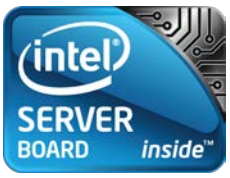
Intel® Server Board S2600WT

Technical Product Specification (TPS)



Revision 1.0

August 2014

Intel® Server Boards and Systems

Revision History

Date	Revision Number	Modifications
August 2014	1..0	1st External Public Release

Disclaimers

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, lifesaving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Server Board S2600WT 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.

This document and the software described in it are furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Copies of documents which have an order number and are referenced in this document, or other Intel® Literature, may be obtained by calling 1-800-548-4725, or go to: http://www.intel.com/design/Literature.htm.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation.

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

Copyright © 2014 Intel Corporation. All rights reserved.

.

Table of Contents

1. Introduction 1
1.1 Chapter Outline 1
1.2 Server Board Use Disclaimer 1
2. Product Features Overview 2
2.1 Server Board Component/Feature Identification 4
2.2 Product Architecture Overview 8
2.3 System Software Overview 9
2.3.1 System BIOS 9
2.3.2 Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data 13
2.3.3 Baseboard Management Controller (BMC) & Management Engine (ME) Firmware 13
3. Processor Support 14
3.1 Processor Socket Assembly 14
3.2 Processor Thermal Design Power (TDP) Support 15
3.3 Processor Population Rules 15
3.4 Processor Initialization Error Summary 16
3.5 Processor Function Overview 18
3.5.1 Processor Core Features: 18
3.5.2 Supported Technologies: 18
4. System Memory 21
4.1 Memory Sub-system Architecture 21
4.2 IMC Modes of operation 22
4.3 Memory RASM Features 22
4.4 Supported Memory 23
4.5 NVDIMM Support 24
4.6 Memory Slot Identification and Population Rules 24
4.6.1 Memory Interleaving Support 27
4.6.2 NUMA Configuration Support 27
4.7 System Memory Sizing and Publishing 27
4.7.1 Effects of Memory Configuration on Memory Sizing 27
4.7.2 Publishing System Memory 28
4.8 Memory Initialization 29
4.8.1 DIMM Discovery 29
4.8.2 DIMM Population Validation Check 29
4.8.3 Channel Training 30
5. System I/O 32
5.1 PCIe* Support 32

5.2 PCIe* Enumeration and Allocation 33
5.3 PCIe* Non-Transparent Bridge (NTB) 33
5.4 Add-in Card Support 34
5.4.1 Riser Card Support 35
5.4.2 I/O Module Support 38
5.4.3 Intel® Integrated RAID Option 39
5.5 Serial ATA (SATA) Support 40
5.5.1 Staggered Disk Spin-Up 42
5.6 Embedded SATA SW-RAID support 42
5.6.1 Intel® Rapid Storage Technology (RSTe) 4.0 42
5.6.2 Intel® Embedded Server RAID Technology 2 (ESRT2) 43
5.7 Network Interface 44
5.7.1 Intel® Ethernet Controller Options 45
5.7.2 Factory Programmed MAC Address Assignments 45
5.8 Video Support 45
5.8.1 Dual Video and Add-In Video Adapters 46
5.8.2 Setting Video Configuration Options using the BIOS Setup Utility 48
5.9 USB Support 50
5.9.1 Low Profile eUSB SSD Support 50
5.10 Serial Ports 51
6. System Security 53
6.1 BIOS Setup Utility Security Options Menu 53
6.1.1 Password Setup 53
6.1.2 System Administrator Password Rights 54
6.1.3 Authorized System User Password Rights and Restrictions 54
6.1.4 Front Panel Lockout 54
6.2 Trusted Platform Module (TPM) Support 55
6.2.1 TPM security BIOS 55
6.2.2 Physical Presence 56
6.2.3 TPM Security Setup Options 56
6.3 Intel® Trusted Execution Technology 57
7. Platform Management 58
7.1 Management Feature Set Overview 58
7.1.1 IPMI 2.0 Features Overview 58
7.1.2 Non IPMI Features Overview 59
7.2 Platform Management Features and Functions 61
7.2.1 Power Sub-system 61
7.2.2 Advanced Configuration and Power Interface (ACPI) 61

7.2.3 System Initialization 61
7.2.4 Watchdog Timer 62
7.2.5 System Event Log (SEL) 62
7.3 Sensor Monitoring 62
7.3.1 Sensor Scanning 63
7.3.2 Sensor Rearm Behavior 63
7.3.3 BIOS Event-Only Sensors 64
7.3.4 Margin Sensors 64
7.3.5 IPMI Watchdog Sensor 64
7.3.6 BMC Watchdog Sensor 64
7.3.7 BMC System Management Health Monitoring 64
7.3.8 VR Watchdog Timer 64
7.3.9 System Airflow Monitoring 64
7.3.10 Thermal Monitoring 65
7.3.11 Processor Sensors 68
7.3.12 Voltage Monitoring 70
7.3.13 Fan Monitoring 70
7.3.14 Standard Fan Management 72
7.3.15 Power Management Bus (PMBus*) 78
7.3.16 Power Supply Dynamic Redundancy Sensor 78
7.3.17 Component Fault LED Control 79
7.3.18 NMI (Diagnostic Interrupt) Sensor 80
7.3.19 LAN Leash Event Monitoring 80
7.3.20 Add-in Module Presence Sensor 80
7.3.21 CMOS Battery Monitoring 80
8. Intel® Intelligent Power Node Manager (NM) Support Overview 81
8.1 Hardware Requirements 81
8.2 Features 81
8.3 ME System Management Bus (SMBus*) interface 81
8.4 PECI 3.0 82
8.5 NM “Discovery” OEM SDR 82
8.6 SmaRT/CLST 82
8.6.1 Dependencies on PMBus*-compliant Power Supply Support 83
9. Basic and Advanced Server Management Features 84
9.1 Dedicated Management Port 85
9.2 Embedded Web Server 85
9.3 Advanced Management Feature Support (RMM4 Lite) 87
9.3.1 Keyboard, Video, Mouse (KVM) Redirection 87

9.3.2 Remote Console 88
9.3.3 Performance 88
9.3.4 Security 89
9.3.5 Availability 89
9.3.6 Usage 89
9.3.7 Force-enter BIOS Setup 89
9.3.8 Media Redirection 89
10. On-board Connector/Header Overview 91
10.1 Power Connectors 91
10.1.1 Main Power 91
10.1.2 Hot Swap Backplane Power Connector 92
10.1.3 Peripheral Drive Power Connector 93
10.1.4 Riser Card Supplemental 12V Power Connectors 93
10.2 Front Panel Headers and Connectors 94
10.2.1 Front Panel Button and LED Support 94
10.2.2 Front Panel LED and Control Button Features Overview 95
10.2.3 Front Panel USB 2.0 Connector 96
10.2.4 Front Panel USB 3.0 Connector 97
10.2.5 Front Panel Video Connector 97
10.2.6 Intel® Local Control Panel Connector 97
10.3 On-Board Storage Option Connectors 98
10.3.1 Single Port SATA Only Connectors 98
10.3.2 Internal Type-A USB Connector 99
10.3.3 Internal 2mm Low Profile eUSB SSD Connector 99
10.4 System Fan Connectors 100
10.5 Other Connectors and Headers 101
10.5.1 Chassis Intrusion Header 101
10.5.2 Storage Device Activity LED Header 101
10.5.3 Intelligent Platform Management Bus (IPMB) Connector 101
10.5.4 Hot Swap Backplane I2C* Connectors 102
10.5.5 SMBus Connector 102
11. Reset and Recovery Jumpers 103
11.1 BIOS Default Jumper Block 103
11.2 Serial Port 'A' Configuration Jumper 104
11.3 Password Clear Jumper Block 104
11.4 Management Engine (ME) Firmware Force Update Jumper Block 104
11.5 BMC Force Update Jumper Block 105
11.6 BIOS Recovery Jumper 106

12. Light Guided Diagnostics 107
12.1 System ID LED 108
12.2 System Status LED 108
12.3 BMC Boot/Reset Status LED Indicators 111
12.4 Post Code Diagnostic LEDs 111
12.5 Fan Fault LEDs 111
12.6 Memory Fault LEDs 111
12.7 CPU Fault LEDs 111
13. Power Supply Specification Guidelines 112
13.1 Power Supply DC Output Connector 112
13.2 Power Supply DC Output Specification 113
13.2.1 Output Power/Currents 113
13.2.2 Standby Output 113
13.2.3 Voltage Regulation 113
13.2.4 Dynamic Loading 113
13.2.5 Capacitive Loading 114
13.2.6 Grounding 114
13.2.7 Closed loop stability 114
13.2.8 Residual Voltage Immunity in Standby mode 114
13.2.9 Common Mode Noise 114
13.2.10 Soft Starting 114
13.2.11 Zero Load Stability Requirements 114
13.2.12 Hot Swap Requirements 114
13.2.13 Forced Load Sharing 114
13.2.14 Ripple/Noise 115
13.2.15 Timing Requirements 115
Appendix A: Integration and Usage Tips 117
Appendix B: Integrated BMC Sensor Tables 118
Appendix C: Management Engine Generated SEL Event Messages 132
Appendix D: POST Code Diagnostic LED Decoder 134
Appendix E: POST Code Errors 141
Appendix F: Supported Intel® Server Systems 147

List of Figures

Figure 1. Server Board Component/Features Identification 4
Figure 2. Intel® Server Board S2600WT External I/O Connector Layout 5
Figure 3. Intel® Light Guided Diagnostics - DIMM Fault LEDs 5
Figure 4. Intel® Light Guided Diagnostic LED Identification 6
Figure 5. Jumper Block Identification 7
Figure 6. Intel® Server Board S2600WT Architectural Block Diagram 8
Figure 7. Processor Socket Assembly 14
Figure 8. LGA2011-3 ILM (Narrow) 14
Figure 9. Memory Sub-system Block Diagram 21
Figure 10. Memory Slots Definition 24
Figure 11. Intel® Server Board S2600WT Memory Slot Layout 25
Figure 12. On-board Add-in Card Support 34
Figure 13. 1U one slot PCIe* riser card (iPN – F1UL16RISER2) 36
Figure 14. 2U three PCIe* slot riser card (iPN – A2UL8RISER2) 36
Figure 15. 2U two PCIe* slot riser card (iPN – A2UL16RISER2) 37
Figure 16. 2U three PCIx/PCIe* slot riser (iPN - A2ULPCIXRISER2) 37
Figure 17. 2U two PCIe* slot (Low Profile) PCIe* Riser card (iPN – A2UX8X4RISER) – Riser Slot #3 compatible only 38
Figure 18. Server Board Layout - I/O Module Connector 38
Figure 19. Server Board Layout – Intel® Integrated RAID Module Option Placement 39
Figure 20. Onboard SATA Features 40
Figure 21. SATA RAID 5 Upgrade Key 44
Figure 22. Network Interface Connectors 44
Figure 23. External RJ45 NIC Port LED Definition 45
Figure 24. BIOS Setup Utility - Video Configuration Options 48
Figure 25. Onboard USB Port Support 50
Figure 26. Low Profile eUSB SSD Support 50
Figure 27. High-level Fan Speed Control Process 75
Figure 28. Intel® RMM4 Lite Activation Key Installation 85
Figure 29. High Power Add-in Card 12V Auxiliary Power Cable Option 93
Figure 30. System Fan Connector Pin-outs 100
Figure 31. System Fan Connector Placement 100
Figure 32. Reset and Recovery Jumper Block Location 103
Figure 33. On-Board Diagnostic LED Placement 107
Figure 34. DIMM Fault LED Placement 108
Figure 35. Turn On/Off Timing (Power Supply Signals) 116

Figure 36. POST Diagnostic LED Location 134
Figure 37. Intel® Server System R1000WT 147
Figure 38. Intel® Server System R2000WT 150

List of Tables

Table 1. Intel® Server Board S2600WT Feature Set 2
Table 2. POST Hot-Keys 11
Table 3. Mixed Processor Configurations Error Summary 16
Table 4. DDR4 RDIMM & LRDIMM Support 23
Table 5. Intel® Server Board S2600WT Memory Slot Identification 25
Table 6. DIMM Population Matrix 26
Table 7. PCIe* Port Routing CPU #1 32
Table 8. PCIe* Port Routing – CPU #2 33
Table 9. Riser Card #1 - PCIe* Root Port Mapping 35
Table 10. Riser Card #2 - PCIe* Root Port Mapping 35
Table 11. Riser Slot #3 - PCIe* Root Port Mapping 36
Table 12. Supported Intel® I/O Module Options 39
Table 13. SATA and sSATA Controller BIOS Utility Setup Options 41
Table 14. SATA and sSATA Controller Feature Support 41
Table 15. Video Modes 46
Table 16. Serial A Connector Pin-out 51
Table 17. Serial-B Connector Pin-out 52
Table 18. TPM Setup Utility – Security Configuration Screen Fields 57
Table 19. Server Board Power Control Sources 61
Table 20. ACPI Power States 61
Table 21. Processor Sensors 68
Table 22. Processor Status Sensor Implementation 68
Table 23. Component Fault LEDs 79
Table 24. Intel® Remote Management Module 4 (RMM4) Options 84
Table 25. Basic and Advanced Server Management Features Overview 84
Table 26. Main Power (Slot 1) Connector Pin-out (“MAIN PWR 1”) 91
Table 27. Main Power (Slot 2) Connector Pin-out ("MAIN PWR 2”) 92
Table 28. Hot Swap Backplane Power Connector Pin-out (“HSBP PWR") 92
Table 29. Peripheral Drive Power Connector Pin-out ("Peripheral_PWR") 93
Table 30. Riser Slot Auxiliary Power Connector Pin-out ("OPT_12V_PWR”) 93
Table 31. Front Panel Features 94
Table 32. Front Panel Connector Pin-out ("Front Panel” and “Storage FP”) 94
Table 33. Power/Sleep LED Functional States 95
Table 34. NMI Signal Generation and Event Logging 96
Table 35. Front Panel USB 2.0 Connector Pin-out ("FP_USB_2.0_5-6 ") 96
Table 36. Front Panel USB 2.0/3.0 Connector Pin-out (“FP_USB_2.0/ 3.0”) 97

Table 37. Front Panel Video Connector Pin-out ("FP VIDEO") 97
Table 38. Intel Local Control Panel Connector Pin-out ("LCP") 98
Table 39. Single Port SATA Connector Pin-out ("SATA 4" & "SATA 5") 98
Table 40. SATA SGPIO Connector Pin-out ("SATA_SGPIO") 99
Table 41. Internal Type-A USB Connector Pin-out ("USB 2.0") 99
Table 42. Internal eUSB Connector Pin-out ("eUSB SSD") 99
Table 43. Chassis Intrusion Header Pin-out ("CHAS_INTR") 101
Table 44. Hard Drive Activity Header Pin-out ("HDD_LED") 101
Table 45. IPMB Connector Pin-out 101
Table 46. Hot-Swap Backplane I2C* Connector Pin-out 102
Table 47. SMBus Connector Pin-out 102
Table 48. System Status LED State Definitions 109
Table 49. BMC Boot/Reset Status LED Indicators 111
Table 50. Power Supply DC Power Output Connector Pinout 112
Table 51. Minimum Load Ratings 113
Table 52. Voltage Regulation Limits 113
Table 53. Transient Load Requirements 113
Table 54. Capacitive Loading Conditions 114
Table 55. Ripples and Noise 115
Table 56. Timing Requirements 115
Table 57. BMC Core Sensors 120
Table 58. Server Platform Services Firmware Health Event 132
Table 59. Node Manager Health Event 133
Table 60. POST Progress Code LED Example 135
Table 61. MRC Progress Codes 135
Table 62. MRC Fatal Error Codes 136
Table 63. POST Progress Codes 138
Table 64. POST Error Codes and Messages 141
Table 65. POST Error Beep Codes 146
Table 66. Integrated BMC Beep Codes 146
Table 67. Intel® Server System R1000WT Product Family Feature Set 147
Table 68. Intel® Server System R2000WT Product Family Feature Set 150

<This page is intentionally left blank.>

1. Introduction

This Technical Product Specification (TPS) provides board-specific information detailing the features, functionality, and high-level architecture of the Intel® Server Board S2600WT.

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

1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Product Features Overview
- Chapter 3 – Processor Support
- Chapter 4 – System Memory
- Chapter 5 – System I/O
- Chapter 6 – System Security
- Chapter 7 – Platform Management
- Chapter 8 – Intel® Intelligent Power Node Manager (NM) Support Overview
- Chapter 9 – Basic and Advanced Server Management Features
- Chapter 10 – On-Board Connector and Header Overview
- Chapter 11 – Reset and Recovery Jumpers
- Chapter 12 – Light-Guided Diagnostics
- Chapter 13 – Power Supply Specification Guidelines
- Appendix A – Integration and Usage Tips
- Appendix B – Integrated BMC Sensor Tables
- Appendix C – Management Engine Generated SEL Event Messages
- Appendix D – POST Code Diagnostic LED Decoder
- Appendix E – POST Code Errors
- Appendix F – Supported Intel® Server Systems

1.2 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 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.

2. Product Features Overview

The Intel® Server Board S2600WT 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 E5-2600 v3 product family. Previous generation Intel® Xeon® processors are not supported.

The server board is offered with either of the two following on-board networking options:

- Intel® Ethernet Controller X540, supporting 10 GbE (Intel Server Board Product Code - S2600WTT)
- Intel® Ethernet Controller I350, supporting 1 GbE (Intel Server Board Product Code – S2600WT2)

All other onboard features will be identical.

Table 1. Intel® Server Board S2600WT Feature Set

Feature	Description
Processor Support	• Two LGA2011-3 (Socket R3) processor sockets • Support for one or two Intel® Xeon® processors E5-2600 v3 product family • Maximum supported Thermal Design Power (TDP) of up to 145 W
Memory	• 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor • Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) • Memory data transfer rates: o DDR4 RDIMM: 1600 MT/s (3DPC), 1866 MT/s (2DPC) and 2133 MT/s (1DPC) o DDR4 LRDIMM: 1600 MT/s (3DPC), 2133 MT/s (2DPC & 1DPC) • DDR4 standard I/O voltage of 1.2V
Chipset	Intel® C612 chipset
External (Back Panel) I/O connections	• DB-15 Video connector • RJ-45 Serial Port A connector • Dual RJ-45 Network Interface connectors supporting either : o 10 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WTT) or o 1 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WT2) • Dedicated RJ-45 server management port • Three USB 2.0 / 3.0 connectors
Internal I/O connectors/headers	• One Type-A USB 2.0 connector • One 2x5 pin connector providing front panel support for two USB 2.0 ports • One 2x10 pin connector providing front panel support for two USB 2.0 / 3.0 ports • One 2x15 pin SSI-EEB compliant Standard Front Panel header • One 2x15 high density Storage Front Panel connector • One 2x7pin Front Panel Video connector • One 1x7pin header for optional Intel® Local Control Panel support • One DH-10 Serial Port B connector
PCIe* Support	• PCIe* 3.0 (2.5, 5, 8 GT/s) – backwards compatible with PCIe* Gen 1 and Gen 2 devices
1U Server – Riser Card Support	• Server board includes two PCIe* 3.0 compatible riser card only slots o Riser #1 – PCIe* 3.0 x24 – 1 PCIe* Full Height / Half Length add-in card support in 1U o Riser #2 – PCIe* 3.0 x24 – 1 PCIe* Full Height / Half Length add-in card support in 1U

Feature	Description
2U Server – Riser Card Support	• Server board includes three PCIe* 3.0 compatible riser card only slots: o Riser #1 – PCIe* 3.0 x24 – up to 3 PCIe* slots in 2U o Riser #2 – PCIe* 3.0 x24 – up to 3 PCIe* slots in 2U o Riser #3 – PCIe* 3.0 x8 + DMI x4 (PCIe* 2.0 compatible) – up to 2 PCIe* slots in 2U • With three riser cards installed, up to 8 possible add-in cards can be supported: o 4 Full Height / Full Length + 2 Full Height / Half Length add-in cards via Risers #1 and #2 o 2 low profile add-in cards via Riser #3
Available I/O Module Options	The server board includes a proprietary on-board connector allowing for the installation of a variety of available I/O modules. An installed I/O module can be supported in addition to standard on-board features and add-in PCIe* cards. • AXX4P1GBPWLIOM – Quad port RJ45 1 GbE based on Intel® Ethernet Controller I350 • AXX10GBTWLIOM – Dual port RJ-45 10GBase-T I/O Module based on Intel® Ethernet Controller x540 • AXX10GBNIAIOM – Dual port SFP+ 10 GbE module based on Intel® 82599 10 GbE controller • AXX1FDRIBIOM – Single port QSFP FDR 56 GT/S speed InfiniBand* module • AXX2FDRIBIOM – Dual port QSFP FDR 56 GT/S speed infiniband* module • AXX1P40FRTIOM – Single port QSFP+ 40 GbE module • AXX2P40FRTIOM – Dual port QSFP+ 40 GbE module
System Fan Support	• Six system fans supported in two different connector formats: hot swap (2U) and cabled (1U) o Six 10-pin managed system fan headers (Sys_Fan 1-6) – Used for 1U system configuration o Six 6-pin hot swap capable managed system fan connectors (Sys_Fan 1-6) – Used for 2U system configuration
Video	• Integrated 2D Video Controller • 16 MB DDR3 Memory
On-board storage controllers and options	• 10x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) o Two 7-pin single port SATA connectors capable of supporting up to 6 Gb/sec o Two 4-port mini-SAS HD (SFF-8643) connectors capable of supporting up to 6 Gb/sec SATA • One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices • Optional SAS IOC/ROC support via on-board Intel® Integrated RAID module connector • Embedded Software SATA RAID o Intel® Rapid Storage RAID Technology (RSTe) 4.0 o Intel® Embedded Server RAID Technology 2 (ESRT2) with optional RAID 5 key support
Security	• Intel® Trusted Platform Module (TPM) - AXXTPME5 (Accessory Option)
Server Management	• Integrated Baseboard Management Controller, IPMI 2.0 compliant • Support for Intel® Server Management Software • On-board RJ45 management interface • Intel® Remote Management Module 4 Lite support (Accessory Option)

2.1 Server Board Component/Feature Identification

The following illustration provides a general overview of the server board, identifying key feature and component locations.

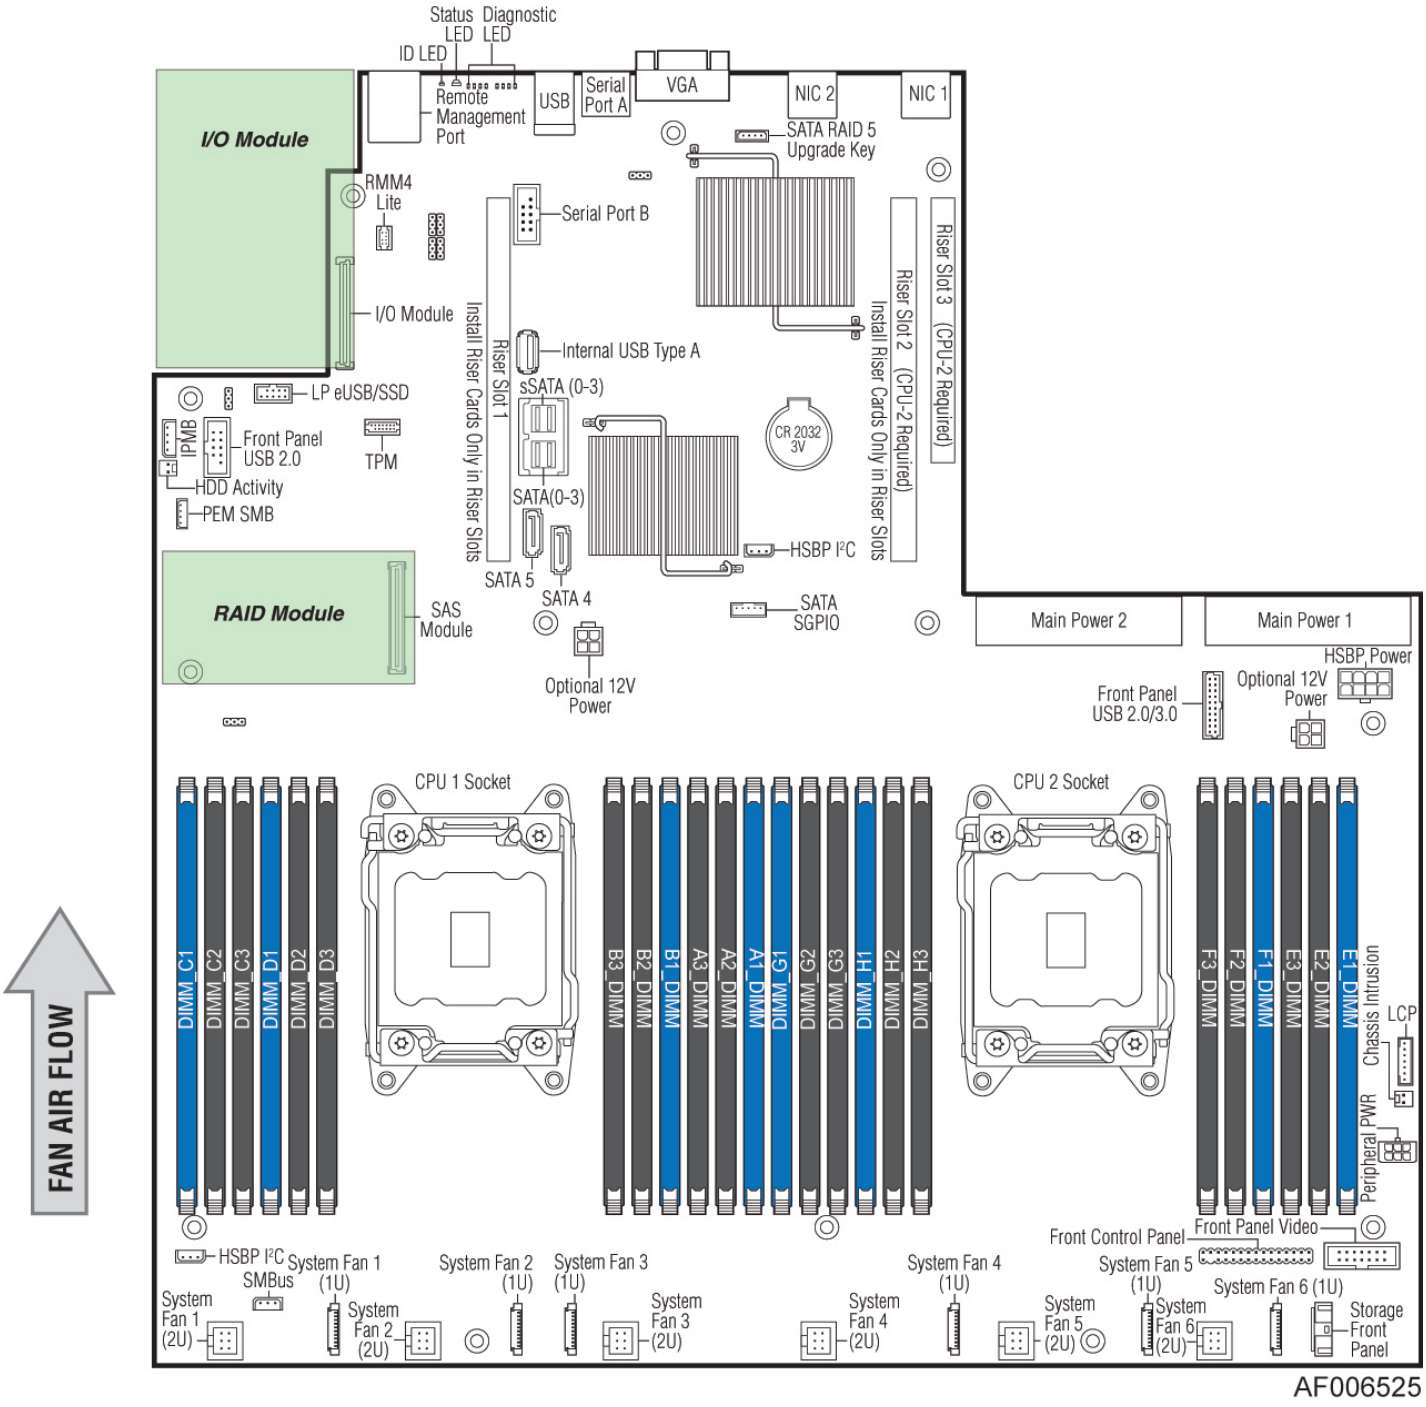
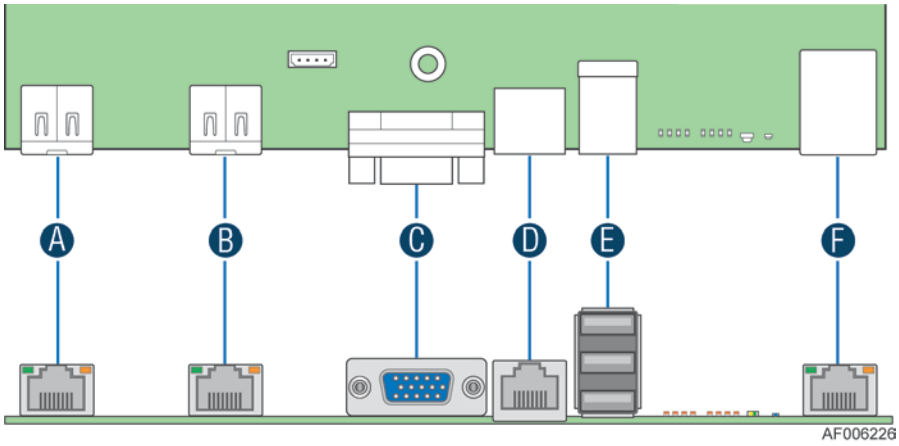


Figure 1. Server Board Component/Features Identification

The back edge of the server board includes several external connectors to support the following features:



A – RJ45 Networking Port – NIC #1
B – RJ45 Networking Port – NIC #2
C – Video
D – RJ45 Serial 'A' Port
E – Stacked 3-port USB 2.0 / 3.0
F – RJ45 Dedicated Management Port

