_2 signals.
 2. DC specification for: PROC_AUDOUT, PROC_AUDIN, PROC_AUDCLK.

13.2.2.4 Single Ended DC Specification

Signal Name	Operation Voltage	Input				Output				Hysteresis
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
CATERR#	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
THERMTRIP#										
PROC_C10_GATE#										
PS_ON#										
GPP_SA14										
GPP_SA15										
GPP_SA16										
GPP_SA22/RSVD	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SA23/RSVD										
DIR_ESPI_CS0#										
DIR_ESPI_RCLK										
DIR_ESPI_RESET#										
SLP_S0#										
DMI_PERST#										
GPP_SB00/ DDPA_CTRLCLK/ SRCCCLKREQ2#	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SB01/ DDPA_CTRLDATA/ SRCCCLKREQ3#										
GPP_SB12/ DDP0_CTRLCLK/ TBT_LX0_TXD										
GPP_SB13/ DDP0_CTRLDATA/ TBT_LX0_RXD										
GPP_SB14/ SRCCCLKREQ1#										
GPP_SB15/ DDP1_CTRLCLK/ TBT_LX1_TXD										
GPP_SB16/ DDP1_CTRLDATA/ TBT_LX1_RXD										
PROCHOT#	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SB02/ DDP2_CTRLCLK										
GPP_SB03/ DDP2_CTRLDATA										
GPP_SB04/ DDP3_CTRLCLK										
GPP_SB05/ DDP3_CTRLDATA										
GPP_SB07/ SRCCCLKREQ0#										
GPP_SB08/ I_SRCCCLKREQ0#										
GPP_SB09/ I_SRCCCLKREQ1#	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SD00/TIME_SYNC0										
GPP_SD01/TIME_SYNC1										

continued...

Signal Name	Operation Voltage	Input				Output				Vhysteresis
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
GPP_SD10/BKLTEN										
GPP_SD11/BKLTCTL	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SD12/DDDSP_HPDA/DISP_MISC4										
GPP_SD13/DDDSP_HPDA/DISP_MISC3										
GPP_SD14/DDDSP_HPDA/DISP_MISC4										
GPP_SD15/DDDSP_HPDA/DISP_MISC1										
GPP_SD16/DDDSP_HPDA/DISP_MISC2										
GPP_SD17/PCI_E_LINK_DOWN										
GPP_SD18/BOOTHALT#										
GPP_SD02										
GPP_SD21/AUDCLK										
GPP_SD22/AUDIN										
GPP_SD23/AUDOUT	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
GPP_SD03										
GPP_SD04										
GPP_SD05										
GPP_SD06/SRCLKREQ2#										
GPP_SD07/SRCLKREQ3#										
GPP_SD08										
GPP_SD09/VDDEN										
PLT_PWROK	1.8 V	0.65 X VCC	0.35 X VCC	- 5 uA (min) 5 uA (max)	2 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.34 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	3.5 kohm (min), 5 kohm (nom), 6.52 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	0.1XVCC
RESET_SYNC#										
RTC_CLK_IN										

Signal Name	Operation Voltage	Input				Output				Vhysteresis
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
DIR_ESPI_I00	1.8 V	0.65 X VCC	0.35 X VCC	15.15 uA (max)	5.25 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.303 kohm (max) 3.5 kohm (min), 5 kohm (nom), 6.5 kohm (max)	3.5 kohm (min), 5 kohm (nom), 6.5 kohm (max) 14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3	0.1XVCC
DIR_ESPI_I01										
DIR_ESPI_I02										
DIR_ESPI_I03										

continued...

Signal Name	Operation Voltage	Input				Output				V _{hysteresis}
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
								14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	X VCC / 0.5 X VCC / 0.7 X VCC	
DIR_ESPI_CLK	1.8 V	0.65 X VCC	0.35 X VCC	15.15 uA (max)	5.25 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.303 kohm (max)	3.5 kohm (min), 5 kohm (nom), 6.5 kohm (max)	0.1XVCC
GPP_SB18/RSVD/BSSB_LS0_RX								3.5 kohm (min), 5 kohm (nom), 6.5 kohm (max)		
GPP_SB19/RSVD/BSSB_LS0_TX								14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC		

13.2.2.4.1 CMOS DC Specifications

Table 34. CMOS Signal Group DC Specifications

Signal Name	Operation Voltage	Input				Output				V _{hysteresis}
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
PECI	1.25 V	0.55 ~ 0.7 X VCC	0.3 ~ 0.5 X VCC	- 20 uA (min) 20 uA (max)	5 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.1XVCC
DBG_PMODE								14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC		

13.2.2.4.2 GTL and OD DC Specification

Table 35. GTL Signal Group DC Specifications

Signal Name	Operation Voltage	Input				Output				V _{hysteresis}
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
New GTL										
PROC_ITAG_TCK	1.25 V	0.5 X VCC	0.2 X VCC	- 20 uA (min) 20 uA (max)	5 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	80mV
PROC_ITAG_TRST								14 kohm (min), 20 kohm (nom), 26 kohm (max) Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC		
Legacy GTL										
BPM[0:3]	1.25 V	0.8 X VCC	0.477 X VCC	- 20 uA (min) 20 uA (max)	5 pF (max)	VCC - 0.20 (min) Test Load: 1 mA	0.20 (max) Test Load: -1 mA	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	60mV
PROC_ITAG_TDI										

continued...

Signal Name	Operation Voltage	Input				Output				V _{hysteresis}
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
PROC_JTAG_TDO								14 kohm (min), 20 kohm (nom), 26 kohm (max)	14 kohm (min), 20 kohm (nom), 26 kohm (max)	
PROC_JTAG_TMS								Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	
PRDY										
PREQ										

13.2.2.4.3 SVID DC Specifications

Table 36. SVID Signal Group DC Specifications

Signal Name	Operation Voltage	Input				Output				V _{hysteresis}
		Input High Voltage Threshold (V _{IH})	Input Low Voltage Threshold (V _{IL})	Input Leakage Current (I _{IL})	Input Pin Capacitance (C _{IN})	Output High Voltage Threshold (V _{OH})	Output Low Voltage Threshold (V _{OL})	Pull-up Resistance (R _{PU})	Pull-down Resistance (R _{PD})	
VIDALERT	1.25 V	0.55 ~ 0.7 X VCC 0.7 X VCC	0.45 X VCC	- 20 uA (min) 20 uA (max)	5 pF (max)	VCC - 0.45 (min) Test Load: 2 mA	0.45 (max) Test Load: -2 mA	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.7 kohm (min), 1 kohm (nom), 1.3 kohm (max)	0.1XVCC
VIDSCK								14 kohm (min), 20 kohm (nom), 26 kohm (max)	14 kohm (min), 20 kohm (nom), 26 kohm (max)	
VIDSOUT								Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	Condition: 0.3 X VCC / 0.5 X VCC / 0.7 X VCC	

14.0 Thermal Management

Table 37. Definitions/Acronyms

Acronyms	Description
Max Operating Temperature	<p>This is the maximum operating temperature allowed as reported by temperature sensors. Instantaneous temperature may exceed this value for short durations.</p> <p><i>Note:</i> Maximum observable temperature is configurable by system vendor and can be design specific.</p>

14.1 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 operating temperature specification at the maximum Processor Base Power.
- Conforms to system constraints, such as system acoustics, system skin-temperatures, and exhaust-temperature requirements.

CAUTION

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

14.1.1 Thermal Considerations

The Processor Base Power is the assured sustained power that should be used for the design of the processor thermal solution, design to a higher thermal capability will get more Turbo residency. Processor Base Power is the time-averaged power dissipation that the processor is validated to not exceed during manufacturing while executing an Intel-specified high complexity workload at Base Frequency and at the maximum operating temperature as specified in the Datasheet for the SKU segment and configuration.

NOTE

The System on Chip processor integrates multiple compute cores and I/O on a single package. Platform support for specific usage experiences may require additional concurrency power to be considered when designing the power delivery and thermal sustained system capability

The processor integrates multiple processing IA cores, graphics cores and for some SKUs a chipset 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 Processor Base Power 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 Processor Base Power 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.

14.1.1.1 Package Power Control

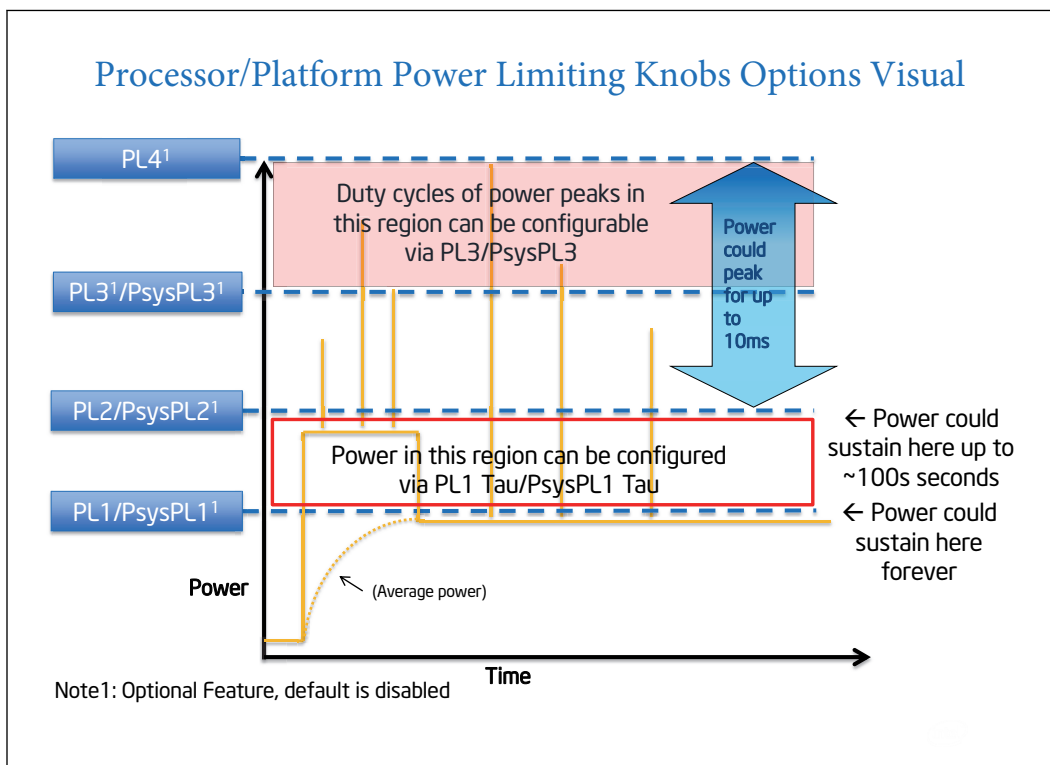
The package power control settings of PL1, PL2, PL3, PL4, and Tau allow the designer to configure Intel® Turbo Boost Technology 2.0 to match the platform power delivery and package thermal solution limitations.

- **Power Limit 1 (PL1):** A threshold for average power that will not exceed - recommend to set to equal Processor Base Power 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.
3. The Intel Dynamic Tuning (DTT) is recommended for performance improvement in mobile platforms. Dynamic Tuning is configured by system manufacturers dynamically optimizing the processor power based on the current platform thermal and power delivery conditions. Contact Intel Representatives for enabling details.

Figure 8. Package Power Control



14.1.1.2 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 IMVP (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 in [Package Power Control](#) on page 84.

- **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 Technology (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 current platform thermal and power delivery conditions. Contact Intel Representatives for enabling details.

14.1.1.3 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.

14.1.2 Assured Power (cTDP)

Assured Power (cTDP) form a design option where the processor's behavior and package Processor Base Power are dynamically adjusted to a desired system performance and power envelope. Assured Power (cTDP) 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. Assured Power (cTDP) are designed to be configured dynamically and do not require an operating system reboot.

NOTES

- PROCHOT events should be triggered after BIOS active. Triggering PROCHOT after BIOS is active should be ensured as it is essential for system stability.
 - Assured Power technologies are not battery life improvement technologies.
-

14.1.2.1 Assured Power (cTDP) Modes

NOTE

Assured Power (cTDP) availability may vary between the different SKUs.

With cTDP, the processor is now capable of altering the maximum sustained power with an alternate processor IA core base frequency. Assured Power (cTDP) allows operation in situations where extra cooling is available or situations where a cooler and quieter mode of operation is desired.

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

Table 38. Assured Power (cTDP)

Processor Power Characteristic	Description of Characteristic	Processor Design Considerations
Maximum Turbo Power	The maximum sustained (>1s) power dissipation of the processor as limited by current and/or temperature controls. Instantaneous power may exceed Maximum Turbo Power for short durations (<=10ms). Maximum Turbo Power is configurable by system vendor and can be system specific.	Intel performance advocacy for power delivery and transient thermal solution design (PL2)
Base Power	The time-averaged power dissipation that the processor is validated to not exceed during manufacturing while executing an Intel-specified high complexity workload at Base Frequency and at the operating temperature as specified in the Datasheet for the SKU segment and configuration.	Intel reference performance advocacy for sustained thermal solution design (PL1)
Minimum Assured Power	Min Assured Power is a performance advocacy determined by running a complex scenario defined by Intel. Every product SKU stack has guidance on Min Assured Power for thermal chassis design. Represents Intel specified min PL1 that needs to be taken for thermal design to get the advocated performance experience. Min Assured Power is the performance vs power cross-over point across the product SKU stack.	Intel minimum performance advocacy for sustained thermal solution design (PL1)
High Concurrency Power	High Concurrency Power is a functional characteristic determined by running a complex scenario defined by Intel. Every processor consumes a minimum power during a maximum connected case. This is a functional characteristic, not intended to indicate performance floor. Scenario takes into consideration IO ports, compute	Thermal solution capability required to support the high concurrency scenario as described

continued...

Processor Power Characteristic	Description of Characteristic	Processor Design Considerations
	IPs concurrency (CPU, GPU, IPU) along with memory BW. Temperature assumption of spec limit. Manufacturing screening is done to exclude parts that don't meet the target power for the high concurrency scenario. The processor may not honor PL1 values set lower than high concurrency power during the high concurrency scenario.	

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 Processor Base Power operation by adjusting processor PL1 dynamically. The cTDP mode does not change the maximum per-processor IA core turbo frequency.

14.1.3 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.

14.1.3.1 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.

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.

Maximum Operating Temperature is factory calibrated and is not user configurable. The default value is software visible in the TEMPERATURE_TARGET (1A2h) 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 = Processor Base Power. The system design should provide a thermal solution that can maintain normal operation when PL1 = Processor Base Power within the intended usage range.

Adaptive Thermal Monitor protection is always enabled.

TCC Activation Offset

TCC Activation Offset can be set as an offset from maximum operating temperature 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 (1A2h) 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 Maximum Operating Temperature 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 (1A2h) MSR, bits [29:24], and the time window (Tau) is written to the TEMPERATURE_TARGET (1A2h) 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 Operating Temperature 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 Maximum Operating Temperature thermal management. That is, whether the TCC activation offset is 0 or not, TCC Activation will occur at Maximum Operating Temperature.

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.

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.

Thermal Throttling

As the processor approaches Maximum Operating Temperature 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 Throttling 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.

14.1.3.2 Digital Thermal Sensor

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

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 (1B1h) MSR and IA32_THERM_STATUS (19Ch) 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, 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 (1A2h) MSR. The temperature returned by the DTS is an implied negative integer indicating the relative offset from Maximum Operating Temperature. The DTS does not report temperatures greater than Maximum Operating Temperature. 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 0h, 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. Refer to the *Intel 64 Architectures Software Developer's Manual* for specific register and programming details.

Digital Thermal Sensor Accuracy (T_{accuracy})

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

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 Maximum Operating Temperature.

14.1.3.3 PROCHOT# Signal

Intel recommends using PROCHOT# as an input signal to avoid Power, Thermal, and Performance implications.

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

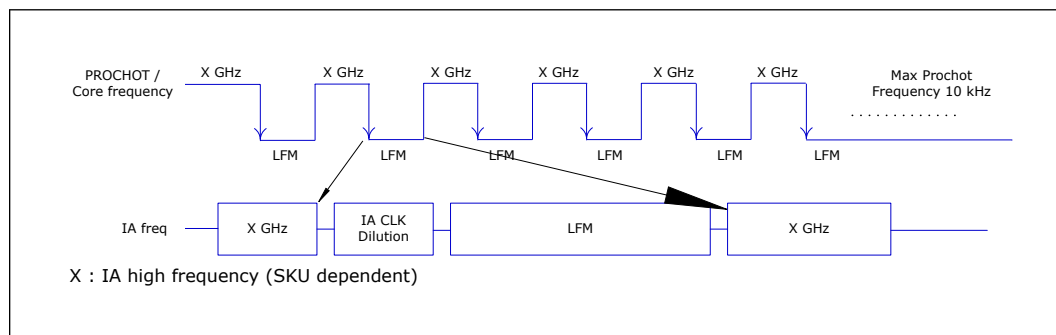
PROCHOT Input Only

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

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

- **Fast PROCHOT:** Respond to PROCHOT# within 1us of PROCHOT# pin assertion, reducing the processor power.
- **PROCHOT Demotion Algorithm:** Designed to improve system performance during multiple PROCHOT assertions.

Figure 9. PROCHOT Demotion Description



14.1.3.4 PROCHOT Output Only

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

14.1.3.5 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.

14.1.3.6 PROCHOT Demotion Algorithm

PROCHOT demotion algorithm is designed to improve system performance following multiple Platform PROCHOT consecutive assertions. During each PROCHOT assertion processor will eventually 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.

14.1.3.7 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, results in power reduction. Power reduction down to LFM and duration of the platform PROCHOT# assertion. 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 Adaptive Thermal Monitor protection is always enabled.

NOTE

During PROCHOT demotion, the core frequency may be reduced below LFM for several uSec.

14.1.3.8 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 operating 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).

14.1.3.9 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.

14.1.3.10 THERMTRIP# Signal

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

14.1.3.11 Critical Temperature Detection

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

14.1.3.12 Software Controlled Clock Modulation (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.

14.1.4 Intel® Memory Thermal Management

DRAM Thermal Aggregation

P-Unit 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.

DRAM Refresh Rate Control

The MRC will natively interface with MR4 or MPR readings to adjust DRAM refresh rate as needed to maintain data integrity. This capability is enabled by default and occurs automatically. Direct override of this capability is available for debug purposes, but this cannot be adjusted during runtime.

DRAM Bandwidth Throttling (Change to DDR Bandwidth Throttling)

Control for bandwidth throttling is available through the memory controller. Software may program a percentage bandwidth target at the current operating frequency and that used to throttle read and write commands based on the maximum memory MPR/MR4 reading.

