



Intel® Server Board S2600BP Intel® Compute Module HNS2600BP Product Family

Technical Product Specification

An overview of product features, functions, architecture, and support specifications

Revision 1.0

JULY 2017

Intel® Server Products and Solutions

<Blank Page>

Document Revision History

Date	Revision	Description of changes
July 2017	1.0	Initial Release

Document Disclaimer Statements

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at Intel.com, or from the OEM or retailer.

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

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

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

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

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

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

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

Copyright © Intel Corporation. All rights reserved.

Table of Contents

1. Introduction	1
1.1 Chapter Outline	1
1.2 Intel Server Board Use Disclaimer	2
1.3 Product Errata	2
2. Product Family Overview	3
2.1 Environmental Limits Specification	11
2.2 Product Weight and Packaging Dimensions	12
2.3 Intel® Server Board S2600BP Family Overview	13
2.3.1 Server Board Feature Identification	14
2.3.2 Server Board Mechanical Dimensional Diagrams	19
2.3.3 Server Board Architecture Overview	24
2.4 Intel® Compute Module HNS2600BP Overview	27
2.4.1 Intel Compute Module Feature Identification	28
2.4.2 Compute Module Mechanical Dimensional Diagrams	29
2.5 Compute Module Components Overview	30
2.5.1 Power Docking Boards	30
2.5.2 Bridge Board Options	32
2.5.3 Compute Module Riser Card Support	34
2.5.4 Compute Module Air Flow Support	36
2.6 System Software Stack	40
2.6.1 Hot Keys Supported During POST	41
2.6.2 Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data	43
3. Processor Support	44
3.1 Processor Socket and Processor Heat Sink Module (PHM) Assembly	44
3.1.1 Bolster Plate Insulator for CPU 1	46
3.2 Processor Thermal Design Power (TDP) Support	47
3.2.1 Processor Heat Sink	47
3.2.2 Processor Features Overview	48
3.3 Processor Population Rules	52
3.4 Processor Initialization Error Summary	53
3.5 Intel® Xeon® processor Scalable Family with Integrated Intel® Omni-Path Host Fabric Interface	55
4. Memory Support	60
4.1 Memory Subsystem Architecture	60
4.2 Supported Memory	61
4.3 Memory Slot Identification and Population Rules	61
4.3.1 DIMM Population Guidelines for Best Performance	63
4.4 Memory RASM Support	64
4.4.1 DIMM Populations Rules and BIOS Setup for Memory RAS	65
5. PCIe* Support	66

5.1	PCIe* Enumeration and Allocation	68
5.2	PCIe* Non-Transparent Bridge (NTB)	68
6.	Server Board I/O	70
6.1	Quick Assist Technology	70
6.2	PCIe* Riser Card Support.....	70
6.2.1	Compute Module – Riser Card Support.....	72
6.3	Onboard Storage Sub-System	72
6.3.1	M.2 SSD Support	72
6.3.2	Embedded Serial ATA (SATA) Support.....	74
6.3.3	Embedded SATA RAID Support.....	76
6.3.4	Intel® Volume Management Device (Intel® VMD) for NVMe	79
6.3.5	Intel® Virtual RAID on CPU (Intel® VROC) for NVMe.....	82
6.4	Video Support	84
6.5	Universal Serial Bus (USB) Ports	85
6.6	Serial Port.....	85
6.7	Network Interface.....	86
6.7.1	Network ports 1 and 2 on the Intel® Server Board S2600BPB and S2600BPQ	86
6.7.2	Network ports 1 and 2 on the Intel® Server Board S2600BPS	87
6.7.3	Server Board Management Networking	87
6.7.4	MAC Address Definition	87
7.	Server Board Connectors and Headers	89
7.1	Power Connectors	89
7.1.1	Main Power Connectors	89
7.1.2	Backup Power Control Connector.....	89
7.2	System Management Headers.....	90
7.2.1	Intel® Remote Management Module 4 (Intel® RMM4) Lite Connector	91
7.2.2	IPMB Header.....	91
7.2.3	Front Control Panel Connector	91
7.2.4	Bridge Board Connector.....	91
7.3	I/O Connectors	93
7.3.1	PCI Express* Slot Connectors	93
7.3.2	VGA Connector.....	101
7.3.3	NIC Connectors	101
7.3.4	SATA Connectors	102
7.3.5	Intel® RAID C600 Upgrade Key Connector	105
7.3.6	Serial Port Connectors.....	105
7.3.7	USB Connectors.....	106
7.3.8	AUX Front Panel	106
7.4	Fan Headers.....	107
7.4.1	Intel Compute Module System Fan Header	107
7.4.2	Non-Intel Chassis Fan Support	107
7.5	Power Docking Board Connectors.....	108

8. Configuration and Recovery Jumpers	109
8.1 BMC Force Update (J6B3).....	110
8.2 ME Force Update (J4B1).....	111
8.3 Password Clear (J4B2)	111
8.4 BIOS Recovery Mode (J4B3).....	112
8.5 BIOS Default (J4B4).....	113
9. Intel® Light-Guided Diagnostics	114
9.1 Status LED	114
9.2 ID LED	117
9.3 BMC Boot/Reset Status LED Indicators	117
9.4 POST Code Diagnostic LEDs	118
10. Basic and Advanced Server Management.....	119
10.1.1 Dedicated Management Port	120
10.1.2 Embedded Web Server.....	120
10.1.3 Advanced Management Feature Support (RMM4 Lite).....	121
11. Intel Compute Module Thermal Management.....	126
12. System Security	128
12.1 Password Setup	128
12.1.1 System Administrator Password Rights	129
12.1.2 Authorized System User Password Rights and Restrictions.....	130
12.2 Front Panel Lockout.....	130
12.3 Trusted Platform Module	130
12.3.1 TPM 2.0.....	130
Appendix A. Integration and Usage Tips.....	132
Appendix B. POST Code Diagnostic LED Decoder.....	133
Appendix C. POST Code Errors	141
Appendix D. Product Family Statements of Volatility	146
Appendix E. Glossary of Terms.....	149
Appendix F. Reference Documents.....	152

List of Figures

Figure 1. Intel® Server Board S2600BP Board	4
Figure 2. Intel® Compute Module HNS2600BP	6
Figure 3. Bridge Board Sample.....	8
Figure 4. Intel® Server Chassis H2204XXLRE	10
Figure 5. Intel® Server Chassis H2312XXLR3	10
Figure 6. Intel® Server Chassis H2224XXLR3	10
Figure 7. Intel® Server Chassis H2000G - Rear View	10
Figure 8. Server Board Feature Identification (S2600BPB, S2600BPQ models)	14
Figure 9. Server Board and Feature Identification (S2600BPS model)	15
Figure 10. Server Board Rear Connectors, (S2600BPB, S2600BPQ).....	16
Figure 11. Server Board Rear Connectors (S2600BPS).....	16
Figure 12. Jumper Block Identification.....	17
Figure 13. Intel® Light Guided Diagnostics LEDs Identification	18
Figure 14. Intel® Server Board S2600BP Product Family Board Dimensions.....	19
Figure 15. Locations of the through-holes on the Intel® Server Board S2600BPB, S2600BPQ.....	19
Figure 16. Locations of the through holes on the Intel® Server Board S2600BPS.....	20
Figure 17. Intel® Server Board S2600BPB, S2600BPQ – Primary Side Keepout Zone	20
Figure 18. Intel® Server Board S2600BPB, S2600BPQ – Secondary Side Keepout Zone	21
Figure 19. Intel® Server Board S2600BPS – Primary Side Keepout Zone	21
Figure 20. Intel® Server Board S2600BPS – Secondary Side Keepout Zone	22
Figure 21. Intel® Server Board S2600BPB, S2600BPQ – Primary Side Height Restrictions	22
Figure 22. Intel® Server Board S2600BPS – Primary Side Height Restrictions	23
Figure 23. Intel® Server Board S2600BPB Architectural Block Diagram.....	24
Figure 24. Intel® Server Board S2600BPQ Architectural Block Diagram.....	25
Figure 25. Intel® Server Board S2600BPS Architectural Block Diagram.....	26
Figure 26. Intel® Compute Module HNS2600BP.....	27
Figure 27. Intel Server Chassis H2000G	27
Figure 28. Intel® Compute Module HNS2600BPB and HNS2600BPQ – Rear Connectors.....	28
Figure 29. Intel® Compute Module HNS2600BPS – Rear Connectors	28
Figure 30. Intel® Compute Module HNS2600BPx – Dimension.....	29
Figure 31. Intel® Compute Module HNS2600BPx24 – Dimension	29
Figure 32. Standard Power Docking Board	30
Figure 33. SAS/PCIe* SFF Combo Power Docking Board Top View	31
Figure 34. 6G SATA Bridge Board Overview.....	33
Figure 35. 12G SAS Bridge Board Overview	33
Figure 36. SAS/PCIe* SFF Combo Bridge Board Overview.....	33
Figure 37. SAS/PCIe* SFF Combo Bridge Board Overview.....	34
Figure 38. Pass Through Bridge Board Overview.....	34
Figure 39. Riser Card for Riser Slot 1	35
Figure 40. Riser Card for Riser Slot 2	35

Figure 41. Compute Module System Fans and Air Duct.....	36
Figure 42. Compute Module Air Flow	36
Figure 43. Compute Module Fan Placement	37
Figure 44. Intel Compute Module Air Duct Placement.....	38
Figure 45. Air Baffle Addition	39
Figure 46. PHM Components and Processor Socket Reference Diagram.....	44
Figure 47. PHM to CPU Socket Orientation and Alignment Features	45
Figure 48. Processor Socket Assembly and Protective Cover	46
Figure 49. Bolster Insulator Plate for CPU1	46
Figure 50. Processor Heat Sinks placement overview.....	47
Figure 51. Processor Host Fabric Interface	55
Figure 52. Intel® Omni-Path Fabric Interconnect.....	56
Figure 53. Fabric Interconnect Configurations	56
Figure 54. Intel Accessory Kit AHWBPFABKIT	57
Figure 55. Server Board Fabric Side Band Connectors	57
Figure 56. IFT Carrier Board – Rear View.....	58
Figure 57. Integrated Memory Controller Functional Block Diagram	60
Figure 58. Intel® Server Board S2600BP Product Family DIMM Slot Layout.....	62
Figure 59. Server Board Riser Slots.....	70
Figure 60. Intel® Server Board S2600BPB & BPS – PCIe* Slot Connector Architectural Block Diagram.....	71
Figure 61. Intel® Server Board S2600BPQ – PCIe* Slot Connector Architectural Block Diagram.....	71
Figure 62. On-board M.2 SSD Support.....	72
Figure 63. Riser Card #2 M.2 SSD Support.....	73
Figure 64. SATA RAID 5 Upgrade Key.....	78
Figure 65. Intel® VMD.....	79
Figure 66. Intel® VMD Interface	80
Figure 67. VMD Support Disabled in <F2> BIOS Setup	81
Figure 68. VMD Support Enabled in <F2> BIOS Setup	81
Figure 69. Intel® VROC Upgrade Key.....	82
Figure 70. USB Ports Block Diagram.....	85
Figure 71. Serial Port A Location	85
Figure 72. Network Interface Connectors S2600BPB, S2600BPQ (Port 1 and Port 2 – RJ45).....	86
Figure 73. RJ45 Connector LEDs.....	86
Figure 74. Network Interface Connectors S2600BPS (Port 1 and Port 2 – SFP+)	87
Figure 75. Server Board Riser Slot Identification.....	93
Figure 76. Mini-SAS HD (SFF-8643) pin-out.....	102
Figure 77. M.2/NGFF Connector	103
Figure 78. M.2/NGFF Mated Module and Connector	103
Figure 79. SW RAID Activation Key (SAK) and Connector.....	105
Figure 80. Server Board System Fan Headers.....	107
Figure 81. Configuration and Recovery Jumpers	109
Figure 82. Status LED (A), ID LED (B), Diagnostic LEDs (C)	114

Figure 83. Rear Panel Diagnostic LEDs 118

Figure 84. Optional RMM4 Lite..... 122

Figure 85. Air Flow and Fan Identification 126

Figure 86. Security Tab of BIOS Setup Utility..... 128

Figure 87. POST Diagnostic LED Location 133

Figure 88. Correct Order Reading POST Diagnostic Codes..... 134

Figure 89. POST Diagnostic Error (LED Location)..... 145

List of Tables

Table 1. Intel® Server Board S2600BP Product Family Codes and Intel® Compute Module HNS2600BP.....	3
Table 2. Server Board Product Family Feature Set.....	5
Table 3. Intel® Compute Module HNS2600BP Product Family Feature Set.....	7
Table 4. Supported Bridge Board Options.....	8
Table 5. 2U Intel® Server Chassis H2000G Product Family Feature Set.....	9
Table 6. Server Board Environmental Limits.....	11
Table 7. Product weight and package dimension information.....	12
Table 8. Server Board Mechanical Specifications.....	13
Table 9. Compute Module Specifications.....	27
Table 10. Computer Module and Bridge board support matrix.....	32
Table 11. POST Hot-Keys.....	41
Table 12. Processor Heat Sinks.....	47
Table 13. Intel® Xeon® Processor Scalable Family - Feature Comparison Table.....	48
Table 14. Mixed Processor Configurations Error Summary.....	54
Table 15. Intel® Xeon® Processor Scalable Family w/ Integrated Intel Omni-Path Fabric – Features Table.....	55
Table 16. IFT Carrier LED Functionality.....	58
Table 17. Power Level Classification for QSFP+ Modules.....	58
Table 18. Fabric Carrier Thermal.....	58
Table 19. Supported Processor Mixing – Fabric vs Non-Fabric Processors.....	59
Table 20. DDR4 DIMM Support Guidelines.....	61
Table 21. DIMM Nomenclature.....	62
Table 22. Memory RASM Features.....	64
Table 24. CPU – PCIe* Port Routing.....	66
Table 25. Chipset PCH PCIe* Connectivity.....	66
Table 26. High-Speed I/O Port Mapping (HSIO).....	67
Table 27. SATA and sSATA Controller Feature Support.....	74
Table 28. SATA and sSATA Controller BIOS Utility Setup Options.....	75
Table 29. Intel® VROC Upgrade Key Option.....	83
Table 30. Supported Intel VROC Activation Key.....	83
Table 31. Onboard Video Resolution and Refresh Rate (Hz).....	84
Table 32. External RJ45 NIC Port LED Definition.....	86
Table 33. Main Power Supply Connector (6-pin 2x3 Connector).....	89
Table 34. Backup Power Control Connector.....	90
Table 35. Intel® RMM4 Lite Connector Pin-out.....	91
Table 36. IPMB Header 4-pin (J6B1).....	91
Table 37. Control Panel Connector.....	91
Table 38. Bridge Board Connector.....	91
Table 39. PCIe* Slot Connector 1 (PCIe_X16_CPU_1).....	94
Table 39. PCIe* Slot Connector 2 (PCIe_X24_CPU_1).....	96

Table 41. PCIe* Slot Connector 3 (PCIe_X24_CPU_2).....	97
Table 42. PCIe* x16 Riser Slot 4 Connector.....	99
Table 43. VGA External Video Connector.....	101
Table 44. RJ-45 100Mb/1Gb/10Gb NIC Connector Pin-out.....	101
Table 45. RJ-45 10/100/1000 Dedicated Management Port NIC Connector Pin-out	101
Table 46. Mini-SAS HD (SFF-8643) SGPIO (SATA sideband signals).....	102
Table 47. M.2/NGFF Socket 3 Connector Pinout (Mechanical Key M).....	104
Table 48. Summary of RAID Keys.....	105
Table 49. Storage Upgrade Key Connector	105
Table 50. Internal 10-pin Serial A	105
Table 51. External USB port Connector.....	106
Table 52. Internal USB Connector	106
Table 53. AUX Front Panel Connector	106
Table 54. Server board Fan Connector.....	107
Table 55. SYS_FAN_# (1-3) Connector Pinout	108
Table 56 . SYS_FAN_6 Connector Pinout	108
Table 57. SYS_FAN_7 Connector Pinout	108
Table 58. Main Power Input Connector	108
Table 59. Force Integrated BMC Update Jumper (J6B3).....	110
Table 60. Force ME Update Jumper (J4B1)	111
Table 61. Password Clear Jumper (J4B2).....	111
Table 62. BIOS Recovery Mode Jumper (J4B3)	112
Table 63. BIOS Default Jumper.....	113
Table 64. Status LED State Definitions	115
Table 65. ID LED	117
Table 66. BMC Boot/Reset Status LED Indicators	118
Table 67. Intel® Remote Management Module 4 (RMM4) Options	119
Table 68. Basic and Advanced Server Management Features Overview.....	119
Table 69. Air Flow	126
Table 70. POST Progress Code Decoding LED Example.....	134
Table 71. MRC Progress Codes	135
Table 72. MRC Fatal Error Codes.....	136
Table 73. POST Progress Codes.....	137
Table 74. POST Error Codes and Messages.....	141
Table 75. POST Error LED Codes	144
Table 76. Integrated BMC Error (LED) Codes	144
Table 77. Server Board Component.....	147
Table 78. System Boards Components	147
Table 79. Server Chassis Components	148

1. Introduction

This Technical Product Specification (TPS) provides a high level overview of the features, functions, and architecture of the Intel® Server Board S2600BP product family and the Intel® Compute Module HNS2600BP product family, which includes: the chassis layout, system boards, power subsystem, cooling subsystem, storage subsystem options, and available installable options. Note that some features are provided as configurable options and may not be included as a standard. Please reference to *the Intel® Server Board S2600BP Product Family Configuration Guide* for a list of configurable options.

Throughout this Document:

- The Intel® Server Board S2600BPS, S2600BPB and S2600BPQ will be collectively referred to as the Intel® Server Board S2600BP.
- The Intel® Compute Module HNS2600BPS, HNS2600BPB and HNS2600BPQ will be collectively referred to as the Intel® Compute Module HNS2600BP.
- The Intel® Compute Module HNS2600BPS24, HNS2600BPB24 and HNS2600BPQ will be collectively referred to as the Intel® Compute Module HNS2600BP24.

In addition, design-level information related to specific server board components/subsystems can be obtained by ordering External Product Specifications (EPS) or External Design Specifications (EDS) related to this server generation. EPS and EDS documents are made available under NDA with Intel and must be ordered through your local Intel representative. See the Reference Documents section at the end of this document for a list of available documents.

1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Product Family Overview
- Chapter 3 – Processor Support
- Chapter 4 – Memory Support
- Chapter 5 – PCIe* Support
- Chapter 6 – Server Board I/O
- Chapter 7 – Connectors and Headers
- Chapter 8 – Configuration and Recovery Jumpers
- Chapter 9 – Intel® Light-Guided Diagnostics
- Chapter 10 – Basic and Advanced Server Management
- Chapter 11 – Thermal Management
- Chapter 12 – System Security
- Appendix A – Integration and Usage Tips
- Appendix B – POST Code Diagnostic LED Decoder
- Appendix C – POST Code Errors
- Appendix D – Product Family Statements of Volatility
- Appendix E – Glossary of Terms
- Appendix F – Reference Documents

1.2 Intel Server Board Use Disclaimer

Intel Corporation server boards support add-in peripherals and contain a number of high-density Very Large Scale Integration (VLSI) and power delivery components that need adequate airflow to cool. Intel ensures through its own chassis development and testing that when Intel server building blocks are used together, the fully integrated system will meet the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel-developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of their published operating or non-operating limits.

1.3 Product Errata

Shipping product may have features or functionality that may deviate from published specifications. These deviations are generally discovered after the product has gone into formal production. Intel terms these deviations as product Errata. Known product Errata will be published in the Monthly Specification Update for the given product family which can be downloaded from the following Intel web site:

<https://www.intel.com/content/www/us/en/support/server-products.html>

2. Product Family Overview

The Density-optimized Intel® Server S2600BP product family offers a variety of building block options to meet the varied configuration requirements of high-density high-performance computing environments. Building block options include board only server board SKUs which provide OEMs and other system integrators the option to develop a custom enclosure for unique server environments, or the server boards are offered integrated onto 1U compute modules, designed and tested to support specific SKUs within the Intel® Server Chassis H2000G product family.

Available Intel® server boards and the Intel® compute modules that make up the product family are listed in the following table.

Table 1. Intel® Server Board S2600BP Product Family Codes and Intel® Compute Module HNS2600BP

Intel Product Code	Product Code Feature Differentiator
BBS2600BPB	Intel® Server Board S2600BPB, dual 10GBaseT Support for Intel Xeon processor Scalable Family and 16 DIMMs
BBS2600BPS	Intel® Server Board S2600BPS, dual 10 GbE SFP+ Support for Intel Xeon processor Scalable Family and 16 DIMMs
BBS2600BPQ	Intel® Server Board S2600BPQ, dual 10GBaseT, with Intel® QuickAssist Technology (QAT) support Support for Intel Xeon processor Scalable Family and 16 DIMMs
HNS2600BPB	Module integrated with an Intel® Server Board S2600BPB. Compatible with Intel® Server Chassis H2312XXLR3 & H2204XXLRE
HNS2600BPS	Module integrated with an Intel® Server Board S2600BPS. Compatible with Intel® Server Chassis H2312XXLR3 & H2204XXLRE
HNS2600BPQ	Module integrated with an Intel® Server Board S2600BPQ. Compatible with Intel® Server Chassis H2312XXLR3 & H2204XXLRE
HNS2600BPB24	Module integrated with an Intel® Server Board S2600BPB. Compatible with Intel® Server Chassis H2224XXLR3
HNS2600BPS24	Module integrated with an Intel® Server Board S2600BPS. Compatible with Intel® Server Chassis H2224XXLR3
HNS2600BPQ24	Module integrated with an Intel® Server Board S2600BPQ. Compatible with Intel® Server Chassis H2224XXLR3

The Intel® Server Board S2600BP product family is a purpose built, rack-optimized server board ideal for use in hyper-converged, data analytics, storage, cloud and high performance computing applications.




Figure 1. Intel® Server Board S2600BP Board

The architecture of the server board is developed around the features and functions of the Intel® Xeon® processor Scalable family. Previous generation Intel® Xeon® processors are not supported. This server board product family includes three board options, each supporting a unique feature set as identified in Table 2.

Table 2. Server Board Product Family Feature Set

Server Board Features	Detail		
Server Board SKU	S2600BPB	S2600BPS	S2600BPQ
Processor	<ul style="list-style-type: none"> Up to two (2) Intel® Xeon® processor Scalable Family Support for Intel® Xeon® processor Scalable Family with Intel® Omni-path Integrated Fabric Connectors – One 100Gb/s port per processor 		
Maximum Processor TDP	Board & Compute Module: 165W (Compute Module in Intel Chassis: 140W / 165W**) ** Note: Supported in enhanced air cooled Intel® Server Chassis H2204XXLRE only		
Processor Socket	Dual Socket-P 3647		
Chipset	Intel® C621 chipset	Intel® C622 chipset	Intel® C628 chipset
Intel® QuickAssist Technology (QAT)	No	No	Yes
Memory	16 DIMM Slots (8 per CPU) – DDR4 RDIMM/LRDIMM, Up to 266** MT/s, 1.2V ** The maximum memory speed supported is dependent on the installed processor SKU		
SATA Support	Four (4) SATA 6Gbps ports via Mini-SAS HD (SFF-8643) connector (from PCH)		No
M.2	One (1) 42mm M.2 SATA/PCIe* x4		
Networking	Dual 10GBase-T ports	Dual 10GbE SFP+ ports support	Dual 10GBase-T ports
USB	Dual stacked rear panel USB 3.0 ports		
Expansion Options			
Bridge Board Slot	See supported Bridge Board Options under Compute Module Specs		
Riser Slot 1 (CPU1)	x16 PCIe* 3.0 lanes	x16 PCIe* 3.0 lanes	No PCIe*, Power Only
Riser Slot 2 (CPU1)	x24 PCIe* 3.0 lanes	x24 PCIe* 3.0 lanes	x24 PCIe* 3.0 lanes
Riser Slot 3 (CPU2)	x24 PCIe* 3.0 lanes	x24 PCIe* 3.0 lanes	x24 PCIe* 3.0 lanes
Riser Slot 4 (CPU2)	x16 PCIe* 3.0 lanes	x16 PCIe* 3.0 lanes	x16 PCIe* 3.0 lanes
Server Management	<ul style="list-style-type: none"> Integrated Baseboard Management Controller, IPMI 2.0 compliant One Dedicated RJ45 Management Port One 2x4 pin header for optional Intel® Remote Management Module 4 Lite (Intel® RMM4 Lite2) with support for Remote KVM On-Board LEDs: System Status, System ID, POST Code Diagnostic, BMC Error Code 		
BIOS	uEFI based BIOS		
Onboard Connectors and Headers	Detail		
Video	One 1x12 pin video header, 16MB of DDR4 video memory		
IPMB	One 1x4 pin IPMB header		
USB	One USB 2.0 connector		
Serial	One DH-10 Serial Port A connector		
Control Panel	One 1x12 pin control panel header		
System Fan Support	One managed 2x7 pin header for system fan module (Intel Chassis) Three managed 1x8 pin System Fan connectors Two 1x4 pin System Fan connectors		
Power	One 1x8 pin Backup Power connector Two 2x3 Main Power connectors		

Note: All riser slots on the server board are designed for riser card support ONLY. Plugging a PCIe* card directly into a riser slot may cause permanent server board and/or PCIe* card damage.

Server boards from the Intel® Server Board S2600BP product family are offered as fully integrated density optimized compute modules known as the Intel® Compute Module HNS2600BP product family.

The Intel® Compute Module HNS2600BP product family is a purpose build 1U density optimized compute module designed to operate as a single system node within a multi-node chassis.




Figure 2. Intel® Compute Module HNS2600BP

The Intel® Compute Module HNS2600BP product family includes six compute module options. Table 3 identifies the feature set associated with each.

Table 3. Intel® Compute Module HNS2600BP Product Family Feature Set

Compute Module Features	Detail		
Compute Module SKUs	HNS2600BPB HNS2600BPB24	HNS2600BPS HNS2600BPS24	HNS2600BPQ HNS2600BPQ24
Processor Support	<ul style="list-style-type: none">Up to two (2) Intel® Xeon® processor Scalable FamilySupport for Intel® Xeon® processor Scalable Family with Intel® Omni-path Integrated Fabric Connectors – One 100Gb/s port per processor		
Maximum Processor TDP	Board & Compute Module: 165W Compute Module installed in Intel Chassis: 140W / 165W** ** Note: Supported in enhanced air cooled Intel® Server Chassis H2204XXLRE only		
Chipset	Intel® C621 chipset	Intel® C622 chipset	Intel® C628 chipset
Memory	16 DIMM Slots (8 per CPU) – DDR4 RDIMM/LRDIMM, Up to 266** MT/s, 1.2V ** The maximum memory speed supported is dependent on the installed processor SKU		
Storage Support	Dependent on optional Bridge Board accessory. See supported Bridge Board options		
M.2	One (1) on-board 42mm M.2 SATA/PCIe* x4 + One (1) Riser 2 80mm M.2 PCIe* x4		
Networking	Dual 10GBase-T ports	Dual 10GbE SFP+ ports support	Dual 10GBase-T ports
Intel® Quick Assist Technology (QAT)	No	No	Yes
USB	Dual stacked rear panel USB 3.0 ports		
Video	DB-15 Rear Panel VGA Video connector on Riser 2, 16MB of DDR4 video memory		
Expansion Options			
Bridge Board Slot	See supported Bridge Board Options		
Riser Slot 1 (CPU1)	Default: Video Adapter 1U PCIe* x16 riser card for supporting in-board low-profile PCIe card Optional: Support for Intel® Omni-path fabric through carrier card	Default: Video Adapter Optional: Support for Intel® Omni-path fabric through carrier card option only	Default: Video Adapter Optional: Support for Intel® Omni-path fabric through carrier card option only
Riser Slot 2 (CPU1)	1U PCIe* x16 riser card for supporting in-board low-profile PCIe card (on slot 2) and one PCIe 80mm M.2 device. Optional: Support for Intel® Omni-path fabric through carrier card	1U PCIe* x16 riser card for supporting in-board low-profile PCIe card (on slot 2) and one PCIe 80mm M.2 device. Optional: Support for Intel® Omni-path fabric through carrier card	1U PCIe* x16 riser card for supporting in-board low-profile PCIe card (on slot 2) and one PCIe 80mm M.2 device. Optional: Support for Intel® Omni-path fabric through carrier card
Riser Slot 3	Not Available with Bridge Board installed	Not Available with Bridge Board installed	Not Available with Bridge Board installed
Riser Slot 4	Not Available with Bridge Board installed	Not Available with Bridge Board installed	Not Available with Bridge Board installed
Server Management	Integrated Baseboard Management Controller, IPMI 2.0 compliant One Dedicated RJ45 Management Port One 2x4 pin header for optional Intel® Remote Management Module 4 Lite (Intel® RMM4 Lite2) with support for Remote KVM On-Board LEDs: System Status, System ID, POST Code Diagnostic, BMC Error Code		
Fans	Three 40x56mm dual rotor system fans		

The following table identifies all available Bridge Board options supported by the Intel Compute Module HNS2600BP product family.

Note: Beyond the feature set associated with each bridge board option, bridge board selection must also be based on the specific Intel compute module and Intel chassis into which it will be installed.

Table 4. Supported Bridge Board Options

Bridge Board Product Code	Description	SATA / SAS	RAID 0, 1, 10	RAID 5	Compatible Intel Products
AHWBPBGB	4-Port Bridge Board	6G SATA PCH	ESRT2 SW RAID 0/1/10	RAID 5 with optional key	Supported Intel Compute Modules: HNS2600BPB HNS2600BPS HNS2600BPQ Supported Intel Chassis: H2312XXLR3 H2204XXLRE
AHWBP12GBGB	4-Port iMR Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS	RAID 0/1/10	No	
AHWBP12GBGBR5	4-Port iMR Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS	RAID 0/1/10	RAID 5	
AHWBPBGB24	6-Port IT Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS NVMe	No RAID	No RAID	Compute Module: HNS2600BPB24 HNS2600BPS24 HNS2600BPQ24 Supported Intel Chassis: H2224XXLR3
AHWBPBGB24R	6-Port IT iMR Bridge Board w/ LSI 3408 IOC	6G SATA 12G SAS NVMe	RAID 0/1/10	RAID 5 with optional key (AXXRPFKHY5)	
AHWBPBGB24P	6-Port Pass-Through Bridge Board	Pass-Through	RAID 0, 1, 10 with optional add-in card	RAID 5 with optional add-in card	




Figure 3. Bridge Board Sample

The Intel® Compute Module HNS2600BP product family is compatible with the following Intel server chassis.

- Intel® Server Chassis H2000G Product Family. See table 5 for supported chassis SKUs

Note: The Intel® Server Chassis H2000G product family consists of several chassis SKUs compatible with different Intel server compute module families. Only the chassis SKUs identified in the following table are compatible with the Intel® Compute Module HNS2600BP product family.

Table 5. 2U Intel® Server Chassis H2000G Product Family Feature Set

Chassis Feature / Chassis SKUs	H2312XXLR3	H2224XXLR3	H2204XXLRE
Intel® Compute Module support	HNS2600BPB HNS2600BPS HNS2600BPQ	HNS2600BPB24 HNS2600BPS24 HNS2600BPQ24	HNS2600BPB HNS2600BPS HNS2600BPQ
Number of Compute Modules	Up to 4	Up to 4	Up to 4
Chassis Dimensions	3.46" x 17.24" x 30.35" 87.9 x 438 x 771mm	3.46 x 17.24" x 28.86" 87.9 x 438 x 733mm 86.9 x 438 x 733mm	3.46" x 17.24" x 30.35" 87.9 x 438 x 771mm
Package Dimensions	983 x 577 x 260mm 983 x 577 x 260mm 983 x 577 x 260mm		
Chassis Weight	21.5kg	20.6kg	20.6kg
Package Weight	29.5kg	28.9kg	28.9kg
Maximum Supported Processor TDP	145 Watts	145 Watts	165 Watts
Power Supply	Two (2) x 2130W AC Common Redundant Power Supply (CRPS)		
Power Supply Efficiency Rating	80 Plus Platinum		
Power Configuration	1+0 – No power redundancy 1+1 – Redundant power 2+0 – Combined power, no redundancy		
Fans	Three (3) system fans per module One (1) fan per power supply		
2.5" PCIe* NVMe SSD Support	None	8 total, max 2 per module	None
Hot Swap Drive Bays	<ul style="list-style-type: none"> • Twelve (12) x 3.5" bays • 6Gbps SATA / 12Gbps SAS • With optional Bridge Board 	<ul style="list-style-type: none"> • Twenty-four (24) x 2.5" bays • 6Gbps SATA / 12Gbps SAS / PCIe* NVMe • With optional Bridge Board 	<ul style="list-style-type: none"> • Four (4) x 2.5" bays • 6Gbps SATA / 12Gbps SAS • With optional Bridge Board

For additional details on chassis features, refer to the *Intel® Server Chassis H2000G Product Family Technical Product Specification*.




Figure 4. Intel® Server Chassis H2204XXLR02



Figure 5. Intel® Server Chassis H2312XXLR3001



Figure 6. Intel® Server Chassis H2224XXLR3001




Figure 7. Intel® Server Chassis H2000G - Rear View

2.1 Environmental Limits Specification

Operation of the server board at conditions beyond those identified in the following table may cause permanent damage to the system. Exposure to absolute maximum rating conditions for extended periods may affect long term system reliability.

NOTE: The Energy Star compliance is available at system level only. Use of Intel® Server Boards as stand alone does not guarantee Energy Star compliance.

Table 6. Server Board Environmental Limits

Parameter	Support Limits
Operating Temperature	+10°C to +35°C with the maximum rate of change not to exceed 10°C per hour
Non-Operating Temperature	-40°C to +70°C
Non-Operating Humidity	90%, non-condensing at 35°C
Acoustic noise	Sound power: 7.0BA with hard disk drive stress only at room ambient temperature (23 +/- 2°C)
Shock, operating	Half sine, 2g peak, 11 mSec
Shock, unpackaged	Trapezoidal, 25g, velocity change 205 inches/second (80 lbs to < 100 lbs)
Vibration, unpackaged	5 Hz to 500 Hz, 2.20 g RMS random
Shock and vibration, packaged	ISTA (International Safe Transit Association) Test Procedure 3A
ESD	+/-12 KV except I/O port +/- 8 KV per Intel® Environmental Test Specification
System Cooling Requirement in BTU/Hr	2130 Watt Max – 7272 BTU/hour

Disclaimer Notes:

- Through its own chassis development and system testing, Intel ensures that the server board meets the specified unpackaged shock and vibration limits identified in Environmental Limits table. It is the responsibility of the system integrator who chooses to use an Intel server board in a non-Intel chassis to perform the necessary validation to ensure specified environmental limits are supported.
- Intel server boards contain a number of high-density VLSI and power delivery components that require adequate airflow to cool. Intel ensures through its own chassis development and testing that when Intel server building blocks are used together, the fully-integrated system will meet the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel-developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of their published operating or non-operating limits.