Figure 2. Intel® Server Board S2600WT External I/O Connector Layout

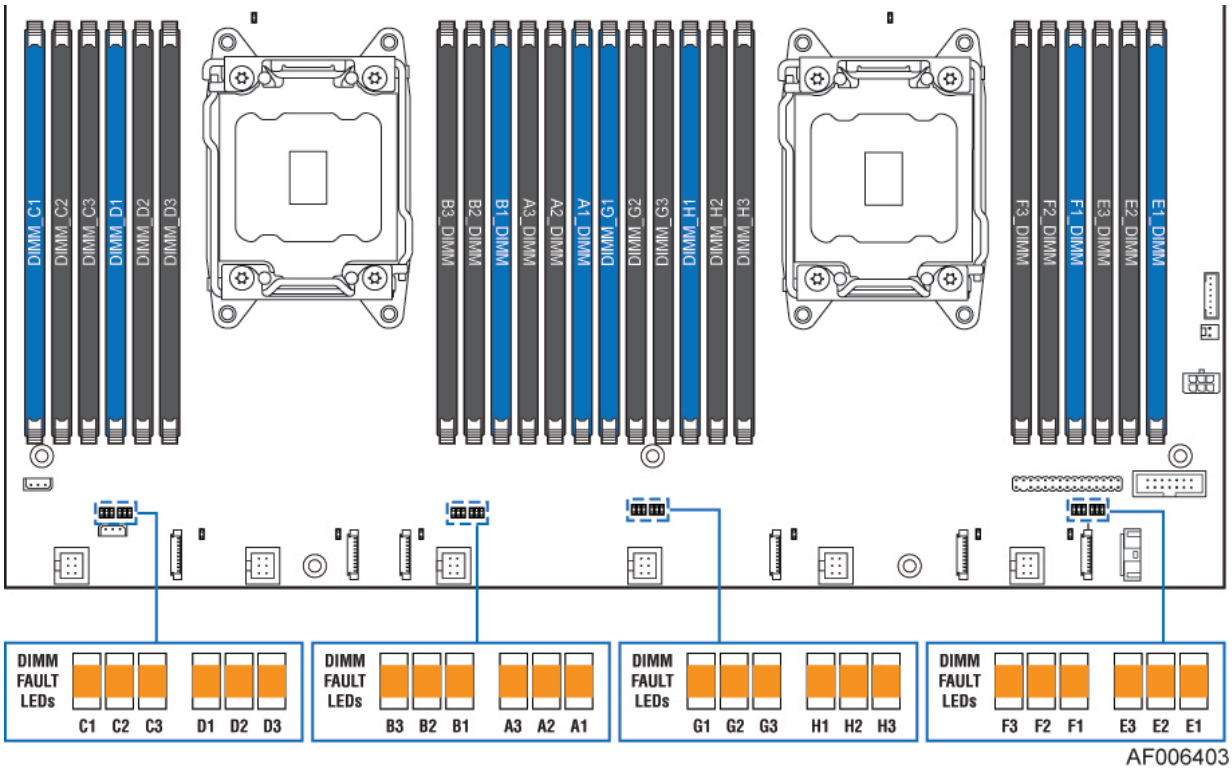


Figure 3. Intel® Light Guided Diagnostics - DIMM Fault LEDs

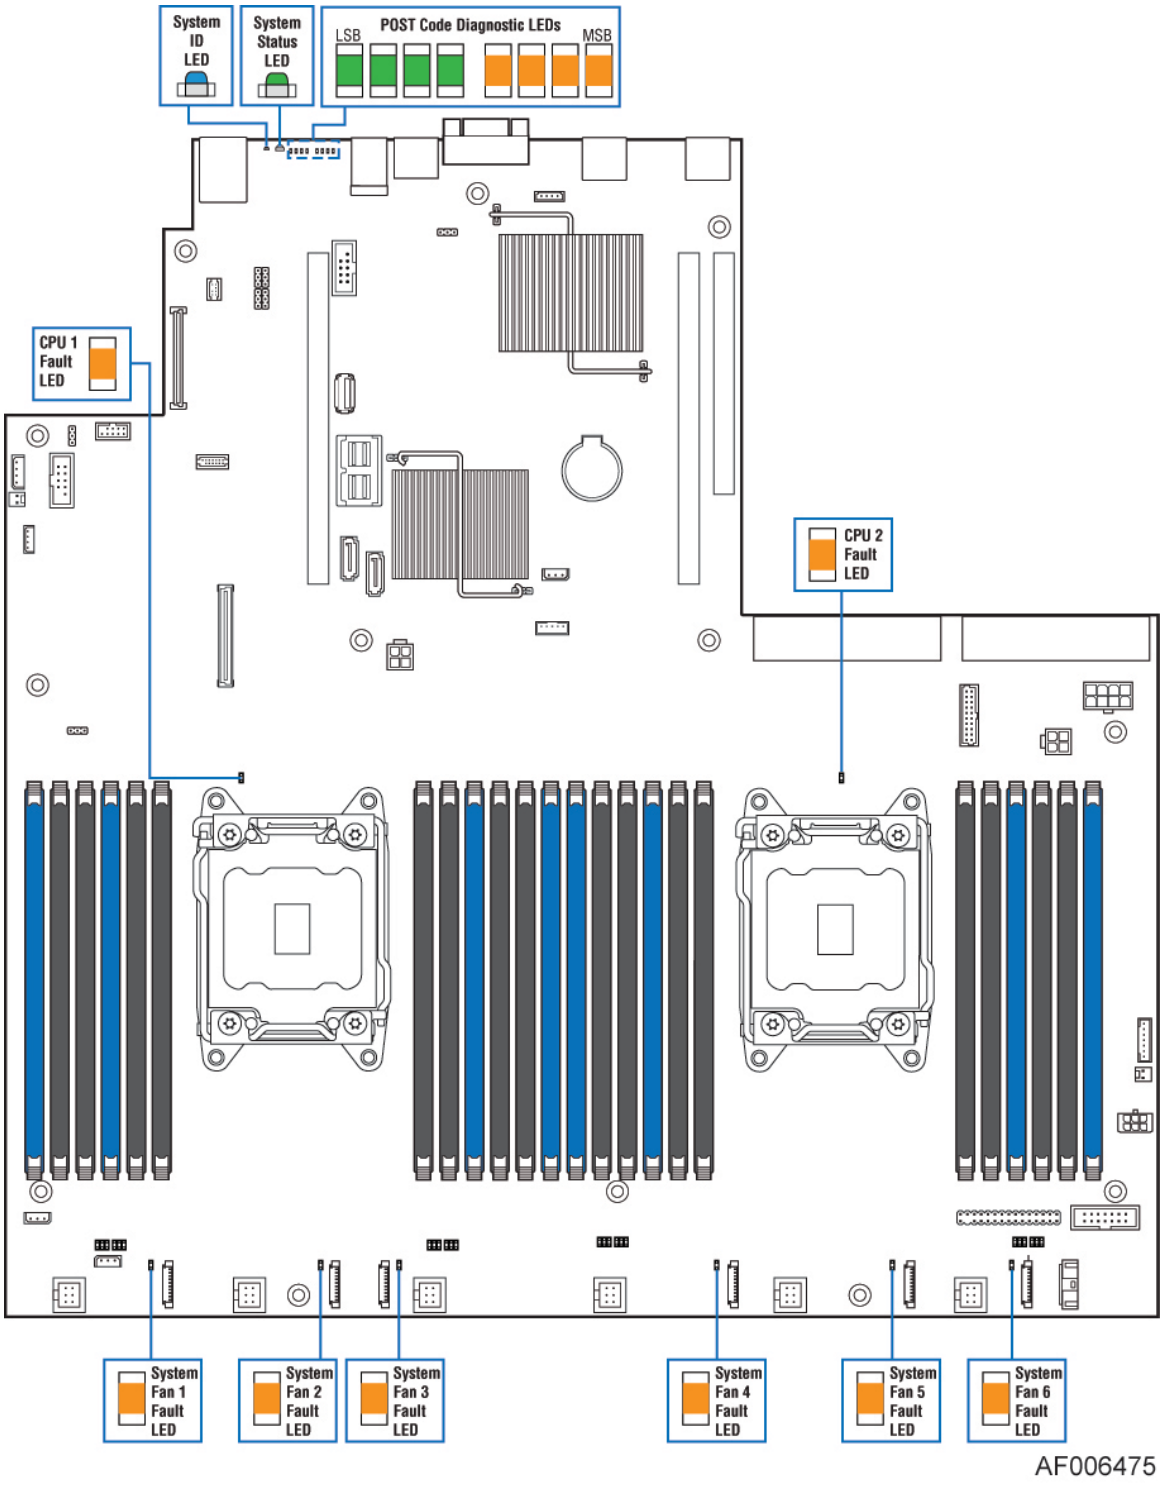


Figure 4. Intel® Light Guided Diagnostic LED Identification

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

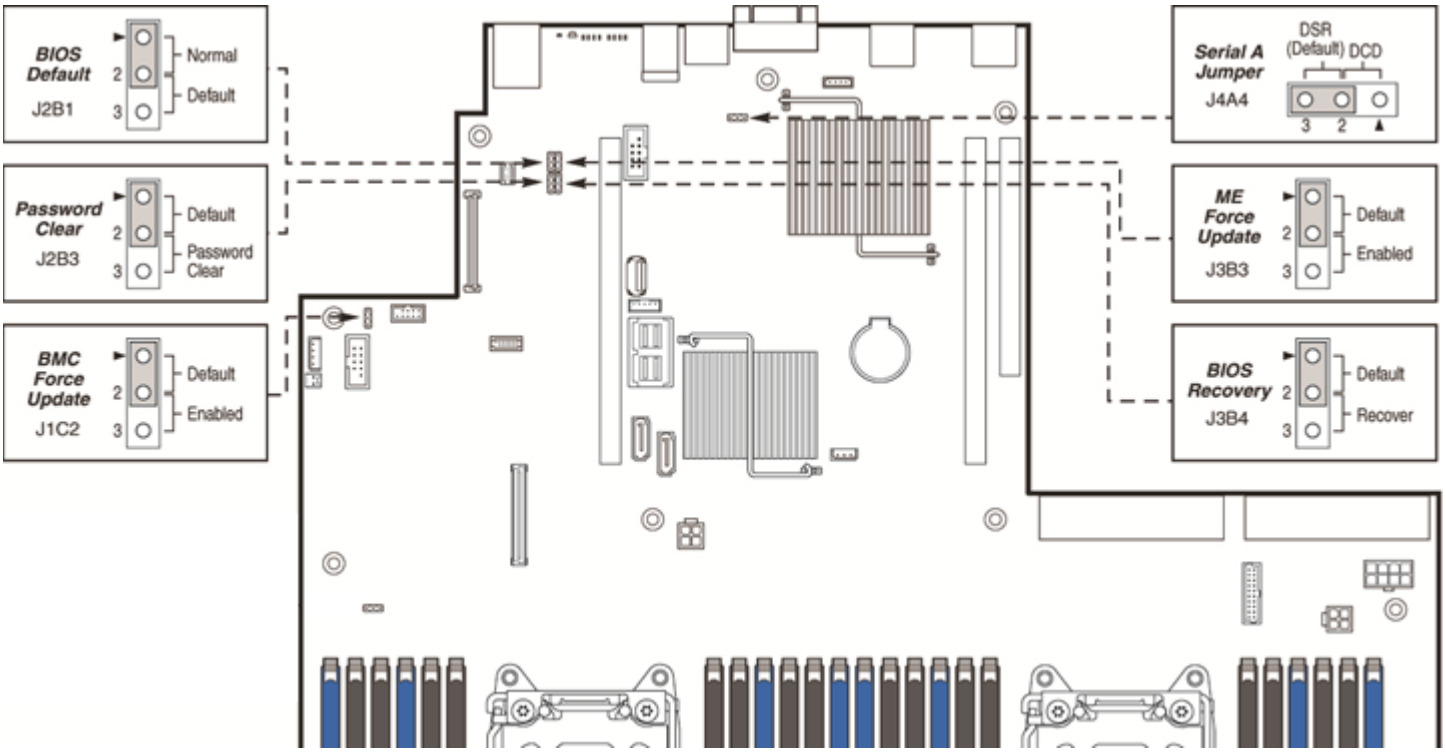


Figure 5. Jumper Block Identification

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

2.2 Product Architecture Overview

The architecture of Intel® Server Board S2600WT is developed around the integrated features and functions of the Intel® Xeon® processor E5-2600 v3 product family, the Intel® C612 chipset, Intel® Ethernet Controllers I350 1 GbE or X540 10 GbE, and the Emulex* Pilot-III Baseboard Management Controller.

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

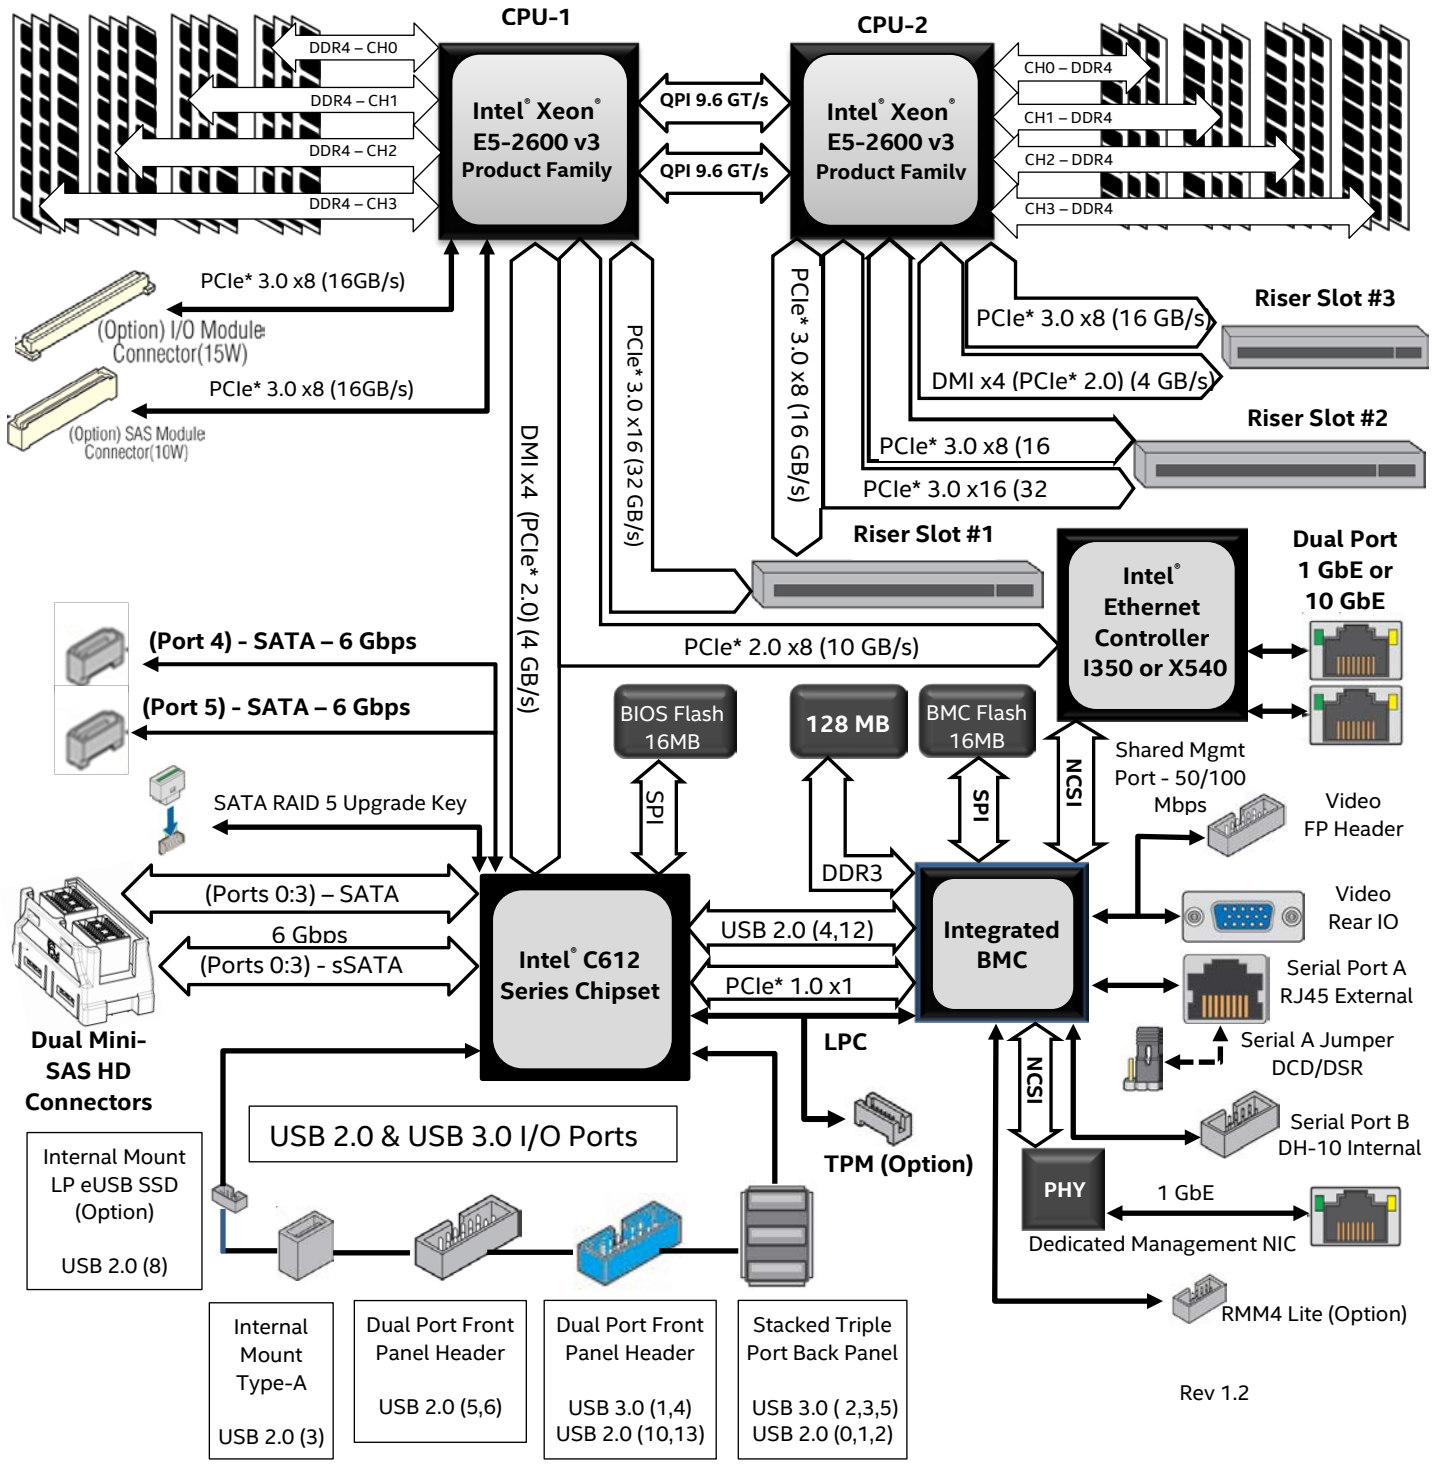


Figure 6. Intel® Server Board S2600WT Architectural Block Diagram

2.3 System Software Overview

The server board includes an embedded software stack to enable, configure, and support various system functions. This software stack includes the System BIOS, Baseboard Management Controller (BMC) Firmware, Management Engine (ME) Firmware, and management support data including Field Replaceable Unit (FRU) data, and Sensor Data Record (SDR) data.

The system software is pre-programmed on the server board during factory assembly, making the server board functional at first power on after system integration. Typically, as part of the initial system integration process, FRU and SDR data will have to be installed onto the server board by the system integrator to ensure the embedded platform management subsystem is able to provide best performance and cooling for the final system configuration. It is also not uncommon for the system software stack to be updated to later revisions to ensure the most reliable system operation. Intel makes periodic system software updates available for download at the following Intel website:

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).

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 E5-2600 v3 product family*
- *Intel® Server System BMC Firmware External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 v3 product family*

2.3.1 System BIOS

The system BIOS is implemented as firmware that resides in flash memory on the server board. The BIOS provides hardware-specific initialization algorithms and standard compatible basic input/output services, and standard Intel® Server Board features. The flash memory also contains firmware for certain embedded devices.

This BIOS implementation is based on the Extensible Firmware Interface (EFI), according to the Intel® Platform Innovation Framework for EFI architecture, as embodied in the industry standards for Unified Extensible Firmware Interface (UEFI).

The implementation is compliant with all Intel® Platform Innovation Framework for EFI architecture specifications, as further specified in the *Unified Extensible Firmware Interface Reference Specification*, Version 2.3.1.

In the UEFI BIOS design, there are three primary components: the BIOS itself, the Human Interface Infrastructure (HII) that supports communication between the BIOS and external programs, and the Shell which provides a limited OS-type command-line interface. This BIOS system implementation complies with HII Version 2.3.1, and includes a Shell.

2.3.1.1 BIOS Revision Identification

The BIOS Identification string is used to uniquely identify the revision of the BIOS being used on the server. The BIOS ID string is displayed on the Power On Self Test (POST) Diagnostic Screen and in the <F2> BIOS Setup Main Screen, as well as in System Management BIOS (SMBIOS) structures.

The BIOS ID string for S2600 series server boards is formatted as follows:

BoardFamilyID.OEMID.MajorVer.MinorVer.RelNum.BuildDateTime

Where:

- **BoardFamilyID** = String name to identify board family.
  - "***SE5C610***" is used to identify BIOS builds for Intel® S2600 series Server Boards, based on the Intel® Xeon® Processor E5-2600 v3 product families and the Intel® C612 chipset.
- **OEMID** = Three-character OEM BIOS Identifier, to identify the board BIOS "owner".
  - "***86B***" is used for Intel PCSD Commercial BIOS Releases.
- **MajorVer** = Major Version, two decimal digits 01-99 which are changed only to identify major hardware or functionality changes that affect BIOS compatibility between boards.
  - "***01***" is the starting BIOS Major Version for all platforms.
- **MinorVer** = Minor Version, two decimal digits 00-99 which are changed to identify less significant hardware or functionality changes which do not necessarily cause incompatibilities but do display differences in behavior or in support of specific functions for the board.
- **RelNum** = Release Number, four decimal digits which are changed to identify distinct BIOS Releases. BIOS Releases are collections of fixes and/or changes in functionality, built together into a BIOS Update to be applied to a Server Board. However, there are "Full Releases" which may introduce many new fixes/functions, and there are "Point Releases" which may be built to address very specific fixes to a Full Release.

  The Release Numbers for Full Releases increase by 1 for each release. For Point Releases, the first digit of the Full Release number on which the Point Release is based is increased by 1. That digit is always 0 (zero) for a Full Release.
- **BuildDateTime** = Build timestamp – date and time in MMDDYYYYHHMM format:
  - **MM** = Two-digit month.
  - **DD** = Two-digit day of month.
  - **YYYY** = Four-digit year.
  - **HH** = Two-digit hour using 24-hour clock.
  - **MM** = Two-digit minute.

An example of a valid BIOS ID String is as follows:

SE5C610.86B.01.01.0003.081320110856

The BIOS ID string is displayed on the POST diagnostic screen for BIOS Major Version 01, Minor Version 01, Full Release 0003 that is generated on August 13, 2011 at 8:56 AM.

The BIOS version in the <F2> BIOS Setup Utility Main Screen is displayed without the time/date timestamp, which is displayed separately as "Build Date":

SE5C610.86B.01.01.0003

2.3.1.2 Hot Keys Supported During POST

Certain “Hot Keys” are recognized during POST. A Hot Key is a key or key combination that is recognized as an unprompted command input, that is, the operator is not prompted to press the Hot Key and typically the Hot Key will be recognized even while other processing is in progress.

The BIOS recognizes a number of Hot Keys during POST. After the OS is booted, Hot Keys are the responsibility of the OS and the OS defines its own set of recognized Hot Keys.

The following table provides a list of available POST Hot Keys along with a description for each.

Table 2. 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

2.3.1.3 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 is displayed 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.

2.3.1.4 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, where a User password allows only booting in the order previously defined by the Administrator.

2.3.1.5 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 is displayed 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 is entered, the Main screen is displayed initially. However, in the event 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 E5-2600 V3 product family

2.3.1.6 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 when update packages are downloaded from the Intel web site.

2.3.1.7 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 provides three mechanisms to start the BIOS recovery process, which is called Recovery Mode:

- Recovery Mode Jumper – this jumper causes the BIOS to boot in Recovery Mode
- The Boot Block detects partial BIOS update and automatically boots in Recovery Mode
- The BMC asserts Recovery Mode GPIO in case of partial BIOS update and FRB2 time-out

The BIOS Recovery takes place without any external media or Mass Storage device as it utilizes a Backup BIOS image 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 *BIOS Recovery Jumper* is set (See Figure 5. Jumper Block Identification), the BIOS begins by logging a 'Recovery Start' event to the System Event Log (SEL). It then loads and boots with a Backup BIOS image residing in the BIOS flash device. This process takes place before any video or console is available. The system boots to the embedded uEFI shell, and a 'Recovery Complete' event is logged to the SEL. 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 web site. Once the update has completed, the recovery jumper is switched back to its default position and the system is power cycled.

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

2.3.2 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. The BMC supports automatic configuration of the manageability subsystem after changes have been made to the system's hardware configuration. Once the system integrator has performed an initial SDR/CFG 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.

- Processors
- I/O Modules (dedicated slot modules)
- Storage modules, such as a SAS module (dedicated slot modules)
- Power supplies
- Fans
- Fan options (e.g. upgrade from non-redundant cooling to redundant cooling)
- Intel® Xeon Phi™ co-processor cards
- 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.

2.3.2.1 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 Intel web site.

2.3.3 Baseboard Management Controller (BMC) & Management Engine (ME) Firmware

See Chapters 7, 8, and 9 for features and functions associated with the BMC firmware and ME firmware.

3. Processor Support

The server board includes two Socket-R3 (LGA2011-3) processor sockets and can support one or two of the following processors:

- Intel® Xeon® processor E5-2600 v3 product family
- Supported Thermal Design Power (TDP) of up to 145W.

Note: Previous generation Intel® Xeon® processors are not supported on the Intel server boards described in this document.

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

3.1 Processor Socket Assembly

Each processor socket of the server board is pre-assembled with an Independent Latching Mechanism (ILM) and Back Plate which allow for secure placement of the processor and processor heat sink to the server board.

The illustration below identifies each sub-assembly component:

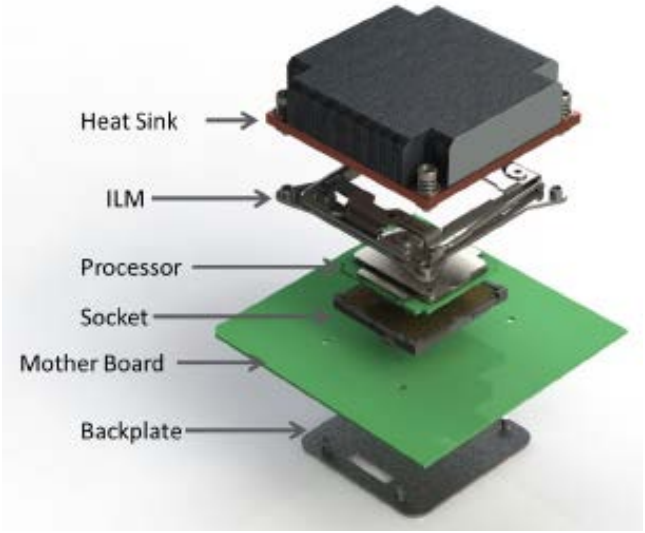


Figure 7. Processor Socket Assembly

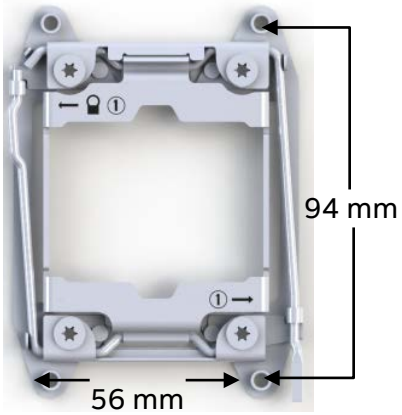


Figure 8. LGA2011-3 ILM (Narrow)

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 E5-2600 v3 product family TDP guidelines up to and including 145W.

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 Processor Population Rules

Note: The server board may support dual-processor configurations consisting of different processors that meet the defined criteria below, however, Intel does not perform validation testing of this configuration. 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-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 board features may not be functional without having a second processor installed. See Figure 6. Intel® Server Board S2600WT Architectural Block Diagram.

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

- Both processors must be of the same processor family
- Both processors must have the same number of cores
- Both processors must have the same cache sizes for all levels of processor cache memory

Processors with different core frequencies can be mixed in a system, given 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.

Processors which have different Intel® Quickpath (QPI) Link Frequencies may operate together if they are otherwise compatible and if a common link frequency can be selected. The common link frequency would be the highest link frequency that all installed processors can achieve.

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

3.4 Processor Initialization Error Summary

The following table describes mixed processor conditions and recommended actions for all Intel® server boards and Intel server systems designed around the Intel® Xeon® processor E5-2600 v3 product family and Intel® C612 chipset architecture. The errors fall into one of the following categories:

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

"Unrecoverable fatal error found. System will not boot until the error is resolved

Press <F2> to enter setup"

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

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

If the system is not able to boot, the system will generate a beep code consisting of 3 long beeps and 1 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 3. Mixed Processor Configurations Error Summary

Error	Severity	System Action
Processor family not Identical	Fatal	The BIOS detects the error condition and responds as follows: 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	The BIOS detects the error condition and responds as follows: 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.

Error	Severity	System Action
Processor cores/threads not identical	Fatal	The BIOS detects the error condition and responds as follows: 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	The BIOS detects the error condition and responds as follows: 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	The BIOS detects the processor frequency difference, and responds as follows: 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. If the frequencies for all processors cannot be adjusted to be the same**,** then this is an error, and the BIOS responds as follows: 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.
Processor Intel® QuickPath Interconnect link frequencies not identical	Fatal	The BIOS detects the QPI link frequencies and responds as follows: Adjusts all QPI interconnect link frequencies to highest common frequency. No error is generated – this is not an error condition. Continues to boot the system successfully. If the link frequencies for all QPI links cannot be adjusted to be the same**,** then this is an error, and the BIOS responds as follows: 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) QPI 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	The BIOS detects the error condition and responds as follows: 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.

Error	Severity	System Action
Processor microcode update missing	Minor	The BIOS detects the error condition and responds as follows: 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.

3.5 Processor Function Overview

The Intel® Xeon® processor E5-2600 v3 product family combines several key system components into a single processor package, including the CPU cores, Integrated Memory Controller (IMC), and Integrated IO Module (IIO). In addition, each processor package includes two Intel® QuickPath Interconnect point-to-point links capable of up to 9.6 GT/s, up to 40 lanes of PCI express** 3.0 links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI express** 2.0 interface with a peak transfer rate of 4.0 GT/s. The processor supports up to 46 bits of physical address space and 48 bits of virtual address space.

The following sections will provide an overview of the key processor features and functions that help to define the architecture, performance, and supported functionality of the server board. Available features may vary between different processor models.

3.5.1 Processor Core Features:

- Up to 18 execution cores (Intel® Xeon® processor E5-2600 v3 product family)
- When enabled, each core can support two threads (Intel® Hyper-Threading Technology)
- 46-bit physical addressing and 48-bit virtual addressing
- 1 GB large page support for server applications
- A 32 KB instruction and 32 KB data first-level cache (L1) for each core
- A 256 KB shared instruction/data mid-level (L2) cache for each core
- Up to 2.5 MB per core instruction/data last level cache (LLC)

3.5.2 Supported Technologies:

- Intel® Virtualization Technology (Intel® VT) for Intel® 64 and IA-32 Intel® Architecture (Intel® VT-x)
- Intel® Virtualization Technology for Directed I/O (Intel® VT-d)
- Intel® Trusted Execution Technology for servers (Intel® TXT)
- Execute Disable
- Advanced Encryption Standard (AES)
- Intel® Hyper-Threading Technology
- Intel® Turbo Boost Technology
- Enhanced Intel® Speed Step Technology
- Intel® Advanced Vector Extensions 2 (Intel® AVX2)
- Intel® Node Manager 3.0
- Intel® Secure Key
- Intel® OS Guard
- Intel® Quick Data Technology
- Trusted Platform Module (TPM) 1.2

3.5.2.1 Intel® Virtualization Technology (Intel® VT) for Intel® 64 and IA-32 Intel® Architecture (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.

3.5.2.2 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.

3.5.2.3 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.

3.5.2.4 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.

3.5.2.5 Advanced Encryption Standard (AES)

These instructions enable fast and secure data encryption and decryption, using the Advanced Encryption Standard (AES)

3.5.2.6 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.

3.5.2.7 Intel® Turbo Boost Technology

Intel® Turbo Boost Technology is a feature that allows the processor to opportunistically and automatically run faster than its rated operating frequency if it is operating below power, temperature, and current limits. The result is increased performance in multi-threaded and single threaded workloads. It should be enabled in the BIOS for the processor to operate with maximum performance.

3.5.2.8 Enhanced Intel® SpeedStep Technology

The processor supports Enhanced Intel SpeedStep Technology (EIST) as an advanced means of enabling very high performance while also meeting the power conservation needs of the platform.

Enhanced Intel SpeedStep Technology builds upon that architecture using design strategies that include the following:

- Separation between Voltage and Frequency changes. By stepping voltage up and down in small increments separately from frequency changes, the processor is able to reduce periods of system unavailability (which occur during frequency change). Thus, the system is able to transition between voltage and frequency states more often, providing improved power/performance balance.
- Clock Partitioning and Recovery. The bus clock continues running during state transition, even when the core clock and Phase-Locked Loop are stopped, which allows logic to remain active. The core clock is also able to restart more quickly under Enhanced Intel SpeedStep Technology.

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

Intel® Advanced Vector Extensions 2.0 (Intel® AVX2) is the latest expansion of the Intel instruction set. Intel® AVX2 extends the Intel® Advanced Vector Extensions (Intel® AVX) with 256-bit integer instructions, floating-point fused multiply add (FMA) instructions and gather operations. The 256-bit integer vectors benefit math, codec, image and digital signal processing software. FMA improves performance in face detection, professional imaging, and high performance computing. Gather operations increase vectorization opportunities for many applications. In addition to the vector extensions, this generation of Intel processors adds new bit manipulation instructions useful in compression, encryption, and general purpose software.

3.5.2.10 Intel® Node Manager 3.0

Intel® Node Manager 3.0 enables the PTAS-CUPS (Power Thermal Aware Scheduling - Compute Usage Per Second) feature of the Intel Server Platform Services 3.0 Intel ME FW. This is a grouping of separate platform functionalities that provide Power, Thermal, and Utilization data that together offer an accurate, real time characterization of server workload. These functionalities include the following:

- Computation of Volumetric Airflow
- New synthesized Outlet Temperature sensor
- CPU, memory, and I/O utilization data (CUPS).

This PTAS-CUPS data can then be used in conjunction with the Intel® Server Platform Services 3.0 Intel® Node Manager power monitoring/controls and a remote management application (such as the Intel® Data Center Manager [Intel® DCM]) to create a dynamic, automated, closed-loop data center management and monitoring system.

3.5.2.11 Intel® Secure Key

The Intel® 64 and IA-32 Architectures instruction RDRAND and its underlying Digital Random Number Generator (DRNG) hardware implementation is useful for providing large entropy random numbers for which high quality keys for cryptographic protocols are created.

3.5.2.12 Intel® OS Guard

Protects a supported operating system (OS) from applications that have been tampered with or hacked by preventing an attack from being executed from application memory. Intel OS Guard also protects the OS from malware by blocking application access to critical OS vectors.

3.5.2.13 Trusted Platform Module (TPM)

Trusted Platform Module is bound to the platform and connected to the PCH via the LPC bus or SPI bus. The TPM provides the hardware-based mechanism to store or ‘seal’ keys and other data to the platform. It also provides the hardware mechanism to report platform attestations

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

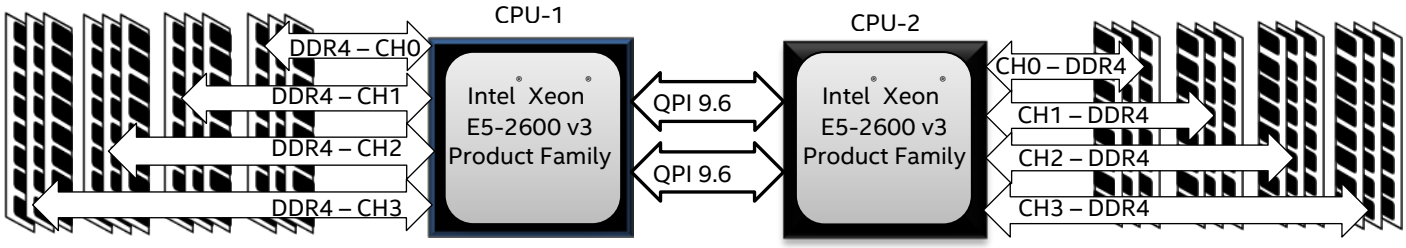


Figure 9. Memory Sub-system Block Diagram

Note: This generation server board has support for DDR4 DIMMs only. DDR3 DIMMs are not supported on this generation server board.

Each installed processor includes two integrated memory controllers (IMC) capable of supporting two memory channels each. Each memory channel is capable of supporting up to three DIMMs. The processor IMC supports the following:

- Registered DIMMs (RDIMMs), and Load Reduced DIMMs (LRDIMMs) are supported
- DIMMs of different types may not be mixed – this is a Fatal Error in memory initialization
- DIMMs composed of 4 Gb or 8 Gb Dynamic Random Access Memory (DRAM) technology
- DIMMs using x4 or x8 DRAM technology
- DIMMs organized as Single Rank (SR), Dual Rank (DR), or Quad Rank (QR)
- DIMM sizes of 4 GB, 8 GB, 16 GB, or 32 GB depending on ranks and technology
- DIMM speeds of 1333, 1600, 1866, or 2133 MT/s (MegaTransfers/second)
- Only Error Correction Code (ECC) enabled RDIMMs or LRDIMMs are supported
- Only RDIMMs and LRDIMMs with integrated Thermal Sensor On Die (TSOD) are supported
- Memory RASM Support:
    - DRAM Single Device Data Correction (SDDCx4)
    - Memory Disable and Map out for FRB
    - Data scrambling with command and address
    - DDR4 Command/Address parity check and retry
    - Intra-socket memory mirroring
    - Memory demand and patrol scrubbing
    - HA and IMC corrupt data containment
    - Rank level memory sparing
    - Multi-rank level memory sparing
    - Failed DIMM isolation

4.2 IMC Modes of operation

A memory controller can be configured to operate in one of two modes, and each IMC operates separately.

Independent Mode: This is also known as performance mode. In this mode each DDR channel is addressed individually via burst lengths of 8 bytes.

- All processors support SECDED ECC with x8 DRAMs in independent mode.
- All processors support SDDC with x4 DRAMs in independent mode.

Lockstep mode: This is also known as RAS mode. Each pair of channels shares a Write Push Logic unit to enable lockstep. The memory controller handles all cache lines across two interfaces on an IMC. The DRAM controllers in the same IMC share a common address decode and DMA engines for the mode. The same address is used on both channels, such that an address error on any channel is detectable by bad ECC.

- All processors support SDDC with x4 or x8 DRAMs in lockstep mode.

For Lockstep Channel Mode and Mirroring Mode, processor channels are paired together as a “Domain”.

- CPU1 Mirroring/Lockstep Domain 1 = Channel A + Channel B
- CPU1 Mirroring/Lockstep Domain 2 = Channel C + Channel D
- CPU2 Mirroring/Lockstep Domain 1 = Channel E + Channel F
- CPU2 Mirroring/Lockstep Domain 2 = Channel G + Channel H

The schedulers within each channel of a domain will operate in lockstep, they will issue requests in the same order and time and both schedulers will respond to an error in either one of the channels in a domain. Lockstep refers to splitting cache lines across channels. The same address is used on both channels, such that an address error on any channel is detectable by bad ECC. The ECC code used by the memory controller can correct 1/18th of the data in a code word. For x8 DRAMs, since there are 9 x8 DRAMs on a DIMM, a code word must be split across 2 DIMMs to allow the ECC to correct all the bits corrupted by a x8 DRAM failure.

For RAS modes that require matching populations, the same slot positions across channels must hold the same DIMM type with regards to number of ranks, number of banks, number of rows, and number of columns. DIMM timings do not have to match but timings will be set to support all DIMMs populated (that is, DIMMs with slower timings will force faster DIMMs to the slower common timing modes).

4.3 Memory RASM Features

DRAM Single Device Data Correction (SDDC): SDDC provides error checking and correction that protects against a single x4 DRAM device failure (hard-errors) as well as multibit faults in any portion of a single DRAM device on a DIMM (require lockstep mode for x8 DRAM device based DIMM).

Memory Disable and Map out for FRB: Allows memory initialization and booting to OS even when a memory fault occurs.

Data Scrambling with Command and Address: Scrambles the data with address and command in "write cycle" and unscrambles the data in "read cycle". This feature addresses reliability by improving signal integrity at the physical layer, and by assisting with detection of an address bit error.

