



Intel® Server Board S2600WF 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 Change
July 2017	1.0	Production 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.

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

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

© Intel Corporation. All Rights Reserved.

Table of Contents

1. Introduction	13
1.1 Document Outline	14
1.2 Intel Server Board Use Disclaimer	14
1.3 Product Errata.....	14
2. Server Board Family Overview	15
2.1 Server Board Family Feature Set.....	17
2.2 Server Board Component / Feature Identification.....	19
2.3 Server Board Mechanical Drawings.....	23
2.4 Product Architecture Overview	27
2.5 System Software Stack	28
2.5.1 Hot Keys Supported During POST.....	29
2.5.2 Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data.....	31
3. Processor Support.....	32
3.1 Processor Socket and Processor Heat Sink Module (PHM) Assembly	32
3.2 Processor Thermal Design Power (TDP) Support	34
3.3 Intel® Xeon® Processor Scalable Family Overview	35
3.3.1 Intel® 64 Instruction Set Architecture (Intel® 64)	36
3.3.2 Intel® Hyper-Threading Technology	36
3.3.3 Enhanced Intel® SpeedStep Technology.....	36
3.3.4 Intel® Turbo Boost Technology 2.0	36
3.3.5 Intel® Virtualization Technology (Intel® VT-x)	36
3.3.6 Intel® Virtualization Technology for Directed I/O (Intel® VT-d)	36
3.3.7 Execute Disable Bit	37
3.3.8 Intel® Trusted Execution Technology for servers (Intel® TXT)	37
3.3.9 Intel® Advanced Vector Extension (Intel AVX-512)	37
3.3.10 Advanced Encryption Standard New Instructions (AES-NI).....	37
3.3.11 Intel® Node Manager 4.0	37
3.4 Processor Population Rules.....	38
3.5 Processor Initialization Error Summary	39
3.6 Intel® Xeon® processor Scalable Family with Integrated Intel® Omni-Path Host Fabric Interface 42	
3.6.1 Intel®Omni-Path IFT Carrier Accessory Kits	44
4. System Memory	47
4.1 Memory Sub-system Architecture	47
4.2 Supported Memory	48
4.3 Memory Slot Identification and Population Rules	48
4.3.1 DIMM Population Guidelines for Best Performance.....	50
4.4 Memory RAS Features.....	51
4.4.1 DIMM Populations Rules and BIOS Setup for Memory RAS	52
5. PCIe* Support.....	53

5.1.1	PCIe* Enumeration and Allocation	53
5.1.2	Non-Transparent Bridge	54
6.	System I/O	55
6.1	PCIe* Add-in Card Support	55
6.1.1	Riser Card Support	57
6.1.2	Intel® OCP Module Support	59
6.1.3	Intel® Integrated RAID Module Support	60
6.2	Onboard Storage Sub-System	61
6.2.1	M.2 SSD Support	61
6.2.2	On Board PCIe* OCuLink* Connectors	62
6.2.3	Intel® Volume Management Device (Intel® VMD) for NVMe	63
6.2.4	Intel® Virtual RAID on Chip (Intel® VROC) For NVMe	65
6.2.5	Onboard SATA Support	67
6.2.6	On-Board SATA RAID Options	69
6.3	Rear External RJ45 Connector Overview	72
6.3.1	RJ45 Dedicated Management Port	73
6.3.2	RJ45 Network Interface Connectors (S2600WFT only)	73
6.3.3	Serial Port Support	73
6.4	USB Support	75
6.4.1	External USB 3.0 Support	75
6.4.2	Internal USB 2.0 Type-A Connector	75
6.4.3	Front Panel USB 3.0 Support	76
6.4.4	Front Panel USB 2.0 Connector	76
6.5	Video Support	77
6.5.1	Onboard Video Connectors	78
6.5.2	Onboard Video and Add-in Video Adapter Support	79
6.5.3	Dual Monitor Support	79
7.	On-board Connector/Header Pin-Out Definition	80
7.1	Power Connectors	80
7.1.1	Main Power	80
7.1.2	Hot Swap Backplane Power Connector	82
7.1.3	Riser Card Supplemental 12V Power Connectors	83
7.1.4	Peripheral Power Connector	84
7.2	Front Control Panel Headers and Connectors	85
7.2.1	Front Panel LED and Control Button Features Overview	86
7.3	System Fan Connectors	88
7.4	Management Connectors	88
8.	Standard and Advanced Server Management Features	90
8.1	Dedicated Management Port	91
8.2	Embedded Web Server	92
8.3	Advanced Management Feature Support (Intel® RMM4 Lite)	93

8.3.1	Keyboard, Video, Mouse (KVM) Redirection	94
8.3.2	Remote Console	94
8.3.3	Performance.....	95
8.3.4	Availability	95
8.3.5	Security.....	95
8.3.6	Usage.....	95
8.3.7	Force-enter BIOS Setup	95
8.3.8	Media Redirection	95
9.	Light Guided Diagnostics.....	97
9.1	System ID LED	98
9.2	System Status LED.....	98
9.3	BMC Boot/Reset Status LED Indicators	100
9.4	Post Code Diagnostic LEDs	100
9.5	Fan Fault LEDs.....	101
9.6	Memory Fault LEDs	101
9.7	CPU Fault LEDs.....	101
10.	System Security	102
10.1	Password Setup	103
10.1.1	System Administrator Password Rights	103
10.1.2	Authorized System User Password Rights and Restrictions.....	104
10.2	Front Panel Lockout.....	104
10.3	Trusted Platform Module (TPM) Support	104
10.3.1	TPM Security BIOS.....	105
10.3.2	Physical Presence	105
10.3.3	TPM Security Setup Options	105
10.4	Intel® Trusted Execution Technology	106
11.	Reset and Recovery Jumpers.....	107
11.1	BIOS Default Jumper Block	107
11.2	Password Clear Jumper Block.....	108
11.3	Management Engine (ME) Firmware Force Update Jumper Block.....	108
11.4	BMC Force Update Jumper Block	109
11.5	BIOS Recovery Jumper	110
12.	Platform Management.....	111
12.1	Management Feature Set Overview	111
12.1.1	IPMI 2.0 Features Overview	111
12.1.2	Non-IPMI Features Overview	112
12.2	Platform Management Features and Functions.....	113
12.2.1	Power Subsystem	113
12.2.2	Advanced Configuration and Power Interface (ACPI)	114
12.2.3	Watchdog Timer	115
12.2.4	System Event Log (SEL)	115
12.3	Sensor Monitoring	115

12.3.1	Sensor Rearm Behavior	116
12.3.2	Thermal Monitoring	116
12.3.3	Standard Fan Management.....	117
12.3.4	Memory Thermal Management.....	119
12.3.5	Power Management Bus (PMBus*).....	120
12.3.6	Component Fault LED Control	120
Appendix A – Integration and Usage Tips		122
Appendix B – POST Code Diagnostic LED Decoder.....		123
Appendix C – POST Code Errors.....		130
Appendix D – Statement of Volatile Memory Components		135
Appendix E – Supported Intel Server Systems		137

List of Figures

Figure 1. Intel® Server Board S2600WF.....	15
Figure 2. Intel® Server Board S2600WF with Available On-Board Options.....	16
Figure 3. Server Board Component / Feature Identification.....	19
Figure 4. Intel® Server Board S2600WF External I/O Connector Layout.....	20
Figure 5. Intel® Light Guided Diagnostics - DIMM Fault LEDs.....	20
Figure 6. Intel® Light Guided Diagnostic LED Identification.....	21
Figure 7. Board Configuration and Recovery Jumpers	22
Figure 8. Intel® Server Board S2600WF – Primary Side Keepout Zone.....	23
Figure 9. Intel® Server Board S2600WF – Hole and Component Positions.....	24
Figure 10. Intel® Server Board S2600WF – Secondary Side Keepout Zone.....	25
Figure 11. Intel® Server Board S2600WF – Primary Side Height Restrictions.....	26
Figure 12. Intel® Server Board S2600WF Product Family Architectural Block Diagram.....	27
Figure 13. PHM Components and Processor Socket Reference Diagram.....	32
Figure 14. PHM to CPU Socket Orientation and Alignment Features	33
Figure 15. Processor Socket Assembly and Protective Cover	34
Figure 16. Intel®Omni-Path IFT Carrier Accessory Kit Components.....	44
Figure 17. Server Board Sideband Connectors	45
Figure 18. IFT Carrier Board – Rear View.....	45
Figure 19. Memory Sub-system Architecture	47
Figure 20. Intel® Server Board S2600WF Memory Slot Layout	48
Figure 21. PCIe* Add-in Card Support	55
Figure 22. 1U One Slot PCIe* Riser Card (iPC – F1UL16RISER3APP).....	57
Figure 23. 2U Three PCIe* Slot Riser Card (iPC – A2UL8RISER2).....	57
Figure 24. 2U Two PCIe* Slot Riser Card (iPC – A2UL16RISER2).....	58
Figure 25. 2U Two PCIe* Slot (Low Profile) PCIe* Riser Card (iPC – A2UX8X4RISER).....	58
Figure 26. Intel® OCP Module Connector.....	59
Figure 27. Intel® Integrated RAID Module.....	60
Figure 28. M.2 Storage Device Connectors	61
Figure 29. On-Board OCuLink Connectors.....	62
Figure 30. VMD Support Disabled in <F2> BIOS Setup	64
Figure 31. VMD Support Enabled in <F2> BIOS Setup	65
Figure 32. Intel® VROC Upgrade Key.....	66
Figure 33. On-Board SATA Port Connector Identification.....	68
Figure 34. ESRT2 SATA RAID-5 Upgrade Key	71
Figure 35. Rear External RJ45 Connectors.....	72
Figure 36. RJ45 Connector LEDs.....	72
Figure 37. RJ45 Serial-A Pin Orientation	73
Figure 38. J4A2 Jumper Block for Serial A Pin 7 Configuration	74
Figure 39. Serial-B Connector (Internal).....	74

Figure 40. External USB 3.0 Ports	75
Figure 41. Internal USB 2.0 Type-A Connector	75
Figure 42. Front Panel USB 3.0 Connector.....	76
Figure 43. Front Panel USB 2.0 Connector.....	77
Figure 44. Rear External Video Connector	78
Figure 45. Front Panel Video Connector.....	78
Figure 46. MAIN PWR 1” and “MAIN PWR 2” Connectors	80
Figure 47. HSBP PWR Connector	82
Figure 48. 12V Power Connectors	83
Figure 49. High Power Add-in Card 12V Auxiliary Power Cable Option.....	84
Figure 50. Peripheral Power Connector	84
Figure 51. Front Control Panel Connectors	85
Figure 52. Example - Front Control Panel View (For reference purposes only)	86
Figure 53. System Fan Connector Pin-outs	88
Figure 54. RMM 4 Lite Key Placement.....	90
Figure 55. Dedicated Managment Port	91
Figure 56. On-Board Diagnostic LED Placement	97
Figure 57.DIMM Fault LEDs.....	98
Figure 58. BIOS Setup Utility Security Tab.....	102
Figure 59. Reset and Recovery Jumper Block Location.....	107
Figure 60. High-level Fan Speed Control Process.....	119
Figure 61. On-board POST Diagnostic LEDs	123
Figure 62. Intel® Server System R1000WF Product Family.....	137
Figure 63. Intel® Server System R2000WF Product Family.....	139

List of Tables

Table 1. Reference Documents	13
Table 2. Server Board Product Family Feature Set.....	17
Table 3. POST Hot-Keys.....	29
Table 4. Intel® Xeon® Processor Scalable Family - Feature Comparison Table.....	35
Table 5. Mixed Processor Configurations Error Summary.....	40
Table 6. Intel® Xeon® Processor Scalable Family w/ Integrated Intel® Omni-Path Fabric.....	42
Table 7. IFT Carrier LED Functionality.....	45
Table 8. Power Level Classification for QSFP+ Modules.....	45
Table 9. Supported Processor Mixing – Fabric vs Non-Fabric Processors	46
Table 10. DDR4 RDIMM and LRDIMM Support.....	48
Table 11. Memory RASM Features.....	51
Table 12. CPU - PCIe* Port Routing.....	53
Table 13. Riser Slot #1 – PCIe* Root Port Mapping.....	56
Table 14. Riser Slot #2 – PCIe* Root Port Mapping.....	56
Table 15. Riser Slot #3 – PCIe* Root Port Mapping.....	56
Table 16. Supported Intel® OCP Modules	59
Table 17. Intel® VROC Upgrade Key Options	67
Table 18. SATA and sSATA Controller Feature Support.....	68
Table 19. SATA and sSATA Controller BIOS Utility Setup Options	69
Table 20. External RJ45 NIC Port LED Definition	72
Table 21. Serial-A Connector Pin-out	73
Table 22. Serial-B Connector Pin-out	74
Table 23. Front Panel USB 2.0/3.0 Connector Pin-out ("FP_USB_2.0/ 3.0")	76
Table 24. Front Panel USB 2.0 Connector Pin-out ("FP_USB_2.0_5-6 ").....	77
Table 25. Front Panel Video Connector Pin-out ("FP VIDEO").....	78
Table 26. Main Power (Slot 1) Connector Pin-out ("MAIN PWR 1")	81
Table 27. Main Power (Slot 2) Connector Pin-out ("MAIN PWR 2")	81
Table 28. Hot Swap Backplane Power Connector Pin-out.....	82
Table 29. Riser Slot Auxiliary Power Connector Pin-out ("OPT_12V_PWR").....	83
Table 30. Peripheral Drive Power Connector Pin-out ("Peripheral_PWR").....	84
Table 31. Front Panel Control Button and LED Support.....	85
Table 32. 30-pin Front Panel Connector Pin-outs.....	86
Table 33. Power/Sleep LED Functional States	86
Table 34. NMI Signal Generation and Event Logging	87
Table 35. Hot Swap Backplane I2C Connector – SMBUS 3-pin (J5C3)	89
Table 36. Hot Swap Backplane I2C Connector – SMBUS 4-pin (J1K1)	89
Table 37. IPMB – SMBUS 4-pin (J1C3).....	89
Table 38. Intel® Remote Management Module 4 (RMM4) Options	90
Table 39. Standard and Advanced Server Management Features.....	91

Table 40. System Status LED States	99
Table 41. BMC Boot/Reset Status LED Indicators	100
Table 42. Power Control Sources.....	113
Table 43. ACPI Power States.....	114
Table 44. Component Fault LEDs.....	121
Table 45. POST Progress Code LED Example.....	123
Table 46. MRC Progress Codes	124
Table 47. MRC Fatal Error Codes.....	125
Table 48. POST Progress Codes.....	126
Table 49. POST Error Messages and Handling.....	131
Table 50. POST Error Beep Codes.....	134
Table 51. Integrated BMC Beep Codes	134
Table 52. Volatile and Non-volatile Components.....	135
Table 53. Intel® Server System R1000WF Product Family Feature Set	138
Table 54. Intel® Server System R2000WF Product Family Feature Set	140

1. Introduction

This Technical Product Specification (TPS) provides a high level overview of the features, functions, architecture and support specifications of the Intel® Server Board S2600WF product family.

Note: This document includes several references to Intel websites where additional product information can be downloaded. However, these public Intel sites will not include content for products in development. Content for these products will be available on the public Intel web sites after their public launch.

Note: Some of the documents listed in the following table are classified as “Intel Confidential”. These documents are made available under a Non-Disclosure Agreement (NDA) with Intel and must be ordered through your local Intel representative.

*All highlighted items maybe subject to change.

For more in-depth technical information, the documents in Table 1 should also be referenced.

Table 1. Reference Documents

Document Title	Document Classification
Intel® Servers System BMC Firmware EPS for Intel® Xeon® processor Scalable Family	Intel Confidential
Intel® Server System BIOS EPS for Intel® Xeon® processor Scalable Family	Intel Confidential
“Lewisburg” Platform Controller Hub External Design Specification	Intel Confidential
“Skylake” Server Processor External Design Specification Doc ID: 546831, 546833, 546834, 546832	Intel Confidential
Intel® Ethernet Connection X557-AT2 Product Brief	Public
Add other docs as required	

1.1 Document Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Server Board Overview
- Chapter 3 – Processor Support
- Chapter 4 – PCIe* Support
- Chapter 5 – System Memory
- Chapter 6 – System I/O
- Chapter 7 – On-Board Connector and Header Pin-out Definition
- Chapter 8 – Basic and Advanced Server Management Features
- Chapter 9 – Light-Guided Diagnostics
- Chapter 10 – System Security
- Chapter 11 – Reset and Recovery Jumpers
- Chapter 12 – Platform Management
- Appendix A – Integration and Usage Tips
- Appendix B – POST Code Diagnostic LED Decoder
- Appendix C – POST Code Errors
- Appendix D – Statement of Volatility
- Appendix E – Supported Intel® Server Systems

1.2 Intel Server Board Use Disclaimer

Intel Corporation server boards support add-in peripherals and contain a number of high-density 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 operating environment. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of its 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:

<http://www.intel.com/support>

2. Server Board Family Overview

The Intel® Server Board S2600WF is a monolithic printed circuit board assembly with features that are intended for high density 1U and 2U rack mount servers. This server board is designed to support the Intel® Xeon® processor Scalable family. Previous generation Intel® Xeon® processors are not supported.




Figure 1. Intel® Server Board S2600WF




Figure 2. Intel® Server Board S2600WF with Available On-Board Options

2.1 Server Board Family Feature Set

Table 2 lists the server board product family feature set.

Table 2. Server Board Product Family Feature Set

Intel Server Board Product Code	S2600WFT	S2600WF0
Processor Support	<ul style="list-style-type: none"> Two LGA3647-0 (Socket P) processor sockets Support for one or two Intel® Xeon® processor Scalable family (Platinum, Gold, Silver, and Bronze) <ul style="list-style-type: none"> Previous generation Intel® Xeon® processors are not supported Maximum supported Thermal Design Power (TDP) of up to 205W (Board Only) <p>Note: Intel Server Systems based on this server board family may support a lower maximum Thermal Design Power (TDP). See appropriate Intel System TPS for max supported TDP</p>	
Memory	<ul style="list-style-type: none"> 24 Total DIMM slots (12 DIMMs per processor) <ul style="list-style-type: none"> 6 Memory Channels per processor / 2 DIMMs per Channel Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) Memory Capacity <ul style="list-style-type: none"> Up to 1.5TB for Gold and Platinum CPUs; Up to 768GB for Silver and Bronze CPUs Memory data transfer rates: <ul style="list-style-type: none"> Up to 2666 MT/s @ 1 DPC and 2 DPC (Processor SKU Dependent) DPC = DIMMs Per Channel DDR4 standard voltage of 1.2V 	
Intel® Chipset	Intel® C624 Chipset	Intel® C624 Chipset
Intel® Quick Assist Technology (QAT)	No	No
Intel® Omni-Path Fabric Support	Yes	Yes
On-board LAN	Dual Port RJ45 10GbE	No
OCP Module Support iPC = Intel Product Code	<ul style="list-style-type: none"> iPC 557T2OCPG1P5 – Dual Port 10Gb RJ45 iPC 527DA2OCPG1P5 – Dual Port SFP+ 	<ul style="list-style-type: none"> iPC I357T4OCPG1P5 – Quad Port 1Gb RJ45 iPC X527DA4OCPG1P5 – Quad Port SFP+ iPC X557T2OCPG1P5 – Dual Port 10Gb RJ45 iPC X527DA2OCPG1P5 – Dual Port SFP+
Intel® Integrated SAS Module Support	Yes	Yes
Onboard PCIe* NVMe Support	4 – PCIe OCuLink Connectors Intel® VMD Support Intel® RSTe VROC Support – Acc. Option	4 – PCIe OCuLink Connectors Intel® VMD Support Intel® RSTe VROC Support – Acc. Option
Onboard SATA Support	<ul style="list-style-type: none"> 12 x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) <ul style="list-style-type: none"> Two single port 7-pin SATA connectors Two M.2 connectors – SATA / PCIe* Two 4-port mini-SAS HD (SFF-8643) connectors Embedded SATA Software RAID <ul style="list-style-type: none"> Intel® Rapid Storage RAID Technology (RSTe) 5.0 Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 with optional RAID 5 key support <ul style="list-style-type: none"> NOTE: ESRT2 is only supported on S2600WFT and S2600WF0 boards 	
Riser Card Support	<p>Concurrent support for up to three riser cards</p> <ul style="list-style-type: none"> Riser #1 – PCIe* 3.0 x24 (CPU1 x16, CPU2 x8) – 2 and 3 slot riser card options available Riser #2 – PCIe* 3.0 x24 (CPU2 x24) – 2 and 3 slot riser card options available Riser #3 (2U systems only) – PCIe* 3.0 (CPU 2 x12) – 2 slot riser card available 	

Intel Server Board Product Code	S2600WFT	S2600WF0
Video	<ul style="list-style-type: none"> • Integrated 2D Video Controller • 16MB of DDR4 Video Memory • One DB-15 External Connector • One 14-Pin Internal connector for optional Front Panel Video support 	
USB Support	<ul style="list-style-type: none"> • Three external USB 3.0 ports • One internal Type-A USB 2.0 port • One internal 20-pin connector for optional 2x USB 3.0 port Front Panel support • One Internal 10-pin connector for optional 2x USB 2.0 port Front Panel support 	
Serial Port Support	<ul style="list-style-type: none"> • One external RJ-45 Serial-A port connector • One internal DH-10 Serial-B port header for optional front or rear serial port support 	
Server Management	<ul style="list-style-type: none"> • Integrated Baseboard Management Controller, IPMI 2.0 compliant • Support for Intel® Server Management Software • On-board dedicated RJ45 management port • Support for Advanced Server Management features via an Intel® Remote Management Module 4 Lite Accessory Option (iPC – AXXRMM4LITE2) 	
Security	<ul style="list-style-type: none"> • Intel® Trusted Platform Module 2.0 (Rest of World) – iPC- AXXTPMENC8 (Accessory Option) • Intel® Trusted Platform Module 2.0 (China Version) – iPC- AXXTPMCHNE8 (Accessory Option) 	
System Fan Support	<ul style="list-style-type: none"> • Six System fans supported in two different connector formats hot swap (2U) and cabled (1U) <ul style="list-style-type: none"> ○ Six 10-pin managed system fan headers (Sys_Fan 1-6) – Used for 1U system configuration ○ Six 6-pin hot swap capable managed system fan connectors (Sys_Fan 1-6) – Used for 2U system Configuration 	

2.2 Server Board Component / Feature Identification

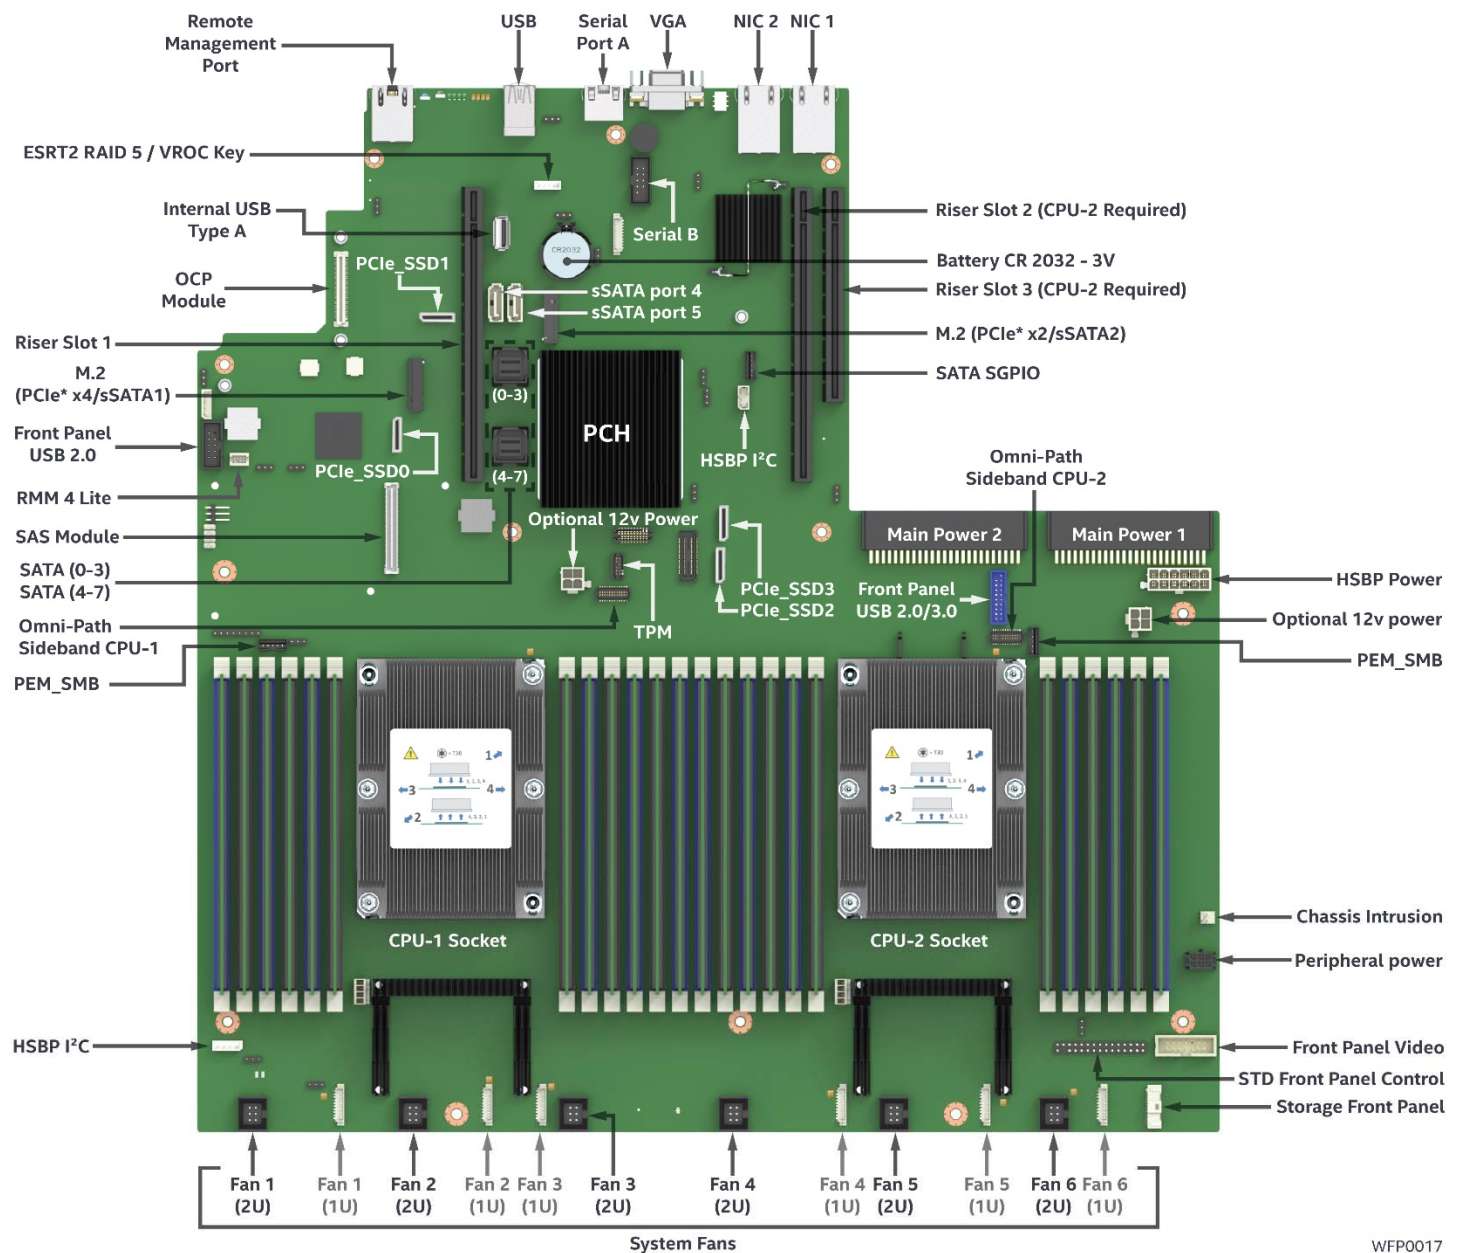



Figure 3. Server Board Component / Feature Identification

Note: Intel® Server Board S2600WFT shown. Some features may not be present on Intel® Server Boards S2600WF0 and/or S2600WFQ.