2.2 Product Weight and Packaging Dimensions

The following table provides information related to product weight and package dimensions

Table 7. Product weight and package dimension information


Product Code	Quantity per Box	Box Dimension (mm)	Net Weight (kg)	Package Weight (kg)
BBS2600BPB	10 in 1	610 X 535 X 270 mm	1.25	16.75
BBS2600BPQ	10 in 1	610 X 535 X 270 mm	1.25	16.75
BBS2600BPS	10 in 1	610 X 535 X 270 mm	1.25	16.75
HNS2600BPB	1 in 1	720 X 272 X 160 mm	4.32	5.6
HNS2600BPQ	1 in 1	720 X 272 X 160 mm	4.32	5.6
HNS2600BPS	1 in 1	720 X 272 X 160 mm	4.32	5.6
HNS2600BPB24	1 in 1	720 X 272 X 160 mm	4.32	5.6
HNS2600BPQ24	1 in 1	720 X 272 X 160 mm	4.32	5.6
HNS2600BPS24	1 in 1	720 X 272 X 160 mm	4.32	5.6

2.3 Intel® Server Board S2600BP Family Overview

The Intel® Server Board S2600BP product family is a purpose built, rack-optimized server board ideal for use in hyper-converged, data analytics, storage, cloud and high performance computing applications. The three server board options share a common form factor and support a common base feature set. However, each board includes features and/or functions that are unique to it. See Table 2 for a complete feature set associated with each board option.

Table 8. Server Board Mechanical Specifications


Server Board Specifications	Detail
Board Dimensions	6.8" x 19.1"
Board Weight	16.75 Kg, 10 boards per box (1.25 Kg per board)
Packaging Dimensions	610 x 535 x 270 mm, 10 boards per box
Packaging Weight	20.7 Kg, 21.4 Kg



BNP094

2.3.1 Server Board Feature Identification

The following figures identify all board features associated with each board option. Features that are unique to a specific board option will be identified.



BNP039

Figure 8. Server Board Feature Identification (S2600BPB, S2600BPQ models)






Figure 9. Server Board and Feature Identification (S2600BPS model)



Label	Description	Label	Description
A	Dedicated Management Port (RJ45)	E	Status LED
B	Dual-port USB 3.0	F	ID LED
C	NIC port 1 (RJ45)	G	POST Code LEDs (8 LEDs)
D	NIC port 2 (RJ45)		


Figure 10. Server Board Rear Connectors, (S2600BPB, S2600BPQ)



Label	Description	Label	Description
A	Dedicated Management Port (RJ45)	E	Status LED
B	Dual-port USB 3.0	F	ID LED
C	NIC port 1 (SFP+)	G	POST Code LEDs (8 LEDs)
D	NIC port 2 (SFP+)		

Figure 11. Server Board Rear Connectors (S2600BPS)

The server board includes several jumper blocks which can be used to configure, protect, or recover specific features of the server board. Figure 12 identifies the location of each jumper block on the server board. Pin 1 of each jumper block can be identified by the arrowhead (▼) silkscreened on the server board next to the pin.



BNP044

Figure 12. Jumper Block Identification

For additional details, see Chapter 8, Configuration and recovery jumpers.

The server board includes several LEDs to identify system status. Figure 13 shows the supported LEDs and identifies their locations. LED locations and support are common to all boards in the Intel® Server Board S2600BP product family. For detailed information, see Chapter 9, Intel® Light-Guided Diagnostics.




Figure 13. Intel® Light Guided Diagnostics LEDs Identification

NOTE: See Appendix B for POST Code Diagnostic LED decoder information.

2.3.2 Server Board Mechanical Dimensional Diagrams

The following figures provide the board and module dimensional data and identify the on-board placement and keep out zones of the server board.




Figure 14. Intel® Server Board S2600BP Product Family Board Dimensions




Figure 15. Locations of the through-holes on the Intel® Server Board S2600BPB, S2600BPQ




Figure 16. Locations of the through holes on the Intel® Server Board S2600BPS




Figure 17. Intel® Server Board S2600BPB, S2600BPQ – Primary Side Keepout Zone




Figure 18. Intel® Server Board S2600BPB, S2600BPQ – Secondary Side Keepout Zone




Figure 19. Intel® Server Board S2600BPS – Primary Side Keepout Zone




Figure 20. Intel® Server Board S2600BPS – Secondary Side Keepout Zone

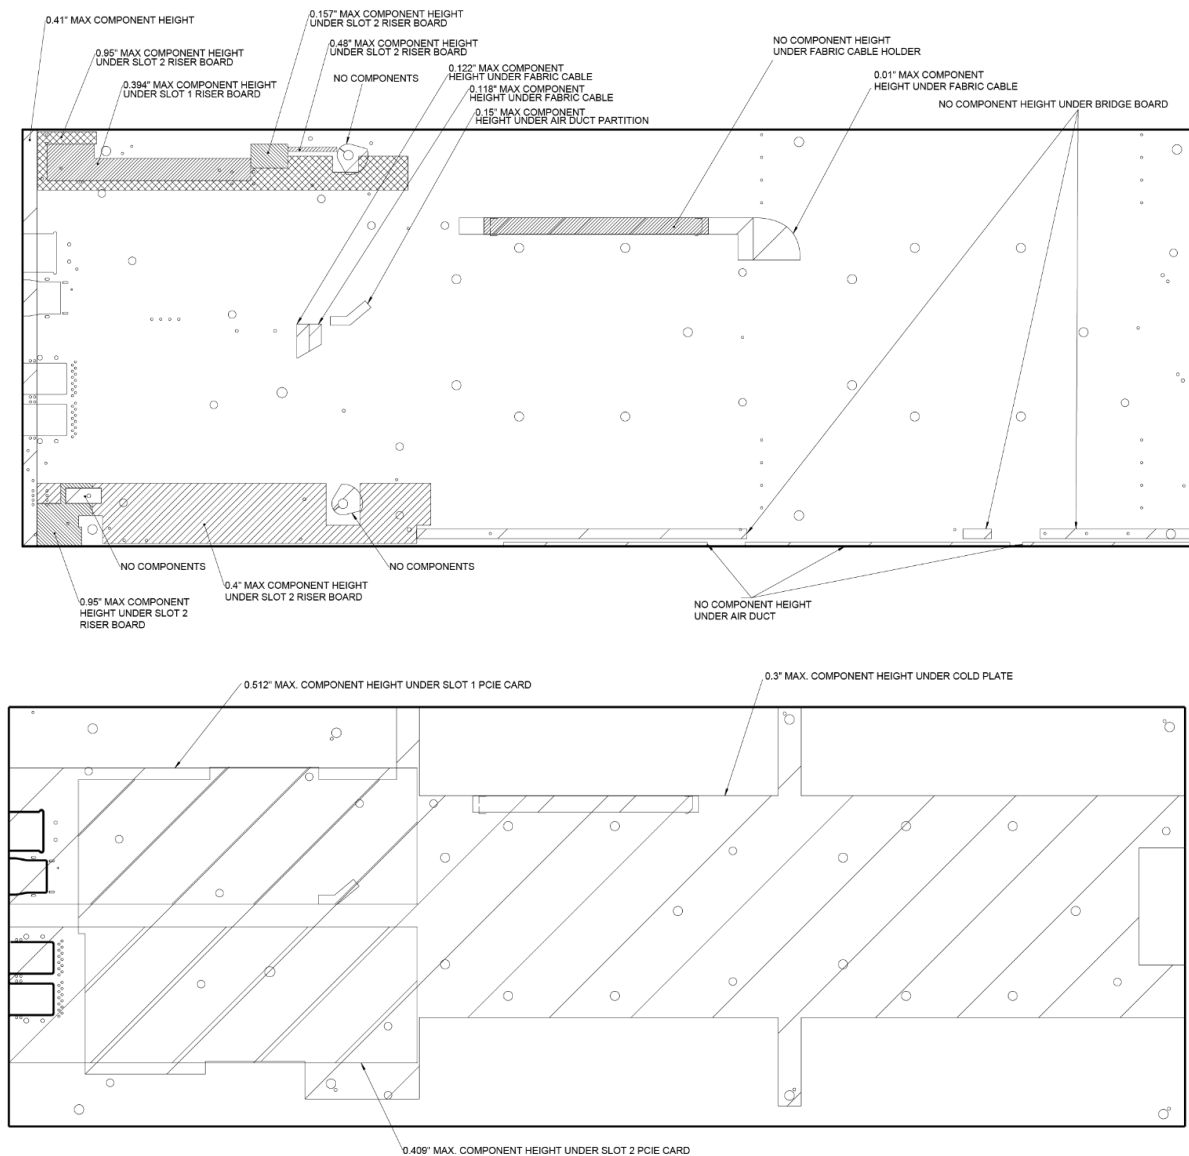


Figure 21. Intel® Server Board S2600BPB, S2600BPQ – Primary Side Height Restrictions

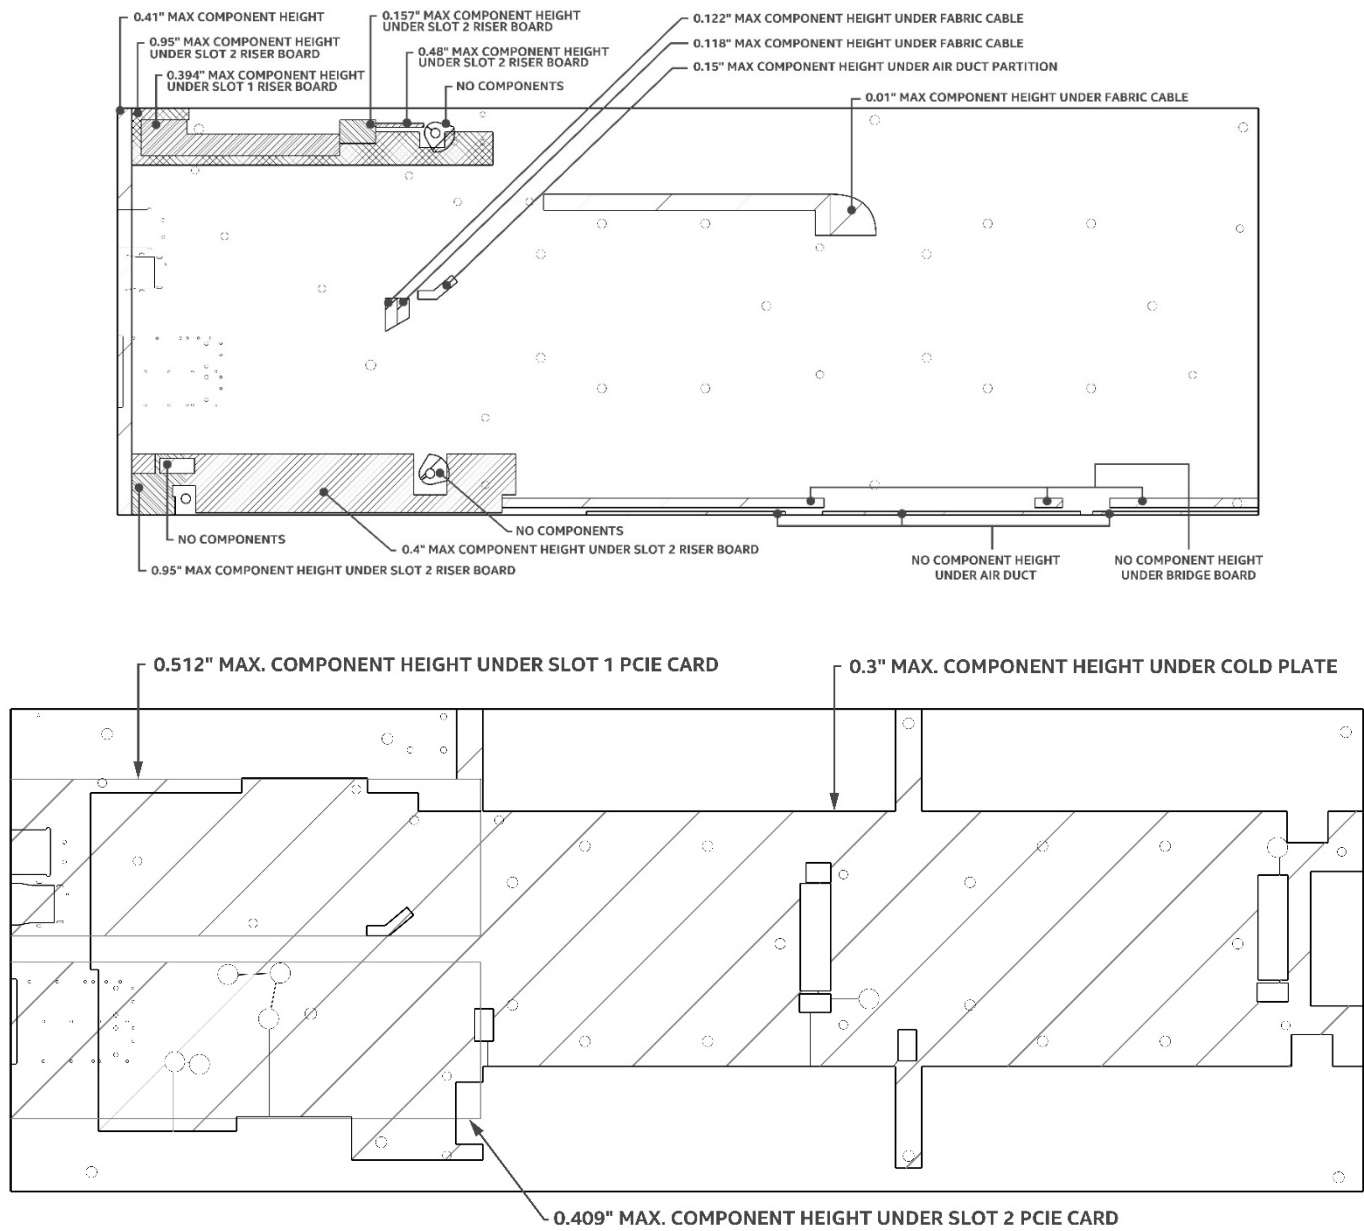


Figure 22. Intel® Server Board S2600BPS – Primary Side Height Restrictions

2.3.3 Server Board Architecture Overview

The architecture of Intel® Server Board S2600BP is developed around the integrated features and functions of the Intel® Xeon® processor Scalable family, the Intel® C62x Series Chipset Family, Intel® Ethernet Controller X550, and the ASPEED® AST2500 Server Board Management Controller. Note that Intel® Server Board **S2600BPS** (SFP+) model utilizes embedded Intel® Ethernet Controller X722 from the Intel® C622 Chipset. See **Figure 25**.

The following diagrams provide an overview of the server board architecture, showing the features and interconnects of each of the major sub-system components.

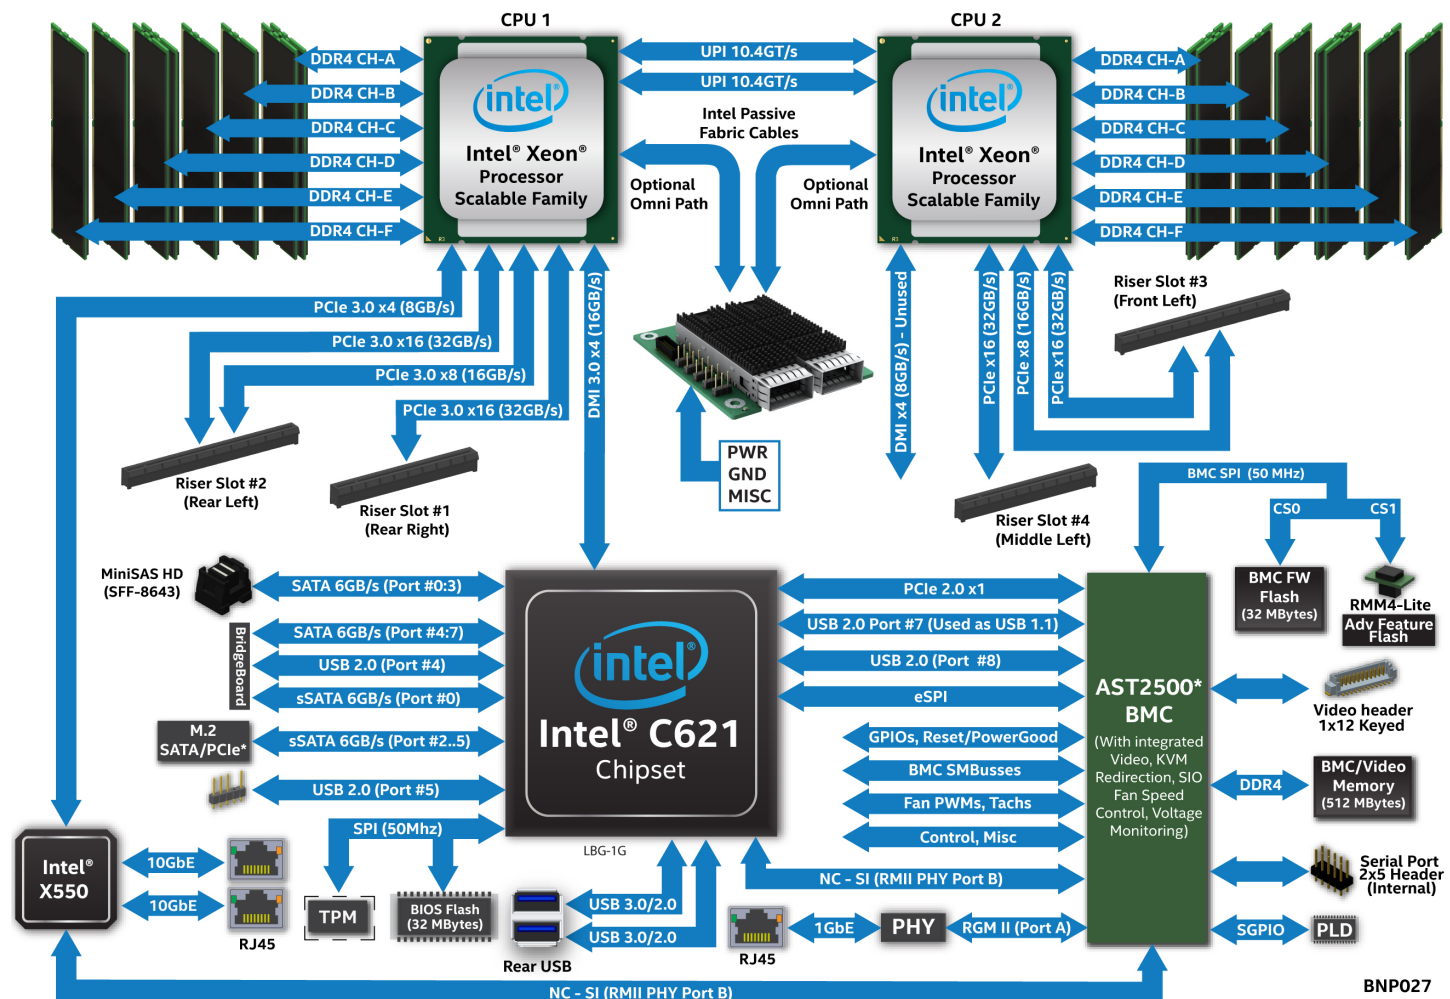


Figure 23. Intel® Server Board S2600BPB Architectural Block Diagram

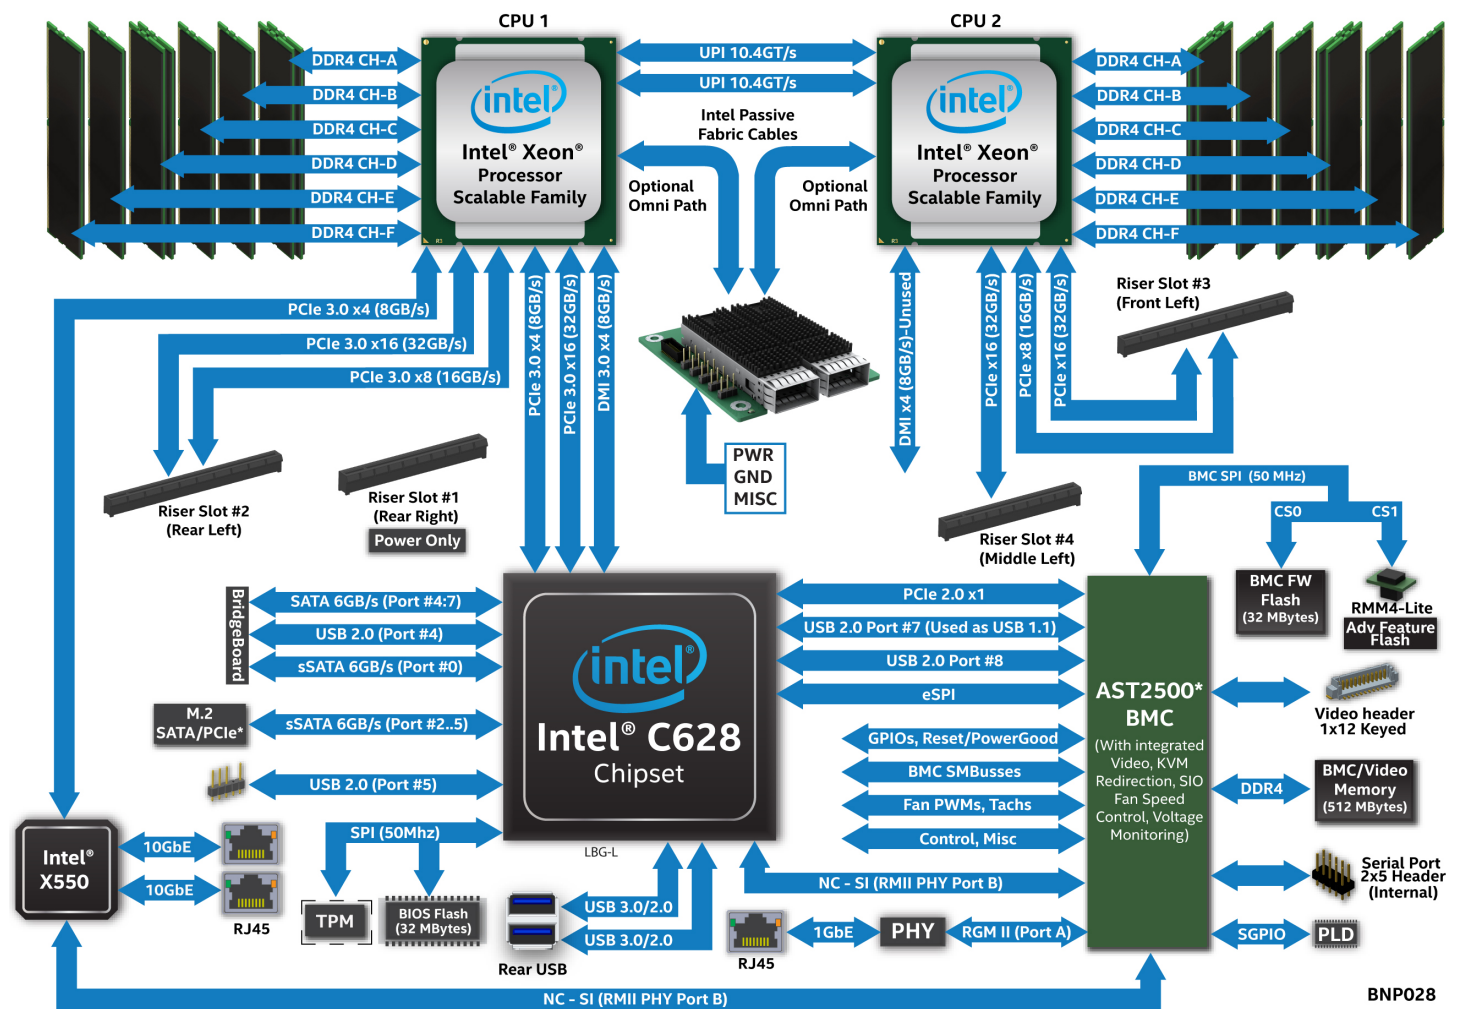


Figure 24. Intel® Server Board S2600BPQ Architectural Block Diagram

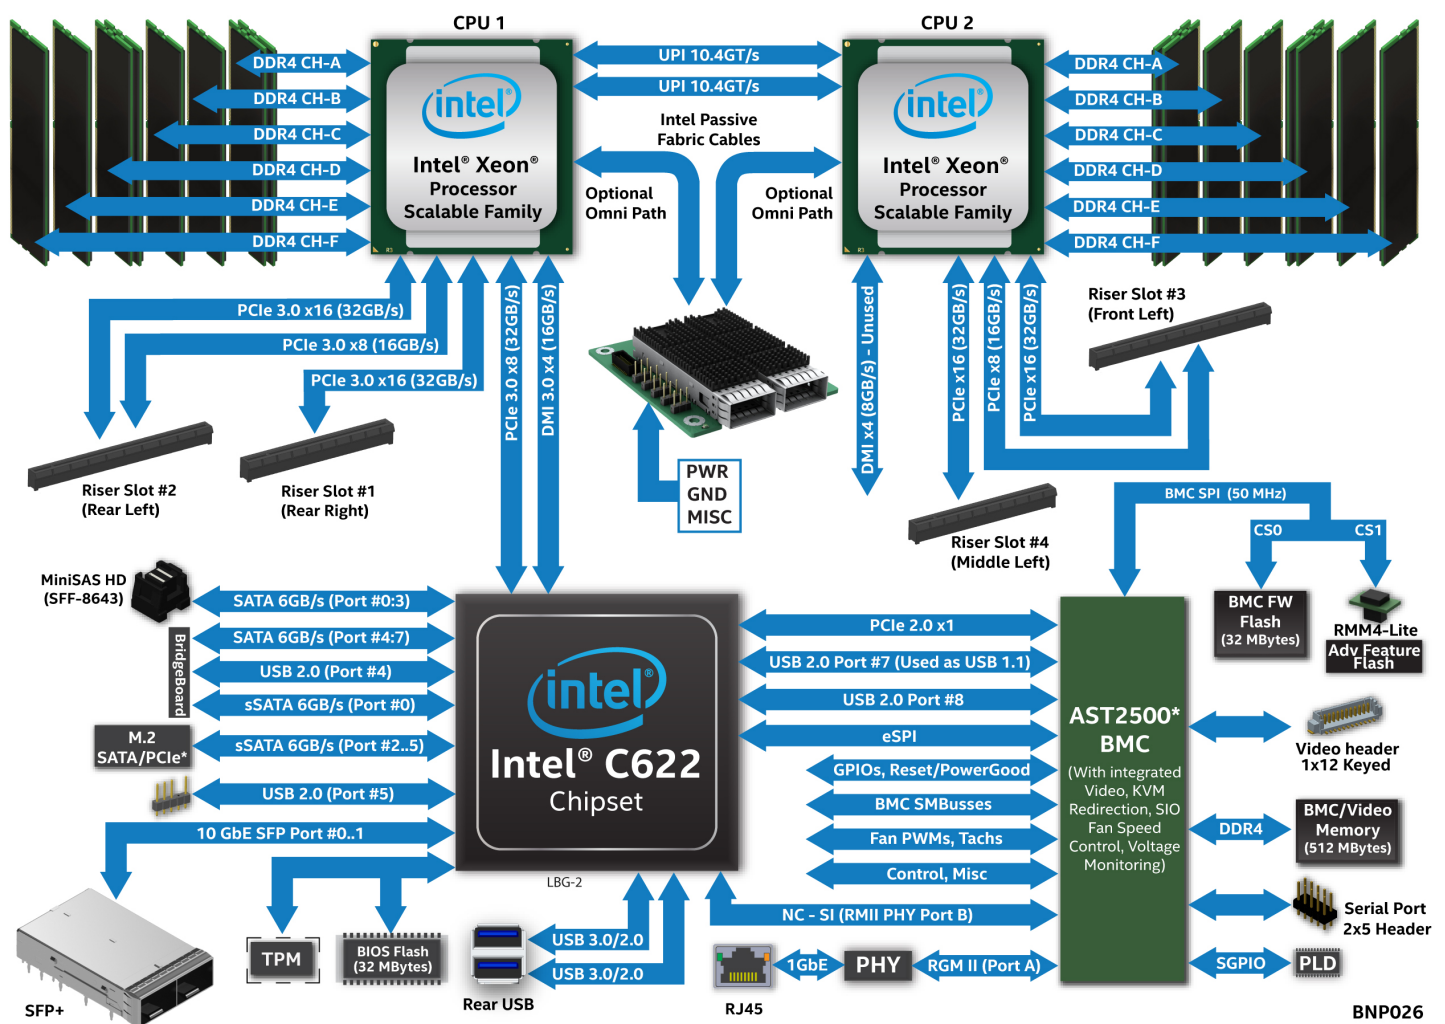


Figure 25. Intel® Server Board S2600BPS Architectural Block Diagram

2.4 Intel® Compute Module HNS2600BP Overview

The Intel® Compute Module HNS2600BP product family is a purpose build 1U density optimized compute module designed to operate as a single system node within a multi-node chassis. This Intel compute module family is compatible with select SKUs within the Intel® Server Chassis H2000G product family. See

Table 5. 2U Intel® Server Chassis H2000G Product Family Feature Set.

Table 9. Compute Module Specifications

Compute Module Specifications	Detail
Form Factor	1U Module
Module Weight	~ 3.6kg
Packaging Dimensions	716 x 269 x 158mm
Packaging Weight	~ 4.8kg




Figure 26. Intel® Compute Module HNS2600BP




Figure 27. Intel Server Chassis H2000G

2.4.1 Intel Compute Module Feature Identification

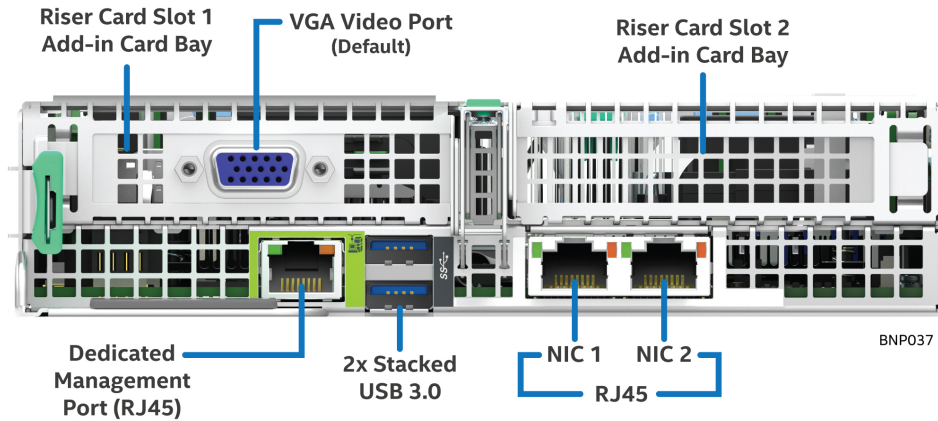



Figure 28. Intel® Compute Module HNS2600BPB and HNS2600BPQ – Rear Connectors




Figure 29. Intel® Compute Module HNS2600BPS – Rear Connectors

2.4.2 Compute Module Mechanical Dimensional Diagrams




Figure 30. Intel® Compute Module HNS2600BPx – Dimension




Figure 31. Intel® Compute Module HNS2600BPx24 – Dimension

2.5 Compute Module Components Overview

The Intel® Compute Module HNS2600BP product family includes several components to provide the module with the appropriate air flow and interconnects when installed within an Intel server chassis. The following sections provide an overview of each component and system interface board.


2.5.1 Power Docking Boards

Depending on the compute module model, one of the following power docking boards is used to enable hot swap support of the compute module into or out of the Intel server chassis.

2.5.1.1 Standard Power Docking Board

The power docking board provides hot swap docking of 12V main power between the compute module and the server. It supports three dual-rotor fan connections, a 12V main power hot swap controller, and current sensing. The standard power docking board is used in the following Intel compute modules:

- HNS2600BPB
- HNS2600BPS
- HNS2600BPQ




Label	Description
A	8 pin connector for fan 1
B	12 pin connector for main power input
C	8 pin connector for fan 3
D	2x6 pin Main Power Output connector
E	8 pin connector for fan 2
F	2x7 pin fan control connector

Figure 32. Standard Power Docking Board

2.5.1.2 SAS/NVMe Combo Power Docking Board

The SAS/NVMe Combo Power Docking Board provides hot swap docking of 12V main power between the compute module and the server. It supports three dual-rotor fan connections, a 12V main power hot swap controller, current sensing, and routes PCIe signals from the server board to the chassis backplane. The SAS/NVMe Combo Power Docking Board is used only with the following Intel® Compute Modules.

- HNS2600BPB24
- HNS2600BPS24
- HNS2600BPQ24



BNP049

Label	Description
A	40 pin Misc. Signal connector (to bridge board)
B	2x7 pin fan control connector
C	8 pin connector for Fan 1
D	2x6 pin Main Power Output connector
E	8 pin connector for Fan 3
F	8 pin connector for Fan 2
G	40 pin Misc. Signal Card Edge connector (to BIB)
H	Power Blade Card Edge connector (to BIB)

Figure 33. SAS/PCIe* SFF Combo Power Docking Board Top View

2.5.2 Bridge Board Options

There are six optional bridge board options that offer different features and functions on the system:

- AHWBPBGB: 4-port, 6G SATA bridge board
- AHWBP12GBGB: 4-port, 12G SAS bridge board (w/LSI* 3008 IOC)
- AHWBP12GBGBR5: 4-port, 12G SAS bridge board (w/LSI 3008 IOC)
- AHWBP12GBGBIT: 4 –port, 12G SAS bridge board (w/LSI* 3008 IOC)
- AHWBPBGB24R : 6-port, 12G SAS/PCIe* SFF Combo Bridge Board (w/LSI 3408 IOC)
- AHWBPBGB24: 6-port, 12G SAS/PCIe* SFF Combo Bridge Board (w/LSI 3008 IOC)
- AHWBPBGB24P: 6-port, Pass-through Bridge Board (Requires RAID AIC to be fully functional)

Table 10. Computer Module and Bridge board support matrix

Product Code	Description	SATA / SAS	RAID 0, 1, 10	RAID 5	Compatible Compute Module and Chassis
AHWBPBGB	4-Port Bridge Board	6G SATA PCH	ESRT2 SW RAID 0/1/10	RAID 5 with optional key	Compute Module: HNS2600BPB HNS2600BPS HNS2600BPQ
AHWBP12GBGB	4-Port iMR Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS	RAID 0/1/10	No	
AHWBP12GBGBR5	4-Port iMR Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS	RAID 0/1/10	RAID 5o	
AHWBP12GBGBIT	4-Port IT Bridge Board w/LSI 3008 IOC	6G SATA 12G SAS	JBOD MODE		Chassis: H2312XXLR3
AHWBPBGB24	6-Port IT Bridge Board w/ LSI 3008 IOC	6G SATA 12G SAS NVMe	JBOD MODE		Compute Module: HNS2600BPB24 HNS2600BPS24 HNS2600BPQ24
AHWBPBGB24R	6-Port IT iMR Bridge Board w/ LSI 3408 IOC	6G SATA 12G SAS NVMe	RAID 0/1/10	RAID 5 with optional key	
AHWBPBGB24P	6-Port Pass-Through Bridge Board (AIC)	Pass-Through	RAID 0, 1, 10 with optional add-in card	RAID 5 with optional add-in card	

Bridge Board Support Notes:

- Intel compute modules do not ship with a bridge board. Bridge boards must be ordered separately from the compute module.
- A bridge board is required in system configurations that need HDD/SSD support from any of the front drive bays in an Intel chassis.
- When ordering a bridge board, ensure it is compatible with the given compute module. See the table above.
- Dual processor system configurations are required to support a bridge board with 12G SAS support. 12G SAS bridge boards will not be functional in a single processor system configuration.
- Each compute module installed within an Intel chassis supports its own bridge board

2.5.2.1 6G SATA Bridge Board (iPC – AHWBPBGB)

The 6 GB SATA bridge board provides data lanes for up four SATA ports to the backplane of the server chassis.