DDR4 Command/Address Parity Check and Retry: DDR4 technology based CMD/ADDR parity check and retry with following attributes:

- CMD/ADDR Parity error address logging
- CMD/ADDR Retry

Intra-Socket Memory Mirroring: Memory Mirroring is a 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. Dynamic (without reboot) failover to the mirrored DIMMs is transparent to the OS and applications. Note that with Memory Mirroring enabled, only half of the memory capacity of both memory channels is available.

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. It prevents accumulation of single-bit errors.

HA and IMC Corrupt Data Containment: Corrupt Data Containment is a process of signaling memory patrol scrub uncorrected data errors synchronous to the transaction, which enhances the containment of the fault and improving the reliability of the system.

Rank Level / Multi Rank Level Memory Sparing: Dynamic fail-over of failing ranks to spare ranks behind the same memory controller. With Multi Rank, up to four ranks out of a maximum of eight ranks can be assigned as spare ranks. Memory mirroring is not supported when memory sparing is enabled.

Failed DIMM Isolation: The 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.

4.4 Supported Memory

Table 4. DDR4 RDIMM & LRDIMM Support

Type	Ranks Per DIMM and Data Width	DIMM Capacity (GB)		Speed (MT/s); Voltage (V); Slot per Channel (SPC) and DIMM per Channel (DPC)		
				3 Slots per Channel		
				1 DPC	2 DPC	3 DPC
		4 Gb	8 Gb	1.2V	1.2V	1.2V
RDIMM	SRx4	8GB	16GB	2133	1866	1600
RDIMM	SRx8	4GB	8GB	2133	1866	1600
RDIMM	DRx8	8GB	16GB	2133	1866	1600
RDIMM	DRx4	16GB	32GB	2133	1866	1600
LRDIMM	QRx4	32GB	64GB	2133	2133	1600

4.5 NVDIMM Support

Future enhancement

4.6 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.

- Each installed processor provides four channels of memory. On the Intel® Server Board S2600WT each memory channel supports three memory slots, for a total possible 24 DIMMs installed.
- System memory is organized into physical slots on DDR4 memory channels that belong to processor sockets.
- The memory channels from processor socket 1 are identified as Channel A, B, C and D. The memory channels from processor socket 2 are identified as Channel E, F, G, and H.
- Each memory slot on the server board is identified by channel and slot number within that channel. For example, DIMM_A1 is the first slot on Channel A on processor 1; DIMM_E1 is the first DIMM socket on Channel E on processor 2.
- 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.
- The BLUE memory slots on the server board identify the first memory slot for a given memory channel.

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. In addition, 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.

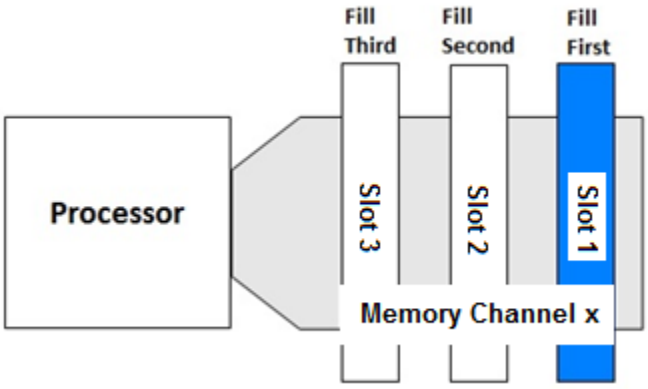


Figure 10. Memory Slots Definition

On the Intel® Server Board S2600WT, a total of 24 DIMM slots is provided – 2 CPUs, 4 Memory Channels/CPU, 3 DIMMs/Channel. The nomenclature for memory slots is detailed in the following table:

Table 5. Intel® Server Board S2600WT Memory Slot Identification

Processor Socket 1												Processor Socket 2											
(0) Channel A			(1) Channel B			(2) Channel C			(3) Channel D			(0) Channel E			(1) Channel F			(2) Channel G			(3) Channel H		
A1	A2	A3	**B1**	B2	B3	**C1**	C2	C3	**D1**	D2	D3	**E1**	E2	E3	**F1**	F2	F3	**G1**	G2	G3	**H1**	H2	H3

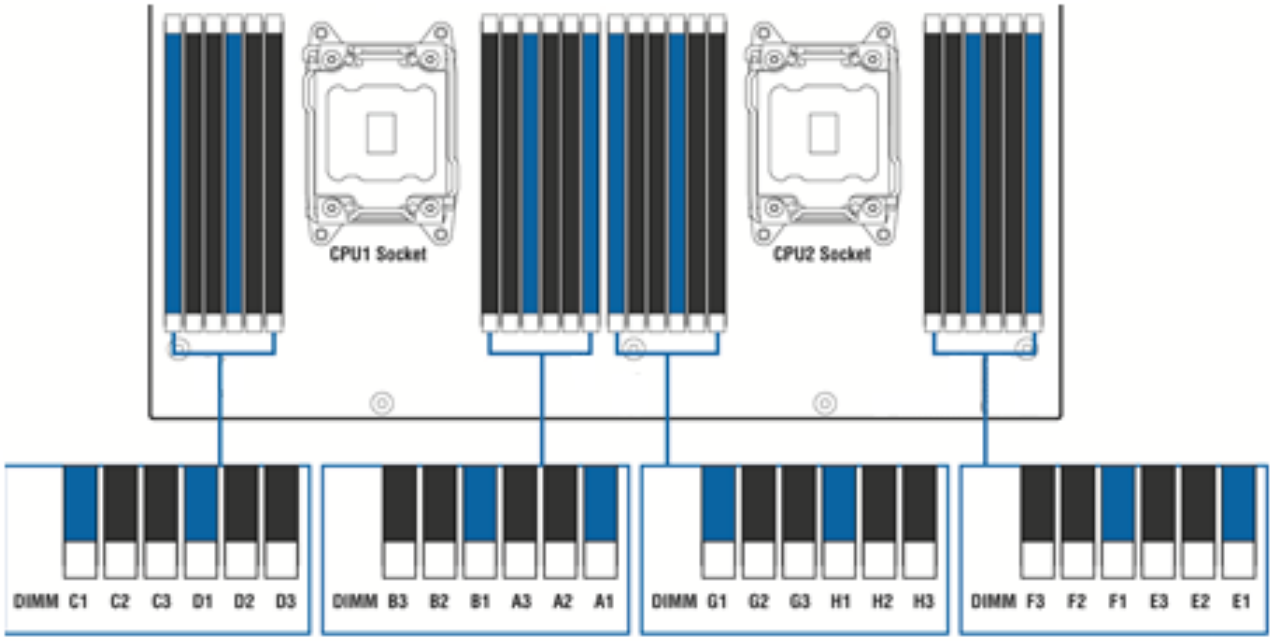


Figure 11. Intel® Server Board S2600WT Memory Slot Layout

The following are the DIMM population requirements

- All DIMMs must be DDR4 DIMMs
- Only Error Correction Code (ECC) enabled RDIMMs and LRDIMMs are supported
- Only RDIMMs and LRDIMMs with integrated on die thermal sensors (TROD) are supported
- DIMM slots on any memory channel must be filled following the “farthest fill first” rule.
- The DIMM slot farthest away from the processor socket must be filled first on any channel. This will always be designated on the board as Slot 1 for the channel.
- When one DIMM is used, it must be populated in the BLUE DIMM slot (farthest away from the CPU) of a given channel.
- A maximum of 8 ranks can be installed on any one channel, counting all ranks in each DIMM on the channel.
- DIMM types (RDIMM, LRDIMM) must not be mixed within or across processor sockets. This is a Fatal Error Halt in Memory Initialization.
- Mixing DIMMs of different frequencies and latencies is not supported within or across processor sockets. If a mixed configuration is encountered, the BIOS will attempt to operate at the highest common frequency and the lowest latency possible.

- LRDIMM Rank Multiplication Mode and Direct Map Mode must not be mixed within or across processor sockets. This is a Fatal Error Halt in Memory Initialization.
- 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.
- RAS Modes Lockstep, Rank Sparing, and Mirroring are mutually exclusive in this BIOS. Only one operating mode may be selected, and it will be applied to the entire system.
- If a RAS Mode has been configured, and the memory population will not 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 meet the requirement of having at least 2 SR or DR DIMM installed, or at least one QR DIMM installed, on each populated channel.
- Lockstep or Mirroring Modes require that for any channel pair that is populated with memory, the memory population on both channels of the pair must be identically sized.

The following table identifies possible DIMM population configurations

Table 6. DIMM Population Matrix

# of DIMMs	Processor Socket 1 = Populated												Processor Socket 2 = Populated												M
	A1	A2	A3	B1	B2	B3	C1	C2	C3	D1	D2	D3	E1	E2	E3	F1	F2	F3	G1	G2	G3	H1	H2	H3	
1	X																								N
2	X	X																							N
2	X			X																					Y
2	X												X												N
3	X	X	X																						N
3	X			X			X																		N
3	X			X									X												N
3	X												X			X									N
4	X	X		X	X																				Y
4	X	X		X			X																		N
4	X			X			X			X															Y
4	X			X									X			X									Y
4	X	X											X	X											N
4													X	X		X	X								Y
5	X	X		X	X		X																		N
5	X			X			X			X			X												N
6	X	X	X	X	X	X																			Y
6							X	X	X	X	X	X													Y
6	X	X		X	X		X	X																	N
6	X			X			X			X			X			X									N
8	X	X		X	X		X	X		X	X														Y
8	X	X		X	X								X	X		X	X								Y
8	X			X			X			X			X			X			X			X			Y
8	X	X	X	X	X	X	X	X																	N
12	X	X	X	X	X	X	X	X	X	X	X	X													Y
12	X	X		X	X		X	X		X	X		X	X		X	X								N

# of DIMMs	Processor Socket 1 = Populated												Processor Socket 2 = Populated												M
	A1	A2	A3	B1	B2	B3	C1	C2	C3	D1	D2	D3	E1	E2	E3	F1	F2	F3	G1	G2	G3	H1	H2	H3	
16	X	X		X	X		X	X		X	X		X	X		X	X		X	X		X	X		Y
16	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X									N
24	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Y

M – Indicates whether the configuration supports the Mirrored Channel Mode of operation.

4.6.1 Memory Interleaving Support

The Intel® Xeon® Processor E5-4600/2600/2400/1600 v3 Product Families support multiple levels of memory interleaving. Memory interleaving is an optimization technique which tries to locate successive data across different memory channels, to allow for overlapping memory access.

The processors and BIOS support inter-socket interleaving across 1, 2, or 4 processor sockets, channel interleaving across 1, 2, 3, or 4 memory channels per processor, and rank interleaving in 1, 2, 4, and 8 way arrangements.

The BIOS will choose an interleave scheme based on the processor population and the DIMM population. If the NUMA option is enabled, then all interleaving is strictly intra-socket to allow for locality to be controlled by the OS. The actual locality is described in ACPI Tables.

4.6.2 NUMA Configuration Support

This BIOS includes support for Non-Uniform Memory Access (NUMA) when more than one processor is installed in a board or one Cluster-on-Die (COD) capable processor installed.

When NUMA support is enabled, interleaving is intra socket only, and the SRAT and SLIT ACPI tables are provided that show the locality of systems resources, especially memory, which allows a “NUMA Aware” OS to optimize which processor threads are used by processes which can benefit by having the best access to those resources.

NUMA support and COD support are enabled/disabled (enabled by default) by an option on the Memory RAS and Performance screen in BIOS setup.

4.7 System Memory Sizing and Publishing

The address space configured in a system depends on the amount of actual physical memory installed, on the RAS configuration, and on the PCIe* configuration. RAS configurations reduce the memory space available in return for the RAS features. PCIe* devices which require address space for Memory Mapped IO (MMIO) with 32-bit or 64- bit addressing, increase the address space in use, and introduce discontinuities in the correspondence between physical memory and memory addresses.

The discontinuities in addressing physical memory revolve around the 4GB 32-bit addressing limit. Since the system reserves memory address space just below the 4GB limit, and 32-bit MMIO is allocated just below that, the addresses assigned to physical memory go up to the bottom of the PCI allocations, then “jump” to above the 4GB limit into 64-bit space. See the comments below about Memory reservations.

4.7.1 Effects of Memory Configuration on Memory Sizing

The system BIOS supports 4 memory configurations – Independent Channel Mode and 3 different RAS Modes. In some modes, memory reserved for RAS functions reduce the amount of memory available.

- **Independent Channel mode:** In Independent Channel Mode, the amount of installed physical memory is the amount of effective memory available. There is no reduction.
- **Lockstep Mode:** For Lockstep Mode, the amount of installed physical memory is the amount of effective memory available. There is no reduction. Lockstep Mode only changes the addressing to address two channels in parallel.
- **Rank Sparing Mode:** In Rank Sparing mode, the largest rank on each channel is reserved as a spare rank for that channel. This reduces the available memory size by the sum of the sizes of the reserved ranks.

  Example: if a system has 2 16GB Quad Rank DIMMS on each of 4 channels on each of 2 processor sockets, the total installed memory will be (((2 * 16GB) * 4 channels) * 2 CPU sockets) = 256GB.

  For a 16GB QR DIMM, each rank would be 4GB. With one rank reserved on each channel, that would 32GB reserved. So the available effective memory size would be 256GB - 32GB, or 224GB.

- **Mirroring Mode:** Mirroring creates a duplicate image of the memory that is in use, which uses half of the available memory to mirror the other half. This reduces the available memory size to half of the installed physical memory.

  Example: if a system has 2 16GB Quad Rank DIMMS on each of 4 channels on each of 2 processor sockets, the total installed memory will be (((2 * 16GB) * 4 channels) * 2 CPU sockets) = 256GB.

  In Mirroring Mode, since half of the memory is reserved as a mirror image, the available memory size would be 128GB.

4.7.2 Publishing System Memory

There are a number of different situations in which the memory size and/or configuration are displayed. Most of these displays differ in one way or another, so the same memory configuration may appear to display differently, depending on when and where the display occurs.

- The BIOS displays the “Total Memory” of the system during POST if Quiet Boot is disabled in BIOS setup. This is the total size of memory discovered by the BIOS during POST, and is the sum of the individual sizes of installed DDR4 DIMMs in the system.
- The BIOS displays the “Effective Memory” of the system in the BIOS Setup. The term Effective Memory refers to the total size of all DDR4 DIMMs that are active (not disabled) and not used as redundant units (see Note below).
- The BIOS provides the total memory of the system in the main page of BIOS setup. This total is the same as the amount described by the first bullet above.
- If Quiet Boot is disabled, the BIOS displays the total system memory on the diagnostic screen at the end of POST. This total is the same as the amount described by the first bullet above.
- The BIOS provides the total amount of memory in the system by supporting the EFI Boot Service function, GetMemoryMap().
- The BIOS provides the total amount of memory in the system by supporting the INT 15h, E820h function. For details, see the Advanced Configuration and Power Interface Specification.

Note: Some server operating systems do not display the total physical memory installed. What is displayed is the amount of physical memory minus the approximate memory space used by system BIOS components. These BIOS components include but are not limited to:

- ACPI (may vary depending on the number of PCI devices detected in the system and the size of memory included on them)
- ACPI NVS table
- Processor microcode

- Memory Mapped I/O (MMIO)
- Manageability Engine (ME)
- BIOS flash

4.8 Memory Initialization

Memory Initialization at the beginning of POST includes multiple functions, including:

- DIMM discovery
- Channel training
- DIMM population validation check
- Memory controller initialization and other hardware settings
- Initialization of RAS configurations (as applicable)

There are several errors which can be detected in different phases of initialization. During early POST, before system memory is available, serious errors that would prevent a system boot with data integrity will cause a System Halt with a beep code and a memory error code to be displayed via the POST Code Diagnostic LEDs.

Less fatal errors will cause a POST Error Code to be generated as a Major Error. This POST Error Code will be displayed in the BIOS Setup Error Manager screen, and will also be logged to the System Event Log (SEL).

4.8.1 DIMM Discovery

Memory initialization begins by determining which DIMM slots have DIMMs installed in them. By reading the Serial Presence Detect (SPD) information from an SEEPROM on the DIMM, the type, size, latency, and other descriptive parameters for the DIMM can be acquired.

Potential Error Cases:

- Memory is locked by Intel® TXT and is inaccessible – This will result in a Fatal Error Halt 0xE9.
- DIMM SPD does not respond – The DIMM will not be detected, which could result in a “No usable memory installed” Fatal Error Halt 0xE8 if there are no other detectable DIMMs in the system. The undetected DIMM could result later in an invalid configuration if the “no SPD” DIMM is in Slot 1 or 2 ahead of other DIMMs on the same channel.
- DIMM SPD read error – This DIMM will be disabled. POST Error Codes 856x “SPD Error” and 854x “DIMM Disabled” will be generated. If all DIMMs are failed, this will result in a Fatal Error Halt 0xE8.
- All DIMMs on the channel in higher-numbered sockets behind the disabled DIMM will also be disabled with a POST Error Code 854x “DIMM Disabled” for each. This could also result in a “No usable memory installed” Fatal Error Halt 0xE8.
- No usable memory installed – If no usable (not failed or disabled) DIMMs can be detected as installed in the system, this will result in a Fatal Error Halt 0xE8. Other error conditions which cause DIMMs to fail or be disabled so they are mapped out as unusable may result in causing this error when no usable DIMM remains in the memory configuration.

4.8.2 DIMM Population Validation Check

Once the DIMM SPD parameters have been read they are checked to verify that the DIMMs on the given channel are installed in a valid configuration. This includes checking for DIMM type, DRAM type and organization, DRAM rank organization, DIMM speed and size, ECC capability, and in which memory slots the DIMMs are installed. An invalid configuration may cause the system to halt.

Potential Error Cases:

- **Invalid DIMM (type, organization, speed, size)** – If a DIMM is found that is not a type supported by the system, the following error will be generated: POST Error Code 8501 "DIMM Population Error", and a "Population Error- Fatal Error Halt 0xED".
- **Invalid DIMM Installation** – The DIMMs are installed incorrectly on a channel, not following the "Fill Farthest First" rule (Slot 1 must be filled before Slot 2, Slot 2 before Slot 3). This will result in a POST Error Code 8501 "DIMM Population Error" with the channel being disabled, and all DIMMs on the channel will be disabled with a POST Error Code 854x "DIMM Disabled" for each. This could also result in a "No usable memory installed" Fatal Error Halt 0xE8.
- **Invalid DIMM Population** – A QR LRDIMM in Direct Map mode which is installed in Slot3 on a 3 DIMM per channel server board is not allowed. This will result in a POST Error Code 8501 "DIMM Population Error" and a "Population Error" Fatal Error Halt 0xED.
- **Mixed DIMM Types** – A mixture of RDIMMs and/or LRDIMMs is not allowed. A mixture of LRDIMMs operating in Direct Map mode and Rank Multiplication mode is also not allowed. This will result in a POST Error Code 8501 "DIMM Population Error" and "Population Error" Fatal Error Halt 0xED.
- **Mixed DIMM Parameters** – Within an RDIMM or LRDIMM configuration, mixtures of valid DIMM technologies, sizes, speeds, latencies, etc., although not supported, will be initialized and operated on a best effort basis, if possible.
- **No usable memory installed** – If no enabled and available memory remains in the system, this will result in a Fatal Error Halt 0xE8.

4.8.3 Channel Training

The Integrated Memory Controller registers are programmed at the controller level and the memory channel level. Using the DIMM operational parameters, read from the SPD of the DIMMs on the channel, each channel is trained for optimal data transfer between the integrated memory controller (IMC) and the DIMMs installed on the given channel.

Potential Error Cases:

- Channel Training Error – If the Data/Data Strobe timing on the channel cannot be set correctly so that the DIMMs can become operational, this results in a momentary Error Display 0xEA, and the channel is disabled. All DIMMs on the channel are marked as disabled, with POST Error Code 854x "DIMM Disabled" for each. If there are no populated channels which can be trained correctly, this becomes a Fatal Error Halt 0xEA.

4.8.3.1 Thermal (CLTT) and power throttling

Potential Error Cases:

- CLTT Structure Error – The CLTT initialization fails due to an error in the data structure passed in by the BIOS. This results in a *Fatal Error Halt **0xEF***. See chapter 7 for information describing CLTT.

4.8.3.2 Built-In Self Test (BIST)

Once the memory is functional, a memory test is executed. This is a hardware-based Built In Self Test (BIST) which confirms minimum acceptable functionality. Any DIMMs which fail are disabled and removed from the configuration.

Potential Error Cases:

- Memory Test Error – The DIMM has failed BIST and is disabled. POST Error Codes **852x** "Failed test/initialization" and **854x** "DIMM Disabled" will be generated for each DIMM that fails. Any DIMMs

installed on the channel behind the failed DIMM will be marked as disabled, with POST Error Code **854x** "DIMM Disabled". This results in a momentary Error Display **0xEB**, and if all DIMMs have failed, this will result in a Fatal Error Halt **0xE8**.

- No usable memory installed – If no enabled and available memory remains, this will result in a Fatal Error Halt **0xE8**.

The ECC functionality is enabled after all of memory has been cleared to zeroes to make sure that the data bits and the ECC bits are in agreement.

4.8.3.3 RAS Mode Initialization

If configured, the DIMM configuration is validated for the specified RAS mode. If the enabled DIMM configuration is compliant for the RAS mode selected, then the appropriate register settings are set and the RAS mode is started.

Potential Error Cases:

- RAS Configuration Failure – If the DIMM configuration is not valid for the RAS mode which was selected, then the operating mode falls back to Independent Channel Mode, and a POST Error Code **8500** "Selected RAS Mode could not be configured" is generated. In addition, a "RAS Configuration Disabled" SEL entry for "RAS Configuration Status" (BIOS Sensor 02/Type 0Ch/Generator ID 01) is logged.

5. 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® E5-2600 v3 processor family, the Intel® C612 chipset, the Intel® Ethernet controller I350 or X540, and the I/O controllers embedded within the Emulex* Pilot-III Management Controller.

See *Figure 6. Intel® Server Board S2600WT Architectural Block Diagram* for an overview of the features and interconnects of each of the major sub-system components

5.1 PCIe* Support

The Integrated I/O (IIO) module of the Intel® Xeon® Processor E5-2600 v3 product family provides the PCI express* interface for general purpose PCI express* devices at up to PCI express* 3.0 speeds.

The IIO module provides the following PCIe* Features:

- Compliant with the PCI express* Base Specification, Revision 2.0 and Revision 3.0
- 2.5 GHz (Gen1) and 5 GHz (Gen2) and 8 GHz (Gen3)
- x16 PCI-Express 3.0 interface supports up to four x4 controllers and is configurable to 4x4 links, 2x8, 2x4\1x8, or 1x16
- x8 PCI-Express 3.0 interface supports up to 2 x4 controllers and is configurable to 2x4 or 1x8
- Full peer-to-peer support between PCI express* interfaces
- Full support for software-initiated PCI express* power management
- x8 Server I/O Module support
- TLP Processing Hints (TPH) for data push to cache
- Address Translation Services (ATS 1.0)
- PCIe* Atomic Operations Completer Capability
- Autonomous Linkwidth
- x4 DMI2 interface
    - All processors support a x4 DMI2 lane which can be connected to a PCH, or operate as a x4 PCIe* 2.0 port.

The following tables provide the PCIe* port routing information:

Table 7. PCIe* Port Routing CPU #1

CPU 1			
PCI Ports	Device (D)	Function (F)	On-board Device
Port DMI 2/PCIe* x4	0		Chipset
Port 1A - x4	D1	F0	SAS Module
Port 1B - x4	D1	F1	SAS Module
Port 2A - x4	D2	F0	IO Module
Port 2B - x4	D2	F1	IO Module
Port 2C - x4	D2	F2	NIC - I350/X540
Port 2D - x4	D2	F3	NIC - I350/X540
Port 3A - x4	D3	F0	Riser Slot #1
Port 3B - x4	D3	F1	Riser Slot #1
Port 3C - x4	D3	F2	Riser Slot #1
Port 3D -x4	D3	F3	Riser Slot #1

Table 8. PCIe* Port Routing – CPU #2

CPU 2			
PCI Ports	Device (D)	Function (F)	On-board Device
Port DMI 2/PCIe* x4	0	F0	Riser Slot #3
Port 1A - x4	D1	F1	Riser Slot #1
Port 1B - x4	D1	F0	Riser Slot #1
Port 2A - x4	D2	F0	Riser Slot #2
Port 2B - x4	D2	F1	Riser Slot #2
Port 2C - x4	D2	F2	Riser Slot #2
Port 2D - x4	D2	F3	Riser Slot #2
Port 3A - x4	D3	F0	Riser Slot #3
Port 3B - x4	D3	F1	Riser Slot #3
Port 3C - x4	D3	F2	Riser Slot #2
Port 3D -x4	D3	F3	Riser Slot #2

Note: See section 5.4.1 for details of root port to PCIe* slot mapping for each supported riser card.

5.2 PCIe* Enumeration and Allocation

The BIOS assigns PCI bus numbers in a depth-first hierarchy, in accordance with the PCI Local Bus Specification, Revision 2.2. 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.

5.3 PCIe* Non-Transparent Bridge (NTB)

PCI express* Non-Transparent Bridge (NTB) acts as a gateway that enables high performance, low overhead communication between two intelligent subsystems, the local and the remote subsystems. The NTB allows a local processor to independently configure and control the local subsystem, 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 PCI express* Port 3A of Intel® Xeon® Processor E5-2600 v3 Product Families can be configured to be a transparent bridge or an NTB with x4/x8/x16 link width and Gen1/Gen2/Gen3 link speed. This NTB port could be attached to another NTB port or PCI express* Root Port on another subsystem. NTB supports three 64bit BARs as configuration space or prefetchable memory windows that can access both 32bit and 64bit address space through 64bit BARs.

There are 3 NTB supported configurations:

- NTB Port to NTB Port Based Connection (Back-to-Back)
- NTB Port to Root Port Based Connection – Symmetric Configuration. The NTB port on the first system is connected to the root port of the second. The second system's NTB port is connected to the root port on the first system making this a fully symmetric configuration.
- NTB Port to Root Port Based Connection – Non-Symmetric Configuration. The root port on the first system is connected to the NTB port of the second system. It is not necessary for the first system to be an Intel® Xeon® Processor E5-2600 v3 Product Families system.

Note: When NTB is enabled, Spread Spectrum Clocking (SSC) is required to be disabled at each NTB link.

5.4 Add-in Card Support

The server board includes features for concurrent support of several add-in card types including: PCIe* add-in cards via three riser card slots, Intel® I/O module options via a proprietary high density 80 pin connector, and Intel® Integrated RAID Modules via a proprietary high density 80 pin connector. The following illustration identifies the location of the onboard connector features and general board placement for add-in modules and riser cards.

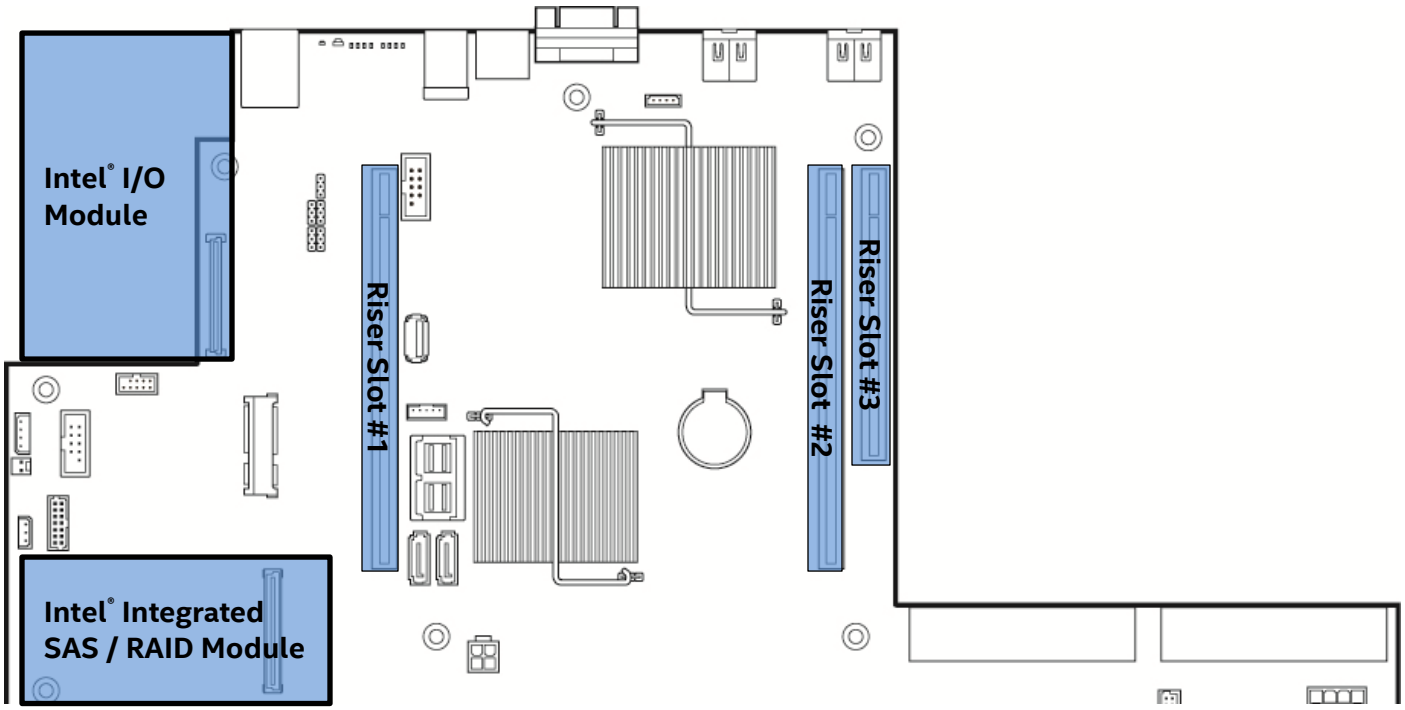


Figure 12. On-board Add-in Card Support

5.4.1 Riser Card Support

The server board provides three riser card slots identified as: Riser Slot #1, Riser Slot #2, and Riser Slot #3.

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.

The PCIe* bus interface for each riser card slot is supported by each of the two installed processors. The following tables provide the PCIe* bus routing for all supported risers cards.

Note: The bottom add-in card slot of 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.

Table 9. Riser Card #1 - PCIe* Root Port Mapping

Riser Slot #1 – Riser Card Options			
2U – PCI-X Riser Card iPN - A2ULPCIXRISER2	2U - 3-Slot Riser Card iPN – A2UL8RISER2	2U - 2-Slot Riser Card iPN – A2UL16RISER2	1U - 1-Slot Riser Card iPN – F1UL16RISER2
Top PCI-X Slot CPU#1 – Port 3C (133MHz – Independent bus)	Top PCIe* Slot CPU #1 – Port 3C (x8 elec, x16 mech)	Top PCIe* Slot CPU #1 – Port 3A (x16 elec, x16 mech)	PCIe* Slot CPU #1 – Port 3A (x16 elec, x16 mech)
Middle PCI-X Slot CPU #2 – Port 1B (133MHz – Independent bus)	Middle PCIe* Slot CPU #1 – Port 3A (x8 elec, x16 mech)		
Bottom PCIe* Slot CPU #1 – Port 3A (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Port 1A (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Port 1A (x8 elec, x8 mech)	

Table 10. Riser Card #2 - PCIe* Root Port Mapping

Riser Slot #2 – Riser Card Options			
2U – PCI-X Riser Card iPN - A2ULPCIXRISER2	2U - 3-Slot Riser Card iPN – A2UL8RISER2	2U - 2-Slot Riser Card iPN – A2UL16RISER2	1U - 1-Slot Riser Card iPN – F1UL16RISER2
Top PCI-X Slot CPU #2 – Port 2C (133MHz – Independent bus)	Top PCIe* Slot CPU #2 – Port 2C (x8 elec, x16 mech)	Top PCIe* Slot CPU #2 – Port 2A (x16 elec, x16 mech)	Top PCIe* Slot CPU #2 – Port 2A (x16 elec, x16 mech)
Middle PCI-X Slot CPU #2 – Port 3C (133MHz - Independent bus)	Middle PCIe* Slot CPU #2 – Port 2A (x8 elec, x16 mech)		
Bottom PCIe* Slot CPU #2 – Port 2A (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Port 3C (x8 elec, x8 mech)	Bottom PCIe* Slot CPU #2 – Port 3C (x8 elec, x8 mech)	

Table 11. Riser Slot #3 - PCIe* Root Port Mapping

Riser Slot #3 - Riser Card Options	
2U - Low Profile Riser Card iPN – A2UX8X4RISER	Notes
Top PCIe* Slot CPU #2 – Port DMI 2 (x4 elec, x8 mech)	PCIe* 2.0 Support Only
Bottom PCIe* Slot CPU #2 – Port 3A (x8 elec, x8 mech)	

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

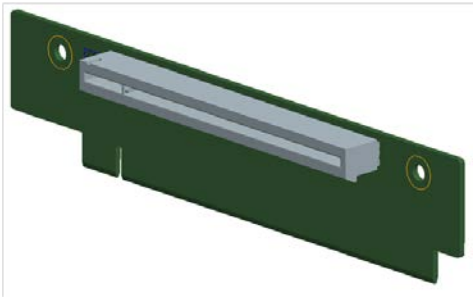
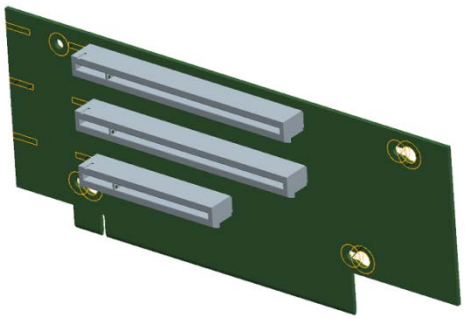


Figure 13. 1U one slot PCIe* riser card (iPN – F1UL16RISER2)

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

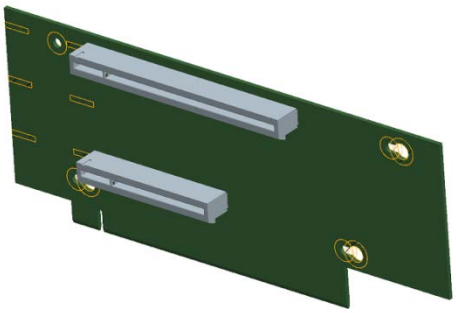


Slot #	Description
Slot-1 (Top)	PCIe* x8 elec, x16 mechanical
Slot-2 (Middle)	PCIe* x8 elec, x16 mechanical
Slot-3 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 14. 2U three PCIe* slot riser card (iPN – 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

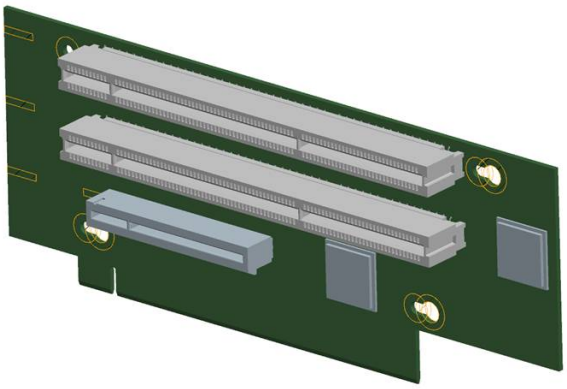


Slot #	Description
Slot-1 (Top)	PCIe* x16 elec, x16 mechanical
Slot-2 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 15. 2U two PCIe* slot riser card (iPN – 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).