- A – RJ45 Networking Port – NIC #1
- B – RJ45 Networking Port – NIC #2
- C – Video
- D – RJ45 Serial 'A' Port
- E – Stacked 3-port USB 3.0
- F – RJ45 Dedicated Management Port

Figure 4. Intel® Server Board S2600WF External I/O Connector Layout

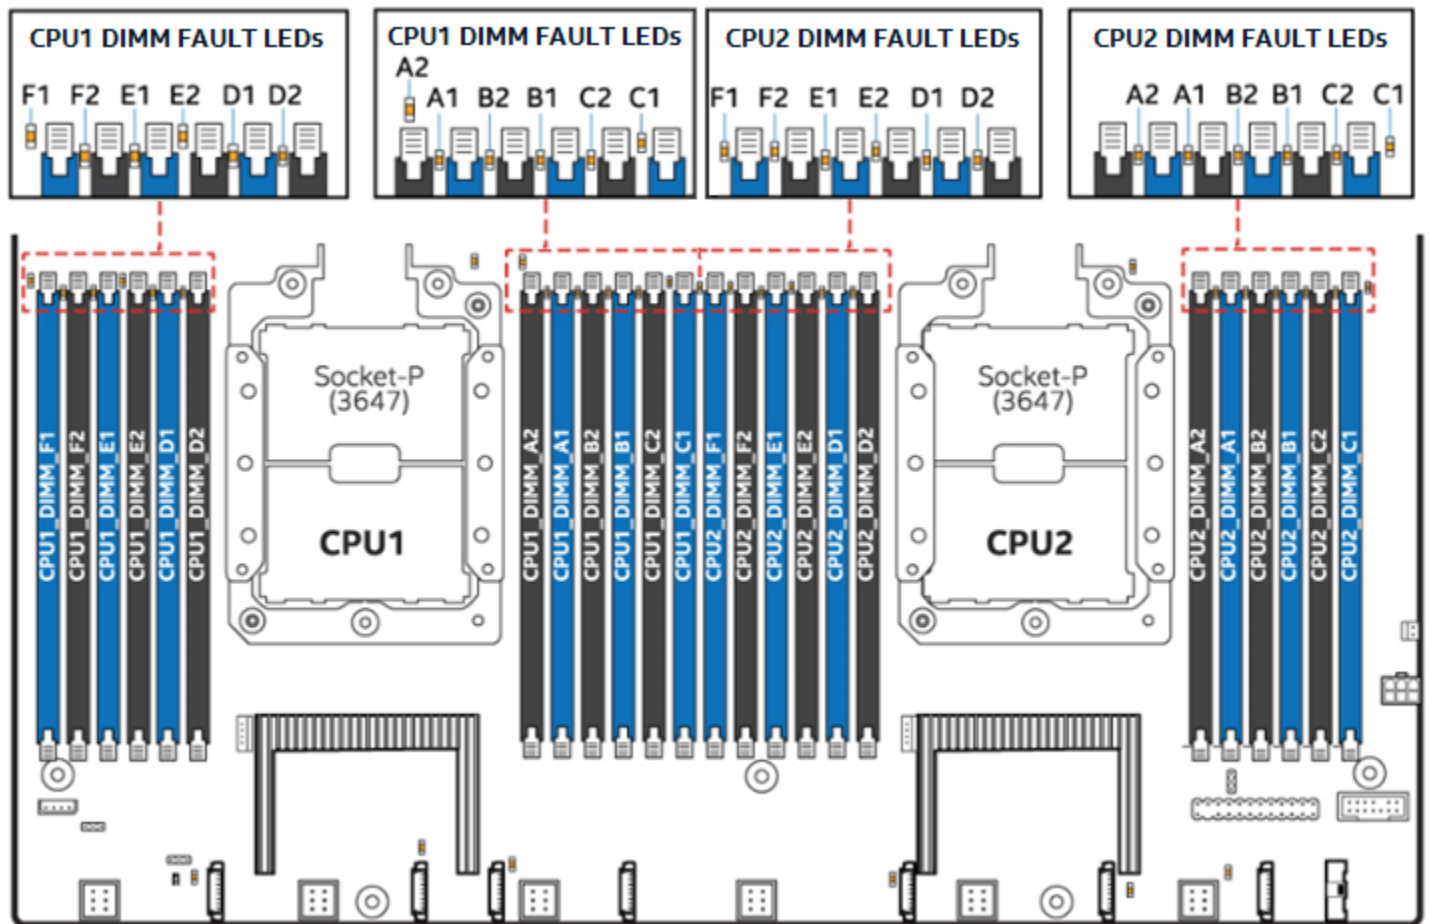


Figure 5. Intel® Light Guided Diagnostics - DIMM Fault LEDs

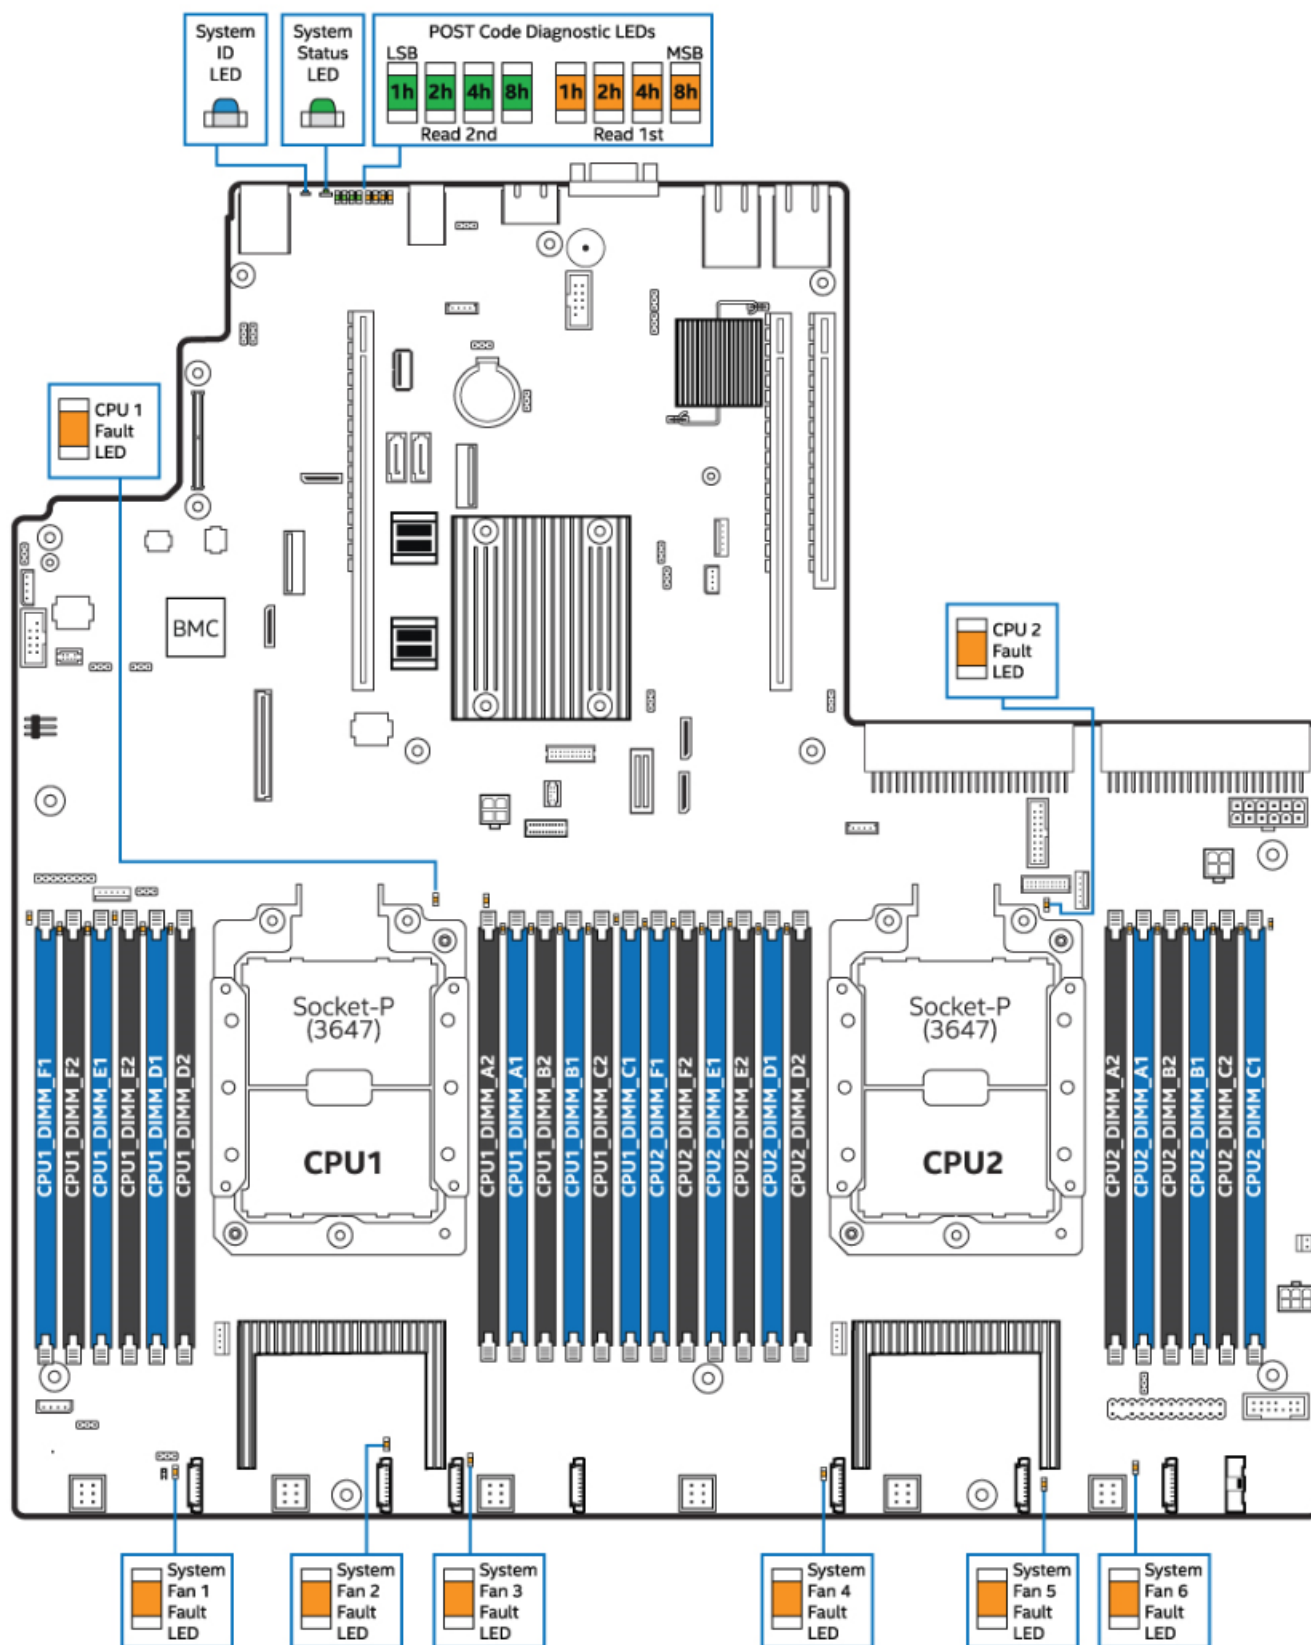


Figure 6. Intel® Light Guided Diagnostic LED Identification

Note: See Appendix D for POST Code Diagnostic LED decoder information




Figure 7. Board Configuration and Recovery Jumpers

See Chapter 11 - Reset and Recovery Jumpers for additional details.