Figure 34. 6G SATA Bridge Board Overview

2.5.2.2 2G SAS Bridge Boards (iPCs – AHWBP12GBGB / AHWBP12GBGBR5 / AHWBP12GBGBIT)

Bridge boards that include support for 12 GB SAS include an LSI* SAS 3008 controller to support up to four SAS/SATA ports. Separate 12 GB SAS bridge board options are available to provide the following RAID levels.

- AHWBP12GBGBIT – JBOD Mode
- AHWBP12GBGB – RAID levels 0, 1, and 10
- AHWBP12GBGBR5 – RAID levels 0, 1, 5, and 10




Figure 35. 12G SAS Bridge Board Overview

2.5.2.3 12G SAS/PCIe* NVMe Combo Bridge Board - 3008 (iPC – AHWBPBGB24)

The 12G SAS/NVMe combo bridge board has one embedded LSI* SAS 3008 controller to support up to six 12Gb/s SAS ports and up to two x4 PCIe* 3.0 lanes to support up to two PCIe* NVMe drives. This bridge board has no embedded RAID support.




Figure 36. SAS/PCIe* SFF Combo Bridge Board Overview

Note: This bridge board is only supported in the 24 drive Intel chassis. See Intel Chassis features tables for supported drive configurations.

2.5.2.4 12G SAS/PCIe* SFF Combo Bridge Board – 3408 (iPC – AHWBPGGB24R)

The 12G SAS/PCIe* SFF combo bridge board has one embedded LSI* SAS 3408 controller to support up to six 12Gb/s SAS ports and two x4 PCIe* 3.0 lanes to support up to two PCIe* SFF devices. This bridge board includes support for RAID Levels 0, 1, and 10. RAID 5 can be supported with the addition of an optional RAID 5 Key.




Figure 37. SAS/PCIe* SFF Combo Bridge Board Overview

2.5.2.5 12G SAS Pass-through Bridge Board (iPC – AHWBPGGB24P)

The 12G SAS Pass-through bridge board provide the I/O connectivity to support up to six 12Gb/s SAS ports between an add-in Host Bus Adapter (HBA) card and the backplane.




Figure 38. Pass Through Bridge Board Overview

2.5.3 Compute Module Riser Card Support

The server board includes four riser card slots. Within the Intel compute modules, Riser Slot #3 and Riser Slot #4 are used to support the various bridge board options. Riser Slots #1 and #2 each have support for riser cards to support different add-in options based on the model of the compute module.

By default, all Intel compute models within the product family will have an add-in VGA video connector bracket installed into the Riser #1 add-in card slot on the back of the compute module. To use the add-in card slot on Riser #1 for any other purpose, the VGA video connector bracket must be removed, thus losing video support from the compute module.

2.5.3.1 Riser Slot #1 Riser Card

Riser Slot 1 functionality is dependent on the specific model of the compute module.

- Intel® Compute Modules **HNS2600BPB & HNS2600BPB24**

Riser slot 1 for these Intel compute modules can be used for one of the following add-in options:

- Riser card supporting one PCIe* 3.0 x16 Elec, x16 Mech add-in card slot. Able to support low profile PCIe add-in cards only
- Support for optional Intel® Omni-path fabric through carrier card
- VGA video connector bracket (Standard Default)

- Intel® Compute Modules **HNS2600BPS, HNS2600BPS24, HNS2600BPQ, HNS2600BPQ24**

The intended use for Riser slot #1 in these compute modules is to support the following options only.

- Support for optional Intel® Omni-path fabric through carrier card
- VGA video connector bracket (Standard Default)

Note: in these compute module models, Riser slot 1 cannot support any PCIe* add-in cards.




Figure 39. Riser Card for Riser Slot 1

2.5.3.2 Riser Slot 2 Riser Card

The riser card for Riser Slot #2 includes concurrent support for the following add-in options in all available Intel compute modules:

- Riser card with support for the following features
 - One PCIe* 3.0 x16 elec, x16 mech add-in card slot. Able to support low profile PCIe add-in cards only.
 - One M.2 PCIe* connector located at the back side of the riser card. Able to support 80mm M.2 drives




Figure 40. Riser Card for Riser Slot 2

2.5.4 Compute Module Air Flow Support

For a compute module to operate at its best performance and meet long term reliability goals, it must remain at or below the thermal limits identified in section 2.4 Environmental Limits. To support the necessary air flow, the Intel compute module includes the following components: three 40 x 40 x 56 dual rotor fans and one air duct.




Figure 41. Compute Module System Fans and Air Duct

2.5.4.1 System Fans

The Intel compute module includes three dual rotor 40 x 40 x 56 system managed fans providing front to back air flow through the compute module.




Figure 42. Compute Module Air Flow

Each fan is mounted within a metal housing on the compute module base as shown in the following figure.




Figure 43. Compute Module Fan Placement

Each system fan is cabled to separate 8-pin connectors on the power docking board. See section 2.5.1. Fan control signals for each system fan are then routed to the server board through a single 2x7 connector on the power docking board, which is cabled to a matching fan controller header on the server board. See section 7.4 for the fan control header pinout.

Each fan within the compute module can support variable speeds. Fan speed may change automatically when any temperature sensor reading changes. Each fan connector within the module supplies a tachometer signal that allows the BMC to monitor the status of each fan. The fan speed control algorithm is programmed into the server board's integrated Baseboard Management Controller (BMC).

Intel compute modules do not support fan redundancy. Should a single rotor stop working, the following events will most likely occur:

- The integrated baseboard management controller (BMC) detects the fan failure
- The event is logged to the System Event Log (SEL)
- The System Status LED on the server board and chassis front panel will turn flashing Green, indicating a system is operating at a degraded state and may fail at some point
- In an effort to keep the compute module at or below pre-programmed maximum thermal limits monitored by the BMC, the remaining functional system fans will operate at 100%

As system thermal continue to rise:

- To minimize system thermals, processors and memory within the compute module will begin to throttle, effecting system performance.
- The System Status LED will likely change to blinking Amber, indicating the system is operational but in a critical state, system failure is likely to occur
- Should the compute module thermals continue above pre-programmed thermal limits, the compute module will shut down.
- The System Status LED will change state to solid Amber. A fatal error has occurred.

Fans are not hot swappable. Should a fan fail, the compute module must be powered down and removed from the chassis before the faulty fan can be replaced.

WARNING: The Intel® Compute Module HNS2600BP product family does not support redundant cooling. If one of the compute module fans fails, it is recommended to replace the failed fan as soon as possible.

2.5.4.2 Air Duct

To ensure proper airflow over critical components within the Intel compute module, the plastic air duct must be installed and properly seated whenever the compute module is operational.





Figure 44. Intel Compute Module Air Duct Placement

WARNING: Before sliding the compute module into the chassis, make sure the air duct is installed properly. For detailed instructions regarding air duct installation refer to the *Intel® Server Board S2600BP and Intel® Compute Module HNS2600BP Product Family Integration and Service Guide*.

In system configurations where CPU 1 is configured with a processor SKU that supports an Integrated Intel® Omni-Path Host Fabric Interface, an additional plastic air baffle is attached to the bottom side of the air duct as shown in the following figure.



BNP122

Figure 45. Air Baffle Addition

The air baffle must be attached to the air duct to ensure proper airflow to the chip set and the Intel Fabric Through (IFT) carrier when installed.

The air baffle does not ship as a standard option with the compute module. The air baffle is made available as part of the Intel® Omni-Path Fabric Processor accessory kit (iPC - **AHWBPFABKIT**).

2.6 System Software Stack

The server board includes a system software stack that consists of the System BIOS, BMC firmware, ME Firmware, and FRU and SDR data. Together, they configure and manage features and functions of the server system.

Many features and functions of the server system are managed jointly by the System BIOS and the BMC firmware, this include:

- IPMI Watchdog timer
- Messaging support, including command bridging and user/session support
- BIOS boot flags support
- Event receiver device: The BMC receives and processes events from the BIOS
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Integrated KVM
- Integrated Remote Media Redirection
- Intel® Intelligent Power Node Manager support
- Sensor and SEL logging additions/enhancements (e.g., additional thermal monitoring capability)
- Embedded platform debug feature, which allows capture of detailed data for later analysis by Intel engineering.

A complete system software stack is pre-programmed by Intel on the server board during the board assembly process, making the server board functional at first power on. However, to ensure the most reliable system operation, it is highly recommended that you check the following Intel website for the latest available system updates: <https://downloadcenter.intel.com/product/93091/>

System updates can be performed in a number of operating environments, including the UEFI Shell using the uEFI only System Update Package (SUP), or under different operating systems using the Intel® One Boot Flash Update Utility (OFU).

As part of the initial system integration process, system integrators must program system configuration data onto the server board using the *FRUSDR Utility* to ensure the embedded platform management subsystem is able to provide the best performance and cooling for the final system configuration. The FRUSDR Utility is included in the SUP and OFU packages. See section 2.5.2 for additional information.

You can reference the following Intel documents for more in-depth information about the system software stack and their functions:

- *Intel® Server System BIOS External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® Processor Scalable Family* – Intel NDA Required
- *Intel® Server System BMC Firmware External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® Processor Scalable Family product family* – Intel NDA Required

2.6.1 Hot Keys Supported During POST

Certain “Hot Keys” are recognized during the system Power On Self Test (POST). The POST process occurs after system power on and before the operating system starts to load. A Hot Key is a key or key combination that is recognized as an unprompted command input, where the operator is not prompted to press the Hot Key. In most cases Hot Keys will be recognized even while other processing is in progress.

The BIOS supported Hot Keys are only recognized by the system BIOS during the system boot time POST process. Once the POST process has completed and hands off the system boot process to the operating system, BIOS supported Hot Keys are no longer recognized.

The following table provides a list of BIOS supported Hot Keys.

Table 11. POST Hot-Keys

HotKey Combination	Function
<F2>	Enter the BIOS Setup Utility
<F6>	Pop-up BIOS Boot Menu
<F12>	Network boot
<Esc>	Switch from Logo Screen to Diagnostic Screen
<Pause>	Stop POST temporarily

2.6.1.1 POST Logo/Diagnostic Screen

The Logo/Diagnostic Screen appears in one of two forms:

- If Quiet Boot is enabled in the <F2> BIOS setup, a “splash screen” is displayed with a logo image, which may be the standard Intel Logo Screen or a customized OEM Logo Screen. By default, Quiet Boot is enabled in BIOS setup, so the Logo Screen is the default POST display. However, if the logo is displayed during POST, the user can press <Esc> to hide the logo and display the Diagnostic Screen instead.
- If a customized OEM Logo Screen is present in the designated Flash Memory location, the OEM Logo Screen will be displayed, overriding the default Intel Logo Screen.
- If a logo is not present in the BIOS Flash Memory space, or if Quiet Boot is disabled in the system configuration, the POST Diagnostic Screen appears with a summary of system configuration information. The POST Diagnostic Screen is purely a Text Mode screen, as opposed to the Graphics Mode logo screen.
- If Console Redirection is enabled in Setup, the Quiet Boot setting is disregarded and the Text Mode Diagnostic Screen is displayed unconditionally. This is due to the limitations of Console Redirection, which transfers data in a mode that is not graphics-compatible.

2.6.1.2 BIOS Boot Pop-Up Menu

The BIOS Boot Specification (BBS) provides a Boot pop-up menu that can be invoked by pressing the <F6> key during POST. The BBS pop-up menu displays all available boot devices. The boot order in the pop-up menu is not the same as the boot order in the BIOS setup. The pop-up menu simply lists all of the available devices from which the system can be booted, and allows a manual selection of the desired boot device.

When an Administrator password is installed in Setup, the Administrator password will be required in order to access the Boot pop-up menu using the <F6> key. If a User password is entered, the Boot pop-up menu will not even appear – the user will be taken directly to the Boot Manager in the Setup Utility, where a User password allows only booting in the order previously defined by the Administrator.

2.6.1.3 Entering BIOS Setup

To enter the BIOS Setup Utility using a keyboard (or emulated keyboard), press the <F2> function key during boot time when the OEM or Intel Logo screen or the POST Diagnostic screen is displayed.

The following instructional message appears on the Diagnostic Screen or under the Quiet Boot Logo screen:

Press <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot

Note: With a USB keyboard, it is important to wait until the BIOS “discovers” the keyboard and beeps – until the USB Controller has been initialized and the USB keyboard activated, key presses will not be read by the system.

When the Setup Utility starts, the Main screen is displayed initially. However, in the event that a serious error occurs during POST, the system will enter the BIOS Setup Utility and display the Error Manager screen instead of the Main screen.

Reference the following Intel document for additional BIOS Setup information:

- *Intel® Server System BIOS Setup Guide for Intel® Servers Systems supporting the Intel® Xeon® processor Scalable family*

2.6.1.4 BIOS Update Capability

In order to bring BIOS fixes or new features into the system, it will be necessary to replace the current installed BIOS image with an updated one. The BIOS image can be updated using a standalone IFLASH32 utility in the uEFI shell, or can be done using the OFU utility program under a supported operating system. Full BIOS update instructions are provided with update packages downloaded from the Intel website.

2.6.1.5 BIOS Recovery

If a system is completely unable to boot successfully to an OS, hangs during POST, or even hangs and fails to start executing POST, it may be necessary to perform a BIOS Recovery procedure, which can replace a defective copy of the Primary BIOS

The BIOS introduces three mechanisms to start the BIOS recovery process, which is called Recovery Mode:

- At power on, the BIOS Boot Block detects a partial BIOS update was performed and automatically boots in Recovery Mode.
- The BMC asserts the Recovery Mode GPIO in case of partial BIOS update and FRB2 time-out.
- The Recovery Mode Jumper causes the BIOS to boot in Recovery Mode.

The BIOS Recovery takes place without any external media or Mass Storage device as it utilizes the Backup BIOS inside the BIOS flash in Recovery Mode.

The Recovery procedure is included here for general reference. However, if in conflict, the instructions in the BIOS Release Notes are the definitive version.

When the Recovery Mode Jumper is set, the BIOS begins with a 'Recovery Start' event logged to the SEL, then loads and boots with the Backup BIOS image inside the BIOS flash itself. This process takes place before any video or console is available. The system boots up directly into the Shell while a 'Recovery Complete' SEL event is logged. From the uEFI Shell, the BIOS can then be updated using a standard BIOS update procedure, defined in Update Instructions provided with the system update package downloaded from the Intel website. After the update is complete, there will be a message displayed stating that the "BIOS has been updated successfully," indicating that the BIOS update process is finished. The User should then switch the recovery jumper back to normal operation and restart the system by performing a power cycle.

If the BIOS detects a partial BIOS update or the BMC asserts Recovery Mode GPIO, the BIOS will boot up in Recovery Mode. The difference is that the BIOS boots up to the Error Manager Page in the BIOS Setup Utility. In the BIOS Setup Utility, boot device, Shell, or Linux for example, could be selected to perform the BIOS update procedure under Shell or OS environment.

NOTE: Before attempting a Recovery Boot, it is highly advisable to reference the BIOS Release Notes to verify the proper Recovery procedure.

2.6.2 Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data

As part of the initial system integration process, the server board/system must have the proper FRU and SDR data loaded. This ensures that the embedded platform management system is able to monitor the appropriate sensor data and operate the system with best cooling and performance. Once the system integrator has performed an initial FRU SDR package update, subsequent auto-configuration occurs without the need to perform additional SDR updates or provide other user input to the system when any of the following components are added or removed:

- Processor
- Memory
- Integrated SAS Raid module
- Power supply
- Fan
- Hot Swap Backplane
- Front Panel

NOTE: The system may not operate with best performance or best/appropriate cooling if the proper FRU and SDR data is not installed.

2.6.2.1 Loading FRU and SDR Data

The FRU and SDR data can be updated using a standalone FRUSDR utility in the uEFI shell, or can be done using the OFU utility program under a supported operating system. Full FRU and SDR update instructions are provided with the appropriate system update package (SUP) or OFU utility which can be downloaded from the following Intel website. <http://downloadcenter.intel.com>

3. Processor Support

The server board includes two Socket-P LGA3647 processor sockets compatible with the following Intel processors:

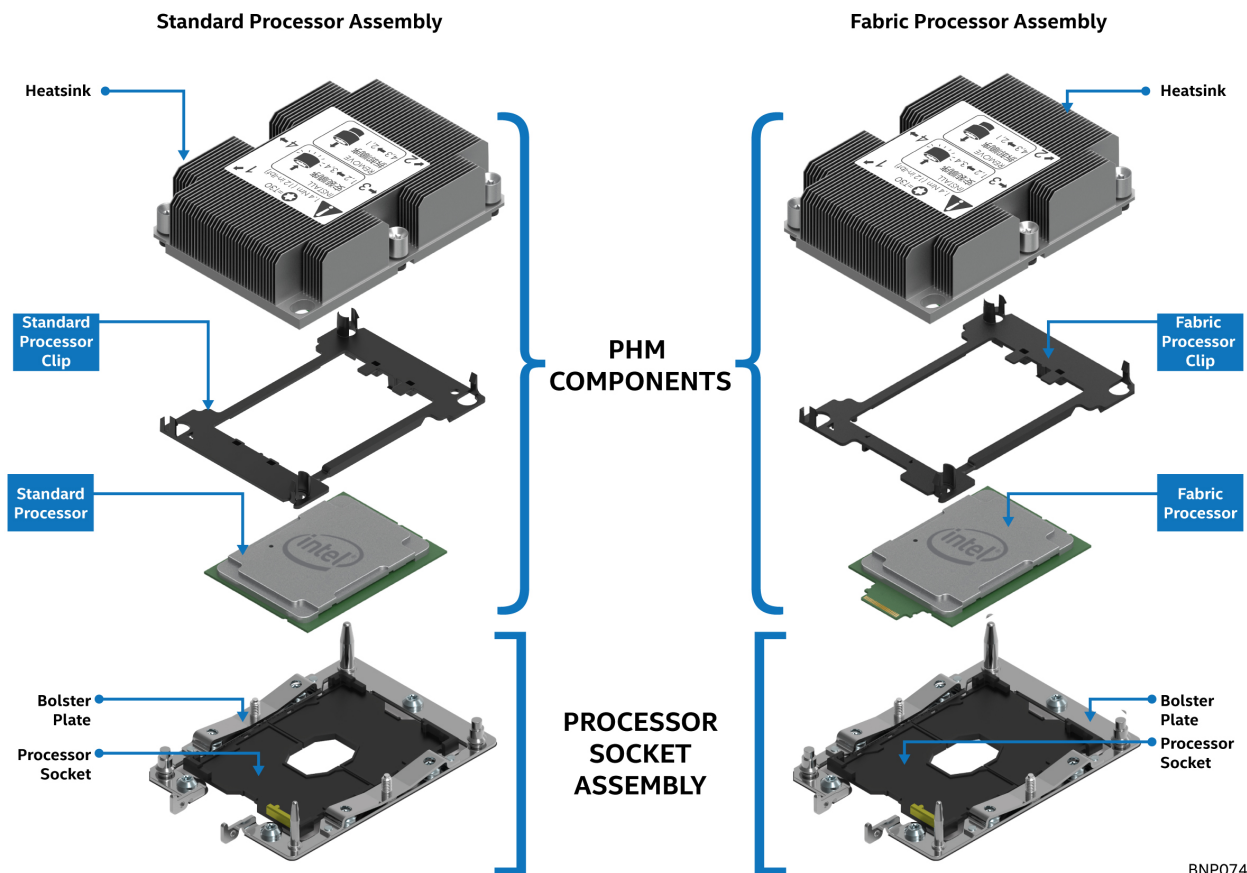
- Intel® Xeon® processor Scalable family with a thermal design power (TDP) of and up to **165W**
 - Intel® compute modules can only support processors with 165W TDP when installed within the high air flow Intel® Server Chassis H2204XXLRE
- Support for Intel® Xeon® processor Scalable family with Intel® Omni-path Integrated Fabric Connectors
 - One 100Gb/s port per processor

WARNING: Previous-generation Intel® Xeon® processors and their supported CPU heat sinks are not compatible with the Intel® Server Board S2600BP product family.

For detailed instructions regarding processor installation, please refer to the *Intel® Server Board S2600BP and Intel® Compute Module HNS2600BP Product Family Integration and Service Guide*.

3.1 Processor Socket and Processor Heat Sink Module (PHM) Assembly


This generation server board introduces the concept of the PHM (Processor Heat Sink module). The following illustration identifies each component associated with the processor assembly. Note that the illustration does NOT represent the processor installation process.



BNP074

Figure 46. PHM Components and Processor Socket Reference Diagram

Processor installation requires that **the processor be attached to the processor heat sink prior to installation onto the server board.**



Two Bolster Plate guide pins of different sizes allows the PHM to be installed only one way onto the processor socket assembly. See Figure 47




Figure 47. PHM to CPU Socket Orientation and Alignment Features

The PHM is properly installed when it is securely seated over the two Bolster Plate guide pins and it sits evenly over the processor socket. Once the PHM is properly seated over the processor socket assembly, the four heat sink torx screws must be tightened in the order specified on the label affixed to the top side of the processor heat sink.

Caution: Failure to tighten / untighten the heat sink screws in the specified order may cause damage to the processor socket assembly. Heat sink screws should be tightened to 12 In-Lbs Torque.

Note: For detailed processor assembly and installation instructions, refer to the appropriate Intel product family *System Integration and Service Guides*.

To protect the pins within a processor socket from being damaged, server boards with no processor or heat sink installed must have a plastic cover installed over each processor socket, as shown in Figure 48.




Figure 48. Processor Socket Assembly and Protective Cover

Processor socket covers must be removed before processor installation.

3.1.1 Bolster Plate Insulator for CPU 1

The Intel® HNS2600BP Compute Module includes a factory installed bolster insulator plate to prevent potential contact between a PCIe* add in card (when installed) and the metal bolster plate of the CPU #1 processor socket.




Figure 49. Bolster Insulator Plate for CPU1

NOTE: The insulator should only be removed when installing a processor SKU that supports the Intel® Omni-Path Fabric Host Fabric Interconnect (HFI). **The insulator must be re-installed after the setup is completed.** Do not operate the compute module when configured with a PCIe add-in card and with the insulator removed. Doing so may critically damage the PCIe add-in card, the server board, or both.

Refer to the *Intel® Server Board and Intel® Compute Module Service and Integration Guide* for detailed installation and removal instructions

3.2 Processor Thermal Design Power (TDP) Support


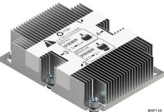
To allow optimal operation and long-term reliability of Intel processor-based systems, the processor must remain within the defined minimum and maximum case temperature (T_{CASE}) specifications. Thermal solutions not designed to provide sufficient thermal capability may affect the long-term reliability of the processor and system. The server boards and compute modules described in this document are designed to support the Intel® Xeon® processor Scalable Family TDP guidelines up to and including 165W.


NOTE: Intel Server Chassis compatible with the compute modules described in this document may support a lower maximum Thermal Design Power (TDP) than that of the compute module. See Table 5.

3.2.1 Processor Heat Sink

There are two types of heat sinks specially designed for optimal cooling performance, as indicated in Table 12. These are NOT interchangeable and must be installed in the indicated order. Figure 50 shows the proper placement.

Table 12. Processor Heat Sinks

Intel Product Code (iPC)	MM#	Description
iPC – FXXHP78X108HS 	MM# 956548	1U Standard Cu/Al 78mm x 108mm x 25.5mm Heat Sink (Rear Heat Sink) – for CPU 1 use only
iPC – FXXE78X108HS 	MM# 948936	1U Standard Ex-Al 78mm x 108mm x 25.5mm Heat Sink (Front Heat Sink) – for CPU 2 use only



BNP052

Figure 50. Processor Heat Sinks placement overview

Intel® Xeon® Processor Scalable Family Overview

The Intel® Server Board S2600BP product family has support for the Intel® Xeon® processor Scalable family:

- Intel® Xeon® Bronze XXXX processor
- Intel® Xeon® Silver XXXX processor
- Intel® Xeon® Gold XXXX processor
- Intel® Xeon® Platinum XXXX processor

XXXX = Intel defined processor SKUs

Table 13. Intel® Xeon® Processor Scalable Family - Feature Comparison Table

Feature	81xx Platinum	61xx Gold	51xx Gold	41xx Silver	31xx Bronze
# of Intel® UPI Links	3	3	2	2	2
UPI Speed	10.4 GT/s	10.4 GT/s	10.4 GT/s	9.6 GT/s	9.6 GT/s
Supported Topologies	2S-2UPI 2S-3UPI 4S-2UPI 4S-3UPI 8S- 3UPI	2S-2UPI 2S-3UPI 4S-2UPI 4S-3UPI	2S-2UPI 4S-2UPI	2S-2UPI	2S-2UPI
Node Controller Support	Yes	Yes	No	No	No
# of Memory Channels	6	6	6	6	6
Max DDR4 Speed	2666	2666	2400	2400	2133
Memory Capacity	768GB 1.5TB (Select SKUs)	768GB 1.5TB (Select SKUs)	768GB 1.5TB (Select SKUs)	768 GB	768 GB
RAS Capability	Advanced	Advanced	Advanced	Standard	Standard
Intel® Turbo Boost	Yes	Yes	Yes	Yes	No
Intel® Hyper-Threading	Yes	Yes	Yes	Yes	No
Intel® AVX-512 ISA Support	Yes	Yes	Yes	Yes	Yes
Intel® AVX-512 - # of 512b FMA Units	2	2	1	1	1
# of PCIe* Lanes	48	48	48	48	48

Additional processor SKUs with an Integrated Intel® Omni-Path Host Fabric Interface are also supported. See Section 3.5.

3.2.2 Processor Features Overview

The Intel® Xeon® processor Scalable Family combines several key system components into a single processor package, including the CPU cores, Integrated Memory Controller (IMC), and Integrated IO Module (IIO).

The processor includes many core and uncore features and technologies including:

Core Features:

- Intel® Ultra Path Interconnect (UPI) – up to 10.4 GT/s
- Intel® Speed Shift Technology
- Intel® 64 Architecture
- Enhanced Intel® SpeedStep Technology
- Intel® Turbo Boost Technology 2.0
- Intel® Hyper-Threading Technology
- Intel® Virtualization Technology (Intel® VT-x)
- Intel® Virtualization Technology for Directed I/O (Intel® VT-d)
- Execute Disable Bit

- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® Advanced Vector Extensions (Intel® AVX-512)
- Advanced Encryption Standard New Instructions (AES-NI)

Uncore Features:

- Up to 48 PCIe* lanes 3.0 lanes per CPU – 79GB/s bi-directional pipeline
- 6 Channels DDR4 memory support per CPU
- On package integration of next generation Intel Omni-Path Fabric Controller – Select SKUs
- DMI3/PCI express* 3.0 interface with a peak transfer rate of 8.0 GT/s.
- Non-Transparent Bridge (NTB) Enhancements – 3 full duplex NTBs and 32 MSI-X vectors
- Intel® Volume Management Device (Intel® VMD) – Manages CPU Attached NVMe SSDs
- Intel® Quick Data Technology
- Support for Intel® Node Manager 4.0

3.2.2.1 Intel® 64 Instruction Set Architecture (Intel® 64)

64-bit memory extensions to the IA-32 architecture. Further details on Intel 64 architecture and programming model can be found at <http://developer.intel.com/technology/intel64/>

3.2.2.2 Intel® Hyper-Threading Technology

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

3.2.2.3 Enhanced Intel® SpeedStep Technology

Processors in the Fifth Generation Intel® Core™ Processor Family support Enhanced Intel SpeedStep® Technology (EIST). The processors support multiple performance states, which allows the system to dynamically adjust processor voltage and core frequency as needed to enable decreased power consumption and decreased heat production. All controls for transitioning between states are centralized within the processor, allowing for an increased frequency of transitions for more effective operation.

The Enhanced Intel Speedstep® Technology feature may be enabled/disabled by an option on the Processor Configuration Setup screen. By default EIST is enabled. If EIST is disabled, then the processor speed is set to the processor's Max TDP Core Frequency (nominal rated frequency).

3.2.2.4 Intel® Turbo Boost Technology 2.0

Intel® Turbo Boost Technology is featured on all processors in the Fifth Generation Intel® Core™ Processor Family. Intel® Turbo Boost Technology opportunistically and automatically allows the processor to run faster than the marked frequency if the processor is operating below power, temperature, and current limits. This results in increased performance for both multi-threaded and single-threaded workloads.

3.2.2.5 Intel® Virtualization Technology (Intel® VT-x)

Hardware support in the core, to improve performance and robustness for virtualization. Intel VT-x specifications and functional descriptions are included in the Intel® 64 and IA-32 Architectures Software Developer's Manual.

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

Hardware support in the core and uncore implementations to support and improve I/O virtualization performance and robustness.

3.2.2.7 Execute Disable Bit

Intel's Execute Disable Bit functionality can help prevent certain classes of malicious buffer overflow attacks when combined with a supporting operating system. This allows the processor to classify areas in memory by where application code can execute and where it cannot. When malicious code attempts to insert code in the buffer, the processor disables code execution, preventing damage and further propagation.

3.2.2.8 Intel® Trusted Execution Technology for servers (Intel® TXT)

Intel TXT defines platform-level enhancements that provide the building blocks for creating trusted platforms. The Intel TXT platform helps to provide the authenticity of the controlling environment such that those wishing to rely on the platform can make an appropriate trust decision. The Intel TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

3.2.2.9 Intel® Advanced Vector Extension (Intel AVX-512)

The base of the 512-bit SIMD instruction extensions are referred to as Intel® AVX-512 foundation instructions. They include extensions of the AVX family of SIMD instructions but are encoded using a new encoding scheme with support for 512-bit vector registers, up to 32 vector registers in 64-bit mode, and conditional processing using opmask registers.

3.2.2.10 Advanced Encryption Standard New Instructions (AES-NI)

Intel® Advanced Encryption Standard New Instructions (AES-NI) is a set of instructions implemented in all processors in the Fifth Generation Intel® Core™ Processor Family. This feature adds AES instructions to accelerate encryption and decryption operations used in the Advanced Encryption Standard. The Intel® AES-NI feature includes 6 additional Single Instruction Multiple Data (SIMD) instructions in the Intel® Streaming SIMD Extensions (SSE) instruction set.

The BIOS is responsible in POST to detect whether the processor has the AES-NI instructions available. Some processors may be manufactured without AES-NI instructions.

The AES-NI instructions may be enabled or disabled by the BIOS. AES-NI instructions will be in an enabled state unless the BIOS has explicitly disabled them.

3.2.2.11 Intel® Node Manager 4.0

The Intel® C620 series chipset Management Engine (ME) supports Intel® Intelligent Power Node Manager (NM) technology. The ME/NM combination is a power and thermal control capability on the platform, which exposes external interfaces that allow IT (through external management software) to query the ME about platform power capability and consumption, thermal characteristics, and specify policy directives (that is, set a platform power budget). The ME enforces these policy directives by controlling the power consumption of underlying subsystems using available control mechanisms (such as processor P/T states). The determination of the policy directive is done outside of the ME either by intelligent management software or by the IT operator.

Below are some of the applications of Intel® Intelligent Power Node Manager technology.

- **Platform Power Monitoring and Limiting:** The ME/NM monitors platform power consumption and holds average power over duration. It can be queried to return actual power at any given instance. The power limiting capability is to allow external management software to address key IT issues by setting a power budget for each server.
- **Inlet Air Temperature Monitoring:** The ME/NM monitors server inlet air temperatures periodically. If there is an alert threshold in effect, then ME/NM issues an alert when the inlet (room) temperature exceeds the specified value. The threshold value can be set by policy.

- **Memory Subsystem Power Limiting:** The ME/NM monitors memory power consumption. Memory power consumption is estimated using average bandwidth utilization information.
- **Processor Power monitoring and limiting:** The ME/NM monitors processor or socket power consumption and holds average power over duration. It can be queried to return actual power at any given instant. The monitoring process of the ME will be used to limit the processor power consumption through processor P-states and dynamic core allocation.
- **Core allocation at boot time:** Restrict the number of cores for OS/VMM use by limiting how many cores are active at boot time. After the cores are turned off, the CPU will limit how many working cores are visible to the BIOS and OS/VMM. The cores that are turned off cannot be turned on dynamically after the OS has started. It can be changed only at the next system reboot.
- **Core allocation at run-time:** This particular use case provides a higher level processor power control mechanism to a user at runtime, after booting. An external agent can dynamically use or not use cores in the processor subsystem by requesting ME/NM to control them, specifying the number of cores to use or not use.

NOTE: For additional information on Intel® Intelligent Power Node Manager usage and support, visit the following Intel Website:

<http://www.intel.com/content/www/us/en/data-center/data-center-management/node-manager-general.html?wapkw=node+manager>

3.3 Processor Population Rules

NOTES: The server board may support mixed processor configurations that meet the defined criteria below. However, Intel does not perform mixed processor validation testing. In addition, Intel does not guarantee that a server system configured with unmatched processors will operate reliably. The system BIOS will attempt to operate with processors which are not matched but are generally compatible.

For optimal system performance in dual non-fabric processor configurations, Intel recommends that identical processors be installed.

Clearing and resetting the server board CMOS is required after every processor configuration change. Clearing and resetting the CMOS is achieved by setting the “BIOS DEFAULT” jumper on the server board. See Chapter 8, Configuration and recovery jumpers.

It is mandatory to populate the first processor socket (labeled as “CPU_1”) for single processor configurations.

Some board features may not be functional if a second processor is not installed. See section 2.3.3 for details.

When two processors are installed, the following population rules apply:

- Both processors must have the same number of cores
- Both processors must have the same cache sizes for all levels of processor cache memory
- Both processors must support identical DDR4 frequencies
- Both processors must have identical extended family, extended model, processor type, family code and model number
- No mixing of processors with FPGA and processors with Intel® Omni-Path Fabric

Processors with different core frequencies can be mixed in a system, given that the prior rules are met. If this condition is detected, all processor core frequencies are set to the lowest common denominator (highest common speed) and an error is reported.

Processor stepping within a common processor family can be mixed as long as it is listed in the processor specification updates published by Intel Corporation. Mixing of steppings is only validated and supported between processors that are plus or minus one stepping from each other.

3.4 Processor Initialization Error Summary

The table on this section describes mixed processor error conditions and recommended actions for all Intel® Server Boards and Intel® Server Systems designed around the Intel® Xeon® processor Scalable Family and Intel® C620 Series Chipset architecture. The errors fall into one of the following categories:

- **Fatal:** If the system can boot, POST will halt and display the following message:

“Unrecoverable fatal error found. System will not boot until the error is resolved
Press <F2> to enter setup”

When the <F2> key on the keyboard is pressed, the error message is displayed on the Error Manager screen, and an error is logged to the System Event Log (SEL) with the POST Error Code.