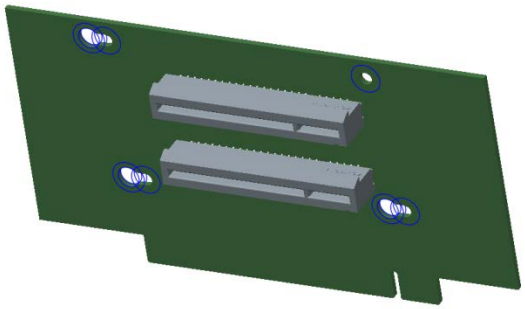
2U – Three add-in card slots(2x PCI-X (133MHz) + 1x PCIe*) –
Note: Due to a mechanical interference, concurrent installation of the PCI-X Riser card in Riser Slot #1 and an on-board SAS module is not supported.



Slot #	Description
Slot-1 (Top)	PCI-X 133 MHz
Slot-2 (Middle)	PCI-X 133 MHz
Slot-3 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 16. 2U three PCIx/PCIe* slot riser (iPN - A2ULPCIXRISER2)

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 (PCIe* 2.0 support only)
Slot-2 (Bottom)	PCIe* x8 elec, x8 mechanical

Figure 17. 2U two PCIe* slot (Low Profile) PCIe* Riser card (iPN – A2UX8X4RISER) – Riser Slot #3 compatible only

5.4.2 I/O Module Support

To broaden the standard on-board feature set, the server board provides support for one of several available Intel I/O Module options. The I/O module attaches to a high density 80-pin connector on the server board labeled “IO_Module” and is supported by x8 PCIe* 3.0 signals from the IIO module of the CPU 1 processor.

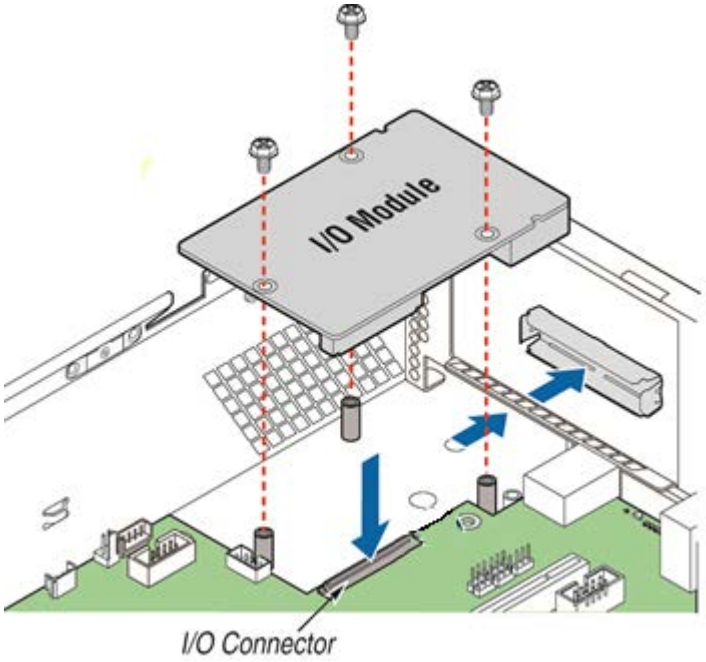


Figure 18. Server Board Layout - I/O Module Connector

Supported I/O modules include:

Table 12. Supported Intel® I/O Module Options

Description	Intel Product Code
Quad Port Intel® I350 GbE I/O Module	AXX4P1GBPWLIOM
Dual Port Intel® X540 10GbE I/O Module	AXX10GBTWLIOM
Dual Port Intel® 82599 10GbE I/O Module	AXX10GBNIAIOM
Single Port FDR InfiniBand* ConnectX*-3 I/O Module	AXX1FDRIBIOM
Dual Ports FDR InfiniBand* ConnectX*-3 I/O Module	AXX2FDRIBIOM
Single port 40GbE I/O Module	AXX1P40FRTIOM
Dual Port 40GbE I/O Module	AXX2P40FRTIOM

5.4.3 Intel® Integrated RAID Option

The server board provides support for Intel® Integrated RAID modules. These optional modules attach to a high density 80-pin connector labeled “SAS Module” on the server board and are supported by x8 PCIe* 3.0 signals from the IIO module of the CPU 1 processor.

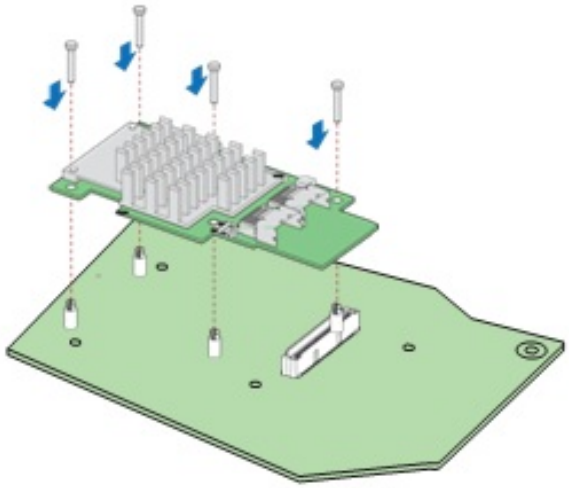


Figure 19. Server Board Layout – Intel® Integrated RAID Module Option Placement

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

https://serverconfigurator.intel.com

5.5 Serial ATA (SATA) Support

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

The AHCI ***SATA*** controller provides support for up to six SATA ports on the server board

- Four SATA ports from the Mini-SAS HD (SFF-8643) connector labeled "**SATA Ports 0-3**" on the server board
- Two SATA ports accessed via two white single port connectors labeled "**SATA-4**" and "**SATA-5**" on the server board

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

- Four SATA ports from the Mini-SAS HD (SFF-8643) connector labeled "**sSATA Ports 0-3**" on the server board

The following diagram identifies the location of all on-board SATA features.

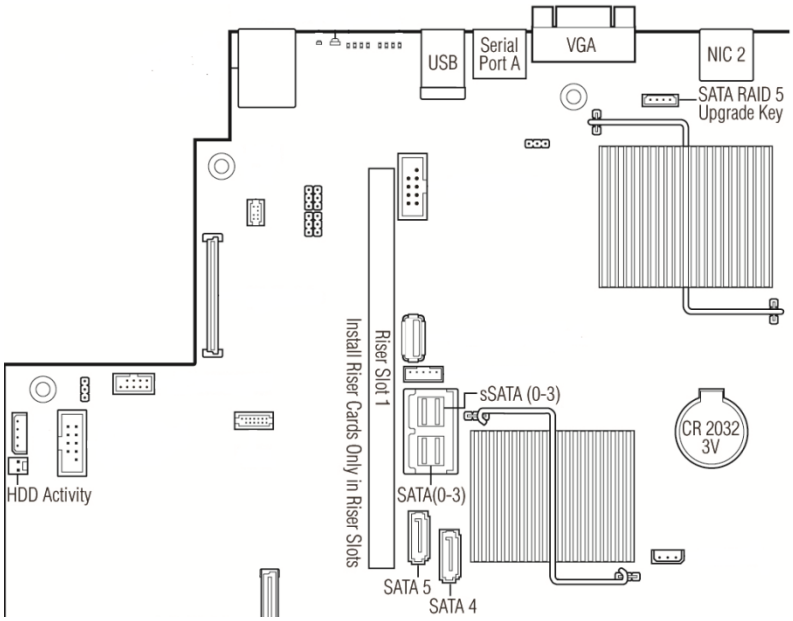


Figure 20. Onboard SATA Features

The SATA controller and the sSATA controller can be independently enabled and disabled and configured through the <F2> BIOS Setup Utility under the "Mass Storage Controller Configuration" menu screen. The following table identifies supported setup options.

Table 13. SATA and sSATA Controller BIOS Utility Setup Options

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

Table 14. SATA and sSATA Controller Feature Support

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

Feature	Description	AHCI / RAID Disabled	AHCI / RAID Enabled
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

5.5.1 Staggered Disk Spin-Up

Because of the high density of disk drives that can be attached to the C612 Onboard AHCI SATA Controller and the sSATA Contoller, the combined startup power demand surge for all drives at once can be much higher than the normal running power requirements and could require a much larger power supply for startup than for normal operations.

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

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

5.6 Embedded SATA SW-RAID support

The server board has embedded support for two SATA SW-RAID options:

- Intel® Rapid Storage Technology (RSTe) 4.0
- Intel® Embedded Server RAID Technology 2 (ESRT2) based on LSI* MegaRAID SW RAID technology

Using the <F2> BIOS Setup Utility, accessed during system POST, options are available to enable/disable SW RAID, and select which embedded software RAID option to use.

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

5.6.1 Intel® Rapid Storage Technology (RSTe) 4.0

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.

- **RAID Level 0** – Non-redundant striping of drive volumes with performance scaling of up to 6 drives, enabling higher throughput for data intensive applications such as video editing.
- Data security is offered through **RAID Level 1**, which performs mirroring.
- **RAID Level 10** provides high levels of storage performance with data protection, combining the fault-tolerance of RAID Level 1 with the performance of RAID Level 0. By striping RAID Level 1 segments, high I/O rates can be achieved on systems that require both performance and fault-tolerance. RAID Level 10 requires 4 hard drives, and provides the capacity of two drives.
- **RAID Level 5** provides highly efficient storage while maintaining fault-tolerance on 3 or more drives. By striping parity, and rotating it across all disks, fault tolerance of any single drive is achieved while only consuming 1 drive worth of capacity. That is, a 3 drive RAID 5 has the capacity of 2 drives, or a 4 drive RAID 5 has the capacity of 3 drives. RAID 5 has high read transaction rates, with a medium write

rate. RAID 5 is well suited for applications that require high amounts of storage while maintaining fault tolerance.

Note: RAID configurations cannot span across the two embedded AHCI SATA controllers.

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 SW-RAID option must be enable 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 levels 5 or 10

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

5.6.2 Intel® Embedded Server RAID Technology 2 (ESRT2)

Features of ESRT2 include the following:

- Based on LSI* MegaRAID Software Stack
- Software RAID with system providing memory and CPU utilization
- **RAID Level 0** - Non-redundant striping of drive volumes with performance scaling up to 6 drives, enabling higher throughput for data intensive applications such as video editing.
- Data security is offered through **RAID Level 1**, which performs mirroring.
- **RAID Level 10** provides high levels of storage performance with data protection, combining the fault-tolerance of RAID Level 1 with the performance of RAID Level 0. By striping RAID Level 1 segments, high I/O rates can be achieved on systems that require both performance and fault-tolerance. RAID Level 10 requires 4 hard drives, and provides the capacity of two drives
- Optional support for **RAID Level 5**
    - Enabled with the addition of an optionally installed ESRT2 SATA RAID 5 Upgrade Key (**iPN – RKSATA4R5**)
    - **RAID Level 5** provides highly efficient storage while maintaining fault-tolerance on 3 or more drives. By striping parity, and rotating it across all disks, fault tolerance of any single drive is achieved while only consuming 1 drive worth of capacity. That is, a 3 drive RAID 5 has the capacity of 2 drives, or a 4 drive RAID 5 has the capacity of 3 drives. RAID 5 has high read transaction rates, with a medium write rate. RAID 5 is well suited for applications that require high amounts of storage while maintaining fault tolerance

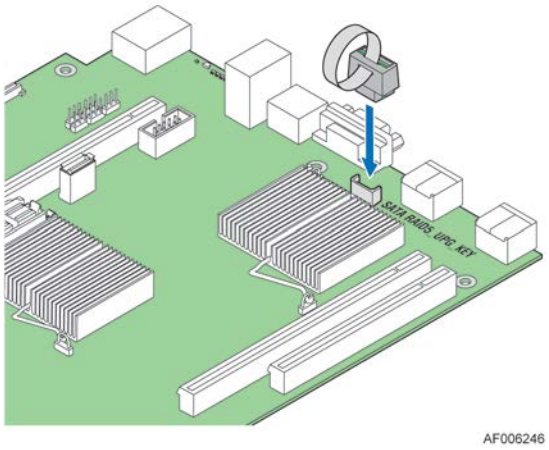


Figure 21. SATA RAID 5 Upgrade Key

- Maximum drive support = 6 (Maximum on-board SATA port support)
- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*.

Note: RAID configurations cannot span across the two embedded AHCI SATA controllers.

5.7 Network Interface

On the back edge of the server board are three RJ45 networking ports; “NIC #1”, “NIC #2”, and a Dedicated Management Port.

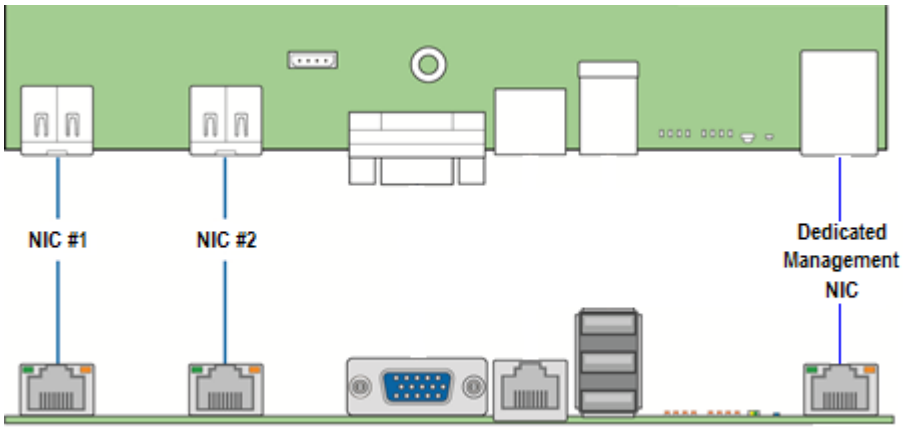


Figure 22. Network Interface Connectors

Each ethernet port drives two LEDs located on each network interface connector. The LED at the left of the connector is the link/activity LED and indicates network connection when on, and transmit/receive activity when blinking.



The LED at the right of the connector indicates link speed as defined in the following table.

LED	Color	LED State	NIC State
Left	Green	Off	LAN link not established
		On	LAN link is established
		Blinking	Transmit / Receive Activity
Right		Off	Lowest supported data rate
	Amber	On	Mid-range supported data rate
	Green	On	Highest supported data rate

Figure 23. External RJ45 NIC Port LED Definition

NOTE: Lowest, Mid-range, and Highest supported data rate is dependent on which onboard networking controller option is present. See section 5.7.1 for details on available onboard network controller options.

5.7.1 Intel® Ethernet Controller Options

The server board is offered with the following Intel® Ethernet Controller options:

- Intel® Ethernet Controller X540 10 GbE (Server board product code - S2600WTT)
- Intel® Ethernet Controller I350 1 GbE (Server board product code - S2600WT2)

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

5.7.2 Factory Programmed MAC Address Assignments

Depending on which onboard ethernet controller is present, the server board may have 5 or 7 MAC addresses programmed at the factory. MAC addresses are assigned as follows:

- NIC # 1 MAC address = Base #
- NIC # 2 MAC address = Base # + 1
- BMC LAN channel 0 MAC address = Base # + 2
- BMC LAN channel 1 MAC address = Base # + 3
- Dedicated On-board Management Port MAC address = Base # + 4

The following MAC address assignments are used for FCoE support on server boards with an on-board Intel® Ethernet Controller X540:

- NIC #1 SAN MAC address = Base # + 5
- NIC #2 SAN MAC address = Base # + 6

The base MAC address will be printed on a label and affixed to the server board and/or Intel server system. Factory programmed MAC addresses can also be viewed in the <F2> BIOS Setup Utility.

5.8 Video Support

The graphics controller of the integrated baseboard management controller provides support for the following features as implemented on the server board:

- Integrated Graphics Core with 2D Hardware accelerator
- DDR-3 memory interface with 16 MB of memory allocated and reported for graphics memory
- High speed Integrated 24-bit RAMDAC
- Single lane PCI-Express host interface running at Gen 1 speed

The integrated video controller supports all standard IBM* VGA modes. The following table shows the 2D modes supported for both CRT and LCD:

Table 15. Video Modes

2D Mode	2D Video Mode Support			
	8 bpp	16 bpp	24 bpp	32 bpp
640x480	X	X	X	X
800x600	X	X	X	X
1024x768	X	X	X	X
1152x864	X	X	X	X
1280x1024	X	X	X	X
1600x1200**	X	X		

** Video resolutions at 1600x1200 and higher are only supported through the external video connector located on the rear I/O section of the server board. Utilizing the optional front panel video connector may result in lower video resolutions.

The server board provides two onboard video interfaces. The primary video interface is accessed using a standard 15-pin VGA connector found on the back edge of the server board. In addition, video signals are routed to a 14-pin header labeled "FP_Video", allowing for the option of cabling to a front panel video connector. Attaching a monitor to the front panel video connector will disable the primary external video connector on the back edge of the board.

5.8.1 Dual Video and Add-In Video Adapters

There are enable/disable options in the <F2> BIOS Setup PCI Configuration screen for "Add-in Video Adapter" and "Onboard Video".

- When Onboard Video is *Enabled*, and Add-in Video Adapter is also *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 be active under an OS environment with the video driver support.
- When Onboard Video is *Enabled*, and Add-in Video Adapter is *Disabled*, then only the onboard video would be active.
- When Onboard Video is *Disabled*, and Add-in Video Adapter is *Enabled*, then only the add-in video adapter would be active.

Configurations with add-in video cards can get more complicated on server boards that have two or more CPU sockets. 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:

- This option appears only on boards which have the possibility of an add-in video adapter in a PCIe* slot on a CPU socket other than socket 1.
- When present, the option is grayed out and unavailable unless an add-in video card is actually installed in a PCIe* slot connected to the other socket.

- Because the Onboard Video is “hardwired” to CPU Socket 1, whenever Legacy VGA Socket is set to a CPU Socket other than Socket 1, that disables both Onboard Video ports.

5.8.1.1 Dual Monitor Video

The BIOS supports single and dual video on the S2600 family of Server Board when add-in video adapters are installed. Although there is no enable/disable option in BIOS screen for Dual Video, it works when both “Onboard video” and “Add-in Video Adapter” are enabled.

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

5.8.1.2 Configuration Cases – Multi-CPU Socket Boards and Add-In Video Adapters

Because this combination of CPU Socket and PCIe* topology is complicated and somewhat confusing, the following set of “Configuration Cases” was generated to clarify the design.

- ***When there are no add-in video cards installed…***

  Case 1: Onboard Video only active display.

  Onboard Video = Enabled (grayout, can't change)
  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)
  Add-in Video Adapter = Disabled (grayout, can't change)

- ***When there is one add-in video card connected to CPU Socket 1…***

  Case 2: Onboard video active display, add-in video doesn't display.

  Onboard Video = Enabled
  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)
  Add-in Video Adapter = Disabled

  Case 3: Add-in video active display, onboard video doesn't display.

  Onboard Video = Disabled,
  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)
  Add-in Video Adapter = Enabled

  Case 4: Both onboard video and add-in video are active displays. But only onboard could be the active display during BIOS POST (Dual Monitor).

  Onboard Video = Enabled
  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)
  Add-in Video Adapter = Enabled

- ***When there is one add-in video card connected to CPU Socket 2…***

  Case 5: Onboard video active display, add-in doesn't display.

  Onboard Video = Enabled
  Legacy VGA Socket = CPU Socket 1
  Add-in Video Adapter = Disabled (grayout, can't change)

  Case 6: Add-in video active display, onboard video doesn't display.

  Onboard Video = Disabled (grayout, can't change)
  Legacy VGA Socket = CPU Socket 2
  Add-in Video Adapter = Enabled (grayout, can't change)

- ***When there are add-in video cards connected to both CPU Socket 1 & 2…***

  Case 7: Onboard video active display, add-in video on Socket 1 and Add-in video on Socket 2 don't actively display.

  Onboard Video = Enabled
  Legacy VGA Socket = CPU Socket 1
  Add-in Video Adapter = Disabled

  Case 8: Add-in video on Socket 1 active display, onboard video and Add-in video on Socket 2 don't actively display.

  Onboard Video = Disabled
  Legacy VGA Socket = CPU Socket 1
  Add-in Video Adapter = Enabled

  Case 9: Both onboard video active and CPU Socket 1 add-in video active display. But only onboard could actively display during BIOS POST.

  Onboard Video = Enabled
  Legacy VGA Socket = CPU Socket 1
  Add-in Video Adapter = Enabled

  Case 10: Only CPU Socket 2 add-in video active display, neither onboard video nor CPU Socket 1 add-in video display.

  Onboard Video = Disabled (grayout, can't change)
  Legacy VGA Socket = CPU Socket 2
  Add-in Video Adapte = Enabled (grayout, can't change)

5.8.2 Setting Video Configuration Options using the BIOS Setup Utility

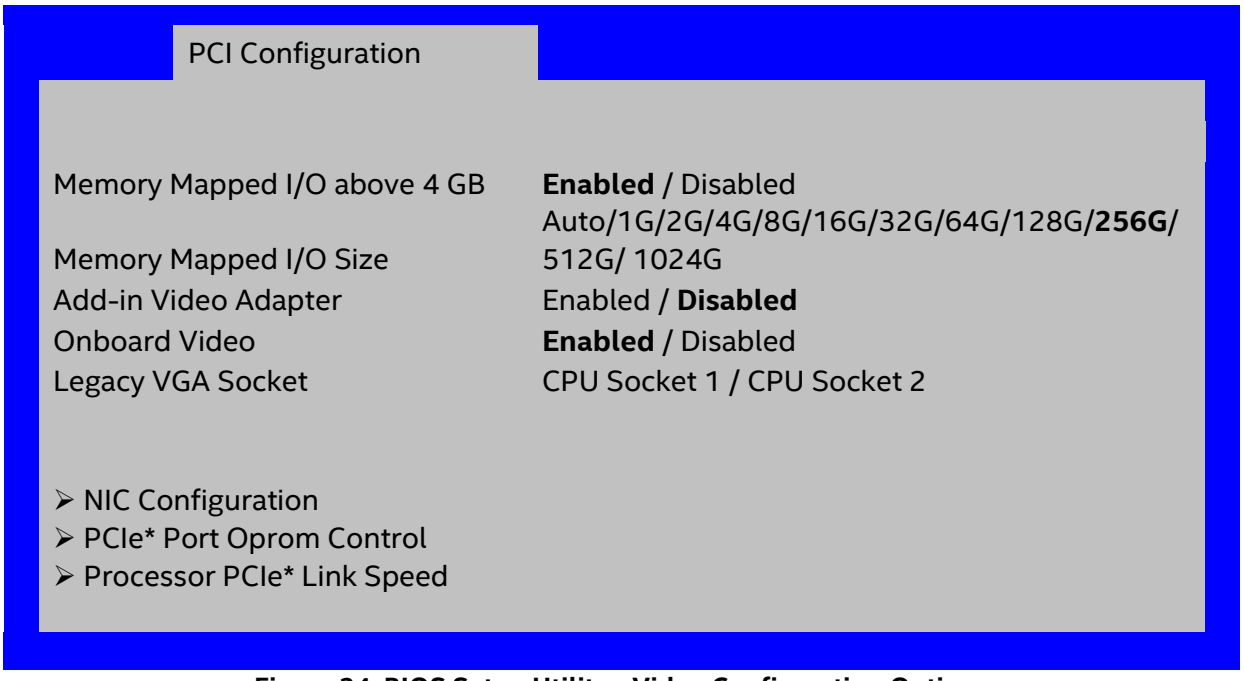


Figure 24. BIOS Setup Utility - Video Configuration Options

1. **Add-in Video Adapter**

Option Values: Enabled
Disabled

Help Text:
If enabled, the Add-in video adapter works as primary video device during POST if installed. If disabled, the on-board video controller becomes the primary video device.

Comments: This option must be enabled to use an add-in card as a primary POST Legacy Video device.
If there is no add-in video card in any PCIe* slot connected to CPU Socket 1 with the Legacy VGA Socket option set to *CPU Socket 1*, this option is set to *Disabled* and grayed out and unavailable.
If there is no add-in video card in any PCIe* slot connected to CPU Socket 2 with the Legacy VGA Socket option set to *CPU Socket 2*, this option is set to *Disabled* and grayed out and unavailable.
If the Legacy VGA Socket option is set to CPU Socket 1 with both Add-in Video Adapter and Onboard Video Enabled, the onboard video device works as primary video device while add-in video adapter as secondary.

2. **Onboard Video**

Option Values: Enabled
Disabled

Help Text:
On-board video controller.
Warning: System video is completely disabled if this option is disabled and an add-in video adapter is not installed.

Comments: When disabled, the system requires an add-in video card for the video to be seen. When there is no add-in video card installed, Onboard Video is set to *Enabled* and grayed out so it cannot be changed.
If there is an add-in video card installed in a PCIe* slot connected to CPU Socket 1, and the Legacy VGA Socket option is set to *CPU Socket 1*, then this Onboard Video option is available to be set and default as Disabled.
If there is an add-in video card installed on a PCIe* slot connected to CPU Socket 2, and the Legacy VGA Socket option is set to *CPU Socket 2*, this option is grayed out and unavailable, with a value set to *Disabled*. This is because the Onboard Video is connected to CPU Socket 1, and is not functional when CPU Socket 2 is the active path for video. When Legacy VGA Socket is set back to *CPU Socket 1*, this option becomes available again and is set to its default value of *Enabled*.

3. **Legacy VGA Socket**

Option Values: CPU Socket 1
CPU Socket 2

Help Text:
Determines whether Legacy VGA video output is enabled for PCIe* slots attached to Processor Socket 1 or 2. Socket 1 is the default.

Comments: This option is necessary when using an add-in video card on a PCIe* slot attached to CPU Socket 2, due to a limitation of the processor IIO. The Legacy video device can be connected through either socket but there is a setting that must be set on only one of the two. This option allows the switch to using a video card in a slot connected to CPU Socket 2.
This option does not appear unless the BIOS is running on a board which has one processor installed on CPU Socket 2 and can potentially have a video card installed in a PCIe* slot connected to CPU Socket 2.
This option is grayed out as unavailable and set to *CPU Socket 1* unless there is a processor installed on CPU Socket 2 and a video card installed in a PCIe* slot connected to CPU Socket 2. When this option is active and is set to *CPU Socket 2*, then both Onboard Video and Dual Monitor Video are set to *Disabled* and grayed out as unavailable. This is because the Onboard Video is a PCIe* device connected to CPU Socket 1, and is unavailable when the Legacy VGA Socket is set to Socket 2.

5.9 USB Support

The server board provides support for both USB 2.0 (up to 480 Mb/sec) and USB 3.0 (up to 5 Gb/sec).

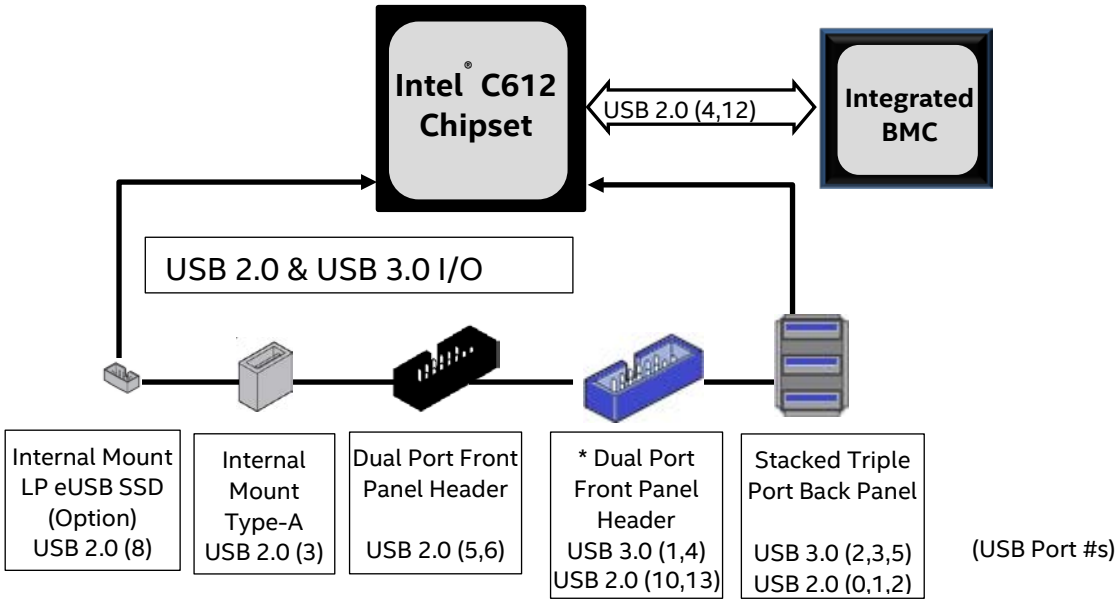


Figure 25. Onboard USB Port Support

* Note: Due to signal strength limits associated with USB 3.0 ports cabled to a front panel, some marginally compliant USB 3.0 devices may not be supported from these ports. In addition, server systems based on the Intel® Server Board S2600WT cannot be USB 3.0 certified with USB 3.0 ports cabled to a front panel.

5.9.1 Low Profile eUSB SSD Support

The server board provides support for a low profile eUSB SSD storage device. A 2mm 2x5-pin connector labeled “eUSB SSD” near the rear I/O section of the server board is used to connect this small flash storage device to the system.

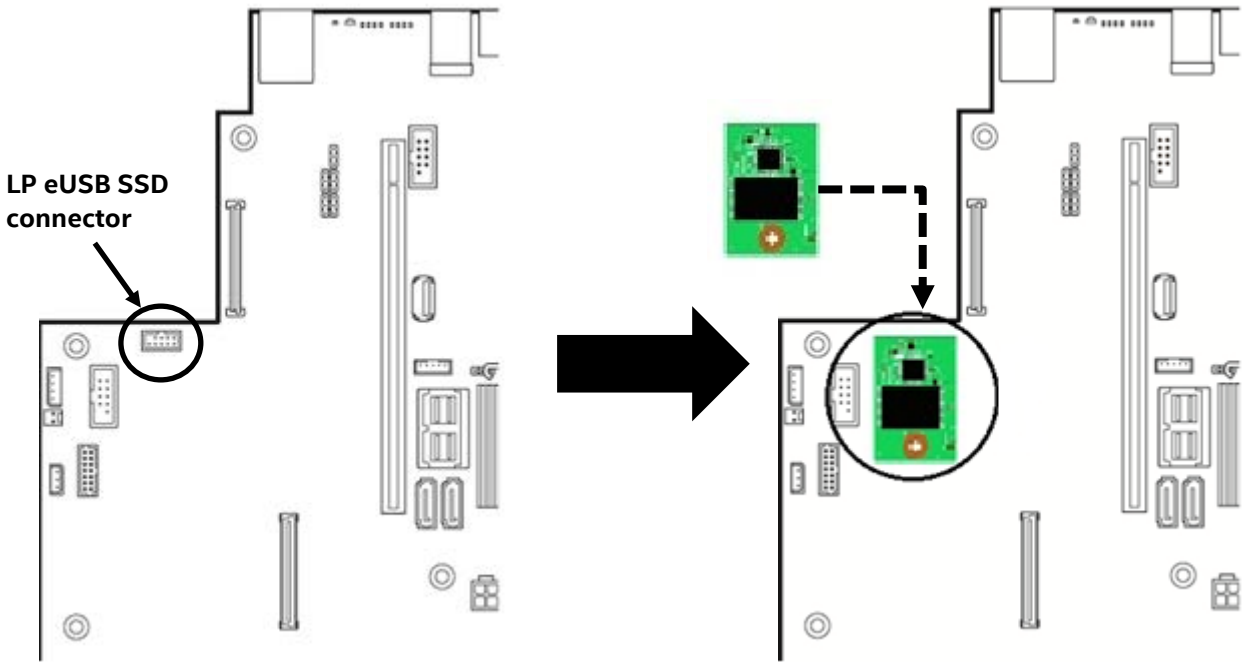


Figure 26. Low Profile eUSB SSD Support

eUSB SSD features include:

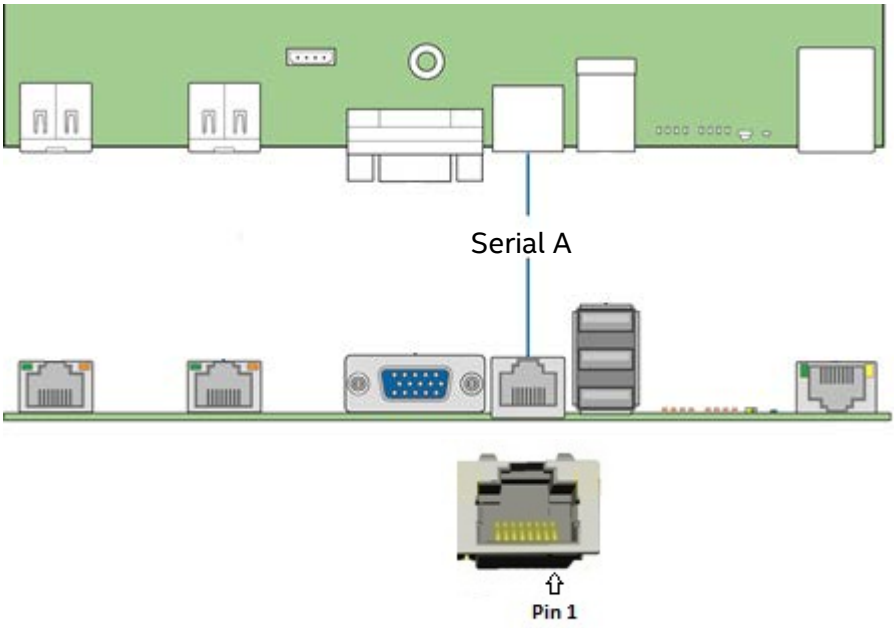
- 2 wire small form factor Universal Serial Bus 2.0 (Hi-Speed USB) interface to host
- Read Speed up to 35 MB/s and write Speed up to 24 MB/s.
- Capacity range from 256 MB to 32 GB.
- Support USB Mass Storage Class requirements for Boot capability.

Please visit the Intel® Server Configurator Tool at the following website for a list of supported eUSB SSD devices.

https://serverconfigurator.intel.com

5.10 Serial Ports

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.



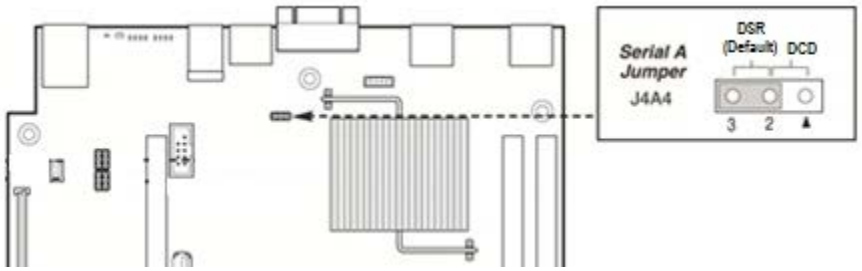
The Serial A connector has the following pin-out configuration.

Table 16. Serial A Connector Pin-out

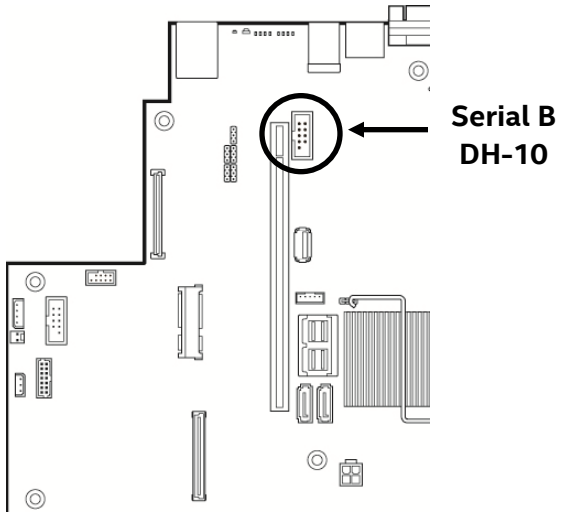
Signal Description	Pin#
RTS	1
DTR	2
SOUT	3
GROUND	4
RI	5
SIN	6
DCD or DSR	7**
CTS	8

** 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 “J4A4”, located behind the connector, from pins 1-2 (default) to pins 2-3.