14.2 Processor Base Power Thermal and Power Specifications

The following notes apply to [Processor Base Power Specifications \(S Processor Line\)](#), [Table 41](#) on page 97.

Table 39. General Notes

Note	Definition
1	The Processor Base Power and Assured Power (cTDP) values are the average power dissipation in operating temperature 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	Processor Base Power 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 Platform Power Control on page 85 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. 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 10ms.
<i>continued...</i>	

Note	Definition
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 Minimum Assured Power (cTDP Down) and/or Maximum Assured Power (cTDP Up) modes. Default power limits can be found in the PKG_PWR_SKU MSR (614h).
11	The processor tile do not reach maximum sustained power simultaneously since the sum of all active circuit's estimated power budget is controlled to be equal to or less than the specified PL1 limit. For additional information, refer to the appropriate Mobile TMDG for more information.
12	Minimum Assured Power(cTDP Down) is based on 128EU equivalent graphics configuration. Minimum Assured Power(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. PL2- Processor opportunistic higher Average Power with limited duration controlled by Tau_PL1 setting, the larger the Tau, the longer the PL2 duration.
15	Processor Base Power workload does not reflect various I/O connectivity cases such as Thunderbolt.
16	Hardware default of PL1 Tau=1s, By including the benefits available from power and thermal management features the recommended is to use PL1 Tau=28s.

Table 40. Processor Base Power Specifications (S Processor Line)

Segment and Package	Processor P/E cores, Graphics Configuration and Processor Base Power	Configuration			Processor P/ E core Frequency [GHz]	General Notes
		IA Core Frequency	Processor Base power	P-Core		
S-Processor Line LGA	8P+16E Core 125W ¹	IA Core Frequency	Processor Base power	E-Core	3.2 GHz	1,9,10,11,12, 15
				Low Frequency Mode - LFM	0.8 GHz	
		Graphics Core Frequency	Graphics Frequency	0.3 GHz		
			Low Frequency Mode - LFM	0.1 GHz		
S-Processor Line LGA	8P+12E Core 125W	IA Core Frequency	Processor Base power	P-Core	3.9 GHz	1,9,10,11,12, 15
				E-Core	3.3 GHz	
		Low Frequency Mode - LFM	0.8 GHz			
		Graphics Core Frequency	Graphics Frequency	0.3 GHz		
Low Frequency Mode - LFM	0.1 GHz					
S-Processor Line LGA	6P+8E Core 125W	IA Core Frequency	Processor Base power	P-Core	4.2 GHz	1,9,10,11,12, 15
				E-Core	3.6 GHz	
		Low Frequency Mode - LFM	0.8 GHz			
		Graphics Core Frequency	Graphics Frequency	0.3 GHz		
Low Frequency Mode - LFM	0.1 GHz					

Note: Refer to [General Notes](#)

14.3 Processor Line Thermal and Power Specifications

Table 41. Package Turbo Specifications (S Processor Lines)

No Specifications for Min/Max PL1/PL2 values.

Hardware default of PL1 Tau=1s, By including the benefits available from power and thermal management features the recommended is to use PL1 Tau=28s.

PL2- Processor opportunistic higher Average Power – Reactive, Limited Duration controlled by Tau_PL1 setting.

PL1 Tau - PL1 average power is controlled via PID algorithm with this Tau, The larger the Tau, the longer the PL2 duration.

System cooling solution and designs found to not being able to support the Performance TauPL1, adjust the TauPL1 to cooling capability.

Segment and Package	Processor IA Cores, Graphics, Configuration and Processor Base Power	Parameter	Minimum	Tau MSR Max Value	Recommended Value	Units	General Notes
S-Processor Line LGA	8P+16E Core 125W ⁶	Power Limit 1 Time (PL1 Tau)	0.1	448	56	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	N/A	125	W	
		Power Limit 2 (PL2)	N/A	N/A		W	
S-Processor Line LGA	8P+12E Core 125W	Power Limit 1 Time (PL1 Tau)	0.1	448	56	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	N/A	125	W	
		Power Limit 2 (PL2)	N/A	N/A		W	
S-Processor Line LGA	6P+8E Core 125W	Power Limit 1 Time (PL1 Tau)	0.1	448	56	S	3,4,5,6,7,8,14,16,17
		Power Limit 1 (PL1)	N/A	N/A	125	W	
		Power Limit 2 (PL2)	N/A	N/A		W	

Note: Refer to [General Notes](#)

Table 42. Operating Temperature Specifications (S Processor Line)

Segment	Package Turbo Parameter	Temperature Range		Processor Base Power Specification Temperature Range		Units	Notes
		Minimum	Maximum	Minimum	Maximum		
S-Processor Line LGA	Operating temperature	0	105	0	105	°C	1, 2

Notes: 1. The thermal solution needs to ensure that the processor temperature does not exceed the Processor Base Power Specification Temperature.
 2. The processor operating temperature is monitored by Digital Temperature Sensors (DTS). For DTS accuracy, refer to [Digital Thermal Sensor](#) on page 91.

Table 43. Low Power and TMTV Specifications (S Processor Line LGA)

Processor IA Cores, Graphics Configuration and Processor Base Power	PCG7	Maximum Power Package C8 (W) ^{1,4,5}	TMTV Processor Base Power (W) ^{6,7}	Min TCASE (°C)	Maximum TMTV TCASE (°C)
8P+16E Core 125W	2020A	N/A	125	0	61.8
8P+16E Core 65W	2022C	N/A	65	0	71.1

Notes: 1. The package C-state power is the worst case power in the system configured as follows:
 a. DMI and PCIe links are at L1
 2. Specification at DTS = 50 °C and minimum voltage loadline.
 3. Specification at DTS = 35 °C and minimum voltage loadline.
 4. These DTS values in Notes 2 - 3 are based on the TCC Activation MSR having a value of 100.
 5. These values are specified at VCC_MAX and VNOM for all other voltage rails for all processor frequencies. Systems should be designed to ensure the processor is not to be subjected to any static VCC and ICC combination wherein VCCP exceeds VCCP_MAX at specified ICCP. Refer to the loadline specifications.
 6. Processor Base Power should be used for processor thermal solution design targets. Processor Base Power is not the maximum power that the processor can dissipate. Processor Base Power is measured at DTS = -1. Processor Base Power is achieved with the Memory configured for DDR.
 7. Platform Compatibility Guide (PCG) (previously known as FMB) provides a design target for meeting all planned processor frequency requirements.

Table 44. TCONTROL Offset Configuration (S Processor Line LGA - Client)

Segment	TEMP_TARGET (TCONTROL) [°C]
8P+16E Core 125W	20

Notes: 1. Digital Thermal Sensor (DTS) based fan speed control is recommended to achieve optimal thermal performance.
 2. Intel recommends full cooling capability at approximately the DTS value of -1, to minimize TCC activation risk.
 3. For example, if TCONTROL = 20 °C, Fan acceleration operation will start at 80 °C (100 °C - 20 °C).

14.4 Error and Thermal Protection Signals

Table 45. Error and Thermal Protection Signals

Signal Name	Description	Dir.	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	O	SE	S Processor Series


continued...

Signal Name	Description	Dir.	Link Type	Availability
	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.			
PECI	Platform Environment Control Interface: A serial sideband interface to the processor. It is used primarily for thermal, power, and error management.	I/O	SE	S Processor Series
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	SE	S Processor Series
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 operating temperature exceeds approximately 125 °C. This is signaled to the system by the THERMTRIP# pin.	O	SE	S Processor Series

14.5 Thermal Metrology

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

Figure 10. Thermal Test Vehicle (TMTV) Case Temperature (T_{CASE}) Measurement Location



The following supplier can machine the groove and attach a thermocouple to the IHS. The following supplier is listed as a convenience to Intel's general customers and may be subject to change without notice.

THERM-X OF CALIFORNIA, 3200 Investment Blvd,

Hayward, Ca 94544. George Landis +1-510-441-7566 Ext. 368 george@therm-x.com.

The vendor part number is XTMS1565.

14.6 Fan Speed Control Scheme with DTS

With Digital Thermal Sensor (DTS) 1.1

To correctly use DTS 1.1, the designer must first select a worst case scenario $T_{AMBIENT}$, and ensure that the Fan Speed Control (FSC) can provide a Ψ_{CA} that is equivalent or greater than the Ψ_{CA} specification.

The DTS 1.1 implementation consists of two points:

- a Ψ_{CA} at T CONTROL
- a Ψ_{CA} at DTS = -1
- The Ψ_{CA} point at DTS = -1 defines the minimum Ψ_{CA} required at Processor Base Power considering the worst case system design $T_{AMBIENT}$ design point:

$$\Psi_{CA} = (T_{CASE-MAX} - T_{AMBIENT-TARGET} - 1) / \text{Processor Base Power}$$

For example, for a 125 W Processor Base Power part, the T_{CASE} maximum is 62.0 °C and at a worst case design point of 40 °C local ambient this will result in:

$$\Psi_{CA} = (62.0 - 40 - 1) / 125 = 0.168 \text{ } ^\circ\text{C/W}$$

Similarly for a system with a design target of 45 °C ambient, the Ψ_{CA} at DTS = -1 needed will be 0.128 °C/W.
- The second point defines the thermal solution performance (Ψ_{CA}) at T CONTROL . The following table lists the required Ψ_{CA} for the various Processor Base Power processors.

These two points define the operational limits for the processor for DTS 1.1 implementation. At T CONTROL the fan speed must be programmed such that the resulting Ψ_{CA} is better than or equivalent to the required Ψ_{CA} listed in the following table. Similarly, the fan speed should be set at DTS = -1 such that the thermal solution performance is better than or equivalent to the Ψ_{CA} requirements at T AMBIENT-MAX .

The fan speed controller must linearly ramp the fan speed from processor DTS = T CONTROL to processor DTS = -1.

Figure 11. Digital Thermal Sensor (DTS) 1.1 Definition Points

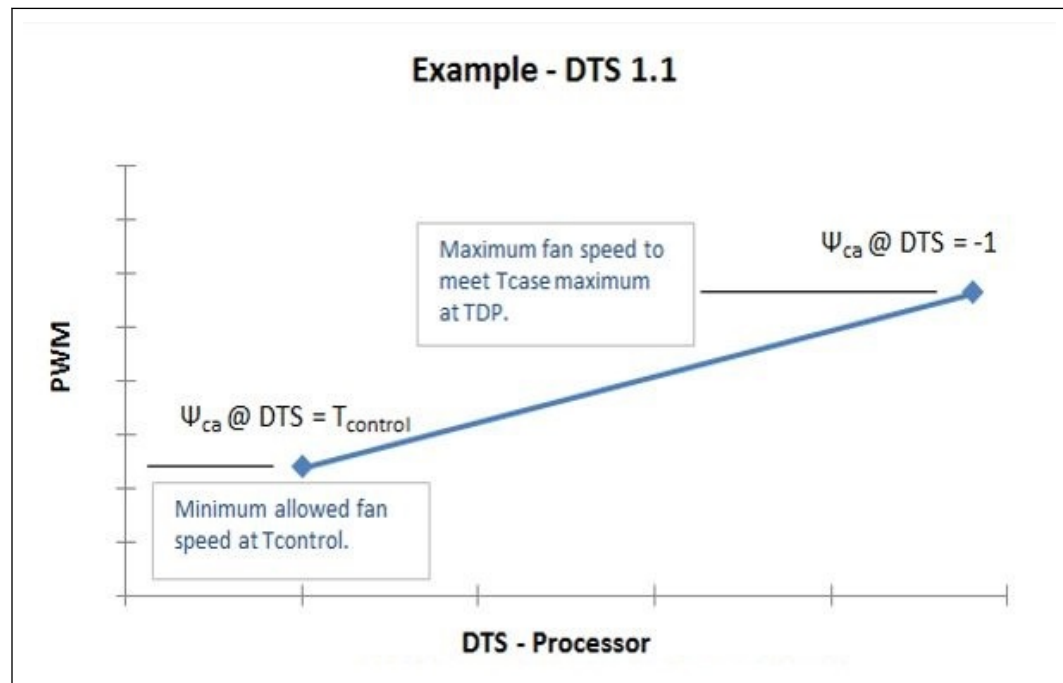


Table 46. Digital Thermal Sensor (DTS) 1.1 Thermal Solution Performance Above T_{CONTROL}

PCG	Die Configuration (Cores/GT)	Processor Base Power [W]	Ψ_{CA} at DTS = T _{CONTROL} ^{1, 2} At System T _{AMBIENT_MAX} = 30 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT_MAX} = 40 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT_MAX} = 45 °C	Ψ_{CA} at DTS = -1 At System T _{AMBIENT_MAX} = 50 °C
2020A	8P+16E	125	0.26	0.17	0.13	0.09
2022C	8P+16E	65	0.63	0.46	0.39	0.31
2022C	6P+8E	65	0.65	0.47	0.39	0.31
2022D	8P+16E	35	1.01	0.70	0.55	0.41
2022D	6P+8E	35	1.02	0.70	0.56	0.42

Notes: 1. Ψ_{CA} at "DTS = T_{CONTROL}" is applicable to systems that have an internal T_{RISE} (T_{ROOM} temperature to Processor cooling fan inlet) of less than 10 °C. In case the expected T_{RISE} is greater than 10 °C, a correction factor should be used as explained below. For each 1 °C T_{RISE} above 10 °C, the correction factor (CF) is defined as CF = 1.7 / (Processor Base Power).
 2. The table data match for GT0.

With Digital Thermal Sensor (DTS) 2.0

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

NOTE

TCC Activation Offset is 0 for the processors.

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

Figure 12. Digital Thermal Sensor (DTS) 2.0 Definition Points




Table 47. Thermal Margin Slope

PCG	Die Configuration (Cores/GT)	Processor Base Power [W]	TCC Activation [°C]	Temperature Control Offset	Thermal Margin Slope [°C/W]
2020A	8P+16E	125	105	20	0.39
2022C	8P+16E	65	105	25	0.59
2022C	6P+8E	65	105	20	0.69
2022D	8P+16E	35	105	37	0.65
2022D	6P+8E	35	105	31	0.78

Note: 1. The default BIOS settings for this SKU is 10C TCC offset.

14.7 Thermal Sensor

The processor incorporates an on-tile Digital Thermal Sensor (DTS) for thermal management.

14.7.1 Modes of Operation

The DTS has two usages when enabled:

1. One use is to provide the temperature of the CPU in units of 1 °C. There is a 9 bit field for the temperature, with a theoretical range from -256 °C to +256 °C. Practically the operational range for TS would be between -40 °C and 110 °C.
2. The second use is to allow programmed trip points to cause alerts to SW or in the extreme case shutdown. Temperature may be provided without having any SW alerts set.

There are two thermal alert capabilities. One is for the catastrophic event (thermal runaway) which results in an immediate system power down (S5 state). The other alert provides an indication to the platform that a particular temperature has been caused. This second alert needs to be routed to SMI or SCI based on SW programming.

14.7.2 Temperature Trip Point

The internal thermal sensor reports three trip points: Cool, Hot, and Catastrophic trip points in the order of increasing temperature.

Crossing the cool trip point when going from higher to lower temperature may generate an interrupt. Crossing the hot trip point going from lower to higher temp may generate an interrupt. Each trip point has control register bits to select what type of interrupt is generated.

Crossing the cool trip point while going from low to higher temperature or crossing the hot trip point while going from high to lower temperature will not cause an interrupt.

When triggered, the catastrophic trip point will transition the system to S5 unconditionally.

14.7.3 Thermal Sensor Accuracy (Taccuracy)

The processor thermal sensor accuracy is:

- ± 5 °C over the temperature range from 50 °C to 110 °C.
- ± 7 °C over the temperature range from 30 °C to 50 °C.
- ± 10 °C over the temperature range from -10 °C to 30 °C.
- No accuracy is specified for temperature range beyond 110 °C or below -10 °C.

14.7.4 Thermal Reporting to EC

To support a platform EC that is managing the system thermals, the processor provides the ability for the EC to read the processor temperature over SMBus and/or over eSPI. If enabled, Power Management will drive the temperature directly to the SMBus and eSPI units. The EC will issue an SMBus read or eSPI OOB Channel request and receives a single byte of data, indicating a temperature between 0°C and 127°C, where 255 (0xFF) indicates that the sensor isn't enabled yet. The EC must be connected to either SMLink1 or eSPI for thermal reporting support.

15.0 Clock Topology

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

- 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

15.1 Integrated Reference Clock PLL

Figure 13. System Clock Block Diagram



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 Overclocking capabilities.

15.2 Processor Clocking Signals

Table 48. Signal Description

Signal Name	Type	Description
CLKOUT_I_GEN4_N1	O	PCI Express* Clock Output: Serial Reference 100 MHz PCIe* specification compliant differential output clocks to PCIe* devices interfacing IOE. CLKOUT_I_GEN5_P/N [1] can be used for IOE PCIe* Gen4 support.
CLKOUT_I_GEN4_P1		
CLKOUT_I_GEN5_N0		
CLKOUT_I_GEN5_P0		
CLKOUT_S_GEN5_N1	O	PCI Express* Clock Output: Serial Reference 100 MHz PCIe* specification compliant differential output clocks to PCIe* devices interfacing SOC-S. CLKOUT_S_GEN5_P/N [3:0] can be used for SOC-S PCIe* Gen5 support
CLKOUT_S_GEN5_P1		
CLKOUT_S_GEN5_N2		
CLKOUT_S_GEN5_P2		
CLKOUT_S_GEN5_N3		
CLKOUT_S_GEN5_P3		
CLKOUT_S_GEN5_N0	O	DMI Clock Output: Serial Reference differential output clocks to PCH DMI interfacing SOC-S. CLKOUT_S_GEN5_P/N [0] is dedicated to be used as Intel® 800 Series Chipset DMI REFCLK only
CLKOUT_S_GEN5_P0		
GPP_SB08/ I_SRCLKREQ0#	IOD	Clock Request: Serial Reference Clock request signals for PCIe* 100 MHz differential clocks. The I_SRCLKREQ # signals can be configured to map to any of the IOE PCI Express* Root Ports while using any of the IOE CLKOUT differential pairs
GPP_SB09/ I_SRCLKREQ1#		
GPP_SD06/ SRCLKREQ2#	IOD	Clock Request: Serial Reference Clock request signals for PCIe* 100 MHz differential clocks. The SRCLKREQ # signals can be configured to map to any of the SOC-S PCI Express* Root Ports while using any of the SOC-S CLKOUT differential pairs. SRCLKREQ2# pins (muxed on GPP_SD06 and GPP_SB00) and SRCLKREQ3# pins (muxed on GPP_SD07 and GPP_SB01) are alternate signals; only one pin can be used at a time.
GPP_SB14/ SRCLKREQ1#		
GPP_SB00/ SRCLKREQ2#		
GPP_SD07/ SRCLKREQ3#		
GPP_SB01/ SRCLKREQ3#		
GPP_SB07/ SRCLKREQ0#	IOD	DMI Clock Request : Serial Reference Clock request signals for PCH DMI interfacing SOC-S. The SRCLKREQ0# signals is dedicated to be used as Intel® 800 Series Chipset DMI CLKREQ only
EXT_INJ_BCLK_N	I	100 MHz Differential bus clock input to the processor

continued...