This operation will occur regardless of whether the BIOS Setup option “Post Error Pause” is set to Enable or Disable.

If the system is not able to boot, the system will generate a beep code consisting of three long beeps and one short beep. The system cannot boot unless the error is resolved. The faulty component must be replaced.

The System Status LED will be set to a steady amber color for all Fatal errors that are detected during processor initialization. A steady amber System Status LED indicates that an unrecoverable system failure condition has occurred.

- **Major:** If the BIOS Setup option for “Post Error Pause” is Enabled, and a Major error is detected, the system will go directly to the Error Manager screen in BIOS Setup to display the error, and logs the POST Error Code to SEL. Operator intervention is required to continue booting the system.

If the BIOS Setup option for “POST Error Pause” is Disabled, and a Major error is detected, the Post Error Code may be displayed to the screen, will be logged to the BIOS Setup Error Manager, an error event will be logged to the System Event Log (SEL), and the system will continue to boot.

- **Minor:** An error message may be displayed to the screen, the error will be logged to the BIOS Setup Error Manager, and the POST Error Code is logged to the SEL. The system continues booting in a degraded state. The user may want to replace the erroneous unit. The POST Error Pause option setting in the BIOS setup does not have any effect on this error.

Table 14. Mixed Processor Configurations Error Summary

Error	Severity	System Action when BIOS Detects the Error Condition
Processor family not identical	Fatal	Halts at POST code 0xE6. Halts with three long beeps and one short beep. Takes fatal error action (see above) and does not boot until the fault condition is remedied.
Processor model not identical	Fatal	Logs the POST error code into the SEL. Alerts the BMC to set the system status LED to steady amber. Displays 0196: Processor model mismatch detected message in the error manager. Takes fatal error action (see above) and does not boot until the fault condition is remedied.
Processor cores/threads not identical	Fatal	Halts at POST code 0xE5. Halts with three long beeps and one short beep. Takes fatal error action (see above) and does not boot until the fault condition is remedied.
Processor cache or home agent not identical	Fatal	Halts at POST code 0xE5. Halts with three long beeps and one short beep. Takes fatal error action (see above) and does not boot until the fault condition is remedied.
Processor frequency (speed) not identical	Fatal	If the frequencies for all processors can be adjusted to be the same: Adjusts all processor frequencies to the highest common frequency. Does not generate an error – this is not an error condition. Continues to boot the system successfully. If the frequencies for all processors cannot be adjusted to be the same: Logs the POST error code into the SEL. Alerts the BMC to set the system status LED to steady amber. Does not disable the processor. Displays 0197: Processor speeds unable to synchronize message in the error manager. Takes fatal error action (see above) and does not boot until the fault condition is remedied
Processor Intel® UPI link frequencies not identical	Fatal	If the link frequencies for all Intel® Ultra Path Interconnect (Intel® UPI) links can be adjusted to be the same: Adjusts all Intel UPI interconnect link frequencies to highest common frequency. Does not generate an error – this is not an error condition. Continues to boot the system successfully. If the link frequencies for all Intel UPI links cannot be adjusted to be the same: Logs the POST error code into the SEL. Alerts the BMC to set the system status LED to steady amber. Does not disable the processor. Displays 0195: Processor Intel(R) UPII link frequencies unable to synchronize message in the error manager. Takes fatal error action (see above) and does not boot until the fault condition is remedied.
Processor microcode update failed	Major	Logs the POST error code into the SEL. Displays 816x: Processor 0x unable to apply microcode update message in the error manager or on the screen. Takes major error action. The system may continue to boot in a degraded state, depending on the “POST Error Pause” setting in setup, or may halt with the POST error code in the error manager waiting for operator intervention.
Processor microcode update missing	Minor	Logs the POST error code into the SEL. Displays 818x: Processor 0x microcode update not found message in the error manager or on the screen. The system continues to boot in a degraded state, regardless of the “POST Error Pause” setting in setup.

3.5 Intel® Xeon® processor Scalable Family with Integrated Intel® Omni-Path Host Fabric Interface

Intel® Omni-Path Fabric is Intel's new Host Fabric Interconnect (HFI). The Intel® Xeon® processor Scalable Family includes SKUs with an integrated Intel® Omni-Path Host Fabric Interface (HFI) connector.

Table 15. Intel® Xeon® Processor Scalable Family w/ Integrated Intel Omni-Path Fabric – Features Table

Feature	81xxF Platinum	61xxF Gold
# of Cores	≥ 24	< 24
# of Omni-Path Fabric Ports	1	1
# of Intel® UPI Links	2	2
UPI Speed	10.4 GT/s	10.4 GT/s
Supported Topologies	2S-2UPI	2S-2UPI
Node Controller Support	No	No
# of Memory Channels	6	6
Max DDR4 Speed	2666	2666
Memory Capacity	768GB 1.5TB (Select SKUs)	768GB 1.5TB (Select SKUs)
RAS Capability	Standard	Standard
Intel® Turbo Boost	Yes	Yes
Intel® Hyper-Threading	Yes	Yes
Intel® AVX-512 ISA Support	Yes	Yes
Intel® AVX-512 - # of 512b FMA Units	2	2
# of PCIe* Lanes	48	48

The current fabric port count is one fabric port per processor socket. Each Omni-Path port supports four lanes of 25Gbps, providing 100Gbps of bandwidth in a single direction.




Figure 51. Processor Host Fabric Interface

Fabric processor support is a Multi-Chip Package (MCP) option, where the processor Host Fabric Interface (HFI) connector is cabled to an Intel Fabric Through (IFT) carrier board installed into either Riser Card 1 or Riser Card 2. Note, the IFT carrier board has no PCIe interface. When installed into a PCIe add-in slot, the card only draws power and SMBUS signals from the PCIe connector. A second cable carrying Omni-path side band signals is connected between the IFT carrier board and sideband connectors on the server board. External cables attach the IFT carrier board to an external Omni-Path Switch.

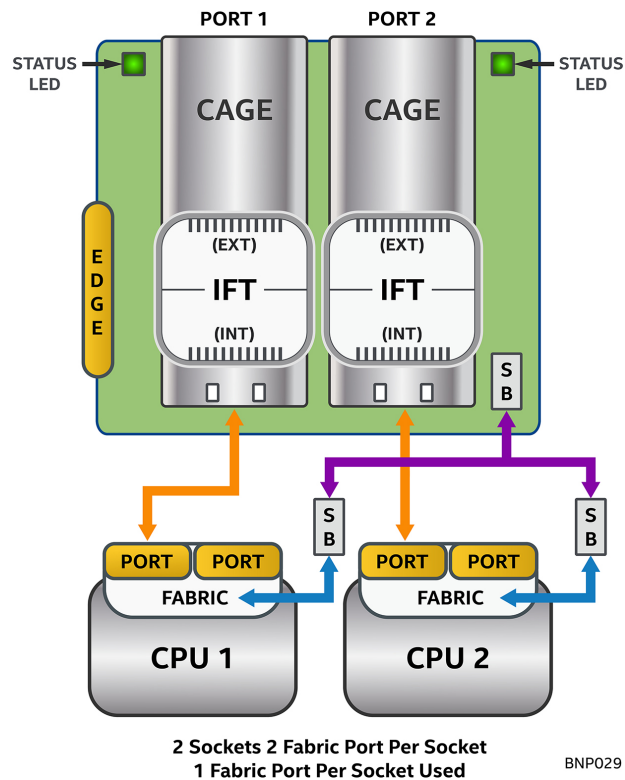


Figure 52. Intel® Omni-Path Fabric Interconnect

The following figure represent two supported configurations for Fabric Processors and Fabric Kit Interconnection, Riser 1 and Riser 2. Refer to the *Intel® Server Board S2600BP and Intel® Compute Module Service and Integration Guide* for a complete set of detailed instructions.




Figure 53. Fabric Interconnect Configurations

All necessary components to support up to two fabric processors are included in an orderable Fabric Accessory Kit (**AHWBPFABKIT**). Refer to the *Intel® S2600BP Server Board Product Family Configuration Guide* for additional information.




Figure 54. Intel Accessory Kit AHWBPFABKIT

Sideband “Y” Cable: Connects from the IFT carrier board to each fabric processor sideband connector on the server board.




Figure 55. Server Board Fabric Side Band Connectors

Each IFT carrier port has one green Status LED.




Figure 56. IFT Carrier Board – Rear View

Table 16. IFT Carrier LED Functionality

LED Color	LED State	Description
Green	OFF	No Link
	Blinking at Slow Rate	Link established but not activated by management
	ON	Link activated by management; but no traffic is present
	Blinking at Steady Rate	Traffic is present

For external connection, the IFT carrier will include two QSFP+28 style connectors. The signal definition of these QSFP+28 style connectors consists of the high speed diff pairs, miscellaneous side band signals, and 3.3V power. The 3.3V power is used for the active logic within the QSFP+ modules. QSFP+ modules have four power classes which control how much power the active logic in the cable can consume as noted in Table 17.

Table 17. Power Level Classification for QSFP+ Modules

Power Level Class	Max Power (W)
1	1.5
2	2.0
3	2.5
4	3.5

Table 18. Fabric Carrier Thermal

Front drive configuration	System ambient (C)	Max processor Power	Estimated QSFP cable Power support	
			Riser 1 slot	Riser 2 slot
24x2.5 HDD	27	135W	1.5W	1.5W
24x2.5 HDD	35	135W	Passive	Passive
16x2.5 HDD with riser 1 air duct	27	135W	2.5W	2.5W
16x2.5 HDD with riser 1 air duct	35	135W	1.5W	1.5W
16x2.5 HDD with riser 1 air duct	27	145W	2.2W	2.2W

Combining fabric processors with non-fabric processors on the Intel® Server Board S2600BP is allowed. In dual processor configurations, with at least one processor having support for Intel® OP HFI, the following population rules apply:

- The base SKU number of both processor types **must** be the same:
 - Example) Intel® Xeon® Platinum **8160F** (Intel® OP HFI) + Intel® Xeon® Platinum **8160** (non-fabric)
 - Example) Intel® Xeon® Gold **6140F** (Intel® OP HFI) + Intel® Xeon® Gold **6140F** (Intel® OP HFI)

There is no restriction on which processor socket is populated with the fabric processor and which processor socket is populated with the matching non-fabric processor.

WARNING: The FPGA Processor is not supported on the Intel® Server Board S2600BP product family. Attempting to install a FPGA processor onto this server board family will incur damage to the server board, the processor, or both.

Supported configurations are listed in Table 19.

Table 19. Supported Processor Mixing – Fabric vs Non-Fabric Processors

CPU Socket 1	CPU Socket 2	Platform Expected Behavior
Processor	Processor	Boot to OS
Processor	Fabric Processor	Boot to OS
Fabric Processor	Processor	Boot to OS
Fabric Processor	Fabric Processor	Boot to OS

4. Memory Support

This chapter describes the architecture that drives the memory subsystem, supported memory types, memory population rules, and supported memory RAS features.

4.1 Memory Subsystem Architecture




Figure 57. Integrated Memory Controller Functional Block Diagram

NOTE: Intel® Server Board S2600BP supports DDR4 DIMMs only.

The Intel® Server Board S2600BP has support for up to 16 DDR4 DIMMs. Each installed processor supports 6 memory channels via two Integrated Memory Controllers (IMC). Each memory channel is assigned an identifier letter A thru F. On the Intel® Server Board S2600BP, channels A and D support two DIMM slots each and channels B, C, E and F support one DIMM slot each.

The server board supports the following:

- Only DDR4 DIMMs are supported
- Only Error Correction Code (ECC) enabled RDIMMs or LRDIMMs are supported
- Registered DIMMs (RDIMMs), Load Reduced DIMMs (LRDIMMs), and NVDIMMs (Non-Volatile Dual Inline Memory Module):
- Only RDIMMs and LRDIMMs with integrated Thermal Sensor On Die (TSOD) are supported
- DIMM sizes of 4 GB, 8 GB, 16 GB, 32 GB, 64 GB and 128 GB depending on ranks and technology
- Maximum supported DIMM speeds will be dependent on the processor SKU installed in the system.
 - Intel® Xeon® Platinum 81xx processor – Max. 2666 MT/s (Mega Transfers / second)
 - Intel® Xeon® Gold 61xx processor – Max. 2666 MT/s
 - Intel® Xeon® Gold 51xx processor – Max. 2400 MT/s
 - Intel® Xeon® Silver processor – Max. 2400 MT/s
 - Intel® Xeon® Bronze processor – Max. 2133 MT/s
- DIMMs organized as Single Rank (SR), Dual Rank (DR), or Quad Rank (QR)
 - RDIMMS – Registered DIMMS – SR/DR/QR, ECC only
 - LRDIMMs – Load Reduced DIMMs – QR only, ECC only
 - Maximum of 8 logical ranks per channel
 - Maximum of 10 physical ranks loaded on a channel

4.2 Supported Memory

The following table lists the DIMM support guidelines

Table 20. DDR4 DIMM Support Guidelines

Type	Ranks Per DIMM and Data Width	DIMM Capacity (GB)		Speed (MT/s); Voltage (V); Slots per Channel (SPC) & DIMMs per Channel (DPC)		
				1 Slot per Channel	2 Slots per Channel	
		DRAM Density		1DPC	1DPC	2DPC
		4Gb	8Gb	1.2V	1.2V	1.2V
RDIMM	SRx4	8GB	16GB	2666	2666	2666
RDIMM	SRx8	4GB	8GB			
RDIMM	DRx8	8GB	16GB			
RDIMM	DRx4	16GB	32GB			
RDIMM 3DS	QRx4	N/A	2H-64GB			
	8Rx4	N/A	4H-128GB			
LRDIMM	QRx4	32GB	64GB			
LRDIMM 3DS	QRx4	N/A	2H-64GB			
	8Rx4	N/A	4H-128GB			

4.3 Memory Slot Identification and Population Rules

NOTE: Although mixed DIMM configurations are supported, Intel performs platform validation only on systems that are configured with identical DIMMs installed.

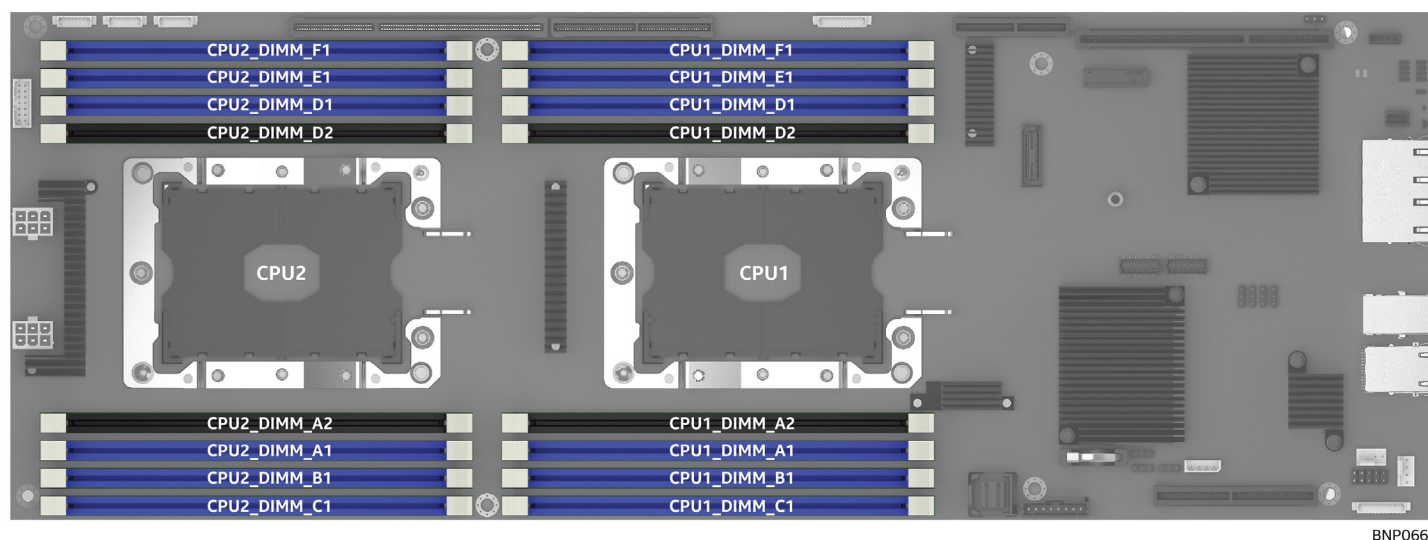
- Each installed processor provides six channels of memory. Memory channels from each processor are identified as Channels A – F.
 - On the server board, each DIMM slot is labeled by CPU #, memory channel, and slot # as shown in the following examples: CPU1_DIMM_A2; CPU2_DIMM_A2
- The Intel® Server Board S2600BP uses a “2-1-1” DIMM slot configuration. For memory channels A and D that include two DIMM slots, the Blue DIMM slot 1 must be populated before the Black DIMM slot 2
- For multiple DIMMs per channel:
 - For RDIMM, LRDIMM, 3DS RDIMM, or 3DS LRDIMM, always populate DIMMs with higher electrical loading in the first slot of a channel (blue slot) followed by the second slot.
- A maximum of 8 logical ranks can be used on any one channel, as well as a maximum of 10 physical ranks loaded on a channel.
- In order to install 3 QR LRDIMMs on the same channel, they must be operated with Rank Multiplication as RM = 2, this will make each LRDIMM appear as a DR DIMM with ranks twice as large.
- A processor can be installed without populating the associated memory slots, so long as a second processor is installed along with its associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.

- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS and Error Management) in the BIOS setup is applied commonly across processor sockets.
- DIMM types (RDIMM, LRDIMM) must not be mixed within or across processor sockets, all DIMMs must be DDR4 DIMMs. Mixed configuration will result in a Fatal Error Halt in Memory initialization.
- Mixing DIMMs of different frequencies and latencies is not supported within or across processor sockets. If a mixed configuration is encountered, the BIOS will attempt to operate at the highest common frequency and the lowest latency possible.
- If DIMMs of different sizes and number of ranks are mixed, although unsupported, the BIOS memory initialization algorithm will make a best-effort attempt to establish settings that will allow all installed DIMMs to operate together.

Table 21. DIMM Nomenclature

Processor Socket 1					
(0) Channel A	(1) Channel B	(2) Channel C	(3) Channel D	(4) Channel E	(5) Channel F
A2	B1	C1	D2	E1	F1
A1			D1		
Processor Socket 2					
(0) Channel A	(1) Channel B	(2) Channel C	(3) Channel D	(4) Channel E	(5) Channel F
A2	B1	C1	D2	E1	F1
A1			D1		

Figure 58 shows the physical location of each DIMM slot.




BNP066

Figure 58. Intel® Server Board S2600BP Product Family DIMM Slot Layout

4.3.1 DIMM Population Guidelines for Best Performance

Processors within the Intel® Xeon® processor Scalable family include two integrated memory controllers (IMC), each supporting three memory channels.



For best performance, DIMMs should be populated using the following guidelines:

- Each installed processor should have matching DIMM configurations
- The following DIMM population guidelines should be followed for each installed processor
 - **1 DIMM to 3 DIMM Configurations** – DIMMs should be populated to DIMM Slot 1 (**Blue Slots**) of Channels **A** thru **C**
 - **4 DIMM Configurations** – DIMMs should be populated to DIMM Slot 1 (**Blue Slots**) of Channels **A**, **B**, **D**, and **E**
 - **5 DIMM Configurations** – NOT Recommended. This is an unbalanced configuration which will yield less than optimal performance
 - **6 DIMM Configurations** – DIMMs should be populated to DIMM Slot 1 (**Blue Slots**) of all Channels
 - **7 DIMM Configurations** – NOT Recommended. This is an unbalanced configuration which will yield less than optimal performance
 - **8 DIMM Configurations** – DIMMs should be populated to DIMM Slots 1 and 2 of Channels **A**, **B**, **C**, **D**, **E** and **F**.

4.4 Memory RASM Support

Supported memory RAS features are dependent on the processor SKU installed. Each processor SKU within the Intel® Xeon® processor Scalable family has support for either Standard or Advanced memory RAS features as defined in the following table.

Table 22. Memory RASM Features

RASM Feature	Description	Standard	Advanced
Device Data Correction	x8 Single Device Data Correction (SDDC) via static virtual lock-step (Applicable to x8 DRAM DIMMs)	✓	✓
	Adaptive Data Correction (SR) (Applicable to x4 DRAM DIMMs)	✓	✓
	x8 Single Device Data Correction + 1 bit (SDDC+1) (Applicable to x8 DRAM DIMMs)		✓
	SDDDC + 1, and ADDDC (MR) + 1 (Applicable to x4 DRAM DIMMs)		✓
DDR4 Command/Address Parity Check and Retry	DDR4 Command/Address Parity Check and Retry: Is a DDR4 technology based CMD/ADDR parity check and retry with following attributes: <ul style="list-style-type: none"> CMD/ADDR Parity error "address" logging CMD/ADDR Retry 	✓	✓
DDR4 Write Data CRC Protection	DDR4 Write Data CRC Protection detects DDR4 data bus faults during write operation.	✓	✓
Memory Demand and Patrol Scrubbing	Demand scrubbing is the ability to write corrected data back to the memory once a correctable error is detected on a read transaction. Patrol scrubbing proactively searches the system memory, repairing correctable errors. Prevents accumulation of singlebit errors.	✓	✓
Memory Mirroring	Full Memory Mirroring: An intra IMC method of keeping a duplicate (secondary or mirrored) copy of the contents of memory as a redundant backup for use if the primary memory fails. The mirrored copy of the memory is stored in memory of the same processor socket's IMC. Dynamic (without reboot) failover to the mirrored DIMMs is transparent to the OS and applications.	✓	✓
	Address Range/Partial Memory Mirroring: Provides further intra socket granularity to mirroring of memory by allowing the firmware or OS to determine a range of memory addresses to be mirrored, leaving the rest of the memory in the socket in non-mirror mode.		✓
Sparing <ul style="list-style-type: none"> Rank Level Memory Sparing Multi-rank Level Memory Sparing 	Dynamic fail-over of failing Ranks to spare Ranks behind the same memory controller DDR ranks.	✓	✓
	With Multi Rank up to two ranks out of a maximum of eight ranks can be assigned as spare ranks.	✓	✓
iMC's Corrupt Data Containment	Corrupt Data Containment is a process of signaling error along with the detected UC data. iMC's patrol scrubber and sparing engine have the ability to poison the UC data.	✓	✓
Failed DIMM Isolation	Ability to identify a specific failing DIMM thereby enabling the user to replace only the failed DIMM(s). In case of uncorrected error and lockstep mode, only DIMM-pair level isolation granularity is supported.	✓	✓
Memory Disable and Map Out for FRB	Allows memory initialization and booting to OS even when memory fault occurs.	✓	✓

RASM Feature	Description	Standard	Advanced
Post Package Repair	Starting with DDR4 technology there is an additional capability available known as PPR (Post Package Repair). PPR offers additional spare capacity within the DDR4 DRAM that can be used to replace faulty cell areas detected during system boot time.	✓	✓
Note: RAS Features may not be supported on all SKUs of a processor type			

4.4.1 DIMM Populations Rules and BIOS Setup for Memory RAS

- Memory Sparing and Memory Mirroring options are enabled in <F2> BIOS Setup
- RAS Modes Rank Sparing, and Mirroring are mutually exclusive in this BIOS. Only one operating mode may be selected, and it will be applied to the entire system.
- If a RAS Mode has been enabled, and the memory configuration is not able to support it during boot, the system will fall back to Independent Channel Mode and log and display errors.
- Rank Sparing Mode is only possible when all channels that are populated with memory that meet the requirement of having at least 2 SR or DR DIMMs installed, or at least one QR DIMM installed, on each populated channel.
- Memory Mirroring Mode requires that for any channel pair that is populated with memory, the memory population on both channels of the pair must be identically sized

Refer to the *Intel® Server System BIOS External Product Specification* for additional details regarding to memory sizing, memory publishing and memory initialization.

5. PCIe* Support

The Integrated I/O (IIO) module of the Intel® Xeon® processor Scalable Family provides the PCI express interface for general purpose PCI Express* (PCIe) devices at up to PCIe* 3.0 speeds.

The IIO module provides the following PCIe Features:

- Compliant with the PCI Express* Base Specification, Revision 2.0 and Revision 3.0
- 2.5 GHz (Gen1) and 5 GHz (Gen2) and 8 GHz (Gen3)
- x16 PCI Express* 3.0 interface supports up to four x4 controllers and is configurable to 4x4 links, 2x8, 2x4\1x8, or 1x16
- x8 PCI Express* 3.0 interface supports up to two x4 controllers and is configurable to 2x4 or 1x8
- Full peer-to-peer support between PCI Express* interfaces
- Full support for software-initiated PCI Express* power management

The following table provide the PCIe* port routing information from each processor:

Table 23. CPU – PCIe* Port Routing

CPU	Port	IOU	Width	Connection
CPU1	DMI3	IOU2	x4	Chipset
CPU1	PE1	IOU2	x4	Intel® X550
CPU1	PE1	IOU2	X8	Chipset Uplink (on S2600BPS model)
CPU1	PE2	IOU0	x16	<ul style="list-style-type: none"> • S2600BPB/S: Riser Slot 1 • S2600BPQ: Chipset
CPU1	PE3	IOU1	x16	Riser Slot 2 Lane Reversal
CPU1	PE1	IOU2	X8	Riser Slot 2 Lane Reversal
CPU2	DMI3	IOU2	x4	Unused
CPU2	PE1	IOU2	x16	Riser 3
CPU2	PE2	IOU0	x16	Riser4 Lane Reversal
CPU2	PE3	IOU1	x8	Riser 3 Lane Reversal

Chipset PCH PCI Express

Chipset PCH supports 20 PCIe Gen 3 Downstream Ports, these ports are MUXed with various other High Speed I/O Ports to minimize pin count and package size.

Table 24. Chipset PCH PCIe* Connectivity

PCH	Device	Lane	Electrical Width	Mode
Chipset	Mini SAS HD, Shared with SATA1-4	12-15	x4	Gen 3t
Chipset	M.2 Shared with SSATA2	8-11	x4	Gen 3
Chipset	Aspeed AST 2500 BMC Video	5	X1	Gen 3

The following table lists the High-Speed I/O Port Mapping (HSIO).

Table 25. High-Speed I/O Port Mapping (HSIO)

Port	Lane	Usage	USB	PCIE	SATA	QAT	LAN
1	0	USB 3.0 Rear	USB3_1				
2	1	USB 3.0 Rear	USB3_2				
3	2		USB3_3				
4	3		USB3_4				
5	4		USB3_5				
6	5		USB3_6				
7	6		USB3_7	PCIE_0	0		
8	7		USB3_8	PCIE_1	1		
9	8		USB3_9	PCIE_2	2		
10	9		USB3_10	PCIE_3	3		GbE
11	10			PCIE_4	4		GbE
12	11	BMC		PCIE_5	5		GbE
13	12	Bridge Board		PCIE_6	SSATA_0		
14	13			PCIE_7	SSATA_1		
15	14	M.2		PCIE_8	SSATA_2		GbE
16	15	M.2		PCIE_9	SSATA_3		
17	16	M.2		PCIE_10	SSATA_4		
18	17	M.2		PCIE_11	SSATA_5		GbE
19	18	MiniSAS HD		PCIE_12	SATA_0	PCIE_UP_0	
20	19	MiniSAS HD		PCIE_13	SATA_1	PCIE_UP_1	
21	20	MiniSAS HD		PCIE_14	SATA_2	PCIE_UP_2	
22	21	MiniSAS HD		PCIE_15	SATA_3	PCIE_UP_3	
23	22	Bridge Board		PCIE_16	SATA_4	PCIE_UP_4	
24	23	Bridge Board		PCIE_17	SATA_5	PCIE_UP_5	
25	24	Bridge Board		PCIE_18	SATA_6	PCIE_UP_6	
26	25	Bridge Board		PCIE_19	SATA_7	PCIE_UP_7	

NOTES:

- Default port usage indicated in Bold.
- PCIe* ports cannot be bifurcated across Quad boundaries [3:0], [7:4], [11:8], [15:12], and [19:16]
- Ports may be bifurcated within their own Quad (4x1, 2x2, 1x2 + 2x1, 1 x4)
For example: A PCIe x2 port cannot come from PCIe[4:3] since this crosses a Quad boundary
- PCH configuration 2x1 + 1x2 (i.e. X1, x1, x2) is not allowed
For example: PCIe_4 = x1, PCIe_5 = x1, PCIe[7:6] = x2.

5.1 PCIe* Enumeration and Allocation

The BIOS assigns PCI bus numbers in a depth-first hierarchy, in accordance with the PCI Local Bus Specification, Revision 3.0. The bus number is incremented when the BIOS encounters a PCI-PCI bridge device. Scanning continues on the secondary side of the bridge until all subordinate buses are assigned numbers. PCI bus number assignments may vary from boot to boot with varying presence of PCI devices with PCI-PCI bridges.

If a bridge device with a single bus behind it is inserted into a PCI bus, all subsequent PCI bus numbers below the current bus are increased by one. The bus assignments occur once, early in the BIOS boot process, and never change during the pre-boot phase.

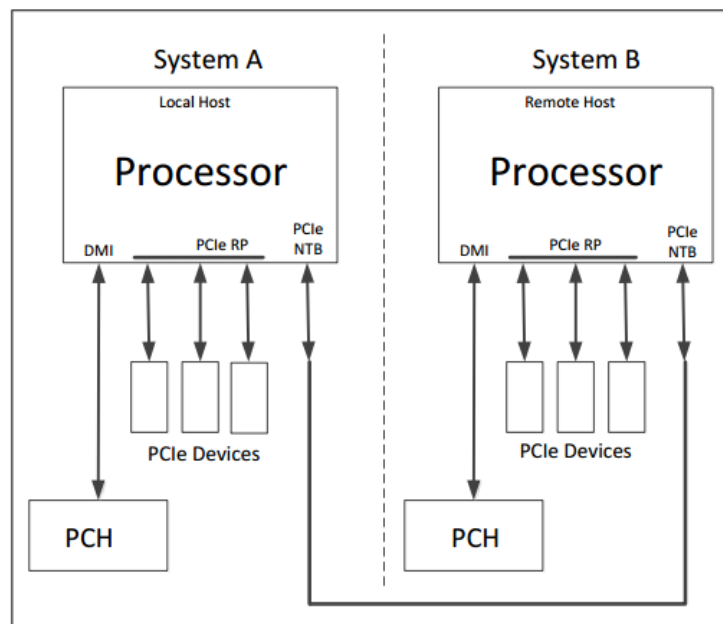
The BIOS resource manager assigns the PIC-mode interrupt for the devices that are accessed by the legacy code. The BIOS ensures that the PCI BAR registers and the command registers for all devices are correctly set up to match the behavior of the legacy BIOS after booting to a legacy OS. Legacy code cannot make any assumption about the scan order of devices or the order in which resources are allocated to them. The BIOS automatically assigns IRQs to devices in the system for legacy compatibility. A method is not provided to manually configure the IRQs for devices.

See Section 6.2 Add-in Card Support for PCIe* Add-in card support and configurations.

5.2 PCIe* Non-Transparent Bridge (NTB)

The PCI Express Non-Transparent Bridge (NTB) acts as a gateway that enables high performance, low latency communication between two PCIe Hierarchies, such as a local and remote system. The NTB allows a local processor to independently configure and control the local system and provides isolation of the local Host memory domain from the remote Host memory domain, while enabling status and data exchange between the two domains. The NTB is discovered by the local processor as a RootComplex Integrated Endpoint (RCiEP).

The figure below shows two systems which are connected through an NTB. Each system is a completely independent PCIe Hierarchy. The width of the NT Link can be x16, x8, or x4 at the expense of other PCIe Root Ports. Only Port A can be configured as an NT Port.



The specified processor family supports the following NTB features.

The NTB only supports one configuration/connection model:

- NT Port attached to another NT Port of the same component type and generation
- The NTB provides Direct Address Translation between the two PCIe Hierarchies through two separate regions in Memory Space. Accesses targeting these Memory addresses are allowed to pass through the NTB to the remote system. This mechanism enables the following transactions flows through the NTB:
 - Both Posted Mem Writes and Non-Posted Mem Read transactions across the NTB
 - Peer-to-Peer Mem Read and Write transactions to and from the NTB

In addition, the NTB provides the ability to interrupt a processor in the remote system through a set of Doorbell registers. A write to a Doorbell register in the local side of the NTB will generate an interrupt to the remote processor. Since the NTB is designed to be symmetric, the converse is also true.

For additional information, refer to the Processor Family External Design Specification (EDS).

6. Server Board I/O

The server board input/output features are provided via the embedded features and functions of several onboard components including: Intel® Omni-path Fabric for the Intel® Xeon® Processor Scalable Family, the Intel® C620 Chipset Series, the Intel® Ethernet controller X550, and the I/O controllers embedded within the Aspeed® AST2500® Management Controller. See the Product Architecture Overview in Chapter 2 for features and interconnects of each of the major sub-system components.

6.1 Quick Assist Technology

Intel® QuickAssist Technology (QAT) provides security and compression acceleration capabilities used to improve performance and efficiency across the data center. The QAT feature is supported on the Intel® Server Board S2600BPQ model only. For more information about Intel® QuickAssist Technology, visit the following website:

<http://www.intel.com/content/www/us/en/embedded/technology/quickassist/overview.html>

6.2 PCIe® Riser Card Support

The server board includes several PCIe® riser slot connectors allowing OEMs and other system integrators to develop custom PCIe® expansion options. These are identified on the server board as:

- Slot_1 (PCIe_X16_CPU_1)
- Slot_2 (PCIe_X24_CPU_1)
- Slot_3 (PCIe_X24_CPU_2)
- Slot_4 (PCIe_X16_CPU_2)




Figure 59. Server Board Riser Slots

Note: PCIe® add-in cards cannot be installed directly into any of the onboard PCIe riser slot connectors. The onboard PCIe riser slot connectors are designed to support PCIe riser cards or other custom PCIe interface cards only. Attempting to install a PCIe add-in card directly into any of the onboard PCIe riser slot connectors may critically damage the PCIe add-in card, the PCIe riser slot connector, or both.

The following figures identify the PCIe bus architecture supporting each of the onboard PCIe riser slot connectors for each server board option.