2.3 Server Board Mechanical Drawings

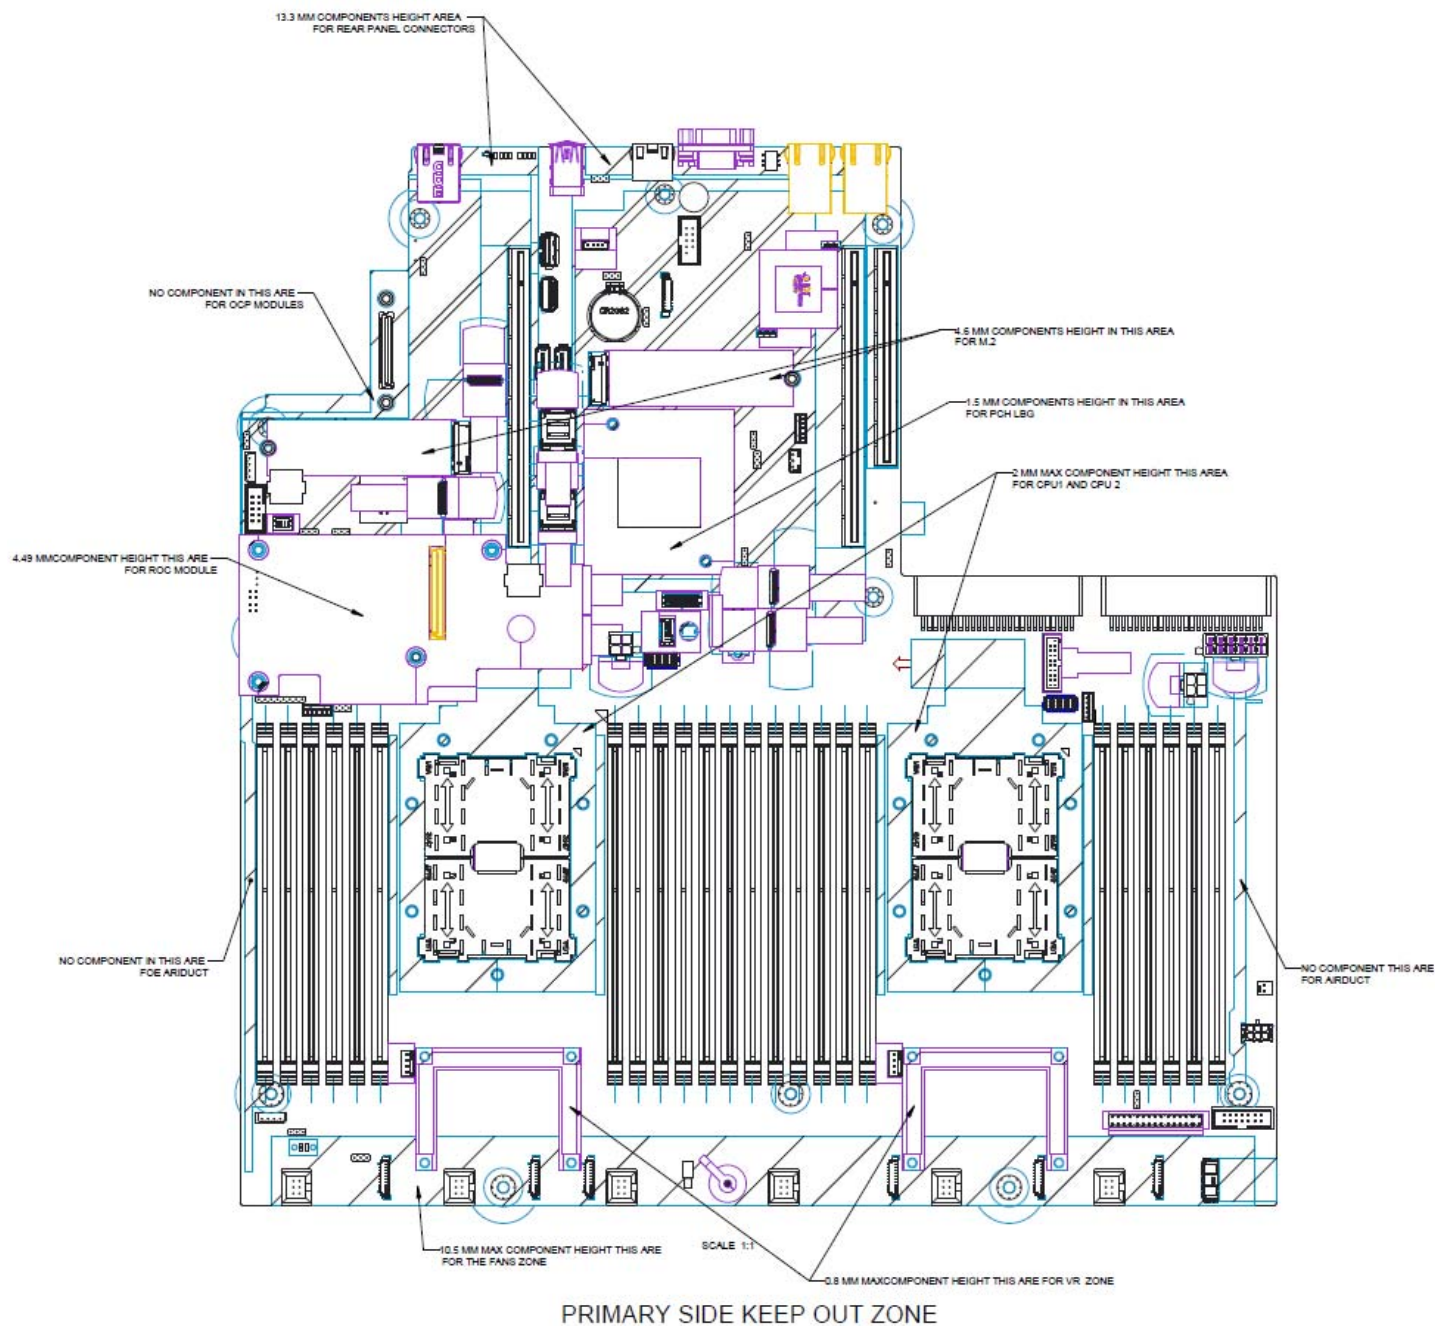


Figure 8. Intel® Server Board S2600WF – Primary Side Keepout Zone

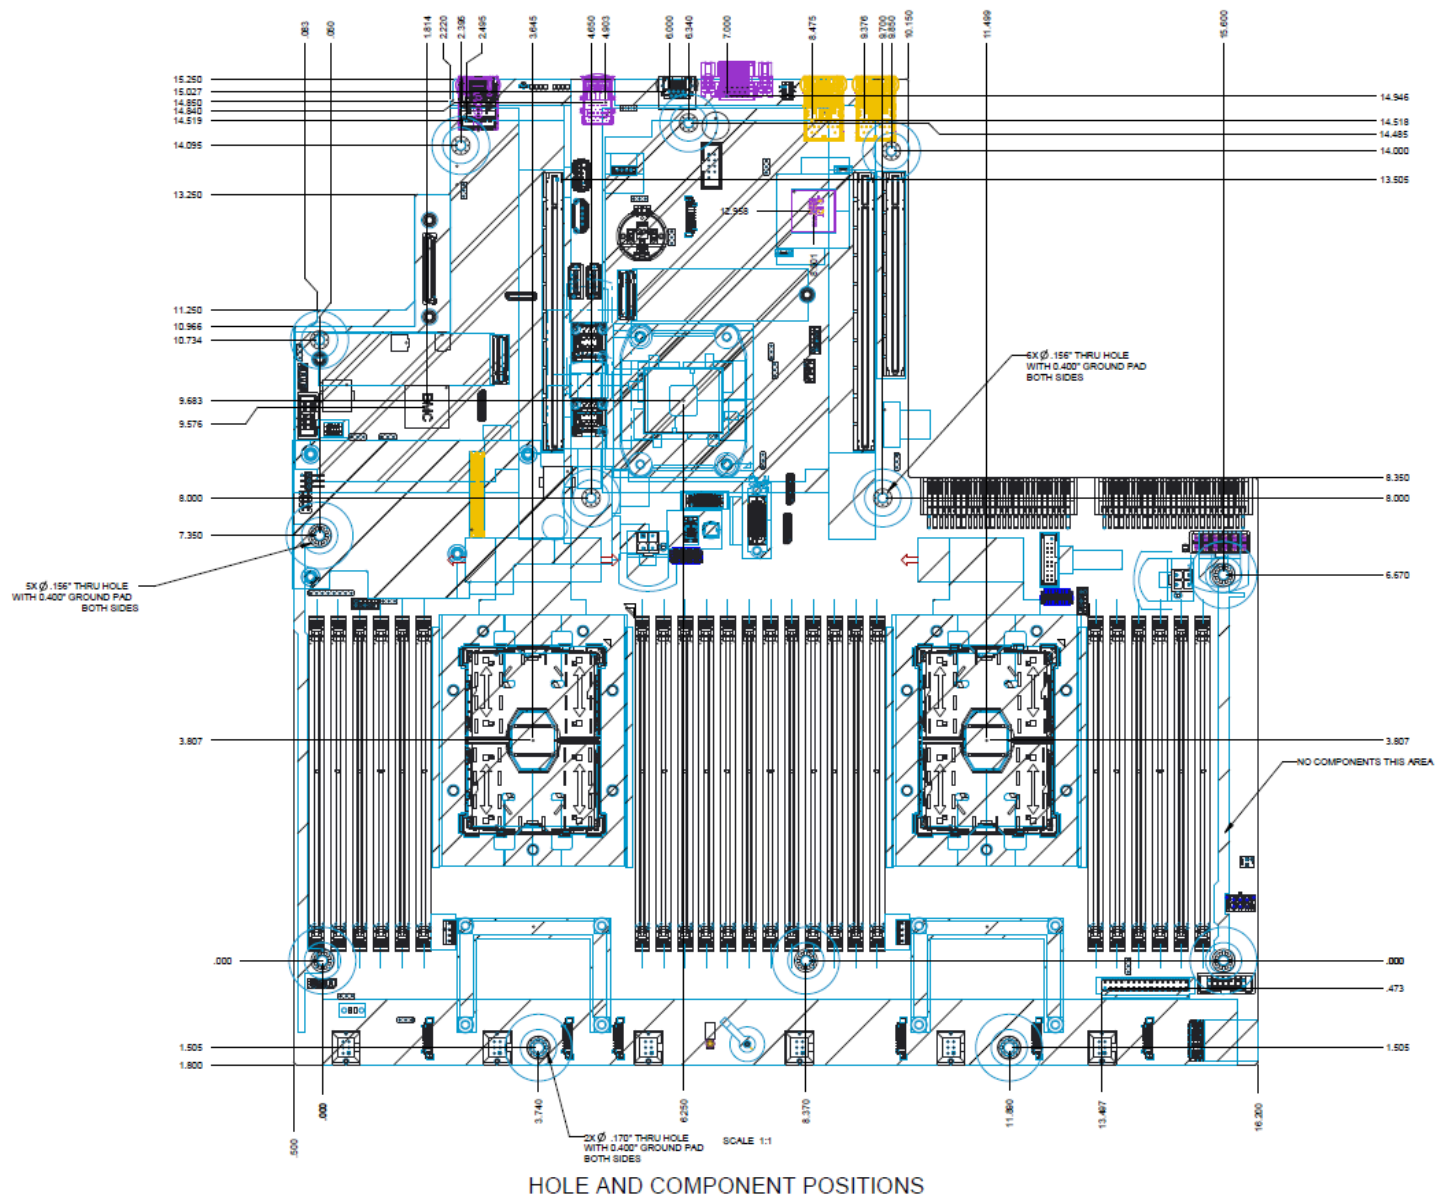


Figure 9. Intel® Server Board S2600WF – Hole and Component Positions

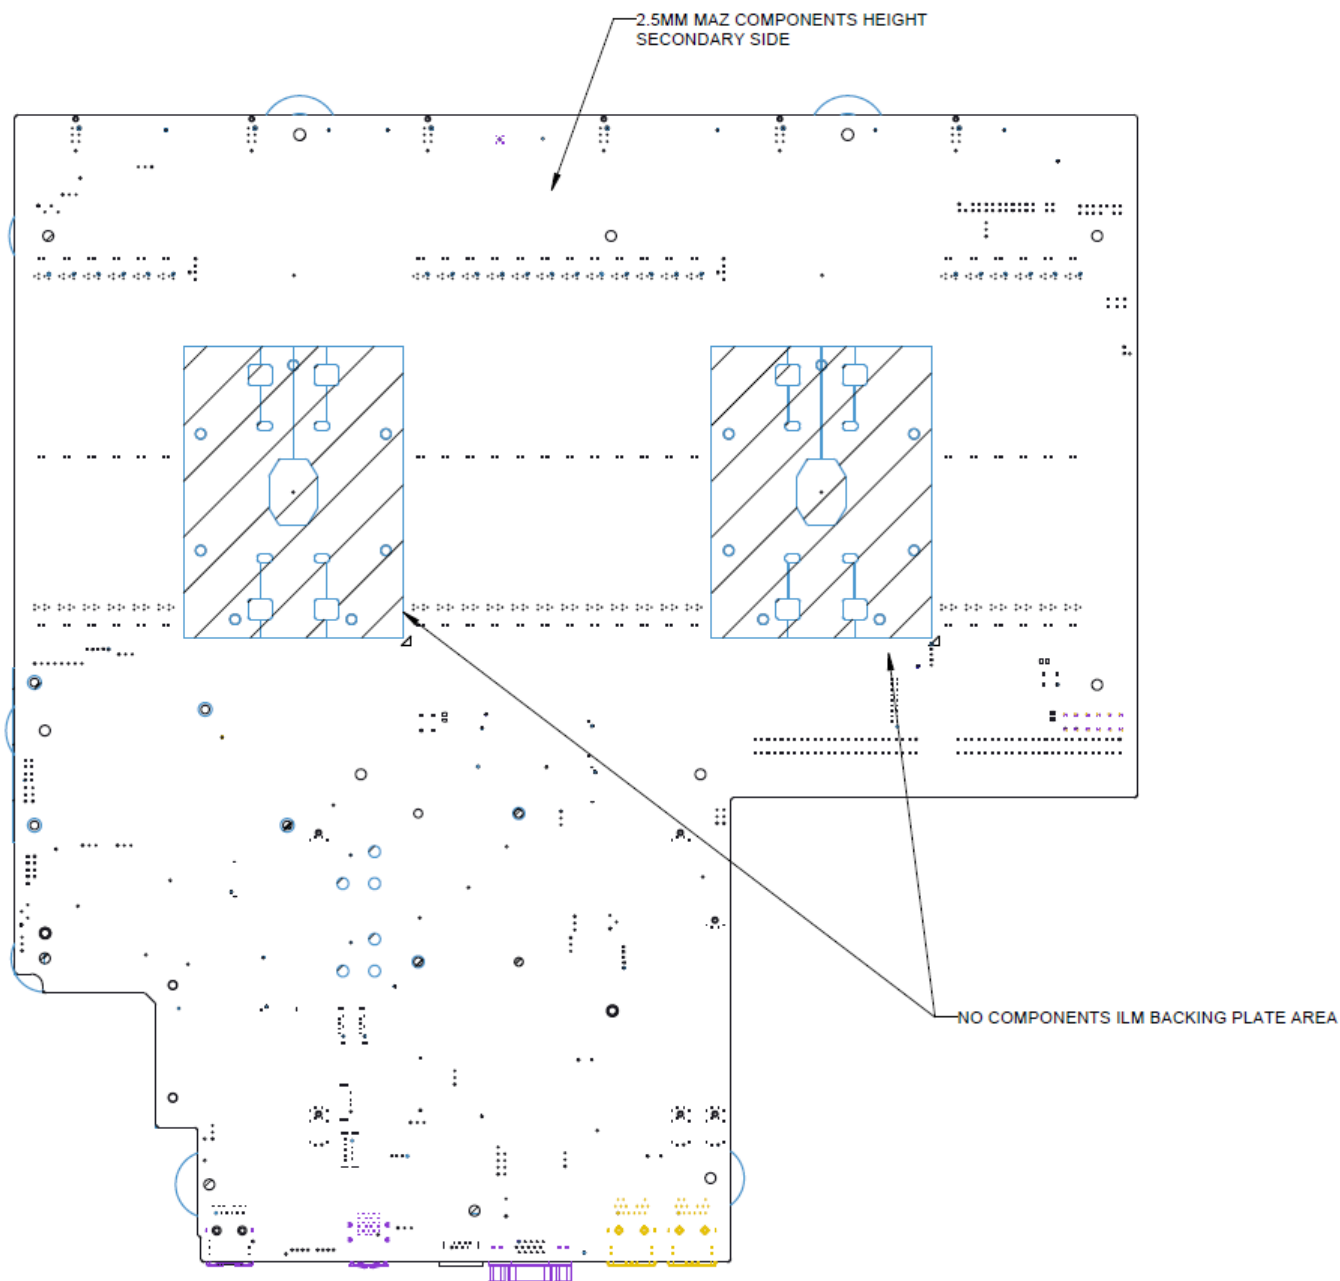


Figure 10. Intel® Server Board S2600WF – Secondary Side Keepout Zone

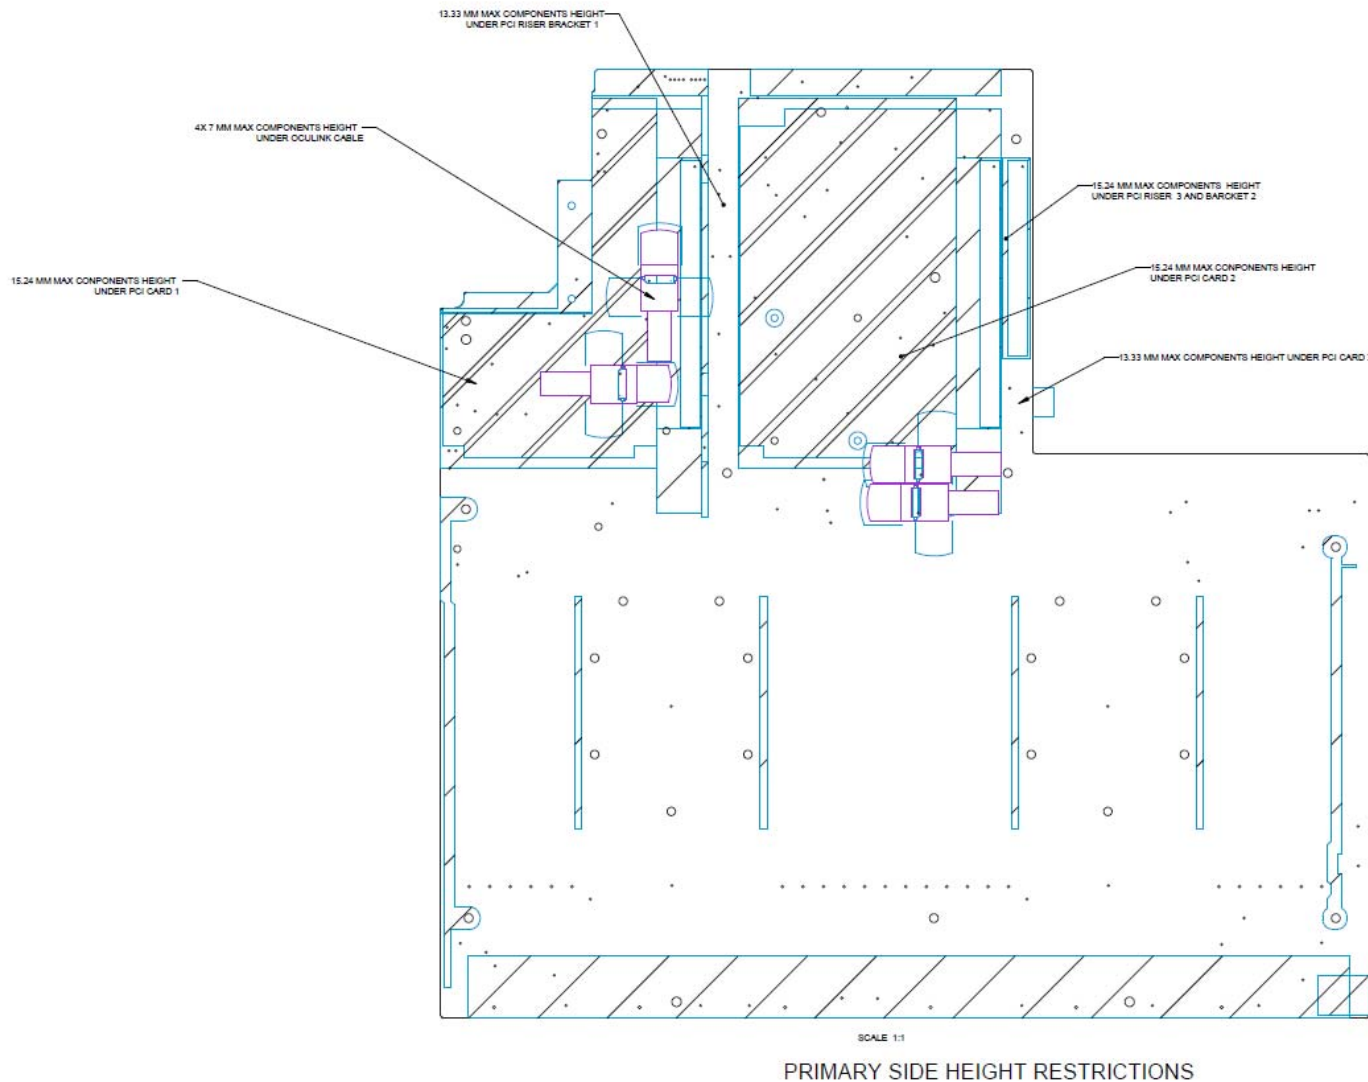


Figure 11. Intel® Server Board S2600WF – Primary Side Height Restrictions

2.4 Product Architecture Overview

The architecture of Intel® Server Board S2600WF product family is developed around the integrated features and functions of the Intel® Xeon® processor Scalable family, the Intel® C620 series chipset (PCH), Intel® Ethernet Controller X557-AT2 (S2600WFT only), and the ASPEED® AST2500 Baseboard Management Controller.

Figure 12 provides an overview of the server board architecture, showing the features and interconnects of each of the major sub-system components.

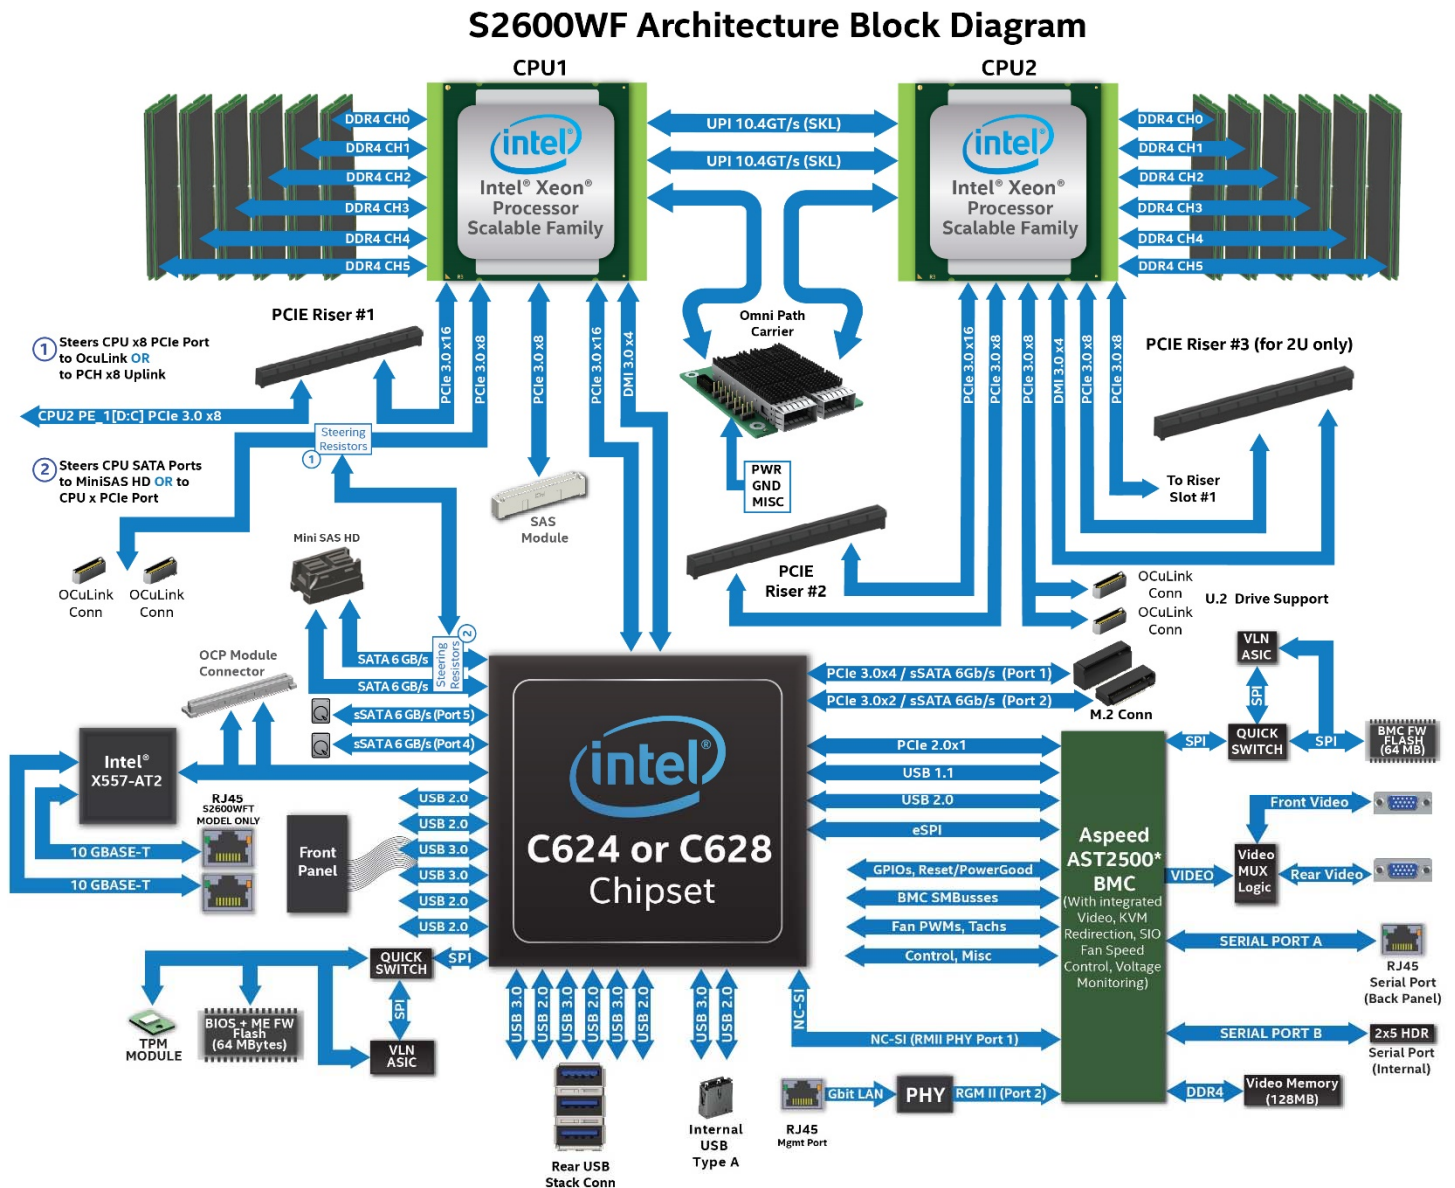


Figure 12. Intel® Server Board S2600WF Product Family Architectural Block Diagram

2.5 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: <http://downloadcenter.intel.com>

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

2.5.1 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 3. 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

2.5.1.1 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.

2.5.1.2 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.

2.5.1.3 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*

2.5.1.4 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.

2.5.1.5 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.

2.5.2 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
- OCP Module
- Integrated SAS Raid module
- Power supply
- Fan
- Intel® Xeon Phi™ co-processor PCIe card
- 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.

2.5.2.1 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 (Standard and Fabric options)
 - Supported processor Thermal Design Power (TDP) of up to 205W

Note: Previous-generation Intel® Xeon® processors and their supported CPU heat sinks are not compatible on server boards described in this document.

Note: The server board is capable of supporting processors with a maximum 205W TDP (Thermal Design Power). However, TDP support may vary depending on the cooling capabilities of the chosen server chassis. Check the server chassis or server system product specifications to determine maximum supported processor TDP.

Visit <http://www.intel.com/support> for a complete list of supported processors.

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.

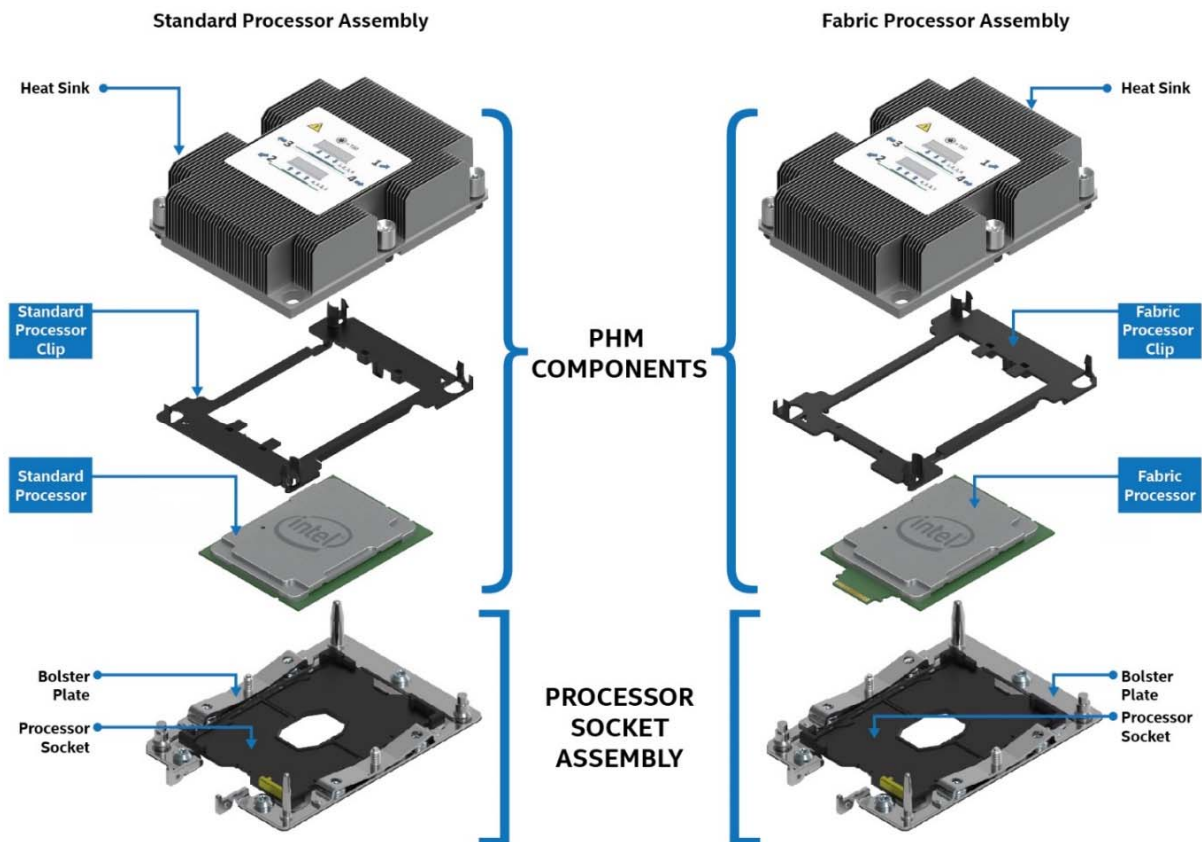



Figure 13. PHM Components and Processor Socket Reference Diagram

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 14).





Figure 14. 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 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 15.




Figure 15. Processor Socket Assembly and Protective Cover

Processor socket covers must be removed before processor installation.

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 board described in this document is designed to support the Intel® Xeon® processor Scalable family TDP guidelines up to and including 205 Watts.

Disclaimer Note: Intel Corporation server boards contain a number of high-density 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 its published operating or non-operating limits.

3.3 Intel® Xeon® Processor Scalable Family Overview

The Intel® Server Board S2600WF 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 4. 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

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 – See Chapter **Error! Reference source not found.**

3.3.1 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/>

3.3.2 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.

3.3.3 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).

3.3.4 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.

3.3.5 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.

3.3.6 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.

3.3.7 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.

3.3.8 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.

3.3.9 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.

3.3.10 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.

3.3.11 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 the 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.

For additional information, visit the following Intel website:

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

3.4 Processor Population Rules

Note: 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.

When using a single processor configuration, the processor must be installed into the processor socket labeled "CPU_1".

Note: Some server board features may not be functional unless a second processor is installed. See Figure 12.

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 memory 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.

3.5 Processor Initialization Error Summary

Table 5 describes mixed processor conditions and recommended actions for all Intel® server boards and Intel server systems designed around the Intel® Xeon® processor E5-2600 v5 product family and Intel® C620 series chipset architecture. The errors fall into one of the following categories:

Fatal: If the system cannot 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 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 or 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 5. Mixed Processor Configurations Error Summary

Error	Severity	System Action
Processor family not Identical	Fatal	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> • Halts at POST Code 0xE6 • Halts with 3 long beeps and 1 short beep • Takes Fatal Error action (see above) and will not boot until the fault condition is remedied
Processor model not Identical	Fatal	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> • Logs the POST Error Code into the System Event Log (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 will not boot until the fault condition is remedied
Processor cores/threads not identical	Fatal	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> • Halts at POST Code 0xE5 • Halts with 3 long beeps and 1 short beep • Takes Fatal Error action (see above) and will not boot until the fault condition is remedied
Processor cache or home agent not identical	Fatal	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> • Halts at POST Code 0xE5 • Halts with 3 long beeps and 1 short beep • Takes Fatal Error action (see above) and will not boot until the fault condition is remedied
Processor frequency (speed) not identical	Fatal	<p>The BIOS detects the processor frequency difference, and responds as follows:</p> <ul style="list-style-type: none"> • Adjusts all processor frequencies to the highest common frequency • No error is generated – this is not an error condition • Continues to boot the system successfully <p>If the frequencies for all processors cannot be adjusted to be the same, then this is an error, and the BIOS responds as follows:</p> <ul style="list-style-type: none"> • 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 will not boot until the fault condition is remedied

Error	Severity	System Action
Processor Intel® UltraPath Interconnect link frequencies not identical	Fatal	<p>The BIOS detects the UPI link frequencies and responds as follows:</p> <ul style="list-style-type: none"> Adjusts all UPI interconnect link frequencies to highest common frequency No error is generated – this is not an error condition Continues to boot the system successfully <p>If the link frequencies for all UPI links cannot be adjusted to be the same, then this is an error, and the BIOS responds as follows:</p> <ul style="list-style-type: none"> 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 will not boot until the fault condition is remedied
Processor microcode update failed	Major	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> 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 setting of POST Error Pause in Setup, or may halt with the POST Error Code in the Error Manager waiting for operator intervention.
Processor microcode update missing	Minor	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> 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 setting of POST Error Pause in the Setup.

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

The Intel® Xeon® processor Scalable family includes SKUs which include an integrated Intel® Omni-Path Host Fabric Interface (HFI) connector.


Table 6. Intel® Xeon® Processor Scalable Family w/ Integrated Intel® Omni-Path Fabric

Host Fabric Interface – Features Table

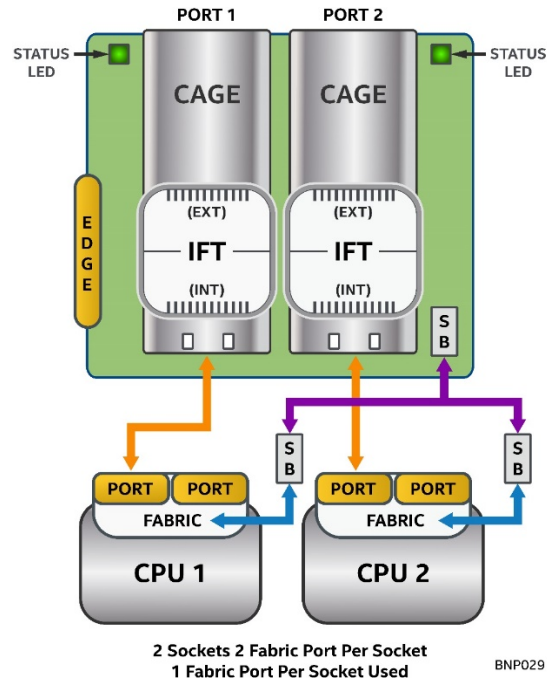
Feature	81xxF Platinum	61xxF Gold
# of Cores	≥ 24	< 24
# of Omni-Path Fabric Ports	1	1
# of Intel® UPI Links	2	2
UPI Speed	10.4 GT/s	10.4 GT/s
Supported Topologies	2S-2UPI	2S-2UPI
Node Controller Support	No	No
# of Memory Channels	6	6
Max DDR4 Speed	2666	2666
Memory Capacity	768GB 1.5TB (Select SKUs)	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.

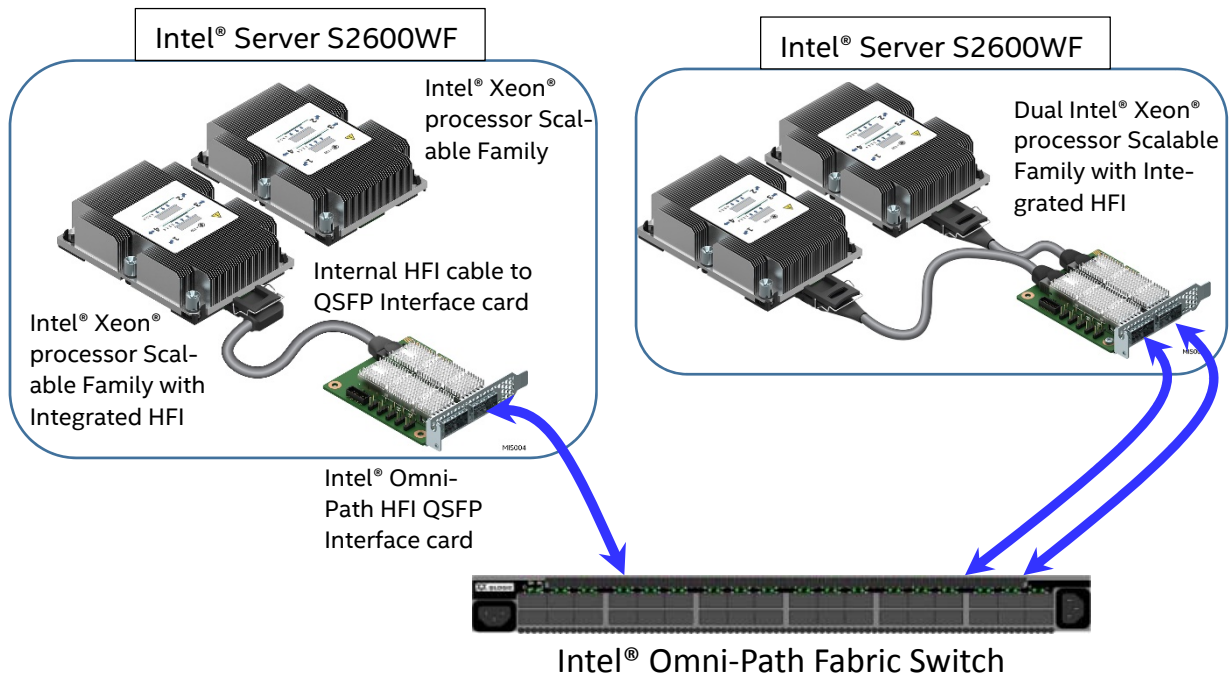
1 port x 100 Gbps Host Fabric Interface (HFI) Connector



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 in any available PCIe add-in card slot or within the OCP module bay. 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.