Signal Name	Type	Description
EXT_INJ_BCLK_P		
EXT_INJ_PHYREF_N	I	100 MHz Differential bus clock input for PCIe and DMI OC
EXT_INJ_PHYREF_P		
XTAL_INJ_N	I	38.4 MHz crystal input
XTAL_INJ_P		
RTC_CLK_IN	I	32.768Khz RTC CLK from PCH to CPU

16.0 Memory

16.1 System Memory Interface

16.1.1 Processor SKU Support Matrix

Table 49. DDR Support Matrix Table

Technology	DDR5	
Processor	S	S
Configuration	1DPC ¹	2DPC ^{7,9}
Maximum Frequency [MT/s]	S-Processor ⁵ : UDIMM/SODIMM 5600 ¹¹ CUDIMM/CSODIMM 6400 ^{11, 12}	S-Processor UDIMM/CUDIMM : 1DIMM 1R/2R 5600 ¹¹ 2DIMM 1R 4800/2R 4400 ^{8, 11}
VDDQ [V] ⁵	5, 1.1	5, 1.1
VDD2 [V] ⁵	1.1	1.1
Maximum RPC ²	2	4
Die Density [Gb]	16, 24, 32 ¹¹	16, 24, 32 ¹²
Ballmap Mode	IL	IL
<p><i>Notes:</i> 1. 1DPC refers to when only 1DIMM slot per channel is routed. 2. RPC = Rank Per Channel 3. VDD2 is Processor and DRAM voltage, and VDDQ is DRAM voltage. 4. 5V is DIMM voltage, 1.1V is DRAM input voltage. 5. Speed is QDF dependant. 6. DDR5 ECC is supported only when all memory populated in system supports ECC 7. Far memory slot to be populated, in case, single DIMM is placed on 2DPC channel. 8. Maximum 2DPC frequency supported when same DIMM part number populated Within channel. Frequency is not guaranteed when mix DIMM's populated. 9. DDR5 CSODIMM and SODIMM or CUDIMM and UDIMM Mixed Configuration is not allowed. 10. DDR POR speed refers to Processor top SKU. Other SKUs may use lower memory speed, refer to ark.intel.com for top memory speed. 11. 64 GB DIMM's and 256 GB max System Capacity are dependent on ecosystem readiness and Intel validation. 12. DDR5 6400 requires to define a main clock out of 2 clock differential pairs for each 32-bit channel.</p>		

Table 50. DDR Technology Support Matrix

Technology	Form Factor	Ball Count	Processor
DDR5	SoDIMM/CSODIMM	262	S
DDR5	UDIMM/CUDIMM	288	S

16.1.2 Supported Memory Modules and Devices

Table 51. Supported DDR5 Non-ECC SoDIMM/CSoDIMM Module Configurations (S - Series Processor)

Raw Card Version	DIMM Capacity [GB]	DRAM Device Technology [Gb]	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/Col Address Bits	# of Banks Inside DRAM	Page Size [K]
A, F	16	16	2048M x 8	8	1	17/10	16	8
C, H	8	16	1024M x 16	4	1	17/10	8	8
B, G	32	16	2048M x 8	16	2	17/10	16	8
A, F	24	24	3072M x 8	8	1	17/10	32	8
C, H	12	24	1536M x 16	4	1	17/10	16	8
B, G	48	24	3072M x 8	16	2	17/10	32	8
A, F	32	32	4096M x 8	8	1	17/10	32	8
C, H	16	32	2048M x 16	4	1	17/10	16	8
B, G	64	32	4096M x 8	16	2	17/10	32	8

Table 52. Supported DDR5 ECC SoDIMM/CSoDIMM Module Configurations (S-Series Processor)

Raw Card Version	DIMM Capacity [GB]	DRAM Device Technology [Gb}	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/Col Address Bits	# of Banks Inside DRAM	Page Size [K]
D, J	16	16	2048M x 8	10	1	17/10	16	8
E, K	32	16	2048M x 8	20	2	17/10	16	8
D, J	24	24	3072M x 8	10	1	17/10	32	8
E, K	48	24	3072M x 8	20	2	17/10	32	8
D, J	32	32	4096M x 8	10	1	17/10	32	8
E, K	64	32	4096M x 8	20	2	17/10	32	8

Table 53. Supported DDR5 Non-ECC UDIMM/CUDIMM Module Configurations (S-Series Processor)

Raw Card Version	DIMM Capacity [GB]	DRAM Device Technology [Gb]	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/Col Address Bits	# of Banks Inside DRAM	Page Size [K]
A, F	16	16	2048M x 8	8	1	17/10	16	8
C, H	8	16	1024M x 16	4	1	17/10	8	8
B, G	32	16	2048M x 8	16	2	17/10	16	8
A, F	24	24	3072M x 8	8	1	17/10	32	8
C, H	12	24	1536M x 16	4	1	17/10	16	8
B, G	48	24	3072M x 8	16	2	17/10	32	8

continued...

Raw Card Version	DIMM Capacity [GB]	DRAM Device Technology [Gb]	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/Col Address Bits	# of Banks Inside DRAM	Page Size [K]
A , F	32	32	4096M x 8	8	1	17/10	32	8
C , H	16	32	2048M x 16	4	1	17/10	16	8
B , G	64	32	4096M x 8	16	2	17/10	32	8

Table 54. Supported DDR5 ECC UDIMM/CUDIMM Module Configurations (S-Series Processor)

Raw Card Version	DIMM Capacity [GB]	DRAM Device Technology [Gb]	DRAM Organization	# of DRAM Devices	# of Ranks	# of Row/Col Address Bits	# of Banks Inside DRAM	Page Size [K]
D , J	16	16	2048M x 8	10	1	17/10	16	8
E , K	32	16	2048M x 8	20	2	17/10	16	8
D , J	24	24	3072M x 8	10	1	17/10	32	8
E , K	48	24	3072M x 8	20	2	17/10	32	8
D , J	32	32	4096M x 8	10	1	17/10	32	8
E , K	64	32	4096M x 8	20	2	17/10	32	8

16.1.3 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 DDR5 command may be issued every 2 clocks
 - 1N indicates a new DDR5 command may be issued every clock

Table 55. DDR5 System Memory Timing Support

DRAM Device	Transfer Rate (MT/s)	tCL (tCK)	tRCD (ns)	tRP (ns)	CWL (tCK)	CMD Mode
DDR5	4800	40	16.00	16.00	38	2N
DDR5	5600	46	16.00	16.00	44	2N
DDR5	6000	48	16.00	16.00	46	2N
DDR5	6400	52	16.00	16.00	50	2N

16.1.3.1 SAGV Points

SAGV (System Agent Geyserville) is a way by which the processor 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.

Table 56. SA Speed Enhanced Speed Steps (SA-GV) and Gear Mode Frequencies

Processor	Technology	Rank Config	DDR Maximum Rate [MT/s]	SAGV-LowBW	SAGV-MedBW	SAGV-HighBW	SAGV- High Performance
S	DDR5		5600	3200	4800	5200	5600
	DDR5		6400	3200	4800	6000	6400

Notes: 1. Intel® Core™ Ultra 200S Series Processors supports dynamic gearing technology where the Memory Controller can run at 1:2 (Gear-2 mode) or 1:4 (Gear-4 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.

2. Frequency points may change depending on system validation.

3. SA-GV modes:

- 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** - Tuned for balance between power & performance.
- HighBW** - Characterized by high power, low latency, moderate BW also used as RFI mitigation point.
- MaxBW/Lowest latency** Lowest Latency point, peak BW and highest power.

DDR Frequency Shifting

DDR interfaces emit electromagnetic radiation which can couple to the antennas of various radios that are integrated in the system, and cause radio frequency interference (RFI). The DDR Radio Frequency Interference Mitigation (DDR RFIM) feature is primarily aimed at resolving narrowband RFI from DDR5 technologies for the Wi-Fi* high and ultra-high bands (~5-7 GHz) . By changing the DDR data rate, the harmonics of the clock can be shifted out of a radio band of interest, thus mitigating RFI to that radio. This feature is working with SAGV on, the 3rd SAGV point is used as RFI mitigation point

16.1.4 Memory Controller (MC)

The integrated memory controller is responsible for transferring data between the processor and the DRAM as well as the DRAM maintenance. There are two instances of MC, one per memory slice. Each controller is capable of supporting up to two channels of DDR5.

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

In a symmetric memory population, each controller provides access to half of the total physical memory address space.

16.1.5 System Memory Controller Organization Mode

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 accesses are directed to a single Memory Controller. Single-Channel mode is used when either the MC0 or MC1 are populated in any order, but not both.

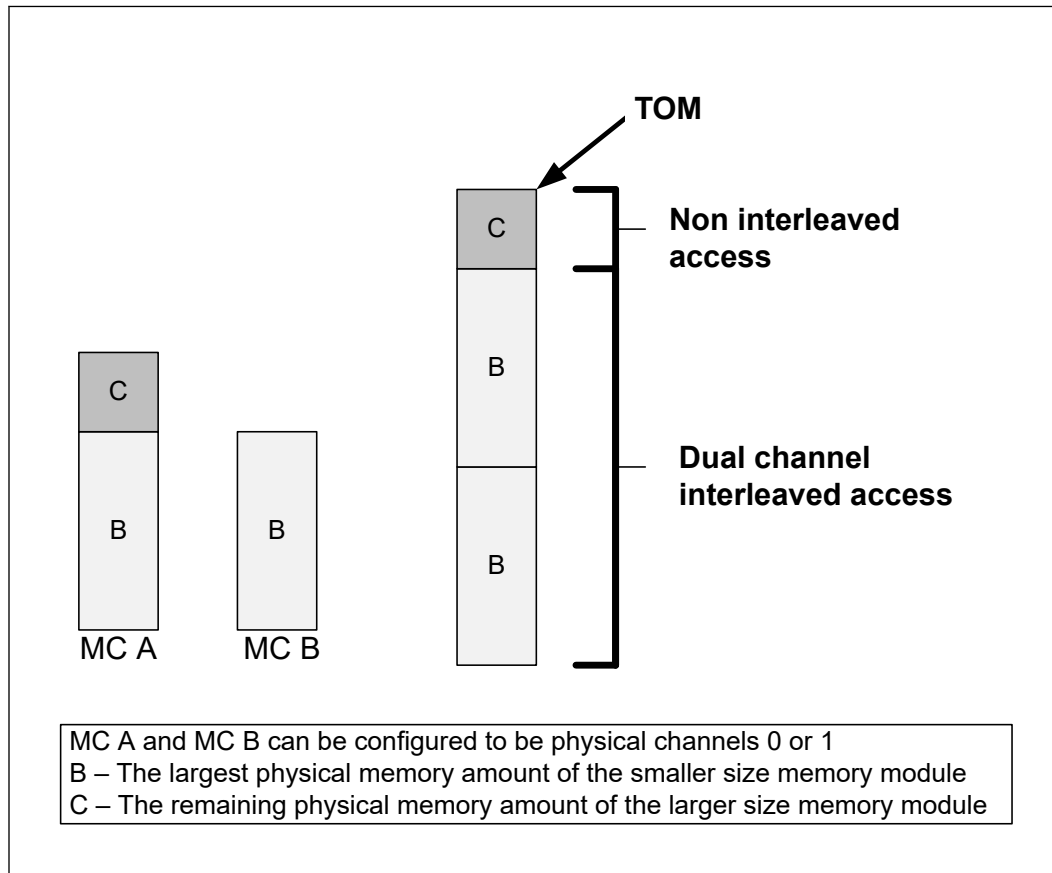
Dual-Channel Mode – Intel® Flex Memory Technology Mode (DDR5 Only)

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 MC 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

MC A and MC B can be mapped for physical MC0 and MC1 respectively or vice versa; however, Channel A size should be greater or equal to Channel B size.

Figure 14. Intel® DDR5 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. 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. Use Dual-Channel Symmetric mode when both MC0 and MC1 are populated in any order, with the total amount of memory in each channel being the same.

When both MCs 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

- The DDR5 DRAM device technology and width may vary from one channel to another.
- Different memory size between channels are relevant to DDR5 only.

16.1.6 System Memory Frequency

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

16.1.7 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.

16.1.8 Data Scrambling

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

16.1.9 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 S-Processor ECC supported SKUs.

16.1.10 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.
- DDR5 x32 sub-channels can be swizzle within their x64 MC.
- DDR5: Byte swapping is allowed within each x32 Channel.
- ECC bits swap is allowed within ECC byte/nibble: DDR5 ECC[3..0].

NOTE

All DRAM devices sharing ZQ resistor must be connected to the same MC channel.

16.1.11 DDR I/O Interleaving

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

16.1.12 DRAM Clock Generation

Each support rank has a differential clock pair for DDR5.

16.1.13 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. In all cases, it has small step sizes and is trained by MRC.

16.1.14 Data Swizzling

All Processor Series have no die-to-package DDR swizzling.

16.1.15 Post Package Repair (PPR)

PPR is supported according to JEDEC Spec.

BIOS can identify a single Row failure per Bank in DRAM and perform Post Package Repair (PPR) to exchange failing Row with spare Row.

PPR can be supported only with DRAM that supports PPR according to Jedec spec.

16.2 Integrated Memory Controller (IMC) Power Management

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

16.2.1 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/CS) are not driven.

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

16.2.2 DRAM Power Management and Initialization

The processor implements extensive support for power management on the memory interface.

The DRAM Powerdown is one of the power-saving means. When DRAM is in Powerdown state, 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 different power-down modes supported are:

- **No power-down:**
- **Pre-charged Power-down (PPD):** This mode is entered if all banks in DDR are pre-charged when entering Powerdown state. Power-saving in this mode is intermediate. Power consumption is defined by IDD2P. Exiting this mode is defined by tXP. In this mode when waking-up, all page-buffers are empty.

The Powerdown state 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. 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.

16.2.2.1 Conditional Self-Refresh

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

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

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

16.2.2.2 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 or pre-charge power-down. 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.

16.2.2.3 DRAM I/O Power Management

Unused signals should be disabled to save power and reduce electromagnetic interference. Clocks 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).

16.2.3 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 in the appropriate state.

In C8 or deeper power state, the processor internally gates VCCSA for all non-critical state to reduce idle power.

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

16.2.4 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.

16.3 Signal Description

Table 57. DDR5 Memory Interface

Signal Name	Description	Dir.	Buffer Type	Link Type	Availability
DDR0_DQ[3:0][7:0] DDR0_DQ[4][7:4] DDR1_DQ[3:0][7:0] DDR1_DQ[4][7:4] DDR2_DQ[3:0][7:0] DDR2_DQ[4][3:0] DDR3_DQ[3:0][7:0] DDR3_DQ[4][3:0]	<p>Data Buses: Data signals interface to the SDRAM data buses.</p> <p>Example: DDR0_DQ[2][5] refers to DDR channel 0, Byte 2, Bit 5.</p>	I/O	DDR5	SE	S-Series Processor
DDR0_DQS_P[4:0] DDR0_DQS_N[4:0] DDR1_DQS_P[4:0] DDR1_DQS_N[4:0] DDR2_DQS_P[4:0] DDR2_DQS_N[4:0] DDR3_DQS_P[4:0] DDR3_DQS_N[4:0]	<p>Data Strobes: Differential data strobe pairs. The data is captured at the crossing point of DQS during reading and write transactions.</p> <p>Example: DDR0_DQSP0 refers to DQSP of DDR channel 0, Byte 0.</p>	I/O	DDR5	Diff	S-Series Processor
DDR0_CLK[3:0]_P DDR0_CLK[3:0]_N DDR1_CLK[3:0]_P DDR1_CLK[3:0]_N DDR2_CLK[3:0]_P DDR2_CLK[3:0]_N DDR3_CLK[3:0]_P DDR3_CLK[3:0]_N	<p>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.</p>	O	DDR5	Diff	S-Series Processor
DDR0_CS[3:0] DDR1_CS[3:0] DDR2_CS[3:0] DDR3_CS[3:0]	<p>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.</p> <p>The Chip select signal is Active Low.</p>	O	DDR5	SE	S-Series Processor
DDR0_CA[12:0] DDR1_CA[12:0] DDR2_CA[12:0] DDR3_CA[12:0]	<p>Command Address: These signals are used to provide the multiplexed command and address to the SDRAM.</p>	O	DDR5	SE	S-Series Processor
DRAM_RESET#	Memory Reset	O	CMOS	SE	S-Series Processor

17.0 USB Type-C* Sub System

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

The USB Type-C sub-system supports USB 3.2, USB4*, DPoC (DisplayPort over Type-C) protocols. The USB Type-C sub-system can also be configured as native DisplayPort 1.4a/2.1 or HDMI 2.1 interfaces, for more information refer to [Display](#) on page 142.

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

- USB 2.0, USB 3.2 (2x10 Gb/s), USB 3.2/DP implemented at the connector.
- In addition, 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 Gb/s) only solutions are implemented, Thunderbolt 3 compatibility as defined by USB4/USB-PD specs and 15 W of bus power are still recommended.