Figure 60. Intel® Server Board S2600BPB & BPS – PCIe* Slot Connector Architectural Block Diagram

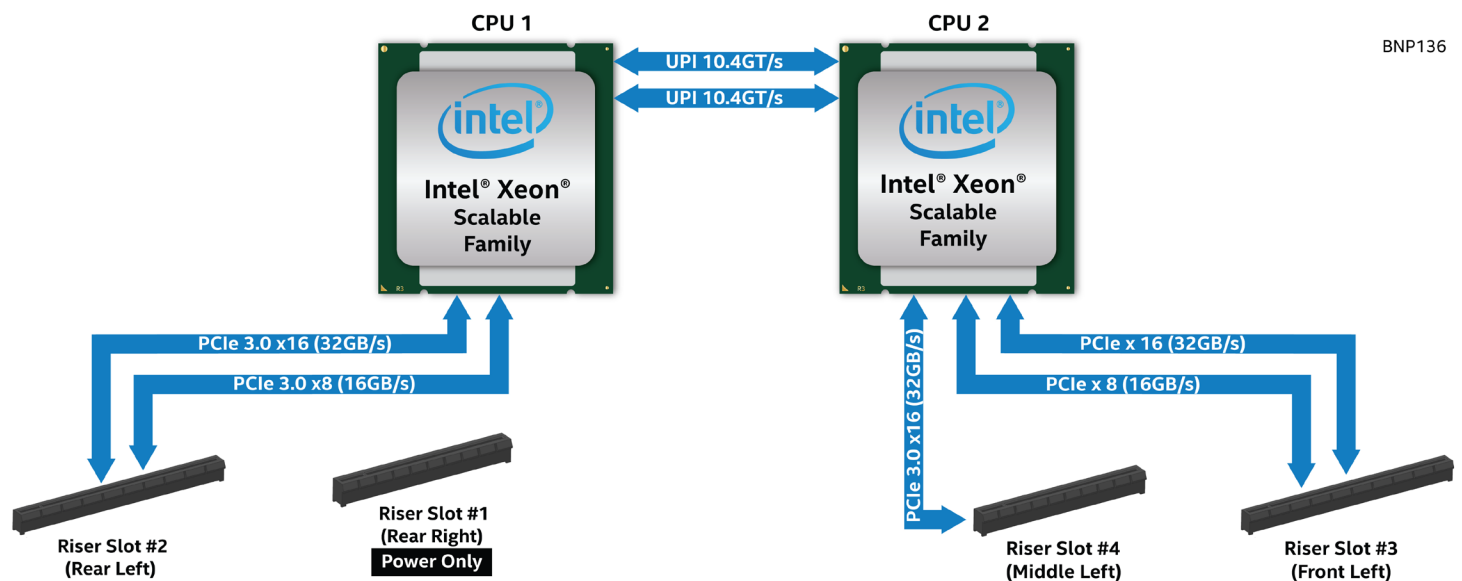
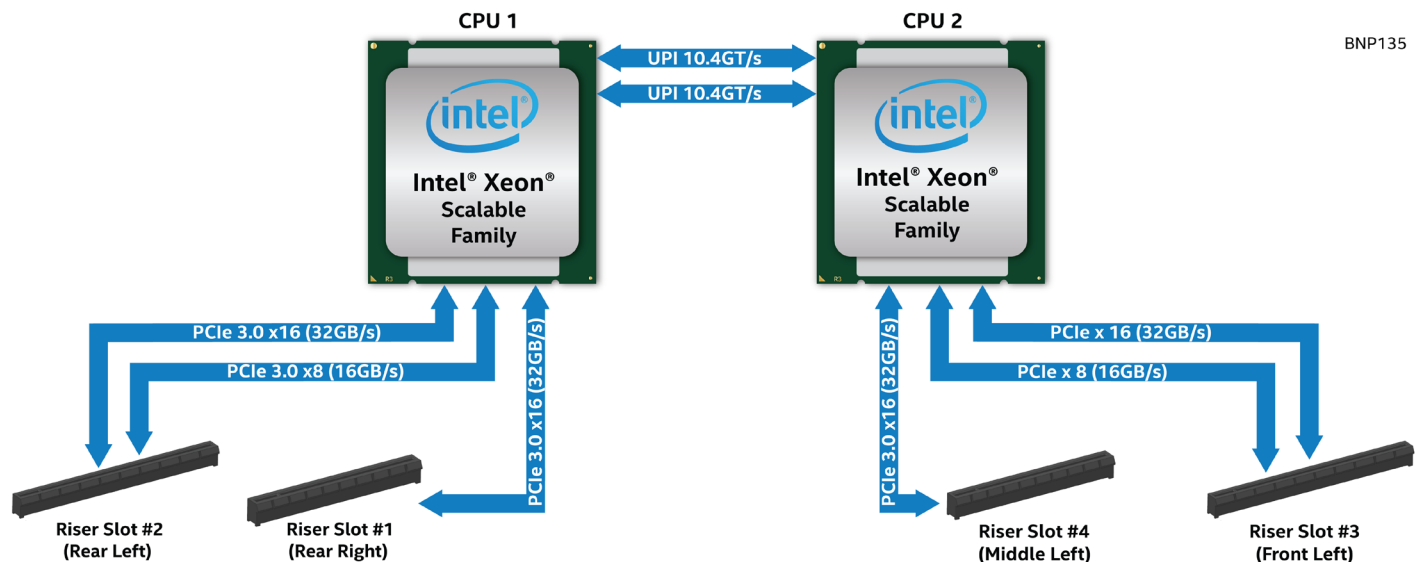


Figure 61. Intel® Server Board S2600BPQ – PCIe* Slot Connector Architectural Block Diagram

See section 7.3.1 for the electrical pinout definition for each PCIe slot connector.

See Table 23. CPU – PCIe* Port Routing

PCIe* Slot support notes:

- Riser Slot #1 on the Intel® Server Board S2600BPQ has NO PCIe bus lanes routed to it. It can only be used to provide power to add-in card options. See section 7.3.1 for connector pin-out definition
- In a 1U chassis or 1U mounting plate implementation, Riser Slot #1 on the Intel® Server Board **S2600BPS** **cannot** support a PCIe* add-in card due to an interference with components on the server board. In a 1U design implementation, this riser slot is compatible with the optional Intel® Omni-path fabric through carrier card available within the Intel® Omni-Path Fabric Processor accessory kit (iPC - AHWBPFABKIT).
- Riser Slots #3 & #4 are not accessible with an optional Intel Bridge Board installed

6.2.1 Compute Module – Riser Card Support

See section 2.5.3 for riser card support details.

6.3 Onboard Storage Sub-System

The Intel® Server S2600BP product family has support for a variety of storage controller and storage device options including:

- 1 – M.2 PCIe*/SATA SSD – Server board feature
- 1 – M.2 PCIe* - Intel compute module feature
- Embedded SATA support
- Intel® Rapid Storage Technology 5.0 (Intel® RSTe) for SATA
- Intel® Embedded Server RAID Technology 2 v1.60 (Intel® ESRT2 1.60) for SATA
- Intel® Volume Management Device (Intel® VMD) for NVMe
- Intel® Virtual RAID on CPU (Intel® VROC) for NVMe

The following sections provide an overview of each option.

6.3.1 M.2 SSD Support**6.3.1.1 Server Board M.2 Support**

The server board includes one (1) onboard M.2 SSD connector capable of supporting a PCIe* or SATA SSD that conforms to a 2260 (42mm) form factor. The sSATA controller embedded within the Intel chipset is provides this connector with the SATA port. X4 PCIe* lanes are routed from the Intel chipset and can be supported in single processor configurations. Circuitry within the server board will auto detect the type of device installed into the M.2 connector.




Figure 62. On-board M.2 SSD Support

6.3.1.2 Intel Compute Module SSD Support

Intel compute modules have the option to support up to two M.2 SSD drives: One on the server board as described in section 6.3.1.1, and one located on the backside of Riser Card 2.

The M.2 connector located on the riser card is capable of supporting PCIe* M.2 SSDs that conform to the 2280 (80mm) form factor. X4 PCIe lanes to the M.2 connector are routed through the PCI Riser slot and are supported from CPU #1.




Figure 63. Riser Card #2 M.2 SSD Support

6.3.1.3 Embedded RAID Support

RAID support from embedded RAID options for the M.2 SSDs is defined as follows:

- The onboard M.2 connector has no embedded RAID support
 - When creating a RAID volume using either of the embedded RAID options, all devices within the volume must be driven from a common SATA controller. On this server board family, there are no other SATA ports used from the chipset sSATA controller beyond the onboard M.2 connector, therefore the onboard M.2 connector has no embedded RAID support when configured with a SATA device on this server board.
 - NVMe RAID support using Intel® RSTe VROC requires that the PCIe bus lanes be routed directly from the CPU. Therefore, since the PCIe lanes are routed from the chipset, the onboard M.2 connector cannot support Intel® RSTe RAID.
 - Intel® ESRT2 RAID has no support for PCIe* NVMe drives
- Embedded RAID support for M.2 SSD on PCIe Riser Slot #2 is as follows
 - Intel® RSTe RAID is supported when using a PCIe* SSD in the M.2 connector on Riser Card #2 along with other NVMe SSDs configured when using the appropriate bridge board.
 - Intel® ESRT2 RAID has no support for PCIe* NVMe drives

Note: Storage devices used to create a single RAID volume created using either RSTe or ESRT2, cannot span across the two embedded SATA controllers nor is mixing both SATA and NVMe devices within a single RAID volume supported.

- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*

6.3.2 Embedded Serial ATA (SATA) Support

The Intel chipset includes two embedded AHCI SATA controllers identified as **SATA** and **sSATA** (“s” is for secondary). On the Intel® Server Board S2600BP, these controllers provide the following SATA support:

The AHCI **SATA** controller provides support for up to eight (8) 6 Gb/sec Serial ATA (SATA3) ports:

- Four ports routed to one Mini-SAS HD (SFF-8643) connector on the server board
 - Intel® **S2600BPB** & **S2600BPS** server boards only
- Four ports routed to the bridge board connector for use with select bridge board options when used within an Intel chassis.

The AHCI **sSATA** controller provides up to two (2) 6 Gb/sec Serial ATA (SATA3) ports:

- One port routed to the on-board M.2 SSD connector. See section 6.3.1.1
- One port routed to the on-board bridge connector (OEM use only)

The following table lists supported features of the SATA and sSATA controllers.

Table 26. SATA and sSATA Controller Feature Support

Feature	Description	AHCI / RAID Disabled	AHCI / RAID Enabled
Native Command Queuing (NCQ)	Allows the device to reorder commands for more efficient data transfers	N/A	Supported
Auto Activate for DMA	Collapses a DMA Setup then DMA Activate sequence into a DMA Setup only	N/A	Supported
Hot Plug Support	Allows for device detection without power being applied and ability to connect and disconnect devices without prior notification to the system	N/A	Supported
Asynchronous Signal Recovery	Provides a recovery from a loss of signal or establishing communication after hot plug	N/A	Supported
6 Gb/s Transfer Rate	Capable of data transfers up to 6 Gb/s	Supported	Supported
ATAPI Asynchronous Notification	A mechanism for a device to send a notification to the host that the device requires attention	N/A	Supported
Host & Link Initiated Power Management	Capability for the host controller or device to request Partial and Slumber interface power states	N/A	Supported
Staggered Spin-Up	Enables the host the ability to spin up hard drives sequentially to prevent power load problems on boot	Supported	Supported
Command Completion Coalescing	Reduces interrupt and completion overhead by allowing a specified number of commands to complete and then generating an interrupt to process the commands		N/A

The SATA controller (AHCI Capable Controller 1) and the sSATA controller (AHCI Capable Controller 2) can be independently enabled, disabled, and configured through the <F2> BIOS Setup Utility under the “Mass Storage Controller Configuration” menu screen as shown in Table 27.

Table 27. SATA and sSATA Controller BIOS Utility Setup Options

SATA Controller	sSATA Controller	Supported
AHCI	AHCI	Yes
AHCI	Enhanced	Yes
AHCI	Disabled	Yes
AHCI	RSTe	Yes
AHCI	ESRT2	Microsoft* Windows Only
Enhanced	AHCI	Yes
Enhanced	Enhanced	Yes
Enhanced	Disabled	Yes
Enhanced	RSTe	Yes
Enhanced	ESRT2	Yes
Disabled	AHCI	Yes
Disabled	Enhanced	Yes
Disabled	Disabled	Yes
Disabled	RSTe	Yes
Disabled	ESRT2	Yes
RSTe	AHCI	Yes
RSTe	Enhanced	Yes
RSTe	Disabled	Yes
RSTe	RSTe	Yes
RSTe	ESRT2	No
ESRT2	AHCI	Microsoft* Windows Only
ESRT2	Enhanced	Yes
ESRT2	Disabled	Yes
ESRT2	RSTe	No
ESRT2	ESRT2	Yes

6.3.2.1 Staggered Disk Spin-Up

Because of the high density of disk drives that can be attached to the onboard AHCI SATA controller and the sSATA controller, the combined startup power demand surge for all drives at once can be much higher than the normal running power requirements and could require a much larger power supply for startup than for normal operations.

In order to mitigate this and lessen the peak power demand during system startup, both the AHCI SATA controller and the sSATA controller implement a Staggered Spin-Up capability for the attached drives. This means that the drives are started up separately, with a certain delay between disk drives starting.

For the onboard SATA controller, Staggered Spin-Up is an option – AHCI HDD Staggered Spin-Up – in the Setup Mass Storage Controller Configuration screen found in the <F2> BIOS Setup Utility.

6.3.3 Embedded SATA RAID Support

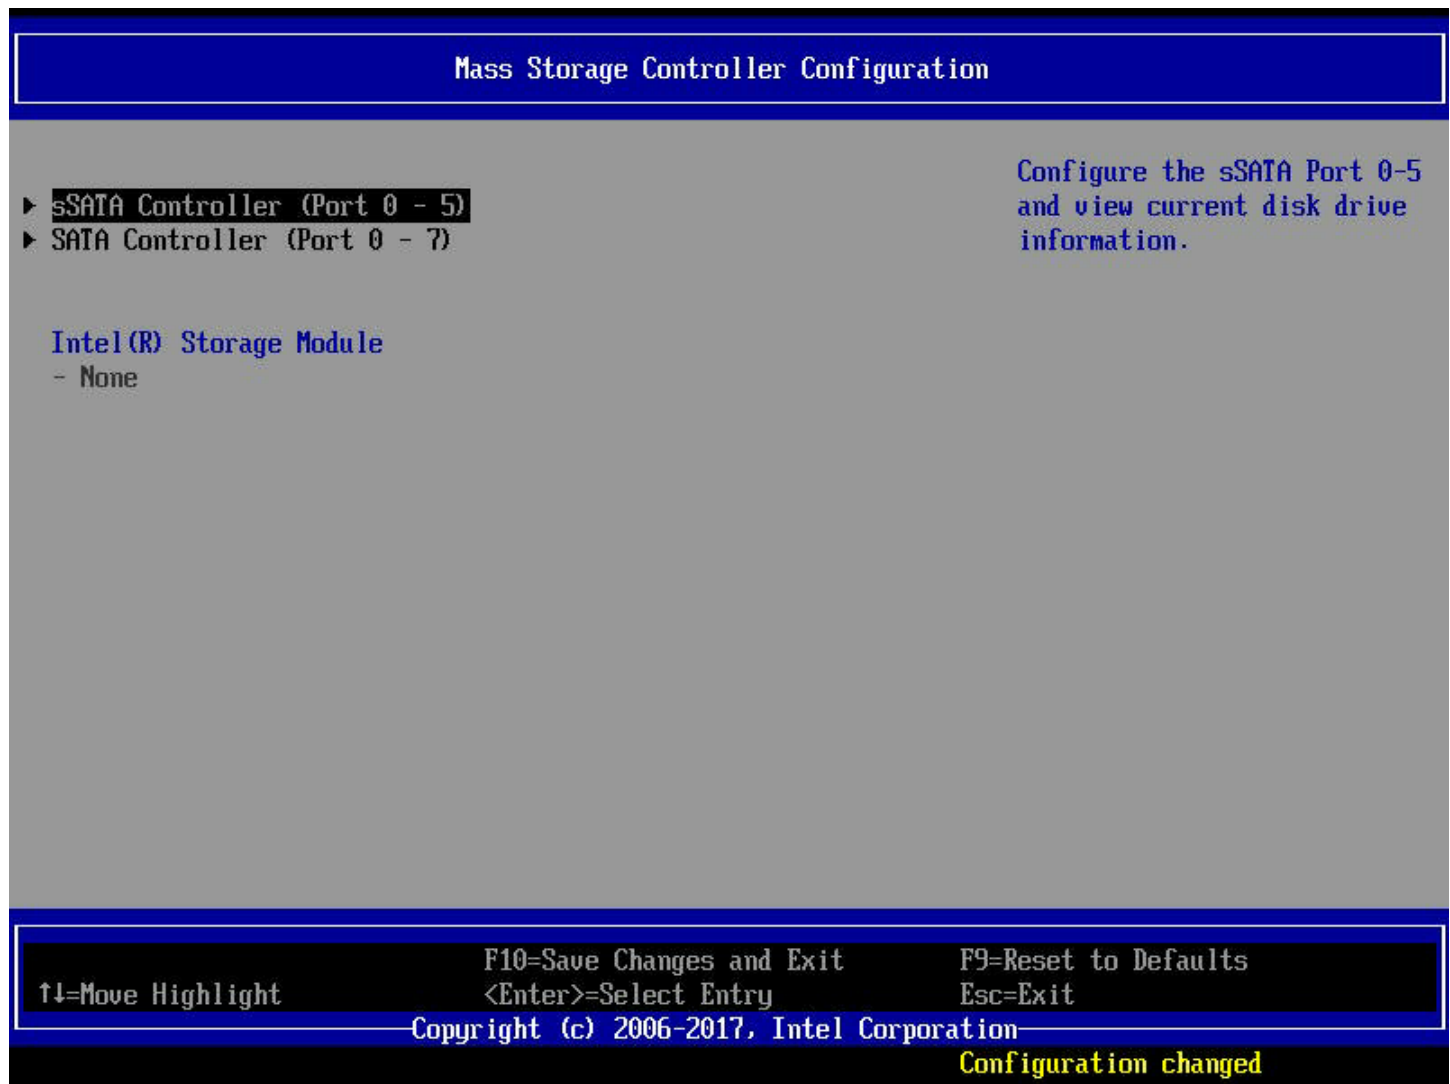
The Intel® Server Board has embedded support for two SATA RAID options:

- Intel® Rapid Storage Technology (RSTe) 5.0
- Intel® Embedded Server RAID Technology 2 (ESRT2) based on LSI* MegaRAID technology

By default, on-board RAID options are set to DISABLED in <F2> BIOS Setup. To enable on-board RAID support, access the <F2> BIOS setup utility during system POST.

The on-board RAID options can be found under the following <F2> BIOS Setup menu options:

ADVANCED > MASS STORAGE CONTROLLER CONFIGURATION >



NOTE: RAID partitions created using either RSTe or ESRT2 cannot span across the two embedded SATA controllers. Only drives attached to a common SATA controller can be included in a RAID partition.

6.3.3.1 Intel® Rapid Storage Technology (RSTe) 5.0 for SATA

Intel® Rapid Storage Technology offers several diverse options for RAID (Redundant Array of Independent Disks) to meet the needs of the end user. AHCI support provides higher performance and alleviates disk bottlenecks by taking advantage of the independent DMA engines that each SATA port offers in the chipset.

- **RAID 0:** Uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.
- **RAID 1:** Uses mirroring so that data written to one disk drive simultaneously writes to another disk drive. This is good for small databases or other applications that require small capacity but complete data redundancy.
- **RAID 5:** Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.
- **RAID 10:** A combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy but uses a larger number of spans.

By using Intel® RSTe, there is no loss of PCI resources (request/grant pair) or add-in card slot. Intel® RSTe functionality requires the following:

- The embedded RAID option must be enabled in <F2> BIOS Setup.
- Intel® RSTe option must be selected in <F2> BIOS Setup.
- Intel® RSTe drivers must be loaded for the installed operating system.
- At least two SATA drives needed to support RAID levels 0 or 1.
- At least three SATA drives needed to support RAID level 5.
- At least four SATA drives needed to support RAID level 10.
- Intel® RSTe does not support mixing of NVMe SSDs and SATA drives within a single RAID volume

With Intel® RSTe SW-RAID enabled, the following features are made available:

- A boot-time, pre-operating-system environment, text-mode user interface that allows the user to manage the RAID configuration on the system. Its feature set is kept simple to keep size to a minimum, but allows the user to create and delete RAID volumes and select recovery options when problems occur. The user interface can be accessed by pressing the <CTRL-I> keys during system POST.
- Provides boot support when using a RAID volume as a boot disk. It does this by providing Int13 services when a RAID volume needs to be accessed by MS-DOS applications (such as NTLDR) and by exporting the RAID volumes to the System BIOS for selection in the boot order.
- At each boot-up, provides the user with a status of the RAID volumes.

6.3.3.2 Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 for SATA

Intel® Embedded Server RAID Technology II (Intel® ESRT2) (Powered by LSI*) is a driver-based RAID solution for SATA that is compatible with previous generation Intel® server RAID solutions. Intel® ESRT2 provides RAID levels 0, 1, and 10, with an optional RAID 5 capability depending on whether a RAID Upgrade Key is installed or not.

Note: The embedded Intel® ESRT2 RAID option has no RAID support for PCIe* NVMe SSDs

Features of ESRT2 include the following:

- Based on LSI* MegaRAID Software Stack
- Software RAID with system providing memory and CPU utilization
- **RAID 0:** Uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.

- **RAID 1:** Uses mirroring so that data written to one disk drive simultaneously writes to another disk drive. This is good for small databases or other applications that require small capacity but complete data redundancy
- **RAID 10:** A combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy but uses a larger number of spans.
- Optional support for **RAID Level 5:**
 - Enabled with the addition of an optionally installed ESRT2 SATA RAID 5 Upgrade Key (**iPN – RKSATA4R5**)
 - RAID 5: Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.




Figure 64. SATA RAID 5 Upgrade Key

- Maximum drive support = 8 (Maximum on-board SATA port support)
- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*

NOTES:

1. RAID partitions created using either RSTe or ESRT2 cannot span across the two embedded SATA controllers
 2. Only drives attached to a common SATA controller can be included in a RAID partition
 3. RAID configurations cannot span across the two embedded AHCI SATA controllers
 4. ESRT2 SW architecture supports only continuous SATA ports starting at port 0
 6. UEFI boot mode requires to be set in order to enable ESRT2
 7. ESRT2 SWRAID is supported only on 8 port SATA controller on non-QAT server boards
 8. The embedded Intel® ESRT2 RAID option has no RAID support for PCIe* NVMe SSDs
-

6.3.4 Intel® Volume Management Device (Intel® VMD) for NVMe

Intel® Volume Management Device (Intel® VMD) is hardware logic inside the processor Root Complex to help manage PCIe* NVMe SSDs. It provides robust **Hot Plug** support and **Status LED** management. This allows servicing of storage system NVMe SSD media without fear of system crashes or hangs when ejecting or inserting NVMe SSD devices on the PCIe* bus.




Figure 65. Intel® VMD

Intel® VMD handles the physical management of NVMe storage devices as a standalone function but can be enhanced when Intel® VROC support options are enabled to implement RAID based storage systems. See Section 6.3.2.

- Hardware is integrated inside the processor PCIe* root complex.
- Maps entire PCIe* trees into its own address space (a domain)
- Each domain manages x16 PCIe* lanes
- Can be enable/disable in <F2> BIOS SETUP at x4 lane granularity
- Driver sets up/manages the domain (enumerate, event/error handling), but out of fast I/O Path
- May load an additional child device driver that is Intel VMD aware
- Hot Plug support - Hot insert array of PCIe* SSDs
- Support for PCIe* SSDs and Switches only (No NICs, graphics cards, etc...)
- Max 128 PCIe* bus numbers per domain
- Support for MCTP over SMBus only
- MMIO only (no port-mapped I/O)
- Does not have support for NTB, Quick Data Tech, Omni-path, SR-IOV
- Correctable errors will not bring system down
- Intel® VMD will only manage devices on PCIe* lanes routed directly from the processor. Intel® VMD cannot provide device management on PCI lanes routed from the chipset (PCH).
- When Intel® VMD is enabled, the BIOS will not enumerate devices that are behind Intel VMD. The Intel® VMD-enabled driver is responsible for enumerating these devices and exposing them to the host
- Intel® VMD must be enabled for the Hot Plug feature to properly work on PCIe* SSDs.
- When Intel® VMD is disabled, the hot plug feature may or may not work with PCIe*SSD drives.

Intel® VMD provided with Intel VROC supports the following features for PCIe* CPU connected to PCIe* NVMe SSDs and PCIe* Switch devices:

- LED Management (VMD Method of LED Management)
- Error Handling
- Surprise HotPlug

Each CPU has 3 VMD domains. Each VMD domain manages x16 lanes. Intel VMD can be turned on/off on x4 lane granularity and supports either NVMe SSD device or PCIe* switch device.

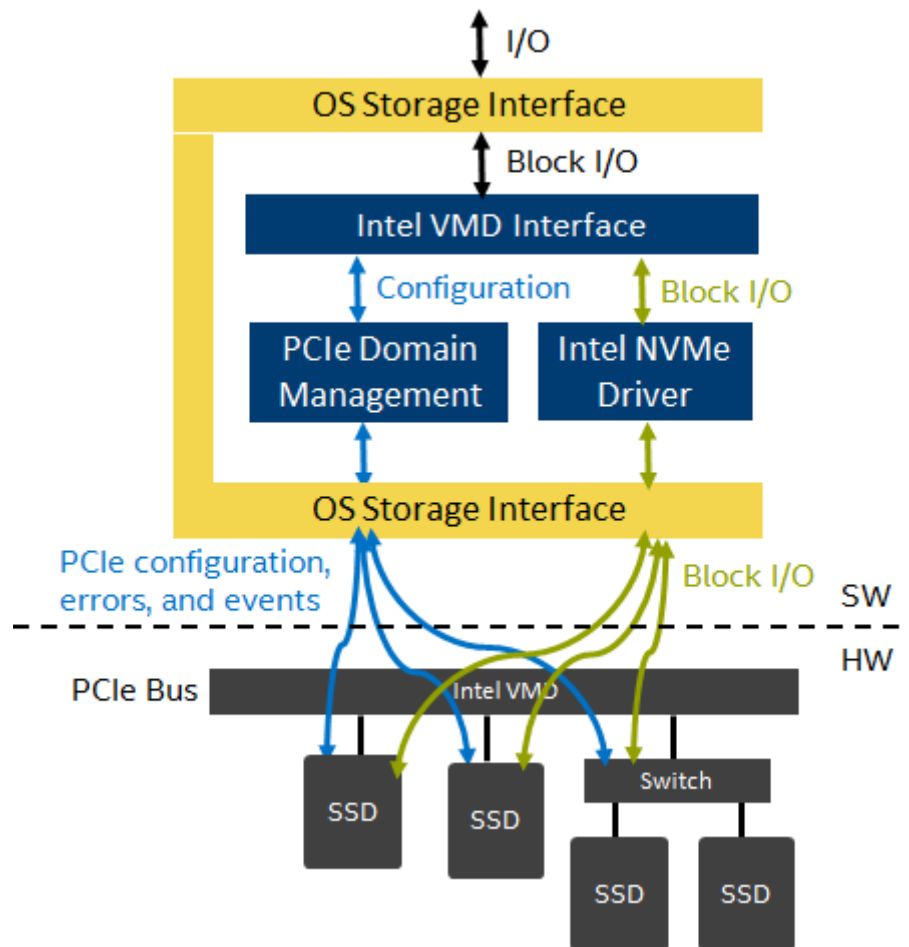


Figure 66. Intel® VMD Interface

6.3.4.1 Enabling Intel® VMD support

In order for installed NVMe devices to utilize the VMD features of the server board, VMD must be **ENABLED** on the appropriate CPU PCIe* Root Ports in <F2> BIOS Setup. By default, VMD support is **DISABLED** on all CPU PCIe* root ports in <F2> BIOS Setup.

NOTE: The Intel® Server Chassis H2000G supports up to two PCIe* SSDs per compute module on a 24 x 2.5” drive chassis. PCIe root ports supporting the NVMe drives are supported by CPU2 Ports 1B and 1D (IOU1).

In <F2> BIOS Setup, the VMD support menu can be found under the following BIOS Setup menu options:

ADVANCED -> PCI CONFIGURATION -> VOLUME MANAGEMENT DEVICE




Figure 67. VMD Support Disabled in <F2> BIOS Setup



Figure 68. VMD Support Enabled in <F2> BIOS Setup

6.3.5 Intel® Virtual RAID on CPU (Intel® VROC) for NVMe

Intel® VROC enables NVMe boot on RAID and volume management (Intel® RSTe 5.0 + Intel® VMD)



- I/O processor w/controller (ROC) and DRAM
- No need for battery back-up / RMFBU
- Protected Write Back Cache – SW and HW that will allow recovery from a double fault
- Isolate storage devices from OS – error handling
- Protect R5 data from OS crash or BSOD
- Boot on NVMe RAID Volumes within a single Intel VMD Domain
- NVMe Hot Plug and Surprise Removal on CPU PCIe* lanes
- LED Management for CPU PCIe* attached storage
- RAID / Storage management using RESTful APIs
- GUI for Linux
- 4K native NVMe SSD support

6.3.5.1 Optional Intel® VROC Activation Key for Intel NVMe SSDs and 3rd Party NVMe Support

Enabling Intel® VROC support requires installation of an optional upgrade key on to the server board as shown in the following illustration.




Figure 69. Intel® VROC Upgrade Key

The following table identifies the Intel® Server Board S2600BP product family supported Intel® VROC upgrade key option.

Table 28. Intel® VROC Upgrade Key Option

	Major features	iPC VROCSTANMOD
		Standard Intel® VROC
NVMe RAID	CPU attached NVMe – high perf.	X
	Boot on RAID Volume	X
	3rd Party vendor SSD support	X
	RSTe 5.0 RAID 0/1/10	X
	RSTe 5.0 RAID 5	-
	RAID Write Hole closed (BBU replacement)	-
	Hot Plug/ Surprise Removal (2.5" SSD form factor only; AIC not supported)	X
	Enclosure LED management	X

Note: Intel® VROC Upgrade Key referenced in Table 24 is used for PCIe* NVMe SSDs only. For SATA RAID support, see Section 6.3.2

Intel® VROC Standard Key: "VROCSTANMOD"

The Activation Key is a small PCB board that has a security EEPROM that is read by the Intel VROC UEFI driver to enable different versions of the Intel VROC software stack to be loaded when VMD is enabled.

The Intel® Server Board S2600BP with Intel® VROC supports the following configuration:

1. No VROC Upgrade Key – Only Intel NVMe SSDs will be enumerated and exposed to the platform BIOS. There is no HII interface and no RAID support in this configuration
2. VROC Standard Upgrade Key - Selected 3rd party NVMe SSD will be enumerated and exposed to the BIOS. RAID HII will be enabled in the BIOS set to support RAID 0/1/10 management.

Table 29. Supported Intel VROC Activation Key

MM#	Item Name	Item Description	Configuration
951605	VROCSTANMOD	Upgrade Module – Standard	Intel VROC RAID 0/1/10

6.4 Video Support

Note: By default, all Intel compute models within the product family will have an add-in VGA video connector bracket installed on the back of the compute module within the Riser #1 add-in card location. To use the add-in card slot on Riser #1 for any other purpose, the VGA video connector bracket must be removed, thus losing video support from the compute module.

The graphics controller of the ASpeed® AST2500 BMC is a VGA-compliant controller with 2D hardware acceleration and full bus master support. With 16MB of memory reserved, the video controller can support the following resolutions:

Table 30. Onboard Video Resolution and Refresh Rate (Hz)

2D Mode	2D Video Mode Support (Color Bit)			
Resolution	8 bpp	16 bpp	24 bpp	32 bpp
640x480	60, 72, 75, 85	60, 72, 75, 85	Not supported	60, 72, 75, 85
800x600	60, 72, 75, 85	60, 72, 75, 85	Not supported	60, 72, 75, 85
1024x768	60, 70, 75, 85	60, 70, 75, 85	Not supported	60, 70, 75, 85
1152x864	75	75	75	75
1280x800	60	60	60	60
1280x1024	60	60	60	60
1440x900	60	60	60	60
1600x1200	60	60	Not Supported	Not Supported
1680x1050	60	60	Not Supported	Not Supported
1920x1080	60	60	Not Supported	Not Supported
1920x1200	60	60	Not Supported	Not Supported

For system configurations that require an add-in video adapter, <F2> BIOS setup includes options to enable /disable the on-board video controller.

6.5 Universal Serial Bus (USB) Ports

The server board provides support for up to 4 USB 3.0/2.0 ports. The USB port distribution is as follows:

- Two external USB 2.0 / USB 3.0 ports located on the back edge of server board
- One internal USB 2.0 port 5-pin header for optional front-panel USB port support (OEM use only)
- One internal USB 2.0 port via the server board Bridge Board (OEM use only).



Figure 70. USB Ports Block Diagram

6.6 Serial Port

The server board has support for one Serial-A port via an internal DH-10 header as shown below.




Figure 71. Serial Port A Location

6.7 Network Interface

On the back edge of the server board are located one RJ45 Dedicated Management Port and two networking ports identified as Port 1 and Port 2.

6.7.1 Network ports 1 and 2 on the Intel® Server Board S2600BPB and S2600BPQ

Two RJ45 networking ports on the back of the server board are supported by an onboard 10GbE Intel® X550 Networking controller which is a dual-port, compact component with two fully-integrated 10GbE Media Access Control (MAC) and Physical Layer (PHY) ports.




Figure 72. Network Interface Connectors S2600BPB, S2600BPQ (Port 1 and Port 2 – RJ45)

The Intel® X550 LAN Controller provides the server board with support for dual LAN ports designed for 10Gb/s, 1Gb/s, and 100Mb/s operation. Refer to the *Intel® X550 Gigabit Ethernet Controller Datasheet* for full details of the NIC feature set.

RJ45 connectors used for the Dedicated Management Port and Network Interface connectors include two LEDs. The LED on the left side of the connector is the link/activity LED and indicates network connection when on, and transmit/receive activity when blinking. The LED on the right side of the connector indicates link speed as defined in the following table.



Figure 73. RJ45 Connector LEDs

Table 31. External RJ45 NIC Port LED Definition

LED	Color	LED State	NIC State
Left	Green	Off	LAN link not established
		On	LAN link is established
		Blinking	LAN activity is occurring
Right	N/A	Off	100 Mbit/sec data rate is selected
	Amber	On	1 Gbit/sec data rate is selected.
	Green	On	10 Gbit/sec data rate is selected

6.7.2 Network ports 1 and 2 on the Intel® Server Board S2600BPS

Two SFP+ networking ports on the back of the server board are supported by the Intel® C622 chipset embedded Intel® Ethernet Controller X722.




Figure 74. Network Interface Connectors S2600BPS (Port 1 and Port 2 – SFP+)

6.7.3 Server Board Management Networking

The BMC implements both the IPMI 1.5 and IPMI 2.0 messaging models. These provide out-of-band local area network (LAN) communication between the BMC and the network.

The BMC supports three RMII/RGMII ports that can be used for communicating with Ethernet devices. Two ports are used for communication with the on-board NICs and one is used for communication with an Ethernet PHY located on the server board.

6.7.3.1 On-board NICs

On server boards that include the onboard Intel® X550 NIC, NIC Port 1 and Port 2 can be used by the BMC firmware to send management traffic in standby. The LAN controller will be used in conjunction with BMC for out of band Management traffic. The BMC will communicate with the LAN device over Port B NC-SI interface (RMII physical). This provides a sideband high-speed connection for manageability traffic to the BMC while still allowing for a simultaneous host access to the OS if desired. The LAN controller will be on Standby power so that the BMC can send management traffic over the NC-SI interface to the network during sleep states S4 and S5. When on Standby, the link can drop to 100Mbit.