Serial-A configuration jumper block (J4A4) setting:



Serial-B is an internal 10-pin DH-10 connector labeled “Serial_B”.



The Serial B connector has the following pin-out.

Table 17. 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. 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

6.1 BIOS Setup Utility Security Options Menu

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

Security

Administrator Password Status	<Installed/Not Installed>
User Password Status	<Installed/Not Installed>
Set Administrator Password	[123aBcDeFgH$#@]
Set User Password	[123aBcDeFgH$#@]
Power On Password	Enabled/**Disabled**
Front Panel Lockout	Enabled/**Disabled**
TPM State	<Displays current TPM Device State>
TPM Administrative Control	**No Operation/**Turn On/Turn Off/Clear Ownership

6.1.1 Password Setup

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

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

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

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

Passwords are case sensitive.

The Administrator and User passwords must be different from each other. An error message will be displayed and a different password must be entered if there is an attempt to enter the same password for both. The use of "Strong Passwords" is encouraged, but not required. In order to meet the criteria for a strong password, the password entered must be at least 8 characters in length, and must include at least one each of alphabetic, numeric, and special characters. If a weak password is entered, a warning message will be displayed, and the weak password will be accepted. Once set, a password can be cleared by changing it to a null string. This requires the Administrator password, and must be done through BIOS Setup. Clearing the Administrator password will also clear the User password. Passwords can also be cleared by using the Password Clear jumper on the server board. See Chapter 10 – Reset and Recovery Jumpers.

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

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

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

6.1.2 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
- Configure all BIOS setup options in the <F2> BIOS Setup Utility
- Clear both the Administrator and User passwords
- Access the <F6> Boot Menu during POST
- If the Power On Password function is enabled in BIOS Setup, the BIOS will halt early in POST to request a password (Administrator or User) before continuing POST.

6.1.3 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

Configuring an Administrator password imposes restrictions on booting the system, and configures most Setup fields to read-only if the Administrator password is not provided. The F6 Boot popup menu requires the Administrator password to function, and the USB Reordering is suppressed as long as the Administrator password is enabled. Users are restricted from booting in anything other than the Boot Order defined in Setup by an Administrator.

6.1.4 Front Panel Lockout

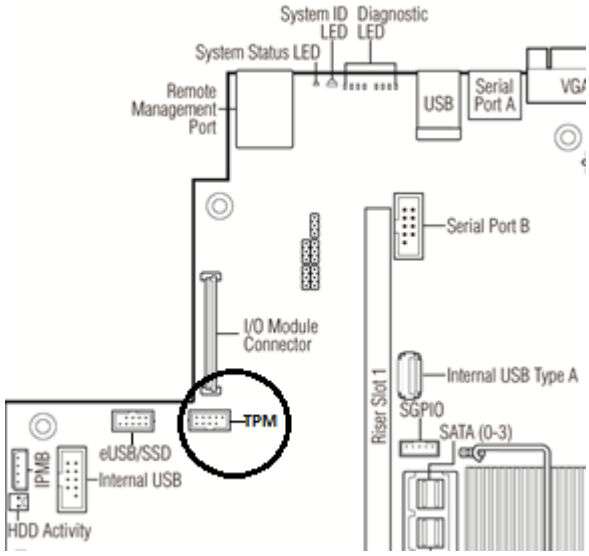
If enabled in BIOS setup, this option disables the following front panel features:

- The OFF function of the Power button
- System Reset button
- NMI Diagnostic Interrupt button

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

6.2 Trusted Platform Module (TPM) Support

The server board has the option to support a Trusted Platform Module (TPM) which plugs into a high density 14-pin connector labeled “TPM”.



A TPM is a hardware-based security device that addresses the growing concern on boot process integrity and offers better data protection. TPM protects the system start-up 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 by the Trusted Computing Group (TCG).

A TPM device 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.

6.2.1 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 EFI and legacy interfaces to a TPM-enabled operating system for using TPM.
- Produces 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 BitLocker* Requirement documents.

6.2.2 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:

1. User makes a TPM administrative request through the operating system's security software.
2. The operating system requests the BIOS to execute the TPM administrative command through TPM ACPI methods and then resets the system.
3. The BIOS verifies the physical presence and confirms the command with the operator.
4. The BIOS executes TPM administrative command(s), inhibits BIOS Setup entry and boots directly to the operating system which requested the TPM command(s).

6.2.3 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. TPM administrative options are only shown in the Security Menu screen when a TPM is physically installed on the board.

Using 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 independent 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.

TPM Setup Options using the BIOS Setup Utility

Table 18. TPM Setup Utility – Security Configuration Screen Fields

Setup Item	Options	Help Text	Comments
TPM State	Enabled and Activated Enabled and Deactivated Disabled and Activated Disabled and Deactivated		Information only. Shows the current TPM device state. A disabled TPM device will not execute commands that use TPM functions and TPM security operations will not be available. An enabled and deactivated TPM is in the same state as a disabled TPM except setting of TPM ownership is allowed if not present already. An enabled and activated TPM executes all commands that use TPM functions and TPM security operations will be available.
TPM Administrative Control	No Operation Turn On Turn Off Clear Ownership	[No Operation] - No changes to current state. [Turn On] - Enables and activates TPM. [Turn Off] - Disables and deactivates TPM. [Clear Ownership] - Removes the TPM ownership authentication and returns the TPM to a factory default state. Note: The BIOS setting returns to [No Operation] on every boot cycle by default.	Any Administrative Control operation selected will require the system to perform a Hard Reset in order to become effective.

6.3 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5-4600/2600/2400/1600 v3 Product Families support 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 from a BIOS Setup option. For general information about Intel® TXT, visit the Intel® Trusted Execution Technology website

http://www.intel.com/technology/security/

7. Platform Management

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

- Control systems functions – power system, ACPI, system reset control, system initialization, front panel interface, system event log
- Monitor various board and system sensors, regulate platform thermals and performance in order to maintain (when possible) server functionality in the event of component failure and/or environmentally stressed conditions
- Monitor and report system health
- Provide 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 v3 product families should be referenced for more in-depth and design level platform management information.

7.1 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.

7.1.1 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 run-time self-tests and makes results available to external entities.

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

7.1.2 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)
- Basic fan control using Control version 2 SDRs
- 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 can act as a DHCP client on all on-board LAN interfaces
- Platform environment control interface (PECI) thermal management support
- E-mail 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 on-line help

- Integrated KVM (with Intel® RMM4 Lite option installed)
- Enhancements to KVM redirection (with Intel® RMM4 Lite option installed)
    - 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
    - Password protected files are created which are accessible by Intel only
- Provisioning and inventory enhancements:
    - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.5 compliance
- Management support for PMBus* rev 1.2 compliant power supplies
- BMC Data Repository (Managed Data Region Feature)
- Support for an Intel® Local Control Display Panel
- 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.

7.2 Platform Management Features and Functions

7.2.1 Power Sub-system

The server board supports several power control sources which can initiate power-up or power-down activity.

Table 19. Server Board 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 *POWER_ON*)	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

7.2.2 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states:

Table 20. ACPI Power States

State	Supported	Description
S0	Yes	Working ▪ 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	Not supported
S4	No	Not supported
S5	Yes	Soft off ▪ 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.

7.2.3 System Initialization

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

7.2.3.1 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.

7.2.3.2 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.

7.2.3.3 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 deactivates POST LEDs after POST had completed. Refer to Appendix D for a complete list of supported POST Code Diagnostic LEDs.

7.2.4 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.

7.2.5 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 95231 bytes (approximately 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.

7.3 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.

7.3.1 Sensor Scanning

The value of many of the BMC’s sensors is derived by the BMC FW periodically polling physical sensors in the system to read temperature, voltages, and so on. Some of these physical sensors are built in to the BMC component itself and some are physically separated from the BMC. Polling of physical sensors for support of IPMI sensor monitoring does not occur until the BMC’s operational code is running and the IPMI FW subsystem has completed initialization. IPMI sensor monitoring is not supported in the BMC boot code. Additionally, the BMC selectively polls physical sensors based on the current power and reset state of the system and the availability of the physical sensor when in that state. For example, non-standby voltages are not monitored when the system is in a S5 power state.

7.3.2 Sensor Rearm Behavior

7.3.2.1 Manual versus Re-arm Sensors

Sensors can be either manual or automatic re-arm. 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 auto-rearm.

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 – Only applies 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.

7.3.2.2 Re-arm and Event Generation

All BMC-owned sensors that show an asserted event status generate a de-assertion SEL event when the sensor is re-armed, provided that the associated SDR is configured to enable a de-assertion event for that condition. This applies regardless of whether the sensor is a threshold/analog sensor or a discrete sensor.

To manually re-arm the sensors, the sequence is outlined below:

1. A failure condition occurs and the BMC logs an assertion event.
2. If this failure condition disappears, the BMC logs a de-assertion event (if so configured.)
3. The sensor is re-armed by one of the methods described in the previous section.
4. The BMC clears the sensor status.
5. The sensor is put into "reading-state-unavailable" state until it is polled again or otherwise updated.
6. The sensor is updated and the “reading-state-unavailable” state is cleared. A new assertion event will be logged if the fault state is once again detected.

All auto-rearm sensors that show an asserted event status generate a de-assertion SEL event at the time the BMC detects that the condition causing the original assertion is no longer present; and the associated SDR is configured to enable a de-assertion event for that condition.

7.3.3 BIOS Event-Only Sensors

BIOS-owned discrete sensors are used for event generation only and are not accessible through IPMI sensor commands like the *Get Sensor Reading* command. Note that in this case the sensor owner designated in the SDR is not the BMC.

An example of this usage would be the SELs logged by the BIOS for uncorrectable memory errors. Such SEL entries would identify a BIOS-owned sensor ID.

7.3.4 Margin Sensors

There is sometimes a need for an IPMI sensor to report the difference (margin) from a non-zero reference offset. For the purposes of this document, these type sensors are referred to as margin sensors. For instance, for the case of a temperature margin sensor, if the reference value is 90 degrees and the actual temperature of the device being monitored is 85 degrees, the margin value would be -5.

7.3.5 IPMI Watchdog Sensor

The BMC supports a Watchdog Sensor as a means to log SEL events due to expirations of the IPMI 2.0 compliant Watchdog Timer.

7.3.6 BMC Watchdog Sensor

The BMC supports an IPMI sensor to report that a BMC reset has occurred due to action taken by the BMC Watchdog feature. A SEL event will be logged whenever either the BMC FW stack is reset or the BMC CPU itself is reset.

7.3.7 BMC System Management Health Monitoring

The BMC tracks the health of each of its IPMI sensors and report failures by providing a “BMC FW Health” sensor of the IPMI 2.0 sensor type Management Subsystem Health with support for the Sensor Failure offset. Only assertions should be logged into the SEL for the Sensor Failure offset. The BMC Firmware Health sensor asserts for any sensor when 10 consecutive sensor errors are read. These are not standard sensor events (that is, threshold crossings or discrete assertions), these are BMC Hardware Access Layer (HAL) errors. This means the BMC is unable to get a reading from the sensor. If a successful sensor read is completed, the counter resets to zero.

7.3.8 VR Watchdog Timer

The BMC FW monitors that the power sequence for the board VR controllers is completed when a DC power-on is initiated. Incompletion of the sequence indicates a board problem, in which case the FW powers down the system.

The BMC FW supports a discrete IPMI sensor for reporting and logging this fault condition.

7.3.9 System Airflow Monitoring

The BMC provides an IPMI sensor to report the volumetric system airflow in CFM (cubic feet per minute). The air flow in CFM is calculated based on the system fan Pulse Width Modulation (PWM) values. The specific PWM or PWMs, used to determine the CFM is SDR configurable. The relationship between PWM and CFM is based on a lookup table in an OEM SDR.

The airflow data is used in the calculation for exit air temperature monitoring. It is exposed as an IPMI sensor to allow a datacenter management application to access this data for use in rack-level thermal management.

7.3.10 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 back plane)
- Board ambient thermal sensors
- Processor temperature
- Memory (DIMM) temperature
- CPU VRD Hot monitoring
- Power supply (only supported for PMBus*-compliant PSUs)

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

7.3.10.1 Absolute Value versus Margin Sensors

Thermal monitoring sensors fall into three basic categories:

- Absolute temperature sensors – These are analog/threshold sensors that provide a value that corresponds to an absolute temperature value.
- Thermal margin sensors – These are analog/threshold sensors that provide a value that is relative to some reference value.
- Thermal fault indication sensors – These are discrete sensors that indicate a specific thermal fault condition.

7.3.10.2 Processor DTS-Spec Margin Sensor(s)

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family incorporate a DTS based thermal spec. This allows a much more accurate control of the thermal solution and will enable lower fan speeds and lower fan power consumption. The main usage of this sensor is as an input to the BMC’s fan control algorithms. The BMC implements this as a threshold sensor. There is one DTS sensor for each installed physical processor package. Thresholds are not set and alert generation is not enabled for these sensors.

7.3.10.3 Processor Thermal Margin Sensor(s)

Each processor supports a physical thermal margin sensor per core that is readable through the PECI interface. This provides a relative value representing a thermal margin from the core’s throttling thermal trip point. Assuming that temperature controlled throttling is enabled; the physical core temperature sensor reads ‘0’, which indicates the processor core is being throttled.

The BMC supports one IPMI processor (margin) temperature sensor per physical processor package. This sensor aggregates the readings of the individual core temperatures in a package to provide the hottest core temperature reading. When the sensor reads ‘0’, it indicates that the hottest processor core is throttling.

Due to the fact that the readings are capped at the core's thermal throttling trip point (reading = 0), thresholds are not set and alert generation is not enabled for these sensors.

7.3.10.4 Processor Thermal Control Monitoring (Prochot)

The BMC FW monitors the percentage of time that a processor has been operationally constrained over a given time window (nominally six seconds) due to internal thermal management algorithms engaging to reduce the temperature of the device. When any processor core temperature reaches its maximum operating temperature, the processor package PROCHOT# (processor hot) signal is asserted and these management algorithms, known as the Thermal Control Circuit (TCC), engage to reduce the temperature, provided TCC is enabled. TCC is enabled by BIOS during system boot. This monitoring is instantiated as one IPMI analog/threshold sensor per processor package. The BMC implements this as a threshold sensor on a per-processor basis.

Under normal operation, this sensor is expected to read '0' indicating that no processor throttling has occurred.

The processor provides PECI-accessible counters, one for the total processor time elapsed and one for the total thermally constrained time, which are used to calculate the percentage assertion over the given time window.

7.3.10.5 Processor Voltage Regulator (VRD) Over-Temperature Sensor

The BMC monitors processor VRD_HOT# signals. The processor VRD_HOT# signals are routed to the respective processor PROCHOT# input in order to initiate throttling to reduce processor power draw, therefore indirectly lowering the VRD temperature.

The processor VRD_HOT# signals are routed to the respective processor PROCHOT# input in order to initiate throttling to reduce processor power draw, therefore indirectly lowering the VRD temperature. There is one processor VRD_HOT# signal per CPU slot.

The memory VRD_HOT# signals are routed to the respective processor MEMHOT# inputs in order to throttle the associated memory to effectively lower the temperature of the VRD feeding that memory. For Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family there are 2 memory VRD_HOT# signals per CPU slot.

The BMC instantiates one discrete IPMI sensor for each processor and memory VRD_HOT# signal.

7.3.10.6 Inlet Temperature Sensor

Each platform supports a thermal sensor for monitoring the inlet temperature. In most cases, ME firmware will issue Get Sensor Reading IPMI command to the BMC to get the Inlet temperature. ME firmware determines which of the BMC thermal sensors to use for inlet temperature.

7.3.10.7 Baseboard Ambient Temperature Sensor(s)

The server baseboard provides one or more physical thermal sensors for monitoring the ambient temperature of a board location. This is typically to provide rudimentary thermal monitoring of components that lack internal thermal sensors.

7.3.10.8 Server South Bridge (SSB) Thermal Monitoring

The BMC monitors the SSB temperature. This is instantiated as an analog (threshold) IPMI thermal sensor.

7.3.10.9 Exit Air Temperature Monitoring

The BMC synthesizes a virtual sensor to approximate system exit air temperature for use in fan control. This is calculated based on the total power being consumed by the system and the total volumetric air flow provided by the system fans. Each system shall be characterized in tabular format to understand total volumetric flow versus fan speed. The BMC calculates an average exit air temperature based on the total system power, front panel temperature, and the volumetric system air flow (cubic feet per meter or CFM).

This sensor is only available on systems in an Intel® chassis. The Exit Air temp sensor is only available when PMBus* power supplies are installed.

7.3.10.10 Ethernet Controller Thermal Monitoring

The Intel® Ethernet Controller I350-AM4 and Intel® Ethernet Controller 10 Gigabit X540 support an on-die thermal sensor. For baseboard Ethernet controllers that use these devices, the BMC will monitor the sensors and use this data as input to the fan speed control. The BMC will instantiate an IPMI temperature sensor for each device on the baseboard.

7.3.10.11 Memory VRD-Hot Sensor(s)

The BMC monitors memory VRD_HOT# signals. The memory VRD_HOT# signals are routed to the respective processor MEMHOT# inputs in order to throttle the associated memory to effectively lower the temperature of the VRD feeding that memory.

For Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family there are 2 memory VRD_HOT# signals per CPU slot. The BMC instantiates one discrete IPMI sensor for each memory VRD_HOT# signal.

7.3.10.12 Add-in Module Thermal Monitoring

Some boards have dedicated slots for an IO module and/or a SAS module. For boards that support these slots, the BMC will instantiate an IPMI temperature sensor for each slot. The modules themselves may or may not provide a physical thermal sensor (a TMP75 device). If the BMC detects that a module is installed, it will attempt to access the physical thermal sensor and, if found, enable the associated IPMI temperature sensor.

7.3.10.13 Processor ThermTrip

When a Processor ThermTrip occurs, the system hardware will automatically power down the server. If the BMC detects that a ThermTrip occurred, then it will set the ThermTrip offset for the applicable processor status sensor.

7.3.10.14 Server South Bridge (SSB) ThermTrip Monitoring

The BMC supports SSB ThermTrip monitoring that is instantiated as an IPMI discrete sensor. When a SSB ThermTrip occurs, the system hardware will automatically power down the server and the BMC will assert the sensor offset and log an event.

7.3.10.15 DIMM ThermTrip Monitoring

The BMC supports DIMM ThermTrip monitoring that is instantiated as one aggregate IPMI discrete sensor per CPU. When a DIMM ThermTrip occurs, the system hardware will automatically power down the server and the BMC will assert the sensor offset and log an event.

This is a manual re-arm sensor that is rearmed on system resets and power-on (AC or DC power on transitions).

7.3.11 Processor Sensors

The BMC provides IPMI sensors for processors and associated components, such as voltage regulators and fans. The sensors are implemented on a per-processor basis.

Table 21. Processor Sensors

Sensor Name	Per Processor Socket	Description
Processor Status	Yes	Processor presence and fault state
Digital Thermal Sensor	Yes	Relative temperature reading by means of PECI
Processor VRD Over-Temperature Indication	Yes	Discrete sensor that indicates a processor VRD has crossed an upper operating temperature threshold
Processor Voltage	Yes	Threshold sensor that indicates a processor power-good state
Processor Thermal Control (PROCHOT#)	Yes	Percentage of time a processor is throttling due to thermal conditions

7.3.11.1 Processor Status Sensors

The BMC provides an IPMI sensor of type processor for monitoring status information for each processor slot. If an event state (sensor offset) has been asserted, it remains asserted until one of the following happens:

1. A Rearm Sensor Events command is executed for the processor status sensor.
2. AC or DC power cycle, system reset, or system boot occurs.

The BMC provides system status indication to the front panel LEDs for processor fault conditions as listed in following table.

CPU Presence status is not saved across AC power cycles and therefore will not generate a de-assertion after cycling AC power.

Table 22. Processor Status Sensor Implementation

Offset	Processor Status	Detected By
0	Internal error (IERR)	Not Supported
1	Thermal trip	BMC
2	FRB1/BIST failure	Not Supported
3	FRB2/Hang in POST failure	BIOS[1]
4	FRB3/Processor startup/initialization failure (CPU fails to start)	Not Supported
5	Configuration error (for DMI)	BIOS[1]
6	SM BIOS uncorrectable CPU-complex error	Not Supported
7	Processor presence detected	BMC
8	Processor disabled	Not Supported
9	Terminator presence detected	Not Supported

Note:

1. Fault is not reflected in the processor status sensor.

7.3.11.2 Processor Population Fault (CPU Missing) Sensor

The BMC supports a *Processor Population Fault* sensor. This is used to monitor for the condition in which processor sockets are not populated as required by the platform HW to allow power-on of the system.

At BMC startup, the BMC will check for the fault condition and set the sensor state accordingly. The BMC also checks for this fault condition at each attempt to DC power-on the system. At each DC power-on attempt, a beep code is generated if this fault is detected.

The following steps are used to correct the fault condition and clear the sensor state:

1. AC power down the server
2. Install a processor into the CPU _1 socket
3. AC power on the server

7.3.11.3 ERR2 Timeout Monitoring

The BMC supports an ERR2 Timeout Sensor (1 per CPU) that asserts if a CPU's ERR[2] signal has been asserted for longer than a fixed time period (> 90 seconds). ERR[2] is a processor signal that indicates when the IIO (Integrated IO module in the processor) has a fatal error which could not be communicated to the core to trigger SMI. ERR[2] events are fatal error conditions, where the BIOS and OS will attempt to gracefully handle error, but may not always be able to do so reliably. A continuously asserted ERR[2] signal is an indication that the BIOS cannot service the condition that caused the error. This is usually because that condition prevents the BIOS from running.

When an ERR2 timeout occurs, the BMC asserts/de-asserts the ERR2 Timeout Sensor, and logs a SEL event for that sensor. The default behavior for BMC core firmware is to initiate a system reset upon detection of an ERR2 timeout. The BIOS setup utility provides an option to disable or enable system reset by the BMC for detection of this condition.

7.3.11.4 CATERR Sensor

The BMC supports a CATERR sensor for monitoring the system CATERR signal.

The CATERR signal is defined as having 3 states:

- high (no event)
- pulsed low (possibly fatal may be able to recover)
- low (fatal).

All processors in a system have their CATERR pins tied together. The pin is used as a communication path to signal a catastrophic system event to all CPUs. The BMC has direct access to this aggregate CATERR signal. The BMC only monitors for the "CATERR held low" condition. A pulsed low condition is ignored by the BMC. If a CATERR-low condition is detected, the BMC logs an error message to the SEL against the CATERR sensor and the default action after logging the SEL entry is to reset the system. The BIOS setup utility provides an option to disable or enable system reset by the BMC for detection of this condition.

The sensor is rearmed on power-on (AC or DC power on transitions). It is not rearmed on system resets in order to avoid multiple SEL events that could occur due to a potential reset loop if the CATERR keeps recurring, which would be the case if the CATERR was due to an MSID mismatch condition.

When the BMC detects that this aggregate CATERR signal has asserted, it can then go through PECI to query each CPU to determine which one was the source of the error and write an OEM code identifying the CPU slot into an event data byte in the SEL entry. If PECI is non-functional (functionality is not guaranteed in this situation), then the OEM code should indicate that the source is unknown.

Event data byte 2 and byte 3 for CATERR sensor SEL events

ED1 – 0xA1

ED2 - CATERR type.
0: Unknown
1: CATERR
2: CPU Core Error (not supported on Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family)
3: MSID Mismatch
4: CATERR due to CPU 3-strike timeout

ED3 - CPU bitmap that causes the system CATERR.
[0]: CPU1
[1]: CPU2
[2]: CPU3
[3]: CPU4

When a CATERR Timeout event is determined to be a CPU 3-strike timeout, The BMC shall log the logical FRU information (e.g. bus/dev/func for a PCIe* device, CPU, or DIMM) that identifies the FRU that caused the error in the extended SEL data bytes. In this case, Ext-ED0 will be set to 0x70 and the remaining ED1-ED7 will be set according to the device type and info available.

7.3.11.5 MSID Mismatch Sensor

The BMC supports a *MSID Mismatch* sensor for monitoring for the fault condition that will occur if there is a power rating incompatibility between a baseboard and a processor. The sensor is rearmed on power-on (AC or DC power on transitions).

7.3.12 Voltage Monitoring

The BMC provides voltage monitoring capability for voltage sources on the main board and processors such that all major areas of the system are covered. This monitoring capability is instantiated in the form of IPMI analog/threshold sensors.

7.3.12.1 Discrete Voltage Sensors

The discrete voltage sensor monitors multiple voltages from sensors around the baseboard and then asserts a bit in the SEL event data for each sensor that is out of range. The sensor name for the asserted bit can be retrieved via the Get Voltage Name IPMI function.

7.3.13 Fan Monitoring

BMC fan monitoring support includes monitoring of fan speed (RPM) and fan presence.

7.3.13.1 Fan Tach Sensors

Fan Tach sensors are used for fan failure detection. The reported sensor reading is proportional to the fan's RPM. This monitoring capability is instantiated in the form of IPMI analog/threshold sensors.

Most fan implementations provide for a variable speed fan, so the variations in fan speed can be large. Therefore the threshold values must be set sufficiently low as to not result in inappropriate threshold crossings.

Fan tach sensors are implemented as manual re-arm sensors because a lower-critical threshold crossing can result in full boosting of the fans. This in turn may cause a failing fan's speed to rise above the threshold and can result in fan oscillations.

As a result, fan tach sensors do not auto-rearm when the fault condition goes away but rather are rearmed for either of the following occurrences:

a. The system is reset or power-cycled.
b. The fan is removed and either replaced with another fan or re-inserted. This applies to hot-swappable fans only. This re-arm is triggered by change in the state of the associated fan presence sensor.

After the sensor is rearmed, if the fan speed is detected to be in a normal range, the failure conditions shall be cleared and a de-assertion event shall be logged.

7.3.13.2 Fan Presence Sensors

Some chassis and server boards provide support for hot-swap fans. These fans can be removed and replaced while the system is powered on and operating normally. The BMC implements fan presence sensors for each hot swappable fan. These are instantiated as IPMI discrete sensors.

Events are only logged for fan presence upon changes in the presence state after AC power is applied (no events logged for initial state).

7.3.13.3 Fan Redundancy Sensor

The BMC supports redundant fan monitoring and implements fan redundancy sensors for products that have redundant fans. Support for redundant fans is chassis-specific.

A fan redundancy sensor generates events when its associated set of fans transition between redundant and non-redundant states, as determined by the number and health of the component 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.

There is a fan redundancy sensor implemented for each redundant group of fans in the system. Assertion and de-assertion event generation is enabled for each redundancy state.

7.3.13.4 Power Supply Fan Sensors

Monitoring is implemented through IPMI discrete sensors, one for each power supply fan. The BMC polls each installed power supply using the PMBus* fan status commands to check for failure conditions for the power supply fans. The BMC asserts the “performance lags” offset of the IPMI sensor if a fan failure is detected.

Power supply fan sensors are implemented as manual re-arm sensors because a failure condition can result in boosting of the fans. This in turn may cause a failing fan’s speed to rise above the “fault” threshold and can result in fan oscillations. As a result, these sensors do not auto-rearm when the fault condition goes away but rather are rearmed only when the system is reset or power-cycled, or the PSU is removed and replaced with the same or another PSU.

After the sensor is rearmed, if the fan is no longer showing a failed state, the failure condition in the IPMI sensor shall be cleared and a de-assertion event shall be logged.

7.3.13.5 Monitoring for “Fans Off” Scenario

On Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family, it is likely that there will be situations where specific fans are turned off based on current system conditions. BMC Fan monitoring will comprehend this scenario and not log false failure events. The recommended method is for

the BMC FW to halt updates to the value of the associated fan tach sensor and set that sensor's IPMI sensor state to "reading-state-unavailable" when this mode is active. Management software must comprehend this state for fan tach sensors and not report these as failure conditions.

The scenario for which this occurs is that the BMC Fan Speed Control (FSC) code turns off the fans by setting the PWM for the domain to 0. This is done when based on one or more global aggregate thermal margin sensor readings dropping below a specified threshold.

By default the fans-off feature will be disabled. There is a BMC command and BIOS setup option to enable/disable this feature.

The SmaRT/CLST system feature will also momentarily gate power to all the system fans to reduce overall system power consumption in response to a power supply event (for example, to ride out an AC power glitch). However, for this scenario, the fan power is gated by HW for only 100ms, which should not be long enough to result in triggering a fan fault SEL event.

7.3.14 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 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.

7.3.14.1 Hot-Swap Fans

Hot-swap 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 rearmed. 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.

7.3.14.2 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.

7.3.14.3 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.

7.3.14.4 Nominal Fan Speed

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.

OEM SDR records are used to configure which temperature sensors are associated with which fan control domains and the algorithmic relationship between the temperature and fan speed. Multiple OEM SDRs can reference or control the same fan control domain; and multiple OEM SDRs can reference the same temperature sensors.

The PWM duty-cycle value for a domain is computed as a percentage using one or more instances of a stepwise linear algorithm and a clamp algorithm. The transition from one computed nominal fan speed (PWM value) to another is ramped over time to minimize audible transitions. The ramp rate is configurable by means of the OEM SDR.

Multiple stepwise linear and clamp controls can be defined for each fan domain and used simultaneously. For each domain, the BMC uses the maximum of the domain's stepwise linear control contributions and the sum of the domain's clamp control contributions to compute the domain's PWM value, except that a stepwise linear instance can be configured to provide the domain maximum.

Hysteresis can be specified to minimize fan speed oscillation and to smooth fan speed transitions. If a Tcontrol SDR record does not contain a hysteresis definition, for example, an SDR adhering to a legacy format, the BMC assumes a hysteresis value of zero.

7.3.14.5 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 buffered DIMMs.

7.3.14.6 Thermal Sensor Input to Fan Speed Control

The BMC uses various IPMI sensors as 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 input to fan speed control:

- Front Panel Temperature Sensor[1]
- CPU Margin Sensors[2,4,5]
- DIMM Thermal Margin Sensors[2,4]
- Exit Air Temperature Sensor[1, 7, 9]
- PCH Temperature Sensor[3,5]
- On-board Ethernet Controller Temperature Sensors[3, 5]
- Add-In Intel SAS Module Temperature Sensors[3, 5]
- PSU Thermal Sensor[3, 8]
- CPU VR Temperature Sensors[3, 6]
- DIMM VR Temperature Sensors[3, 6]
- BMC Temperature Sensor[3, 6]
- Global Aggregate Thermal Margin Sensors [7]
- Hot Swap Backplane Temperature Sensors
- I/O Module Temperature Sensor (With option installed)
- Intel® SAS Module (With option installed)
- Riser Card Temperature Sensors
- Intel® Xeon Phi™ coprocessor (With option installed)

Notes:
1. For fan speed control in Intel chassis
2. Temperature margin from throttling threshold
3. Absolute temperature
4. PECI value or margin value
5. On-die sensor

6. On-board sensor
7. Virtual sensor
8. Available only when PSU has PMBus
9. Calculated estimate

A simple model is shown in the following figure which gives a high level representation of how the fan speed control structure creates the resulting fan speeds.

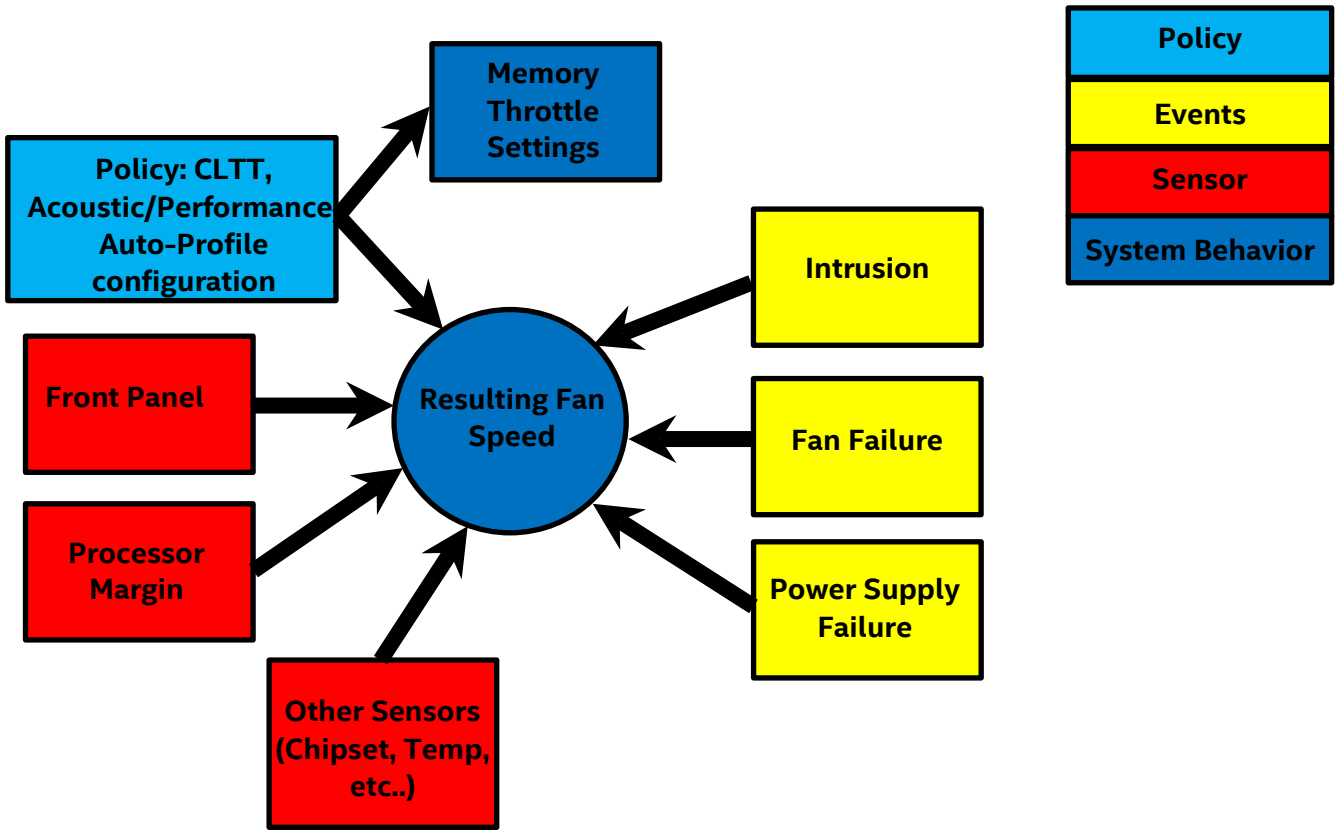


Figure 27. High-level Fan Speed Control Process

7.3.14.6.1 Processor Thermal Management

Processor thermal management utilizes clamp algorithms for which the Processor DTS-Spec margin sensor is a controlling input. This replaces the use of the (legacy) raw DTS sensor reading that was utilized on previous generation platforms. The legacy DTS sensor is retained only for monitoring purposes and is not used as an input to the fan speed control.

7.3.14.6.2 Memory Thermal Management

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

7.3.14.6.2.1 Memory Thermal Throttling

The system shall support thermal management through open loop thermal throttling (OLTT) and closed loop thermal throttling (CLTT) of system memory based on the platform as well as availability of valid temperature sensors on the installed memory DIMMs. 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. Support for CLTT on mixed-mode DIMM populations (that is, some installed DIMMs have valid temp sensors and some do not) is not supported. 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 Open Loop Thermal Throttling (Static-OLTT):** OLTT control registers that are configured by BIOS MRC remain fixed after post. The system does not change any of the throttling control registers in the embedded memory controller during runtime.
- **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 Open Loop Thermal Throttling (Dynamic-OLTT):** OLTT control registers are configured by BIOS MRC during POST. Adjustments are made to the throttling during runtime based on changes in system cooling (fan speed).
- **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 v3 product family introduce a new type of CLTT which is referred to as 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 v3 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.