The following figure illustrates two supported dual processor configurations with one or two Fabric processors. In the diagram, each processor HFI connector is cabled to a QSFP28 interface card (See section 3.6.1).



3.6.1 Intel®Omni-Path IFT Carrier Accessory Kits

All necessary components to support up to two fabric processors are included in orderable Fabric Accessory Kits (**AWF1PFABKITM** and **AWF1PFABKITP**).

Intel Product Code (iPC)	Description	Accessory Kit Contents
AWF1PFABKITM	Intel IFT Carrier Kit – Mezzanine	1 – Dual port IFT Carrier Mezzanine Card 1 – Internal Omni-Path Cable (CPU1) 1 – Internal Omni-Path Cable (CPU2) 1 – Internal Omni-Path Sideband Cable 2 – Fabric Processor Carriers
AWF1PFABKITP	Intel IFT Carrier Kit – PCIe	1 – Dual Port IFT Carrier PCIe Add-in Card 1 – Internal Omni-Path Cable (CPU1) 1 – Internal Omni-Path Cable (CPU2) 1 – Internal Omni-Path Sideband Cable 2 – Fabric Processor Carriers

** Table Contents Subject to change

Two options for the Intel Fabric Through (IFT) carrier card are offered:

- Mezzanine – mounted directly to the server board in the designated OCP Module mounting location
- PCIe Add-in card – installed to any available Riser Slot 2 PCIe add-in slot




Figure 16. Intel®Omni-Path IFT Carrier Accessory Kit Components

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




Figure 17. Server Board Sideband Connectors

Each IFT carrier port has one green Status LED.




Figure 18. IFT Carrier Board – Rear View

Table 7. 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 8.

Table 8. Power Level Classification for QSFP+ Modules

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

The server board has support for processor configurations where one or two installed processors may have an Intel® Omni-Path Host Fabric Interface. 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.

Table 9. 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. System Memory

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

4.1 Memory Sub-system Architecture



Figure 19. Memory Sub-system Architecture

Note: This generation server board only has support for DDR4 memory.

The Intel® Server Board S2600WF has support for up to 24 DDR4 DIMMs, 12 per processor. Each installed processor supports 6 memory channels via two Integrated Memory Controllers (IMC). On the server board memory channels are assigned an identifier letter A thru F, with each memory channel supporting two DIMM slots.

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

Table 10.DDR4 RDIMM and LRDIMM Support

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

*Subject To change

4.3 Memory Slot Identification and Population Rules

Note: Although mixed DIMM configurations may be functional, Intel only supports and performs platform validation on systems that are configured with identical DIMMs installed.

On the Intel® Server Board S2600WF, a total of 24 DIMM slots are provided – 2 CPUs, 6 Memory Channels/CPU, 2 DIMMs/Channel. Figure 20. Identifies all DIMM slots on the server board.





Figure 20. Intel® Server Board S2600WF Memory Slot Layout

- Each installed processor provides six channels of memory. Memory channels from each processor are identified as Channels A – F.
- On the Intel® Server Board S2600WF, each memory channel supports two DIMM slots, identified as slots 1 and 2.
 - 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**
- DIMM population rules require that DIMMs within a channel be populated starting with the BLUE DIMM slot or DIMM farthest from the processor in a “fill-farthest” approach.
- When only one DIMM is used for a given memory channel, it must be populated in the BLUE DIMM slot (furthest from the CPU).
- Mixing of DDR4 DIMM Types (RDIMM, LRDIMM, 3DS RDIMM, 3DS LRDIMM, NVDIMM) within a channel socket or across sockets produces a Fatal Error Halt during 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.
- When populating a Quad-rank DIMM with a Single- or Dual-rank DIMM in the same channel, the Quad-rank DIMM must be populated farthest from the processor. Intel MRC will check for correct DIMM placement. 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.
- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- A processor may be installed without populating the associated memory slots, provided a second processor is installed with 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.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS, Error Management,) in the BIOS setup are applied commonly across processor sockets.
- For multiple DIMMs per channel:
 - For RDIMM, LRDIMM, 3DS RDIMM, 3DS LRDIMM; Always populate DIMMs with higher electrical loading in slot1, followed by slot 2.

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 Slot1 (**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, D, and E
 - **9 DIMM, 10, DIMM, and 11 DIMM Configurations - NOT Recommended.** These are an unbalanced configurations which will yield less than optimal performance
 - **12 DIMM Configurations** – DIMMs are populated to ALL DIMM Slots

4.4 Memory RAS Features

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

Table 11. Memory RASM Features

RASM Feature	Description	Standard	Advanced
Device Data Correction	x8 Single Device Data Correction (SDDC) via static virtual lockstep (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.	✓	✓

RASM Feature	Description	Standard	Advanced
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.	✓	✓
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
- Memory Sparing and Memory Mirroring options are mutually exclusive. Only one operating mode may be selected in BIOS Setup
- 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.

5. PCIe* Support

The PCI Express interface of the Intel® Server Board S2600WF product family is fully compliant with the PCI Express Base Specification, Revision 3.0 supporting the following PCIe bit rates: Gen 3.0 (8.0 GT/s), Gen 2.0 (5.0 GT/s), and Gen 1.0 (2.5 GT/s).

For specific board features and functions supported by the PCIe sub-system, see Chapter 6. The following table provide the PCIe* port routing information from each processor:

Table 12. CPU - PCIe* Port Routing

CPU 1		CPU 2	
PCI Ports	On-board Device	PCI Ports	On-board Device
Port DMI 3 - x4	Chipset	Port DMI 3 - x4	Riser Slot #3
Port 1A - x4	Riser Slot #1	Port 1A - x4	Riser Slot #2
Port 1B - x4	Riser Slot #1	Port 1B - x4	Riser Slot #2
Port 1C - x4	Riser Slot #1	Port 1C - x4	Riser Slot #1
Port 1D - x4	Riser Slot #1	Port 1D - x4	Riser Slot #1
Port 2A - x4	Chipset (PCH) - uplink	Port 2A - x4	Riser Slot #2
Port 2B - x4	Chipset (PCH) - uplink	Port 2B - x4	Riser Slot #2
Port 2C - x4	Chipset (PCH) - uplink	Port 2C - x4	Riser Slot #2
Port 2D - x4	Chipset (PCH) - uplink	Port 2D - x4	Riser Slot #2
Port 3A - x4	SAS Module	Port 3A - x4	OCuLink PCIe_SSD2
Port 3B - x4	SAS Module	Port 3B - x4	OCuLink PCIe_SSD3
Port 3C - x4	OCuLink PCIe_SSD0	Port 3C - x4	Riser Slot #3
Port 3D - x4	OCuLink PCIe_SSD1	Port 3D - x4	Riser Slot #3

5.1.1 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.

5.1.2 Non-Transparent Bridge

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. System I/O

The server board Input/Output features are provided via the embedded features and functions of several onboard components including: the Integrated I/O Module (IIO) of the Intel® Xeon processor, the Intel® C620 series chipset (PCH), and the I/O controllers embedded within the Aspeed® AST2500 Management Controller. See Figure 12. Intel® Server Board S2600WF Product Family Architectural Block Diagram for an overview of the features and interconnects of each of the major sub-system components. Server Board I/O features include:

- QAT support (S2600WFQ only)
- PCIe® Riser Card and Add-in Card Support
- Intel® OCP Module Support
- Intel® Integrated RAID Module Support
- On-board Storage Sub-system
- External I/O Port Support

6.1 PCIe® Add-in Card Support

The server board provides three riser card slots identified as: Riser Slot #1, Riser Slot #2, and Riser Slot #3. Per the PCIe specification, each riser card slot can support a maximum 75W of power. The PCIe® bus lanes for each riser card slot is supported by each of the two installed processors.

Note: 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.

Note: A dual processor configuration is required when using Riser Slot #2 and Riser Slot #3, as well as the bottom add-in card slot for 2U riser cards installed in Riser Slot #1.

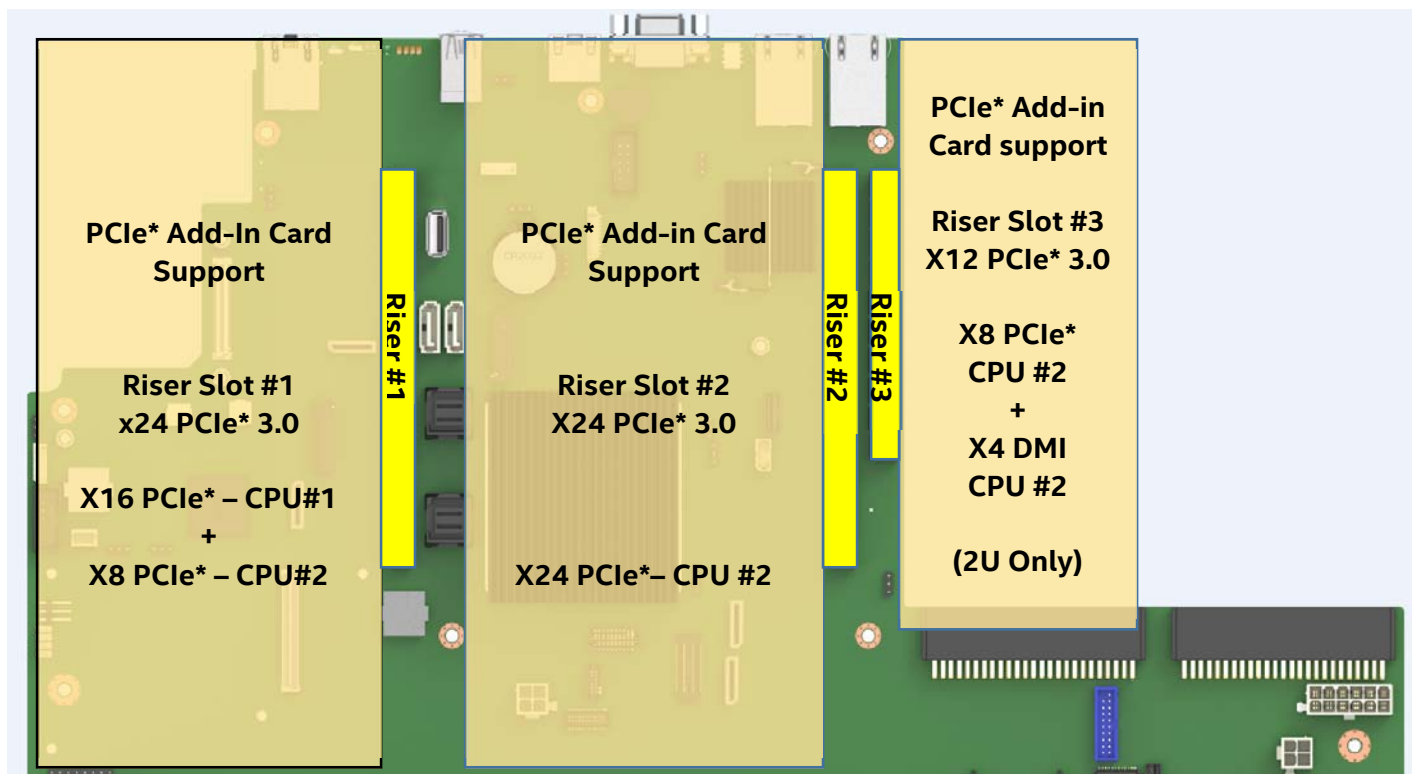


Figure 21. PCIe® Add-in Card Support

The following tables provide the PCIe* bus routing for all supported risers cards.

Table 13. Riser Slot #1 – PCIe* Root Port Mapping

Riser Slot #1 – Riser Card Options		
2U - 3-Slot Riser Card iPC – A2UL8RISER2	2U - 2-Slot Riser Card iPC – A2UL16RISER2	1U - 1-Slot Riser Card iPC – F1UL16RISER3APP
Top PCIe* Slot CPU #1 – Ports 1A and 1B (x8 elec, x16 mech)	Top PCIe* Slot CPU #1 – Ports 1A thru 1D (x16 elec, x16 mech)	PCIe* Slot CPU #1 – Ports 1A thru 1D (x16 elec, x16 mech)
Middle PCIe* Slot CPU #1 – Ports 1C and 1D (x8 elec, x16 mech)		
Bottom PCIe* Slot CPU #2 – Ports 1C and 1D (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Ports 1C and 1D (x8 elec, x8 mech)	

Table 14. Riser Slot #2 – PCIe* Root Port Mapping

Riser Slot #2 – Riser Card Options		
2U - 3-Slot Riser Card iPC – A2UL8RISER2	2U - 2-Slot Riser Card iPC – A2UL16RISER2	1U - 1-Slot Riser Card iPC – F1UL16RISER3APP
Top PCIe* Slot CPU #2 – Ports 2A and 2B (x8 elec, x16 mech)	Top PCIe* Slot CPU #2 – Ports 2A thru 2D (x16 elec, x16 mech)	Top PCIe* Slot CPU #2 – Ports 2A thru 2D (x16 elec, x16 mech)
Middle PCIe* Slot CPU #2 – Ports 2C and 2D (x8 elec, x16 mech)		
Bottom PCIe* Slot CPU #2 – Ports 1A and 1B (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Ports 1A and 1B (x8 elec, x8 mech)	

Table 15. Riser Slot #3 – PCIe* Root Port Mapping

Riser Slot #3 – Riser Card Options	
2U - Low Profile Riser Card iPC – A2UX8X4RISER	Notes
Top PCIe* Slot CPU #2 – DMI x4 (x4 elec, x8 mech)	Low profile cards only
Bottom PCIe* Slot CPU #2 – Ports 3C and 3D (x8 elec, x8 mech)	Low profile cards only

6.1.1 Riser Card Support

Available riser cards for Riser Slots #1 and #2 are common between the two slots.

- **1U** – One PCIe* add-in card slot – PCIe* x16, x16 mechanical





Figure 22. 1U One Slot PCIe* Riser Card (iPC – F1UL16RISER3APP)

Each riser card assembly has support for a single full height, ½ length PCIe* add-in card. However, riser card #2 may be limited to ½ length, ½ height add-in cards if either of the two mini-SAS HD connectors on the server board are used.

Note: Add-in cards that exceed the PCI specification for ½ length PCI add-in cards (167.65mm or 6.6in) may interfere with other installed devices on the server board.

- **2U** – Three PCIe* add-in card slots




Slot #	Description
Slot-1 (Top)	PCIe* x8 elec, x16 mechanical
Slot-2 (Middle)	PCIe* x8 elec, x16 mechanical
Slot-3 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 23. 2U Three Slot PCIe* Riser Card (iPC – A2UL8RISER2)

Each riser card assembly has support for up to two full height full length add-in cards (top and middle slots) and one full height ½ length add-in card (bottom slot).

- **2U** – Two PCIe* add-in card slots




Slot #	Description
Slot-1 (Top)	PCIe* x16 elec, x16 mechanical
Slot-2 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 24. 2U Two PCIe* Slot Riser Card (iPC – A2UL16RISER2)

Each riser card assembly has support for one full height full length add-in card (top slot) and one full height ½ length add-in card (bottom slot).

Riser Slot #3 is provided to support up to two additional PCIe* add-in card slots for 2U server configurations. The available riser card option is designed to support low profile add-in cards only.



Slot #	Description
Slot-1 (Top)	PCIe* x4 elec, x8 mechanical
Slot-2 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 25. 2U Two PCIe* Slot (Low Profile) PCIe* Riser Card (iPC – A2UX8X4RISER)

6.1.2 Intel® OCP Module Support

The Intel® Server Board S2600WF Product Family offers a line of LAN KR OCP mezzanine modules that follows the OCP 2.0 form factor.

The optional OCP mezzanine module can be installed onto the connector (labeled “OCP_IO_Module”) on the server board, as shown in the following illustration.




Figure 26. Intel® OCP Module Connector

The following table lists the supported Intel® OCP Modules.

Table 16. Supported Intel® OCP Modules

Description	Intel Product Code
Quad Port, 1GB, RJ45	I357T4OCPG1P5
Quad Port, SFP+	X527DA4OCPG1P5
Dual Port, SFP+	X527DA2OCPG1P5
Dual Port, 10Gb	X557T2OCPG1P5

Note: Only dual-port SFP+ and dual-port 10GB RJ45 OCP modules are supported on S2600WFT.

6.1.3 Intel® Integrated RAID Module Support

The server board has support for many Intel and third-party PCIe add-in 12Gb RAID adapters which can be installed in available PCIe add-in cards slots. For system configurations with limited add-in card slot availability, an optional Intel® Integrated RAID mezzanine module can be installed onto a high density 80-pin connector (labeled “SAS Module”) on the server board.





Figure 27. Intel® Integrated RAID Module

Please visit the Intel® Server Configurator Tool at the following website for a list of supported Intel® Integrated RAID options:

<https://serverconfigurator.intel.com>

6.2 Onboard Storage Sub-System

The Intel® Server Board S2600WF product family includes support for many storage related technologies and on-board features to support a wide variety of storage options. These include:

- 2 – M.2 PCIe* / SATA
- 4 – PCIe* OCuLink*
- Intel® Volume Management Device (Intel® VMD) for NVMe
- Intel® Virtual RAID on CPU (Intel® VROC) for NVMe
- 2 – 7-pin single port SATA
- 2 – Mini-SAS HD (SFF-8643) 4-port SATA (S2600WFT & S2600WF0 boards only)
- On-Board SATA RAID Options
 - Intel® Rapid Storage Technology 5.0 (Intel® RSTe) for SATA
 - Intel® Embedded Server RAID Technology 2 v1.60 (Intel® ESRT2 1.60) for SATA

The following sections provide an overview of each option. See Chapter 7 for M.2 connector pin-out definition.

6.2.1 M.2 SSD Support

The Intel® Server Board S2600WF product family includes two M.2 SSD connectors labeled “M2_x4PCIE/sSATA_1” and “M2_x2PCIE/sSATA_2” on the server board as shown below.




Figure 28. M.2 Storage Device Connectors

Each M.2 connector can support PCIe or SATA modules that conform to a 2280 (80mm) form factor.

PCIe bus lanes for each connector are routed from the Intel chipset and can be supported in single processor configurations.

The M.2 connector to the left of Riser Slot #1 is supported by PCIe x4 bus lanes and **sSATA-1** from the chipset embedded **sSATA** controller.

The M.2 connector to the right of Riser Slot #1 is supported by PCIe x2 bus lanes and **sSATA-2** from the chip-set embedded **sSATA** controller.

6.2.1.1 Embedded RAID Support

RAID support from embedded RAID options for server board mounted M.2 SSDs is defined as follows:

- Neither Intel® ESRT2 nor Intel® RSTe have RAID support for PCIe* M.2 SSDs when installed to the M.2 connectors on the server board.
 - **Note:** NVMe RAID support using Intel® RSTe VROC requires that the PCIe bus lanes be routed directly from the CPU. On this server board, the PCIe bus lanes routed to the on-board M.2 connectors are routed from the Intel chipset (PCH).
 - **Note:** The Intel® ESRT2 onboard RAID option does not support PCIe devices..
- Both Intel® ESRT2 and Intel® RSTe provide RAID support for SATA devices (See section 6.3.6)
- Neither embedded RAID option supports mixing of M.2 SATA SSDs and SATA hard drives within a single RAID volume

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*

6.2.2 On Board PCIe* OCuLink* Connectors

Depending on the model of the server board installed, the server board will have two (S2600WFQ) or four (S2600WFO & S2600WFT) PCIe* OCuLink connectors to provide the PCIe interface for NVMe* SSDs installed to the front hot swap backplane. PCIe signals for OCuLink connectors “PCIe_SSD0” and “PCIe_SSD1” are routed directly from CPU_1 and PCIe signals for OCuLink connectors “PCIe_SSD2” and “PCIe_SSD3” are directly routed from CPU_2. See Chapter 7 for OCuLink connector pin-out definition.





Figure 29. On-Board OCuLink Connectors

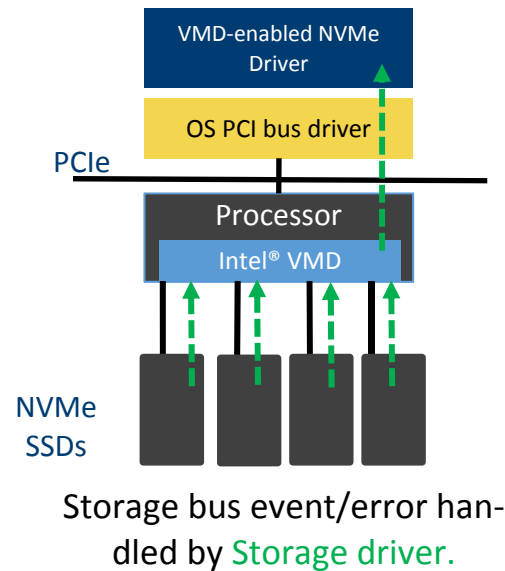
6.2.3 Intel® Volume Management Device (Intel® VMD) for NVMe

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.

NVMe Support w/o Intel® VMD



NVMe Storage with 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.

- 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/disabled 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). (See Figure 12)
- 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 supports hot-plug PCIe* SSDs connected to switch downstream ports. Intel® VMD does not support hot-plug of the switch itself

6.2.3.1 Enabling 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.

See Chapter 4, Table 12. CPU - PCIe* Port Routing, to determine which specific CPU PCIe Root Ports are used to supply the PCIe* bus lanes for on-board OCuLink connectors.

For NVMe devices attached to a riser card via a PCIe Switch or plugged directly into a PCIe add-in card slot, see the following tables in section 6.2 to determine CPU PCIe Root Ports supporting each add-in card slot.

The following tables provide the PCIe* bus routing for all supported risers cards.

Table 13. Riser Slot #1 – PCIe* Root Port Mapping

Table 14. Riser Slot #2 – PCIe* Root Port Mapping

Table 15. Riser Slot #3 – PCIe* Root Port Mapping

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

ADVANCED -> PCI CONFIGURATION -> VOLUME MANAGEMENT DEVICE

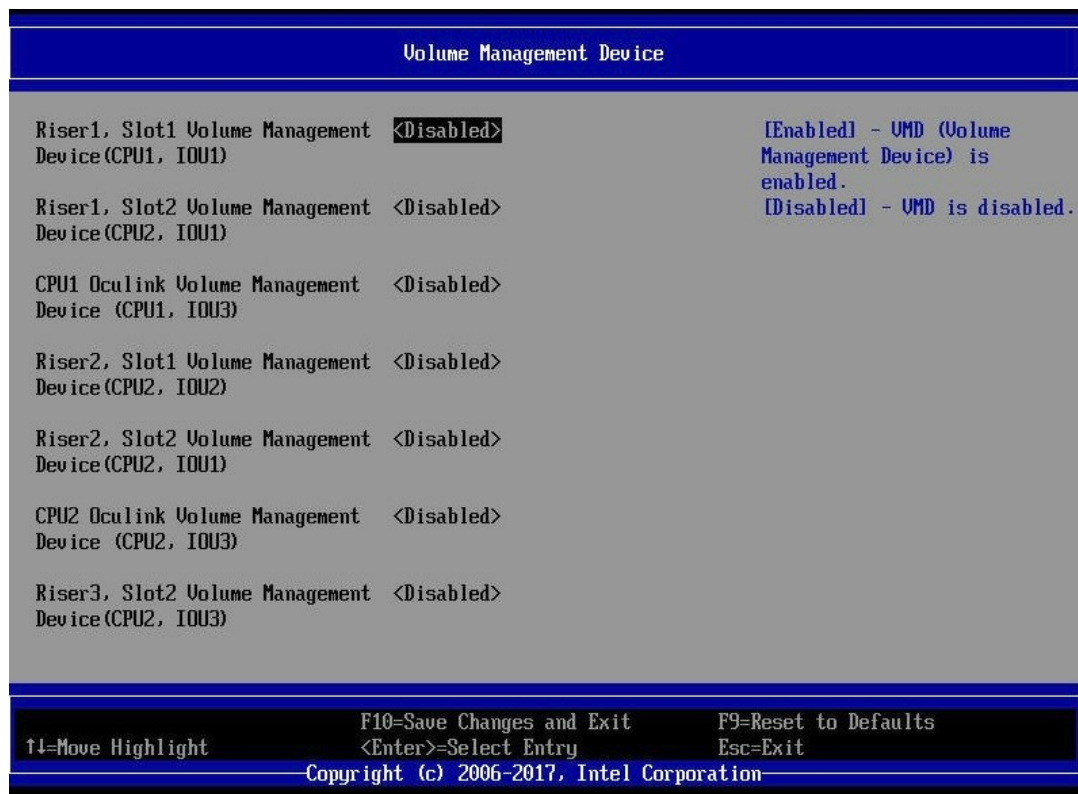


Figure 30. VMD Support Disabled in <F2> BIOS Setup

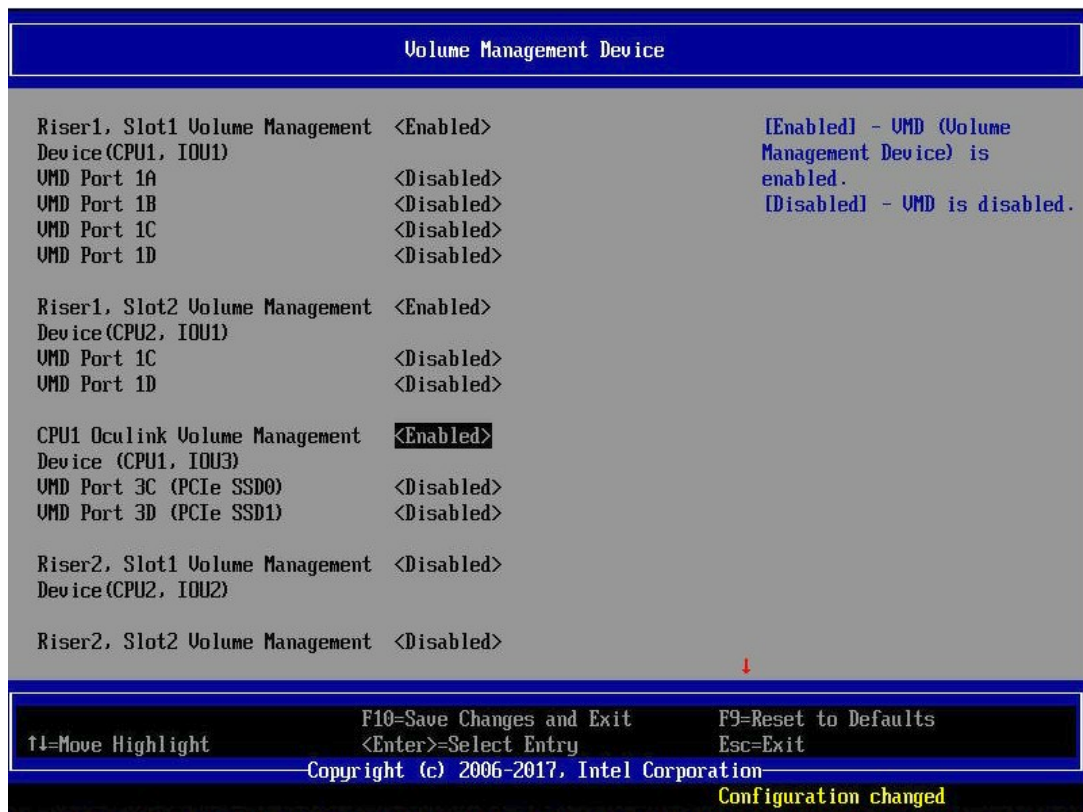



Figure 31. VMD Support Enabled in <F2> BIOS Setup

6.2.4 Intel® Virtual RAID on Chip (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

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




Figure 32. Intel® VROC Upgrade Key

The following table identifies available Intel® VROC upgrade key options.