6.7.3.2 Dedicated Management Port

An additional LAN channel dedicated to BMC usage and not available to host SW is supported through an extra PHY on the server board. The PHY on the board connects to the BMC's other RMII/RGMII interface. This BMC port is configured for RGMII usage.

6.7.4 MAC Address Definition

The Intel® Server Board S2600BP products have the following four MAC addresses assigned to it at the Intel factory:

- NIC 1 MAC address (for OS usage)
- NIC 2 MAC address = NIC 1 MAC address + 1 (Server Management & WOL)
- BMC LAN Channel 1 MAC address = NIC 1 MAC address + 2
- BMC LAN Channel 2 MAC address = NIC 1 MAC address + 3

- BMC LAN Channel 3 (Intel® Dedicated Management NIC (Intel® DNM)) MAC address = NIC 1 MAC address + 4

The BMC queries a server board NIC over the NC-SI for a host MAC address and then derives up to four MAC addresses to allocate for manageability usage for the platform. The BMC FW assigns the MAC addresses to specific LAN ports according to which ones are enabled for manageability.

The Intel® Server Board S2600BP has a white MAC address sticker included with the board. The sticker displays the NIC1 Port1 MAC address in both bar code and alphanumeric formats.

7. Server Board Connectors and Headers

7.1 Power Connectors


7.1.1 Main Power Connectors

To facilitate custom OEM designs that require a cabled power harness from a power supply, the server board supports two 2x3-pin Minifit Jr* connectors, which can be used to deliver 12amps per pin or 60+Amps total.

Note: No 240VA protective circuits exist on the server board, power supply, or power distribution boards.

Table 32. Main Power Supply Connector (6-pin 2x3 Connector)

Pin	Signal Name	Pin	Signal Name
1	GND	4	+12V
2	GND	5	+12V
3	GND	6	+12V



7.1.2 Backup Power Control Connector

A 1x8-pin backup (or auxiliary) power connector on the server board is used for power control when used in a non-Intel Chassis. Connector pins 7 or 8 (defined in the table below) should be used to provide the server board with stand-by power. The connector is capable of delivering up to 3 amps. Connector type is an AMPMODU MTE Interconnection System or equivalent.





Table 33. Backup Power Control Connector

Pin	Signal Name
1	SMB_PMBUS_CLK
2	SMB_PMBUS_DAT
3	IRQ_PMBUS_ALERT_N
4	GND
5	PWROK
6	PSON_N
7	5V STBY
8	12V STBY

7.2 System Management Headers



7.2.1 Intel® Remote Management Module 4 (Intel® RMM4) Lite Connector

A 7-pin Intel® RMM4 Lite connector (J2A1) is included on the server board to support advanced management features.

Table 34. Intel® RMM4 Lite Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	P3V3_AUX	2	SPI_RMM4_LITE_DI
3	Key Pin	4	SPI_RMM4_LITE_CLK
5	SPI_RMM4_LITE_DO	6	GND
7	SPI_RMM4_LITE_CS_N	8	GND

7.2.2 IPMB Header

Table 35. IPMB Header 4-pin (J6B1)

Pin	Signal Name	Description
1	SMB_IPMB_5VSB_DAT	BMC IPMB 5V standby data line
2	GND	Ground
3	SMB_IPMB_5VSB_CLK	BMC IPMB 5V standby clock line
4	P5V_STBY	+5V standby power

7.2.3 Front Control Panel Connector

The server board includes a 2x6-pin Control Panel connector for use in non-Intel chassis. It is designed to use either discrete 2-pin connectors (similar to the ATX chassis) or a ribbon cable.

Table 36. Control Panel Connector

Pin	Signal Name	Pin	Signal Name
1	FP_ID_LED_N	2	FP_P5V_AUX_0
3	FP_HD_ACT_LED_N	4	FP_P5V_AUX_1
5	FP_PWR_LED_N	6	FP_P5V_AUX_2
7	GND	8	FP_PWR_BTN_N
9	GND	10	FP_ID_BTN_N
11	FP_RST_BTN_N	12	Key

7.2.4 Bridge Board Connector

The server board includes a bridge board connector that delivers SATA signals, Disk backplane management signals, and BMC SMBUS's, as well as front-panel and miscellaneous Node-specific signals.

Table 37. Bridge Board Connector

Pin	Signal Name	Pin	Signal Name
80	FM_SEL_SAS_N_SATA	79	GND
78	GND	77	GND
76	SATA6G_S0_RX_C_DP	75	SATA6G_P0_TX_DN
74	SATA6G_S0_RX_C_DN	73	SATA6G_P0_TX_DP
72	GND	71	GND
70	SATA6G_S1_TX_C_DP	69	SATA6G_P1_RX_DN
68	SATA6G_S1_TX_C_DN	67	SATA6G_P1_RX_DP

Pin	Signal Name	Pin	Signal Name
66	GND	65	GND
64	SATA6G_S2_RX_C_DP	63	SATA6G_P2_TX_DN
62	SATA6G_S2_RX_C_DN	61	SATA6G_P2_TX_DP
60	GND	59	GND
58	SATA6G_S3_TX_C_DP	57	SATA6G_P3_RX_DN
56	SATA6G_S3_TX_C_DN	55	SATA6G_P3_RX_DP
54	GND	53	GND
52	SGPIO_SATA_CLOCK_R2	51	PWRGD_PSU
50	BMC_NODE_ID1	49	SGPIO_SATA_LOAD
48	BMC_NODE_ID2	47	SGPIO_SATA_DATAOUT0
46	BMC_NODE_ID3	45	SGPIO_SATA_DATAOUT1
KEY			
44	BMC_NODE_ID4	43	PS_EN_PSU_N
42	SPA_COM_SIN_N	41	IRQ_SML1_PMBUS_ALERT_N
40	SPA_COM_OUT_N	39	GND
38	FP_NMI_BTN_N	37	SMB_PMBUS_CLK
36	FP_PWR_BTN_N	35	SMB_PMBUS_DATA
34	FP_RST_BTN_N	33	GND
32	FP_ID_BTN_N	31	SMB_HSBP_STBY_LVC3_CLK
30	FP_ID_LED_N	29	SMB_HSBP_STBY_LVC3_DATA
28	FP_PWR_LED_N	27	GND
26	FP_LED_STATUS_GREEN_N	25	SMB_CHAS_SENSOR_STBY_LVC3_CLK
24	FP_LED_STATUS_AMBER_N	23	SMB_CHAS_SENSOR_STBY_LVC3_DATA
22	FP_Activity_LED_N	21	GND
20	FP_HDD_ACT_LED_N	19	SMB_IPMB_5VSTBY_CLK
18	GND	17	SMB_IPMB_5VSTBY_DATA
16	USB2_04_BB_DN	15	GND
14	USB2_FP_DP	13	Spare
12	GND	11	FM_PS_ALL_NODE_OFF
10	SATA6G_S4_RX_C_DP	9	FM_NODE_PRESENT_N (GND)
8	SATA6G_S4_RX_C_DN	7	GND
6	GND	5	SATA6G_P4_TX_DP
4	FM_USB_OC_FP_N	3	SATA6G_P4_TX_DN
2	P5V Aux	1	P5V Aux

7.3 I/O Connectors

7.3.1 PCI Express* Slot Connectors

The Intel® Server Board S2600BP includes four PCI Express* Riser slots identified on the server board as follows:

- Slot_1 (PCIe_X16_CPU_1)
- Slot_2 (PCIe_X24_CPU_1)
- Slot_3 (PCIe_X24_CPU_2)
- Slot_4 (PCIe_X16_CPU_2)

Note: PCIe* add-in cards cannot be installed directly into any of the onboard PCIe riser slot connectors. The onboard PCIe riser slot connectors are designed to support PCIe riser cards or other custom PCIe interface cards only. Attempting to install a PCIe add-in card directly into any of the onboard PCIe riser slot connectors may critically damage the PCIe add-in card, the PCIe slot connector, or both.




Figure 75. Server Board Riser Slot Identification

7.3.1.1 PCIe* Riser Slot Connector 1 (PCIe_X16_CPU_1)

120PIN High Density Connector – 3M*

NOTE: Riser Slot 1 on the Intel® S2600BPQ Server Board is used as a power source only, no PCIe* lanes are routed to it.



Table 38. PCIe* Slot Connector 1 (PCIe_X16_CPU_1)

Pin	Signal Name	Pin	Signal Name
1	P12V	2	P12V
3	RST_RISER1_PERST_N	4	P12V
5	GND	6	P3V3_AUX
7	NC_CLK_100M_RISER1_PE2_DP	8	GND
9	NC_CLK_100M_RISER1_PE2_DN	10	PE_RX_DP<0>
11	GND	12	PE_RX_DN<0>
13	CLK_100M_RISER1_PE1_DP	14	GND
15	CLK_100M_RISER1_PE1_DN	16	NC_PE_RISER1_P16
17	GND	18	IRQ_LVC3_WAKE_N
19	PE_TX_DP<0>	20	RISER_INTERLOCK_N
21	PE_TX_DN<0>	22	GND
23	GND	24	PE_RX_DP<1>
25	PE_TX_DP<1>	26	PE_RX_DN<1>
27	PE_TX_DN<1>	28	GND
29	GND	30	PE_RX_DP<2>
31	PE_TX_DP<2>	32	PE_RX_DN<2>
33	PE_TX_DN<2>	34	GND
35	GND	36	PE_RX_DP<3>
37	PE_TX_DP<3>	38	PE_RX_DN<3>
39	PE_TX_DN<3>	40	GND
41	GND	42	PE_RX_DP<4>
43	PE_TX_DP<4>	44	PE_RX_DN<4>
45	PE_TX_DN<4>	46	GND
47	GND	48	PE_RX_DP<5>
49	PE_TX_DP<5>	50	PE_RX_DN<5>
51	PE_TX_DN<5>	52	GND
53	GND	54	PE_RX_DP<6>
55	PE_TX_DP<6>	56	PE_RX_DN<6>
57	PE_TX_DN<6>	58	GND
59	GND	60	PE_RX_DP<7>
61	NC_PE_RISER1_P61	62	PE_RX_DN<7>
63	FM_LINK_WIDTH_RISER1_ID0	64	GND
65	GND	66	FM_THROTTLE_RISER1_N
67	PE_TX_DP<7>	68	GND
69	PE_TX_DN<7>	70	PE_RX_DP<8>
71	GND	72	PE_RX_DN<8>
73	PE_TX_DP<8>	74	GND
75	PE_TX_DN<8>	76	PE_RX_DP<9>
77	GND	78	PE_RX_DN<9>
79	PE_TX_DP<9>	80	GND
81	PE_TX_DN<9>	82	PE_RX_DP<10>
83	GND	84	PE_RX_DN<10>
85	PE_TX_DP<10>	86	GND
87	PE_TX_DN<10>	88	PE_RX_DP<11>
89	GND	90	PE_RX_DN<11>
91	PE_TX_DP<11>	92	GND
93	PE_TX_DN<11>	94	PE_RX_DP<12>
95	GND	96	PE_RX_DN<12>
97	PE_TX_DP<12>	98	GND
99	PE_TX_DN<12>	100	PE_RX_DP<13>
101	GND	102	PE_RX_DN<13>
103	PE_TX_DP<13>	104	GND
105	PE_TX_DN<13>	106	PE_RX_DP<14>
107	GND	108	PE_RX_DN<14>
109	PE_TX_DP<14>	110	GND
111	PE_TX_DN<14>	112	PE_RX_DP<15>

Pin	Signal Name	Pin	Signal Name
113	GND	114	PE_RX_DN<15>
115	PE_TX_DP<15>	116	GND
117	PE_TX_DN<15>	118	SMB_PCI_RISER1_CLK
119	GND	120	SMB_PCI_RISER1_DATA

7.3.1.2 PCIe* Riser Slot Connector 2 (PCIe_X24_CPU_1)

Provide PCIe* x24 to Riser in order to support an x16 bus and an x8 bus, (x16 interface can be configured as a two x8 if required)

Uses a 200-pin High Speed Edge Connector 8mm (HSEC8) – Edgeline*




Table 39. PCIe* Slot Connector 2 (PCIe_X24_CPU_1)

Pin	Signal Name	Pin	Signal Name
1	12V	2	12V
3	12V	4	12V
5	GND	6	GND
7	GND	8	GND
9	3.3V	10	3.3V
11	3.3V	12	3.3V
13	GND	14	GND
15	3.3V AUX	16	5V AUX
17	GND	18	GND
19	SPARE	20	SPARE
21	SPARE	22	SPARE
23	SPARE	24	SPARE
25	Riser2 ID1	26	SPARE
27	GND	28	SPARE
29	SPARE	30	THROTTLE_N
31	Riser2 ID0	32	GND
33	GND	34	PERST#
35	SMCLK_R2M1	36	WAKE#
37	SMDATA_R2M1	38	GND
39	GND	40	PE3_CLK3+
41	PE1_R00-	42	PE3_CLK3-
43	PE1_R00+	44	GND
45	GND	46	PE3_CLK3+
47	PE1_R01-	48	PE3_CLK3-
49	PE1_R01+	50	GND
51	GND	52	PE1_T00-
53	PE1_R02-	54	PE1_T00+
55	PE1_R02+	56	GND
57	GND	58	PE1_T01-
59	PE1_R03-	60	PE1_T01+
61	PE1_R03+	62	GND
63	GND	64	GND
65	SPARE	66	GND
67	SPARE	68	PE1_T03-
69	GND	70	PE1_T03+
71	PE1_R04-	72	GND
73	PE1_R04+	74	PE1_T04-
75	GND	76	PE1_T04+
77	PE1_R05-	78	GND
79	PE1_R05+	80	PE1_T05-

Pin	Signal Name	Pin	Signal Name
101	GND	102	PE2_CLK1-
103	R00-	104	GND
105	R00+	106	T00-
107	GND	108	T00+
109	R01-	110	GND
111	R01+	112	T01-
113	GND	114	T01+
115	R02-	116	GND
117	R02+	118	T02-
119	GND	120	T02+
121	R03-	122	GND
123	R03+	124	T03-
125	GND	126	T03+
127	R04-	128	GND
129	R04+	130	T04-
131	GND	132	T04+
133	R05-	134	GND
135	R05+	136	T05-
137	GND	138	T05+
139	R06-	140	GND
141	R06+	142	T06-
143	GND	144	T06+
145	R07-	146	GND
147	R07+	148	T07-
149	GND	150	T07+
151	R08-	152	GND
153	R08+	154	T08-
155	GND	156	T08+
157	R09-	158	GND
159	R09+	160	T09-
161	GND	162	T09+
163	R10-	164	GND
165	R10+	166	T10-
167	GND	168	T10+
169	R11-	170	GND
171	R11+	172	T11-
173	GND	174	T11+
175	R12-	176	GND
177	R12+	178	T12-
179	GND	180	T12+

Pin	Signal Name	Pin	Signal Name
81	GND	82	PE1_T05+
83	PE1_R06-	84	GND
85	PE1_R06+	86	PE1_T06-
87	GND	88	PE1_T06+
89	PE1_R07-	90	GND
91	PE1_R07+	92	PE1_T07-
93	GND	94	PE1_T07+
95	PE2_CLK2+	96	GND
97	PE2_CLK2-	98	GND
99	GND	100	PE2_CLK1+

Pin	Signal Name	Pin	Signal Name
181	R13-	182	GND
183	R13+	184	T13-
185	GND	186	T13+
187	R14-	188	GND
189	R14+	190	T14-
191	GND	192	T14+
193	R15-	194	GND
195	R15+	196	T15-
197	GND	198	T15+
199	Spare	200	GND

7.3.1.3 PCIe* Riser Slot Connector 3 (PCIe_X24_CPU_2)

PCIe* x24 - Supports the following configurations through the use of Port Width ID Bits:

- One x16 + One x8 [OEM – for 1U and 2U Custom designed Risers]
- One x8 + four x4 [OEM – for 1U and 2U Custom designed Risers]
- Six x4 [Providing PCIe NVMe support for 24 Drive Configurations]
- This connector is placed in line with the bridge board and is meant to be used when support for outboard IO boards are required (including a double wide GPU board) and support of 24 Drive Config using the following two types of Bridge Boards:
 - Bridge Board (SAS Controller and PCIe* ReTimer): 24 Drive with 4 SAS and 2 NVMe per Node
 - Bridge Board (PCIe* ReTimer): 24 Drive with 6 NVMe per Node

200-pin High Speed Edge Connector 8mm (HSEC8) – Edgeline*



Table 40. PCIe* Slot Connector 3 (PCIe_X24_CPU_2)

Pin	Signal Name		Pin	Signal Name	
1	Riser ID0	0=3x8	2	GND	
3	GND		4	SAS_T07-	LSI 3008
5	SAS_R07-	LSI 3008	6	SAS_T07+	LSI 3008
7	SAS_R07+	LSI 3008	8	GND	
9	GND		10	SAS_T06-	LSI 3008
11	SAS_R06-	LSI 3008	12	SAS_T06+	LSI 3008
13	SAS_R06+	LSI 3008	14	GND	
15	GND		16	SAS_T05-	LSI 3008
17	SAS_R05-	LSI 3008	18	SAS_T05+	LSI 3008
19	SAS_R05+	LSI 3008	20	GND	
21	GND		22	SAS_T04-	LSI 3008
23	SAS_R04-	LSI 3008	24	SAS_T04+	LSI 3008
25	SAS_R04+	LSI 3008	26	GND	
27	GND		28	SAS_T03-	
29	SAS_R03-		30	SAS_T03+	

Pin	Signal Name		Pin	Signal Name	
31	SAS_R03+		32	GND	
33	GND		34	SAS_T02-	
35	SAS_R02-		36	SAS_T02+	
37	SAS_R02+		38	GND	
39	GND		40	SAS_T01-	
41	SAS_R01-		42	SAS_T01+	
43	SAS_R01+		44	GND	
45	GND		46	SAS_T00-	
47	SAS_R00-		48	SAS_T00+	
49	SAS_R00+		50	GND	
51	GND		52	T15-	
53	R15-		54	T15+	
55	R15+		56	GND	
57	GND		58	T14-	
59	R14-		60	T14+	
61	R14+		62	GND	
63	GND		64	SPARE	
KEY	KEY				
65	GND		66	GND	
67	R13-		68	T13-	
69	R13+		70	T13+	
71	GND		72	GND	
73	GND		74	T12-	
75	R12-		76	T12+	
77	R12+		78	GND	
79	GND		80	T11-	NVME Drive 2
81	R11-	NVME Drive 2	82	T11+	NVME Drive 2
83	R11+	NVME Drive 2	84	GND	
85	GND		86	T10-	NVME Drive 2
87	R10-	NVME Drive 2	88	T10+	NVME Drive 2
89	R10+	NVME Drive 2	90	GND	
91	GND		92	T09-	NVME Drive 2
93	R09-	NVME Drive 2	94	T09+	NVME Drive 2
95	R09+	NVME Drive 2	96	GND	
97	GND		98	T08-	NVME Drive 2
99	R08-	NVME Drive 2	100	T08+	NVME Drive 2
101	R08+	NVME Drive 2	102	GND	
103	GND		104	T07-	
105	R07-		106	T07+	
107	R07+		108	GND	
109	GND		110	T06-	
111	R06-		112	T06+	
113	R06+		114	GND	
115	GND		116	T05-	
117	R05-		118	T05+	
119	R05+		120	GND	
121	GND		122	T04-	
123	R04-		124	T04+	
125	R04+		126	GND	
127	GND		128	T03-	NVME Drive1
129	R03-	NVME Drive1	130	T03+	NVME Drive1
131	R03+	NVME Drive1	132	GND	
133	GND		134	T02-	NVME Drive1
135	R02-	NVME Drive1	136	T02+	NVME Drive1
137	R02+	NVME Drive1	138	GND	
139	GND		140	T01-	NVME Drive1
141	R01-	NVME Drive1	142	T01+	NVME Drive1

Pin	Signal Name		Pin	Signal Name	
143	R01+	NVME Drive1	144	GND	
145	GND		146	T00-	NVME Drive1
147	R00-	NVME Drive1	148	T00+	NVME Drive1
149	R00+	NVME Drive1	150	GND	
151	GND		152	PE3_CLK4-	PCIe Retimer
153	PE3_CLK3-	PCIe Clock Buffer	154	PE3_CLK4+	PCIe Retimer
155	PE3_CLK3+	PCIe Clock Buffer	156	GND	
157	GND		158	WAKE#	
159	CLK_SAS-	LSI 3008 PCIe Clock	160	PERST#	
161	CLK_SAS+	LSI 3008 PCIe Clock	162	GND	
163	GND		164	Spare	
165	PE_SCL	VPP reg data BCP	166	Spare	
167	PE_SDA	VPP reg Clock BCP	168	Spare	
169	GND		170	THROTTLE_N	new for PHI cards
171	Spare		172	GND	
173	GND		174	Spare	
175	Spare		176	Spare	
177	Spare		178	Spare	
179	Spare		180	Spare	
181	Spare		182	GND	
183	Spare		184	3.3V	
185	Riser ID1	ID1 00= 3 x8?	186	SMCLK_R2M2	
187	Spare		188	GND	
189	SMDAT_R2M2		190	Spare	
191	GND		192	Spare	
193	Spare		194	GND	
195	GND		196	12V	66W for GPU
197	12V	66W for GPU	198	12V	66W for GPU
199	12V		200	12V	20W 3.3V generate

7.3.1.4 PCIe* Slot Connector 4 (PCIe_X16_CPU_2)

- PCIe x16 from CPU 2
- 120-pin High Speed Edge Connector 8mm (HSEC8) – Edgeline*



Table 41. PCIe* x16 Riser Slot 4 Connector

Pin	Signal Name	Pin	Signal Name
120	12V	119	SMCLK_R2M3
118	12V	117	SMDAT_R2M3
116	12V	115	GND
114	3.3V AUX	113	CLK2+
112	GND	111	CLK2-
110	PERST#	109	GND
108	WAKE#	107	CLK1+

Pin	Signal Name	Pin	Signal Name
106	SPARE	105	CLK1-
104	THROTTLE_N	103	GND
102	GND	101	R00-
100	T00-	99	R00+
98	T00+	97	GND
96	GND	95	R01-
94	T01-	93	R01+
92	T01+	91	GND
90	GND	89	R02-
88	T01-	87	R02+
86	T01+	85	GND
84	GND	83	R03-
82	T01-	81	R03+
80	T01+	79	GND
78	GND	77	R04-
76	T01-	75	R04+
74	T01+	73	GND
72	GND	71	R05-
70	T01-	69	R05+
68	T01+	67	GND
66	GND	65	SPARE
64	SPARE	63	GND
62	GND	61	R06-
60	T01-	59	R06+
58	T01+	57	GND
56	GND	55	R07-
54	T01-	53	R07+
52	T01+	51	GND
50	GND	49	R08-
48	T01-	47	R08+
46	T01+	45	GND
44	GND	43	R09-
42	T01-	41	R09-
40	T01+	39	GND
38	GND	37	R10-
36	T01-	35	R10+
34	T01+	33	GND
32	GND	31	R11-
30	T01-	29	R11+
28	T01+	27	GND
26	GND	25	R12-
24	T01-	23	R12+
22	T01+	21	GND
20	GND	19	R13-
18	T01-	17	R13+
16	T01+	15	GND
14	GND	13	R14-
12	T01-	11	R14+
10	T01+	9	GND
8	GND	7	R15-
6	T01-	5	R15+
4	T01+	3	GND
2	GND	1	Riser ID0

7.3.2 VGA Connector

Table 42 details the pin-out definition of the internal 1x12 Pin video header using a ribbon cable to a standard 15-pin external VGA connector.

Table 42. VGA External Video Connector

Pin	Signal Name
1	SCL
2	SDA
3	GND
4	BLUE
5	B_RETURN
6	GREEN
7	G_RETURN
8	RED
9	R_RETURN
10	HSYNC
11	VSNC
12	IO_0 / RSVD

7.3.3 NIC Connectors

The server board provides three RJ45 networking ports: “Port 1”, “Port 2”, and a Dedicated Management Port. Dual connector pin-outs for NIC Port1 and Port2 are identical and are defined in Table 43. The pin-outs for the Dedicated Management Port are defined in Table 44.

Table 43. RJ-45 100Mb/1Gb/10Gb NIC Connector Pin-out

Pin	Signal Name
1	LAN_NICO_MDI_DP
2	LAN_NICO_MDI_DN
3	LAN_CT01_NICO
4	LAN_NICO_MDI_DP
5	LAN_NICO_MDI_DN
6	LAN_NICO_MDI_DP
7	LAN_NICO_MDI_DN
8	LAN_CT23_NICO
9	LAN_NICO_MDI_DP
10	LAN_NICO_MDI_DN
11	LAN_NICO_MDI_DP
12	LAN_NICO_MDI_DN
13	LED_NICO_1G_R
14	LED_NICO_10G_R

Table 44. RJ-45 10/100/1000 Dedicated Management Port NIC Connector Pin-out

Pin	Signal Name
1	LED_NIC_LINK0_100_N
2	LED_NIC_LINK0_1000_R_N
3	NIC_0_0_DP
4	NIC_0_0_DN
5	NIC_0_1_DP
6	NIC_0_1_DN
7	NIC_CT1
8	NIC_CT2
9	NIC_0_2_DP9
10	NIC_0_2_DN

Pin	Signal Name
11	NIC_0_3_DP
12	NIC_0_3_DN
13	LED_NIC_LINK0_LNKUP_N
14	LED_NIC_LINK0_ACT_R_N

7.3.4 SATA Connectors

7.3.4.1 Mini-SAS HD (SFF-8643) pin-out

The server board provides four SATA 6Gbps ports via a Mini-SAS HD (SFF-8643) connector (on Intel® S2600BPB and S2600BPS models only) and one M.2 SATA/PCIe* on-board connector.




Figure 76. Mini-SAS HD (SFF-8643) pin-out

Table 45. Mini-SAS HD (SFF-8643) SGPIO (SATA sideband signals)

SGPIO (SATA sideband signals)		
SB Pin	SGPIO	I/O
0	Clock	O
1	Load	O
2	Ground	
3	Ground	
4	Data out 0	O
5	Data out 1 (optional)	O
6	PD (optional)	O
7	Test Point (optional)	I

7.3.4.2 M.2 SATA/PCIe*

Intel® Server Board S2600BP supports both a SATA interface and PCIe* x4 interface via the Intel® C62x Series Chipset.

M.2 is a small form factor module supporting SSD/Memory-offloading technology using SATA or PCIe* x4 links.

The M.2 connector located on the server board provides support for a 42mm M.2 SSD and the M.2 connector on Riser Slot 2 Card provides support for a 80mm M.2 SSD.




Figure 77. M.2/NGFF Connector




Figure 78. M.2/NGFF Mated Module and Connector

Table 46. Lists the server board M.2 socket connector pinout:

Table 46. M.2/NGFF Socket 3 Connector Pinout (Mechanical Key M)

Pin	Signal	Signal	Pin
		GND	75
74	3.3V	GND	73
72	3.3V	GND	71
70	3.3V	PEDET (GND-SATA)	69
68	SUSCLK(32kHz) (I) (0/3.3V)	N/C	67
66	Module Key	Module Key	65
64	Module Key	Module Key	63
62	Module Key	Module Key	61
60	Module Key	Module Key	59
58	Reserved/MFG Clock	GND	57
56	Reserved/MFG Data	N/C	55
54	N/C	N/C	53
52	N/C	GND	51
50	N/C	SATA-A+	49
48	N/C	SATA-A+	47
46	N/C	GND	45
44	N/C	SATA-A+	43
42	N/C	SATA-A+	41
40	N/C	GND	39
38	DEVSLP (I) (0/3.3V)	N/C	37
36	N/C	N/C	35
34	N/C	GND	33
32	N/C	N/C	31
30	N/C	N/C	29
28	N/C	GND	27
26	N/C	N/C	25
24	N/C	N/C	23
22	N/C	GND	21
20	N/C	N/C	19
18	3.3V	N/C	17
16	3.3V	GND	15
14	3.3V	N/C	13
12	3.3V	N/C	11
10	DAS/DSS# (0)(0D)	GND	9
8	N/C	N/C	7
6	N/C	N/C	5
4	3.3V	GND	3
2	3.3V	GND	1

Pin	Signal	Signal	Pin
		GND	75
74	3.3V	GND	73
72	3.3V	GND	71
70	3.3V	PEDET (GND-SATA)	69
68	SUSCLK(32kHz) (I) (0/3.3V)	N/C	67
66	Module Key	Module Key	65
64	Module Key	Module Key	63
62	Module Key	Module Key	61
60	Module Key	Module Key	59
58	Reserved/MFG Clock	GND	57
56	Reserved/MFG Data	REFCLKP	55
54	PEWake#(IO)(0/3.3V)	REFCLKN	53
52	CLSREQ#(IO)(0/3.3V)	GND	51
50	PERST#(I)(0/3.3V)	PERp0	49
48	N/C	PERn0	47
46	N/C	GND	45
44	N/C	PERp0	43
42	N/C	PERn0	41
40	N/C	GND	39
38	DEVSLP (I) (0/3.3V)	PERp1	37
36	N/C	PERn1	35
34	N/C	GND	33
32	N/C	PERp1	31
30	N/C	PERn1	29
28	N/C	GND	27
26	N/C	PERp2	25
24	N/C	PERn2	23
22	N/C	GND	21
20	N/C	PERp2	19
18	3.3V	PERn2	17
16	3.3V	GND	15
14	3.3V	PERp3	13
12	3.3V	PERn3	11
10	DAS/DSS# (0)(0D)	GND	9
8	N/C	PERp3	7
6	N/C	PERn3	5
4	3.3V	GND	3
2	3.3V	GND	1

7.3.5 Intel® RAID C600 Upgrade Key Connector

The server board provides support for one Intel® RAID C600 Upgrade Key (storage upgrade key) connector.

The Intel® RAID C600 Upgrade Key is a small PCB board that has up to two security EEPROMs that are read by the system ME to enable different versions of Intel® ESRT2 * RAID 5 software stack.

Table 47. Summary of RAID Keys

Storage Key Type	Description	Key Color
PCH SATA (No Key)	ESRT2 SW RAID 0/1/10	N/A
PCH SATA SW RAID 5 Key	ESRT2 SW RAID 5	Black
Intel® VROC – Standard (for NVMe)	RAID 0/1/10	TBD




Figure 79. SW RAID Activation Key (SAK) and Connector

The pin configuration of connector is defined in the following table.

Table 48. Storage Upgrade Key Connector

Pin	Signal Description
1	GND
2	PU_KEY
3	GND
4	PCH_SATA_RAID_KEY

7.3.6 Serial Port Connectors

The server board provides one internal DH-10, Serial-A header. The DH10 Serial Port header on board adheres to the *DTK wiring specification*. Table 49 defines the pin-outs.

Table 49. Internal 10-pin Serial A

Pin	Signal Name	Pin	Signal Name
1	SPA_DCD	2	SPA_DSR
3	SPA_SIN_N	4	SPA_RTS
5	SPA_SOUT_N	6	SPA_CTS
7	SPA_DTR	8	SPA_RI
9	GND	10	KEY

7.3.7 USB Connectors

Table 50 details the pin-out of the external stack USB port connectors found on the back edge of the server board.

Table 50. External USB port Connector

Pin	Signal Name	Description
1	+5V	USB Power
2	USB_N	Differential data line paired with DATAH0
3	USB_P	Differential data line paired with DATAH0
4	GND	Ground

One 2x5 connector on the server board provides an option to support one additional internal USB 2.0 port. The pin-out is detailed in Table 51.

Table 51. Internal USB Connector

Pin	Signal Name
1	GND
2	USB_P
3	USB_N
4	Key
5	GND

7.3.8 AUX Front Panel

This connector is used for front panel control when the server board is used in a non-Intel chassis.

Table 52. AUX Front Panel Connector

Pin	Signal Name	Pin	Signal Name
1	FP1 ID LED_N	2	R470-5VSB
3	FP1 HD activity LED_N	4	R470-5VSB
5	FP1 PWR LED_N	6	R470-5VSB
7	GND	8	FP1 PWR BTN_N
9	GND	10	FP1 ID BTN_N
11	FP1 RST BTN_N	12	Key

7.4 Fan Headers

The server board can support up to five system fans and includes system fan connector/header options to support both Intel compute module and non-Intel chassis configurations.




Figure 80. Server Board System Fan Headers

Note: System fans 1 thru 3 include common signals that are shared between two fan interface connector types: three 8-pin system fan connectors (OEM use only) and the 14-pin header used for Intel compute modules. Concurrent use of both fan signal interfaces is not supported.

7.4.1 Intel Compute Module System Fan Header

Intel compute modules support three dual rotor system fans. Each fan is cabled to an 8-pin connector on a power docking board which is then cabled to the 14-pin system fan header on the server board. All three compute module system fans are managed by the onboard BMC which uses a shared pulse width modulation (PWM) signal to control fan speed.

Table 53. Server board Fan Connector

Pin	Signal Name	Pin	Signal Name
1	PWM1	2	Reserved
3	Tach0	4	Tach1
5	Tach2	6	Tach3
7	Tach4	8	Tach5
9	NODE_ON	10	GND
11	SMBUS_R4 CLK	12	SMBUS_R4 DAT
13	NODE_ADRO	14	NODE_PWRGD

When the compute module is turned off, the fans will continue to rotate at a preset rate; this rate is selected by Intel and preset by the Fan manufacturer. This is done to stop air recirculation between compute modules. When docking the board to a live 12V rail, the fans could spin up immediately.

7.4.2 Non-Intel Chassis Fan Support

In support of non-Intel chassis, the server board includes three 8-pin system fan connectors labeled as SYS_FAN_#(1-3) with signals to support dual rotor system fans.

Table 54. SYS_FAN_# (1-3) Connector Pinout

Pin	Signal Description
1	GND
2	P12V
3	TACH (0,2,4)
4	PWM1
5	GND
6	P12V
7	TACH (1,3,5)
8	PWM1

These fan connectors use common Fan TACH and PWM signals as those routed to the 14-pin Intel compute module fan connector. For system fan 1-3 support, system OEMs and integrators must choose to use either the three 8-pin connectors or the 14-pin connector. Concurrent use of both connector types is not supported.

An additional two 4-pin system fan connectors, located near the back edge of the server board, provide System OEMs the option of additional system fan support or power for alternate cooling solutions. On the server board these connectors are labeled as “SYS_FAN_6” and “SYS_FAN_7”.

Table 55 . SYS_FAN_6 Connector Pinout

Pin	Signal Description
1	GND
2	P12V
3	TACH6
4	PWM6

Table 56. SYS_FAN_7 Connector Pinout

Pin	Signal Description
1	GND
2	P12V
3	TACH7
4	PWM7


7.5 Power Docking Board Connectors

Table 57 lists the connector type and pin definition on the power docking board.