17.1 General Capabilities

- xHCI (USB 3.2 host controller) and xDCI (USB 3.2 Gen 1x1 device controller) implemented in the processor.
- Intel® AMT/vPro over Thunderbolt™ docking.
- Support power saving when USB Type-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 USB4.
 - Wake capable on each host port from S0i3, Sx.
- Device
 - Aggregate BW through xDCI controller at max 5 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 58. USB Type-C* Port Configuration

	Port	S-Processor IOE-P-Series	S PCH SKUs
Group A	TCP 0	USB4 ⁴ , DisplayPort ¹ , USB 3.2 ³ , HDMI ²	W880, Z890, Q870, B860, H810
	TCP 1		W880, Z890, Q870
<p><i>Note:</i> TCP 1 is available as fixed HDMI/DP only with PCH SKUs B860 and H810.</p> <p><i>Notes:</i> 1. Supported on Type-C or Native connector (Fixed DP up to HBR3 link rate) 2. Supported only on Native connector. 3. USB 3.2 supported link rates: a. USB 3.2 Gen 1x1 (5 Gbps) b. USB 3.2 Gen 2x1 (10 Gbps) c. USB 3.2 Gen 2x2 (20 Gbps) 4. USB4 operating link rates (including both rounded and non-rounded modes for Thunderbolt™ 3 compatibility): a. USB4 Gen 2x2 (20 Gbps) b. USB4 Gen 3x2 (40 Gbps) c. 10.3125 Gbps, 20.625 Gbps per lane - Compatible to Thunderbolt™ 3 non-rounded modes. 5. USB 2.0 interface supported over Type-C connector. 6. Port group is defined as two ports sharing USB4 router, each router supports up to two display interfaces. 7. Display interface can be connected directly to a DP/HDMI/Type-C port or through USB4 router on a Type-C connector. 8. 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.</p>			

Table 59. USB Type-C* Lanes Configuration

Lane1	Lane2	Comments
USB4 / TBT3	USB4 / TBT3	Both lanes operate at same speed, one of (20.6 Gbps/10.3 Gbps/20 Gbps/10 Gbps)
USB4 / TBT3	No connect	20.6g/10.3g/20g/10g
No connect	USB4 / TBT3	
USB 3.2	USB 3.2	Multi-Lane USB 3.2 (Host Only), 2x10G = 20G
USB 3.2	No connect	Any combination of: USB 3.2 Gen 1x1 (5Gb/s) USB 3.2 Gen 2x1 (10Gb/s)
No connect	USB 3.2	
USB 3.2	DPx2	Any of HBR3/HBR2/HBR1/HRBR for DP1.4a, DP2.1 (2x10/20 Gbps) , and USB 3.2 (10 Gbps)
DPx2	USB 3.2	
DPx4	Both lanes at same DP rate - no support for 2x DPx2 USB Type-C connector	Any of HBR3/HBR2/HBR1/HRBR for DP1.4a, DP2.1 (4x10/20 Gbps)

Table 60. USB Type-C* Non-Supported Lane Configuration

Lane1	Lane2	Comments
-	PCIe* Gen3/2/1	No PCIe* native support
PCIe* Gen3/2/1	-	
-	USB4 / TBT3	No support for USB4 / TBT3 with any other protocol
USB4 / TBT3	-	

17.2 USB4* Router

USB4 is a Standard architecture (formerly known as CIO), but with the addition of USB 3.2 (20G) 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* & USB 3.2) and video

(DisplayPort*) on a single cable connection that can daisy-chain up to five 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, USB 3.2 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.

17.2.1 USB4 Host Router Implementation Capabilities

The integrated USB Type-C sub-system implements the following interfaces via USB4:

- Up to 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
- Up to two PCI Express* Root Port interfaces each one capable of:
 - PCI Express* 3.0 x4 compliant @ 8.0 GT/s
- Up to two xHCI Port interfaces each one capable of:
 - USB 3.2 Gen 2x1 (10 Gbps)
 - USB 3.2 Gen 2x2 (20 Gbps)
- USB4 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
- USB4 Time Management Unit (TMU):
- Up to two Interfaces to USB Type-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 per lane.
 - 16 counters per port

17.3 xHCI/xDCI Controllers

The processor supports xHCI/xDCI controllers. The native USB 3.2 path proceeds from the memory directly to PHY.

17.3.1 USB 3 Controllers

17.3.1.1 Extensible Host Controller Interface (xHCI)

Extensible Host Controller Interface (xHCI) is an interface specification that defines Host Controller for a universal Serial Bus (USB 3.2), which is capable of interfacing with USB 1.x, 2.0, and 3.x compatible devices.

In case that a device (example, USB 3.2 Flash Drive) was connected to the computer, the computer will work as Host and the xHCI will be activated inside the processor.

The xHCI controller support link rate of up to USB 3.2 Gen 2x2 (20G).

17.3.1.2 Extensible Device Controller Interface (xDCI)

Extensible Device Controller Interface (xDCI) is an interface specification that defines Device Controller for a universal Serial Bus (USB 3.2), 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).

NOTE

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

17.4 Display Interface

Refer to [Display](#) on page 142.

17.5 USB Type-C Signals

Signal Name	Description	Dir.	Link Type	Availability
TCP[1:0]_TX[1:0]_P TCP[1:0]_TX[1:0]_N	TX Data Lane.	O	Diff	S-Series Processor
TCP[1:0]_TXRX[1:0]_P TCP[1:0]_TXRX[1:0]_N	RX Data Lane, also serves as the secondary TX data lane.	I/O	Diff	S-Series Processor
TCP[1:0]_AUX_P TCP[1:0]_AUX_N	Common Lane AUX-PAD.	I/O	Diff	S-Series Processor

Note: TCP1 is available as fixed HDMI/DP **only** with PCH SKUs B860 and H810.

17.6 LSx

LSx interface supports Four ports. Each port of the LSx controller has two bi-directional signals configured either as Tx (Output) or Rx (Input). Operating voltage of the LSx interface is 1.8 V. LSx controller is responsible for link initialization/management of HSIO in the Thunderbolt subsystem.

17.6.1 LSx Signal Description

Signal Name	Type	Description
GPP_SB13/DDP0_CTRLDATA/ TBT_LSx0_RXD	I	LSx 0 Receive Data
GPP_SB12/DDP0_CTRLCLK/ TBT_LSx0_TXD	O	LSx 0 Transmit Data
GPP_SB16/DDP1_CTRLDATA/ TBT_LSx1_RXD	I	LSx 1 Receive Data
GPP_SB15/DDP1_CTRLCLK/ TBT_LSx1_TXD	O	LSx 1 Transmit Data

17.6.2 Integrated Pull-Ups and Pull-Downs

None.

17.6.3 I/O Signal Planes and States

Signal Name	Power Plane	During Reset ¹	Immediately after Reset ¹	S4/S5
TBT_LSx[0:1]_RXD	Primary	Undriven	Undriven	Undriven
TBT_LSx[0:1]_TXD	Primary	Undriven	Undriven	Undriven

Note: 1. Reset reference for primary well pins is RSMRST#.

17.7 AUX BIAS Control

On processor which support integrated USB Type-C* subsystem, the AUX BIAS control is required on the USB Type-C implementation (without retimer) for orientation connections. The functionality is muxed with certain GPIO pins. Refers to the GPIO implementation document for more information on the muxing and supported GPIO pin on the specific platform. In order to use the GPIO pin correctly for AUX BIAS control, the correct native functionality need to be configured and the correct Virtual Wire Index bit position need to be programmed in the BIOS policy.

Figure 15. GPIO - Virtual Wire Index Bit Mapping

GPIO Pin Group	Virtual Wire Index	Bit Position*
USB-C_GPP_[C07:C00]	10h	[7h:0h]
USB-C_GPP_[C15:C08]	11h	[7h:0h]
USB-C_GPP_[C23:C16]	12h	[7h:0h]
USB-C_GPP_[H07:H00]	13h	[7h:0h]
USB-C_GPP_[H15:H08]	14h	[7h:0h]
USB-C_GPP_[H19:H16]	15h	[3h:0h]

NOTE

1. The bit position corresponds to each corresponding GPIO pin in the group.
For example: the bit position for USB-C_GPP_C0 is bit 0h in Virtual Wire Index 10h.

18.0 Intel® Volume Management Device (Intel® VMD) 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.

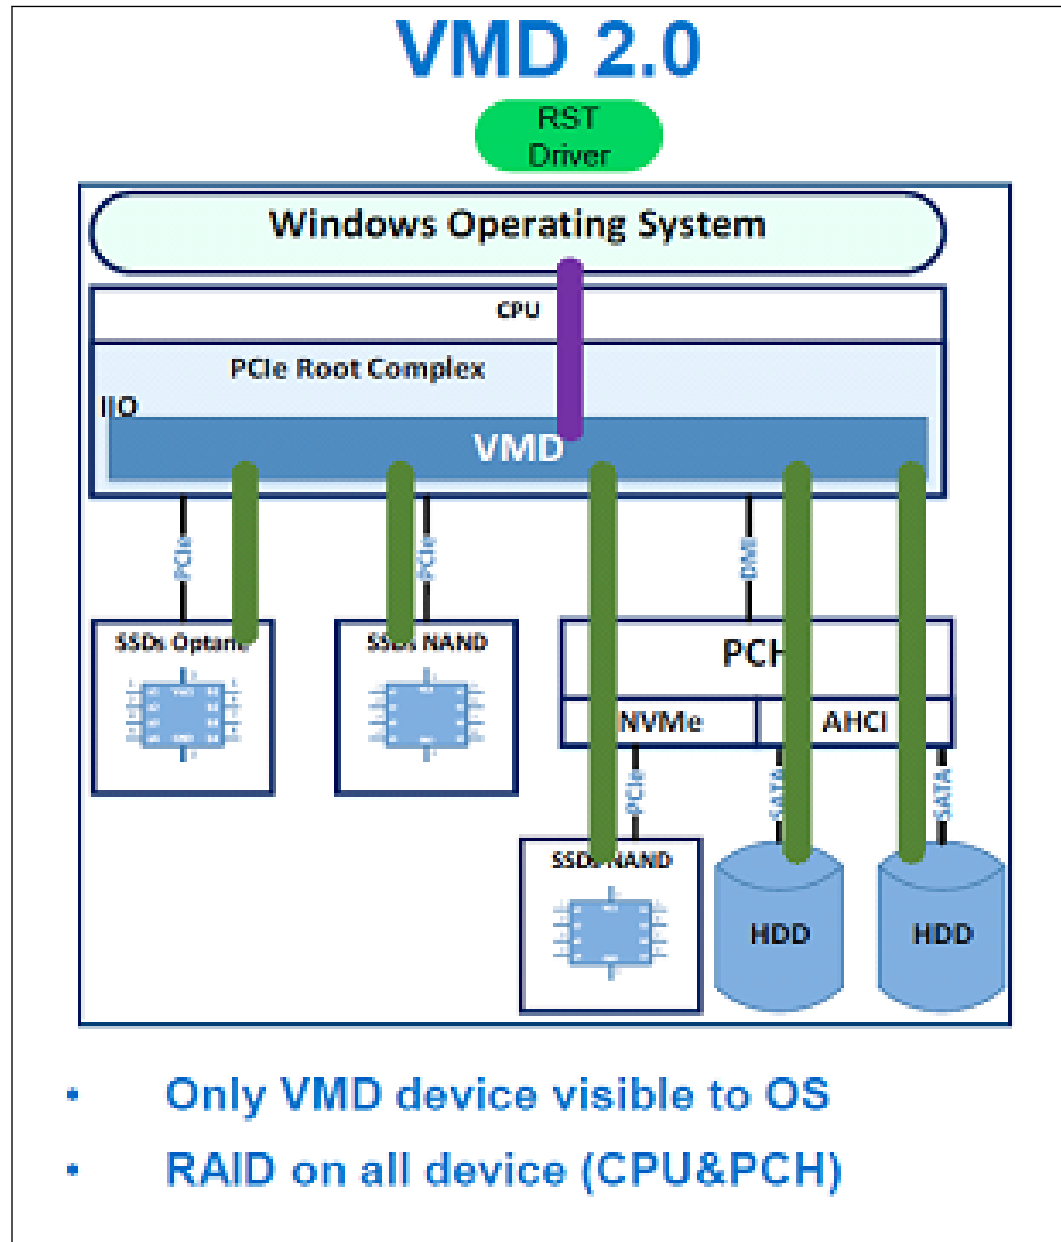
Intel® Volume Management Device (Intel® VMD) technology provides a means to provide volume management across separate PCI Express HBAs and SSDs without requiring operating system support or communication between drivers. For example, the OS will see a single RAID volume instead of multiple storage volumes,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

Technology Description

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.

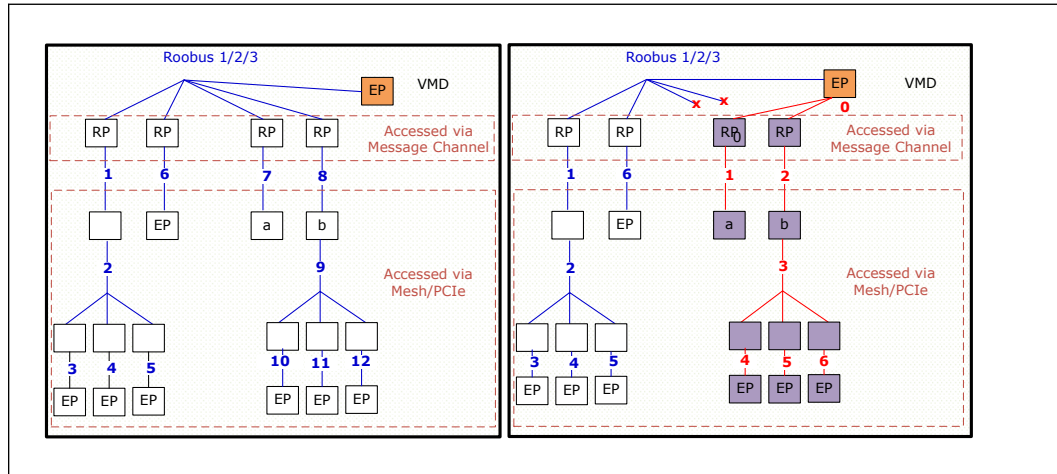
Intel® Volume Management technology requires support in BIOS and driver, memory and configuration space management.

Figure 16. Technology Description



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. 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.



Features Supported

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

19.0 PCI Express* (PCIe*)

Table 61. Acronym

Acronyms	Description
PCIe*	PCI Express* (Peripheral Component Interconnect Express*)

Table 62. Reference Table

Specification	Location
PCI Express* Base Specification Revision 5.0 Version 1.0, 22 May 2019	https://pcisig.com/
PCI Express* M.2 Specification Revision 4.0, Version 1.1, April 14, 2022	https://pcisig.com/
PCI Express* Card Electromechanical Specification, Revision 5.0, Version 1.0, June 9, 2021	https://pcisig.com/

19.1 Functional Description

Table 63. Features Supported

PCIe Controller Feature	PCIe Controllers			
	1	2	3	4
L1 Sub-States (L1.0, L1.1, L1.2)	Yes	Yes	Yes	Yes
L0s Link State (RX/TX)	Yes	Yes	Yes	Yes
S3/S4/S5 Sleep States (Sx)	Yes	Yes	Yes	Yes
Common Clock Mode	Yes	Yes	Yes	Yes
Separate Reference Clock with Independent SSC (SRIS)	No	No	Yes	No
Separate Reference Clock with No SSC (SRNS)	No	No	Yes	No
Precision Time Management (PTM)	Yes	Yes	Yes	Yes
Advanced Error Reporting (AER)	Yes	Yes	Yes	Yes
End-to-End Lane Reversal	Yes	Yes	Yes	No
Latency Tolerance Reporting (LTR)	Yes	Yes	Yes	Yes
PCIe TX Half Swing	No	No	No	No
PCIe TX Full Swing	Yes	Yes	Yes	Yes
Run Time D3 (RTD3)	Yes	Yes	Yes	Yes
Access Control Services (ACS)	Yes	Yes	Yes	Yes
Alternative Routing-ID Interpretation (ARI)	Yes	Yes	Yes	Yes
Port 80h Decode	Yes	Yes	Yes	Yes

continued...

PCIe Controller Feature	PCIe Controllers			
	1	2	3	4
Lane Polarity Inversion	Yes	Yes	Yes	Yes
PCIe Controller Root Port Hot-Plug	Yes	Yes	Yes	Yes
Downstream Port Containment (DPC)	No	No	No	No
Enhanced Downstream Port Containment (eDPC)	No	No	No	No
Virtual Channel (VC)	0	0	0/1	0
NVMe Cycle Router	No	No	No	No
Volume Management Device (Intel® VMD)	Yes	Yes	Yes	Yes
RAID[0] and RAID[1] Mode Support ^{1,2}	Yes	Yes	Yes	Yes
RAID[5] Mode Support ^{1,2}	Yes	Yes	Yes	Yes
RAID[10] Mode Support ^{1,2,3}	No	No	No	No
PCIe Controller (PC) Root Port (RP) Peer-2-Peer (P2P) Mem Write Transactions	RPs between PC2 and PC3 = No RPs between and within PCH PC1/2/3/4/5/6 = No RPs between PC1/2/3/4 and PCH PC1/2/3/4/5/6 = No RPs within PC1 = Yes RPs between PC2/3 and PC1/4 = Yes RPs between PC1 and PC4 = Yes			
PCIe Controller (PC) Root Port (RP) Peer-2-Peer (P2P) Mem Read Transactions	No			
PCIe Controller (PC) Root Port (RP) Peer-2-Peer (P2P) MCTP VDM Transactions	RPs within PC1 = Yes RPs between PC1/2/3/4 = Yes RPs between and within PCH PC1/2/3/4/5/6 = Yes RPs between PC1/2/3/4 and PCH PC1/2/3/4/5/6 = Yes			
PCIe Root Port Initiated Dynamic Width Change	No	No	No	No
PCIe Root Port Initiated Dynamic Speed Change	Yes	Yes	Yes	Yes
End Point Device Initiated Dynamic Width Change	Yes	Yes	Yes	Yes
End Point Device Initiated Dynamic Speed Change	Yes	Yes	Yes	Yes
Flattening Portal Bridge (FPB)	No	No	No	No

NOTES

- 1. No restrictions on PCIe Controller. PCIe RAID is expected to work across all Root Ports within a PCIe Controller and between Root Ports from different PCIe Controllers (Processor and PCH).
- 2. No RAID support between PCIe and SATA storage devices.
- 3. The Intel® Rapid Storage Technology (RST) does not restrict RAID modes so if any unsupported RAID mode is enabled it is up to the motherboard designer to validate any non-supported RAID mode.

19.1.1 PCI Express* Power Management

S4/S5 Sleep State Support

Software initiates the transition to S4/S5 by performing an IO write to the Power Management Controller. After the IO write completion has been returned the Power Management Controller will signal each root port to send a PME_Turn_Off message on the downstream link. The device attached to the link will eventually respond with a PME_TO_Ack followed by sending a PM_Enter_L23 DLLP request to enter L23. The Express ports and Power Management Controller take no action upon receiving a PME_TO_Ack. When all the Express port links are in state L23, the Power Management Controller will proceed with the entry into S4/S5.