Table 17. Intel® VROC Upgrade Key Options

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

Note: Intel® VROC Upgrade Keys referenced in Table 12 are used for PCIe* NVMe SSDs only. For SATA RAID support, see Section 6.3.6.

6.2.5 Onboard SATA Support

The server board utilizes two chipset embedded AHCI SATA controllers, identified as **SATA** and **sSATA**, providing for up to twelve 6 Gb/sec Serial ATA (SATA) ports.

The AHCI **sSATA** controller provides support for up to 4 SATA ports on the server board:

- Two ports accessed via two white single port 7-pin connectors labeled “**sSATA-4**” and “**sSATA-5**” on the server board
- Two ports (**sSATA 1** and **sSATA 2**) via two M.2 SSD connectors

The AHCI **SATA** controller provides support for up to 8 SATA ports on the server board: (S2600WFT & S2600W0 boards only)

- Four ports from the Mini-SAS HD (SFF-8643) connector labeled “**SATA Ports 0-3**” on the server board
- Four ports from the Mini-SAS HD (SFF-8643) connector labeled “**SATA Ports 4-7**” on the server board




Figure 33. On-Board SATA Port Connector Identification

Note: The onboard SATA controllers are not compatible with and cannot be used with SAS Expander Cards.

Table 18. SATA and sSATA Controller Feature Support

Feature	Description	AHCI Mode	RSTe RAID	ESRT2 RAID
Native Command Queuing (NCQ)	Allows the device to reorder commands for more efficient data transfers	Supported	Supported	
Auto Activate for DMA	Collapses a DMA Setup then DMA Activate sequence into a DMA Setup only	Supported	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	Supported	Supported	
Asynchronous Signal Recovery	Provides a recovery from a loss of signal or establishing communication after hot plug	Supported	Supported	
6 Gb/s Transfer Rate	Capable of data transfers up to 6 Gb/s	Supported	Supported	Supported
ATAPI Asynchronous Notification	A mechanism for a device to send a notification to the host that the device requires attention	Supported	Supported	
Host and Link Initiated Power Management	Capability for the host controller or device to request Partial and Slumber interface power states	Supported	Supported	
Staggered Spin-Up	Enables the host the ability to spin up hard drives sequentially to prevent power load problems on boot	Supported	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	Supported	N/A	

The SATA controller and the sSATA controller can be independently enabled and disabled and configured through the BIOS setup utility under the Mass Storage Controller Configuration menu screen. The following table identifies supported setup options.

Table 19. SATA and sSATA Controller BIOS Utility Setup Options

SATA Controller	sSATA Controller	Supported
AHCI	AHCI	Yes
AHCI	Disabled	Yes
AHCI	Intel® RSTe	Yes
AHCI	Intel® Embedded Server RAID Technology 2	Microsoft Windows* only
Disabled	AHCI	Yes
Disabled	Disabled	Yes
Disabled	Intel® RSTe	Yes
Disabled	Intel® Embedded Server RAID Technology 2	Yes
Intel® RSTe	AHCI	Yes
Intel® RSTe	Disabled	Yes
Intel® RSTe	Intel® RSTe	Yes
Intel® RSTe	Intel® Embedded Server RAID Technology 2	No
Intel® Embedded Server RAID Technology 2	AHCI	Microsoft Windows only
Intel® Embedded Server RAID Technology 2	Disabled	Yes
Intel® Embedded Server RAID Technology 2	Intel® RSTe	No
Intel® Embedded Server RAID Technology 2	Intel® Embedded Server RAID Technology 2	Yes

6.2.5.1 Staggered Disk Spin-Up

Hard Disks, or spinning storage media, have a peak power draw when a system is first powered on and the media is spinning up and initializing. With the number of SATA hard disk drives that can be supported from the embedded server board SATA controllers, it is possible for the given system configuration to exceed the power draw of a given system power supply should all hard drives spin up at the same time.

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 controllers, 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.

6.2.6 On-Board SATA RAID Options

The server board includes support for two embedded SATA RAID options:

- Intel® Rapid Storage Technology (RSTe) 5.0 (See Section 6.3.6.1)
- Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 (See Section 6.3.6.2)

By default, onboard RAID options are disabled in BIOS setup. To enable onboard RAID support, access the BIOS setup utility during POST. The onboard RAID options can be found under the **sSATA Controller** or **SATA Controller** options under the following BIOS setup menu:

Advanced -> Mass Storage Controller Configuration



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

Intel® Rapid Storage Technology offers several 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. Supported RAID levels include 0, 1, 5, and 10.

- **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.

6.2.6.2 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**)




Figure 34. ESRT2 SATA RAID-5 Upgrade Key

- **RAID 5:** Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.

6.3 Rear External RJ45 Connector Overview

The back edge of the server board includes several RJ45 connectors providing support for the following onboard features:

- Dedicated server management port
- Network Interface Connectors (S2600WFT Only)
- Serial-A port




Figure 35. Rear External RJ45 Connectors

RJ45 connectors used for the Dedicated Management Port and Network Interface connectors include two LEDs.



Figure 36. RJ45 Connector 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.

Table 20. 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	Transmit / Receive Activity
Right	Amber	On	1Gb data rate
	Green	On	10 Gb data rate

6.3.1 RJ45 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. See **Error! Reference source not found.** See Chapter 7 for additional information about onboard Server Management support.

6.3.2 RJ45 Network Interface Connectors (S2600WFT only)

On the back edge of the Intel Server Board S2600WFT are two RJ45 networking ports, “NIC #1” and “NIC #2” and one RJ45 Dedicated Management Port (all board models).

6.3.2.1 On-board Intel® Ethernet Controller

The 2600WFT model of the server board family includes the following onboard Intel® Ethernet Controller:

- Intel® Ethernet Controller X557-AT2 10 GbE

Refer to the respective product data sheet for a complete list of supported Ethernet Controller features.

6.3.3 Serial Port Support

The server board has support for two serial ports, Serial-A and Serial-B.

Serial A is an external RJ45 type connector located on the back edge of the server board.




Figure 37. RJ45 Serial-A Pin Orientation

Table 21. Serial-A Connector Pin-out

Signal Description	Pin#
RTS	1
DTR	2
SOUT	3
GROUND	4
RI	5
SIN	6
DCD or DSR	7**
CTS	8

** Pin 7 of the RJ45 Serial A connector is configurable to support either a DSR (Default) signal or a DCD signal. Pin 7 signals are changed by moving the jumper on the jumper block labeled “J4A2” from pins 1-2 (default) to pins 2-3.




Figure 38. J4A2 Jumper Block for Serial A Pin 7 Configuration

Serial B is provided through an internal DH-10 header labeled “Serial_B” on the server board




Figure 39. Serial-B Connector (Internal)

Table 22. Serial-B Connector Pin-out

Signal Description	Pin#	Pin#	Signal Description
DCD	1	2	DSR
SIN	3	4	RTS
SOUT	5	6	CTS
DTR	7	8	RI
GROUND	9		KEY

6.4 USB Support

USB support is provided through several on board internal and external connectors as described in the following sections.

6.4.1 External USB 3.0 Support

The server board includes three (1x3 stacked) USB 3.0 ports on the back edge of the server board.




Figure 40. External USB 3.0 Ports

6.4.2 Internal USB 2.0 Type-A Connector

The server board includes one internal Type-A USB 2.0 connector.




Figure 41. Internal USB 2.0 Type-A Connector

6.4.3 Front Panel USB 3.0 Support

A Blue 20-pin (2x10) shrouded connector on the server board (labeled “FP_USB_2.0/3.0”) provides the option of routing two USB 3.0 ports to the front of a given chassis.




Figure 42. Front Panel USB 3.0 Connector


Note: The following USB ports are routed to this connector: USB 3.0 ports 1 and 2, USB 2.0 ports 11 and 13

Table 23. Front Panel USB 2.0/3.0 Connector Pin-out (“FP_USB_2.0/ 3.0”)

Signal Name	Pin#	Pin#	Signal Name
		1	P5V_USB_FP
P5V_USB_FP	19	2	USB3_04_RXN
USB3_01_RXN	18	3	USB3_04_RXP
USB3_01_RXP	17	4	GROUND
GROUND	16	5	USB3_04_TXN
USB3_01_TXN	15	6	USB3_04_TXP
USB3_01_TXP	14	7	GROUND
GROUND	13	8	USB2_13_DN
USB2_10_DN	12	9	USB2_13_DP
USB2_10_DP	11	10	USB3_ID

6.4.4 Front Panel USB 2.0 Connector

The server board includes a 10-pin connector that, when cabled, can provide up to two USB 2.0 ports to a front panel. On the server board, the connector is labeled “FP_USB_2.0_5-6” and is located on the left side, near the I/O module connector. Table 24 provides the connector pin-out.

**Figure 43. Front Panel USB 2.0 Connector****Table 24. Front Panel USB 2.0 Connector Pin-out ("FP_USB_2.0_5-6 ")**

Signal Name	Pin#	Pin#	Signal Name
P5V_USB_FP	1	2	P5V_USB_FP
USB2_P11_F_DN	3	4	USB2_P13_F_DN
USB2_P11_F_DP	5	6	USB2_P13_F_DP
GROUND	7	8	GROUND
		10	TP_USB2_FP_10

6.5 Video Support

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:

2D Mode	2D Video Support (Color Bit)			
Resolution	8 bpp	16 bpp	24 bpp	32 bpp
640 x 480	60, 72, 75, 85	60, 72, 75, 85	Not Supported	60, 72, 75, 85
800 x 600	60, 72, 75, 85	60, 72, 75, 85	Not Supported	60, 72, 75, 85
1024 x 768	60, 72, 75, 85	60, 72, 75, 85	Not Supported	60, 72, 75, 85
1152 x 864	75	75	75	75
1280 x 800	60	60	60	60
1280 x 1024	60	60	60	60
1440 x 900	60	60	60	60
1600 x 1200	60	60	Not Supported	Not Supported
1680 x 1050	60	60	Not Supported	Not Supported
1920 x 1080	60	60	Not Supported	Not Supported
1920 x 1200	60	60	Not Supported	Not Supported

6.5.1 Onboard Video Connectors

The server board includes two options to attached a monitor to the server system:

- A standard 15-pin video connector is located on the back edge of the server board..




Figure 44. Rear External Video Connector

- On the server board near the front right edge, is a connector labeled "FP_VIDEO", which when cabled, can provide video from the front of the server system. When a monitor is attached to the front of the system, the video out the back is disabled.




Figure 45. Front Panel Video Connector

The following table provides the pin-out for this connector.

Table 25. Front Panel Video Connector Pin-out ("FP VIDEO")

Signal Description	Pin#	Pin#	Signal Description
V_IO_FRONT_R_CONN	1	2	GROUND
V_IO_FRONT_G_CONN	3	4	GROUND
V_IO_FRONT_B_CONN	5	6	GROUND
V_BMC_GFX_FRONT_VSYN	7	8	GROUND

Signal Description	Pin#	Pin#	Signal Description
V_BMC_GFX_FRONT_HSYN	9		KEY
V_BMC_FRONT_DDC_SDA_CONN	11	12	V_FRONT_PRES_N
V_BMC_FRONT_DDC_SCL_CONN	13	14	P5V_VID_CONN_FNT

6.5.2 Onboard Video and Add-in Video Adapter Support

Add-in video cards can be used to either replace or complement the onboard video option of the server board.

<F2> BIOS Setup includes options to support the desired video operation when a add-in video card is installed.

- When both the Onboard Video Add-in Video Adapter options are set to Enabled, then both video displays can be active. The onboard video is still the primary console and active during BIOS POST; the add-in video adapter would only be active under an OS environment with video driver support.
- When onboard video is Enabled, and the Add-in Video Adapter is Disabled, then only the onboard video will be active.
- When onboard video is Disabled, and Add-in Video Adapter is Enabled, then only the add-in video adapter will be active.

Configurations with add-in video cards can get more complicated with a dual CPU socket board. Some multi-socket boards have PCIe slots capable of hosting an add-in video card which are attached to the IIOs of CPU sockets other than CPU Socket 1. However, only one CPU Socket can be designated as “Legacy VGA Socket” as required in POST. To provide for this, there is another PCI Configuration option to control “Legacy VGA Socket”. The rules for this are:

- The option in <F2> BIOS Setup is grayed out and unavailable unless an add-in video card is installed in a PCIe slot supported by CPU 2
- Because the Onboard Video is “hardwired” to CPU Socket 1, whenever Legacy VGA Socket is set to a CPU Socket 2, the Onboard Video is disabled.

6.5.3 Dual Monitor Support

The BIOS supports single and dual video when add-in video adapters are installed. Although there is no enable/disable option in <F2> BIOS Setup for Dual Video, it works when both “Onboard video” and “Add-in Video Adapter” options are enabled.

In the single video mode, the onboard video controller or the add-in video adapter is detected during POST.

In dual video mode, the onboard video controller is enabled and is the primary video device while the add-in video adapter is allocated resources and is considered as the secondary video device during POST. The add-in video adapter won't be active until the OS environment is loaded.

7. On-board Connector/Header Pin-Out Definition

This section identifies the location and pin-out for most on-board connectors and headers of the server board. Information for some connectors and headers will be found elsewhere in the document where the feature is described in more detail.

Pinout definition for the following onboard connectors is only made available by obtaining the board schematics directly from Intel (NDA Required).

- All Riser Slots
- OCP Module Connector
- SAS Module Connector
- M.2 SSD Connectors
- DIMM Slots
- Processor Sockets

7.1 Power Connectors

The server board includes several power connectors that are used to provide DC power to various devices.

7.1.1 Main Power

Main server board power is supplied from two slot connectors, which allow for one or two (redundant) power supplies to dock directly to the server board. Each connector is labeled as “MAIN PWR 1” or “MAIN PWR 2” on the server board, as shown in Figure 46. The server board provides no option to support power supplies with cable harnesses. In a redundant power supply configuration, a failed power supply module is hot-swappable. Table 26 provides the pin-out mapping for the “MAIN PWR 1” connector and Table 27 provides the pin-out mapping for the “MAIN PWR 2” connector.




Figure 46. MAIN PWR 1” and “MAIN PWR 2” Connectors

Table 26. Main Power (Slot 1) Connector Pin-out ("MAIN PWR 1")

Signal Name	Pin #	Pin#	Signal Name
GROUND	B1	A1	GROUND
GROUND	B2	A2	GROUND
GROUND	B3	A3	GROUND
GROUND	B4	A4	GROUND
GROUND	B5	A5	GROUND
GROUND	B6	A6	GROUND
GROUND	B7	A7	GROUND
GROUND	B8	A8	GROUND
GROUND	B9	A9	GROUND
P12V	B10	A10	P12V
P12V	B11	A11	P12V
P12V	B12	A12	P12V
P12V	B13	A13	P12V
P12V	B14	A14	P12V
P12V	B15	A15	P12V
P12V	B16	A16	P12V
P12V	B17	A17	P12V
P12V	B18	A18	P12V
P3V3_AUX: PD_PS1_FRU_A0	B19	A19	SMB_PMBUS_DATA_R
P3V3_AUX: PD_PS1_FRU_A1	B20	A20	SMB_PMBUS_CLK_R
P12V_STBY	B21	A21	FM_PS_EN_PSU_N
FM_PS_CR1	B22	A22	IRQ_SML1_PMBUS_ALERTR2_N
P12V_SHARE	B23	A23	ISENSE_P12V_SENSE_RTN
TP_1_B24	B24	A24	ISENSE_P12V_SENSE
FM_PS_COMPATIBILITY_BUS	B25	A25	PWRGD_PS_PWROK

Table 27. Main Power (Slot 2) Connector Pin-out ("MAIN PWR 2")

Signal Name	Pin #	Pin#	Signal Name
GROUND	B1	A1	GROUND
GROUND	B2	A2	GROUND
GROUND	B3	A3	GROUND
GROUND	B4	A4	GROUND
GROUND	B5	A5	GROUND
GROUND	B6	A6	GROUND
GROUND	B7	A7	GROUND
GROUND	B8	A8	GROUND
GROUND	B9	A9	GROUND
P12V	B10	A10	P12V
P12V	B11	A11	P12V
P12V	B12	A12	P12V
P12V	B13	A13	P12V
P12V	B14	A14	P12V
P12V	B15	A15	P12V

Signal Name	Pin #	Pin#	Signal Name
P12V	B16	A16	P12V
P12V	B17	A17	P12V
P12V	B18	A18	P12V
P3V3_AUX: PU_PS2FRU_A0	B19	A19	SMB_PMBUS_DATA_R
P3V3_AUX: PD_PS2_FRU_A1	B20	A20	SMB_PMBUS_CLK_R
P12V_STBY	B21	A21	FM_PS_EN_PSU_N
FM_PS_CR1	B22	A22	IRQ_SML1_PMBUS_ALERTR3_N
P12V_SHARE	B23	A23	ISENSE_P12V_SENSE_RTN
TP_2_B24	B24	A24	ISENSE_P12V_SENSE
FM_PS_COMPATIBILITY_BUS	B25	A25	PWRGD_PS_PWROK

7.1.2 Hot Swap Backplane Power Connector

The server board includes one white 2x6-pin power connector that when cabled provides power for hot swap backplanes, as shown in Figure 47. On the server board, this connector is labeled as “HSBP PWR”.




Figure 47. HSBP PWR Connector

Table 28. Hot Swap Backplane Power Connector Pin-out

SIGNAL NAME	PIN	PIN	SIGNAL NAME
GND	1	7	P12V_240VA3
GND	2	8	P12V_240VA3
GND	3	9	P12V_240VA2
GND	4	10	P12V_240VA2
GND	5	11	P12V_240VA1
GND	6	12	P12V_240VA1

7.1.3 Riser Card Supplemental 12V Power Connectors

The server board includes two white 2x2-pin power connectors labeled “OPT_12V_PWR” that provide supplemental 12V power-out to high power PCIe* x16 add-in cards (Video, GPGPU, Intel® Xeon Phi™) that have power requirements that exceed the 75W maximum power supplied by the riser card slot. These connectors are identified in Figure 48. A cable from these connectors may be routed to a power-in connector on the given add-in card. Maximum power draw for each connector is 225W, but is also limited by available power provided by the power supply and the total power draw of the given system configuration. A power budget for the complete system should be performed to determine how much supplemental power is available to support any high-power add-in cards.




Figure 48. 12V Power Connectors

Table 29 provides the pin-out values for the 12V power connectors.

Table 29. Riser Slot Auxiliary Power Connector Pin-out ("OPT_12V_PWR")

Signal Name	Pin#	Pin#	Signal Name
P12V	3	1	GROUND
P12V	4	2	GROUND

Intel makes available a 12V supplemental power cable that can support both 6 and 8 pin 12V AUX power connectors found on high power add-in cards. The power cable (as shown in Figure 49) is available as a separate orderable accessory kit from Intel using the following Intel product code: *AXXGPGPUCABLE*.




Figure 49. High Power Add-in Card 12V Auxiliary Power Cable Option

7.1.4 Peripheral Power Connector

The server board includes one 6-pin power connector intended to provide power for peripheral devices such as Optical Disk Drives (ODD) and/or Solid State Devices (SSD). On the server board this connector is labeled as "Peripheral_PWR". The following table provides the pin-out for this connector.




Figure 50. Peripheral Power Connector

Table 30. Peripheral Drive Power Connector Pin-out ("Peripheral_PWR")

Signal Name	Pin#	Pin#	Signal Name
P12V	4	1	P5V
P3V3	5	2	P5V
GROUND	6	3	GROUND

7.2 Front Control Panel Headers and Connectors

The server board includes several connectors that provide various possible front panel options. This section provides a functional description and pin-out for each connector.

For Front Panel control button and LED support, the server board includes two connector options: a 30-pin SSI compatible front panel header labeled “FRONT_PANEL”, and a custom high density 30-pin front panel connector, labeled “STORAGE_FP”.




Figure 51. Front Control Panel Connectors

Supported control buttons and LEDs are identified in Table 31.

Table 31. Front Panel Control Button and LED Support

Control Button/LED	Support
Power / Sleep Button	Yes
System ID Button	Yes
System Reset Button	Yes
NMI Button	Yes
NIC Activity LED	Yes
Storage Device Activity LED	Yes
System Status LED	Yes
System ID LED	Yes

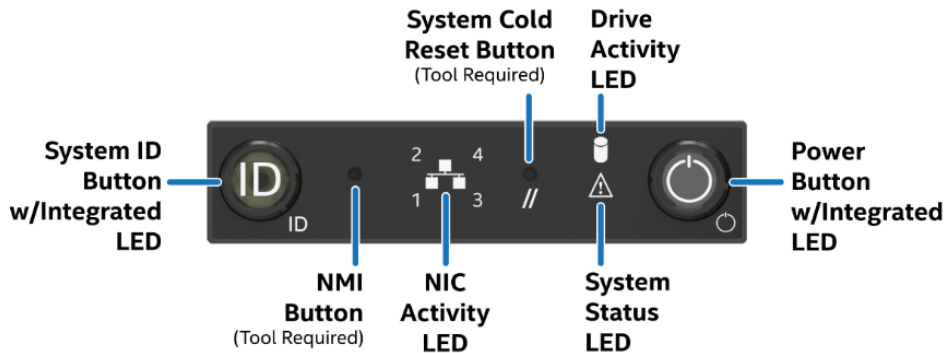


Figure 52. Example - Front Control Panel View (For reference purposes only)

The pinout for both connector types is identical.

Table 32. 30-pin Front Panel Connector Pin-outs

Signal Name	Pin#	Pin#	Signal Name
P3V3_AUX	1	2	P3V3_AUX
KEY		4	P5V_STBY
FP_PWR_LED_BUF_R_N	5	6	FP_ID_LED_BUF_R_N
P3V3	7	8	FP_LED_STATUS_GREEN_R_N
LED_HDD_ACTIVITY_R_N	9	10	FP_LED_STATUS_AMBER_R_N
FP_PWR_BTN_N	11	12	LED_NIC_LINK0_ACT_FP_N
GROUND	13	14	LED_NIC_LINK0_LNKUP_FP_N
FP_RST_BTN_R_N	15	16	SMB_SENSOR_3V3STBY_DATA_R0
GROUND	17	18	SMB_SENSOR_3V3STBY_CLK
FP_ID_BTN_R_N	19	20	FP_CHASSIS_INTRUSION
PU_FM_SIO_TEMP_SENSOR	21	22	LED_NIC_LINK1_ACT_FP_N
FP_NMI_BTN_R_N	23	24	LED_NIC_LINK1_LNKUP_FP_N
KEY			KEY
LED_NIC_LINK2_ACT_FP_N	27	28	LED_NIC_LINK3_ACT_FP_N
LED_NIC_LINK2_LNKUP_FP_N	29	30	LED_NIC_LINK3_LNKUP_FP_N

7.2.1 Front Panel LED and Control Button Features Overview

7.2.1.1 Power/Sleep Button and LED Support

Pressing the Power button will toggle the system power on and off. This button also functions as a sleep button if enabled by an ACPI compliant operating system. Pressing this button will send a signal to the integrated BMC, which will power on or power off the system. The power LED is a single color and is capable of supporting different indicator states as defined in Table 33.

Table 33. Power/Sleep LED Functional States

State	Power Mode	LED	Description
Power-off	Non-ACPI	Off	System power is off, and the BIOS has not initialized the chipset.
Power-on	Non-ACPI	On	System power is on
S5	ACPI	Off	Mechanical is off, and the operating system has not saved any context to the hard disk.
S0	ACPI	Steady on	System and the operating system are up and running.

7.2.1.2 System ID Button and LED Support

Pressing the System ID Button will toggle both the ID LED on the front panel and the Blue ID LED on the back edge of the server board, on and off. The System ID LED is used to identify the system for maintenance when installed in a rack of similar server systems. The System ID LED can also be toggled on and off remotely using the IPMI “Chassis Identify” command which will cause the LED to blink for 15 seconds.

7.2.1.3 System Reset Button Support

When pressed, this button will reboot and re-initialize the system.

7.2.1.4 NMI Button Support

When the NMI button is pressed, it puts the server in a halt state and causes the BMC to issue a non-maskable interrupt (NMI) for generating diagnostic traces and core dumps from the operating system. Once an NMI has been generated by the BMC, the BMC does not generate another NMI until the system has been reset or powered down.