7.3.14.6.3 DIMM Temperature Sensor Input to Fan Speed Control

A clamp algorithm is used for controlling fan speed based on DIMM temperatures. Aggregate DIMM temperature margin sensors are used as the control input to the algorithm.

7.3.14.6.4 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 is does not provide results as accurate 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.

BIOS will have an embedded lookup table that provides thermal offset values for each DIMM type and air velocity range (3 ranges of air velocity are supported). During system boot BIOS will provide 3 offset values (corresponding to the 3 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.

7.3.14.6.5 Fan Profiles

The server system supports multiple fan control profiles to support acoustic targets and American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) compliance. The BIOS Setup utility can be used to choose between meeting the target acoustic level or enhanced system performance. This is accomplished through fan profiles.

The BMC supports eight fan profiles, numbered from 0 to 7. A profile is a defined set of fan control configuration SDRs which determines the behavior of the fans based on various thermal and status inputs in the system. Profile 0 provides the lowest base fan speeds and therefore the lowest noise level. Profile 7 provides the highest fan speeds and therefore the highest noise level. The BMC automatically determines which profile to use based on CPU population and type, memory population and type, hard drive presence and PCIe* card population. The BIOS Setup Utility also provides a “Performance” setting which forces the BMC to use a higher than normal profile.

The BMC provides commands that query for fan profile support and it provides a way to enable a fan profile. Enabling a fan profile determines which Tcontrol SDRs are used for fan management. The BMC only supports enabling a fan profile through the command if that profile is supported on all fan domains defined for the system. It is important to configure the SDRs so that all desired fan profiles are supported on each fan domain. If no single profile is supported across all domains, the BMC, by default, uses profile 0 and does not allow it to be changed.

7.3.14.6.6 Open-Loop Thermal Throttling Fallback

Normal system operation uses closed-loop thermal throttling (CLTT) and DIMM temperature monitoring as major factors in overall thermal and acoustics management. In the event that BIOS is unable to configure the system for CLTT, it defaults to open-loop thermal throttling (OLTT). In the OLTT mode, it is assumed that the DIMM temperature sensors are not available for fan speed control. The BIOS communicates the throttling mode to the BMC along with the fan profile number when it sends the *Set Fan Control Configuration* command.
When OLTT mode is specified, the BMC internally blocks access to the DIMM temperatures, causing the DIMM aggregate margin sensors to be marked as *Reading/State Unavailable*. The BMC then uses the failure-control values for these sensors, if specified in the Tcontrol SDRs, as their fan speed contributions.

7.3.14.6.7 ASHRAE Compliance

System requirements for ASHRAE compliance are defined in the *Common Fan Speed Control & Thermal Management Platform Architecture Specification*.

7.3.14.7 Power Supply Fan Speed Control

This section describes the system level control of the fans internal to the power supply over the PMBus*. Some, but not all, Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family will require that the power supplies be included in the system level fan speed control. For any system that requires either of these capabilities, the power supply must be PMBus*-compliant.

7.3.14.7.1 System Control of Power Supply Fans

Some products require that the BMC control the speed of the power supply fans, as is done with normal system (chassis) fans, except that the BMC cannot reduce the power supply fan any lower than the internal power supply control is driving it. For these products the BMC FW must have the ability to control and monitor the power supply fans through PMBus* commands. The power supply fans are treated as a system

fan domain for which fan control policies are mapped, just as for chassis system fans, with system thermal sensors (rather than internal power supply thermal sensors) used as the input to a clamp algorithm for the power supply fan control. This domain has both piecewise clipping curves and clamped sensors mapped into the power supply fan domain. All the power supplies can be defined as a single fan domain.

7.3.14.7.2 Use of Power Supply Thermal Sensors as Input to System (Chassis) Fan Control

Some products require that the power supply internal thermal sensors are used as control inputs to the system (chassis) fans, in the same manner as other system thermal sensors are used for this purpose. The power supply thermal sensors are included as clamped sensors into one or more system fan domains, which may include the power supply fan domain.

7.3.14.8 Fan Boosting due to Fan Failures

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family introduce additional capabilities for handling fan failure or removal as described in this section.

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, then there will be six different fan fail reactions.

7.3.14.9 Programmable Fan PWM Offset

The system provides a BIOS Setup option to boost the system fan speed by a programmable positive offset or a “Max” setting. Setting the programmable offset causes the BMC to add the offset to the fan speeds to which it would otherwise be driving the fans. The Max setting causes the BMC to replace the domain minimum speed with alternate domain minimums that also are programmable through SDRs.

This capability is offered to provide system administrators the option to manually configure fan speeds in instances where the fan speed optimized for a given platform may not be sufficient when a high end add-in adapter is configured into the system. This enables easier usage of the fan speed control to support Intel as well as third party chassis and better support of ambient temperatures higher than 35°C.

7.3.15 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*, please see the System Management Interface Forum Web site http://www.powersig.org/.

7.3.16 Power Supply Dynamic Redundancy Sensor

The BMC supports redundant power subsystems and implements a Power Unit Redundancy sensor per platform. A Power Unit Redundancy sensor is of sensor type Power Unit (09h) and reading type Availability Status (0Bh). This sensor generates events when a power subsystem transitions between redundant and non-redundant states, as determined by the number and health of the power subsystem’s component power supplies. The BMC implements Dynamic Power Supply Redundancy status based upon current system load requirements as well as total Power Supply capacity. This status is independent of the Cold Redundancy status. This prevents the BMC from reporting Fully Redundant Power supplies when the load required by the system exceeds half the power capability of all power supplies installed and operational. Dynamic Redundancy detects this condition and generates the appropriate SEL event to notify the user of the condition. Power supplies of different power ratings may be swapped in and out to adjust the power capacity

and the BMC will adjust the Redundancy status accordingly. The definition of redundancy is power subsystem dependent and sometimes even configuration dependent.
This sensor is configured as a manual-rearm sensor in order to avoid the possibility of extraneous SEL events that could occur under certain system configuration and workload conditions. The sensor shall rearm for the following conditions:

- PSU hot-add
- system reset
- AC power cycle
- DC power cycle

System AC power is applied but on standby - Power unit redundancy is based on OEM SDR power unit record and number of PSU present.

System is (DC) powered on - The BMC calculates Dynamic Power Supply Redundancy status based upon current system load requirements as well as total Power Supply capacity.

The BMC allows redundancy to be configured on a per power-unit-redundancy sensor basis by means of the OEM SDR records.

7.3.17 Component Fault LED Control

Several sets of component fault LEDs are supported on the server board. See Figure 3. Intel® Light Guided Diagnostics - DIMM Fault LEDs and Figure 4. Intel® Light Guided Diagnostic LED Identification. Some LEDs are owned by the BMC and some by the BIOS.

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

Table 23. Component Fault LEDs

Component	Owner	Color	State	Description
Fan Fault LED	BMC	Amber	Solid On	Fan failed
		Amber	Off	Fan working correctly
DIMM Fault LED	BMC	Amber	Solid On	Memory failure – detected by 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.

- **Fan fault LEDs** – A fan fault LED is associated with each fan. 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.
- **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* owns the HW control for these LEDs and detection of the fault/status conditions that the LEDs reflect.
- **CPU Fault LEDs**. The BMC owns control for these LEDs. An LED is lit if there is an MSID mismatch (that is, CPU power rating is incompatible with the board)

7.3.18 NMI (Diagnostic Interrupt) Sensor

The BMC supports an NMI sensor for logging an event when a *diagnostic interrupt* is generated for the following cases:

- The front panel NMI (diagnostic interrupt) button is pressed
- The BMC receives an IPMI command Chassis Control command that requests this action

Note that the BMC may also generate this interrupt due to an IPMI Watchdog Timer pre-timeout interrupt; however an event for this occurrence is already logged against the Watchdog Timer sensor so it will not log an NMI sensor event.

7.3.19 LAN Leash Event Monitoring

The Physical Security sensor is used to monitor the LAN link and chassis intrusion status. This is implemented as a *LAN Leash* offset in this discrete sensor. This sensor monitors the link state of the two BMC embedded LAN channels. It does not monitor the state of any optional NICs.

The *LAN Leash Lost* offset asserts when one of the two BMC LAN channels loses a previously established link. It de-asserts when at least one LAN channel has a new link established after the previous assertion. No action is taken if a link has never been established.

LAN Leash events do not affect the front panel system status LED.

7.3.20 Add-in Module Presence Sensor

Some server boards provide dedicated slots for add-in modules/boards (for example, SAS, IO, PCIe*-riser). For these boards the BMC provides an individual presence sensor to indicate if the module/board is installed.

7.3.21 CMOS Battery Monitoring

The BMC monitors the voltage level from the CMOS battery, which provides battery backup to the chipset Real Time Clock. This is monitored as an auto-rearm threshold sensor.

Unlike monitoring of other voltage sources for which the Emulex* Pilot III component continuously cycles through each input, the voltage channel used for the battery monitoring provides a SW enable bit to allow the BMC FW to poll the battery voltage at a relatively slow rate in order to conserve battery power.

8. Intel® Intelligent Power Node Manager (NM) Support Overview

Power management deals with requirements to manage processor power consumption and manage power at the platform level to meet critical business needs. Node Manager (NM) is a platform resident technology that enforces power capping and thermal-triggered power capping policies for the platform. These policies are applied by exploiting subsystem knobs (such as processor P and T states) that can be used to control power consumption. NM enables data center power management by exposing an external interface to management software through which platform policies can be specified. It also implements specific data center power management usage models such as power limiting, and thermal monitoring.

The NM feature is implemented by a complementary architecture utilizing the ME, BMC, BIOS, and an ACPI-compliant OS. The ME provides the NM policy engine and power control/limiting functions (referred to as Node Manager or NM) while the BMC provides the external LAN link by which external management software can interact with the feature. The BIOS provides system power information utilized by the NM algorithms and also exports ACPI Source Language (ASL) code used by OS-Directed Power Management (OSPM) for negotiating processor P and T state changes for power limiting. PMBus*-compliant power supplies provide the capability to monitoring input power consumption, which is necessary to support NM.

The NM architecture applicable to this generation of servers is defined by the NPTM Architecture Specification v2.0. NPTM is an evolving technology that is expected to continue to add new capabilities that will be defined in subsequent versions of the specification. The ME NM implements the NPTM policy engine and control/monitoring algorithms defined in the Node Power and Thermal Manager (NPTM) specification.

8.1 Hardware Requirements

NM is supported only on platforms that have the NM FW functionality loaded and enabled on the Management Engine (ME) in the SSB and that have a BMC present to support the external LAN interface to the ME. NM power limiting features requires a means for the ME to monitor input power consumption for the platform. This capability is generally provided by means of PMBus*-compliant power supplies although an alternative model using a simpler SMBus* power monitoring device is possible (there is potential loss in accuracy and responsiveness using non-PMBus* devices). The NM SmaRT/CLST feature does specifically require PMBus*-compliant power supplies as well as additional hardware on the baseboard.

8.2 Features

NM provides feature support for policy management, monitoring and querying, alerts and notifications, and an external interface protocol. The policy management features implement specific IT goals that can be specified as policy directives for NM. Monitoring and querying features enable tracking of power consumption. Alerts and notifications provide the foundation for automation of power management in the data center management stack. The external interface specifies the protocols that must be supported in this version of NM.

8.3 ME System Management Bus (SMBus*) interface

- The ME uses the SMLink0 on the SSB in multi-master mode as a dedicated bus for communication with the BMC using the IPMB protocol. The BMC FW considers this a secondary IPMB bus and runs at 400 kHz.

- The ME uses the SMLink1 on the SSB in multi-master mode bus for communication with PMBus* devices in the power supplies for support of various NM-related features. This bus is shared with the BMC, which polls these PMBus* power supplies for sensor monitoring purposes (for example, power supply status, input power, and so on). This bus runs at 100 KHz.
- The Management Engine has access to the “Host SMBus*”.

8.4 PECI 3.0

- The BMC owns the PECI bus for all Intel server implementations and acts as a proxy for the ME when necessary.

8.5 NM “Discovery” OEM SDR

A NM “discovery” OEM SDR must be loaded into the BMC’s SDR repository if and only if the NM feature is supported on that product. This OEM SDR is used by management software to detect if NM is supported and to understand how to communicate with it.

Since PMBus* compliant power supplies are required in order to support NM, the system should be probed when the SDRs are loaded into the BMC’s SDR repository in order to determine whether or not the installed power supplies do in fact support PMBus*. If the installed power supplies are not PMBus* compliant then the NM “discovery” OEM SDR should not be loaded.

Please refer to the Intel® Intelligent Power Node Manager 2.0 External Architecture Specification using IPMI for details of this interface.

8.6 SmaRT/CLST

The power supply optimization provided by SmaRT/CLST relies on a platform HW capability as well as ME FW support. When a PMBus*-compliant power supply detects insufficient input voltage, an overcurrent condition, or an over-temperature condition, it will assert the SMBAlert# signal on the power supply SMBus* (such as, the PMBus*). Through the use of external gates, this results in a momentary assertion of the PROCHOT# and MEMHOT# signals to the processors, thereby throttling the processors and memory. The ME FW also sees the SMBAlert# assertion, queries the power supplies to determine the condition causing the assertion, and applies an algorithm to either release or prolong the throttling, based on the situation.

System power control modes include:

1. SmaRT: Low AC input voltage event; results in a one-time momentary throttle for each event to the maximum throttle state
2. Electrical Protection CLST: High output energy event; results in a throttling hiccup mode with fixed maximum throttle time and a fix throttle release ramp time.
3. Thermal Protection CLST: High power supply thermal event; results in a throttling hiccup mode with fixed maximum throttle time and a fix throttle release ramp time.

When the SMBAlert# signal is asserted, the fans will be gated by HW for a short period (~100ms) to reduce overall power consumption. It is expected that the interruption to the fans will be of short enough duration to avoid false lower threshold crossings for the fan tach sensors; however, this may need to be comprehended by the fan monitoring FW if it does have this side-effect.

ME FW will log an event into the SEL to indicate when the system has been throttled by the SmaRT/CLST power management feature. This is dependent on ME FW support for this sensor. Please reference the ME FW EPS for SEL log details.

8.6.1 Dependencies on PMBus*-compliant Power Supply Support

The SmaRT/CLST system feature depends on functionality present in the ME NM SKU. This feature requires power supplies that are compliant with the PMBus specification.

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

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

---

9. Basic and Advanced Server Management Features

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

Table 24. Intel® Remote Management Module 4 (RMM4) Options

Intel Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE	Intel® Remote Management Module 4 Lite	RMM4 Lite Activation Key	Enables KVM & media redirection

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 Basic and Advanced server management features.

Table 25. Basic and Advanced Server Management Features Overview

Feature	Basic	Advanced w/RMM4 Lite Key
IPMI 2.0 Feature Support	X	X
In-circuit BMC Firmware Update	X	X
FRB2	X	X
Chassis Intrusion Detection	X	X
Fan Redundancy Monitoring	X	X
Hot-Swap Fan Support	X	X
Acoustic Management	X	X
Diagnostic Beep Code Support	X	X
Power State Retention	X	X
ARP/DHCP Support	X	X
PECI Thermal Management Support	X	X
E-mail Alerting	X	X
Embedded Web Server	X	X
SSH Support	X	X
Integrated KVM		X
Integrated Remote Media Redirection		X
Lightweight Directory Access Protocol (LDAP)	X	X
Intel® Intelligent Power Node Manager Support	X	X
SMASH CLP	X	X

On the server board the Intel® RMM4 Lite key is installed at the following location.

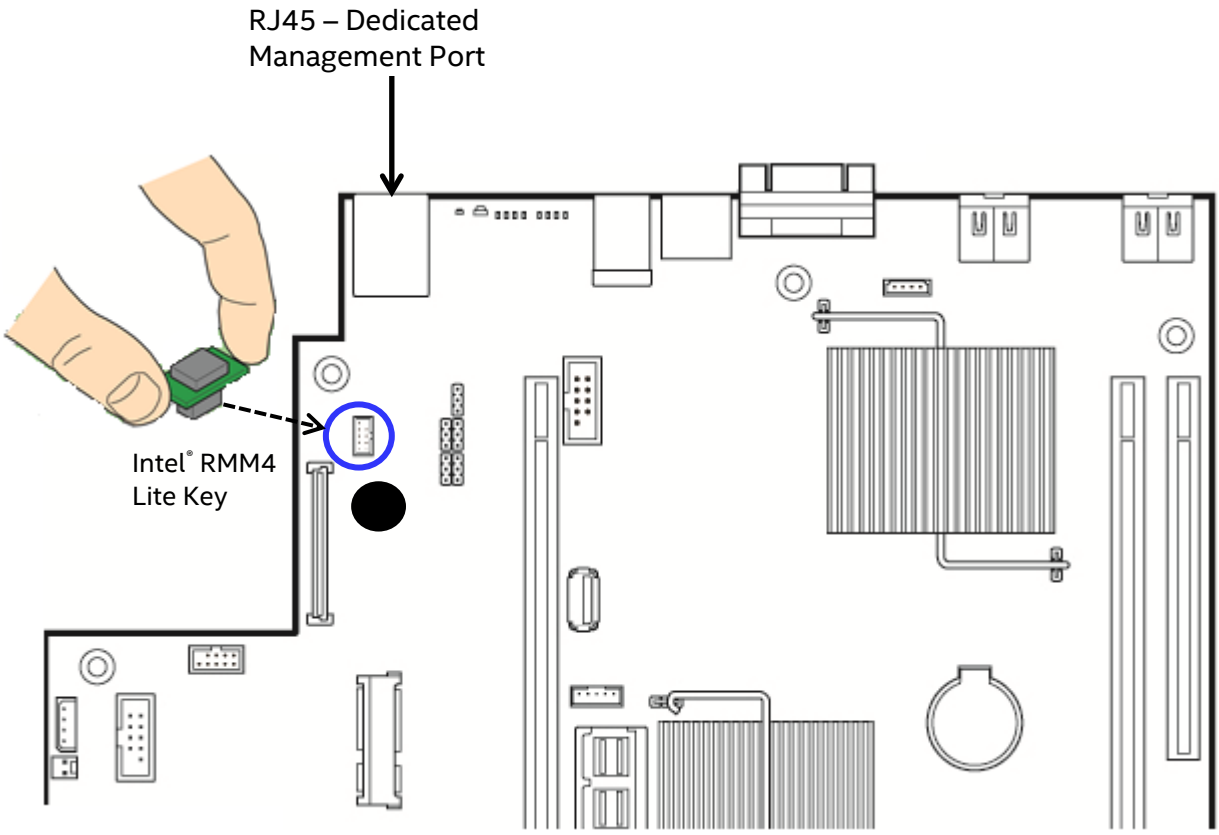


Figure 28. Intel® RMM4 Lite Activation Key Installation

9.1 Dedicated Management Port

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

9.2 Embedded Web Server

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

- Microsoft Internet Explorer 9.0*
- Microsoft Internet Explorer 10.0*
- Mozilla Firefox 24*
- Mozilla Firefox 25*

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

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

Additional features supported by the web GUI include:

- 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.
- Configuration of various IPMI over LAN parameters for both IPV4 and IPV6
- Configuration of alerting (SNMP and SMTP)
- Display system asset information for the product, board, and chassis.
- Display 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)
- Automatic refresh of sensor data with a configurable refresh rate
- Online help
- Display/clear SEL (display is in easily understandable human readable format)
- Support major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*)
- The GUI session automatically times out after a user-configurable inactivity period. By default, this inactivity period is 30 minutes.
- Embedded Platform Debug feature - Allow the user to initiate a "debug dump" to a file that can be sent to Intel® for debug purposes.
- Virtual Front Panel. The Virtual Front Panel provides the same functionality as the local front panel. The displayed LEDs match the current state of the local panel LEDs. The displayed buttons (for example, power button) can be used in the same manner as the local buttons.
- Display of ME sensor data. Only sensors that have associated SDRs loaded will be displayed.
- Ability to save the SEL to a file
- Ability to force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display of processor and memory information that is available over IPMI over LAN.
- Ability to get and set Node Manager (NM) power policies
- Display of power consumed by the server
- Ability to view and configure VLAN settings
- Warn user the reconfiguration of IP address will cause disconnect.
- Capability to block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.

- Server Power Control - Ability to force into Setup on a reset
- System POST results – The web server provides the system's Power-On Self Test (POST) sequence for the previous two boot cycles, including timestamps. The timestamps may be displayed as a time relative to the start of POST or the previous POST code.
- Customizable ports - The web server provides the ability to customize the port numbers used for SMASH, http, https, KVM, secure KVM, remote media, and secure remote media..

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

9.3 Advanced Management Feature Support (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 add-on offers convenient, remote KVM access and control through LAN and internet. It captures, digitizes, and compresses video and transmits it with keyboard and mouse signals to and from a remote computer. Remote access and control software runs in the integrated baseboard management controller, utilizing expanded capabilities enabled by the Intel RMM4 hardware.

Key Features of the RMM4 add-on are:

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

9.3.1 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 60Hx, 72Hz, 75Hz, 85Hz
- 1280x960 at 60Hz
- 1280x1024 at 60Hz
- 1600x1200 at 60Hz
- 1920x1080 (1080p)
- 1920x1200 (WUXGA)
- 1650x1080 (WSXGA+)

9.3.2 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.

9.3.3 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.

9.3.4 Security

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

9.3.5 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. An BMC reset (for example, due to an 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.

9.3.6 Usage

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

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

9.3.7 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.

9.3.8 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.

9.3.8.1 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.

9.3.8.2 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.

10. On-board Connector/Header Overview

This section identifies the location and pin-out for on-board connectors and headers of the server board that provide an interface to system options/features, on-board platform management, or other user accessible options/features.

10.1 Power Connectors

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

10.1.1 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. 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. The following tables provide the pin-out for both “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
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

10.1.2 Hot Swap Backplane Power Connector

The server board includes one 8-pin power connector that can be cabled to provide power for hot swap backplanes. On the server board, this connector is labeled as “HSBP PWR”. The following table provides the pin-out for this connector.

Table 28. Hot Swap Backplane Power Connector Pin-out (“HSBP PWR")

Signal Name	Pin#	Pin#	Signal Name
P12V_240VA1	5	1	GROUND
P12V_240VA1	6	2	GROUND
P12V_240VA2	7	3	GROUND
P12V_240VA2	8	4	GROUND

10.1.3 Peripheral Drive 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.

Table 29. 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

10.1.4 Riser Card Supplemental 12V Power Connectors

The server board includes two white 2x2-pin power connectors that provide supplemental power 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. A cable from this connector may be routed to a power 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 rest of the system. 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.

Table 30. 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 below) is available as a separate orderable accessory kit from Intel using the following Intel product code: *AXXGPGPUCABLE*.

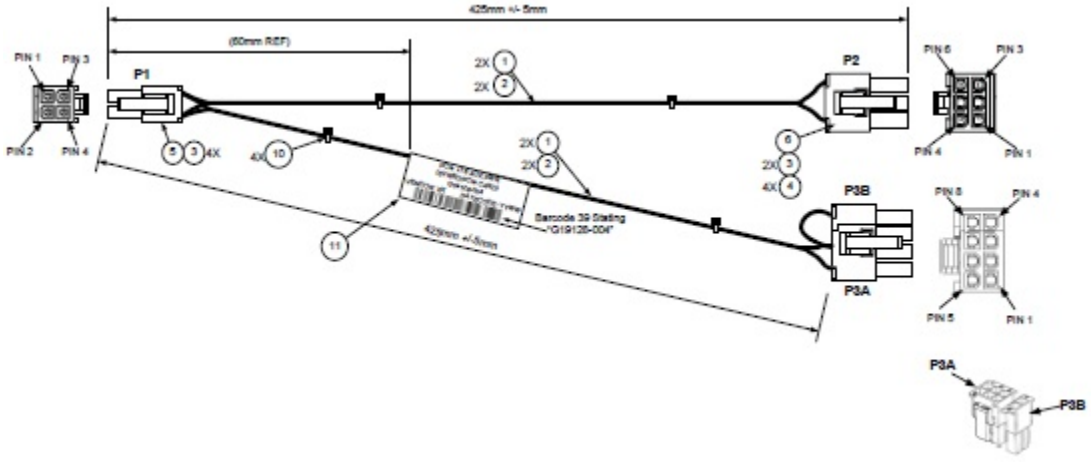


Figure 29. High Power Add-in Card 12V Auxiliary Power Cable Option

10.2 Front 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.

10.2.1 Front Panel Button and LED Support

Included near the right front edge of the server board are two front panel connectors:

- Standard 30-pin header “FRONT_PANEL” - SSI compatible
- Custom 30-pin high density “STORAGE_FP” – Used on storage models of Intel server systems with a rack handle mounted front panel

Each connector provides an interface supporting system control buttons and LEDs. The following table identifies the supported button and LED features supported from each front panel connector.

Table 31. Front Panel Features

	SSI Front Panel	Storage Front Panel
Power / Sleep Button	Yes	Yes
System ID Button	Yes	Yes
System Reset Button	Yes	Yes
NMI Button	Yes	Yes
NIC Activity LED	Yes	Yes
Storage Device Activity LED	Yes	Yes
System Status LED	Yes	Yes
System ID LED	Yes	Yes

The pinout is identical for both front panel connectors.

Table 32. Front Panel Connector Pin-out ("Front Panel" and “Storage FP”)

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

10.2.2 Front Panel LED and Control Button Features Overview

10.2.2.1 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 the following table.

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.

10.2.2.2 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.

10.2.2.3 System Reset Button Support

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

10.2.2.4 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.

The following table 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

10.2.2.5 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 turn on solid. The LED will blink once network activity occurs at a rate that is consistent with the amount of network activity that is occurring.

10.2.2.6 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.

10.2.2.7 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. See section 12.2 for a list of supported System Status LED states.

10.2.3 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 of the server board near the I/O module connector. The following table provides the connector pin-out.

Note: The numbers 5 & 6 in the silk screen label identify the USB ports routed to this connector.

Table 35. 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

10.2.4 Front Panel USB 3.0 Connector

The server board includes a Blue 20-pin connector that, when cabled, can provide up to two USB 2.0 / 3.0 ports to a front panel. On the server board the connector is labeled "FP_USB_2.0/3.0" and is located near the Main Power #1 connector. The following table provides the connector pin-out.

Note: The following USB ports are routed to this connector: USB 3.0 ports 1 and 4, USB 2.0 ports 10 and 13

Table 36. 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

Note: Due to signal strength limits associated with USB 3.0 ports cabled to a front panel, some marginally compliant USB 3.0 devices may not be supported from these ports.

10.2.5 Front Panel Video Connector

The server board includes a 14-pin header that, when cabled, can provide an alternate video connector to the front panel. On the server board this connector is labeled "FP_VIDEO" and is located near the right edge of the board next to the 30-pin front panel connector. When a monitor is attached to the front panel video connector, the external video connector located on the back edge of the board is disabled. The following table provides the pin-out for this connector.

Table 37. 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
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

10.2.6 Intel® Local Control Panel Connector

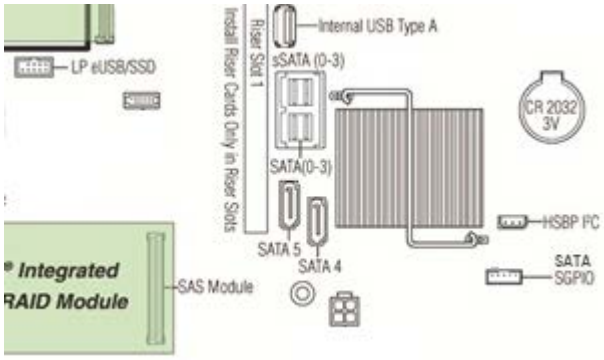
The server board includes a white 7-pin connector that is used when the system is configured with the Intel® Local Control Panel with LCD support. On the server board this connector is labeled "LCP" and is located on the right edge of the server board. The following table provides the pin-out for this connector.

Table 38. Intel Local Control Panel Connector Pin-out ("LCP")

Signal Description	Pin#
SMB_SENSOR_3V3STBY_DATA_R0	1
GROUND	2
SMB_SENSOR_3V3STBY_CLK	3
P3V3_AUX	4
FM_LCP_ENTER_N_R	5
FM_LCP_LEFT_N_R	6
FM_LCP_RIGHT_N_R	7

10.3 On-Board Storage Option Connectors

The server board provides connectors to support several storage device options. This section provides a functional overview and pin-out of each connector.



10.3.1 Single Port SATA Only Connectors

The server board includes two white single port SATA only connectors capable of transfer rates of up to 6Gb/s. On the server board these connectors are labeled as “SATA 4” and “SATA 5”. The following table provides the pin-out for both connectors.

Table 39. Single Port SATA Connector Pin-out ("SATA 4" & "SATA 5")

Signal Description	Pin#
GROUND	1
SATA_TXP	2
SATA_TXN	3
GROUND	4
SATA_RXN	5
SATA_RXP	6
GROUND	7

10.3.1.1 SATA SGPIO Connector

The server board includes a 5-pin SATA SGPIO connector. When cabled to a hot-swap backplane, this connector provides drive fault LED support for the single onboard SATA ports (SATA_4 and SATA_5). The connector has the following pin-out:

Table 40. SATA SGPIO Connector Pin-out ("SATA_SGPIO")

Signal Description	Pin#
SGPIO SATA CLK	1
SGPIO SATA LOAD	2
GROUND	3
SGPIO SATA DATA OUT	4
PU-SGPIO SATA	5

10.3.2 Internal Type-A USB Connector

The server board includes one internal Type-A USB connector labeled “USB 2.0” and is located to the right of Riser Slot #1. The following table provides the pin-out for this connector.

Note: The following USB 2.0 port is routed to this connector: USB 2.0 port 9

Table 41. Internal Type-A USB Connector Pin-out ("USB 2.0")

Signal Description	Pin#
P5V_USB_INT	1
USB2_P2_F_DN	2
USB2_P2_F_DP	3
GROUND	4

10.3.3 Internal 2mm Low Profile eUSB SSD Connector

The server board includes one 10-pin 2mm low profile connector with an intended usage of supporting low profile eUSB SSD devices. On the server board this connector is labeled “eUSB SSD”. The following table provides the pin-out for this connector.

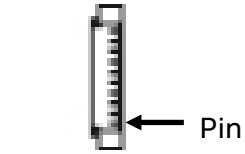
Note: The following USB 2.0 port is routed to this connector: USB 2.0 port 8

Table 42. Internal eUSB Connector Pin-out ("eUSB SSD")

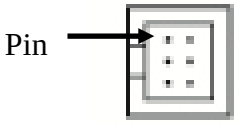
Signal Description	Pin#	Pin#	Signal Description
P5V	1	2	NOT USED
USB2_P0_DN	3	4	NOT USED
USB2_P0_DP	5	6	NOT USED
GROUND	7	8	NOT USED
NOT USED	9	10	LED_HDD_ACT_N

10.4 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.



Fixed Mount Fan



Hot Swap Fan

Dual Rotor Fixed SYS_FAN # (1-6)	
Signal Description	Pin#
LED_FAN	10
LED_FAN_FAULT	9
SYS FAN PRSNT	8
GROUND	7
GROUND	6
FAN_TACH_#	5
P12V_FAN	4
P12V_FAN	3
FAN PWM	2
FAN_TACH_#+1	1

Hot Swap SYS_FAN # (1-6)			
Signal Description	Pin#	Pin#	Signal Description
GROUND	1	2	P12V FAN
FAN TACH	3	4	FAN PWM
SYS FAN PRSNT	5	6	LED FAN FAULT

Figure 30. 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.

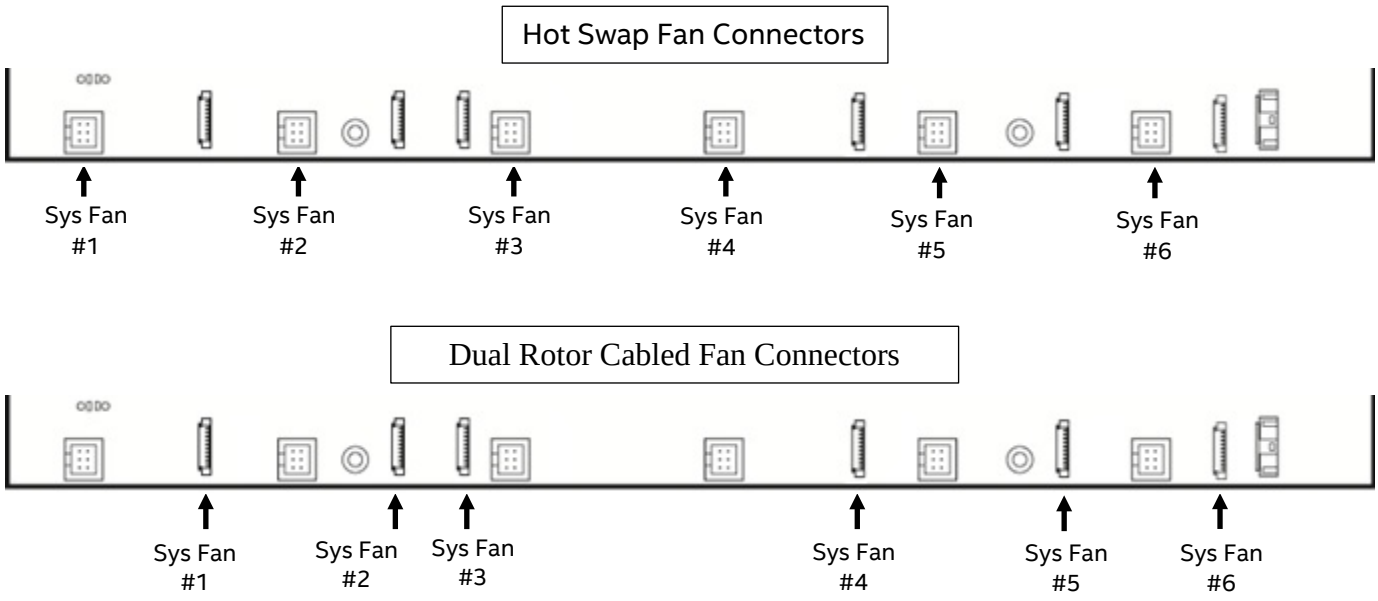


Figure 31. System Fan Connector Placement

10.5 Other Connectors and Headers

10.5.1 Chassis Intrusion Header

The server board includes a 2-pin chassis intrusion header which can be used when the chassis is configured with a chassis intrusion switch and the proper platform management SDR is programmed and installed. On the server board, this header is labeled “CHAS_INTR” and is located on the right edge of the server board. The header has the following pin-out.

Table 43. Chassis Intrusion Header Pin-out ("CHAS_INTR")

Signal Description	Pin#
FP_CHASSIS_INTRUSION	1
GROUND	2

If configured, the BMC can monitor the state of the *Chassis Intrusion* signal and makes the status of the signal available through the *Get Chassis Status* command and the *Physical Security* sensor state. A chassis intrusion state change causes the BMC to generate a *Physical Security* sensor event message with a *General Chassis Intrusion* offset (00h).

The BMC detects chassis intrusion and logs a SEL event when the system is in the on, sleep, or standby state. Chassis intrusion is not detected when the system is in an AC power-off (AC lost) state.

The BMC hardware cannot differentiate between a missing chassis intrusion cable or connector, and a true security violation. If the chassis intrusion cable or connector is removed or damaged, the BMC treats it as if the chassis cover is open, and takes the appropriate actions.

System fans can be set to boost to maintain proper system cooling when a chassis intrusion is detected.

10.5.2 Storage Device Activity LED Header

The server board includes a 2-pin storage device activity LED header used with some SAS/SATA controller add-in cards. On the server board, this header is labeled “HDD LED” and is located on the left edge of the server board. The header has the following pin-out.