Latency Tolerance Reporting (LTR)

The PCIe Controller Root Ports support the extended Latency Tolerance Reporting (LTR) capability. LTR provides a means for device endpoints to dynamically report their service latency requirements for memory reads and write access's to the Root Ports through the Latency Tolerance Reporting messages. Endpoint devices should transmit a new LTR message to the Root Ports initially during boot and each time its latency tolerance changes. This latency information allows the Power Management Controller (PMC) to make effective and accurate decisions to transition the platform to deeper power management states without the cost of making the wrong decision, since deeper power management states are usually associated with longer exit latency.

19.1.2 Port 80h Decode

The PCIe* root ports will explicitly decode and claim I/O cycles within the 80h – 8Fh range when MPC.P8XDE is set. The claiming of these cycles are not subjected to standard PCI I/O Base/Limit and I/O Space Enable fields. This allows a POST-card to be connected to the Root Port either directly as a PCI Express* device or through a PCI Express* to PCI bridge as a PCI card.

Any I/O reads or writes will be forwarded to the link as it is. The device will need to be able to return the previously written value, on I/O read to these ranges. BIOS must ensure that at any one time, no more than one Root Port is enabled to claim Port 80h cycles.

19.1.3 Separate Reference Clock with Independent SSC (SRIS)

The current PCI - SIG "PCI Express* External Cabling Specification" (www.pcisig.com) defines the reference clock as part of the signals delivered through the cable. Inclusion of the reference clock in the cable requires an expensive shielding solution to meet EMI requirements.

The need for an inexpensive PCIe* cabling solution for PCIe* SSDs requires a cabling form factor that supports non Common Clock Mode with spread spectrum enabled, such that the reference clock does not need to be part of the signals delivered through the cable. This clock mode requires the components on both sides of a link to tolerate a much higher ppm tolerance of ~5600 ppm compared to the PCIe* Base Specification defined as 600 ppm.

Soft straps are needed as a method to configure the port statically to operate in this mode. This mode is only enabled if the SSD connector is present on the motherboard, where the SSD connector does not include the reference clock. No change is being made to PCIe* add-in card form factors and solutions.

ASPM L0s is not supported in this form factor. The L1 exit latency advertised to software would be increased to 10 us. The root port does not support Lower SKP Ordered Set generation and reception feature defined in SRIS ECN.

19.1.4 Advanced Error Reporting

The PCI Express* Controller Root Ports each provide basic error handling, as well as Advanced Error Reporting (AER) as described in the latest PCI Express* Base Specification.

19.1.5 Single - Root I/O Virtualization (SR - IOV)

Alternative Routing ID Interpretation (ARI) and Access Control Services (ACS) are supported as part of the complementary technologies to enable SR - IOV capability.

Alternative Routing - ID Interpretation (ARI)

Alternative Routing - ID Interpretation (ARI) is a mechanism that can be used to extend the number of functions supported by a multi - function ARI device connected to the Root Port, beyond the conventional eight functions.

Access Control Services (ACS)

ACS is defined to control access between different Endpoints and between different Functions of a multi - function device. ACS defines a set of control points to determine whether a TLP should be routed normally, blocked, or redirected.

19.1.6 PCI Express* Receiver Lane Polarity Inversion

The PCI Express* Base Specification requires polarity inversion to be supported independently by all receivers across a Link where 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 in the PCIe* Controllers 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".

19.1.7 Precision Time Measurement (PTM)

Hardware protocol for precise coordination of events and timing information across multiple upstream and downstream devices using Transaction Layer Protocol (TLP) Message Requests. Minimizes timing translation errors resulting in the increased coordination of events across multiple components with very fine precision.

All of the PCIe* Controllers and their assigned Root Ports support PTM where each Root Port can have PTM enabled or disabled individually from one another.

19.2 Signal Description

Signal Name	Type	Description	Availability
PCIE_[24:1]_TX_N PCIE_[24:1]_TX_P	O	PCI Express* Differential Transmit Pairs These are the PCI Express* based outbound high-speed differential signals.	S-Series Processor
PCIE_[24:1]_RX_N PCIE_[24:1]_RX_P	I	PCI Express* Differential Receive Pairs These are the PCI Express* based inbound high-speed differential signals.	
GPP_SD17/PCIE_LINK_DOWN	O	PCI Express* Link Down Debug Signal PCIe link failure debug signal. PCIe Root Port(s) will assert this signal when a link down event occurs and is detected. For example when a link fails to train during an L1 sub-state exit event.	

19.3 I/O Signal Planes and States

Table 64. Power Plane and States for PCI Express* Signals

Signal Name	Type	Power Plane	During Reset ²	Immediately After Reset ²	S4/S5	Deep Sx
PCIE_[24:1]_TX_N PCIE_[24:1]_TX_P	O	Primary	Internal Pull-down	Internal Pull-down	Internal Pull-down	OFF
PCIE_[24:1]_RX_N PCIE_[24:1]_RX_P	I	Primary	Internal Pull-down	Internal Pull-down	Internal Pull-down	OFF

Notes: 1. PCIe_[24:1]_RX_N/PCIE_[24:1]_RX_P pins transition from un-driven to Internal Pull-down during Reset.
2. Reset reference for primary well pins is RSMRST#.

19.4 PCI Express* Root Port Support Feature Details

Table 65. PCI Express* Root Port Feature Details

Product	Max Transfer Rate	Max Devices (Root Ports)	Max Lanes	PCIe Gen Type	Encoding	Transfer Rate (MT/s)	Theoretical Max Bandwidth (GB/s)				
							x1	x2	x4	x8	x16
Processor	32 GT/s (Gen5)	5	24 ²	1	8b/10b	2500	0.25	0.50	1.00	2.00	4.00
				2	8b/10b	5000	0.50	1.00	2.00	4.00	8.00
				3	128b/130b	8000	1.00	2.00	3.94	7.88	15.75
				4	128b/130b	16000	1.97	3.94	7.88	15.75	31.51
				5	128b/130b	32000	3.94	7.88	15.75	31.51	63.02

Notes: 1. Theoretical Maximum Bandwidth (GB/s) = ((Transfer Rate * Encoding * # PCIe Lanes) / 8) / 1000
 • Gen5 with 16 PCIe Lanes Example = ((32000 * 128/130 * 16) / 8) / 1000 = 63.02 GB/s
 2. Maximum of 20 PCIe Lanes support up to Gen5 transfer rate (16 SOC-S + 4 IOE)

Figure 17. Supported PCI Express* Link Configurations

S-Processor	SOC (System On Chip) Tile																							IOE (IO Expander) Tile							
	FIA-2																							FIA-3				FIA-4			
	Flex I/O Lanes	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31						
PCIe Controllers	1 and 2																							3				4			
PCIe Max Rate	Gen5																							Gen4				Gen5			
PCIe Lanes	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24							
PCIe Configurations (Bifurcation) ³	> 1	1px16 / 1px16(LR)															1px4/1px4(LR)/Empty				Empty										
	2	1px16 / 1px16(LR)															Empty				1px4										
	3	1px8							1px8								1px4/1px4(LR)/Empty				Empty										
	4	1px8							1px8								Empty				1px4										
	5	1px8(LR)							1px8(LR)								1px4/1px4(LR)/Empty				Empty										
	6	1px8(LR)							1px8(LR)								Empty				1px4										
	7	1px8							1px4				1px4 / 1px4(LR)				1px4/1px4(LR)/Empty				1px4 / Empty										
	8	1px8(LR)							1px4(LR) / 1px4				1px4(LR)				1px4/1px4(LR)/Empty				1px4 / Empty										
	9	1px8 / 1px8(LR)							Empty								1px4 / 1px4(LR)				1px4 / Empty										
	10	1px8							1px4				Empty				1px4/1px4(LR)/Empty				1px4 / Empty										
	11	1px8(LR)							1px4(LR)				Empty				1px4/1px4(LR)/Empty				1px4 / Empty										
	12	1px16 / 1px16(LR)															1px4/1px4(LR)				1px4										
	13	1px8							1px8								1px4/1px4(LR)				1px4										
	14	1px8(LR)							1px8(LR)								1px4/1px4(LR)				1px4										
Logical Link Lanes	> 1	0/15	1/14	2/13	3/12	4/11	5/10	6/9	7/8	8/7	9/6	10/5	11/4	12/3	13/2	14/1	15/0	0/3	1/2	2/1	3/0										
	2	0/15	1/14	2/13	3/12	4/11	5/10	6/9	7/8	8/7	9/6	10/5	11/4	12/3	13/2	14/1	15/0	0/3	1/2	2/1	3/0	0	1	2	3						
	3	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0/3	1/2	2/1	3/0										
	4	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0/3	1/2	2/1	3/0	0	1	2	3						
	5	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	0/3	1/2	2/1	3/0										
	6	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	0/3	1/2	2/1	3/0	0	1	2	3						
	7	0	1	2	3	4	5	6	7	0	1	2	3	0/3	1/2	2/1	3/0	0/3	1/2	2/1	3/0	0	1	2	3						
	8	7	6	5	4	3	2	1	0	3/0	2/1	1/2	0/3	3	2	1	0	0/3	1/2	2/1	3/0	0	1	2	3						
	9	0/7	1/6	2/5	3/4	4/3	5/2	6/1	7/0					0/3	1/2	2/1	3/0	0/3	1/2	2/1	3/0	0	1	2	3						
	10	0	1	2	3	4	5	6	7	0	1	2	3					0/3	1/2	2/1	3/0	0	1	2	3						
	11	7	6	5	4	3	2	1	0	3	2	1	0					0/3	1/2	2/1	3/0	0	1	2	3						
	12	0/15	1/14	2/13	3/12	4/11	5/10	6/9	7/8	8/7	9/6	10/5	11/4	12/3	13/2	14/1	15/0	0/3	1/2	2/1	3/0	0	1	2	3						
	13	0	1	2	3	4	5	6	7	0	1	2	3	4	5	6	7	0/3	1/2	2/1	3/0	0	1	2	3						
	14	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	0/3	1/2	2/1	3/0	0	1	2	3						
Assigned Root Ports	> 1	RP13															RP10				RP12										
	2	RP13																			RP12										
	3	RP13							RP14								RP10				RP12										
	4	RP13							RP14												RP12										
	5	RP14							RP13								RP10														
	6	RP14							RP13												RP12										
	7	RP13							RP14				RP15				RP10				RP12										
	8	RP14							RP15				RP13				RP10				RP12										
	9	RP13/RP14											RP15/RP13				RP10				RP12										
	10	RP13							RP14								RP10				RP12										
	11	RP14							RP15								RP10				RP12										
	12	RP13															RP10				RP12										
	13	RP13							RP14								RP10				RP12										
	14	RP14							RP13								RP10				RP12										
Bus - Dev - Func (BDF) Assignments ²	RP13: Bus = 0h - Dev = 6h - Func = 0h RP14: Bus = 0h - Dev = 6h - Func = 3h RP15: Bus = 0h - Dev = 6h - Func = 4h																							RP	Bus	Dev	Func	RP	Bus	Dev	Func
																								10	0h	6h	1h	12	0h	1h	0h

NOTES

1. Covers all the Processor PCIe Controller hardware supported PCIe Bifurcation Configurations. Refer to the processor SKU breakdowns for SKU specific supported PCIe Bifurcation Configurations
 2. Device (BDF) groupings have multiple functions, the lowest active Root Port within the Device (BDF) grouping will always be assigned Function 0 while any remaining active Root Port within the Device (BDF) grouping will be assigned their mapped Function # as shown.
 3. Reduced Root Port width configurations, within Bi-Furcation configurations, are supported (example: x2 PCIe End Point Device populated in a PCIe Controller set as 1px4/1px16 will result in a 1px2 PCIe Root Port configuration or x1 PCIe End Point Device populated in a PCIe Controller set as 1px4/1px16 will result in a 1px1 PCIe Root Port configuration).
 4. FIA = Flex-IO Adapter
 5. The PCIe* Link Configuration support will vary depending on the SKU. Refer to the SKU details covered in the [Introduction](#) on page 12.
 6. LR = Lane Reversal
 7. PCIe Configuration (#p) x (#) = (Number of PCIe Root Ports) x (Number of Data Lane Pairs per PCIe Root Port)
 8. RP# refers to a specific PCI Express* Root Port #; for example RP3 = PCI Express* Root Port 3
 9. A PCIe* Lane is composed of a single pair of Transmit (TX) and Receive (RX) differential pairs. A connection between two PCIe* devices is known as a PCIe* Link, and is built up from a collection of one or more PCIe* Lanes which make up the width of the link (such as bundling 2 PCIe* Lanes together would make a x2 PCIe* Link). A PCIe* Link is addressed by the lowest number PCIe* Lane it connects to and is known as the PCIe* Root Port (such as a x2 PCIe* Link connected to PCIe* Lanes 3 and 4 would be called x2 PCIe* Root Port 3).
 10. The PCIe* Lanes can be configured independently from one another but the max number of configured Root Ports (Devices) must not be exceeded
 11. Unidentified lanes within a PCIe* Link Configuration are disabled but their physical lanes are used for the identified Root Port
-

Figure 18. PCIe Controller Bifurcation/Configuration Mode Strap Details

Processor		SOC (System On Chip) Tile																							PCIe Controller Configuration Mode		
		FIA-2																									
PCIe I/O Lanes		8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	GPP_SA[16:15] Pin Strap Settings		Soft Config							
PCIe Controllers		1 and 2																									
PCIe Lanes		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16										
PCIe Configurations (Bifurcation)	Configuration # -->	1	1px16 / 1px16(LR)														00b = 1px16										
		2	1px16 / 1px16(LR)														00b = 1px16										
		3	1px8						1px8						10b = 2px8												
		4	1px8						1px8						10b = 2px8												
		5	1px8(LR)						1px8(LR)						10b = 2px8												
		6	1px8(LR)						1px8(LR)						10b = 2px8												
		7	1px8						1px4			1px4 / 1px4(LR)			11b = 1px8 + 1px4 + 1px4		1x8, 2x4										
		8	1px8(LR)						1px4(LR) / 1px4			1px4(LR)			11b = 1px8 + 1px4 + 1px4		1x8, 1x4										
		9	1px8 / 1px8(LR)						Empty			1px4 / 1px4(LR)			11b = 1px8 + 1px4 + 1px4		1x8, 2x4										
		10	1px8						1px4			Empty			11b = 1px8 + 1px4 + 1px4												
		11	1px8(LR)						1px4(LR)			Empty			11b = 1px8 + 1px4 + 1px4		1x8										
		12	1px16 / 1px16(LR)														00b = 1px16										
		13	1px8						1px8						10b = 2px8												
		14	1px8(LR)						1px8(LR)						10b = 2px8												

NOTES

1. The SOC Tile x16 Gen5 PCIe Controller Bifurcation modes, for PCIe Lanes 1 to 16, are configured based off Pin Strap and Soft Strap settings. The Pin Strap settings are done with the GPP_SA[16:15] signal pins where motherboard designs may implement them using board switches, jumpers, or through pull-ups and pull-downs. Both of these Pin Straps and Soft Straps must be set to match the motherboard implementation for the SOC Tile x16 Gen5 PCIe Controller Bifurcation.
 2. One has to understand the impact of these SOC Tile strap configuration settings between the SOC GPP_SA[16:15] Pin Straps and the SOC Soft Straps. Incorrect or mismatched settings could result in a Root Port and link not training to the expected data width or an end point device not getting detected.
 - If the GPP_SA Pin Straps and the Soft Straps do not match then the one with the lowest data width strap value between the two will have priority and set the bi-furcation for the associated PCIe Controller Root Port.
 - If the GPP_SA Pin Straps and the Soft Straps both match then the data width value from them results in the same priority and sets the bi-furcation for the associated PCIe Controller Root Port.
 - The Soft Straps set the Lane Reversal (LR) for the associated PCIe Controller Root Port.
 - Examples:
 - Case1: GPP_SA [16:15] = 10b (2px8) and Soft Straps = 1x16 --> Results in 2px8 Root Port Configuration
 - Case2: GPP_SA [16:15] = 10b (2px8) and Soft Straps = 1x16 LR --> Results in 1px8 LR + 1px8 LR Root Port Configuration
 - Case3: GPP_SA [16:15] = 00b (1px16) and Soft Straps = 2x8 --> Results in 2px8 Root Port Configuration
 - Case4: GPP_SA [16:15] = 11b (1px8 + 1px4 + 1px4) and Soft Straps = 1x8 LR, 1x8 LR --> Results in 1px8 LR + 1px4 LR + 1px4 LR Root Port Configuration
 3. The IOE Tile PCIe Gen4 and Gen5 x4 Controller configuration modes are set only through the IOE Soft Straps.
-

20.0 Graphics

20.1 Processor Graphics

The processor graphics is based on Xe graphics core architecture that enables substantial gains in performance and lower-power consumption over prior generations. Xe architecture supports up to 8 Xe-core 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. Xe 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.

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

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

20.1.1.1 Hardware Accelerated Video Decode

Xe 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:

- Direct3D11 Video API
- Direct3D12 Video API
- Intel Media SDK
- MFT (Media Foundation Transform) filters¹
- Intel VA API ²
- Intel one VPL

NOTES

1. Only for JPEG Decoder
 2. Only for Linux*
-

Xe supports full HW accelerated video decoding for MPEG2/AVC/HEVC/VP9/JPEG/AV1.

Table 66. Hardware Accelerated Video Decoding

Codec	Profile	Level	Maximum Resolution
MPEG2	Main	Main - 15Mbps High - 40Mbps	FHD
AVC/H264	High Main Constrained Baseline	L5.2	4K
	4:2:0 8bit		4K @ 60
JPEG/MJPEG	Baseline	Unified level	16K x16K
HEVC/H265	Main12 420, 422, 444 - 8b/10b/12b SCC 420, 444 - 8b/10b	L6.1	8K @ 60(Decode Only) 8K@30(Decode Playback)
VP9	0 (420 8b) 1 (444 8b) 2 (420 10b/12b) 3 (444 10b/12b)	Unified level	8K @ 60(Decode only) 8K@30 (Decode Playback) 16Kx4K
AV1	Main (420 8-bit/10b)	L6.1	8K @ 60 (Video, Decode only) 8K@30 (Decode Playback) 16K x 16K (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.

20.1.1.2 Hardware Accelerated Video Encode

Xe implements a low-power low-latency fixed function encoder which supports AVC, HEVC, VP9 and AV1.

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

- Direct3D12 Video API
- Intel® one VPL
- MFT (Media Foundation Transform) filters [Only for AVC/HEVC/JPEG/AV1 Encoder]

Xe supports full HW accelerated video encoding for AVC/HEVC/VP9/JPEG/AV1.