The following actions cause the BMC to generate an NMI pulse:

- Receiving a *Chassis Control* command to pulse the diagnostic interrupt. This command does not cause an event to be logged in the SEL.
- Watchdog timer pre-timeout expiration with NMI/diagnostic interrupt pre-timeout action enabled.

Table 34 describes behavior regarding NMI signal generation and event logging by the BMC.

Table 34. NMI Signal Generation and Event Logging

Causal Event	NMI	
	Signal Generation	Front Panel Diag Interrupt Sensor Event Logging Support
Chassis Control command (pulse diagnostic interrupt)	X	–
Front panel diagnostic interrupt button pressed	X	X
Watchdog Timer pre-timeout expiration with NMI/diagnostic interrupt action	X	X

7.2.1.5 NIC Activity LED Support

The Front Control Panel includes an activity LED indicator for each on-board Network Interface Controller (NIC). When a network link is detected, the LED will light up constantly. The LED will begin to blink once network activity occurs at a rate that is consistent with the amount of network activity that is occurring.

7.2.1.6 Storage Device Activity LED Support

The storage device activity LED on the front panel indicates drive activity from the on-board storage controllers. The server board also provides a 2-pin header, labeled “HDD_Activity” on the server board, giving access to this LED for add-in controllers.

7.2.1.7 System Status LED Support

The System Status LED is a bi-color (Green/Amber) indicator that shows the current health of the server system. The system provides two locations for this feature; one is located on the Front Control Panel, the other is located on the back edge of the server board, viewable from the back of the system. Both LEDs are tied together and will show the same state. The System Status LED states are driven by the on-board platform management sub-system.

7.3 System Fan Connectors

The server board is capable of supporting up to a total of six system fans. Each system fan includes a pair of fan connectors; a 1x10 pin connector to support a dual rotor cabled fan, typically used in 1U system configurations, and a 2x3 pin connector to support a single rotor hot swap fan assembly, typically used in 2U system configurations. Concurrent use of both fan connector types for any given system fan pair is not supported.

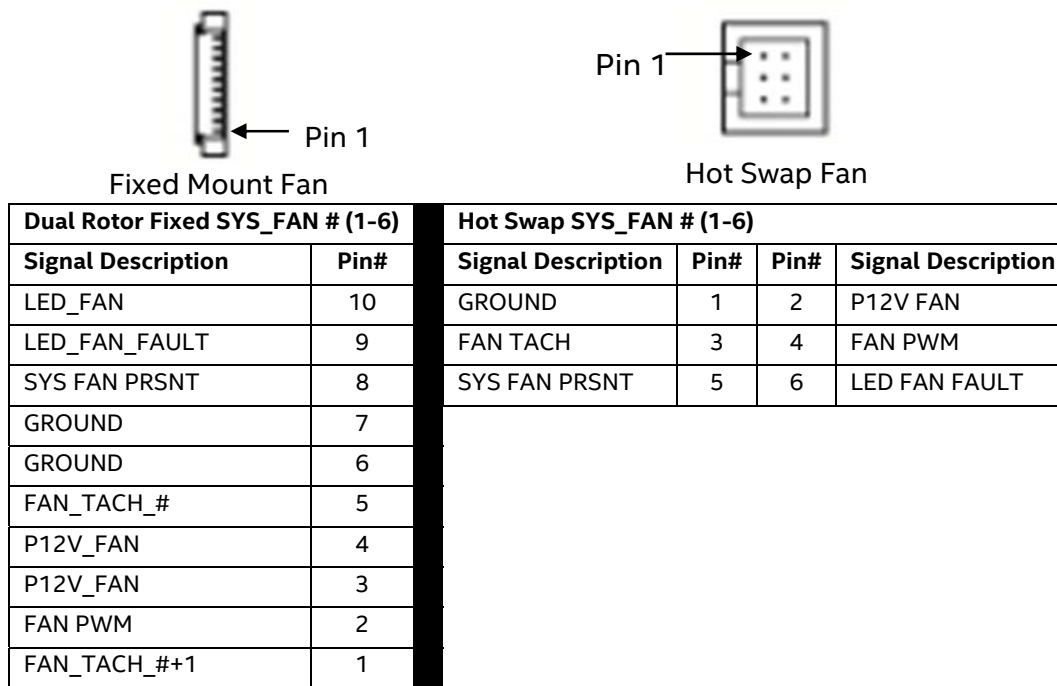



Figure 53. System Fan Connector Pin-outs

Each connector is monitored and controlled by on-board platform management. On the server board, each system fan connector pair is labeled “SYS_FAN #”, where # = 1 – 6. The following illustration shows the location of each system fan connector on the server board.



7.4 Management Connectors

The server board includes several management interface connectors. The following tables provide the pinout definition for each.



4-pin HSBP I2C (J1K1)
(Lower Left Corner)



Table 35. Hot Swap Backplane I2C Connector – SMBUS 3-pin (J5C3)

Pin	Signal
1	SDA
2	Ground
3	SCL

Table 36. Hot Swap Backplane I2C Connector – SMBUS 4-pin (J1K1)

Pin	Signal
1	SDA
2	Ground
3	SCL
4	RST_PCIE_SSD_PERST

Table 37. IPMB – SMBUS 4-pin (J1C3)

Pin	Signal
1	CMOS_SDA
2	Ground
3	CMOS_SCL
4	P5V_AUX

8. Standard and Advanced Server Management Features

The integrated BMC has support for Standard and advanced server management features. Standard management features are available by default. Advanced management features are enabled with the addition of an optionally installed Intel® Remote Management Module 4 Lite (Intel® RMM4 Lite) enablement key, described in Table 38.

Table 38. 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 and media redirection

On the server board, the Intel® RMM4 Lite key is installed at the location indicated in Figure 54.




Figure 54. RMM 4 Lite Key Placement

When the BMC FW initializes, it attempts to access the Intel® RMM4 Lite. If the attempt to access the Intel® RMM4 Lite is successful, then the BMC activates the advanced features.

The following table identifies both standard and advanced server management features.

Table 39. Standard and Advanced Server Management Features

Feature	Standard	RMM4-Lite (Advanced)
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

8.1 Dedicated Management Port

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




Figure 55. Dedicated Management Port

8.2 Embedded Web Server

BMC Base manageability provides an embedded web server and an OEM-customizable web GUI that 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 from the on-board dedicated management port. At least two concurrent web sessions, from up to two different users, is supported. The embedded web user interface shall support the following client web browsers:

- Microsoft Internet Explorer*
- Mozilla Firefox*
- Google Chrome*
- Safari*

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 256-bit secure sockets layer (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 disables 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 a greyed-out font on that user's UI display. The web GUI also provides a launch point for advanced features, KVM and media redirection. These features are also grayed out in the GUI unless the system has been updated to support these advanced features. The embedded web server displays US English or Chinese language output only.

Additional features supported by the web GUI can

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

- Using the Embedded Platform Debug feature, allow the user to initiate a “debug dump” to a file that can be sent to Intel® for debug purposes
- Employ the Virtual Front Panel to provide 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.
- Save the SEL to a file
- Force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display processor and memory information that is available over IPMI over LAN.
- Get and set Node Manager (NM) power policies
- Display power consumed by the server
- View and configure VLAN settings
- Warn users that the reconfiguration of IP address will cause a disconnect
- 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.
- Employ Server Power Control to force boot-up into Setup on a reset
- Report 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 displayed as a time relative to the start of POST or the previous POST code.
- Allow customization of 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.

8.3 Advanced Management Feature Support (Intel® RMM4 Lite)

The integrated baseboard 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-lite option 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 baseboard management controller, utilizing expanded capabilities enabled by the Intel RMM4 hardware.

Key Features of the RMM4-lite enablement key are:

- KVM redirection from either the dedicated management NIC or the server board NICs used for management traffic; up to two KVM sessions. 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.
- 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.

8.3.1 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. USB1.1 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.
- Support user definable keyboard macros.

KVM redirection feature supports the following resolutions and refresh rates:

- 640x480 at 60Hz, 72Hz, 75Hz, 85Hz, 100Hz
- 800x600 at 60Hz, 72Hz, 75Hz, 85Hz
- 1024x768 at 60Hz, 72Hz, 75Hz, 85Hz
- 1280x960 at 60Hz
- 1280x1024 at 60Hz
- 1600x1200 at 60Hz
- 1920x1080 (1080p) at 60Hz
- 1920x1200 (WUXGA+) at 60Hz
- 1650x1080 (WSXGA+) at 60Hz

8.3.2 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.

8.3.3 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.

8.3.4 Availability

The remote KVM session is available even when the server is powered off (in stand-by mode). No restart 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 FW update) will require the session to be re-established. KVM sessions persist across system reset, but not across an AC power loss.

8.3.5 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.

8.3.6 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.

8.3.7 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.

8.3.8 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 usable 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 FW update) will require the session to be re-established
- The mounted device is visible to (and usable 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 install.

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.

8.3.8.1 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.

8.3.8.2 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, reference the Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide.

9. Light Guided Diagnostics

The server board includes several on-board LED indicators to aid troubleshooting various board level faults. The following diagram shows the location for each LED.

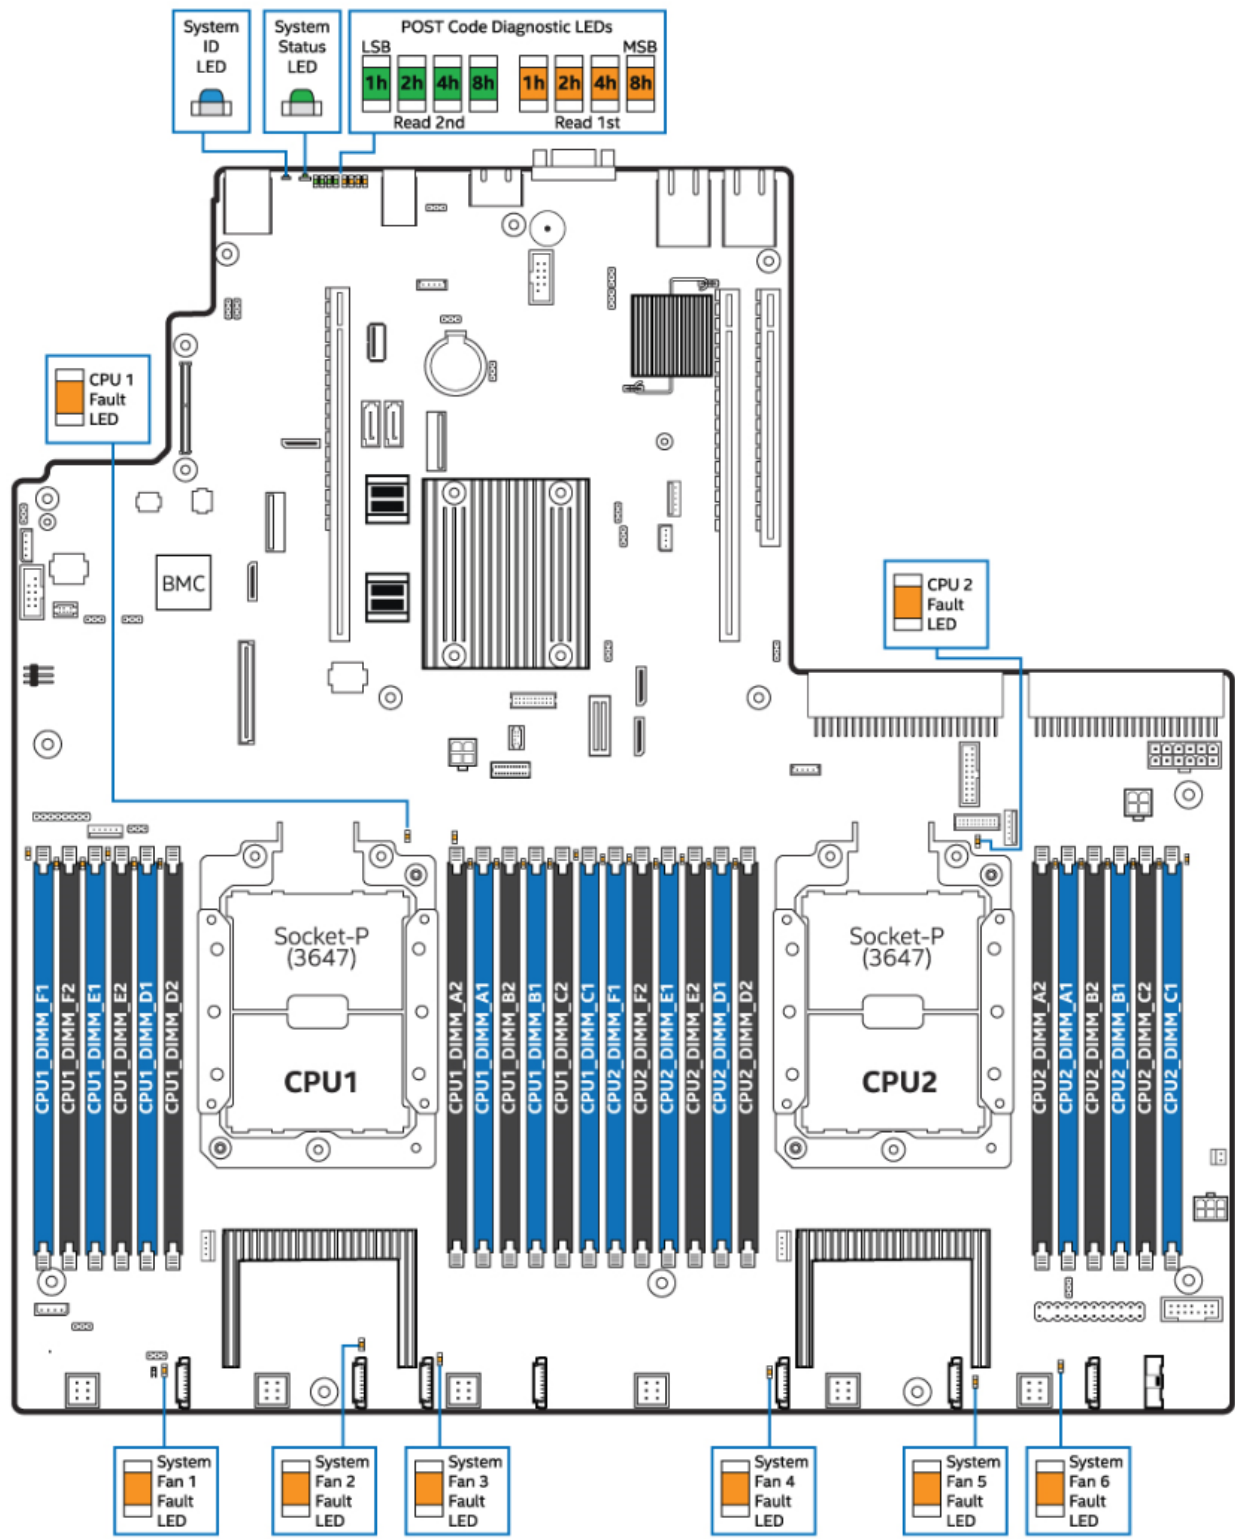


Figure 56. On-Board Diagnostic LED Placement

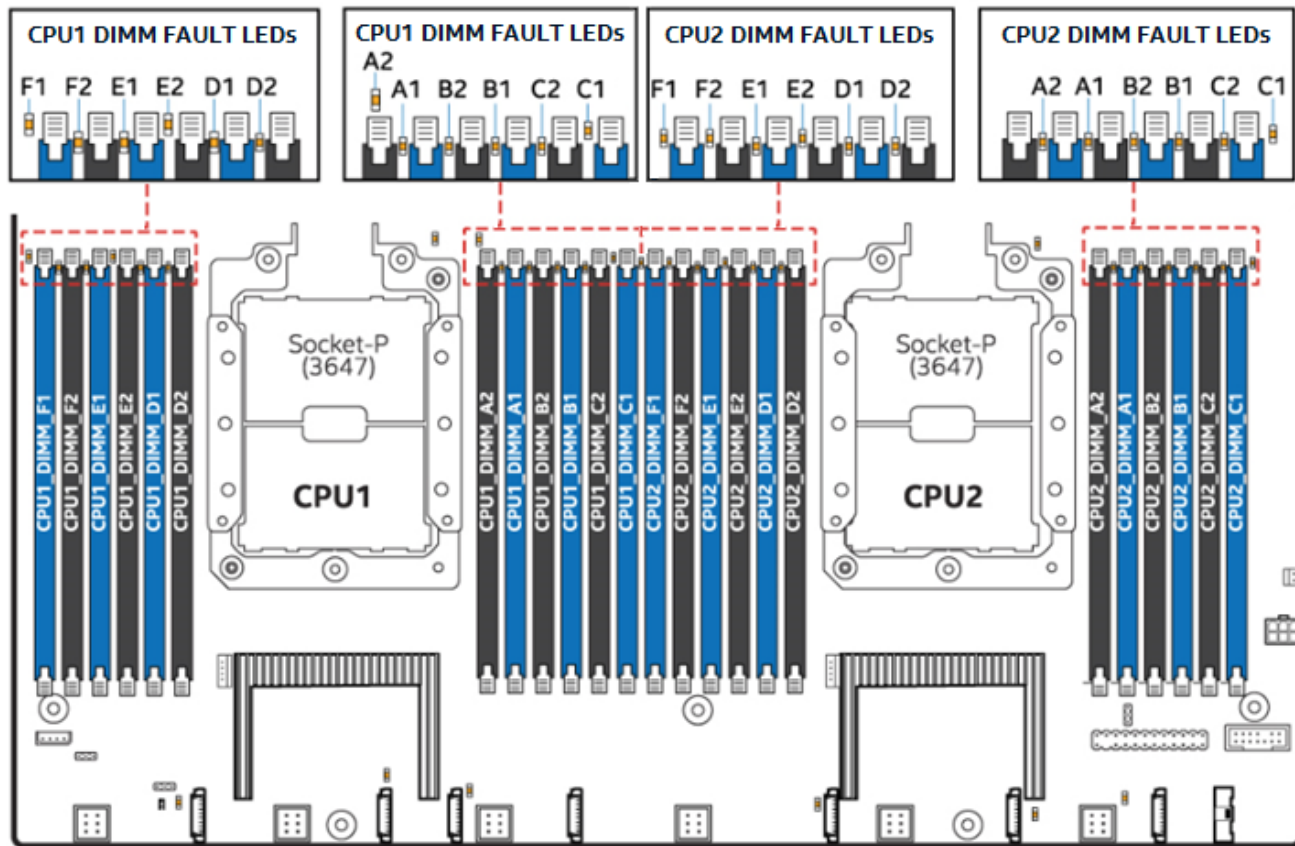


Figure 57.DIMM Fault LEDs

9.1 System ID LED

The server board includes a blue system ID LED which is used to visually identify a specific server installed among many other similar servers. There are two options available for illuminating the System ID LED.

1. The front panel ID LED Button is pushed, which causes the LED to illuminate to a solid on state until the button is pushed again.
2. An IPMI "Chassis Identify" command is entered remotely, which causes the LED to blink.

The System ID LED on the server board is tied directly to the System ID LED on system front panel if present.

9.2 System Status LED

The server board includes a bi-color System Status LED. The System Status LED on the server board is tied directly to the System Status LED on the front panel (if present). This LED indicates the current health of the server. Possible LED states include solid green, blinking green, blinking amber, and solid amber.

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

When AC power is first applied to the system, the status LED turns solid amber and then immediately changes to blinking green to indicate that the BMC is booting. If the BMC boot process completes with no errors, the status LED will change to solid green.

Table 40 lists and describes the states of the system status LEDs.

Table 40. System Status LED States

Color	State	System Status	Description
Green	Solid on	Ok	<p>Indicates that the System 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> <ol style="list-style-type: none"> 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.
Green	~1 Hz blink	Degraded	<p>System Degraded:</p> <ol 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 more 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) ¹. Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the user no longer has spared DIMMs indicating 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. Power Unit sensor offset for configuration error is asserted. HDD HSC is off-line or degraded. Hard drive fault

Color	State	System Status	Description
Amber	~1 Hz blink	Warning	<p>Warning alarm – system is likely to fail:</p> <ol style="list-style-type: none"> 1. 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. 2. VRD Hot asserted. 3. Minimum number of fans to cool the system not present or failed 4. Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present)

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 the AC power is first applied. (DC power on/off will not reset BMC.) BMC reset will occur after a BMC firmware update, on receiving a BMC cold reset command, and following a reset initiated by the BMC Watchdog. The following table defines the LED states during the BMC Boot/Reset process.

Table 41. 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

A bank of eight POST code diagnostic LEDs are located on the back edge of the server next to the stacked USB connectors. See Figure 56. 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. The purpose of these LEDs is to assist in troubleshooting a system hang condition during the POST process. The diagnostic LEDs can be used to identify the last POST process to be executed. See Appendix B for a complete description of how these LEDs are read, and for a list of all supported POST codes

9.5 Fan Fault LEDs

The server board includes a Fan Fault LED next to each of the six system fan (See Figure 56). The LED has two states: On and Off. The BMC lights a fan fault LED if the associated fan-tach sensor has a lower critical threshold event status asserted. Fan-tach sensors are manual re-arm sensors. Once the lower critical threshold is crossed, the LED remains lit until the sensor is rearmed. These sensors are rearmed at system DC power-on and system reset.

9.6 Memory Fault LEDs

The server board includes a Memory Fault LED for each DIMM slot (See Figure 57). When the BIOS detects a memory fault condition, it sends an IPMI OEM command (*Set Fault Indication*) to the BMC to instruct the BMC to turn on the associated Memory Slot Fault LED. These LEDs are only active when the system is in the 'on' state. The BMC will not activate or change the state of the LEDs unless instructed by the BIOS.

9.7 CPU Fault LEDs

The server board includes a CPU fault LED for each CPU socket. The CPU Fault LED is lit if there is an MSID mismatch error is detected (that is, CPU power rating is incompatible with the board).

10. 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
- Trusted Platform Module (TPM) support
- Intel® Trusted Execution Technology (Intel® TXT)

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 58. BIOS Setup Utility Security Tab

10.1 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. It 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.

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 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.

10.1.1 System Administrator Password Rights

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

- Access the <F2> BIOS Setup Utility
- Has the ability to configure all BIOS setup options in the <F2> BIOS Setup Utility
- Has the ability to 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.

10.1.2 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 any order other than the Boot Order defined in the setup by an Administrator.

10.2 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.

10.3 Trusted Platform Module (TPM) Support

The Trusted Platform Module (TPM) option is a hardware-based security device that addresses the growing concern about boot process integrity and offers better data protection. TPM protects the system startup process by ensuring it is tamper-free before releasing system control to the operating system. A TPM device provides secured storage to store data, such as security keys and passwords. In addition, a TPM device has encryption and hash functions. The server board implements TPM as per *TPM PC Client Specifications revision 1.2*, published by the Trusted Computing Group (TCG).

A TPM device is optionally installed on a high-density 14-pin connector labeled “TPM” on the server board, and is secured from external software attacks and physical theft. A pre-boot environment, such as the BIOS and operating system loader, uses the TPM to collect and store unique measurements from multiple factors within the boot process to create a system fingerprint. This unique fingerprint remains the same unless the pre-boot environment is tampered with. Therefore, it is used to compare to future measurements to verify the integrity of the boot process.

After the system BIOS completes the measurement of its boot process, it hands off control to the operating system loader and, in turn, to the operating system. If the operating system is TPM-enabled, it compares the BIOS TPM measurements to those of previous boots to make sure the system was not tampered with before continuing the operating system boot process. Once the operating system is in operation, it optionally uses TPM to provide additional system and data security (for example, Microsoft Windows 10* supports Bitlocker drive encryption).

10.3.1 TPM Security BIOS

The BIOS TPM support conforms to the TPM PC Client Implementation Specification for Conventional BIOS the TPM Interface Specification, and the Microsoft Windows BitLocker* Requirements. The role of the BIOS for TPM security includes the following:

- Measures and stores the boot process in the TPM microcontroller to allow a TPM-enabled operating system to verify system boot integrity.
- Produces extensible firmware interface (EFI) and legacy interfaces to a TPM-enabled operating system for using TPM.
- Produces Advanced Configuration and Power Interface (ACPI) TPM device and methods to allow a TPM-enabled operating system to send TPM administrative command requests to the BIOS.
- Verifies operator physical presence. Confirms and executes operating system TPM administrative command requests.
- Provides BIOS setup options to change TPM security states and to clear TPM ownership.

For additional details, refer to the *TCG PC Client Specific Implementation Specification*, the *TCG PC Client Specific Physical Presence Interface Specification*, and the *Microsoft Windows* BitLocker* Requirements* documents.

10.3.2 Physical Presence

Administrative operations to the TPM require TPM ownership or physical presence indication by the operator to confirm the execution of administrative operations. The BIOS implements the operator presence indication by verifying the setup Administrator password.

A TPM administrative sequence invoked from the operating system proceeds as follows:

- A user makes a TPM administrative request through the operating system's security software.
- The operating system requests the BIOS to execute the TPM administrative command through TPM ACPI methods and then resets the system.
- The BIOS verifies the physical presence and confirms the command with the operator.

The BIOS executes TPM administrative command(s), inhibits BIOS Setup entry, and boots directly to the operating system which requested the TPM command(s).

10.3.3 TPM Security Setup Options

The BIOS TPM Setup allows the operator to view the current TPM state and to carry out rudimentary TPM administrative operations. Performing TPM administrative options through the BIOS setup requires TPM physical presence verification.

Using the BIOS TPM Setup, the operator can turn ON or OFF TPM functionality and clear the TPM ownership contents. After the requested TPM BIOS Setup operation is carried out, the option reverts to No Operation.

The BIOS TPM Setup also displays the current state of the TPM, whether TPM is enabled or disabled and activated or deactivated. Note that while using TPM, a TPM-enabled operating system or application may change the TPM state independently of the BIOS setup. When an operating system modifies the TPM state, the BIOS Setup displays the updated TPM state.

The BIOS Setup TPM Clear option allows the operator to clear the TPM ownership key and allows the operator to take control of the system with TPM. You use this option to clear security settings for a newly initialized system or to clear a system for which the TPM ownership security key was lost.

10.4 Intel® Trusted Execution Technology

The Intel® Xeon® processor Scalable product family supports Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment. Designed to help protect against software-based attacks, Intel® Trusted Execution Technology integrates new security features and capabilities into the processor, chipset, and other platform components. When used in conjunction with Intel® Virtualization Technology, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications.