Table 44. Hard Drive Activity Header Pin-out ("HDD_LED")

Signal Description	Pin#
LED_HDD_ACT_N	1
TP_LED_HDD_ACT	2

10.5.3 Intelligent Platform Management Bus (IPMB) Connector

The Intelligent Platform Management Bus (IPMB) is designed to be incorporated into mission critical server platforms for the main purpose of supporting Server Platform Management. The server board includes a 4-pin Intelligent Platform Management Bus (IPMB) connector located on the left edge of the server board. The connector has the following pin-out.

Table 45. IPMB Connector Pin-out

Signal Description	Pin#
IPMB Data	1
Ground	2
IPMB Clock	3
5V AUX	4

10.5.4 Hot Swap Backplane I2C* Connectors

The server board includes two 3-pin hot swap backplane I2C* connectors. These are located near the center of the board near the chipset heat sink, and towards the front left side of the board. Each is labeled as "HSBP I2C". When cabled, these connectors provide a communication path for the onboard BMC to a hot swap backplane, allowing for firmware updates and other platform management functions. These connectors have the following pin-out.

Table 46. Hot-Swap Backplane I2C* Connector Pin-out

Signal Description	Pin#
HSBP Data	1
Ground	2
HSBP Clock	3

10.5.5 SMBus Connector

The server board includes a 3-pin SMBus connector. This connector is located near the front left corner of the server board and is labeled "SMBus". When cabled, this connector is used as an interface to the embedded server management bus.

Table 47. SMBus Connector Pin-out

Signal Description	Pin#
SMB Data	1
Ground	2
SMB Clock	3

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. The following diagram identifies the location of each jumper block on the server board. Pin 1 of each jumper block can be identified by the "▼" silkscreened on the server board next to the pin.

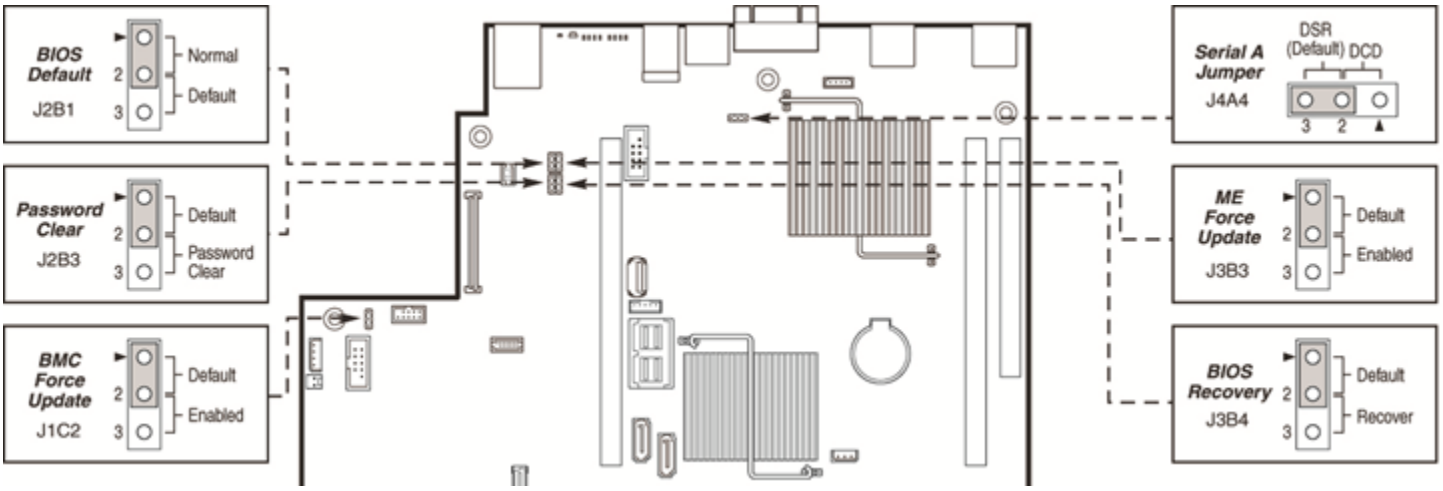


Figure 32. Reset and Recovery Jumper Block Location

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

Note: The system will automatically power on after AC is applied to the system.

6. During POST, access the <F2> BIOS Setup utility to configure and save desired BIOS options

Note:
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

Note also that the system time and date may need to be reset.

11.2 Serial Port 'A' Configuration Jumper

See section 5.10 for details

11.3 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

11.4 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. The following procedure should be followed.

Note: System Update files are included in the System Update Packages (SUP) posted to Intel's Download center web site. 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 successfully completed, 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 system

11.5 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. The following procedure should be followed:

Note: System Update files are included in the System Update Packages (SUP) posted to Intel's Download center web site. 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

11.6 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 web site. 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 web site. 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. 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.

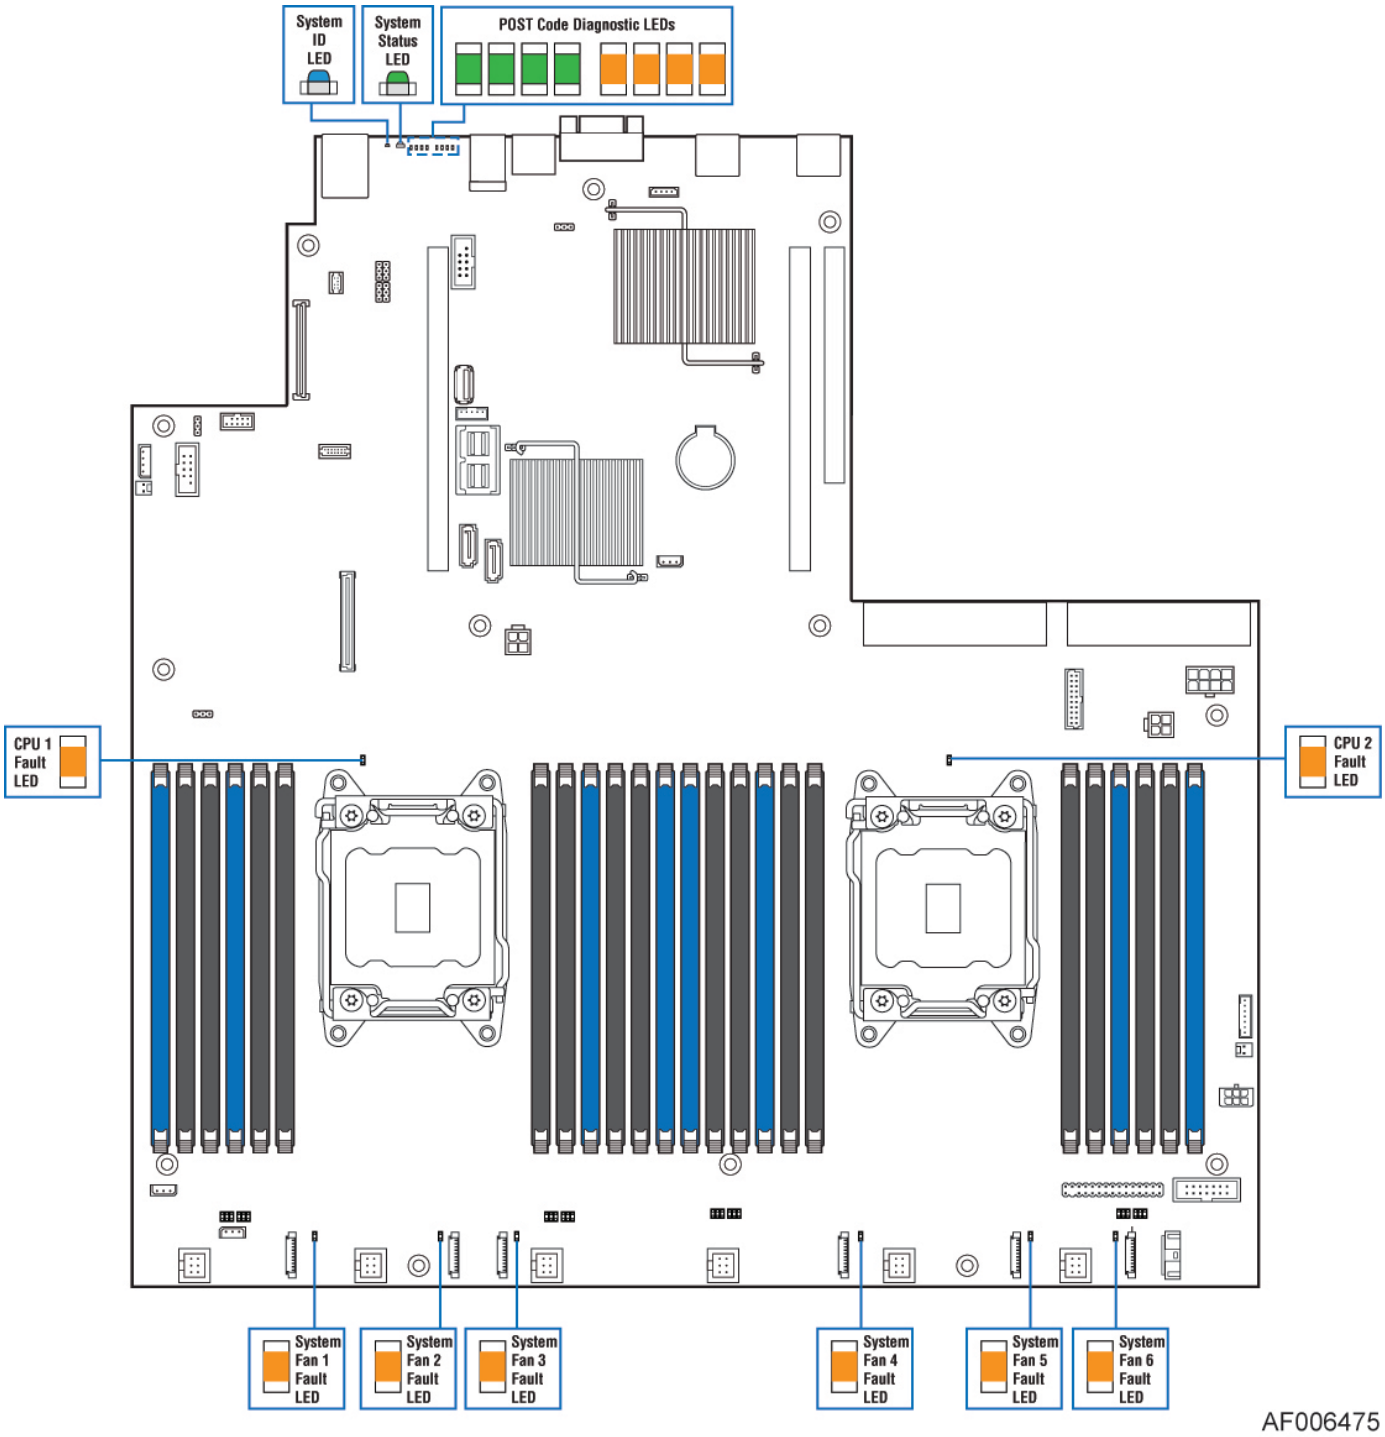


Figure 33. On-Board Diagnostic LED Placement

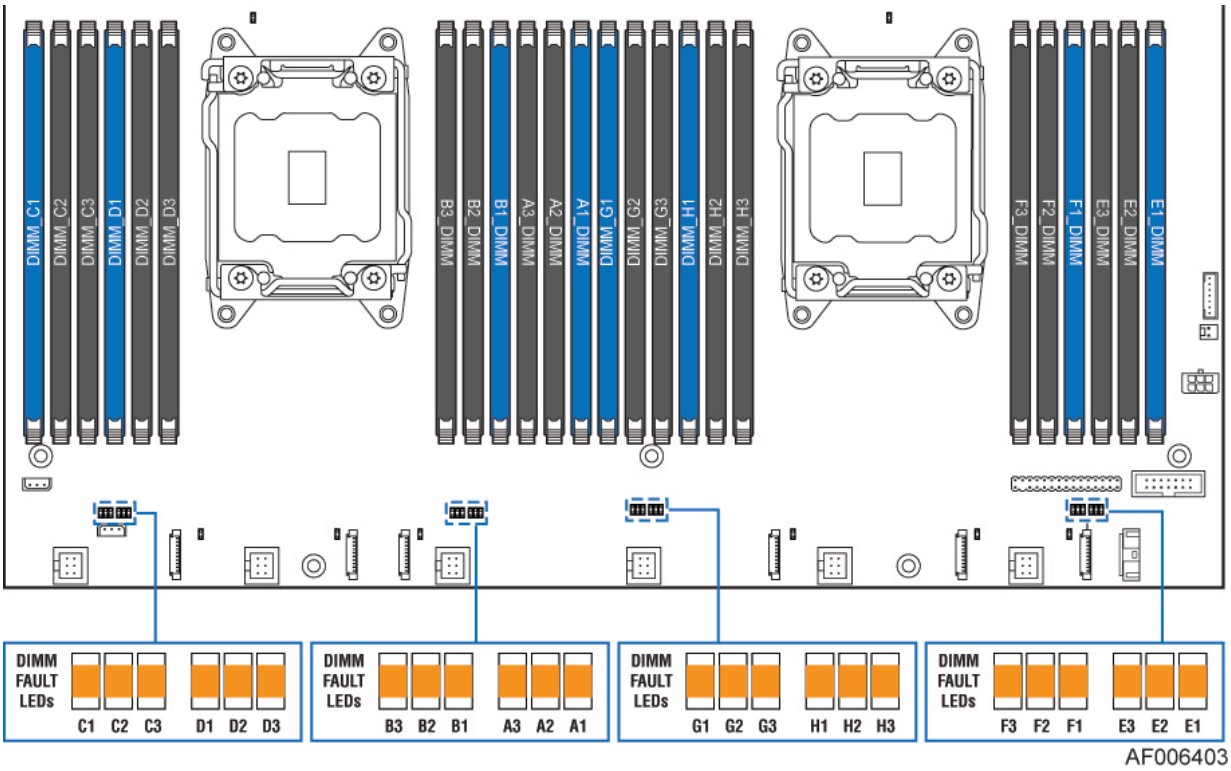


Figure 34. DIMM Fault LED Placement

12.1 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 remotely entered, 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.

12.2 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 48. System Status LED State Definitions

Color	State	Criticality	Description
Off	System is not operating	Not ready	• System is powered off (AC and/or DC). • System is in EuP Lot6 Off Mode. • System is in S5 Soft-Off State.
Green	Solid on	Ok	Indicates that the System is running (in S0 State) and its status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running. After a BMC reset, and in conjuction 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 - system is operating in a degraded state although still functional, *or* system is operating in a redundant state but with an impending failure warning	System degraded: • Redundancy loss such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. • Fan warning or failure when the number of fully operational fans is less than minimum number needed to cool the system. • Non-critical threshold crossed – Temperature (including HSBP temp), voltage, input power to power supply, output current for main power rail from power supply and Processor Thermal Control (Therm Ctrl) sensors. • Power supply predictive failure occurred while redundant power supply configuration was present. • Unable to use all of the installed memory (more than 1 DIMM installed). • Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the system no longer has spared DIMMs (a redundancy lost condition). Corresponding DIMM LED lit. • In mirrored configuration, when memory mirroring takes place and system loses memory redundancy. • Battery failure. • BMC executing in uBoot. (Indicated by Chassis ID blinking at 3Hz). System in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux*. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux* image into flash. • BMC Watchdog has reset the BMC. • Power Unit sensor offset for configuration error is asserted. • HDD HSC is off-line or degraded.
Amber	~1 Hz blink	Non-critical - System is operating in a degraded state	Non-fatal alarm – system is likely to fail: • Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply

Color	State	Criticality	Description
		with an impending failure warning, although still functioning	and PROCHOT (Therm Ctrl) sensors. • VRD Hot asserted. • Minimum number of fans to cool the system not present or failed • Hard drive fault • Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present) • In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window
Amber	Solid on	Critical, non-recoverable – System is halted	Fatal alarm – system has failed or shutdown: • CPU CATERR signal asserted • MSID mismatch detected (CATERR also asserts for this case). • CPU 1 is missing • CPU Thermal Trip • No power good – power fault • DIMM failure when there is only 1 DIMM present and hence no good memory present. • Runtime memory uncorrectable error in non-redundant mode. • DIMM Thermal Trip or equivalent • SSB Thermal Trip or equivalent • CPU ERR2 signal asserted • BMC/Video memory test failed. (Chassis ID shows blue/solid-on for this condition) • Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition) • 240VA fault • Fatal Error in processor initialization: o Processor family not identical o Processor model not identical o Processor core/thread counts not identical o Processor cache size not identical o Unable to synchronize processor frequency o Unable to synchronize QPI link frequency • Uncorrectable memory error in a non-redundant mode

12.3 BMC Boot/Reset Status LED Indicators

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

Table 49. 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.

12.4 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. 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 D for a complete description of how these LEDs are read, and for a list of all supported POST codes

12.5 Fan Fault LEDs

The server board includes a Fan Fault LED next to each of the six system fan. 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.

12.6 Memory Fault LEDs

The server board includes a Memory Fault LED for each DIMM slot. 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.

12.7 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).

13. Power Supply Specification Guidelines

This section provides power supply specification guidelines recommended for providing the specified server platform with stable operating power requirements.

Note: The power supply data provided in this section is for reference purposes only. It reflects Intel's own DC power out requirements for a 750W power supply as used in an Intel designed 2U server platform. The intent of this section is to provide customers with a guide to assist in defining and/or selecting a power supply for custom server platform designs that utilize the server board detailed in this document.

13.1 Power Supply DC Output Connector

The server board includes two main power slot connectors allowing for power supplies to attach directly to the server board. Power supplies must utilize a card edge output connection for power and signal that is compatible with a 2x25 Power Card Edge connector (equivalent to 2x25 pin configuration of the FCI power card connector 10035388-102LF).

Table 50. Power Supply DC Power Output Connector Pinout

Pin	Name	Pin	Name
A1	GND	B1	GND
A2	GND	B2	GND
A3	GND	B3	GND
A4	GND	B4	GND
A5	GND	B5	GND
A6	GND	B6	GND
A7	GND	B7	GND
A8	GND	B8	GND
A9	GND	B9	GND
A10	+12V	B10	+12V
A11	+12V	B11	+12V
A12	+12V	B12	+12V
A13	+12V	B13	+12V
A14	+12V	B14	+12V
A15	+12V	B15	+12V
A16	+12V	B16	+12V
A17	+12V	B17	+12V
A18	+12V	B18	+12V
A19	PMBus SDA	B19	A0 (SMBus address)
A20	PMBus SCL	B20	A1 (SMBus address)
A21	PSON	B21	12V stby
A22	SMBAlert#	B22	Cold Redundancy Bus
A23	Return Sense	B23	12V load share bus
A24	+12V remote Sense	B24	No Connect
A25	PWOK	B25	Compatibility Check pin*

13.2 Power Supply DC Output Specification

13.2.1 Output Power/Currents

The following tables define the minimum power and current ratings. The power supply must meet both static and dynamic voltage regulation requirements for all conditions.

Table 51. Minimum Load Ratings

Parameter	Min	Max.	Peak [2, 3]	Unit
12V main	0.0	62.0	70.0	A
12Vstby [1]	0.0	2.1	2.4	A

Notes:

1. 12Vstby must provide 4.0A with two power supplies in parallel. The Fan may start to work when stby current >1.5A
2. Peak combined power for all outputs shall not exceed 850W.
3. Length of time peak power can be supported is based on thermal sensor and assertion of the SMBAlert# signal. Minimum peak power duration shall be 20 seconds without asserting the SMBAlert# signal at maximum operating temperature.

13.2.2 Standby Output

The 12VSB output shall be present when an AC input greater than the power supply turn on voltage is applied. There should be load sharing in the standby rail. Two PSU modules should be able to support 4A standby current.

13.2.3 Voltage Regulation

The power supply output voltages must stay within the following voltage limits when operating at steady state and dynamic loading conditions. These limits include the peak-peak ripple/noise. These shall be measured at the output connectors.

Table 52. Voltage Regulation Limits

PARAMETER	TOLERANCE	MIN	NOM	MAX	UNITS
+12V	- 5%/+5%	+11.40	+12.00	+12.60	V_{rms}
+12V stby	- 5%/+5%	+11.40	+12.00	+12.60	V_{rms}

13.2.4 Dynamic Loading

The output voltages shall remain within limits specified for the step loading and capacitive loading specified in the table below. The load transient repetition rate shall be tested between 50Hz and 5kHz at duty cycles ranging from 10%-90%. The load transient repetition rate is only a test specification. The Δ step load may occur anywhere within the MIN load to the MAX load conditions.

Table 53. Transient Load Requirements

Output	Δ Step Load Size	Load Slew Rate	Test capacitive Load
+12VSB	1.0A	0.25 A/μsec	20 μF
+12V	60% of max load	0.25 A/μsec	2000 μF

Note: For dynamic condition +12V MIN loading is 1A.

13.2.5 Capacitive Loading

The power supply shall be stable and meet all requirements with the following capacitive loading ranges.

Table 54. Capacitive Loading Conditions

Output	MIN	MAX	Units
+12VSB	20	3100	μF
+12V	500	25000	μF

13.2.6 Grounding

The output ground of the pins of the power supply provides the output power return path. The output connector ground pins shall be connected to the safety ground (power supply enclosure). This grounding should be well designed to ensure passing the max allowed Common Mode Noise levels.

The power supply shall be provided with a reliable protective earth ground. All secondary circuits shall be connected to protective earth ground. Resistance of the ground returns to chassis shall not exceed 1.0 mΩ. This path may be used to carry DC current.

13.2.7 Closed loop stability

The power supply shall be unconditionally stable under all line/load/transient load conditions including specified capacitive load ranges. A minimum of: **45 degrees phase margin** and -**10dB-gain margin** is required. Closed-loop stability must be ensured at the maximum and minimum loads as applicable.

13.2.8 Residual Voltage Immunity in Standby mode

The power supply should be immune to any residual voltage placed on its outputs (Typically a leakage voltage through the system from standby output) up to **500mV**. There shall be no additional heat generated, nor stressing of any internal components with this voltage applied to any individual or all outputs simultaneously. It also should not trip the protection circuits during turn on.
The residual voltage at the power supply outputs for no load condition shall not exceed **100mV** when AC voltage is applied and the PSON# signal is de-asserted.

13.2.9 Common Mode Noise

The Common Mode noise on any output shall not exceed **350mV pk-pk** over the frequency band of 10Hz to 20MHz.

13.2.10 Soft Starting

The Power Supply shall contain a control circuit which provides monotonic soft start for its outputs without overstress of the AC line or any power supply components at any specified AC line or load conditions.

13.2.11 Zero Load Stability Requirements

When the power subsystem operates in a no load condition, it does not need to meet the output regulation specification, but it must operate without any tripping of over-voltage or other fault circuitry. When the power subsystem is subsequently loaded, it must begin to regulate and source current without fault.

13.2.12 Hot Swap Requirements

Hot swapping a power supply is the process of inserting and extracting a power supply from an operating power system. During this process the output voltages shall remain within the limits with the capacitive load specified. The hot swap test must be conducted when the system is operating under static, dynamic, and zero loading conditions.

13.2.13 Forced Load Sharing

The +12V output will have active load sharing. The output will share within 10% at full load. The failure of a power supply should not affect the load sharing or output voltages of the other supplies still operating. The

supplies must be able to load share in parallel and operate in a hot-swap/redundant **1+1** configurations. The 12VSBoutput is not required to actively share current between power supplies (passive sharing). The 12VSBoutput of the power supplies are connected together in the system so that a failure or hot swap of a redundant power supply does not cause these outputs to go out of regulation in the system.

13.2.14 Ripple/Noise

The maximum allowed ripple/noise output of the power supply is defined in the following table. This is measured over a bandwidth of 10Hz to 20MHz at the power supply output connectors. A 10μF tantalum capacitor in parallel with a 0.1μF ceramic capacitor is placed at the point of measurement.

Table 55. Ripples and Noise

+12V main	+12VSB
120mVp-p	120mVp-p

13.2.15 Timing Requirements

These are the timing requirements for the power supply operation. The output voltages must rise from 10% to within regulation limits (T_{vout_rise}) within 5 to 70ms. For 12VSB, it is allowed to rise from 1.0 to 25ms. **All outputs must rise monotonically**. The following table shows the timing requirements for the power supply being turned on and off from the AC input, with PSON held low and the PSON signal, with the AC input applied.

Table 56. Timing Requirements

Item	Description	MIN	MAX	UNITS
T_{vout_rise}	Output voltage rise time	5.0 *	70 *	ms
Tsb_on_delay	Delay from AC being applied to 12VSBbeing within regulation.		1500	ms
T ac_on_delay	Delay from AC being applied to all output voltages being within regulation.		3000	ms
Tvout_holdup	Time 12Vl output voltage stay within regulation after loss of AC.	13		ms
Tpwok_holdup	Delay from loss of AC to de-assertion of PWOK	12		ms
Tpson_on_delay	Delay from PSON# active to output voltages within regulation limits.	5	400	ms
T pson_pwok	Delay from PSON# deactivate to PWOK being de-asserted.		5	ms
Tpwok_on	Delay from output voltages within regulation limits to PWOK asserted at turn on.	100	500	ms
T pwok_off	Delay from PWOK de-asserted to output voltages dropping out of regulation limits.	1		ms
Tpwok_low	Duration of PWOK being in the de-asserted state during an off/on cycle using AC or the PSON signal.	100		ms
Tsb_vout	Delay from 12VSBbeing in regulation to O/Ps being in regulation at AC turn on.	50	1000	ms
T12VSB_holdup	Time the 12VSBoutput voltage stays within regulation after loss of AC.	70		ms

* The 12VSBoutput voltage rise time shall be from 1.0ms to 25ms

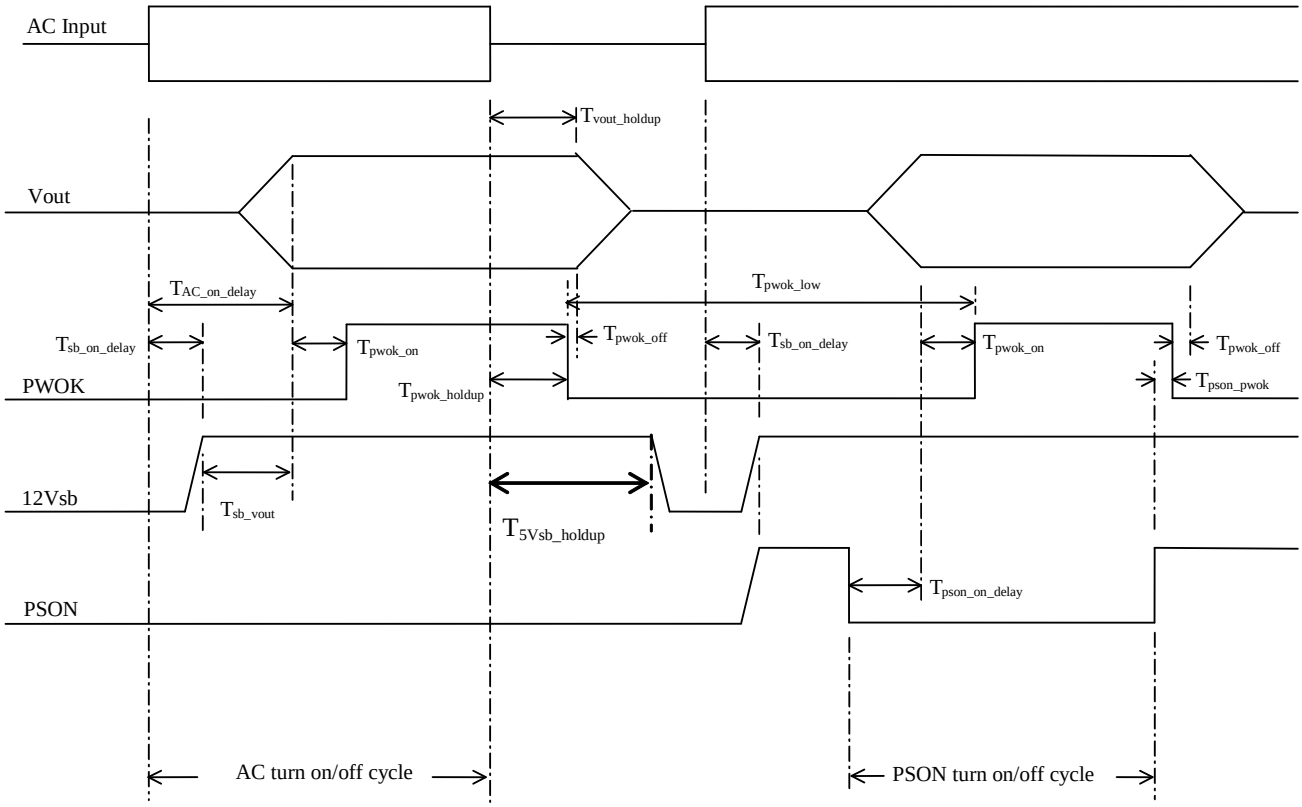


Figure 35. Turn On/Off Timing (Power Supply Signals)

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 E5-2600 v3 product family with a Thermal Design Power (TDP) of up to and including 145 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 bottom add-in card slot of 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.
- This server board only supports DDR4 ECC RDIMM – Registered (Buffered) DIMMS and DDR4 ECC LRDIMM – Load Reduced DIMMs
- 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: Integrated BMC Sensor Tables

This appendix provides BMC core sensor information common to all Intel server boards within this generation of product. Specific server boards and/or server platforms may only implement a sub-set of sensors and/or may include additional sensors. The actual sensor name associated with a sensor number may vary between server boards or systems.

- Sensor Type

  The Sensor Type values are the values enumerated in the *Sensor Type Codes* table in the IPMI specification. The Sensor Type provides the context in which to interpret the sensor, such as the physical entity or characteristic that is represented by this sensor.

- Event/Reading Type

  The Event/Reading Type values are from the *Event/Reading Type Code Ranges* and *Generic Event/Reading Type Codes* tables in the IPMI specification. Digital sensors are a specific type of discrete sensor, which have only two states.

- Event Offset/Triggers

  Event Thresholds are event-generating thresholds for threshold types of sensors.

  - [u,l][nr,c,nc]: upper non-recoverable, upper critical, upper non-critical, lower non-recoverable, lower critical, lower non-critical
  - uc, lc: upper critical, lower critical

  Event Triggers are supported event-generating offsets for discrete type sensors. The offsets can be found in the *Generic Event/Reading Type Codes* or *Sensor Type Codes* tables in the IPMI specification, depending on whether the sensor event/reading type is generic or a sensor-specific response.

- Assertion/De-assertion Enables

  Assertion and de-assertion indicators reveal the type of events the sensor generates:

  - As: Assertions
  - De: De-assertion

- Readable Value/Offsets
  - Readable Value indicates the type of value returned for threshold and other non-discrete type sensors.
  - Readable Offsets indicate the offsets for discrete sensors that are readable with the *Get Sensor Reading* command. Unless otherwise indicated, all event triggers are readable; Readable Offsets consist of the reading type offsets that do not generate events.

- Event Data

  Event data is the data that is included in an event message generated by the sensor. For threshold-based sensors, the following abbreviations are used:

  - R: Reading value
  - T: Threshold value

- Rearm Sensors

  The rearm is a request for the event status for a sensor to be rechecked and updated upon a transition between good and bad states. Rearming the sensors can be done manually or automatically. This column indicates the type supported by the sensor. The following abbreviations are used to describe a sensor:

  - A: Auto-rearm
  - M: Manual rearm

- Default Hysteresis

  The hysteresis setting applies to all thresholds of the sensor. This column provides the count of hysteresis for the sensor, which can be 1 or 2 (positive or negative hysteresis).

- Criticality

  Criticality is a classification of the severity and nature of the condition. It also controls the behavior of the Control Panel Status LED.

- Standby

  Some sensors operate on standby power. These sensors may be accessed and/or generate events when the main (system) power is off, but AC power is present.