Table 57. Main Power Input Connector

Pin	Signal Description	Pin	Signal Description
Lower Blade (Circuit 1)			
1	GND	2	GND
3	GND	4	GND
5	GND	6	GND
Upper Blade (Circuit 2)			
7	P12V	8	P12V
9	P12V	10	P12V
11	P12V	12	P12V

8. Configuration and Recovery Jumpers



BNP044

Figure 81. Configuration and Recovery Jumpers

Jumper Name	Description
BMC Force Update (J6B3)	If pins 2-3 are selected, the Integrated BMC Force Update Mode is enabled. These pins should be selected on 1-2 for normal system operation.
BIOS Default (J4B4)	If pins 2-3 are selected, the BIOS settings are restored to the factory defaults on the next reset. These pins should be selected on 1-2 for normal system operation.
BIOS Recovery (J4B3)	If the system BIOS is corrupted, an onboard backup copy of the BIOS can be loaded using the BIOS Recovery Jumper. To load the backup BIOS image, move the jumper from pins 1-2 (default) to pins 2-3, and power on the system. The system will boot to the backup BIOS image. These pins should be selected on 1-2 for normal system operation.
Password Clear (J4B2)	If pins 2-3 are selected, administrator and user passwords are cleared within five to ten seconds after the system is powered on. These pins should be selected on 1-2 for normal system operation.
ME Force Update (J4B1)	If pins 2-3 are selected, the ME Force Update Mode is enabled. These pins should be selected on 1-2 for normal system operation.

8.1 BMC Force Update (J6B3)

When performing a standard BMC firmware update procedure, the update utility places the BMC into an update mode, allowing the firmware to load safely onto the flash device. In the unlikely event the BMC firmware update process fails due to the BMC not being in the proper update state, the server board provides a BMC Force Update jumper (J6B3) which will force the BMC into the proper update state. The following procedure should be followed in the event the standard BMC firmware update process fails.

Table 58. Force Integrated BMC Update Jumper (J6B3)

Jumper Position	Mode of Operation	Note
1-2	Normal	Normal operation
2-3	Update	BMC in force update mode

To perform a Force BMC Update, follow these steps:

1. Unplug the compute module.
2. Remove the air duct. Refer to the Intel® Server Board S2600BP and Intel® Compute Module Service and Integration Guide for instructions.
3. Move the jumper (J6B3) from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Restore the air duct to the compute module.
5. Insert the compute module back to the chassis.
6. Power on the compute module by pressing the power button on the front panel.
7. Perform the BMC firmware update procedure as documented in the *Release Notes* included in the given BMC firmware update package. After successful completion of the firmware update process, the firmware update utility may generate an error stating the BMC is still in update mode.
8. Power down and plug out the compute module.
9. Remove the air duct.
10. Move the jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
11. Restore the air duct to the compute module.
12. Plug in the compute module back to the chassis and power up the server.

NOTE: Normal BMC functionality is disabled with the Force BMC Update jumper set to the enabled position. You should never run the server with the BMC Force Update jumper set in this position. You should use this jumper setting only when the standard firmware update process fails. This jumper should remain in the default/disabled position when the server is running normally.

The server board has several 3-pin jumper blocks that can be used to configure, protect, or recover specific features of the server board.

8.2 ME Force Update (J4B1)

When this 3-pin jumper is set, it manually puts the ME firmware in update mode, which enables the user to update ME firmware code when necessary.

Table 59. Force ME Update Jumper (J4B1)

Jumper Position	Mode of Operation	Note
1-2	Normal	Normal operation
2-3	Update	ME in force update mode

NOTE: Normal ME functionality is disabled with the Force ME Update jumper set to the enabled position. You should never run the server with the ME Force Update jumper set in this position. You should only use this jumper setting when the standard firmware update process fails. This jumper should remain in the default/disabled position when the server is running normally.

To perform a Force ME Update, follow these steps:

1. Unplug the compute module from the chassis.
2. Remove the air duct. Refer to the Intel® Server Board S2600BP - *Service and Integration Guide* for instructions.
3. Move the jumper (J4B1) from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Restore the air duct back to the compute module.
5. Plug in the compute module back to the chassis.
6. Perform the ME firmware update procedure as documented in the Release Notes file that is included in the given system update package.
7. After update process is done, plug out the compute module out of the chassis.
8. Remove the air duct.
9. Move the jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
10. Restore the compute module back to the chassis.

8.3 Password Clear (J4B2)

The user sets this 3-pin jumper to clear the password.

Table 60. Password Clear Jumper (J4B2)

Jumper Position	Mode of Operation	Note
1-2	Normal	Normal mode, password in protection
2-3	Clear Password	BIOS password is cleared

This jumper causes both the User password and the Administrator password to be cleared if they were set. The operator should be aware that this creates a security gap until passwords have been installed again.

NOTE: No method of resetting BIOS configuration settings to the default values will affect either the Administrator or User passwords.

This is the only method by which the Administrator and User passwords can be cleared unconditionally. Other than this jumper, passwords can only be set or cleared by changing them explicitly in BIOS Setup or by similar means.

The recommended steps for clearing the User and Administrator passwords are

1. Plug out the compute module and remove the air duct.
2. Move the jumper from pins 1-2 to pins 2-3. It is necessary to leave the jumper in place while rebooting the system in order to clear the passwords.
3. Installed the air duct and plug in and power up the compute module.
4. Boot into the BIOS Setup. Check the Error Manager tab for POST Error Codes:
5221 Passwords cleared by jumper
5224 Password clear jumper is set
5. Power down and plug out the compute module and remove the air duct again.
6. Restore the jumper from pins 2-3 to the normal setting of pins 1-2.
7. Install the air duct and plug in and power up the compute module.
8. **Strongly recommended:** Boot into the BIOS Setup immediately, go to the Security tab and set the Administrator and User passwords if you intend to use BIOS password protection.

8.4 BIOS Recovery Mode (J4B3)

If a system is completely unable to boot successfully to an OS, hangs during POST, or even hangs and fails to start executing POST, it may be necessary to perform a BIOS Recovery procedure, which can replace a defective copy of the Primary BIOS.

The BIOS introduces three mechanisms to start the BIOS recovery process, which is called Recovery Mode:

- The Recovery Mode Jumper causes the BIOS to boot in Recovery Mode.
- The BootBlock detects partial BIOS update and automatically boots in Recovery Mode.
- The BMC asserts Recovery Mode GPIO in case of partial BIOS update and FRB2 time-out.

Table 61. BIOS Recovery Mode Jumper (J4B3)

Jumper Position	Mode of Operation	Note
1-2	Normal	Normal mode
2-3	Recovery	BIOS in recovery mode

The BIOS Recovery takes place without any external media or Mass Storage device as it utilizes the Backup BIOS inside the BIOS flash in Recovery Mode. The Recovery procedure is included here for general reference. However, if in conflict, the instructions in the BIOS Release Notes are the definitive version.

When Recovery Mode Jumper is set, the BIOS begins with a “Recovery Start” event logged to the SEL, loads and boots with the Backup BIOS image inside the BIOS flash itself. This process takes place before any video or console is available. The system boots up into the Shell directly while a “Recovery Complete” SEL logged. An external media is required to store the BIOS update package and steps are the same as the normal BIOS update procedures. After the update is complete, there will be a message displayed stating that the “BIOS has been updated successfully” indicating the BIOS update process is finished. The User should then switch the recovery jumper back to normal operation and restart the system by performing a power cycle.

If the BIOS detects partial BIOS update or the BMC asserts Recovery Mode GPIO, the BIOS will boot up with Recovery Mode. The difference is that the BIOS boots up to the Error Manager Page in the BIOS Setup utility. In the BIOS Setup utility, boot device, Shell or Linux for example, could be selected to perform the BIOS update procedure under Shell or OS environment.

Again, before starting to perform a Recovery Boot, be sure to check the BIOS Release Notes and verify the Recovery procedure shown in the Release Notes.

The following steps demonstrate this recovery process:

1. Unplug the compute module and remove the air duct.
2. Move the jumper (J4B3) from the default operating position (covering pins 1 and 2) to the BIOS Recovery position (covering pins 2 and 3).
3. Restore the air duct back to the compute module.
4. Plug in the compute module back to the chassis.
5. Power on the compute module.
6. The BIOS will load and boot with the backup BIOS image without any video or display.
7. When the compute module boots into the EFI shell directly, the BIOS recovery is successful.
8. Power off the compute module.
9. Plug out the compute module from the chassis.
10. Remove the air duct and put the jumper (J4B3) back to the normal position (covering pins 1 and 2).
11. Restore the air duct and put the compute module back to the chassis.
12. A normal BIOS update can be performed if needed.

8.5 BIOS Default (J4B4)

Table 62. BIOS Default Jumper

Jumper Position	Mode of Operation	Note
1-2	Normal	Normal mode
2-3	Clear BIOS settings	BIOS settings are reset to factory default

This jumper causes the BIOS Setup settings to be reset to their default values. On previous generations of server boards, this jumper has been referred to as “Clear CMOS”, or “Clear NVRAM”. Setting this jumper according to the procedure below will clear all current contents of NVRAM variable storage, and then load the BIOS default settings.

Note that this jumper does not reset Administrator or User passwords. In order to reset passwords, the Password Clear jumper must be used.

The recommended steps to reset to the BIOS defaults are:

1. Plug out the compute module and remove the air duct.
2. Move the jumper from pins 1-2 to pins 2-3 momentarily. It is not necessary to leave the jumper in place while rebooting.
3. Restore the jumper from pins 2-3 to the normal setting of pins 1-2.
4. Install the air duct and plug in the compute module, and power up.
5. Boot the system into Setup. Check the Error Manager tab, and you should see POST Error Codes:
0012 System RTC date/time not set
5220 BIOS Settings reset to default settings
6. Go to the Setup Main tab, and set the System Date and System Time to the correct current settings. Make any other changes that are required in Setup – for example, Boot Order.

9. Intel® Light-Guided Diagnostics

Intel® Server Board S2600BP has several onboard diagnostic LEDs to assist in troubleshooting board-level issues. This section provides a description of the location and function of each LED on the server board.

9.1 Status LED

NOTE: The status LED state shows the state for the current, most severe fault. For example, if there was a critical fault due to one source and a non-critical fault due to another source, the status LED state would be solid on (the critical fault state).

The status LED is a bicolor LED. Green (status) shows a normal operation state or a degraded operation. Amber (fault) shows the hardware state and overrides the green status.

The Integrated BMC-detected state and the state from the other controllers, such as the SCSI/SATA hot-swap controller state, are included in the LED state. For fault states monitored by the Integrated BMC sensors, the contribution to the LED state follows the associated sensor state, with the priority going to the most critical state currently asserted.

When the server is powered down (transitions to the DC-off state or S5), the Integrated BMC is still on standby power and retains the sensor and front panel status LED state established prior to the power-down event.

Table 63 maps the server state to the LED state.




Figure 82. Status LED (A), ID LED (B), Diagnostic LEDs (C)

Table 63. Status LED State Definitions

Color	State	Criticality	Description
Off	System is not operating	Not ready	<ul style="list-style-type: none"> System is powered off (AC and/or DC). System is in EuP Lot6 Off Mode. System is in S5 Soft-Off State.
Green	Solid on	OK	<p>Indicates that the System is running (in S0 State) and its status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running.</p> <p>After a BMC reset, and in conjunction with the Chassis ID solid ON, the BMC is booting Linux*. Control has been passed from BMC uBoot to BMC Linux* itself. It will be in this state for ~10~20 seconds.</p>
Green	~1 Hz blink	Degraded – system is operating in a degraded state although still functional, or system is operating in a redundant state but with an impending failure warning	<p>System degraded:</p> <ul style="list-style-type: none"> Redundancy loss such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. Fan warning or failure when the number of fully operational fans is less than minimum number needed to cool the system Non-critical threshold crossed – Temperature (including HSBP temp), voltage, input power to power supply, output current for main power rail from power supply and Processor Thermal Control (Therm Ctrl) sensors Power supply predictive failure occurred while redundant power supply configuration was present Unable to use all of the installed memory (more than 1 DIMM installed) Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the system no longer has spared DIMMs (a redundancy lost condition). Corresponding DIMM LED lit. In mirrored configuration, when memory mirroring takes place and system loses memory redundancy Battery failure BMC executing in uBoot. (Indicated by Chassis ID blinking at 3Hz). System in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux*. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux* image into flash. BMC Watchdog has reset the BMC

Color	State	Criticality	Description
			<ul style="list-style-type: none"> Power Unit sensor offset for configuration error is asserted HDD HSC is off-line or degraded
Amber	~1 Hz blink	Non-critical – System is operating in a degraded state with an impending failure warning, although still functioning	<p>Non-fatal alarm – system is likely to fail:</p> <ul style="list-style-type: none"> Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply and PROCHOT (Therm Ctrl) sensors VRD Hot asserted Minimum number of fans to cool the system not present or failed Hard drive fault Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present) In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window
Amber	Solid on	Critical, non-recoverable – System is halted	<p>Fatal alarm – system has failed or shutdown:</p> <ul style="list-style-type: none"> CPU CATERR signal asserted MSID mismatch detected (CATERR also asserts for this case) CPU 1 is missing CPU Thermal Trip No power good – power fault DIMM failure when there is only 1 DIMM present and hence no good memory present Runtime memory uncorrectable error in non-redundant mode DIMM Thermal Trip or equivalent SSB Thermal Trip or equivalent CPU ERR2 signal asserted BMC/Video memory test failed. (Chassis ID shows blue/solid-on for this condition) Both uBoot BMC firmware images are bad. (Chassis ID shows blue/solid-on for this condition) 240VA fault

Color	State	Criticality	Description
			<ul style="list-style-type: none"> Fatal Error in processor initialization: <ul style="list-style-type: none"> Processor family not identical Processor model not identical Processor core/thread counts not identical Processor cache size not identical Unable to synchronize processor frequency Unable to synchronize QPI link frequency Uncorrectable memory error in a non-redundant mode

9.2 ID LED

The ID LED provides a visual indication of the server board or compute module that may require service. This is useful in a rack environment where there are multiple systems operating in close proximity. The state of the ID LED is affected by the following:

- Toggled by the ID button on the Front Panel
- Controlled by the *Chassis Identify* command (IPMI)

Table 64. ID LED

State	LED State
Identify active through button	Solid on
Identify active through command	~1 Hz blink
Off	Off

There is no precedence or lock-out mechanism for the control sources. When a new request arrives, all previous requests are terminated. For example, if the ID LED is blinking and the chassis ID button is pressed, then the ID LED changes to solid on. If the button is pressed again with no intervening commands, the ID LED turns off.

9.3 BMC Boot/Reset Status LED Indicators

During the BMC boot or BMC reset process, the System Status LED and System ID LED are used to indicate BMC boot process transitions and states. A BMC boot will occur when AC power is first applied to the system. A BMC reset will occur after a BMC firmware update, upon receiving a BMC cold reset command, and upon a BMC watchdog initiated reset. The following table defines the LED states during the BMC Boot/Reset process.


Table 65. BMC Boot/Reset Status LED Indicators

BMC Boot/Reset State	Chassis ID LED	Status LED	Comment
BMC/Video memory test failed	Solid Blue	Solid Amber	Non-recoverable condition. Contact your Intel representative for information on replacing this motherboard.
Both Universal Bootloader (u-Boot) images bad	Blink Blue 6 Hz	Solid Amber	Non-recoverable condition. Contact your Intel representative for information on replacing this motherboard.
BMC in u-Boot	Blink Blue 3 Hz	Blink Green 1Hz	Blinking green indicates degraded state (no manageability), blinking blue indicates u-Boot is running but has not transferred control to BMC Linux. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux image into flash.
BMC Booting Linux	Solid Blue	Solid Green	Solid green with solid blue after an AC cycle/BMC reset, indicates that the control has been passed from u-Boot to BMC Linux itself. It will be in this state for ~10~20 seconds.
End of BMC boot/reset process. Normal system operation	Off	Solid Green	Indicates BMC Linux has booted and manageability functionality is up and running. Fault/Status LEDs operate as per usual.

9.4 POST Code Diagnostic LEDs

Eight amber POST code diagnostic LEDs are located on the back left edge of the server board, as in Figure 83.

During the system boot process, the BIOS executes a number of platform configuration processes, each of which is assigned a specific hex POST code number. As each configuration routine is started, the BIOS displays the given POST code to the POST code diagnostic LEDs on the back edge of the server board. To assist in troubleshooting a system hang during the POST process, you can use the Diagnostic LEDs to identify the last POST process executed. For a complete description of how these LEDs are read and a list of all supported POST codes, refer to Appendix B.



BNP043

Figure 83. Rear Panel Diagnostic LEDs

10. Basic and Advanced Server Management

The integrated BMC has support for basic and advanced server management features. Basic management features are available by default. Advanced management features are enabled with the addition of an optionally installed Remote Management Module 4 Lite (RMM4 Lite) key.

Table 66. Intel® Remote Management Module 4 (RMM4) Options

Intel Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE2	Intel® Remote Management Module 4 Lite	RMM4 Lite Activation Key	Enables KVM & media redirection

When the BMC firmware initializes, it attempts to access the Intel® RMM4 lite. If the attempt to access Intel® RMM4 lite is successful, then the BMC activates the Advanced features.

Table 67 identifies both Basic and Advanced server management features.

Table 67. Basic and Advanced Server Management Features Overview

Feature	Basic	Advanced w/RMM4 Lite Key
IPMI 2.0 Feature Support	X	X
In-circuit BMC Firmware Update	X	X
FRB 2	X	X
Chassis Intrusion Detection	X	X
Fan Redundancy Monitoring	X	X
Hot-Swap Fan Support	X	X
Acoustic Management	X	X
Diagnostic Beep Code Support	X	X
Power State Retention	X	X
ARP/DHCP Support	X	X
PECI Thermal Management Support	X	X
E-mail Alerting	X	X
Embedded Web Server	X	X
SSH Support	X	X
Integrated KVM		X
Integrated Remote Media Redirection		X
Lightweight Directory Access Protocol (LDAP)	X	X
Intel® Intelligent Power Node Manager Support	X	X
SMASH CLP	X	X

10.1.1 Dedicated Management Port

The server board includes a dedicated 1GbE RJ45 Management Port. The management port is active with or without the RMM4 Lite key installed.

10.1.2 Embedded Web Server

BMC Base manageability provides an embedded web server and an OEM-customizable web GUI which exposes the manageability features of the BMC base feature set. It is supported over all on-board NICs that have management connectivity to the BMC as well as an optional dedicated add-in management NIC. At least two concurrent web sessions from up to two different users are supported. The embedded web user interface shall support the following client web browsers:

- Microsoft Internet Explorer*
- Mozilla Firefox*

The embedded web user interface supports strong security (authentication, encryption, and firewall support) since it enables remote server configuration and control. The user interface presented by the embedded web user interface, shall authenticate the user before allowing a web session to be initiated. Encryption using 128-bit SSL is supported. User authentication is based on user id and password.

The GUI presented by the embedded web server authenticates the user before allowing a web session to be initiated. It presents all functions to all users but grays-out those functions that the user does not have privilege to execute. For example, if a user does not have privilege to power control, then the item shall be displayed in grey-out font in that user's UI display. The web GUI also provides a launch point for some of the advanced features, such as KVM and media redirection. These features are grayed out in the GUI unless the system has been updated to support these advanced features. The embedded web server only displays US English or Chinese language output.

Additional features supported by the web GUI includes the following:

- Presents all the Basic features to the users
- Power on/off/reset the server and view current power state
- Displays BIOS, BMC, ME and SDR version information
- Display overall system health
- Configuration of various IPMI over LAN parameters for both IPV4 and IPV6
- Configuration of alerting (SNMP and SMTP)
- Display system asset information for the product, board, and chassis
- Display of BMC-owned sensors (name, status, current reading, enabled thresholds), including color-code status of sensors
- Provides ability to filter sensors based on sensor type (Voltage, Temperature, Fan and Power supply related)
- Automatic refresh of sensor data with a configurable refresh rate
- On-line help
- Display/clear SEL (display is in easily understandable human readable format)
- Support for major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*)
- Automatic time-out of the GUI session after a user-configurable inactivity period (30 minutes, by default)

- Embedded Platform Debug feature – Allow the user to initiate a “debug dump” to a file that can be sent to Intel for debug purposes.
- A Virtual Front Panel provides the same functionality as the local front panel. The displayed LEDs match the current state of the local panel LEDs. The displayed buttons (for example, power button) can be used in the same manner as the local buttons.
- Display of ME sensor data. Only sensors that have associated SDRs loaded will be displayed.
- Ability to save the SEL to a file
- Ability to force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display of processor and memory information as is available over IPMI over LAN.
- Ability to get and set Node Manager (NM) power policies
- Display of power consumed by the server
- Ability to view and configure VLAN settings
- Warn user the reconfiguration of IP address will cause disconnect
- Capability to block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.
- Server Power Control – Ability to force into Setup on a reset
- System POST results – The web server provides the system's Power-On Self Test (POST) sequence for the previous two boot cycles, including timestamps. The timestamps may be viewed in relative to the start of POST or the previous POST code.
- Customizable ports – The web server provides the ability to customize the port numbers used for SMASH, http, https, KVM, secure KVM, remote media, and secure remote media.

For additional information, refer to the *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*.

10.1.3 Advanced Management Feature Support (RMM4 Lite)

The integrated server board management controller has support for advanced management features which are enabled when an optional Intel® Remote Management Module 4 Lite (RMM4 Lite) is installed. The Intel® RMM4 add-on offers convenient, remote KVM access and control through LAN and internet. It captures, digitizes, and compresses video and transmits it with keyboard and mouse signals to and from a remote computer. Remote access and control software runs in the integrated server board management controller, utilizing expanded capabilities enabled by the Intel® RMM4 (iPC – AXXRMM4LITE2) hardware.




Figure 84. Optional RMM4 Lite

Key Features of the RMM4 add-on are

- KVM redirection from either the dedicated management NIC or the server board NICs used for management traffic; up to two KVM sessions.
- Media Redirection – The media redirection feature is intended to allow system administrators or users to mount a remote IDE or USB CDROM, floppy drive, or a USB flash disk as a remote device to the server. Once mounted, the remote device appears just like a local device to the server allowing system administrators or users to install software (including operating systems), copy files, update BIOS, or boot the server from this device.
- KVM – Automatically senses video resolution for best possible screen capture, high performance mouse tracking and synchronization. It allows remote viewing and configuration in pre-boot POST and BIOS setup.

10.1.3.1 Keyboard, Video, Mouse (KVM) Redirection

The BMC firmware supports keyboard, video, and mouse redirection (KVM) over LAN. This feature is available remotely from the embedded web server as a Java applet. This feature is only enabled when the Intel® RMM4 lite is present. The client system must have a Java Runtime Environment (JRE) version 6.0 or later to run the KVM or media redirection applets.

The BMC supports an embedded KVM application (Remote Console) that can be launched from the embedded web server from a remote console. USB 1.1 or USB 2.0 based mouse and keyboard redirection are supported. It is also possible to use the KVM-redirection (KVM-r) session concurrently with media-redirection (media-r). This feature allows a user to interactively use the keyboard, video, and mouse (KVM) functions of the remote server as if the user were physically at the managed server. KVM redirection console supports the following keyboard layouts: English, Dutch, French, German, Italian, Russian, and Spanish.

KVM redirection includes a “soft keyboard” function. The soft keyboard is used to simulate an entire keyboard that is connected to the remote system. The soft keyboard functionality supports the following layouts: English, Dutch, French, German, Italian, Russian, and Spanish.

The KVM-redirection feature automatically senses video resolution for best possible screen capture and provides high-performance mouse tracking and synchronization. It allows remote viewing and configuration in pre-boot POST and BIOS setup, once BIOS has initialized video.

Other attributes of this feature include

- Encryption of the redirected screen, keyboard, and mouse.
- Compression of the redirected screen.
- Ability to select a mouse configuration based on the OS type.
- Supports user definable keyboard macros.

The KVM redirection feature supports the following resolutions and refresh rates:

- 640x480 at 60Hz, 72Hz, 75Hz, 85Hz
- 800x600 at 60Hz, 72Hz, 75Hz, 85Hz
- 1024x768 at 60Hz, 72Hz, 75Hz, 85Hz
- 1152x864 at 75Hz
- 1280x800 at 60Hz
- 1280x1024 at 60Hz
- 1440x900 at 60Hz
- 1600x1200 at 60Hz

10.1.3.2 Remote Console

The Remote Console is the redirected screen, keyboard and mouse of the remote host system. To use the Remote Console window of your managed host system, the browser must include a Java® Runtime Environment plug-in. If the browser has no Java support, such as with a small handheld device, the user can maintain the remote host system using the administration forms displayed by the browser.

The Remote Console window is a Java Applet that establishes TCP connections to the BMC. The protocol that is run over these connections is a unique KVM protocol and not HTTP or HTTPS. This protocol uses ports #7578 for KVM, #5120 for CDROM media redirection, and #5123 for Floppy/USB media redirection. When encryption is enabled, the protocol uses ports #7582 for KVM, #5124 for CDROM media redirection, and #5127 for Floppy/USB media redirection. The local network environment must permit these connections to be made, that is, the firewall and, in case of a private internal network, the NAT (Network Address Translation) settings have to be configured accordingly.

10.1.3.3 Performance

The remote display accurately represents the local display. The feature adapts to changes to the video resolution of the local display and continues to work smoothly when the system transitions from graphics to text or vice-versa. The responsiveness may be slightly delayed depending on the bandwidth and latency of the network.

Enabling KVM and/or media encryption will degrade performance. Enabling video compression provides the fastest response while disabling compression provides better video quality.

For the best possible KVM performance, a 2Mb/sec link or higher is recommended.

The redirection of KVM over IP is performed in parallel with the local KVM without affecting the local KVM operation.

10.1.3.4 Security

The KVM redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.

10.1.3.5 Availability

The remote KVM session is available even when the server is powered-off (in stand-by mode). No re-start of the remote KVM session shall be required during a server reset or power on/off. A BMC reset (for example, due to a BMC Watchdog initiated reset or BMC reset after BMC firmware update) will require the session to be re-established.

KVM sessions persist across system reset, but not across an AC power loss.

10.1.3.6 Usage

As the server is powered up, the remote KVM session displays the complete BIOS boot process. The user is able to interact with BIOS setup, change and save settings as well as enter and interact with option ROM configuration screens.

At least two concurrent remote KVM sessions are supported. It is possible for at least two different users to connect to the same server and start remote KVM sessions.

10.1.3.7 Force-enter BIOS Setup

KVM redirection can present an option to force-enter BIOS Setup. This enables the system to enter F2 setup while booting which is often missed by the time the remote console redirects the video.

10.1.3.8 Media Redirection

The embedded web server provides a Java applet to enable remote media redirection. This may be used in conjunction with the remote KVM feature, or as a standalone applet.

The media redirection feature is intended to allow system administrators or users to mount a remote IDE or USB CD-ROM, floppy drive, or a USB flash disk as a remote device to the server. Once mounted, the remote device appears just like a local device to the server, allowing system administrators or users to install software (including operating systems), copy files, update BIOS, and so on, or boot the server from this device.

The following capabilities are supported:

- The operation of remotely mounted devices is independent of the local devices on the server. Both remote and local devices are useable in parallel.
- Either IDE (CD-ROM, floppy) or USB devices can be mounted as a remote device to the server.
- It is possible to boot all supported operating systems from the remotely mounted device and to boot from disk IMAGE (*.IMG) and CD-ROM or DVD-ROM ISO files. See the Tested/supported Operating System List for more information.
- Media redirection supports redirection for both a virtual CD device and a virtual Floppy/USB device concurrently. The CD device may be either a local CD drive or else an ISO image file; the Floppy/USB device may be either a local Floppy drive, a local USB device, or else a disk image file.
- The media redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.

- A remote media session is maintained even when the server is powered-off (in standby mode). No restart of the remote media session is required during a server reset or power on/off. An BMC reset (for example, due to an BMC reset after BMC firmware update) will require the session to be re-established.
- The mounted device is visible to (and useable by) managed system's OS and BIOS in both pre-boot and post-boot states.
- The mounted device shows up in the BIOS boot order and it is possible to change the BIOS boot order to boot from this remote device.
- It is possible to install an operating system on a bare metal server (no OS present) using the remotely mounted device. This may also require the use of KVM-r to configure the OS during installation.

USB storage devices will appear as floppy disks over media redirection. This allows for the installation of device drivers during OS installation.

If either a virtual IDE or virtual floppy device is remotely attached during system boot, both the virtual IDE and virtual floppy are presented as bootable devices. It is not possible to present only a single-mounted device type to the system BIOS.

10.1.3.9 Availability

The default inactivity timeout is 30 minutes and is not user-configurable. Media redirection sessions persist across system reset but not across an AC power loss or BMC reset.

10.1.3.10 Network Port Usage


The KVM and media redirection features use the following ports:

- 5120 – CD Redirection
- 5123 – FD Redirection
- 5124 – CD Redirection (Secure)
- 5127 – FD Redirection (Secure)
- 7578 – Video Redirection
- 7582 – Video Redirection (Secure)

For additional information, refer to the *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*.

11. Intel Compute Module Thermal Management

The compute module is designed to operate at external ambient temperatures of between 10°C and 35°C with limited excursion-based operation up to 45°C. Working with integrated platform management, several features within the compute module are designed to move air in a front-to-back direction, through the compute module and over critical components to prevent them from overheating and allow the system to operate with best performance.



BNP052

Figure 85. Air Flow and Fan Identification

Table 68 provides air flow data associated with the different product models within this product family, and is provided for reference purposes only. The data was derived from actual wind tunnel test methods and measurements using fully configured system configurations. Lesser system configurations may produce slightly different data results. As such, the CFM data provided using server management utilities that utilize platform sensor data, may vary from the data listed in the table.

Table 68. Air Flow

Model	Single Compute Module Airflow
With Intel® Server Chassis H2312XXLR3	6~44.1
With Intel® Server Chassis H2224XXLR3	6.4~46.7
With Intel® Server Chassis H2204XXLRE	6.4~46.7

To keep the compute module operating within the supported maximum thermal limits, the compute module must meet the following operating and configuration guidelines:

- The compute module operating ambient is designed for sustained operation from 10°C up to 35°C (ASHRAE Class A2)
- The compute module can operate up to 40°C (ASHRAE Class A3) for up to 900 hours per year.
- The compute module can operate up to 45°C (ASHRAE Class A4) for up to 90 hours per year.

Note: There is no long-term system reliability impact when operating at the extended temperature range within the approved duration limits. However, compute module performance may be impacted when operating within the extended operating temperature ranges

- Specific configuration requirements and limitations will be documented in the thermal configuration matrix tables found in the *Power Budget and Thermal Configuration Tool*, available online at <https://www.intel.com/content/www/us/en/support/server-products.html>
- The CPU-1 processor + CPU heat sink must be installed first. The CPU-2 heat sink must be installed at all times, with or without a processor installed
- With the compute module operating, the air duct must be installed at all times
- The Intel® Compute Module HNS2600BP product family does not support redundant cooling. If one of the compute module fans fails, it is recommended to replace the failed fan as soon as possible

12. System Security

The server board supports a variety of system security options designed to prevent unauthorized system access or tampering of server settings. System security options supported include:

- Password Protection
- Front Panel Lockout

The <F2> BIOS Setup Utility, accessed during POST, includes a Security tab where options to configure passwords, and front panel lockout can be found.



Figure 86. Security Tab of BIOS Setup Utility

12.1 Password Setup

The BIOS uses passwords to prevent unauthorized access to the server. Passwords can restrict entry to the BIOS Setup utility, restrict use of the Boot Device popup menu during POST, suppress automatic USB device re-ordering, and prevent unauthorized system power on. It is strongly recommended that an Administrator Password be set. A system with no Administrator password set allows anyone who has access to the server to change BIOS settings.

An Administrator password must be set in order to set the User password.

The maximum length of a password is 14 characters and can be made up of a combination of alphanumeric (a-z, A-Z, 0-9) characters and any of the following special characters:

! @ # \$ % ^ & * () - _ + = ?

Passwords are case sensitive.

The Administrator and User passwords must be different from each other. An error message will be displayed and a different password must be entered if there is an attempt to enter the same password for both. The use of "Strong Passwords" is encouraged, but not required. In order to meet the criteria for a strong password, the password entered must be at least 8 characters in length, and must include at least one each of alphabetic, numeric, and special characters. If a weak password is entered, a warning message will be displayed, and the weak password will be accepted.

Once set, a password can be cleared by changing it to a null string. This requires the Administrator password, and must be done through BIOS Setup or other explicit means of changing the passwords. Clearing the Administrator password will also clear the User password. Passwords can also be cleared by using the Password Clear jumper on the server board. See Chapter 8 – Configuration and Recovery Jumpers

Resetting the BIOS configuration settings to default values (by any method) has no effect on the Administrator and User passwords.

As a security measure, if a User or Administrator enters an incorrect password three times in a row during the boot sequence, the system is placed into a halt state. A system reset is required to exit out of the halt state. This feature makes it more difficult to guess or break a password.

In addition, on the next successful reboot, the Error Manager displays a Major Error code 0048, which also logs a SEL event to alert the authorized user or administrator that a password access failure has occurred.

12.1.1 System Administrator Password Rights

When the correct Administrator password is entered when prompted, the user has the ability to perform the following actions:

- Access the <F2> BIOS Setup Utility
- Configure all BIOS setup options in the <F2> BIOS Setup Utility
- Clear both the Administrator and User passwords
- Access the <F6> Boot Menu during POST

If the Power On Password function is enabled in BIOS Setup, the BIOS will halt early in POST to request a password (Administrator or User) before continuing POST

12.1.2 Authorized System User Password Rights and Restrictions

When the correct User password is entered, the user has the ability to perform the following:

- Access the <F2> BIOS Setup Utility
- View, but not change any BIOS Setup options in the <F2> BIOS Setup Utility
- Modify System Time and Date in the BIOS Setup Utility
- If the Power On Password function is enabled in BIOS Setup, the BIOS will halt early in POST to request a password (Administrator or User) before continuing POST

In addition to restricting access to most Setup fields to viewing only when a User password is entered, defining a User password imposes restrictions on booting the system. In order to simply boot in the defined boot order, no password is required. However, the F6 Boot popup menu prompts for a password, and can only be used with the Administrator password. Also, when a User password is defined, it suppresses the USB Reordering that occurs, if enabled, when a new USB boot device is attached to the system. A User is restricted from booting in anything other than the Boot Order defined in the Setup by an Administrator.

12.2 Front Panel Lockout

If enabled in BIOS setup, this option disables the following front panel features:

- The OFF function of the Power button
- System Reset button

If [Enabled] is selected, system power off and reset must be controlled via a system management interface.

12.3 Trusted Platform Module

The Intel® Server Board S2600BP has incorporated additional features for a trusted platform to include Trusted Execution Technology (TXT) and Bootguard that utilizes the Trusted Platform Module (TPM) 2.0.

12.3.1 TPM 2.0

The Trusted Platform Module (TPM) provides platform security functions such as hash, encryption and secure storage, it works in conjunction with the Processor's TXT functionality. TPM2.0 is the next generation of TPM, it provides multiple benefits over the former TPM1.2 specification. The Intel® Server Board S2600BP uses the SPI version of TPM2.0 down on the board.

The key advantages of TPM2.0 are

- No special provisioning process
- Authorization is unified for more flexibility and relatively easy to revoke keys
- Inclusion of SHA-2 and AES encryption algorithms
- TCM support of Intel Processor TXT features
- Allows a "one-size-fits" all approach to International Security

Trusted Execution Technology (TXT) is a hardware solution that validates the behavior of key components within a server or PC at startup. Known as the "root of trust," the system checks the consistency in behaviors and launch time configurations against a "known good" sequence.