Table 67. Hardware Accelerated Video Encode

Codec	Profile	Level	Maximum Resolution
AVC/H264	High Main Constrained Baseline	L5.2	4K@60
JPEG			16Kx16K
HEVC/H265	Main10 422 - 8b/10b	L5.2	4K@60

continued...

Codec	Profile	Level	Maximum Resolution
	Main Main10 420, 444 - 8b/10b SCC 420 444 - 8b/10b	L6.1	4320p(8K) @60 16Kx12K
VP9	0 (420 8b) 1 (444 8b) 2 (420 10b) 3 (444 10b)	—	8K @30
AV1	Main (4:2:0 8b, 10b)	L6	8K@30

NOTE

Hardware encode for H264 SVC is not supported.

20.1.1.3 Hardware Accelerated Video Processing

There is hardware support for image processing functions such as De-interlacing, Film cadence detection, detail enhancement, gamut compression, Adaptive contrast enhancement, skin tone enhancement, total color control, De-noise, SFC (Scalar and Format Conversion), memory compression, Localized Adaptive Contrast Enhancement (LACE), 16 bpc support for de-noise/de-mosaic, Facial filter, HDR10 and Dolby Vision Tone Mapping HW acceleration.

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

- Direct3D* 11 Video API
- Intel® One VPL
- Intel® Graphics Control Library
- Intel VA API

NOTE

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

20.1.1.4 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.
- Low power Scaler and Format Converter.

20.1.2 Graphics Core Cache

The Xe Graphics Core architecture has a hierarchy of caches which contains first, second and third level caches.

First and Second Level Cache

The first and second level cache is a lower-level Instruction and the Data caches. They are implemented close to the Xe/3D compute elements, decode/encode and media pipelines. These cache units are not shared between the different units.

Third Level Cache

Third level cache is a central memory cache that is implemented as higher cache hierarchy. The device cache is a multi-way set-associative that allow memory pages to be cached either coherently or non-coherently with respect to an external memory system.

20.2 Platform Graphics Hardware Feature

20.2.1 Hybrid Graphics

Microsoft* Windows* 11 operating system enables the Windows*11 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.
 - a. Microsoft* Hybrid DOES NOT support hot-plugging of GPUs
 - b. OEMS should seek further guidance from Microsoft* before designing systems with the concept of hot-plugging
4. Starting with Windows*11 (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.

21.0 Display

This chapter provides information on the following topics:

- Display Technologies Support
- Display Configuration
- Display Features

21.1 Display Technologies Support

Technology	Standard
eDP* 1.4b	VESA* Embedded DisplayPort* Standard 1.4b
DisplayPort* 2.1	VESA* DisplayPort* Standard 2.1
HDMI* 2.1	High-Definition Multimedia Interface Specification Version 2.1

Table 68. Display Ports Availability and Link Rate

Port	S-Series Processor
DDI A	eDP*, DP up to HBR3, HDMI up to 6GHz
DDI 2	DP* up to UHBR20 HDMI* up to 12 Gbps
DDI 3	
TCP 0	
TCP 1	
<i>Note:</i> 1. MIPI DSI can be supported using on-board eDP to DSI bridge.	

21.2 Display Interfaces

This section provides information on the following topics:

- Digital Display Interface (DDI) Signals
- Digital Display Interface TCP Signals

21.2.1 Digital Display Interface DDI Signals

Table 69. Digital Display Interface DDI Signals

Signal Name	Type	Description
DDIA_TX_P[3:0] DDIA_TX_N[3:0]	O	Digital Display Interface A (DDIA): Digital Display Interface main link transmitter lanes.
DDIA_AUX_P DDIA_AUX_N	I/O	Digital Display Interface A (DDIA): DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.
<i>continued...</i>		

Signal Name	Type	Description
GPP_SD12/ DDSP_HPDA /DISP_MISCA ¹ (Option 1)	I	Digital Display Interface A (DDIA): Hot Plug Detect (HPD).
GPP_SA14/ DDSP_HPDA /DISP_MISCA ¹ (Option 2)		
GPP_SD09/ VDDEN ¹ (Option 1)	O	Digital Display Interface A (DDIA): eDP Panel power control enable signal.
GPP_SA21/ VDDEN ¹ (Option 2)		
GPP_SD10/ BKLTEN ¹ (Option 1)	O	Digital Display Interface A (DDIA): eDP Panel back-light control enable signal.
GPP_SA22/RSVD/ BKLTEN ¹ (Option 2)		
GPP_SD11/ BKLTCTL ¹ (Option 1)	O	Digital Display Interface A (DDIA): eDP Panel back-light control Pulse Wide Modulation (PWM) signal.
GPP_SA23/RSVD/ BKLTCTL ¹ (Option 2)		
DDI2_TX_P[3:0] DDI2_TX_N[3:0] DDI3_TX_P[3:0] DDI3_TX_N[3:0]	O	Digital Display Interface x (DDIx): Digital Display Interface main link transmitter lanes.
DDI2_AUX_P DDI2_AUX_N DDI3_AUX_P DDI3_AUX_N	I/O	Digital Display Interface x (DDIx): DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.
GPP_SB02/ DDP2_CTRLCLK GPP_SB04/ DDP3_CTRLCLK GPP_SB03/ DDP2_CTRLDATA GPP_SB05/ DDP3_CTRLDATA	I/O	Digital Display Interface x (DDIx): HDMI Graphics Management Bus (GMBUS).
GPP_SD15/ DDSP_HPDA /DISP_MISC1 ¹ (Option 1) GPP_SD16/ DDSP_HPDA /DISP_MISC2 ¹ (Option 1) GPP_SD13/ DDSP_HPDA /DISP_MISC3 ¹ (Option 1) GPP_SD14/ DDSP_HPDA /DISP_MISC4 ¹ (Option 1) GPP_SA17/ DDSP_HPDA /DISP_MISC1 ¹ (Option 2) GPP_SA18/ DDSP_HPDA /DISP_MISC2 ¹ (Option 2) GPP_SA15/ DDSP_HPDA /DISP_MISC3 ¹ (Option 2) GPP_SA16/ DDSP_HPDA /DISP_MISC4 ¹ (Option 2)	I	Digital Display Interface x (DDIx): Hot Plug Detect (HPD).
GPP_SD12/DDSP_HPDA/ DISP_MISCA GPP_SD15/DDSP_HPDA1/ DISP_MISC1 GPP_SD16/DDSP_HPDA2/ DISP_MISC2 GPP_SD13/DDSP_HPDA3/ DISP_MISC3 GPP_SD14/DDSP_HPDA4/ DISP_MISC4	O	DDI Misc signals.
GPP_SB00/ DDPA_CTRLCLK /SRCCLKREQ2#	I/O	(S-Processor only) Digital Display Interface x (DDIA): HDMI Graphics Management Bus (GMBUS). DDC CTRLCLK
GPP_SB01/ DDPA_CTRLDATA /SRCCLKREQ3#	I/O	(S-Processor only) Digital Display Interface x (DDIA): HDMI Graphics Management Bus (GMBUS). DDC CTRLDATA
<p><i>Notes:</i> 1. These signals are alternate signals with identical functionality, either Option 1 or Option 2 should be used. 2. Auxiliary Channel (AUX CH) is a half-duplex bidirectional channel used for link management and device control. AUX CH is an AC coupled differential signal. 3. GMBUS follows I2C Protocol.</p>		

21.2.2 Digital Display Interface TCP Signals

Table 70. Digital Display Interface TCP Signals

Signal Name	Type	Description
TCP0_TXRX[1:0]_P TCP0_TXRX[1:0]_N TCP0_TX[1:0]_P TCP0_TX[1:0]_N	O	Digital Display Interface 0 (TCP0): Digital Display Interface main link transmitter lanes.
TCP0_AUX_P TCP0_AUX_N	I/O	Digital Display Interface 0 (TCP0): DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.
GPP_SB13/DDP0_CTRLDATA/ TBT_LSX0_RXD GPP_SB12/DDP0_CTRLCLK/ TBT_LSX0_TXD	I/O	Digital Display Interface 0 (TCP0): HDMI Graphics Management Bus (GMBUS).
GPP_SD15/DDSP_HPD1/ DISP_MISC1	I	Digital Display Interface 0 (TCP0): Hot Plug Detect (HPD).
TCP1_TXRX[1:0]_P TCP1_TXRX[1:0]_N TCP1_TX[1:0]_P TCP1_TX[1:0]_N	O	Digital Display Interface 1 (TCP1): Digital Display Interface main link transmitter lanes.
TCP1_AUX_P TCP1_AUX_N	I/O	Digital Display Interface 1 (TCP1): DisplayPort Auxiliary: Half-duplex, bidirectional channel consist of one differential pair.
GPP_SB16/DDP1_CTRLDATA/ TBT_LSX1_RXD GPP_SB15/DDP1_CTRLCLK/ TBT_LSX1_TXD	I/O	Digital Display Interface 1 (TCP1): HDMI Graphics Management Bus (GMBUS).
GPP_SD16/DDSP_HPD2/ DISP_MISC2	I	Digital Display Interface 1 (TCP1): Hot Plug Detect (HPD).
<i>Notes:</i> <ul style="list-style-type: none"> • Auxiliary Channel (AUX CH) is a half-duplex bidirectional channel used for link management and device control. AUX CH is an AC coupled differential signal. • GMBUS follows I2C Protocol. 		

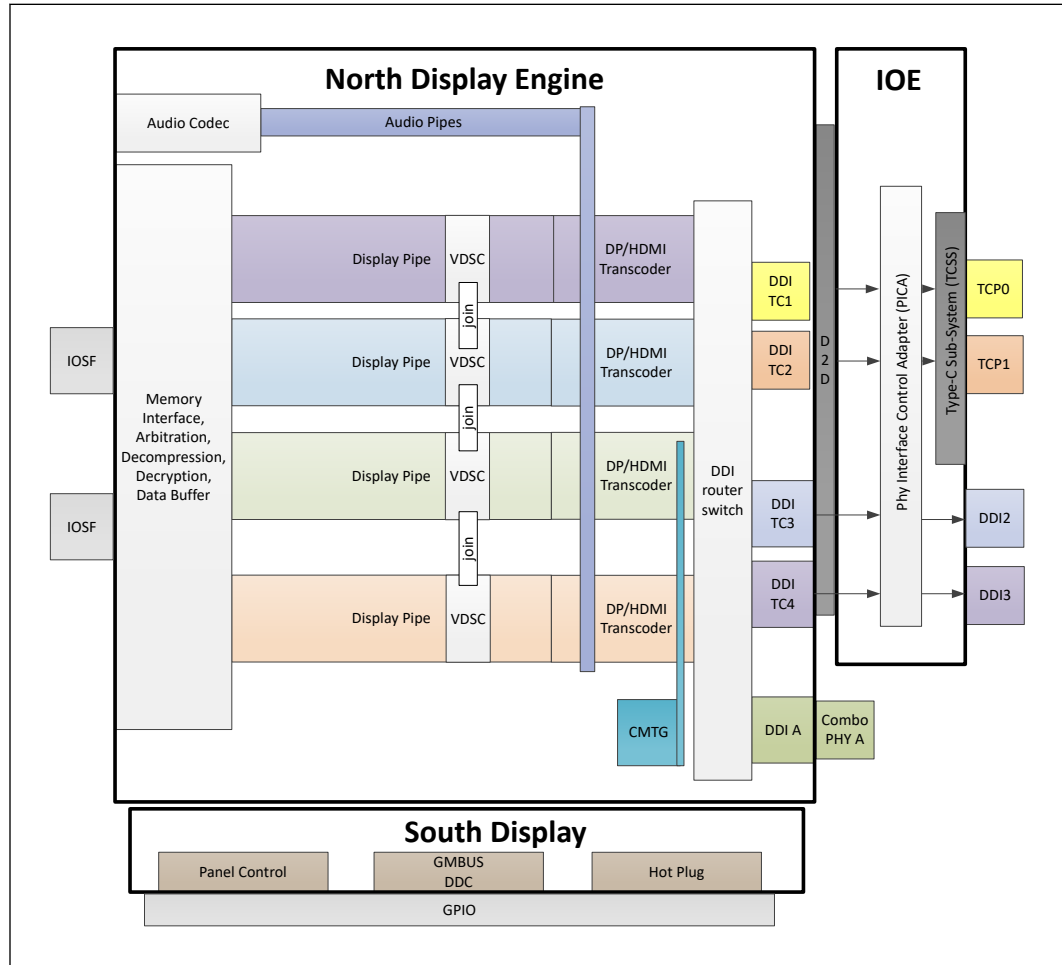
21.3 Display Features

This section provides information on the following topics:

- General Capabilities
- Multiple Display Configurations
- High-bandwidth Digital Content Protection (HDCP)
- DisplayPort*
- High-Definition Multimedia Interface (HDMI*)
- embedded DisplayPort* (eDP*)
- Integrated Audio

21.3.1 General Capabilities

Figure 19. Processor Display Architecture



NOTE

For port availability in each of the processor series, refer to [Table 68](#) on page 142.

- Up to four simultaneous displays, 4K60Hz Embedded panel concurrent with:
 - Single external panel up to 8K60Hz, supported by joining two pipes over single port.
 - Up to 3x4K60Hz External panel.
- Display interfaces supported:
 - DDI interfaces supports DP*, HDMI*, eDP*
 - TCP interfaces supports DP*, HDMI*, Display Alt Mode over Type-C and DP* tunneled.
- End-To-End (E2E) compression, Unified memory compression across GT, media and display.

- Audio stream support on external ports.
- HDR (High Dynamic Range) support.
- Four Display Pipes - Supporting blending, color adjustments, scaling and dithering.
- Transcoder - Containing the Timing generators supporting eDP*, DP*, HDMI* interfaces.
- One Low Power optimized pipes supporting Embedded DisplayPort*
 - 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.

21.3.2 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.

21.3.3 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 content protection over wired displays (HDMI* and DisplayPort*).

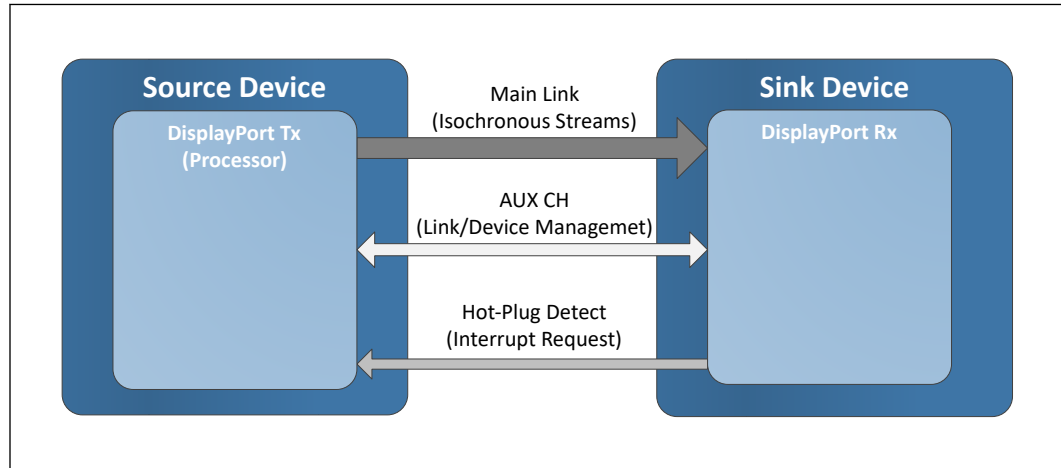
21.3.4 DisplayPort*

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.

Figure 20. DisplayPort* Overview



- Supports main link of 1, 2, or 4 data lanes.
- Link rate supports up to UHBR20 (UHBR13.5 is not supported)
- Aux channel for Link/Device management
- Hot Plug Detect
- Supports up to 36 BPP (Bit Per Pixel)
- Supports SSC
- Supports YCbCR 4:4:4, YCbCR 4:2:0, YCbCR 4:2:2, and RGB color format
- Supports MST (Multi-Stream Transport)
- Supports VESA DSC 1.2b
- Supports panel replay
- Adaptive Sync

21.3.4.1 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 71. Display Resolutions and Link Bandwidth for Multi-Stream Transport Calculations

Pixels per Line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
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
<i>continued...</i>				

Pixels per Line	Lines	Refresh Rate [Hz]	Pixel Clock [MHz]	Link Bandwidth [Gbps]
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:

- All the above is related to bit depth of 24.
- The data rate for a given video mode can be calculated as
Data Rate = Pixel Frequency * Bit Depth
- The bandwidth requirements for a given video mode can be calculated as: Bandwidth = Data Rate * 1.25 (for 8b/10b coding overhead).
- The link bandwidth depends if the standards is reduced blanking or not.
If the standard is not reduced blanking - the expected bandwidth may be higher.
For more details, refer to VESA and Industry Standards and Guidelines for Computer Display Monitor Timing (DMT). Version 1.0, Rev. 13 February 8, 2013
- To calculate what are the resolutions that can be supported in MST configurations, follow the below guidelines:
 - Identify what is the link bandwidth column according to the requested display resolution.
 - 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:
 - Docking two displays: 3840x2160@60 Hz + 1920x1200@60 Hz = 16 + 4.62 = 20.62 Gbps [Supported]
 - Docking three displays: 3840x2160@30 Hz + 3840x2160@30 Hz + 1920x1080@60 Hz = 7.88 + 7.88 + 4.16 = 19.92 Gbps [Supported].
- MST bandwidth number is calculated without VESA Display Stream Compression (VDSC).

Table 72. DisplayPort Maximum Resolution

Standard	S-Series Processor
DP*	8K60Hz compressed, 5K120Hz compressed

Notes:

- bpp - bit per pixel.
- Resolution support is subject to memory BW availability.
- High resolutions will consume two display pipes.

21.3.5 High-Definition Multimedia Interface (HDMI*)

The High-Definition Multimedia Interface (HDMI*) is provided for transmitting 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) or Fixed Rate Link (FRL) to carry audiovisual information through the same HDMI cable.