This hardware-rooted security provides a general-purpose, safer computing environment capable of running a wide variety of operating systems and applications to increase the confidentiality and integrity of sensitive information without compromising the usability of the platform.

Intel® Trusted Execution Technology requires a computer system with Intel® Virtualization Technology enabled (both VT-x and VT-d), an Intel® Trusted Execution Technology-enabled processor, chipset, and BIOS, Authenticated Code Modules, and an Intel® Trusted Execution Technology compatible measured launched environment (MLE). The MLE could consist of a virtual machine monitor, an OS, or an application. In addition, Intel® Trusted Execution Technology requires the system to include a TPM v1.2, as defined by the *Trusted Computing Group TPM PC Client Specifications, Revision 1.2*.

When available, Intel® Trusted Execution Technology can be enabled or disabled in the processor by a BIOS Setup option.

For general information about Intel® TXT, visit the Intel® Trusted Execution Technology website, <http://www.intel.com/technology/security/>.

11. Reset and Recovery Jumpers

The server board includes several jumper blocks which can be used to configure, protect, or recover specific features of the server board. Figure 59 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.




Figure 59. Reset and Recovery Jumper Block Location

The following sections describe how each jumper block is used.

11.1 BIOS Default Jumper Block

This jumper resets BIOS options, configured using the <F2> BIOS Setup Utility, back to their original default factory settings.

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

1. Power down the server and unplug the power cord(s).
2. Remove the system top cover and move the "BIOS DFLT" jumper from pins 1 - 2 (default) to pins 2 - 3 (Set BIOS Defaults).
3. Wait 5 seconds then move the jumper back to pins 1 - 2.
4. Re-install the system top cover.
5. Re-Install system power cords.
6. During POST, access the <F2> BIOS Setup utility to configure and save desired BIOS options.

Notes:

- The system will automatically power on after AC is applied to the system.
 - The system time and date may need to be reset.
 - After resetting BIOS options using the BIOS Default jumper, the Error Manager Screen in the <F2> BIOS Setup Utility will display two errors:
 - 0012 System RTC date/time not set
 - 5220 BIOS Settings reset to default settings
-

11.2 Password Clear Jumper Block

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 through the <F2> BIOS Setup utility. 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. No method of resetting BIOS configuration settings to default values will affect either the Administrator or User passwords.

1. Power down the server. For safety, unplug the power cord(s).
2. Remove the system top cover.
3. Move the “*Password Clear*” jumper from pins 1 – 2 (default) to pins 2 – 3 (password clear position).
4. Re-install the system top cover and re-attach the power cords.
5. Power up the server and access the <F2> BIOS Setup utility.
6. Verify the password clear operation was successful by viewing the Error Manager screen. Two errors should be logged:
 - 5221 Passwords cleared by jumper
 - 5224 Password clear jumper is set
7. Exit the BIOS Setup utility and power down the server. For safety, remove the AC power cords.
8. Remove the system top cover and move the “*Password Clear*” jumper back to pins 1 - 2 (default).
9. Re-install the system top cover and reattach the AC power cords.
10. Power up the server.
11. **Strongly recommended:** Boot into <F2> BIOS Setup immediately, go to the Security tab and set the Administrator and User passwords if you intend to use BIOS password protection.

11.3 Management Engine (ME) Firmware Force Update Jumper Block

When the ME Firmware Force Update jumper is moved from its default position, the ME is forced to operate in a reduced minimal operating capacity. This jumper should only be used if the ME firmware has gotten corrupted and requires re-installation. Use the procedure below.

Note: System Update files are included in the System Update Packages (SUP) posted to Intel's Download Center website, <http://downloadcenter.intel.com>.

1. Turn off the system.
2. Remove the AC power cords.

Note: If the ME FRC UPD jumper is moved with AC power applied to the system, the ME will not operate properly.

3. Remove the system top cover.
4. Move the “*ME FRC UPD*” Jumper from pins 1 – 2 (default) to pins 2 – 3 (Force Update position).

5. Re-install the system top cover and re-attach the AC power cords.
6. Power on the system.
7. Boot to the EFI shell.
8. Change directories to the folder containing the update files.
9. Update the ME firmware using the following command:

```
iflash32 /u /ni <version#>_ME.cap
```

10. When the update has completed successfully, power off the system.
11. Remove the AC power cords.
12. Remove the system top cover.
13. Move the “ME FRC UPD” jumper back to pins 1-2 (default).
14. Re-attach the AC power cords.
15. Power on the system.

11.4 BMC Force Update Jumper Block

The BMC Force Update jumper is used to put the BMC in Boot Recovery mode for a low-level update. It causes the BMC to abort its normal boot process and stay in the boot loader without executing any Linux code.

This jumper should only be used if the BMC firmware has gotten corrupted and requires re-installation. Do the following:

Note: System Update files are included in the System Update Packages (SUP) posted to Intel's Download Center website, <http://downloadcenter.intel.com>

1. Turn off the system.
2. Remove the AC power cords.

Note: If the BMC FRC UPD jumper is moved with AC power applied to the system, the BMC will not operate properly.

3. Remove the system top cover.
4. Move the “BMC FRC UPD” Jumper from pins 1 - 2 (default) to pins 2 - 3 (Force Update position).
5. Re-install the system top cover and re-attach the AC power cords.
6. Power on the system.
7. Boot to the EFI shell.
8. Change directories to the folder containing the update files.
9. Update the BMC firmware using the following command:

```
FWPIAUPD -u -bin -ni -b -o -pia -if=usb <file name.BIN>
```

10. When the update has successfully completed, power off the system.
11. Remove the AC power cords.
12. Remove the system top cover.
13. Move the “BMC FRC UPD” jumper back to pins 1-2 (default).
14. Re-attach the AC power cords.
15. Power on system.
16. Boot to the EFI shell.
17. Change directories to the folder containing the update files.
18. Re-install the board/system SDR data by running the FRUSDR utility.
19. After the SDRs have been loaded, reboot the server.

11.5 BIOS Recovery Jumper

When the BIOS Recovery jumper block is moved from its default pin position (pins 1–2), the system will boot using a backup BIOS image to the uEFI shell, where a standard BIOS update can be performed. See the BIOS update instructions that are included with System Update Packages (SUP) downloaded from Intel's download center website. This jumper is used when the system BIOS has become corrupted and is non-functional, requiring a new BIOS image to be loaded on to the server board.

Note: The BIOS Recovery jumper is ONLY used to re-install a BIOS image in the event the BIOS has become corrupted. This jumper is NOT used when the BIOS is operating normally and you need to update the BIOS from one version to another.

The following procedure should be followed.

Note: System Update Packages (SUP) can be downloaded from Intel's download center website, <http://downloadcenter.intel.com>

1. Turn off the system.
2. For safety, remove the AC power cords.
3. Remove the system top cover.
4. Move the “*BIOS Recovery*” jumper from pins 1 – 2 (default) to pins 2 – 3 (BIOS Recovery position).
5. Re-install the system top cover and re-attach the AC power cords.
6. Power on the system.
7. The system will automatically boot to the EFI shell. Update the BIOS using the standard BIOS update instructions provided with the system update package.
8. After the BIOS update has successfully completed, power off the system. For safety, remove the AC power cords from the system.
9. Remove the system top cover.
10. Move the BIOS Recovery jumper back to pins 1 – 2 (default).
11. Re-install the system top cover and re-attach the AC power cords.
12. Power on the system and access the <F2> BIOS Setup utility.
13. Configure desired BIOS settings.
14. Hit the <F10> key to save and exit the utility.

12. Platform Management

Platform management is supported by several integrated hardware and software components that work together to support the following:

- Controlling system functions like the power system, ACPI, system reset control, system initialization, front panel interface, system event log
- Monitoring various board and system sensors, regulating platform thermals and performance in order to maintain (when possible) server functionality in the event of component failure and/or environmentally stressed conditions
- Monitoring and reporting system health
- Providing an interface for Server Management Software applications

This chapter provides a high level overview of the platform management features and functionality implemented on the server board.

The *Intel® Server System BMC Firmware External Product Specification (EPS)* and the *Intel® Server System BIOS External Product Specification (EPS)* for Intel® Server Products based on the Intel® Xeon® processor E5-2600 v5 product families should be referenced for more in-depth and design level platform management information.

12.1 Management Feature Set Overview

The following sections outline features that the integrated BMC firmware can support. Support and utilization for some features is dependent on the server platform in which the server board is integrated and any additional system level components and options that may be installed.

12.1.1 IPMI 2.0 Features Overview

- Baseboard management controller (BMC)
- IPMI Watchdog timer
- Messaging support, including command bridging and user/session support
- Chassis device functionality, including power/reset control and BIOS boot flags support
- Event receiver device: The BMC receives and processes events from other platform subsystems.
- Field Replaceable Unit (FRU) inventory device functionality: The BMC supports access to system FRU devices using IPMI FRU commands.
- System Event Log (SEL) device functionality: The BMC supports and provides access to a SEL including SEL Severity Tracking and the Extended SEL.
- Sensor Data Record (SDR) repository device functionality: The BMC supports storage and access of system SDRs.
- Sensor device and sensor scanning/monitoring: The BMC provides IPMI management of sensors. It polls sensors to monitor and report system health.
- IPMI interfaces
 - Host interfaces include system management software (SMS) with receive message queue support, and server management mode (SMM)
 - IPMB interface
 - LAN interface that supports the IPMI-over-LAN protocol (RMCP, RMCP+)
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS.
- BMC self-test: The BMC performs initialization and runtime self-tests and makes results available to external entities.

See also the Intelligent Platform Management Interface Specification Second Generation v2.0.

12.1.2 Non-IPMI Features Overview

The BMC supports the following non-IPMI features:

- In-circuit BMC firmware update
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Chassis intrusion detection (dependent on platform support)
- Fan speed control with SDR. Fan redundancy monitoring and support
- Enhancements to fan speed control
- Power supply redundancy monitoring and support
- Hot-swap fan support
- Acoustic management: Support for multiple fan profiles
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval
- 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.
- Power state retention.
- Power fault analysis.
- Intel® Light-Guided Diagnostics.
- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support.
- Email alerting.
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server.
 - Human-readable SEL
 - Additional system configurability
 - Additional system monitoring capability
 - Enhanced online help
- Integrated KVM.
- Enhancements to KVM redirection.
 - Support for higher resolution
- Integrated Remote Media Redirection.
- Lightweight Directory Access Protocol (LDAP) support.
- Intel® Intelligent Power Node Manager support.
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
 - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.5 compliance (product-specific).
- Management support for PMBus* rev1.2 compliant power supplies.
- BMC Data Repository (Managed Data Region Feature).
- System Airflow Monitoring.

- Exit Air Temperature Monitoring.
- Ethernet Controller Thermal Monitoring.
- Global Aggregate Temperature Margin Sensor.
- Memory Thermal Management.
- Power Supply Fan Sensors.
- Energy Star Server Support.
- Smart Ride Through (SmaRT) / Closed Loop System Throttling (CLST).
- Power Supply Cold Redundancy.
- Power Supply FW Update.
- Power Supply Compatibility Check.
- BMC FW reliability enhancements:
 - Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC
 - BMC System Management Health Monitoring

12.2 Platform Management Features and Functions

12.2.1 Power Subsystem

The server board supports several power control sources that can initiate power-up or power-down activity, as listed in Table 42.

Table 42. Power Control Sources

Source	External Signal Name or Internal Subsystem	Capabilities
Power button	Front panel power button	Turns power on or off
BMC watchdog timer	Internal BMC timer	Turns power off, or power cycle
BMC chassis control Commands	Routed through command processor	Turns power on or off, or power cycle
Power state retention	Implemented by means of BMC internal logic	Turns power on when AC power returns
Chipset	Sleep S4/S5 signal (same as <i>POWER_ON</i>)	Turns power on or off
CPU Thermal	Processor Thermtrip	Turns power off
PCH Thermal	PCH Thermtrip	Turns power off
WOL (Wake On LAN)	LAN	Turns power on

12.2.2 Advanced Configuration and Power Interface (ACPI)

The server board has support for the ACPI states described in Table 43.

Table 43. ACPI Power States

State	Supported	Description
S0	Yes	Working. <ul style="list-style-type: none"> ▪ The front panel power LED is on (not controlled by the BMC). ▪ The fans spin at the normal speed, as determined by sensor inputs. ▪ Front panel buttons work normally.
S1	No	Not supported
S2	No	Not supported
S3	No	Supported only on Workstation platforms. See appropriate Platform Specific Information for more information.
S4	No	Not supported
S5	Yes	Soft off <ul style="list-style-type: none"> ▪ The front panel buttons are not locked. ▪ The fans are stopped. ▪ The power-up process goes through the normal boot process. ▪ The power, reset, front panel NMI, and ID buttons are unlocked.

During system initialization, both the BIOS and the BMC initialize the items described in the following sections.

12.2.2.1 Processor Tcontrol Setting

Processors used with this chipset implement a feature called Tcontrol, which provides a processor-specific value that can be used to adjust the fan-control behavior to achieve optimum cooling and acoustics. The BMC reads these from the CPU through PECI Proxy mechanism provided by Manageability Engine (ME). The BMC uses these values as part of the fan-speed-control algorithm.

12.2.2.2 Fault Resilient Booting (FRB)

Fault resilient booting (FRB) is a set of BIOS and BMC algorithms and hardware support that allow a multiprocessor system to boot even if the bootstrap processor (BSP) fails. Only FRB2 is supported using watchdog timer commands.

FRB2 refers to the FRB algorithm that detects system failures during POST. The BIOS uses the BMC watchdog timer to back up its operation during POST. The BIOS configures the watchdog timer to indicate that the BIOS is using the timer for the FRB2 phase of the boot operation.

After the BIOS has identified and saved the BSP information, it sets the FRB2 timer use bit and loads the watchdog timer with the new timeout interval.

If the watchdog timer expires while the watchdog use bit is set to FRB2, the BMC (if so configured) logs a watchdog expiration event showing the FRB2 timeout in the event data bytes. The BMC then hard resets the system, assuming the BIOS-selected reset as the watchdog timeout action.

The BIOS is responsible for disabling the FRB2 timeout before initiating the option ROM scan and before displaying a request for a boot password. If the processor fails and causes an FRB2 timeout, the BMC resets the system.

The BIOS gets the watchdog expiration status from the BMC. If the status shows an expired FRB2 timer, the BIOS enters the failure in the system event log (SEL). In the OEM bytes entry in the SEL, the last POST code generated during the previous boot attempt is written. FRB2 failure is not reflected in the processor status sensor value.

The FRB2 failure does not affect the front panel LEDs.

12.2.2.3 Post Code Display

The BMC, upon receiving standby power, initializes internal hardware to monitor port 80h (POST code) writes. Data written to port 80h is output to the system POST LEDs.

The BMC will deactivate POST LEDs after POST completes.

12.2.3 Watchdog Timer

The BMC implements a fully IPMI 2.0 compatible watchdog timer. For details, see the *Intelligent Platform Management Interface Specification Second Generation v2.0*. The NMI/diagnostic interrupt for an IPMI 2.0 watchdog timer is associated with an NMI. A watchdog pre-timeout SMI or equivalent signal assertion is not supported.

12.2.4 System Event Log (SEL)

The BMC implements the system event log as specified in the *Intelligent Platform Management Interface Specification, Version 2.0*. The SEL is accessible regardless of the system power state through the BMC's in-band and out-of-band interfaces.

The BMC allocates 95,231 bytes (approx. 93 KB) of non-volatile storage space to store system events. The SEL timestamps may not be in order. Up to 3,639 SEL records can be stored at a time. Because the SEL is circular, any command that results in an overflow of the SEL beyond the allocated space will overwrite the oldest entries in the SEL, while setting the overflow flag.

12.3 Sensor Monitoring

The BMC monitors system hardware and reports system health. The information gathered from physical sensors is translated into IPMI sensors as part of the "IPMI Sensor Model." The BMC also reports various system state changes by maintaining virtual sensors that are not specifically tied to physical hardware. This section describes general aspects of BMC sensor management as well as describing how specific sensor types are modeled. Unless otherwise specified, the term "sensor" refers to the IPMI sensor-model definition of a sensor.

- Sensor Scanning
- BIOS Event-Only Sensors
- Margin Sensors
- IPMI Watchdog Sensor
- BMC Watchdog Sensor
- BMC System Management Health Monitoring
- VR Watchdog Timer
- System Airflow Monitoring Sensors - valid for Intel® server chassis only
- Fan monitoring sensors
- Thermal monitoring sensors
- Voltage monitoring sensors
- CATERR sensor

- LAN Leash Event Monitoring
- CMOS Battery Monitoring
- NMI (Diagnostic Interrupt) Sensor

12.3.1 Sensor Rearm Behavior

12.3.1.1 Manual versus Automatic Re-arm Sensors

Sensors can be re-armed either manually or automatically. An automatic re-arm sensor will "re-arm" (clear) the assertion event state for a threshold or offset if that threshold or offset is de-asserted after having been asserted. This allows a subsequent assertion of the threshold or an offset to generate a new event and associated side-effect. An example side-effect would be boosting fans due to an upper critical threshold crossing of a temperature sensor. The event state and the input state (value) of the sensor track each other. Most sensors are of the auto-rearm type.

A manual re-arm sensor does not clear the assertion state even when the threshold or offset becomes de-asserted. In this case, the event state and the input state (value) of the sensor do not track each other. The event assertion state is "sticky". The following methods can be used to re-arm a sensor:

- Automatic re-arm – Applies only to sensors that are designated as "auto-rearm."
- IPMI command – Re-arm Sensor Event.
- BMC internal method – The BMC may re-arm certain sensors due to a trigger condition. For example, some sensors may be re-armed due to a system reset. A BMC reset will re-arm all sensors.
- System reset or DC power cycle – Will re-arm all system fan sensors.

12.3.2 Thermal Monitoring

The BMC provides monitoring of component and board temperature sensing devices. This monitoring capability is instantiated in the form of IPMI analog/threshold or discrete sensors, depending on the nature of the measurement.

For analog/threshold sensors, with the exception of Processor Temperature sensors, critical and non-critical thresholds (upper and lower) are set through SDRs and event generation enabled for both assertion and de-assertion events.

For discrete sensors, both assertion and de-assertion event generation are enabled.

Mandatory monitoring of platform thermal sensors includes:

- Inlet temperature (physical sensor is typically on system front panel or HDD backplane)
- Board ambient thermal sensors
- Processor temperature
- Memory (DIMM) temperature
- CPU VRD Hot monitoring
- Power supply Inlet temperature (only supported for PMBus*-compliant PSUs)

Additionally, the BMC FW may create "virtual" sensors that are based on a combination or aggregation of multiple physical thermal sensors and the application of a mathematical formula to thermal or power sensor readings.

12.3.3 Standard Fan Management

The BMC controls and monitors the system fans. Each fan is associated with a fan speed sensor that detects fan failure and may also be associated with a fan presence sensor for hot-swap support. For redundant fan configurations, the fan failure and presence status determines the fan redundancy sensor state.

The system fans are divided into fan domains, each of which has a separate fan speed control signal and a separate configurable fan control policy. A fan domain can have a set of temperature and fan sensors associated with it. These are used to determine the current fan domain state.

A fan domain has three states:

- The sleep and boost states have fixed (but configurable through OEM SDRs) fan speeds associated with them.
- The nominal state has a variable speed determined by the fan domain policy. An OEM SDR record is used to configure the fan domain policy.

The fan domain state is controlled by several factors. They are listed below in order of precedence, high to low:

- Boost
 - Associated fan is in a critical state or missing. The SDR describes which fan domains are boosted in response to a fan failure or removal in each domain. If a fan is removed when the system is in “Fans-off” mode, it will not be detected and there will not be any fan boost till the system comes out of “Fans-off” mode.
 - Any associated temperature sensor is in a critical state. The SDR describes which temperature-threshold violations cause fan boost for each fan domain.
 - The BMC is in firmware update mode, or the operational firmware is corrupted.
 - If any of the above conditions apply, the fans are set to a fixed boost state speed.
- Nominal
 - A fan domain’s nominal fan speed can be configured as static (fixed value) or controlled by the state of one or more associated temperature sensors.

12.3.3.1 Hot-Swappable Fans

Hot-swappable fans are supported. These fans can be removed and replaced while the system is powered on and operating. The BMC implements fan presence sensors for each hot-swappable fan.

When a fan is not present, the associated fan speed sensor is put into the *reading/unavailable* state, and any associated fan domains are put into the boost state. The fans may already be boosted due to a previous fan failure or fan removal.

When a removed fan is inserted, the associated fan speed sensor is re-armed. If there are no other critical conditions causing a fan boost condition, the fan speed returns to the nominal state. Power cycling or resetting the system re-arms the fan speed sensors and clears fan failure conditions. If the failure condition is still present, the boost state returns once the sensor has re-initialized and the threshold violation is detected again.

12.3.3.2 Fan Redundancy Detection

The BMC supports redundant fan monitoring and implements a fan redundancy sensor. A fan redundancy sensor generates events when its associated set of fans transitions between redundant and non-redundant

states, as determined by the number and health of the fans. The definition of fan redundancy is configuration dependent. The BMC allows redundancy to be configured on a per fan redundancy sensor basis through OEM SDR records.

A fan failure or removal of hot-swap fans up to the number of redundant fans specified in the SDR in a fan configuration is a non-critical failure and is reflected in the front panel status. A fan failure or removal that exceeds the number of redundant fans is a non-fatal, insufficient-resources condition and is reflected in the front panel status as a non-fatal error.

Redundancy is checked only when the system is in the DC-on state. Fan redundancy changes that occur when the system is DC-off or when AC is removed will not be logged until the system is turned on.

12.3.3.3 Fan Domains

System fan speeds are controlled through pulse width modulation (PWM) signals, which are driven separately for each domain by integrated PWM hardware. Fan speed is changed by adjusting the duty cycle, which is the percentage of time the signal is driven high in each pulse.

The BMC controls the average duty cycle of each PWM signal through direct manipulation of the integrated PWM control registers.

The same device may drive multiple PWM signals.

12.3.3.4 Thermal and Acoustic Management

This feature refers to enhanced fan management to keep the system optimally cooled while reducing the amount of noise generated by the system fans. Aggressive acoustics standards might require a trade-off between fan speed and system performance parameters that contribute to the cooling requirements, primarily memory bandwidth. The BIOS, BMC and SDRs work together to provide control over how this trade-off is determined.

This capability requires the BMC to access temperature sensors on the individual memory DIMMs. Additionally, closed-loop thermal throttling is only supported with DIMMs with temperature sensors.

12.3.3.5 Thermal Sensor Input to Fan Speed Control

The BMC uses various IPMI sensors as an input to the fan speed control. Some of the sensors are IPMI models of actual physical sensors whereas some are “virtual” sensors whose values are derived from physical sensors using calculations and/or tabular information.

The following IPMI thermal sensors are used as the input to the fan speed control:

- Baseboard temperature sensors
- CPU DTS-Spec margin sensors
- DIMM thermal margin sensors
- Exit air temperature sensor
- PCH Temperature sensor
- Global aggregate thermal margin sensors
- SSB (Intel® C620 Series Chipset) temperature sensor
- On-board Ethernet controller temperature sensors (support for this is specific to the Ethernet controller being used)
- On-board SAS controller temperature sensors (when available)
- CPU VR Temperature sensor

- DIMM VR Temperature sensor
- BMC Temperature sensor
- DIMM VRM Temperature sensor

Figure 60 shows a high-level illustration of the fan speed control structure that determines fan speed.

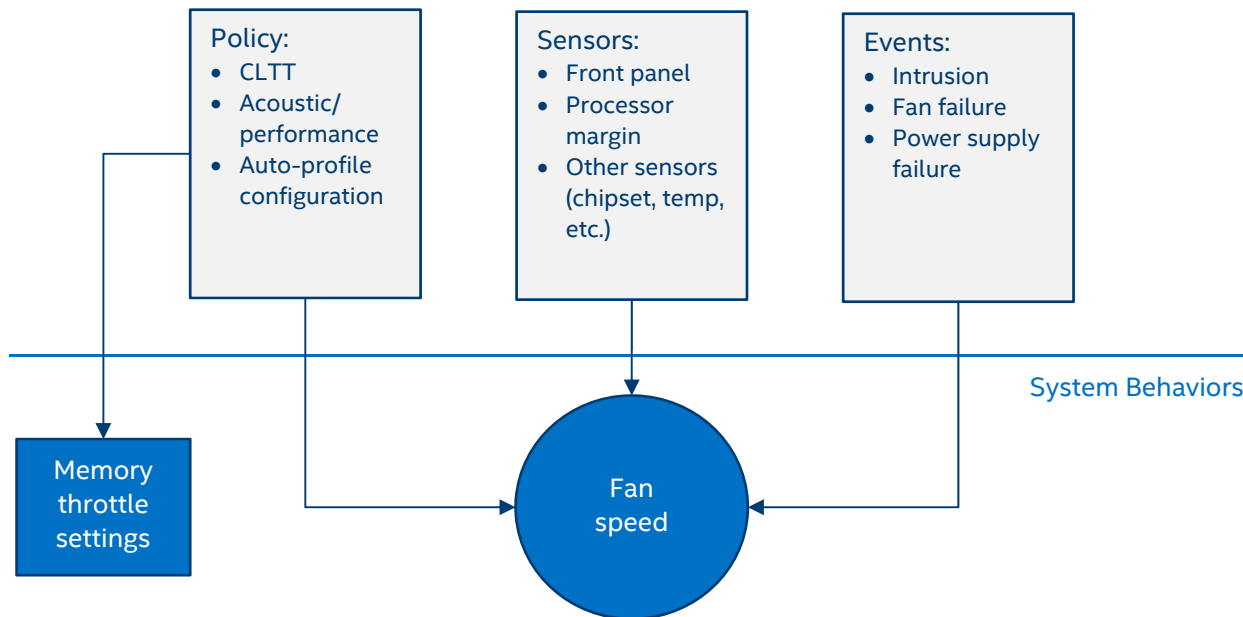


Figure 60. High-level Fan Speed Control Process

12.3.3.6 Fan Boosting due to Fan Failures

Each fan failure shall be able to define a unique response from all other fan domains. An OEM SDR table defines the response of each fan domain based on a failure of any fan, including both system and power supply fans (for PMBus*-compliant power supplies only). This means that if a system has six fans, there will be six different fan fail reactions.