The Boot guard provides a hardware-based Static Root of Trust for Measurement (RTM) and Root of Trust for Verification (RTV) using Intel architectural components. This is accomplished in Boot guard by cryptographically verifying first portion of OEM BIOS code executed out of reset.

NOTES:

Once TPM is enabled, the deactivation of TPM only takes affect after AC power cycle.

BMC will only toggle the TPM_EN_VAR variable based on command from management network connection through SSH connection.

BMC can read the status of TPM disable pin from KCS interface, but cannot modify it.

For additional details, refer to the *Intel® Server System BIOS External Product Specification* at RDC.

Appendix A. *Integration and Usage Tips*

- This server board supports the Intel® Xeon® processor Scalable family with a Thermal Design Power (TDP) of up to and including 165 Watts on selected chassis models. Previous generations of the Intel® Xeon® processors are not supported. Server systems using this server board may or may not meet the TDP design limits of the server board. Validate the TDP limits of the server system before selecting a processor.
- Processors must be installed in order. CPU 1 must be populated for the server board to operate
- The “BIOS Default” jumper must be used to Clear CMOS whenever the processor configuration changes
- **The riser card slots are specifically designed to support riser cards only.** Attempting to install a PCIe* add-in card directly into a riser card slot on the server board may damage the server board, the add-in card, or both. **PCIe* Add-in in cards cannot be installed** in Raiser Slot Card #1 on the Intel® Compute Module HNS2600BPS
- For the best performance, the number of DDR4 DIMMs installed should be balanced across both processor sockets and memory channels
- On the back edge of the server board are eight diagnostic LEDs that display a sequence of amber POST codes during the boot process. If the server board hangs during POST, the LEDs display the last POST event run before the hang.
- The System Status LED will be set to a steady Amber color for all Fatal Errors that are detected during processor initialization. A steady Amber System Status LED indicates that an unrecoverable system failure condition has occurred
- RAID partitions created using either RSTe or ESRT2 cannot span across the two embedded SATA controllers. Only drives attached to a common SATA controller can be included in a RAID partition.


Appendix B. *POST Code Diagnostic LED Decoder*

As an aid to assist in troubleshooting a system hang that occurs during a system's Power-On Self-Test (POST) process, the server board includes a bank of eight POST Code Diagnostic LEDs on the back edge of the server board.

During the system boot process, Memory Reference Code (MRC) and System BIOS execute a number of memory initialization and platform configuration processes, each of which is assigned a hex POST code number. As each routine is started, the given POST code number is displayed to the POST Code Diagnostic LEDs on the back edge of the server board.

During a POST system hang, the displayed post code can be used to identify the last POST routine that was run prior to the error occurring, helping to isolate the possible cause of the hang condition.

Each POST code is represented by eight LEDs; four Green and four Amber. The POST codes are divided into two nibbles, an upper nibble and a lower nibble. The upper nibble bits are represented by Amber Diagnostic LEDs #7, #6, #5, #4. The lower nibble bits are represented by Green Diagnostics LEDs #3, #2, #1 and #0. If the bit is set in the upper and lower nibbles, the corresponding LED is lit. If the bit is clear, the corresponding LED is off (Lit LED = 1, Off LED = 0).



BNP043

Figure 87. POST Diagnostic LED Location

NOTE: When facing the back of the system, all POST Diagnostic codes are read from left to right starting from MSB to LSB in given numerical order (7-6-5-4-3-2-1-0) as show on Figure 88. Failing to follow this instruction will result in decoding the LEDs incorrectly.

In the following example, the BIOS sends a value of “ACh” to the diagnostic LED decoder.




Figure 88. Correct Order Reading POST Diagnostic Codes

NOTE: Diagnostic LEDs are best read and decoded when viewing the LEDs from the back of the system.

Table 69. POST Progress Code Decoding LED Example

LEDs		Upper Nibble AMBER LEDs				Lower Nibble GREEN LEDs			
		MSB							LSB
		LED #7	LED #6	LED #5	LED #4	LED #3	LED #2	LED #1	LED #0
		8h	4h	2h	1h	8h	4h	2h	1h
Status		ON	OFF	ON	OFF	ON	ON	OFF	OFF
Read Value	Binary	1	0	1	0	1	1	0	0
	Hexadecimal	Ah				Ch			
Result		ACh							

Upper nibble bits = 1010b = **Ah**; Lower nibble bits = 1100b = **Ch**; the two Hex Nibble values are combined to create a single **ACh** POST Progress Code.

Early POST Memory Initialization MRC Diagnostic Codes

Memory Initialization at the beginning of POST includes multiple functions, including: discovery, channel training, validation that the DIMM population is acceptable and functional, initialization of the IMC and other hardware settings, and initialization of applicable RAS configurations.

The MRC Progress Codes are displayed via the Diagnostic LEDs that show the execution point in the MRC operational path at each step.

Table 70. MRC Progress Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
MRC Progress Codes									
B0h	1	0	1	1	0	0	0	0	Detect DIMM population
B1h	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2h	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3h	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4h	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5h	1	0	1	1	0	1	0	1	Program registers on the channel level
B6h	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7h	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8h	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9h	1	0	1	1	1	0	0	1	Hardware memory test and init
BAh	1	0	1	1	1	0	1	0	Execute software memory init
BBh	1	0	1	1	1	0	1	1	Program memory map and interleaving
BCh	1	0	1	1	1	1	0	0	Program RAS configuration
BFh	1	0	1	1	1	1	1	1	MRC is done

Should a major memory initialization error occur, preventing the system from booting with data integrity, a beep code is generated, the MRC will display a fatal error code on the diagnostic LEDs, and a system halt command is executed. Fatal MRC error halts do NOT change the state of the System Status LED, and they do NOT get logged as SEL events. Figure 89 lists all MRC fatal errors that are displayed to the Diagnostic LEDs.

NOTE: Fatal MRC errors will display POST error codes that may be the same as BIOS POST progress codes displayed later in the POST process. The Fatal MRC codes can be distinguished from the BIOS POST progress codes by the accompanying memory failure beep code of three long beeps, as identified in Table 71.

Table 71. MRC Fatal Error Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
MRC Fatal Error Codes									
E8h	1	1	1	0	1	0	0	0	No usable memory error 01h = No memory was detected from SPD read, or invalid config that causes no operable memory. 02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error. 03h = No memory installed. All channels are disabled.
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel Trusted Execution Technology and is inaccessible
EAh	1	1	1	0	1	0	1	0	DDR4 channel training error 01h = Error on read DQ/DQS (Data/Data Strobe) init 02h = Error on Receive Enable 03h = Error on Write Leveling 04h = Error on write DQ/DQS (Data/Data Strobe)
EBh	1	1	1	0	1	0	1	1	Memory test failure 01h = Software memtest failure. 02h = Hardware memtest failed.
EDh	1	1	1	0	1	1	0	1	DIMM configuration population error 01h = Different DIMM types (RDIMM, LRDIMM) are detected installed in the system. 02h = Violation of DIMM population rules. 03h = The 3rd DIMM slot cannot be populated when QR DIMMs are installed. 04h = UDIMMs are not supported. 05h = Unsupported DIMM Voltage.
EFh	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error

BIOS POST Progress Codes

Table 72 provides a list of all POST progress codes.

Table 72. POST Progress Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	
SEC Phase									
01h	0	0	0	0	0	0	0	1	First POST code after CPU reset
02h	0	0	0	0	0	0	1	0	Microcode load begin
03h	0	0	0	0	0	0	1	1	CRAM initialization begin
04h	0	0	0	0	0	1	0	0	EI Cache When Disabled
05h	0	0	0	0	0	1	0	1	SEC Core at Power on Begin
06h	0	0	0	0	0	1	1	0	Early CPU initialization during Sec Phase.
UPI RC (Fully leverage without platform change)									
A1h	1	0	1	0	0	0	0	1	Collect info such as SBSP, Boot Mode, Reset type etc
A3h	1	0	1	0	0	0	1	1	Setup minimum path between SBSP & other sockets
A7h	1	0	1	0	0	1	1	1	Topology discovery and route calculation
A8h	1	0	1	0	1	0	0	0	Program final route
A9h	1	0	1	0	1	0	0	1	Program final IO SAD setting
AAh	1	0	1	0	1	0	1	0	Protocol layer and other uncore settings
ABh	1	0	1	0	1	0	1	1	Transition links to full speed operation
Ach	1	0	1	0	1	1	0	0	Phy layer setting
ADh	1	0	1	0	1	1	0	1	Link layer settings
A Eh	1	0	1	0	1	1	1	0	Coherency settings
AFh	1	0	1	0	1	1	1	1	UPI initialization done
07h	0	0	0	0	0	1	1	1	Early SB initialization during Sec Phase.
08h	0	0	0	0	1	0	0	0	Early NB initialization during Sec Phase.
09h	0	0	0	0	1	0	0	1	End Of Sec Phase.
0Eh	0	0	0	0	1	1	1	0	Microcode Not Found.
0Fh	0	0	0	0	1	1	1	1	Microcode Not Loaded.
PEI Phase									
10h	0	0	0	1	0	0	0	0	PEI Core
11h	0	0	0	1	0	0	0	1	CPU PEIM
15h	0	0	0	1	0	1	0	1	NB PEIM
19h	0	0	0	1	1	0	0	1	SB PEIM
MRC Progress Codes									
31h	0	0	1	1	0	0	0	1	Memory Installed
32h	0	0	1	1	0	0	1	0	CPU PEIM (CPU Init)
33h	0	1	0	0	1	1	1	1	CPU PEIM (Cache Init)
4Fh	0	1	0	0	1	1	1	1	Dxe IPL started

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	
DXE Phase									
60h	0	1	1	0	0	0	0	0	DXE Core started
61h	0	1	1	0	0	0	0	1	DXE NVRAM Init
62h	0	1	1	0	0	0	1	0	DXE Setup Init
63h	0	1	1	0	0	1	0	1	DXE CPU Init
65h	0	1	1	0	0	1	0	1	DXE CPU BSP Select
66h	0	1	1	0	0	1	1	0	DXE CPU AP Init
68h	0	1	1	0	1	0	0	0	DXE PCI Host Bridge Init
69h	0	1	1	0	1	0	0	1	DXE NB Init
6Ah	0	1	1	0	1	0	1	0	DXE NB SMM Init
70h	0	1	1	1	0	0	0	0	DXE SB Init
71h	0	1	1	1	0	0	0	1	DXE SB SMM Init
72h	0	1	1	1	0	0	1	0	DXE SB devices Init
78h	0	1	1	1	1	0	0	1	DXE ACPI Init
79h	0	1	1	1	1	0	0	1	DXE CSM Init
80h	1	0	0	0	0	0	0	0	DXE BDS Started
81h	1	0	0	0	0	0	0	1	DXE BDS connect drivers
82h	1	0	0	0	0	0	1	0	DXE PCI Bus begin
83h	1	0	0	0	0	0	1	1	DXE PCI Bus HPC Init
84h	1	0	0	0	0	1	0	0	DXE PCI Bus enumeration
85h	1	0	0	0	0	1	0	1	DXE PCI Bus resource requested
86h	1	0	0	0	0	1	1	0	DXE PCI Bus assign resource
87h	1	0	0	0	0	1	1	1	DXE CON_OUT connect
88h	1	0	0	0	1	0	0	0	DXE CON_IN connect
89h	1	0	0	0	1	0	0	1	DXE SIO Init
8Ah	1	0	0	0	1	0	1	0	DXE USB start
8Bh	1	0	0	0	1	0	1	1	DXE USB reset
8Ch	1	0	0	0	1	1	0	0	DXE USB detect
8Dh	1	0	0	0	1	1	0	1	DXE USB enable
91h	1	0	0	1	0	0	0	1	DXE IDE begin
92h	1	0	0	1	0	0	1	0	DXE IDE reset
93h	1	0	0	1	0	0	1	1	DXE IDE detect
94h	1	0	0	1	0	1	0	0	DXE IDE enable
95h	1	0	0	1	0	1	0	1	DXE SCSI begin
96h	1	0	0	1	0	1	1	0	DXE SCSI reset
97h	1	0	0	1	0	1	1	1	DXE SCSI detect
98h	1	0	0	1	1	0	0	0	DXE SCSI enable
99h	1	0	0	1	1	0	0	1	DXE verifying SETUP password
9Bh	1	0	0	1	1	0	1	1	DXE SETUP start
9Ch	1	0	0	1	1	1	0	0	DXE SETUP input wait
9Dh	1	0	0	1	1	1	0	1	DXE Ready to Boot

Checkpoint	Diagnostic LED Decoder								
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	Description
9Eh	1	0	0	1	1	1	1	0	DXE Legacy Boot
9Fh	1	0	0	1	1	1	1	1	DXE Exit Boot Services
C0h	1	1	0	0	0	0	0	0	RT Set Virtual Address Map Begin
C2h	1	1	0	0	0	0	1	0	DXE Legacy Option ROM init
C3h	1	1	0	0	0	0	1	1	DXE Reset system
C4h	1	1	0	0	0	1	0	0	DXE USB Hot plug
C5h	1	1	0	0	0	1	0	1	DXE PCI BUS Hot plug
C6h	1	1	0	0	0	1	1	0	DXE NVRAM cleanup
C7h	1	1	0	0	0	1	1	1	DXE ACPI Enable
0h	0	0	0	0	0	0	0	0	Clear POST Code
S3 Resume									
40h	0	1	0	0	0	0	0	0	S3 Resume PEIM (S3 started)
41h	0	1	0	0	0	0	0	1	S3 Resume PEIM (S3 boot script)
42h	0	1	0	0	0	0	1	0	S3 Resume PEIM (S3 Video Repost)
43h	0	1	0	0	0	0	1	1	S3 Resume PEIM (S3 OS wake)
BIOS Recovery									
46h	0	1	0	0	0	1	1	0	PEIM which detected forced Recovery condition
47h	0	1	0	0	0	1	1	1	PEIM which detected User Recovery condition
48h	0	1	0	0	1	0	0	0	Recovery PEIM (Recovery started)
49h	0	1	0	0	1	0	0	1	Recovery PEIM (Capsule found)
4Ah	0	1	0	0	1	0	1	0	Recovery PEIM (Capsule loaded)
E8h	1	1	1	0	1	0	0	0	No Usable Memory Error:
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel® Trusted Execution Technology and is inaccessible.
EAh	1	1	1	0	1	0	1	0	DDR4 Channel Training Error:
EBh	1	1	1	0	1	0	1	1	Memory Test Failure:
Edh	1	1	1	0	1	1	0	1	DIMM Configuration/Population Error:
EFh	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error.
B0h	1	0	1	1	0	0	0	0	Detect DIMM population
B1h	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2h	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3h	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4h	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5h	1	0	1	1	0	1	0	1	Program registers on the channel level
B6h	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7h	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8h	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9h	1	0	1	1	1	0	0	1	Hardware memory test and init
Bah	1	0	1	1	1	0	1	0	Execute software memory init
BBh	1	0	1	1	1	0	1	1	Program memory map and interleaving
BCh	1	0	1	1	1	1	0	0	Program RAS configuration

Checkpoint	Diagnostic LED Decoder								
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	Description
BFh	1	0	1	1	1	1	1	1	MRC is done

Appendix C. POST Code Errors

Most error conditions encountered during POST are reported using **POST Error Codes**. These codes represent specific failures, warnings, or are informational. POST Error Codes may be displayed in the Error Manager Display screen, and are always logged to the System Event Log (SEL). Logged events are available to System Management applications, including Remote and Out of Band (OOB) management.

There are exception cases in early initialization where system resources are not adequately initialized for handling POST Error Code reporting. These cases are primarily Fatal Error conditions resulting from initialization of processors and memory, and they are handled by a Diagnostic LED display with a system halt.

Table 73 lists the supported POST Error Codes. Each error code is assigned an error type which determines the action the BIOS will take when the error is encountered. Error types include Minor, Major, and Fatal. The BIOS action for each is defined as follows:

- **Minor:** The error message is displayed on the screen or on the Error Manager screen, and an error is logged to the SEL. The system continues booting in a degraded state. The user may want to replace the erroneous unit. The POST Error Pause option setting in the BIOS setup does not have any effect on this error.
- **Major:** The error message is displayed on the Error Manager screen, and an error is logged to the SEL. The POST Error **Pause** option setting in the BIOS setup determines whether the system pauses to the Error Manager for this type of error so the user can take immediate corrective action or the system continues booting.

Note that for 0048, "Password check failed", the system halts, and then after the next reset/reboot will display the error code on the Error Manager screen.

- **Fatal:** The system halts during POST at a blank screen with the text "**Unrecoverable fatal error found. System will not boot until the error is resolved**" and "**Press <F2> to enter Setup**". The POST Error Pause option setting in the BIOS setup does not have any effect with this class of error.

When the operator presses the **F2** key on the keyboard, the error message is displayed on the Error Manager screen, and an error is logged to the SEL with the error code. The system cannot boot unless the error is resolved. The user needs to replace the faulty part and restart the system.

NOTE: The POST Code errors in the following table are common to all current generation Intel server platforms. Features present on a given server board/system will determine which of the listed error codes are supported.

Table 73. POST Error Codes and Messages

Error Code	Error Message	Response
0012	System RTC date/time not set	Major
0048	Password check failed	Major
0140	PCI component encountered a PERR error	Major
0141	PCI resource conflict	Major
0146	PCI out of resources error	Major
0191	Processor core/thread count mismatch detected	Fatal
0192	Processor cache size mismatch detected	Fatal
0194	Processor family mismatch detected	Fatal
0195	Processor Intel(R) QPI link frequencies unable to synchronize	Fatal

Error Code	Error Message	Response
0196	Processor model mismatch detected	Fatal
0197	Processor frequencies unable to synchronize	Fatal
5220	BIOS Settings reset to default settings	Major
5221	Passwords cleared by jumper	Major
5224	Password clear jumper is set	Major
8130	Processor 01 disabled	Major
8131	Processor 02 disabled	Major
8160	Processor 01 unable to apply microcode update	Major
8161	Processor 02 unable to apply microcode update	Major
8170	Processor 01 failed Self Test (BIST)	Major
8171	Processor 02 failed Self Test (BIST)	Major
8180	Processor 01 microcode update not found	Minor
8181	Processor 02 microcode update not found	Minor
8190	Watchdog timer failed on last boot	Major
8198	OS boot watchdog timer failure	Major
8300	Server board management controller failed Self Test	Major
8305	Hot Swap Controller failure	Major
83A0	Management Engine (ME) failed Self Test	Major
83A1	Management Engine (ME) failed to respond.	Major
84F2	Server board management controller failed to respond	Major
84F3	Server board management controller in update mode	Major
84F4	Sensor data record empty	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8501	DIMM Population Error	Major
8520	DIMM_A1 failed test/initialization	Major
8521	DIMM_A2 failed test/initialization	Major
8523	DIMM_B1 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major
852A	DIMM_D2 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852F	DIMM_F1 failed test/initialization	Major
8540	DIMM_A1 disabled	Major
8541	DIMM_A2 disabled	Major
8543	DIMM_B1 disabled	Major
8546	DIMM_C1 disabled	Major
8549	DIMM_D1 disabled	Major
854A	DIMM_D2 disabled	Major
854C	DIMM_E1 disabled	Major
854F	DIMM_F1 disabled	Major
8560	DIMM_A1 encountered a Serial Presence Detection (SPD) failure	Major
8561	DIMM_A2 encountered a Serial Presence Detection (SPD) failure	Major
8563	DIMM_B1 encountered a Serial Presence Detection (SPD) failure	Major
8566	DIMM_C1 encountered a Serial Presence Detection (SPD) failure	Major

Error Code	Error Message	Response
8569	DIMM_D1 encountered a Serial Presence Detection (SPD) failure	Major
856A	DIMM_D2 encountered a Serial Presence Detection (SPD) failure	Major
856C	DIMM_E1 encountered a Serial Presence Detection (SPD) failure	Major
856F	DIMM_F1 encountered a Serial Presence Detection (SPD) failure	Major
8604	POST Reclaim of non-critical NVRAM variables	Minor
8605	BIOS Settings are corrupted	Major
8606	NVRAM variable space was corrupted and has been reinitialized	Major
8607	Recovery boot has been initiated. Note: The Primary BIOS image may be corrupted or the system may hang during POST. A BIOS update is required.	Fatal
92A3	Serial port component was not detected	Major
92A9	Serial port component encountered a resource conflict error	Major
A000	TPM device not detected.	Minor
A001	TPM device missing or not responding.	Minor
A002	TPM device failure.	Minor
A003	TPM device failed self test.	Minor
A100	BIOS ACM Error	Major
A421	PCI component encountered a SERR error	Fatal
A5A0	PCI express* component encountered a PERR error	Minor
A5A1	PCI express* component encountered an SERR error	Fatal
A6A0	DXE Boot Services driver: Not enough memory available to shadow a Legacy Option ROM.	Minor

POST Error Codes LED

Prior to system video initialization, the BIOS uses Post Error LED codes to inform users on error conditions. A user-visible error code is followed by the POST Progress LEDs.

Table 74 lists the POST Error LED codes

Table 74. POST Error LED Codes

POST Error LED Sequence	Error Message	POST Progress Code	Description
1 blink	USB device action	N/A	Short LED blink whenever USB device is discovered in POST, or inserted or removed during runtime.
1 long blink	Intel® TXT security violation	0xAE, 0xAF	System halted because Intel® Trusted Execution Technology detected a potential violation of system security.
3 blinks	Memory error	Multiple	System halted because a fatal error related to the memory was detected.
3 long blinks followed by 1	CPU mismatch error	0xE5, 0xE6	System halted because a fatal error related to the CPU family/core/cache mismatch was detected.
The following POST Error LED Codes are lighted during BIOS Recovery			
2 blinks	Recovery started	N/A	Recovery boot has been initiated.
4 blinks	Recovery failed	N/A	Recovery has failed. This typically happens so quickly after recovery is initiated that it lights like a 2-4 LED code.

The Integrated BMC may generate POST Error Codes upon detection of failure conditions. These codes are translated into visual LED sequences each time the problem is discovered, such as on each power-up attempt, but are not lit continuously. Codes that are common across all Intel server boards and systems that use the same generation of chipset are listed in Table 75. Each digit in the code is represented by a LED lit/off sequence of whose count is equal to the digit.

Table 75. Integrated BMC Error (LED) Codes

Code	Associated Sensors	Reason for Error (LED lit)
1-5-2-1	No CPUs installed or first CPU socket is empty.	CPU1 socket is empty, or sockets are populated incorrectly CPU1 must be populated before CPU2.
1-5-2-4	MSID Mismatch	MSID mismatch occurs if a processor is installed into a system board that has incompatible power capabilities.
1-5-4-2	Power fault	DC power unexpectedly lost (power good dropout) – Power unit sensors report power unit failure offset
1-5-4-4	Power control fault (power good assertion timeout).	Power good assertion timeout – Power unit sensors report soft power control failure offset
1-5-1-2	VR Watchdog Timer sensor assertion	VR controller DC power on sequence was not completed in time.
1-5-1-4	Power Supply Status	The system does not power on or unexpectedly powers off and a Power Supply Unit (PSU) is present that is an incompatible model with one or more other PSUs in the system.




Figure 89. POST Diagnostic Error (LED Location)

Appendix D. *Product Family Statements of Volatility*

This Appendix describes the volatile and non-volatile components on the Intel® Server S2600BP product family. It is not the intention of this document to include any components not directly mounted to the listed Intel® server board family or associated system boards used within the compute modules or supported Intel chassis, these may include: processors, memory, storage devices, or add-in cards.

The tables provide the following data for each identified component.

Component Type

Three types of memory components are used on the server board assembly. These include:

- **Non-volatile:** Non-volatile memory is persistent, and is not cleared when power is removed from the system. Non-Volatile memory must be erased to clear data. The exact method of clearing these areas varies by the specific component. Some areas are required for normal operation of the server, and clearing these areas may render the server board inoperable.
- **Volatile:** Volatile memory is cleared automatically when power is removed from the system.
- **Battery powered RAM:** Battery powered RAM is similar to volatile memory, but is powered by a battery on the server board. Data in Battery powered Ram is persistent until the battery is removed from the server board.

Size

The size of each component includes sizes in bits, Kbits, bytes, kilobytes (KB) or megabytes (MB).

Board Location

The physical location of each component is specified in the Board Location column. The board location information corresponds to information on the server board silkscreen.

User Data

The flash components on the server boards do not store user data from the operating system. No operating system level data is retained in any listed components after AC power is removed. The persistence of information written to each component is determined by its type as described in the table.

Each component stores data specific to its function. Some components may contain passwords that provide access to that device's configuration or functionality. These passwords are specific to the device and are unique and unrelated to operating system passwords. The specific components that may contain password data are:

- **BIOS:** The server board BIOS provides the capability to prevent unauthorized users from configuring BIOS settings when a BIOS password is set. This password is stored in BIOS flash, and is only used to set BIOS configuration access restrictions.

BMC: The server boards support an Intelligent Platform Management Interface (IPMI) 2.0 conformant baseboard management controller (BMC). The BMC provides health monitoring, alerting and remote power control capabilities for the Intel® server board. The BMC does not have access to operating system level data.

The BMC supports the capability for remote software to connect over the network and perform health monitoring and power control. This access can be configured to require authentication by a password. If configured, the BMC will maintain user passwords to control this access. These passwords are stored in the BMC flash.

Intel® Server Board S2600BP Family

Intel® Server boards include several components that can be used to store data. A list of components for the Intel® Server Board S2600BP is included in **Error! Reference source not found.**

Table 76. Server Board Component

Component Type	Size	Board Location	User Data	Name
Non-Volatile	32 MB	U5B2	No (firmware)	Firmware Flash(BMC)
Non-Volatile	32 MB	U6B1	No (BIOS)	BIOS Flash
Non-Volatile	16MB	U6M1	No	NIC EEPROM
Volatile	16 MB	U5B1	No	Firmware SDRAM

Intel® Compute Module HNS2600BP Family

System boards within Intel® compute modules may include components that can be used to store data. The following tables provide a list of components associated with specific system board or system board options supported with this product family. For server board components, see the section above.

Table 77. System Boards Components

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256 Bytes	U4A1	No	Slot1 Riser
Non-Volatile	256 Bytes	U2A2	No	Slot2 M.2 Riser
Non-Volatile	128 Bytes	U6L2	No	3408 Bridge Board
Non-Volatile	8K Bytes	U7A1	No	3408 Bridge Board
Non-Volatile	4K Bytes	U2A1,U3A4	No	3408 Bridge Board
Non-Volatile	128k Bytes	U3B1	No	3408 Bridge Board
Non-Volatile	32M Bytes	U3A3	No	3408 Bridge Board
Non-Volatile	256 Bytes	U1A2	No	24HDD HSBP
Non-Volatile	8K Bytes	U3L1	No	6Port PT Bridge Board
Non-Volatile	1K Bytes	U3L8, U8L1	No	6Port PT Bridge Board
Non-Volatile	8K Bytes	U12,U20	No	6Port 3008 Bridge Board
Non-Volatile	32K Bytes	U502	No	6Port 3008 Bridge Board
Non-Volatile	4K Bytes	U6,U7	No	6Port 3008 Bridge Board
Non-Volatile	16M Bytes	U8	No	6Port 3008 Bridge Board
Non-Volatile	8K Bytes	U2A6	No	4Port 3008 Bridge Board
Non-Volatile	32K Bytes	U8L1	No	4Port 3008 Bridge Board
Non-Volatile	16M Bytes	U2A3	No	4Port 3008 Bridge Board
Non-Volatile	128 Bytes	U2A1	No	4Port 3008 Bridge Board (RAID 0/1/10)
Non-Volatile	128 Bytes	U2A2	No	4Port 3008 Bridge Board (RAID 0/1/10/5)

Intel® Server Chassis H2000G Family

System boards within Intel server chassis contains several components that can be used to store data. A list of components for the HSBP and Power Supply Unit of the Server Chassis is included in the table below. For server board components and compute module components, see the sections above.

Table 78. Server Chassis Components

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256 Bytes	UM801	No	PSU Firmware
Non-Volatile	512 Bytes	U6N2	No	12 x 3.5" HSBP FRU
Non-Volatile	512 Bytes	U504	No	16 x 2.5" HSBP FRU
Non-Volatile	512 Bytes	U1A2	No	24 x 2.5" HSBP FRU
Non-Volatile	512 Bytes	U6N2	No	4 x 3.5" HSBP FRU

Appendix E. Glossary of Terms

This appendix provides a list of Acronyms and Terms used throughout this document.

Acronym	Description
AC	Alternating Current, a type of electrical current in which the current repeatedly changes direction
ACM	Authenticated Code Mode
ACPI	Advanced Configuration Power Interface
AP	Application Processor
BDS	Boot Device Selection
BIB	Burn in Board
BIOS	Basic Input / Output System – Firmware interface to the system hardware
BIST	Built-in Self Test
BMC	Server board Management Controller
BSP	Boot strap processor. The processor selected at boot time to be the primary processor in a multi-processor system.
BTU/hour	A unit of power. 1 watt is approximately 3.41214 BTU/h[, and 1000 BTU/h is approximately 293.071 W
CLTT	Closed Loop Thermal Throttling
CMOS	Complementary Metal-oxide-semiconductor
CPU	Central Processing Unit
CRAM	Configuration RAM - a programmable bit inside an FPGA that controls its behavior
CSM	Compatibility Support Module
DC	Direct current, the flow of electric charge is only in one direction.
DIMM	Dual In-line Memory Module, a plug-in memory module with signal and power pins on both sides of the internal printed circuit board (front and back).
DQ	Data Quality
DQS	Bi-directional Data Strobe
DXE	Driver Execution Environment. Component of Intel® Platform Innovation Framework for EFI architecture
EI	Enhanced Intel
ESD	Electrostatic Discharge
FRU	Field Replaceable Unit
GT/s	GigaTransfers per second
HBA	Hot Bus Adapter
HPC	High Performance Computing
IDE	Integrated Drive Electronics, a disk interface standard
IFT	Intel Fabric Through
IMC	Integrated Memory Controller – memory controller integrated into the processor chip
IPL	Initial Program Load
IPMB	Intelligent Platform Management Bus
ISTA	International Safe Transit Association

Acronym	Description
JEDEC	Joint Electron Device Engineering Council, industry organization for memory standards
KVM	Keyboard, Video, and Mouse – an attachment that mimics those devices, and connects them to a remote I/O user
LAN	Local Area Network
LED	Light Emitting Diode
LOM	LAN on Board
LRDIMM	Load Reduced DIMM memory modules have buffer registers for both address and data between the SDRAM modules and the system's memory controller.
ME	Management Engine
MM#	Material Management number
MRC	Memory Reference Code
MSB	Most Significant Bit
MSID	CPU Icc Mismatch
MT/s	MegaTransfers per second
NB	Northbound
NIC	Network Interface Card
NVRAM	Non-volatile RAM
OEM	Original Equipment Manufacturer
OFU	One-Boot Flash Update
OLTT	Open Loop Thermal Throttling
OOB	Out of Band
OS	Operating System
PCH	Platform Controller Hub
PCI	Peripheral Component Interconnect, or PCI Local Bus Standard – also called “Conventional PCI”
PCIe*	PCI Express* -- an updated form of PCI offering better throughput and better error management
PEI	Pre EFI Initialization. Component of Intel® Platform Innovation Framework for EFI architecture.
PEIM	PEI Module
PERR	Parity Error
PHM	Processor Heatsink Module
POST	Power On Self Test – BIOS activity from the time on Power On until Operating System boot begins.
PSU	Power Supply Unit
QPI	Intel® QuickPath Interconnect
QR	Quad Rank – memory DIMM organization, DRAMs organized in four ranks
RAID	Redundant Array of Inexpensive Disks – provides data security by spreading data over multiple disk drives. RAID 0, RAID 1, RAID 10, and RAID 5 are different patterns of data on varying numbers of disks to provide varying degrees of security and performance.
RAM	Random Access Memory
RAS	Reliability, Availability, and Serviceability
RC	Raw Class

Acronym	Description
RDIMM	Registered DIMM (also called buffered) memory modules have an address buffer register between the SDRAM modules and the system's memory controller.
ROM	Read-Only Memory
RT	Runtime. Component of Intel® Platform Innovation Framework for EFI architecture
RTC	Real Time Clock
SAD	Source Address Decoder
SAS	Serial Attached SCSI, a high speed serial data version of SCSI
SATA	Serial ATA, a high speed serial data version of the disk ATA interface
SB	Southbound
SBSP	System Boot-Strap Processor
SCSI	Small Computer System Interface, a connection usually used for disks of various types
SDR	Sensor Data Record
SEC	Security. Component of Intel® Platform Innovation Framework for EFI architecture
SEL	System Event Log
SERR	System Error
SFF	Small Form Factor
SFP+	The enhanced small form-factor pluggable (SFP+) is an enhanced version of the SFP that supports data rates up to 16 Gbit/s.
SIO	Super I/O
SMM	System Management Mode
SPD	Serial Presence Detect
SUP	System Updated Package
TDP	Thermal Design Power
TIM	Thermal Interface Material
TPM	Trusted Platform Module
TXT	Intel® Trusted Execution Technology
UPI	Intel® UltraPath® Interconnect
USB	Universal Serial Bus, a standard serial expansion bus meant for connecting peripherals.
VGA	Video Graphics Array
VR	Voltage Regulator

Appendix F. *Reference Documents*

Refer to the following documents for additional information:

- Intel® Server Chassis H2000G Product Family Service Guide
- Intel® Server Board S2600BP Product Family and Intel® Server Chassis H2000G Product Family Technical Product Specification and Addendum
- Intel® Server Board S2600BP Product Family and Intel® Compute Module HNS2600BP Product Family Configuration Guide
- Intel® Server Board S2600BP Product Family and Intel® Compute Module HNS2600BP Product Family Service Guide
- Intel® Server Board S2600BP Product Family and Intel® Compute Module HNS2600BP Product Family Technical Update
- Intel® Servers System BMC Firmware EPS for Intel® Xeon® Processor Scalable Family
- Intel® Server System BIOS EPS for Intel® Xeon® Processor Scalable Family
- Platform Controller Hub External Design Specification
- Intel® Xeon® Processor Scalable Family External Design Specification
- Intel® Ethernet Controller X550 Product Info
- Intel® Purley BIOS Setup Design for Intel® Xeon® Processor Scalable Family

Notes

[illegible]

Mouser Electronics

Authorized Distributor

Click to View Pricing, Inventory, Delivery & Lifecycle Information:

[Intel:](#)

[HNS2600BPS24](#) [BBS2600BPQ](#) [HNS2600BPQ24](#) [HNS2600BPB](#) [HNS2600BPB24](#) [BBS2600BPB](#) [HNS2600BPS](#)
[HNS2600BPQ](#) [BBS2600BPS](#) [HNS2600BPB24R](#) [HNS2600BPQ24R](#) [HNS2600BPQR](#) [HNS2600BPS24R](#)
[HNS2600BPSR](#) [HNS2600BPBLC24R](#) [HNS2600BPBLCR](#) [HNS2600BPBR](#) [HNS2600BPB24RX](#)