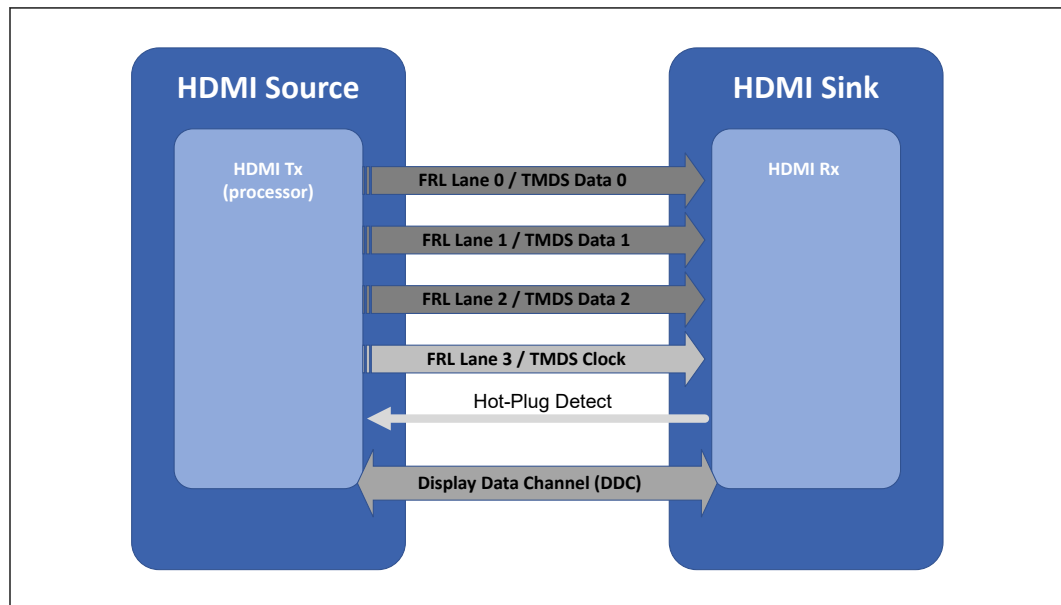
HDMI* includes three separate communications channels: TMDS or FRL, DDC/GMBUS, and the optional CEC (consumer electronics control). CEC is not supported on the processor. As shown in the following figure, the HDMI* cable carries four differential pairs that make up the TMDS data and clock channels or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

TMDS architecture has 3 Data lane and 1 Clock lane

FRL architecture has 4 Data lane, and no clock lane

Figure 21. HDMI* Overview



- Supports up to 6Gbps TMDS link rates on 3 lanes
- Supports up to 12Gbps FRL link rates on 4 lanes
- Support YCbCR 4:4:4, YCbCR 4:2:0, YCbCR 4:2:2, and RGB color format
- Supports up to 36 BPP (Bit Per Pixel)
- Supports VESA DSC 1.2a in FRL mode
- Hot Plug Detect
- Adaptive Sync supported in FRL mode

Table 73. HDMI Maximum Resolution

Standard	S-Series Processor
HDMI 2.1 (Up to 6Gbps)	4K60 Hz 24 bpp
HDMI 2.1 (Up to 12 Gbps)	8K60Hz Compressed, 5K120Hz compressed

continued...

Standard	S-Series Processor
	4K144Hz Compressed
<i>Notes:</i> 1. bpp - bit per pixel. 2. Resolution support is subject to memory BW availability. 3. Compressed mean DSC only	

21.3.6 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.

- Supports Low power optimized pipes
- Supports up to HBR3 link rate
- Supports Backlight PWM control and enable signals, and power enable
- Supports VESA DSC 1.2a
- Supports SSC
- Panel Self Refresh 1
- Panel Self Refresh 2
- MSO 2x2, 4x1(Multi Segment Operation)
- Dedicated Aux channel
- Adaptive Sync

Table 74. Embedded DisplayPort Maximum Resolution

Standard	S-Series Processor ¹
eDP*	4K60Hz HDR 5K60Hz SDR
<i>Notes:</i> 1. Maximum resolution is based on the implementation of 4 lanes at HBR3 link data rate. 2. 5k120Hz cannot work with PSR*. 3. Resolution support is subject to memory BW availability.	

21.3.7 Integrated Audio

- HDMI* and DisplayPort interfaces can carry audio along with video.
- The processor supports up to four High Definition Audio streams on four digital ports simultaneously.

Table 75. Processor Supported Audio Formats over HDMI* and DisplayPort*

Audio Formats	HDMI*	DisplayPort*
AC-3 Dolby* Digital	Yes	Yes
Dolby* Digital Plus	Yes	Yes
<i>continued...</i>		



Audio Formats	HDMI*	DisplayPort*
DTS-HD*	Yes	Yes
LPCM, 32 KHz, 44.1 KHz, 48 KHz, 88.2 KHz, 96 KHz, 176.4 KHz, and 192 KHz, 16/24 bit, 2/4/6/8 channels	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 32 KHz, 44.1 KHz, 48 KHz, 88.2 KHz, 96 KHz, 176.4 KHz, and 192 KHz sampling rates and silent multi-stream support.

22.0 General Purpose Input and Output

The General Purpose Input/Output (GPIO) signals are grouped into multiple groups (GPP_SA, GPP_SB, and GPP_SD). All GPIO groups are powered by the Primary well.

The high level features of GPIO:

- 1.8 V
- Integrated programmable pull-up / pull-down
- Configurable as GPIO input, GPIO output, or native function signal
- SCI (GPE) capable on all GPIOs
- NMI and SMI capable (on selected GPIOs)

Table 76. Acronyms

Acronyms	Description
GPI	General Purpose Input
GPO	General Purpose Output
GPP	General Purpose I/O in Primary Well

22.1 Functional Description

22.1.1 Interrupt / IRQ via GPIO Requirement

A GPIO, as an input, can be used to generate an interrupt / IRQ to the processor. In this case, it is required that the pulse width on the GPIO must be at least 100 us for the processor to recognize the interrupt.

22.1.2 Integrated Pull-ups and Pull-downs

All GPIOs have programmable internal pull-up/pull-down resistors. The internal pull-up/pull-down for each GPIO can be statically enabled/disabled by BIOS programming the corresponding TERM field in PAD_CFG_DW1 register (refer to Volume 2 for more details on the register).

Note that, although most of the GPIOs have the internal pull-up/pull-down disabled by default, some of them have the internal pull-up/pull-down enabled by default. Refer to the GPIO Implementation Summary document for detailed info on default termination.

22.1.3 SCI / SMI and NMI

SCI capability is available on all GPIOs, while SMI and NMI capability is available on only select GPIOs.

Below are the GPIOs that can be routed to generate SMI# or NMI:

- GPP_SB14
- GPP_SD[4:0]

22.1.4 Timed GPIO

The processor supports two Timed GPIOs as native function (TIME_SYNC) that is multiplexed on GPIO pins. The intent usage of the Timed GPIO function is for time synchronization purpose.

Timed GPIO can be an input or an output:

- As an input, a GPIO input event triggers the HW to capture the processor Always Running Timer (ART) time in the Time Capture register. The GPIO input event must be asserted for at least two crystal oscillator clocks period in order for the event to be recognized.
- As an output, a match between the ART time and the software programmed time value triggers the HW to generate a GPIO output event and capture the ART time in the Time Capture register. If periodic mode is enabled, HW generates the periodic GPIO events based on the programmed interval. The GPIO output event is asserted by HW for at least two crystal oscillator clock periods.

NOTE


TIME_SYNC can be set as input when both Direction (DIR) bit and Enable (EN) bit in Timed GPIO Control Register are set to 1 (refer to Datasheet Vol2 for the register info). When EN bit is set to 0, TIME_SYNC will default to output low regardless of DIR bit setting.

Timed GPIO supports event counter. When Timed GPIO is configured as input, event counter increments by one for every input event triggered. When Timed GPIO is configured as output, event counter increments by one for every output event generated. The event counter provides the correlation to associate the Timed GPIO event (the nth event) with the captured ART time. The event counter value is captured when a read to the Time Capture Value register occurs.

NOTE

When Timed GPIO is enabled, the crystal oscillator will not be shut down as crystal clock is needed for the Timed GPIO operation. As a result, SLP_S0# will not be asserted. This has implication to platform power (such as IDLE or S0ix power). Software should only enable Timed GPIO when needed and disable it when Timed GPIO functionality is not required.

22.2 Signal Description

For GPIO pin implementation including multiplexed native functions, default values, signal states, and other characteristics, download the pdf, click  on the navigation pane and refer the spreadsheet, **832586-001_GPIO.xlsx**.

23.0 Interrupt Timer Subsystem (ITSS)

Table 77. Acronyms

Acronym	Description
ITSS	Interrupt Timer Subsystem
HPET	High Precision Event Timer
8254 PIT	Legacy 8254 Programmable Interrupt Timer
INTR	Interrupt
NMI	Non-maskable Interrupt
INIT	Processor Initialization
SERR	System Error

Table 78. References

Specification	Document Number/Location
ACPI Specification, Rev 5.0a	https://uefi.org/acpi/specs

23.1 Feature Overview

ITSS supports following features:

- It houses the HPET, Legacy 8254 Timers and APIC Interrupt Controllers.
- Fully synchronous-based design adopted for 8254 PIT.
- Functions as a simple Internal Host Space Error Collector and Reporting Block.
- 8254 PIT - Consists of 3 16-bit Timers capable of supporting up to 6 different modes.
- APIC - Supports up to 120 IRQs.
- HPET - Contains 8 Timer Blocks and a single always running 64-bit counter. Each Timer is interrupt capable, with option to route to APIC or directly to hose using MSI. Improved resolution, reduced overhead in comparison to Legacy 8254, IOxAPIC & RTC Timers.

23.2 Functional Description

The ITSS (Interrupt Timer Sub System) have below sub blocks:

- **ITSS** : Consists of the HPET, 8254 and APIC.

23.2.1 8254 Timers

There are three counters that have fixed uses. All registers and functions associated with these timers are in the Primary well. The 8254 unit is clocked by a 1.193 MHz periodic timer tick, which is functional only in S0 states. The 1.193 MHz periodic timer tick is generated off the XTAL clock.

Counter 0, System Timer

This counter functions as the system timer by controlling the state of IRQ0 and is typically programmed for Mode 3 operation. The counter produces a square wave with a period equal to the product of the counter period (838 ns) and the initial count value. The counter loads the initial count value 1 counter period after software writes the count value to the counter I/O address. The counter initially asserts IRQ0 and decrements the count value by two each counter period. The counter negates IRQ0 when the count value reaches 0. It then reloads the initial count value and again decrements the initial count value by two each counter period. The counter then asserts IRQ0 when the count value reaches 0, reloads the initial count value, and repeats the cycle, alternately asserting and negating IRQ0.

23.2.1.1 Timer Programming

The counter/timers are programmed in the following fashion:

1. Write a control word to select a counter.
2. Write an initial count for that counter.
3. Load the least and/or most significant bytes (as required by Control Word bits 5, 4) of the 16 bit counter.
4. Repeat with other counters.

Only two conventions need to be observed when programming the counters. First, for each counter, the control word must be written before the initial count is written. Second, the initial count must follow the count format specified in the control word (least significant Byte only, most significant Byte only, or least significant Byte, and then most significant Byte).

A new initial count may be written to a counter at any time without affecting the counter's programmed mode. Counting is affected as described in the mode definitions. The new count must follow the programmed count format.

If a counter is programmed to read/write 2-byte counts, the following precaution applies – a program must not transfer control between writing the first and second Byte to another routine, which also writes into that same counter. Otherwise, the counter will be loaded with an incorrect count.

The Control Word Register at port 43h controls the operation of counter. Several commands are available:

- **Control Word Command:** Specifies which counter to read or write, the operating mode, and the count format (binary or BCD).
- **Counter Latch Command:** Latches the current count so that it can be read by the system. The countdown process continues.
- **Read Back Command:** Reads the count value, programmed mode, the current state of the OUT pins, and the state of the Null Count Flag of the selected counter.

The table below lists the six operating modes for the interval counters:

Table 79. Counter Operating Modes

Mode	Function	Description
0	Out signal on end of count (=0)	Output is 0. When count goes to 0, output goes to 1 and stays at 1 until counter is reprogrammed.
1	Hardware retriggerable one-shot	Output is 0. When count goes to 0, output goes to 1 for one clock time.
2	Rate generator (divide by n counter)	Output is 1. Output goes to 0 for one clock time, then back to 1 and counter is reloaded.
3	Square wave output	Output is 1. Output goes to 0 when counter rolls over, and counter is reloaded. Output goes to 1 when counter rolls over, and counter is reloaded, and so on
4	Software triggered strobe	Output is 1. Output goes to 0 when count expires for one clock time.
5	Hardware triggered strobe	Output is 1. Output goes to 0 when count expires for one clock time.

23.2.1.2 Reading from Interval Timer

It is often desirable to read the value of a counter without disturbing the count in progress. There are three methods for reading the counters—a simple read operation, counter Latch command, and the Read-Back command. Each one is explained below:

With the simple read and counter latch command methods, the count must be read according to the programmed format; specifically, if the counter is programmed for 2-byte counts, 2-bytes must be read. The 2-bytes do not have to be read one right after the other. Read, write, or programming operations for other counters may be inserted between them.

Simple Read

The first method is to perform a simple read operation. The counter is selected through Port 40h (Counter 0).

NOTE

Performing a direct read from the counter does not return a determinate value, because the counting process is asynchronous to read operations.

Counter Latch Command

The Counter Latch command, written to Port 43h, latches the count of a specific counter at the time the command is received. This command is used to ensure that the count read from the counter is accurate, particularly when reading a 2-byte count. The count value is then read from each counter’s Count register as was programmed by the Control register.

The count is held in the latch until it is read or the counter is reprogrammed. The count is then unlatched. This allows reading the contents of the counters on the fly without affecting counting in progress. Multiple Counter Latch Commands may be used to latch more than one counter. Counter Latch commands do not affect the programmed mode of the counter in any way.

If a Counter is latched and then, sometime later, latched again before the count is read, the second Counter Latch command is ignored. The count read is the count at the time the first Counter Latch command was issued.

Read Back Command

The Read Back command, written to Port 43h, latches the count value, programmed mode, and current states of the OUT pin and Null Count flag of the selected counter or counters. The value of the counter and its status may then be read by I/O access to the counter address.

The Read Back command may be used to latch multiple counter outputs at one time. This single command is functionally equivalent to several counter latch commands, one for each counter latched. Each counter's latched count is held until it is read or reprogrammed. Once read, a counter is unlatched. The other counters remain latched until they are read. If multiple count Read Back commands are issued to the same counter without reading the count, all but the first are ignored.

The Read Back command may additionally be used to latch status information of selected counters. The status of a counter is accessed by a read from that counter's I/O port address. If multiple counter status latch operations are performed without reading the status, all but the first are ignored.

Both the count and status of the selected counters may be latched simultaneously. This is functionally the same as issuing two consecutive, separate Read Back commands. If multiple count and/or status Read Back commands are issued to the same counters without any intervening reads, all but the first are ignored.

If both the count and status of a counter are latched, the first read operation from that counter returns the latched status, regardless of which was latched first. The next one or two reads, depending on whether the counter is programmed for one or two type counts, returns the latched count. Subsequent reads return unlatched count.

23.2.2 APIC Advanced Programmable Interrupt Controller

The APIC is accessed via an indirect addressing scheme. These registers are mapped into memory space. These are programmable through PCI Config IOAC register.

23.2.3 High Precision Event Timer (HPET)

This function provides a set of timers that can be used by the operating system. The timers are defined such that the operating system may assign specific timers to be used directly by specific applications. Each timer can be configured to cause a separate interrupt.

The processor provides eight timers. The timers are implemented as a single counter with a set of comparators. Each timer has its own comparator and value register. The counter increases monotonically. Each individual timer can generate an interrupt when the value in its value register matches the value in the main counter.

Timer 0 supports periodic interrupts.

The registers associated with these timers are mapped to a range in memory space (much like the I/O APIC). However, it is not implemented as a standard PCI function. The BIOS reports to the operating system the location of the register space using

ACPI. The hardware can support an assignable decode space; however, BIOS sets this space prior to handing it over to the operating system. It is not expected that the operating system will move the location of these timers once it is set by BIOS.

23.2.3.1 Timer Accuracy

The timers are accurate over any 1 ms period to within 0.05% of the time specified in the timer resolution fields.

Within any 100 us period, the timer reports a time that is up to two ticks too early or too late. Each tick is less than or equal to 100 ns; thus, this represents an error of less than 0.2%.

The timer is monotonic. It does not return the same value on two consecutive reads (unless the counter has rolled over and reached the same value).

The main counter uses the XTAL as its clock. The accuracy of the main counter is as accurate as the crystal that is used in the system.

23.2.3.2 Timer Off-load

The timer off-load feature allows the HPET timers to remain operational during very low power S0 operational modes when the XTAL clock is disabled. The clock source during this off-load is the Real Time Clock's 32.768 kHz clock. This clock is calibrated against the XTAL clock during boot time to an accuracy that ensures the error introduced by this off-load is less than 10 ppb (0.000001%).

When the XTAL clock is active, the 64 bit counter will increment by one each cycle of the XTAL clock when enabled. When the XTAL clock is disabled, the timer is maintained using the RTC clock. The long-term (> 1 ms) frequency drift allowed by the HPET specification is 500 ppm. The off-load mechanism ensures that it contributes < 1 ppm to this, which will allow this specification to be easily met given the clock crystal accuracies required for other reasons.

Timer off-load is prevented when there are HPET comparators active.

The HPET timer runs typically on the XTAL crystal clock and is off-loaded to the 32 kHz clock once the processor enters C10. This is the state where there are no C10 wake events pending and when the off-load calibrator is not running. HPET timer re-uses this 28 bit calibration value calculated by PMC when counting on the 32 kHz clock. During C10 entry, PMC sends an indication to HPET to off-load and keeps the indication active as long as the processor is in C10 on the 32 kHz clock. The HPET counter will be off-loaded to the 32 kHz clock domain to allow the XTAL clock to shut down when it has no active comparators.

Theory of Operation

The Off-loadable Timer Block consists of a 64 bit fast clock counter and an 82 bit slow clock counter. During fast clock mode the counter increments by one on every rising edge of the fast clock. During slow clock mode, the 82 bit slow clock counter will increment by the value provided by the Off-load Calibrator.

The Off-loadable Timer will accept an input to tell it when to switch to the slow RTC clock mode and provide an indication of when it is using the slow clock mode. The switch will only take place on the slow clock rising edge, so for the 32 kHz RTC clock the maximum delay is around 30 us to switch to or from slow clock mode. Both of these flags will be in the fast clock domain.