12.3.4 Memory Thermal Management

The system memory is the most complex subsystem to manage thermally, as it requires substantial interactions between the BMC, BIOS, and the embedded memory controller hardware. This section provides an overview of this management capability from a BMC perspective.

12.3.4.1 Memory Thermal Throttling

The system supports thermal management through closed loop throttling (CLTT) only. Throttling levels are changed dynamically to cap throttling based on memory and system thermal conditions as determined by the system and DIMM power and thermal parameters. The BMC fan speed control functionality is related to the memory throttling mechanism used.

The following terminology is used for the various memory throttling options:

- **Static Closed Loop Thermal Throttling (Static-CLTT):** CLTT control registers are configured by BIOS MRC during POST. The memory throttling is run as a closed-loop system with the DIMM temperature sensors as the control input. Otherwise, the system does not change any of the throttling control registers in the embedded memory controller during runtime.

- **Dynamic Closed Loop Thermal Throttling (Dynamic-CLTT):** CLTT control registers are configured by BIOS MRC during POST. The memory throttling is run as a closed-loop system with the DIMM temperature sensors as the control input. Adjustments are made to the throttling during runtime based on changes in system cooling (fan speed).

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v5 product family support a type of CLTT, called a Hybrid CLTT, for which the integrated Memory controller estimates the DRAM temperature in between actual reads of the TSODs. Hybrid CLTT shall be used on all Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v5 product family that have DIMMs with thermal sensors. Therefore, the terms Dynamic-CLTT and Static-CLTT are really referring to this “hybrid” mode. Note that if the IMC's polling of the TSODs is interrupted, the temperature readings that the BMC gets from the IMC shall be these estimated values.

12.3.4.2 Dynamic (Hybrid) CLTT

The system will support dynamic (memory) CLTT for which the BMC FW dynamically modifies thermal offset registers in the IMC during runtime based on changes in system cooling (fan speed). For static CLTT, a fixed offset value is applied to the TSOD reading to get the die temperature; however this does not provide as accurate results as when the offset takes into account the current airflow over the DIMM, as is done with dynamic CLTT.

In order to support this feature, the BMC FW will derive the air velocity for each fan domain based on the PWM value being driven for the domain. Since this relationship is dependent on the chassis configuration, a method must be used which supports this dependency (for example, through OEM SDR) that establishes a lookup table providing this relationship.

The BIOS will have an embedded lookup table that provides thermal offset values for each DIMM type, altitude setting, and air velocity range (three ranges of air velocity are supported). During system boot the BIOS will provide three offset values (corresponding to the three air velocity ranges) to the BMC for each enabled DIMM. Using this data the BMC FW constructs a table that maps the offset value corresponding to a given air velocity range for each DIMM. During runtime the BMC applies an averaging algorithm to determine the target offset value corresponding to the current air velocity and then the BMC writes this new offset value into the IMC thermal offset register for the DIMM.

12.3.5 Power Management Bus (PMBus*)

The Power Management Bus (PMBus*) is an open standard protocol that is built upon the SMBus* 2.0 transport. It defines a means of communicating with power conversion and other devices using SMBus*-based commands. A system must have PMBus*-compliant power supplies installed in order for the BMC or ME to monitor them for status and/or power metering purposes.

For more information on PMBus*, see the System Management Interface Forum website, <http://www.powersig.org/>.

12.3.6 Component Fault LED Control

Several sets of component fault LEDs are supported on the server board. See the figures for Intel® Light Guided Diagnostics. Some LEDs are owned by the BMC and some by the BIOS.

The BMC owns control of the following FRU/fault LEDs:

- **DIMM fault LEDs** – The BMC owns the hardware control for these LEDs. The LEDs reflect the state of BIOS-owned event-only sensors. When the BIOS detects a DIMM fault condition, it sends an IPMI OEM command (*Set Fault Indication*) to the BMC to instruct the BMC to turn on the associated DIMM Fault LED. These LEDs are only active when the system is in the “on” state. The BMC will not activate or change the state of the LEDs unless instructed by the BIOS.
- **Hard Disk Drive Status LEDs** – The HSBP PSoC* of supported Intel and third-party chassis, owns the HW control for these LEDs if present, and detection of the fault/status conditions that the LEDs reflect.
- **CPU Fault LEDs** – The server board provides a fault LED for each processor socket and are controlled by the BMC. An LED is lit if there is an MSID mismatch where the CPU power rating is incompatible with the board.

Table 44. Component Fault LEDs

Component	Owner	Color	State	Description
DIMM Fault LED	BMC	Amber	Solid On	Memory failure – detected by the BIOS
		Amber	Off	DIMM working correctly
HDD Fault LED	HSBP PSoC*	Amber	On	HDD Fault
		Amber	Blink	Predictive failure, rebuild, identify
		Amber	Off	OK (no errors)
CPU Fault LEDs	BMC	Amber	Off	OK (no errors)
		Amber	On	MSID mismatch

Appendix A – Integration and Usage Tips

- When adding or removing components or peripherals from the server board, power cords must be disconnected from the server. With power applied to the server, standby voltages are still present even though the server board is powered off.
- This server board supports the Intel® Xeon® processor scalable family with a Thermal Design Power (TDP) of up to and including 205 Watts. 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 2U 3-slot riser card and Riser Card Slots #2 and #3 on the server board can only be used in dual processor configurations.
- **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.
- 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, as shown in Figure 61.

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 and the lower nibble bits are represented by Green Diagnostics. 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.



Note: Diag LEDs are best read and decoded when viewing the LEDs from the back of the system

Figure 61. On-board POST Diagnostic LEDs

In the following example, the BIOS sends a value of AC to the diagnostic LED decoder. The LEDs are decoded as shown in Table 45, where the upper nibble bits represented by the amber LEDs equal 1010_b or A_h and the lower nibble bits represented by the green LEDs equal 1100_b or C_h . The two are concatenated as AC_h .

Table 45. POST Progress Code LED Example

	Upper Nibble AMBER LEDs				Lower Nibble GREEN LEDs			
	MSB							LSB
	1	0	1	0	1	1	0	0
Binary Value	ON	OFF	ON	OFF	ON	ON	OFF	OFF
LED State	ON	OFF	ON	OFF	ON	ON	OFF	OFF
Hex Value	8h	4h	2h	1h	8h	4h	2h	1h
Hex Result	Ah				Ch			

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 to the Diagnostic LEDs that show the execution point in the MRC operational path at each step.

Table 46. 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	
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. The following table 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 46.

Table 47. MRC Fatal Error Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Read 1st)				Lower Nibble (Read 2nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
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 48 provides a list of all POST progress codes.

Table 48. POST Progress Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Read 1 st)				Lower Nibble (Read 2 nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
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
AEnh	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

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Read 1 st)				Lower Nibble (Read 2 nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
19h	0	0	0	1	1	0	1	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	0	1	1	0	0	1	1	CPU PEIM (Cache Init)
4Fh	0	1	0	0	1	1	1	1	Dxe IPL started
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	0	1	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	0	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

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Read 1 st)				Lower Nibble (Read 2 nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
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
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

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Read 1 st)				Lower Nibble (Read 2 nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
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
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.

The following table 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.

The following table identifies POST Error Messages and Handling that 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 49. POST Error Messages and Handling

Error Code	Error Message	Action message	Response
0012	System RTC date/time not set		Major
0048	Password check failed	Please put right password.	Major
0140	PCI component encountered a PERR error		Major
0141	PCI resource conflict		Major
0146	PCI out of resources error	Please enable Memory Mapped I/O above 4 GB item at SETUP to use 64bit MMIO.	Major
0191	Processor core/thread count mismatch detected	Please use identical CPU type.	Fatal
0192	Processor cache size mismatch detected	Please use identical CPU type.	Fatal
0194	Processor family mismatch detected	Please use identical CPU type.	Fatal
0195	Processor Intel(R) UPI link frequencies unable to synchronize		Fatal
0196	Processor model mismatch detected	Please use identical CPU type.	Fatal
0197	Processor frequencies unable to synchronize	Please use identical CPU type.	Fatal
5220	BIOS Settings reset to default settings		Major
5221	Passwords cleared by jumper		Major
5224	Password clear jumper is Set	Recommend to remind user to install BIOS password as BIOS admin password is the master keys for several BIOS security features.	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	Baseboard 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	Baseboard management controller failed to respond		Major
84F3	Baseboard management controller in update mode		Major
84F4	Sensor data record empty	Please update right SDR.	Major
84FF	System event log full	Please clear SEL through EWS or SELVIEW utility.	Minor
8500	Memory component could not be configured in the selected RAS mode		Major
8501	DIMM Population Error	Please plug DIMM at right population.	Major
8520	CPU1_DIMM_A1 failed test/initialization	Please remove the disabled DIMM.	Major
8521	CPU1_DIMM_A2 failed test/initialization	Please remove the disabled DIMM.	Major
8523	CPU1_DIMM_B1 failed test/initialization	Please remove the disabled DIMM.	Major
8524	CPU1_DIMM_B2 failed test/initialization	Please remove the disabled DIMM.	Major
8526	CPU1_DIMM_C1 failed test/initialization	Please remove the disabled DIMM.	Major
8527	CPU1_DIMM_C2 failed test/initialization	Please remove the disabled DIMM.	Major
8529	CPU1_DIMM_D1 failed test/initialization	Please remove the disabled DIMM.	Major
852A	CPU1_DIMM_D2 failed test/initialization	Please remove the disabled DIMM.	Major
852C	CPU1_DIMM_E1 failed test/initialization	Please remove the disabled DIMM.	Major
852D	CPU1_DIMM_E2 failed test/initialization	Please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
852F	CPU1_DIMM_F1 failed test/initialization	Please remove the disabled DIMM.	Major
8530	CPU1_DIMM_F2 failed test/initialization	Please remove the disabled DIMM.	Major
8533	CPU1_DIMM_G2 failed test/initialization	Please remove the disabled DIMM.	Major
8538	CPU2_DIMM_A1 failed test/initialization	Please remove the disabled DIMM.	Major
8539	CPU2_DIMM_A2 failed test/initialization	Please remove the disabled DIMM.	Major
853B	CPU2_DIMM_B1 failed test/initialization	Please remove the disabled DIMM.	Major
853C	CPU2_DIMM_B2 failed test/initialization	Please remove the disabled DIMM.	Major
853E	CPU2_DIMM_C1 failed test/initialization	Please remove the disabled DIMM.	Major
853F (Go to 85C0)	CPU2_DIMM_C2 failed test/initialization	Please remove the disabled DIMM.	Major
8540	CPU1_DIMM_A1 disabled	Please remove the disabled DIMM.	Major
8541	CPU1_DIMM_A2 disabled	Please remove the disabled DIMM.	Major
8543	CPU1_DIMM_B1 disabled	Please remove the disabled DIMM.	Major
8544	CPU1_DIMM_B2 disabled	Please remove the disabled DIMM.	Major
8546	CPU1_DIMM_C1 disabled	Please remove the disabled DIMM.	Major
8547	CPU1_DIMM_C2 disabled	Please remove the disabled DIMM.	Major
8549	CPU1_DIMM_D1 disabled	Please remove the disabled DIMM.	Major
854A	CPU1_DIMM_D2 disabled	Please remove the disabled DIMM.	Major
854C	CPU1_DIMM_E1 disabled	Please remove the disabled DIMM.	Major
854D	CPU1_DIMM_E2 disabled	Please remove the disabled DIMM.	Major
854F	CPU1DIMM_F1 disabled	Please remove the disabled DIMM.	Major
8550	CPU1DIMM_F2 disabled	Please remove the disabled DIMM.	Major
8558	CPU2_DIMM_A1 disabled	Please remove the disabled DIMM.	Major
8559	CPU2_DIMM_A2 disabled	Please remove the disabled DIMM.	Major
855B	CPU2_DIMM_B1 disabled	Please remove the disabled DIMM.	Major
855C	CPU2_DIMM_B2 disabled	Please remove the disabled DIMM.	Major
855E	CPU2_DIMM_C1 disabled	Please remove the disabled DIMM.	Major
855F (Go to 85D0)	CPU2_DIMM_C2 disabled	Please remove the disabled DIMM.	Major
8560	CPU1_DIMM_A1 encountered a Serial Presence Detection (SPD) failure		Major
8561	CPU1_DIMM_A2 encountered a Serial Presence Detection (SPD) failure		Major
8563	CPU1_DIMM_B1 encountered a Serial Presence Detection (SPD) failure		Major
8564	CPU1_DIMM_B2 encountered a Serial Presence Detection (SPD) failure		Major
8566	CPU1_DIMM_C1 encountered a Serial Presence Detection (SPD) failure		Major
8567	CPU1_DIMM_C2 encountered a Serial Presence Detection (SPD) failure		Major
8569	CPU1_DIMM_D1 encountered a Serial Presence Detection (SPD) failure		Major
856A	CPU1_DIMM_D2 encountered a Serial Presence Detection (SPD) failure		Major
856C	CPU1_DIMM_E1 encountered a Serial Presence Detection (SPD) failure		Major
856D	CPU1_DIMM_E2 encountered a Serial Presence Detection (SPD) failure		Major
856F	CPU1_DIMM_F1 encountered a Serial Presence Detection (SPD) failure		Major
8570	CPU1_DIMM_F2 encountered a Serial Presence Detection (SPD) failure		Major
8578	CPU2_DIMM_A1 encountered a Serial Presence Detection (SPD) failure		Major
8579	CPU2_DIMM_A2 encountered a Serial Presence Detection (SPD) failure		Major

Error Code	Error Message	Action message	Response
857B	CPU2_DIMM_B1 encountered a Serial Presence Detection (SPD) failure		Major
857C	CPU2_DIMM_B2 encountered a Serial Presence Detection (SPD) failure		Major
857E	CPU2_DIMM_C1 encountered a Serial Presence Detection (SPD) failure		Major
857F (Go to 85E0)	CPU2_DIMM_C2 encountered a Serial Presence Detection (SPD) failure		Major
85C1	CPU2_DIMM_D1 failed test/initialization	Please remove the disabled DIMM.	Major
85C2	CPU2_DIMM_D2 failed test/initialization	Please remove the disabled DIMM.	Major
85C4	CPU2_DIMM_E1 failed test/initialization	Please remove the disabled DIMM.	Major
85C5	CPU2_DIMM_E2 failed test/initialization	Please remove the disabled DIMM.	Major
85C7	CPU2_DIMM_F1 failed test/initialization	Please remove the disabled DIMM.	Major
85C8	CPU2_DIMM_F2 failed test/initialization	Please remove the disabled DIMM.	Major
85D1	CPU2_DIMM_D1 disabled	Please remove the disabled DIMM.	Major
85D2	CPU2_DIMM_D2 disabled	Please remove the disabled DIMM.	Major
85D4	CPU2_DIMM_E1 disabled	Please remove the disabled DIMM.	Major
85D5	CPU2_DIMM_E2 disabled	Please remove the disabled DIMM.	Major
85D7	CPU2_DIMM_F1 disabled	Please remove the disabled DIMM.	Major
85D8	CPU2_DIMM_F2 disabled	Please remove the disabled DIMM.	Major
85E0	CPU2_DIMM_C3 encountered a Serial Presence Detection (SPD) failure		Major
85E1	CPU2_DIMM_D1 encountered a Serial Presence Detection (SPD) failure		Major
85E2	CPU2_DIMM_D2 encountered a Serial Presence Detection (SPD) failure		Major
85E4	CPU2_DIMM_E1 encountered a Serial Presence Detection (SPD) failure		Major
85E5	CPU2_DIMM_E2 encountered a Serial Presence Detection (SPD) failure		Major
85E7	CPU2_DIMM_F1 encountered a Serial Presence Detection (SPD) failure		Major
85E8	CPU2_DIMM_F2 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	Please disable OpRom at SETUP to save runtime memory.	Minor

POST Error Beep Codes

Table 50 lists the POST error beep codes. Prior to system video initialization, the BIOS uses these beep codes to inform users on error conditions. The beep code is followed by a user-visible code on the POST Progress LEDs.

Table 50. POST Error Beep Codes

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

The Integrated BMC may generate beep codes upon detection of failure conditions. Beep codes are sounded each time the problem is discovered, such as on each power-up attempt, but are not sounded continuously. Codes that are common across all Intel server boards and systems that use same generation chipset are listed in Table 51. Each digit in the code is represented by a sequence of beeps whose count is equal to the digit.

Table 51. Integrated BMC Beep Codes

Code	Reason for Beep	Associated Sensors
1-5-1-2	VR Watchdog Timer sensor assertion	VR Watchdog Timer
1-5-1-4	The system does not power on or unexpectedly power off and a power supply unit (PSU) is present that is an incompatible model with one or more other PSUs in the system	PS Status
1-5-2-1	No CPUs installed or first CPU socket is empty	CPU Missing Sensor
1-5-2-2	CPU CAT Error (IERR) assertion	CPU ERR2 Timeout Sensor
1-5-2-3	CPU ERR2 timeout assertion	CPU ERR2 Timeout Sensor
1-5-2-4	CPU lcc max Mismatch	CPU lcc max Mismatch Sensor
1-5-2-5	CPU population error	CPU 0 Status Sensor
1-5-4-2	Power fault: DC power is unexpectedly lost (power good dropout).	Power unit – power unit failure offset
1-5-4-4	Power control fault (power good assertion timeout).	Power unit – soft power control failure offset

Appendix D – Statement of Volatile Memory Components

Table 52 identifies the volatile and non-volatile memory components of the Intel® Server Board S2600WF (Intel Product Codes S2600WFT, S2600WFQ and s2600WF0) server board assembly.

Table 52. Volatile and Non-volatile Components

Component Type	Size	Board Location	User Data	Name
Non-volatile	64MB	U3D1	No(BIOS)	BIOS Flash
Non-volatile	64MB	U1D1	NO(FW)	BMC Flash
Non-Volatile	4MBIT	U5M1	No	10 GB NIC EEPROM (S2600WF)
Non-volatile	N/A	U1E3	No	CPLD
Volatile	512MB	U1D2	No	BMC FW SDRAM
None-volatile	8GB	U8N1	No	BMC eMMC

Note: Table 52 does not identify volatile and non-volatile memory components for devices which may be installed onto or may be used with the server board. These may include: system boards used inside a server system, processors, memory, storage devices, or add-in cards.

Table 52 provides 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 Table 52.

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.

Appendix E – Supported Intel Server Systems

The Intel® Server Board S2600WF product family is designed to be integrated into high density 1U and 2U rack mount servers chassis. Intel Server Systems that include this server board family include the **Intel® Server System R1000WF product family** and the **Intel® Server System R2000WF product family**. The sections below provide a high level overview of the features associated with each. For additional product information, refer to the Technical Product Specification, Integration and Service Guide, and Product Family Configuration Guide, and other marketing material available for each of these server product families. These documents can be downloaded from the Intel website.

Intel® Server System R1000WF product family




Figure 62. Intel® Server System R1000WF Product Family

Table 53. Intel® Server System R1000WF Product Family Feature Set

Feature	Description
Chassis Type	1U Rack Mount Chassis
Server Board	Intel® Server Board S2600WF product family
Maximum Supported Processor Thermal Design Power (TDP)	165 Watts
External I/O Connections	<ul style="list-style-type: none"> • DB-15 Video connectors <ul style="list-style-type: none"> ◦ Front and Back • RJ-45 Serial Port A connector • Dual RJ-45 Network Interface connectors – (S2600WFT based systems only) • Dedicated RJ-45 server management NIC • Three USB 3.0 connectors on back panel • Two USB 3.0 connectors on front panel
Internal I/O Connectors / Headers	<ul style="list-style-type: none"> • One Type-A USB 2.0 connector • One DH-10 Serial Port B connector
System Fans	<ul style="list-style-type: none"> ▪ Six managed 40mm dual rotor system fans ▪ One power supply fan for each installed power supply module
Riser Card Support	<p>Support for two riser cards:</p> <ul style="list-style-type: none"> ▪ Riser #1 – PCIe* Gen3 x24 ▪ Riser #2 – PCIe* Gen3 x24 <p>With two riser cards installed, up to 2 possible add-in cards can be supported:</p> <ul style="list-style-type: none"> ▪ One x16 PCIe* 3.0 Add-in card slot per riser card ▪ 2 Full Height / Half Length add-in cards via Risers #1 and #2
Power Supply Options	<ul style="list-style-type: none"> ▪ The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 Redundant Power, and 2+0 Combined Power ▪ Two power supply options: <ul style="list-style-type: none"> ◦ AC 1100W Platinum ◦ DC 750W Gold
Storage Bay Options	<p>Hot Swap Backplane Options:</p> <p>Note: All available backplane options have support for SAS 3.0 (12 Gb/sec)</p> <ul style="list-style-type: none"> ◦ 4 x 3.5" SAS/SATA backplane ◦ 8 x 2.5" combo backplane – SAS/SATA/NVMe <p>Storage Bay Options:</p> <ul style="list-style-type: none"> ◦ 4 x 3.5" SAS/SATA hot swap drive bays + front panel I/O ◦ 8 x 2.5" SAS/SATA/NVMe hot swap drive bays + front panel I/O
Supported Rack Mount Kit Accessory Options	<ul style="list-style-type: none"> ▪ A1UFULLRAIL – Tool-less rack mount rail kit – 780mm max travel length ▪ A1USHRTRAIL – Tool-less rack mount rail kit – 780mm max travel length – No CMA support ▪ AXXELVRAIL – Enhanced value rack mount rail kit - 424mm max travel length ▪ AXX1U2UCMA – Cable Management Arm – (*supported with A1UFULLRAIL only) ▪ AXX2POSTBRCKT – 2-post fixed mount bracket kit

Intel® Server System R2000WF product family




Figure 63. Intel® Server System R2000WF Product Family

Table 54. Intel® Server System R2000WF Product Family Feature Set

Feature	Description
Chassis Type	2U Rack Mount Chassis
Server Board	Intel® Server Board S2600WF product family
Maximum Supported Processor Thermal Design Power (TDP)	145 Watts 165 Watts only supported in R2208WF models
External I/O connections	<ul style="list-style-type: none"> DB-15 Video connectors <ul style="list-style-type: none"> Front and Back RJ-45 Serial Port A connector Dual RJ-45 Network Interface connectors – (S2600WFT based systems only) Dedicated RJ-45 server management NIC Three USB 3.0 connectors on back panel Two USB 3.0 connectors on front panel
Internal I/O Connectors / Headers	<ul style="list-style-type: none"> One Type-A USB 2.0 connector One DH-10 Serial Port B connector
System Fans	<ul style="list-style-type: none"> Six managed 60mm dual rotor system fans One power supply fan for each installed power supply module
Riser Card Support	<p>Support for three riser cards:</p> <ul style="list-style-type: none"> Riser #1 – PCIe* 3.0 x24 - up to 3PCIe slots Riser #2 – PCIe* 3.0 x24 - up to 3PCIe slots Riser #3 – PCIe* 3.0 x16 – up to 2PCIe slots (optional) <p>With three riser cards installed, up to 8 possible add-in cards can be supported:</p> <ul style="list-style-type: none"> 4 Full Height / Half Length + 2 Full Height / Half Length add-in cards via Risers #1 and #2 2 Low profile add in cards via riser #3 (option) <p>See Chapter 8 for available riser card options.</p>
Power Supply Options	<ul style="list-style-type: none"> The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 Redundant Power, and 2+0 Combined Power Two power supply options: <ul style="list-style-type: none"> AC 1100W Platinum AC 1300W DC 750W Gold
Storage Bay Options	<ul style="list-style-type: none"> Hot Swap Backplane Options: Note: All available backplane options have support for SAS 3.0 (12 Gb/sec) <ul style="list-style-type: none"> 8 x 3.5" sas/sata 8 x 2.5" combo backplane – SAS/SATA/NVMe 12 x 3.5" sas/sata (supports up to 2 NVMe drives) Storage Bay Options: <ul style="list-style-type: none"> 8 x 3.5" SAS/SATA hot swap drive bays + Standard front panel 12 x 3.5" SAS/SATA hot swap drive bays (supports up to 2NVMe drives) + Storage Front Panel 8 x 2.5" SAS/SATA/NVMe hot swap drive bays + Standard front panel 16 x 2.5" SAS/SATA/NVMe hot swap drive bays + Standard front panel 24 x 2.5" SAS/SATA/NVMe swap drive bays + Standard front panel 2 x 2.5" SATA SSD Back of Chassis Hot Swap Drive Bays (accessory Option) 2 x internal fixed mount 2.5" SSDs (all SYSTEM MODELS)
Supported Rack Mount Kit Accessory Options	<ul style="list-style-type: none"> AXXELVRAIL – Enhanced value rack mount rail kit - 424mm max travel length AXXFULLRAIL – 2U Premium tool-less rail with CMA support AXXSHRTRAIL – 2U Premium tool-less rail kit with No CMA support AXXCMA2– Cable Management Arm – (*supported with AXXFULLRAIL only) AXX2POSTBRCKT – 2-post fixed mount bracket kit

Appendix F - Glossary of Terms

Term	Definition
BMC	Baseboard Management Controller
BIOS	Basic Input/Output System
CMOS	Complementary Metal-oxide-semiconductor
CPU	Central Processing Unit
DDR4	Double Data Rate 4th edition
DIMM	Dual In-line Memory Module
DPC	DIMMs per Channel
EDS	External Design Specification
EPS	External Product Specification
FP	Front Panel
FRB	Fault Resilient Boot
FRU	Field Replaceable Unit
GPGPU	General Purpose Graphic Processing Unit
I2C	Inter-integrated Circuit bus
iPC	Intel Product Code
LED	Light Emitting Diode
LRDIMM	Load Reduced DIMM
LSB	Least Significant Bit
MSB	Most Significant Bit
NIC	Network Interface Card
NMI	Non-maskable Interrupt
OCuLink	Optical Copper Link
PCI	Peripheral Component Interconnect
PCIe*	Peripheral Component Interconnect Express*
POST	Power-on Self-Test
PSU	Power Supply Unit
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RDIMM	Registered DIMM
ROC	RAID On Chip
SAS	Serial Attached SCSI
SATA	Serial Advanced Technology Attachment
SCA	Single Connector Attachment
SCSI	Small Computer System Interface
SDR	Sensor Data Record
SSD	Solid State Device
TPS	Technical Product Specification
Intel® TXT	Intel® Trusted Execution Technology for servers
VLSI	Very Large Scale Integration
VSB	Voltage Standby
Intel® VROC	Intel® Virtual RAID on CPU

Mouser Electronics

Authorized Distributor

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

Intel:

[S2600WF0](#) [S2600WFT](#)