When transitioning from fast clock to slow clock, the fast clock value will be loaded into the upper 64 bits of the 82 bit counter, with the 18 LSBs set to zero. The actual transition though happens in two stages to avoid metastability. There is a fast clock sampling of the slow clock through a double flop synchronizer. Following a request to transition to the slow clock, the edge of the slow clock is detected and this causes the fast clock value to park. At this point the fast clock can be gated. On the next rising edge of the slow clock, the parked fast clock value (in the upper 64 bits of an 82 bit value) is added to the value from the Off-load Calibrator. On subsequent edges while in slow clock mode the slow clock counter increments its count by the value from the Off-load Calibrator.

When transitioning from slow clock to fast clock, the fast clock waits until it samples a rising edge of the slow clock through its synchronizer and then loads the upper 64 bits of the slow clock value as the fast count value. It then de-asserts the indication that slow clock mode is active. The 32 kHz clock counter no longer counts. The 64 bit MSB will be over-written when the 32 kHz counter is reloaded once conditions are met to enable the 32 kHz HPET counter but the 18 bit LSB is retained and it is not cleared out during the next reload cycle to avoid losing the fractional part of the counter.

After initiating a transition from fast clock to slow clock and parking the fast counter value, the fast counter no longer tracks. This means if a transition back to fast clock is requested before the entry into off-load slow clock mode completes, the Off-loadable Timer must wait until the next slow clock edge to restart. This case effectively performs the fast clock to slow clock and back to fast clock on the same slow clock edge.

23.2.3.3 Periodic Versus Non-Periodic Modes

Non-Periodic Mode

This mode can be thought of as creating a one-shot timer.

When a timer is set up for non-periodic mode, it will generate an interrupt when the value in the main counter matches the value in the timer's comparator register. Another interrupt will be generated when the main counter matches the value in the timer's comparator register after a wrap around.

During run-time, the value in the timer's comparator value register will not be changed by the hardware. Software can of course change the value.

The Timer 0 Comparator Value register cannot be programmed reliably by a single 64 bit write in a 32 bit environment except if only the periodic rate is being changed during run-time. If the actual Timer 0 Comparator Value needs to be reinitialized, then the following software solution will always work regardless of the environment:

- Set `TIMER0_VAL_SET_CNF` bit
- Set the lower 32 bits of the Timer0 Comparator Value register
- Set `TIMER0_VAL_SET_CNF` bit
- Set the upper 32 bits of the Timer0 Comparator Value register

Timer 0 is configurable to 32 (default) or 64 bit mode, whereas Timers 1:7 only support 32 bit mode.

WARNING

Software must be careful when programming the comparator registers. If the value written to the register is not sufficiently far in the future, then the counter may pass the value before it reaches the register and the interrupt will be missed. The BIOS should pass a data structure to the operating system to indicate that the operating system should not attempt to program the periodic timer to a rate faster than 5 us.

All of the timers support non-periodic mode.

Refer to *IA-PC HPET Specification* for more details of this mode.

Periodic Mode

When a timer is set up for periodic mode, the software writes a value in the timer's comparator value register. When the main counter value matches the value in the timer's comparator value register, an interrupt can be generated. The hardware will then automatically increase the value in the comparator value register by the last value written to that register.

To make the periodic mode work properly, the main counter is typically written with a value of 0 so that the first interrupt occurs at the right point for the comparator. If the main counter is not set to 0, interrupts may not occur as expected.

During run-time, the value in the timer's comparator value register can be read by software to find out when the next periodic interrupt will be generated (not the rate at which it generates interrupts). Software is expected to remember the last value written to the comparator's value register (the rate at which interrupts are generated).

If software wants to change the periodic rate, it should write a new value to the comparator value register. At the point when the timer's comparator indicates a match, this new value will be added to derive the next matching point.

If the software resets the main counter, the value in the comparator's value register needs to reset as well. This can be done by setting the `TIMERn_VAL_SET_CNF` bit. Again, to avoid race conditions, this should be done with the main counter halted. The following usage model is expected:

1. Software clears the `ENABLE_CNF` bit to prevent any interrupts.
2. Software Clears the main counter by writing a value of 00h to it.
3. Software sets the `TIMER0_VAL_SET_CNF` bit.
4. Software writes the new value in the `TIMER0_COMPARATOR_VAL` register.

Software sets the `ENABLE_CNF` bit to enable interrupts.

NOTE

As the timer period approaches zero, the interrupts associated with the periodic timer may not get completely serviced before the next timer match occurs. Interrupts may get lost and/or system performance may be degraded in this case.

Each timer is NOT required to support the periodic mode of operation. A capabilities bit indicates if the particular timer supports periodic mode. The reason for this is that supporting the periodic mode adds a significant amount of gates.

Only timer 0 will support the periodic mode. This saves a substantial number of gates.

23.2.3.4 Enabling the Timers

The BIOS or operating system PnP code should route the interrupts. This includes the Legacy Rout bit, Interrupt Rout bit (for each timer), and interrupt type (to select the edge or level type for each timer).

The Device Driver code should do the following for an available timer:

1. Set the Overall Enable bit (Offset 10h, bit 0).
2. Set the timer type field (selects one-shot or periodic).
3. Set the interrupt enable.
4. Set the comparator value.

23.2.3.5 Interrupt Levels

Interrupts directed to the internal 8259s are active high. Refer to the **Advanced Programmable Interrupt Controller (APIC) (D31:F0)** for information regarding the polarity programming of the I/O APIC for detecting internal interrupts.

If the interrupts are mapped to the 8259 or I/O APIC and set for level-triggered mode, they can be shared with legacy interrupts. They may be shared although it is unlikely for the operating system to attempt to do this.

If more than one timer is configured to share the same IRQ (using the `TIMERn_INT_ROUT_CNF` fields), then the software must configure the timers to level-triggered mode. Edge-triggered interrupts cannot be shared.

For handling interrupts and issues related to 64 bit timers with 32 bit processors, refer to IA-PC HPET Specification.

24.0 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.

The main characteristics are as follows:


- 8 lanes Gen 4 DMI support
- 4 lanes Gen 4 Reduced DMI support
- 16 GT/s point-to-point DMI interface to PCH
- DC coupling - no capacitors between the processor and the PCH
- PCH end-to-end lane reversal across the link
- L0 (Active) and L1 (Low power) states support
- Half-Swing support (low-power/low-voltage)

24.1 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 22. Example for DMI Lane Reversal Connection



1. DMI Lane Reversal is supported only on PCH-S 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.

24.2 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.

24.3 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.

24.4 Signal Description

Signal Name	Type	Description
DMI_TX_P[0] DMI_TX_N[0]	O	DMI transmit lane 0
DMI_RX_P[0] DMI_RX_N[0]	I	DMI receive lane 0
DMI_TX_P[1] DMI_TX_N[1]	O	DMI transmit lane 1
DMI_RX_P[1] DMI_RX_N[1]	I	DMI receive lane 1
DMI_TX_P[2] DMI_TX_N[2]	O	DMI transmit lane 2
DMI_RX_P[2] DMI_RX_N[2]	I	DMI receive lane 2
DMI_TX_P[3] DMI_TX_N[3]	O	DMI transmit lane 3
DMI_RX_P[3] DMI_RX_N[3]	I	DMI receive lane 3
DMI_TX_P[4] DMI_TX_N[4]	O	DMI transmit lane 4
DMI_RX_P[4] DMI_RX_N[4]	I	DMI receive lane 4
DMI_TX_P[5] DMI_TX_N[5]	O	DMI transmit lane 5
DMI_RX_P[5] DMI_RX_N[5]	I	DMI receive lane 5
DMI_TX_P[6] DMI_TX_N[6]	O	DMI transmit lane 6
DMI_RX_P[6] DMI_RX_N[6]	I	DMI receive lane 6
DMI_TX_P[7] DMI_TX_N[7]	O	DMI transmit lane 7
DMI_RX_P[7] DMI_RX_N[7]	I	DMI receive lane 7
DMI_PERST#	O	DMI PERST Strobe
DMI_RCOMP	Analog	Configuration Resistance Compensation

25.0 Direct Enhanced Serial Peripheral Interface (Direct eSPI)

25.1 Functional Description

The Direct eSPI interface is used in order to support the secure eSPI link between the processor and the PCH.

25.1.1 Processor-PCH eSPI Return Clock Support

In order to improve the eSPI bandwidth between the processor and the PCH, an eSPI return clock is added to increase the eSPI clock frequency up to 100 MHz. The processor Direct eSPI controller samples the return data using the return clock. The Direct eSPI interface supports 66 MHz, 80 MHz, and 100 MHz.

25.1.2 PCH Multiple External eSPI Device Support

For the Direct eSPI interface, there is only one eSPI chip select between the processor and the PCH. The processor Direct eSPI controller initializes the PCH Direct eSPI target and the four channels during the initialization flow. However, the processor Direct eSPI controller supports multiple external eSPI devices connected to the PCH eSPI controller including the peripheral channel, OOB channel, and VW channel. The Flash channel is only supported on CS0 on the PCH eSPI controller.

25.1.3 Processor Direct eSPI Channel Support

Peripheral Channel

PCH eSPI Attached External EC/Device IO/Memory Cycles

The processor routes the eSPI peripheral channel fixed IO, generic IO/memory cycles to the processor Direct eSPI controller through the peripheral channel to the PCH Direct eSPI target. Then the BIOS sends these IO/memory cycles to different external eSPI devices attached to the PCH. The PCH eSPI switch will route them to different eSPI devices based on the IOE/IOD/LGMR register setting on the PCH eSPI controller.

The PCH eSPI switch forwards the EC bus mastering memory cycles and LTR messages between the processor Direct eSPI controller and an EC attached to the PCH.

PCH RTC Port

The processor Direct eSPI controller access and sends the IO cycle to the RTC IO port through the PCH on the peripheral channel and the completion is converted back to the processor.

PCH IP Register Access

The PCH IP register space for SPI/eSPI, GPIO, RTC, and IOxAPIC can be accessed by the processor BIOS using the reserved window of the processor Direct eSPI controller peripheral channel device 0 Generic memory range(LGMR). The BIOS access can be locked using a lock bit in the PCH eSPI private register.

Virtual Wire Channel

The processor PMC, ITSS, and GPIO transmit/receive virtual wires to/from the PCH attached EC. The VWs are forwarded to/from the EC by the PCH eSPI switch.

The GPIO VWs are also supported between the processor GPIO controller and the PCH attached EC in both directions.

A new upstream VW is added to support PCH NMI reporting. The processor Direct eSPI controller forwards this VW to the processor ITSS.

The processor Direct eSPI controller does not broadcast the power management VWs (e.g. PLTSRT/SLP_Sx etc) using the PCH Direct eSPI target. The PCH takes responsibility to broadcast the power management VWs if there are multiple eSPI devices attached to the PCH.

OOB Channel

The processor PMC, Intel® CSME, Intel® Silicon Security Engine, TCSS, DfX tracing, and the RTC uses the Direct eSPI to communicate to the PCH via the OOB channel. The processor PMC also uses the legacy OOB interface with PCH attached EC for other OOB services including PECE over eSPI.

Flash Channel

The processor BIOS, CSME, and ESE use the processor SPI controller for flash access in SAF mode. The BIOS also uses the processor SPI controller for RPMC access in SAF mode. The PCH CSME uses PCH SPI controller for RPMC access.

To improve the flash access/boot time in the desktop platform, the processor BIOS initializes the DMA. During the DMA flow, it sends multiple read requests for the processor and the PCH to the flash device to optimize the performance.

The processor descriptor contains the region definitions and descriptor based permission control. The processor SPI controller accesses its descriptor on PCH Flash or EC via the SAF channel between the processor and the PCH.

The PCH CSME share region 2 and the processor ESE and the PCH ESE share region 10. The other region assignment is listed in the table below.

Table 80. Region Entries in the Descriptor

Register Name	Region Name	Flash Master	Special Notes
FLREG0	Descriptor	Processor/PCH descriptor	Located at the first 12KB of the flash.
FLREG1	BIOS	Processor BIOS region	
FLREG2	CSME	PCH CSME region	
FLREG3	GbE	PCH Gbe region	
<i>continued...</i>			

Register Name	Region Name	Flash Master	Special Notes
FLREG4	Platform Data		
FLREG5	Device Expansion #1	Reserved	
FLREG6	Secondary BIOS	Processor 2nd BIOS	
FLREG7	Reserved	Reserved	
FLREG8	Embedded Controller	PCH EC	
FLREG9	2nd Descriptor Region	Processor/PCH SPI controller use this region for the 2nd descriptor region.	For the usages of the 2nd descriptor. Always located at flash physical address 12KB-24KB.
FLREG10	Processor/PCH ESE	Processor/PCH ESE shared region	
FLREG11	Reserved		
FLREG12	Reserved		
FLREG13	Reserved		
FLREG14	Reserved		
FLREG15	Reserved		

Processor BIOS FRACC (Flash Region Access Permission) Control

The BIOS FRACC and SFRACC (Secondary FRACC) grant permission to ESE (in Processor and the PCH) to read/write access to the BIOS/secondary BIOS region in the flash. To provide BIOS to access the PCH SPI controller's control registers FRACC and SFRACC, the PCH SPI controller's BIOS PCI configuration and MMIO registers are mapped to the processor's Direct eSPI peripheral channel generic memory range.

The processor and the PCH expose the PCH SPI controller to BIOS through DMI so that BIOS can directly program the BIOS FRACC from the PCH SPI controller to grant access to the BIOS in the PCH flash.

Processor BIOS Software Sequencing

The processor only supports SAF mode on the Direct eSPI. The BIOS software sequencing is not supported in SAF mode. To use this feature on the platform, the processor and the PCH exposes the PCH SPI controller to BIOS through DMI so that BIOS can directly program the software sequencing from PCH SPI controller to send software sequencing cycles.

25.2 Signal Description

Signal Name	Type	Description
DIR_ESPI_IO0	I/O	Direct eSPI Signal 0 (Processor): Direct eSPI bi-directional command or data between Processor and PCH.
DIR_ESPI_IO1	I/O	Direct eSPI Signal 1 (Processor): Direct eSPI bi-directional command or data between Processor and PCH.
<i>continued...</i>		



Signal Name	Type	Description
DIR_ESPI_IO2	I/O	Direct eSPI Signal 2 (Processor): Direct eSPI bi-directional command or data between Processor and PCH.
DIR_ESPI_IO3	I/O	Direct eSPI Signal 3 (Processor): Direct eSPI bi-directional command or data between Processor and PCH.
DIR_ESPI_CS0#	O	Direct eSPI Chip Select (Processor): Driving CS# signal low to enable eSPI PCH for the transaction.
DIR_ESPI_CLK	O	Direct eSPI Clock (Processor): eSPI clock output from the Processor to PCH.
DIR_ESPI_RESET#	O	Direct eSPI Reset (Processor): Reset signal from the Processor to PCH.
DIR_ESPI_RCLK	I	Direct eSPI Return Clock (Processor): Direct eSPI Clock from PCH to Processor.

26.0 Testability and Monitoring

This section contains information regarding the testability signals that provides access to JTAG, run control, system control, and observation resources.

Table 81. Acronyms

Acronyms	Description
BSDL	Boundary Scan Description Language
DCI	Direct Connect Interface
DbC	Debug Class Devices
DFP	Downward Facing Port, USB Type-C term
IEEE	Institute of Electrical and Electronics Engineers
I/O	Input/Output
I/OD	Input/Output Open Drain
Intel® TH	Intel® Trace Hub
JTAG	Joint Test Action Group
KMD	Kernel Mode Debug
UFP	Upstream Facing Port, USB Type-C term
2W	2-Wire

Table 82. References

Specification	Document Number/Location
Specification IEEE Standard Test Access Port and Boundary Scan Architecture	http://standards.ieee.org/findstds/standard/1149.1-2013.html

26.1 Signal Description

Table 83. Testability Signals

26.2 I/O Signal Planes and States

Table 84. Power Planes and States for Testability Signals

Signal Name	Power Plane ²	Resistors ^{2 3}	During Reset ¹	Immediately after Reset ¹	S4/S5
Processor JTAG signals					
PROC_JTAG_TCK	VCCPRI_M_IO	Strong Internal Pull-Down	Driven Low	Driven Low	Driven Low
PROC_JTAG_TMS	VCCPRI_M_IO	Internal Pull-Up	Driven High	Driven High	Driven High
<i>continued...</i>					

Signal Name	Power Plane ²	Resistors ^{2 3}	During Reset ¹	Immediately after Reset ¹	S4/S5
PROC_JTAG_TDI	VCCPRI M_IO	Internal Pull-Up	Driven High	Driven High	Driven High
PROC_JTAG_TDO	VCCPRI M_IO	External Pull-Up	Undriven	Undriven	Undriven
PROC_JTAG_TRST #	VCCPRI M_IO	Internal Pull-Down	Driven Low	Driven Low	Driven Low
DBG_PMODE	VCCPRI M_IO	Internal Pull-Up	Driven High	Driven High	Driven High
PRDY#	Primary	External Pull-Up	Driven High	Driven High	Undriven
PREQ#	Primary	External Pull-Up	Driven High	Driven High	Driven High
BPM[3:0]	Primary	External Pull-Up	Undriven	Undriven	Undriven

Notes: 1. Reset reference for primary well pins is RSMRST#.

2. It is strongly recommended to reserve pads for PU\PD resistor in parallel to the internal resistor

27.0 Miscellaneous Signals

27.1 Signal Description

Signal Name	Type	Description
CPU_ID	I	A platform indication signal for compatibility option.
EKEY	N/A	Socket Electronic Key: Used to distinguish between packages with different pins assignment. Connect this pin to the Enable signal of the first VR in sequence. Or as appropriate, to shut down complete power to processor/platform when a wrong package is being used.
SKTOCC#	O	Socket Occupied: Pulled down directly (0 Ohms) on the processor package to the ground. There is no connection to the processor silicon for this signal. System board designers may use this signal to determine if the processor is present.
GPP_SD00/TIME_SYNC0	I	Time Synchronization GPIO 0: Timed GPIO event for time synchronization for interfaces that do not support time synchronization natively

27.2 Ground and Reserved 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_PULLDOWN - These signals should be pull down with resistor.

Arbitrary connection of these signals to VCC, VDD2, VSS, or to any other signal (including each other) may result in component malfunction or incompatibility with future processors. Refer to the table below.

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 (VSS). 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, unless otherwise noted.

Table 85. GND, RSVD, and NCTF Signals

Signal Name	Description
VSS	Ground: Processor ground node
VSS_NCTF	Non-Critical To Function: These signals are for package mechanical reliability and should be connected to VSS on the board.
RSVD	Reserved: All signals that are RSVD should not be connected on the board.
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.
RSVD_PULLDOWN	All signals that are RSVD_PULLDOWN should be connected on the board to the ground with a resistor