Note: All sensors listed below may not be present on all platforms. Please reference the BMC EPS for platform applicability. Redundancy sensors will only be present on systems with appropriate hardware to support redundancy (for instance, fan or power supply)

Table 57. BMC Core Sensors

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Unit Status (Pwr Unit Status)	01h	All	Power Unit 09h	Sensor Specific 6Fh	00 - Power down	OK	As and De	–	Trig Offset	A	X
					02 - 240 VA power down	Fatal					
					04 - A/C lost	OK					
					05 - Soft power control failure	Fatal					
					06 - Power unit failure						
Power Unit Redundancy (Pwr Unit Redund)	02h	Chassis-specific	Power Unit 09h	Generic 0Bh	00 - Fully Redundant	OK	As	–	Trig Offset	M	X
					01 - Redundancy lost	Degraded					
					02 - Redundancy degraded	Degraded					
					03 - Non-redundant: sufficient resources. Transition from full redundant state.	Degraded					
					04 – Non-redundant: sufficient resources. Transition from insufficient state.	Degraded					
					05 - Non-redundant: insufficient resources	Fatal					
					06 – Redundant: degraded from fully redundant state.	Degraded					
					07 – Redundant: Transition from non-redundant state.	Degraded					
IPMI Watchdog (IPMI Watchdog)	03h	All	Watchdog 2 23h	Sensor Specific 6Fh	00 - Timer expired, status only	OK	As	–	Trig Offset	A	X
					01 - Hard reset						
					02 - Power down						

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
					03 - Power cycle 08 - Timer interrupt						
Physical Security (Physical Scrty)	04h	Chassis Intrusion is chassis-specific	Physical Security 05h	Sensor Specific 6Fh	00 - Chassis intrusion 04 - LAN leash lost	Degraded OK	As and De	–	Trig Offset	A	X
FP Interrupt (FP NMI Diag Int)	05h	Chassis - specific	Critical Interrupt 13h	Sensor Specific 6Fh	00 - Front panel NMI/diagnostic interrupt	OK	As	–	Trig Offset	A	–
SMI Timeout (SMI Timeout)	06h	All	SMI Timeout F3h	Digital Discrete 03h	01 – State asserted	Fatal	As and De	–	Trig Offset	A	–
System Event Log (System Event Log)	07h	All	Event Logging Disabled 10h	Sensor Specific 6Fh	02 - Log area reset/cleared	OK	As	–	Trig Offset	A	X
System Event (System Event)	08h	All	System Event 12h	Sensor Specific 6Fh	04 – PEF action	OK	As	-	Trig Offset	A	X
Button Sensor *(Button)*	09h	All	Button/Switch 14h	Sensor Specific 6Fh	00 – Power Button 02 – Reset Button	OK	AS	_	Trig Offset	A	X
BMC Watchdog	0Ah	All	Mgmt System Health 28h	Digital Discrete 03h	01 – State Asserted	Degraded	As	–	Trig Offset	A	-
Voltage Regulator Watchdog (VR Watchdog)	0Bh	All	Voltage 02h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Fan Redundancy (Fan Redundancy)	0Ch	Chassis-specific	Fan 04h	Generic 0Bh	00 - Fully redundant 01 - Redundancy lost 02 - Redundancy degraded	OK Degraded Degraded	As and De	–	Trig Offset	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
					03 - Non-redundant: Sufficient resources. Transition from redundant	Degraded					
					04 - Non-redundant: Sufficient resources. Transition from insufficient.	Degraded					
					05 - Non-redundant: insufficient resources.	Non-Fatal					
					06 – Non-Redundant: degraded from fully redundant.	Degraded					
					07 - Redundant degraded from non-redundant	Degraded					
SSB Thermal Trip *(SSB Therm Trip)*	0Dh	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
IO Module Presence (IO Mod Presence)	0Eh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	-
SAS Module Presence (SAS Mod Presence)	0Fh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	X
BMC Firmware Health (BMC FW Health)	10h	All	Mgmt Health 28h	Sensor Specific 6Fh	04 – Sensor Failure	Degraded	As	-	Trig Offset	A	X
System Airflow (System Airflow)	11h	All	Other Units 0Bh	Threshold 01h	–	–	–	Analog	–	–	–
FW Update Status	12h	All	Version Change 2Bh	OEM defined 70h	00h – Update started 01h – Update completed successfully. 02h – Update failure	OK	As	_	Trig Offset	A	_

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
IO Module2 Presence (IO Mod2 Presence)	13h	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	-
Baseboard Temperature 5 *(*Platform Specific*)*	14h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 6 *(*Platform Specific*)*	15h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Module2 Temperature (I/O Mod2 Temp)	16h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 3 Temperature (PCI Riser 3 Temp)	17h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 4 Temperature (PCI Riser 4 Temp)	18h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard +1.05V Processor3 Vccp (BB +1.05Vccp P3)	19h	Platform-specific	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.05V Processor4 Vccp (BB +1.05Vccp P4)	1Ah	Platform-specific	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard Temperature 1 *(*Platform Specific*)*	20h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Front Panel Temperature (Front Panel Temp)	21h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SSB Temperature *(SSB Temp)*	22h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Baseboard Temperature 2 *(Platform Specific)*	23h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 3 *(Platform Specific)*	24h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 4 *(Platform Specific)*	25h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Module Temperature (I/O Mod Temp)	26h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 1 Temperature (PCI Riser 1 Temp)	27h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Riser Temperature (IO Riser Temp)	28h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 1 Temperature (HSBP 1 Temp)	29h	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 2 Temperature (HSBP 2 Temp)	2Ah	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 3 Temperature (HSBP 3 Temp)	2Bh	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 2 Temperature (PCI Riser 2 Temp)	2Ch	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SAS Module Temperature (SAS Mod Temp)	2Dh	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Exit Air Temperature (Exit Air Temp)	2Eh	Chassis and Platform Specific	Temperature 01h	Threshold 01h	This sensor does not generate any events.	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Network Interface Controller Temperature (LAN NIC Temp)	2Fh	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Fan Tachometer Sensors (Chassis specific sensor names)	30h–3Fh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal[Note3]	As and De	Analog	R, T	M	-
Fan Present Sensors (Fan x Present)	40h–4Fh	Chassis and Platform Specific	Fan 04h	Generic 08h	01 - Device inserted	OK	As and De	-	Triggered Offset	Auto	-
Power Supply 1 Status (PS1 Status)	50h	Chassis-specific	Power Supply 08h	Sensor Specific 6Fh	00 - Presence	OK	As and De	–	Trig Offset	A	X
					01 - Failure	Degraded					
					02 – Predictive Failure	Degraded					
					03 - A/C lost	Degraded					
					06 – Configuration error	OK					
Power Supply 2 Status (PS2 Status)	51h	Chassis-specific	Power Supply 08h	Sensor Specific 6Fh	00 - Presence	OK	As and De	–	Trig Offset	A	X
					01 - Failure	Degraded					
					02 – Predictive Failure	Degraded					
					03 - A/C lost	Degraded					
					06 – Configuration error	OK					
Power Supply 1 AC Power Input (PS1 Power In)	54h	Chassis-specific	Other Units 0Bh	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 AC Power Input (PS2 Power In)	55h	Chassis-specific	Other Units 0Bh	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Supply 1 +12V % of Maximum Current Output (PS1 Curr Out %)	58h	Chassis-specific	Current 03h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 +12V % of Maximum Current Output (PS2 Curr Out %)	59h	Chassis-specific	Current 03h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 1 Temperature (PS1 Temperature)	5Ch	Chassis-specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 Temperature (PS2 Temperature)	5Dh	Chassis-specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hard Disk Drive 15 - 23 Status (HDD 15 - 23 Status)	60h – 68h	Chassis-specific	Drive Slot 0Dh	Sensor Specific 6Fh	00 - Drive Presence	OK	As and De	–	Trig Offset	A	X X
					01- Drive Fault	Degraded					
					07 - Rebuild/Remap in progress	Degraded					
Processor 1 Status (P1 Status)	70h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 2 Status (P2 Status)	71h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 3 Status (P3 Status)	72h	Platform-specific	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 4 Status (P4 Status)	73h	Platform-specific	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 1 Thermal Margin (P1 Therm Margin)	74h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 2 Thermal Margin (P2 Therm Margin)	75h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 3 Thermal Margin (P3 Therm Margin)	76h	Platform-specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Processor 4 Thermal Margin (P4 Therm Margin)	77h	Platform-specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 1 Thermal Control % (P1 Therm Ctrl %)	78h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor 2 Thermal Control % (P2 Therm Ctrl %)	79h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor 3 Thermal Control % (P3 Therm Ctrl %)	7Ah	Platform-specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor 4 Thermal Control % (P4 Therm Ctrl %)	7Bh	Platform-specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor ERR2 Timeout (CPU ERR2)	7Ch	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	A	–
Catastrophic Error (CATERR)	80h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	–
MTM Level Change (MTM Lvl Change)	81h	All	Mgmt Health 28h	Digital Discrete 03h	01 – State Asserted	-	As and De	–	Trig Offset	A	-
Processor Population Fault (CPU Missing)	82h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	–
Processor 1 DTS Thermal Margin (P1 DTS Therm Mgn)	83h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 2 DTS Thermal Margin (P2 DTS Therm Mgn)	84h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 3 DTS Thermal Margin (P3 DTS Therm Mgn)	85h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Processor 4 DTS Thermal Margin (P4 DTS Therm Mgn)	86h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Auto Config Status (AutoCfg Status)	87h	All	Mgmt Health 28h	Digital Discrete 03h	01 – State Asserted	-	As and De	–	Trig Offset	A	-
Processor 1 VRD Temperature (VRD Hot)	90h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Power Supply 1 Fan Tachometer 1 *(PS1 Fan Tach 1)*	A0h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	A	-
Power Supply 1 Fan Tachometer 2 *(PS1 Fan Tach 2)*	A1h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	A	-
MIC 1 Status (GPGPU1 Status)	A2h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
MIC 2 Status (GPGPU2 Status)	A3h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
Power Supply 2 Fan Tachometer 1 *(PS2 Fan Tach 1)*	A4h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 2 *(PS2 Fan Tach 2)*	A5h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
MIC 3 Status (GPGPU3 Status)	A6h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
MIC 4 Status (GPGPU4 Status)	A7h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
Processor 1 DIMM Aggregate Thermal Margin 1 (P1 DIMM Thrm Mrgn1)	B0h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 1 DIMM Aggregate Thermal Margin 2 (P1 DIMM Thrm Mrgn2)	B1h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 2 DIMM Aggregate Thermal Margin 1 (P2 DIMM Thrm Mrgn1)	B2h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 2 DIMM Aggregate Thermal Margin 2 (P2 DIMM Thrm Mrgn2)	B3h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 3 DIMM Aggregate Thermal Margin 1 (P3 DIMM Thrm Mrgn1)	B4h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 3 DIMM Aggregate Thermal Margin 2 (P3 DIMM Thrm Mrgn2)	B5h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 4 DIMM Aggregate Thermal Margin 1 (P4 DIMM Thrm Mrgn1)	B6h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 4 DIMM Aggregate Thermal Margin 2 (P4 DIMM Thrm Mrgn2)	B7h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Node Auto-Shutdown Sensor (Auto Shutdown)	B8h	Multi-Node Specific	Power Unit 09h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	A	-
Fan Tachometer Sensors (Chassis specific sensor names)	BAh–BFh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal[2]	As and De	Analog	R, T	M	-

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Processor 1 DIMM Thermal Trip (P1 Mem Thrm Trip)	C0h	All	Memory 0Ch	Sensor Specific 6Fh	0A- Critical overtemperature	Fatal	As and De	–	Trig Offset	M	-
Processor 2 DIMM Thermal Trip (P2 Mem Thrm Trip)	C1h	All	Memory 0Ch	Sensor Specific 6Fh	0A- Critical overtemperature	Fatal	As and De	–	Trig Offset	M	-
Processor 3 DIMM Thermal Trip (P3 Mem Thrm Trip)	C2h	Platform Specific	Memory 0Ch	Sensor Specific 6Fh	0A- Critical overtemperature	Fatal	As and De	–	Trig Offset	M	X
Processor 4 DIMM Thermal Trip (P4 Mem Thrm Trip)	C3h	Platform Specific	Memory 0Ch	Sensor Specific 6Fh	0A- Critical overtemperature	Fatal	As and De	–	Trig Offset	M	X
MIC 1 Temp (GPGPU1 Core Temp)	C4h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 2 Temp (GPGPU2 Core Temp)	C5h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 3 Temp (GPGPU3 Core Temp)	C6h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 4 Temp (GPGPU4 Core Temp)	C7h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
Global Aggregate Temperature Margin 1 (Agg Therm Mrgn 1)	C8h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 2 (Agg Therm Mrgn 2)	C9h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 3 (Agg Therm Mrgn 3)	CAh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 4 (Agg Therm Mrgn 4)	CBh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert /De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Global Aggregate Temperature Margin 5 (Agg Therm Mrgn 5)	CCh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 6 (Agg Therm Mrgn 6)	CDh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 7 (Agg Therm Mrgn 7)	CEh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 8 (Agg Therm Mrgn 8)	CFh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Baseboard +12V (BB +12.0V)	D0h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Voltage Fault (Voltage Fault)	D1h	All	Voltage 02h	Discrete 03h	01 – Asserted	-	-	-	-	A	-
Baseboard CMOS Battery (BB +3.3V Vbat)	DEh	All	Voltage 02h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Hard Disk Drive 0 -14 Status (HDD 0 - 14 Status)	F0h - FEh	Chassis-specific	Drive Slot 0Dh	Sensor Specific 6Fh	00 - Drive Presence	OK	As and De	–	Trig Offset	A	X
					01- Drive Fault	Degraded					
					07 - Rebuild/Remap in progress	Degraded					

Appendix C: Management Engine Generated SEL Event Messages

This appendix lists the OEM System Event Log message format of events generated by the Management Engine (ME). This includes the definition of event data bytes 10-16 of the Management Engine generated SEL records. For System Event Log format information, see the *Intelligent Platform Management Interface Specification, Version 2.0*.

Table 58. Server Platform Services Firmware Health Event

Server Platform Services Firmware Health Event	Request Byte 1 - EvMRev =04h (IPMI2.0 format) Byte 2 – Sensor Type =DCh (OEM) Byte 3 – Sensor Number =23 – Server Platform Services Firmware Health Byte 4 – Event Dir | Event Type [7] – Event Dir =0 Assertion Event [6-0] – Event Type =75h (OEM) Byte 5 – Event Data 1 [7,6]=10b – OEM code in byte 2 [5,4]=10b – OEM code in byte 3 [3..0] – Health Event Type =00h –Firmware Status Byte 6 – Event Data 2 =0 - Forced GPIO recovery. Recovery Image loaded due to MGPIO<n> (default recovery pin is MGPIO1) pin asserted. *Repair action: Deassert MGPIO1 and reset the ME* =1 - Image execution failed. Recovery Image loaded because operational image is corrupted. This may be either caused by Flash device corruption or failed upgrade procedure. *Repair action: Either the Flash device must be replaced (if error is persistent) or the upgrade procedure must be started again.* =2 - Flash erase error. Error during Flash erases procedure probably due to Flash part corruption. *Repair action: The Flash device must be replaced.* =3 – Flash corrupted. Error while checking Flash consistency probably due to Flash part corruption. *Repair action: The Flash device must be replaced (if error is persistent).* =4 – Internal error. Error during firmware execution. *Repair action: FW Watchdog Timeout* *Operational image shall be upgraded to other version or hardware board repair is needed (if error is persistent).* =5..255 – Reserved Byte 7 – Event Data 3 =<Extended error code. Should be used when reporting an error to the support>

Table 59. Node Manager Health Event

Node Manager Health Event	Request
	Byte 1 - EvMRev =04h (IPMI2.0 format) Byte 2 – Sensor Type =DCh (OEM) Byte 3 – Sensor Number (Node Manager Health sensor) Byte 4 – Event Dir | Event Type [0:6] – Event Type = 73h (OEM) [7] – Event Dir =0 Assertion Event Byte 5 – Event Data 1 [0:3] – Health Event Type =02h – Sensor Node Manager [4:5]=10b – OEM code in byte 3 [6:7]=10b – OEM code in byte 2 Byte 6 – Event Data 2 [0:3] – Domain Id (Currently, supports only one domain, Domain 0) [4:7] – Error type =0-9 - Reserved =10 – Policy Misconfiguration =11 – Power Sensor Reading Failure =12 – Inlet Temperature Reading Failure =13 – Host Communication error =14 – Real-time clock synchronization failure =15 – Reserved Byte 7 – Event Data 3 if error indication = 10 <PolicyId> if error indication = 11 <PowerSensorAddress> if error indication = 12 <InletSensorAddress> Otherwise set to 0.

Appendix D: POST Code Diagnostic LED Decoder

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

During the system boot process, Memory Reference Code (MRC) and System BIOS execute a number of memory initialization and platform configuration processes, each of which is assigned a hex POST code number.

As each routine is started, the given POST code number is displayed to the POST Code Diagnostic LEDs on the back edge of the server board.

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

Each POST code is represented by eight LEDs; four Green and four Amber. The POST codes are divided into two nibbles, an upper nibble and a lower nibble. The upper nibble bits are represented by Amber Diagnostic LEDs #4, #5, #6, #7. The lower nibble bits are represented by Green Diagnostics LEDs #0, #1, #2 and #3. 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.

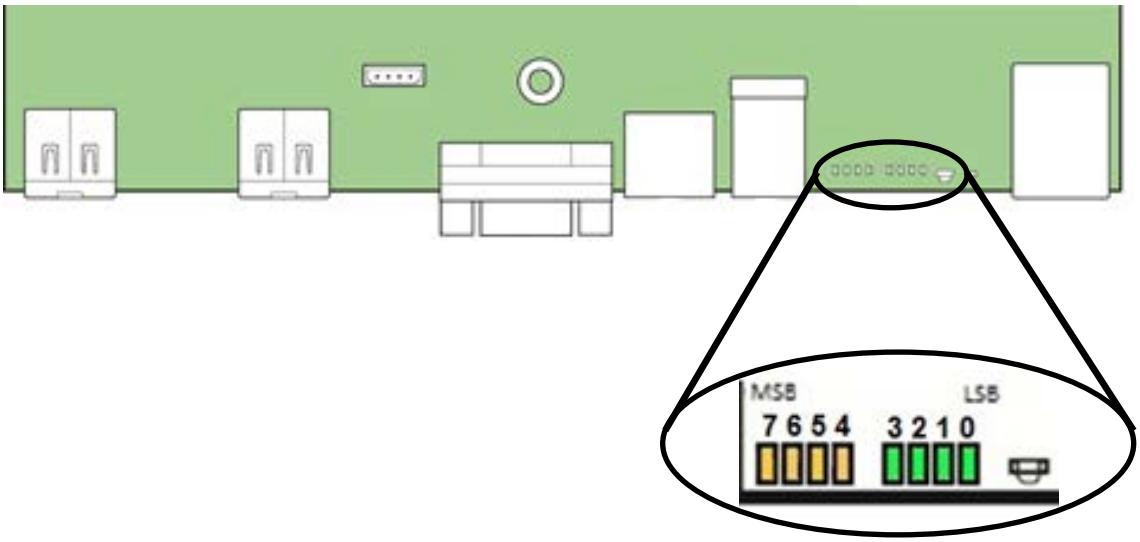


Figure 36. POST Diagnostic LED Location

In the following example, the BIOS sends a value of ACh to the diagnostic LED decoder. The LEDs are decoded as follows:

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

Table 60. POST Progress Code LED Example

LEDs	Upper Nibble AMBER LEDs				Lower Nibble GREEN LEDs			
	MSB							LSB
	LED #7	LED #6	LED #5	LED #4	LED #3	LED #2	LED #1	LED #0
	8h	4h	2h	1h	8h	4h	2h	1h
Status	ON	OFF	ON	OFF	ON	ON	OFF	OFF
Results	1	0	1	0	1	1	0	0
	Ah				Ch			

Upper nibble bits = 1010b = Ah; Lower nibble bits = 1100b = Ch; the two are concatenated as ACh

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 61. MRC Progress Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off								Description
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
MRC Progress Codes									
B0h	1	0	1	1	0	0	0	0	Detect DIMM population
B1h	1	0	1	1	0	0	0	1	Set DDR3 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 DDR3 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 3 long beeps as identified in Table 59.

Table 62. MRC Fatal Error Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off								Description
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
MRC Fatal Error Codes									
E8h	1	1	1	0	1	0	0	0	No usable memory error 01h = No memory was detected from SPD read, or invalid config that causes no operable memory. 02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error. 03h = No memory installed. All channels are disabled.
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel Trusted Execution Technology and is inaccessible
EAh	1	1	1	0	1	0	1	0	DDR3 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. 03h = Hardware Memtest failure in Lockstep Channel mode requiring a channel to be disabled. This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.
EDh	1	1	1	0	1	1	0	1	DIMM configuration population error 01h = Different DIMM types (UDIMM, 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 in the 3rd DIMM slot. 05h = Unsupported DIMM Voltage.

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
EFh	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error

BIOS POST Progress Codes

The following table provides a list of all POST progress codes.

Table 63. POST Progress Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off Upper Nibble MSB 8h				Lower Nibble 8h			LSB 1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	Description
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	PEI 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.
07h	0	0	0	0	0	1	1	1	Early SB initialization during Sec Phase.
08h	0	0	0	0	1	0	0	0	Early NB initialization during Sec Phase.
09h	0	0	0	0	1	0	0	1	End Of SEC Phase.
0Eh	0	0	0	0	1	1	1	0	Microcode Not Found.
0Fh	0	0	0	0	1	1	1	1	Microcode Not Loaded.
PEI Phase									
10h	0	0	0	1	0	0	0	0	PEI Core
11h	0	0	0	1	0	0	0	1	CPU PEIM
15h	0	0	0	1	0	1	0	1	NB PEIM
19h	0	0	0	1	1	0	0	1	SB PEIM
MRC Process Codes – MRC Progress Code Sequence is executed - See Table 56. MRC Progress Codes									
PEI Phase continued…									
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)
34h	0	0	1	1	0	1	0	0	CPU PEIM (BSP Select)
35h	0	0	1	1	0	1	0	1	CPU PEIM (AP Init)
36h	0	0	1	1	0	1	1	0	CPU PEIM (CPU SMM 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	SB RUN Init
63h	0	1	1	0	0	0	1	1	DXE CPU 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

Checkpoint	Diagnostic LED Decoder								
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	Description
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
90h	1	0	0	1	0	0	0	0	DXE BDS Started
91h	1	0	0	1	0	0	0	1	DXE BDS connect drivers
92h	1	0	0	1	0	0	1	0	DXE PCI Bus begin
93h	1	0	0	1	0	0	1	1	DXE PCI Bus HPC Init
94h	1	0	0	1	0	1	0	0	DXE PCI Bus enumeration
95h	1	0	0	1	0	1	0	1	DXE PCI Bus resource requested
96h	1	0	0	1	0	1	1	0	DXE PCI Bus assign resource
97h	1	0	0	1	0	1	1	1	DXE CON_OUT connect
98h	1	0	0	1	1	0	0	0	DXE CON_IN connect
99h	1	0	0	1	1	0	0	1	DXE SIO Init
9Ah	1	0	0	1	1	0	1	0	DXE USB start
9Bh	1	0	0	1	1	0	1	1	DXE USB reset
9Ch	1	0	0	1	1	1	0	0	DXE USB detect
9Dh	1	0	0	1	1	1	0	1	DXE USB enable
A1h	1	0	1	0	0	0	0	1	DXE IDE begin
A2h	1	0	1	0	0	0	1	0	DXE IDE reset
A3h	1	0	1	0	0	0	1	1	DXE IDE detect
A4h	1	0	1	0	0	1	0	0	DXE IDE enable
A5h	1	0	1	0	0	1	0	1	DXE SCSI begin
A6h	1	0	1	0	0	1	1	0	DXE SCSI reset
A7h	1	0	1	0	0	1	1	1	DXE SCSI detect
A8h	1	0	1	0	1	0	0	0	DXE SCSI enable
A9h	1	0	1	0	1	0	0	1	DXE verifying SETUP password
ABh	1	0	1	0	1	0	1	1	DXE SETUP start
ACh	1	0	1	0	1	1	0	0	DXE SETUP input wait
ADh	1	0	1	0	1	1	0	1	DXE Ready to Boot
AEh	1	0	1	0	1	1	1	0	DXE Legacy Boot
AFh	1	0	1	0	1	1	1	1	DXE Exit Boot Services
B0h	1	0	1	1	0	0	0	0	RT Set Virtual Address Map Begin
B1h	1	0	1	1	0	0	0	1	RT Set Virtual Address Map End
B2h	1	0	1	1	0	0	1	0	DXE Legacy Option ROM init
B3h	1	0	1	1	0	0	1	1	DXE Reset system
B4h	1	0	1	1	0	1	0	0	DXE USB Hot plug
B5h	1	0	1	1	0	1	0	1	DXE PCI BUS Hot plug
B6h	1	0	1	1	0	1	1	0	DXE NVRAM cleanup
B7h	1	0	1	1	0	1	1	1	DXE Configuration Reset
00h	0	0	0	0	0	0	0	0	INT19
S3 Resume									

Checkpoint	Diagnostic LED Decoder								
	1 = LED On, 0 = LED Off								
	Upper Nibble				Lower Nibble				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
LED #	#7	#6	#5	#4	#3	#2	#1	#0	Description
E0h	1	1	1	0	0	0	0	0	S3 Resume PEIM (S3 started)
E1h	1	1	1	0	0	0	0	1	S3 Resume PEIM (S3 boot script)
E2h	1	1	1	0	0	0	1	0	S3 Resume PEIM (S3 Video Repost)
E3h	1	1	1	0	0	0	1	1	S3 Resume PEIM (S3 OS wake)
BIOS Recovery									
F0h	1	1	1	1	0	0	0	0	PEIM which detected forced Recovery condition
F1h	1	1	1	1	0	0	0	1	PEIM which detected User Recovery condition
F2h	1	1	1	1	0	0	1	0	Recovery PEIM (Recovery started)
F3h	1	1	1	1	0	0	1	1	Recovery PEIM (Capsule found)
F4h	1	1	1	1	0	1	0	0	Recovery PEIM (Capsule loaded)

Appendix E: 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 handed 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.

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

Table 64. POST Error Codes and Messages

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

Error Code	Error Message	Response
0196	Processor model mismatch detected	Fatal
0197	Processor frequencies unable to synchronize	Fatal
5220	BIOS Settings reset to default settings	Major
5221	Passwords cleared by jumper	Major
5224	Password clear jumper is set	Major
8130	Processor 01 disabled	Major
8131	Processor 02 disabled	Major
8160	Processor 01 unable to apply microcode update	Major
8161	Processor 02 unable to apply microcode update	Major
8170	Processor 01 failed Self Test (BIST)	Major
8171	Processor 02 failed Self Test (BIST)	Major
8180	Processor 01 microcode update not found	Minor
8181	Processor 02 microcode update not found	Minor
8190	Watchdog timer failed on last boot	Major
8198	OS boot watchdog timer failure	Major
8300	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	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8501	DIMM Population Error	Major
8520	DIMM_A1 failed test/initialization	Major
8521	DIMM_A2 failed test/initialization	Major
8522	DIMM_A3 failed test/initialization	Major
8523	DIMM_B1 failed test/initialization	Major
8524	DIMM_B2 failed test/initialization	Major
8525	DIMM_B3 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8527	DIMM_C2 failed test/initialization	Major
8528	DIMM_C3 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major
852A	DIMM_D2 failed test/initialization	Major
852B	DIMM_D3 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852D	DIMM_E2 failed test/initialization	Major
852E	DIMM_E3 failed test/initialization	Major
852F	DIMM_F1 failed test/initialization	Major
8530	DIMM_F2 failed test/initialization	Major
8531	DIMM_F3 failed test/initialization	Major
8532	DIMM_G1 failed test/initialization	Major
8533	DIMM_G2 failed test/initialization	Major

Error Code	Error Message	Response
8534	DIMM_G3 failed test/initialization	Major
8535	DIMM_H1 failed test/initialization	Major
8536	DIMM_H2 failed test/initialization	Major
8537	DIMM_H3 failed test/initialization	Major
8538	DIMM_J1 failed test/initialization	Major
8539	DIMM_J2 failed test/initialization	Major
853A	DIMM_J3 failed test/initialization	Major
853B	DIMM_K1 failed test/initialization	Major
853C	DIMM_K2 failed test/initialization	Major
853D	DIMM_K3 failed test/initialization	Major
853E	DIMM_L1 failed test/initialization	Major
853F (Go to 85C0)	DIMM_L2 failed test/initialization	Major
8540	DIMM_A1 disabled	Major
8541	DIMM_A2 disabled	Major
8542	DIMM_A3 disabled	Major
8543	DIMM_B1 disabled	Major
8544	DIMM_B2 disabled	Major
8545	DIMM_B3 disabled	Major
8546	DIMM_C1 disabled	Major
8547	DIMM_C2 disabled	Major
8548	DIMM_C3 disabled	Major
8549	DIMM_D1 disabled	Major
854A	DIMM_D2 disabled	Major
854B	DIMM_D3 disabled	Major
854C	DIMM_E1 disabled	Major
854D	DIMM_E2 disabled	Major
854E	DIMM_E3 disabled	Major
854F	DIMM_F1 disabled	Major
8550	DIMM_F2 disabled	Major
8551	DIMM_F3 disabled	Major
8552	DIMM_G1 disabled	Major
8553	DIMM_G2 disabled	Major
8554	DIMM_G3 disabled	Major
8555	DIMM_H1 disabled	Major
8556	DIMM_H2 disabled	Major
8557	DIMM_H3 disabled	Major
8558	DIMM_J1 disabled	Major
8559	DIMM_J2 disabled	Major
855A	DIMM_J3 disabled	Major
855B	DIMM_K1 disabled	Major
855C	DIMM_K2 disabled	Major
855D	DIMM_K3 disabled	Major
855E	DIMM_L1 disabled	Major
855F (Go to 85D0)	DIMM_L2 disabled	Major

Error Code	Error Message	Response
8560	DIMM_A1 encountered a Serial Presence Detection (SPD) failure	Major
8561	DIMM_A2 encountered a Serial Presence Detection (SPD) failure	Major
8562	DIMM_A3 encountered a Serial Presence Detection (SPD) failure	Major
8563	DIMM_B1 encountered a Serial Presence Detection (SPD) failure	Major
8564	DIMM_B2 encountered a Serial Presence Detection (SPD) failure	Major
8565	DIMM_B3 encountered a Serial Presence Detection (SPD) failure	Major
8566	DIMM_C1 encountered a Serial Presence Detection (SPD) failure	Major
8567	DIMM_C2 encountered a Serial Presence Detection (SPD) failure	Major
8568	DIMM_C3 encountered a Serial Presence Detection (SPD) failure	Major
8569	DIMM_D1 encountered a Serial Presence Detection (SPD) failure	Major
856A	DIMM_D2 encountered a Serial Presence Detection (SPD) failure	Major
856B	DIMM_D3 encountered a Serial Presence Detection (SPD) failure	Major
856C	DIMM_E1 encountered a Serial Presence Detection (SPD) failure	Major
856D	DIMM_E2 encountered a Serial Presence Detection (SPD) failure	Major
856E	DIMM_E3 encountered a Serial Presence Detection (SPD) failure	Major
856F	DIMM_F1 encountered a Serial Presence Detection (SPD) failure	Major
8570	DIMM_F2 encountered a Serial Presence Detection (SPD) failure	Major
8571	DIMM_F3 encountered a Serial Presence Detection (SPD) failure	Major
8572	DIMM_G1 encountered a Serial Presence Detection (SPD) failure	Major
8573	DIMM_G2 encountered a Serial Presence Detection (SPD) failure	Major
8574	DIMM_G3 encountered a Serial Presence Detection (SPD) failure	Major
8575	DIMM_H1 encountered a Serial Presence Detection (SPD) failure	Major
8576	DIMM_H2 encountered a Serial Presence Detection (SPD) failure	Major
8577	DIMM_H3 encountered a Serial Presence Detection (SPD) failure	Major
8578	DIMM_J1 encountered a Serial Presence Detection (SPD) failure	Major
8579	DIMM_J2 encountered a Serial Presence Detection (SPD) failure	Major
857A	DIMM_J3 encountered a Serial Presence Detection (SPD) failure	Major
857B	DIMM_K1 encountered a Serial Presence Detection (SPD) failure	Major
857C	DIMM_K2 encountered a Serial Presence Detection (SPD) failure	Major
857D	DIMM_K3 encountered a Serial Presence Detection (SPD) failure	Major
857E	DIMM_L1 encountered a Serial Presence Detection (SPD) failure	Major
857F (Go to 85E0)	DIMM_L2 encountered a Serial Presence Detection (SPD) failure	Major
85C0	DIMM_L3 failed test/initialization	Major
85C1	DIMM_M1 failed test/initialization	Major
85C2	DIMM_M2 failed test/initialization	Major
85C3	DIMM_M3 failed test/initialization	Major
85C4	DIMM_N1 failed test/initialization	Major
85C5	DIMM_N2 failed test/initialization	Major
85C6	DIMM_N3 failed test/initialization	Major
85C7	DIMM_P1 failed test/initialization	Major
85C8	DIMM_P2 failed test/initialization	Major
85C9	DIMM_P3 failed test/initialization	Major
85CA	DIMM_R1 failed test/initialization	Major
85CB	DIMM_R2 failed test/initialization	Major
85CC	DIMM_R3 failed test/initialization	Major

Error Code	Error Message	Response
85CD	DIMM_T1 failed test/initialization	Major
85CE	DIMM_T2 failed test/initialization	Major
85CF	DIMM_T3 failed test/initialization	Major
85D0	DIMM_L3 disabled	Major
85D1	DIMM_M1 disabled	Major
85D2	DIMM_M2 disabled	Major
85D3	DIMM_M3 disabled	Major
85D4	DIMM_N1 disabled	Major
85D5	DIMM_N2 disabled	Major
85D6	DIMM_N3 disabled	Major
85D7	DIMM_P1 disabled	Major
85D8	DIMM_P2 disabled	Major
85D9	DIMM_P3 disabled	Major
85DA	DIMM_R1 disabled	Major
85DB	DIMM_R2 disabled	Major
85DC	DIMM_R3 disabled	Major
85DD	DIMM_T1 disabled	Major
85DE	DIMM_T2 disabled	Major
85DF	DIMM_T3 disabled	Major
85E0	DIMM_L3 encountered a Serial Presence Detection (SPD) failure	Major
85E1	DIMM_M1 encountered a Serial Presence Detection (SPD) failure	Major
85E2	DIMM_M2 encountered a Serial Presence Detection (SPD) failure	Major
85E3	DIMM_M3 encountered a Serial Presence Detection (SPD) failure	Major
85E4	DIMM_N1 encountered a Serial Presence Detection (SPD) failure	Major
85E5	DIMM_N2 encountered a Serial Presence Detection (SPD) failure	Major
85E6	DIMM_N3 encountered a Serial Presence Detection (SPD) failure	Major
85E7	DIMM_P1 encountered a Serial Presence Detection (SPD) failure	Major
85E8	DIMM_P2 encountered a Serial Presence Detection (SPD) failure	Major
85E9	DIMM_P3 encountered a Serial Presence Detection (SPD) failure	Major
85EA	DIMM_R1 encountered a Serial Presence Detection (SPD) failure	Major
85EB	DIMM_R2 encountered a Serial Presence Detection (SPD) failure	Major
85EC	DIMM_R3 encountered a Serial Presence Detection (SPD) failure	Major
85ED	DIMM_T1 encountered a Serial Presence Detection (SPD) failure	Major
85EE	DIMM_T2 encountered a Serial Presence Detection (SPD) failure	Major
85EF	DIMM_T3 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

Error Code	Error Message	Response
A003	TPM device failed self test.	Minor
A100	BIOS ACM Error	Major
A421	PCI component encountered a SERR error	Fatal
A5A0	PCI express* component encountered a PERR error	Minor
A5A1	PCI express* component encountered an SERR error	Fatal
A6A0	DXE Boot Services driver: Not enough memory available to shadow a Legacy Option ROM.	Minor

POST Error Beep Codes

The following table 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 65. 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 the following table. Each digit in the code is represented by a sequence of beeps whose count is equal to the digit.

Table 66. Integrated BMC Beep Codes

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

Appendix F: Supported Intel® Server Systems

Two Intel® Server System product families integrate the Intel® Server board S2600WT, they are the 1U rack mount Intel® Server System R1000WT product family and the 2U rack mount Intel® R2000WT product family.

Intel® Server System R1000WT Product Family

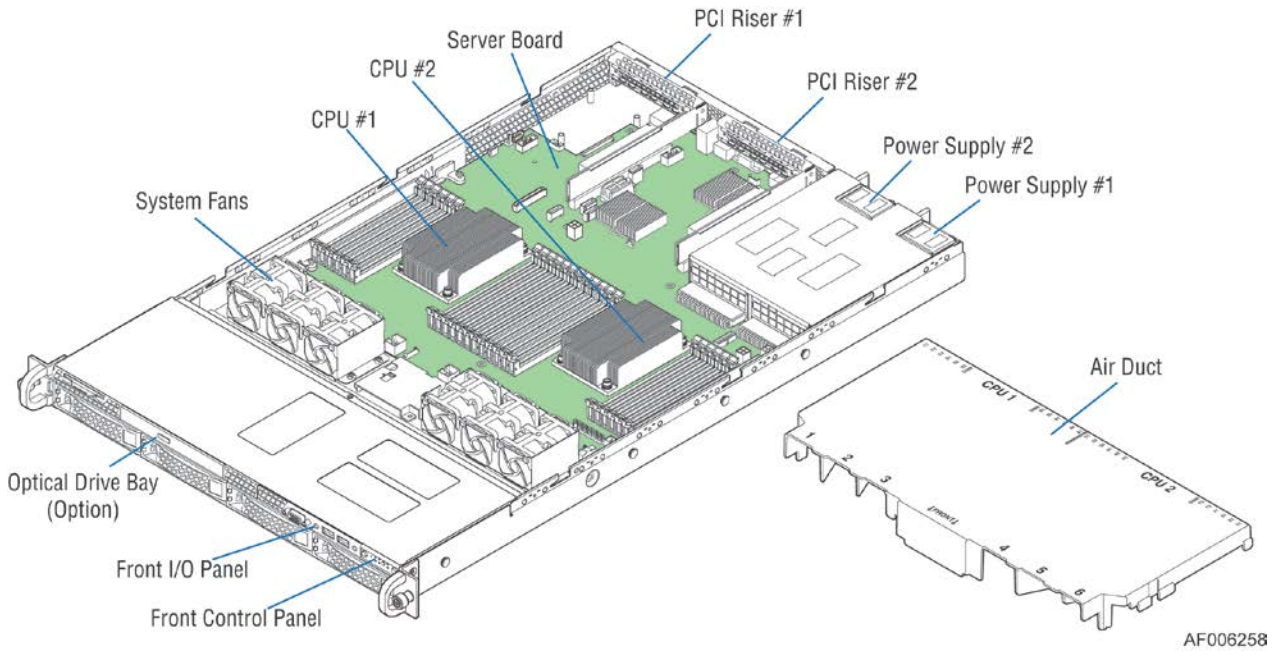


Figure 37. Intel® Server System R1000WT

Table 67. Intel® Server System R1000WT Product Family Feature Set

Feature	Description
Chassis Type	1U Rack Mount Chassis
Server Board Options	• Intel® Server Board S2600WT w/Dual 1GbE ports – S2600WT2 • Intel® Server Board S2600WT w/Dual 10GbE ports – S2600WTT
Processor Support	• Two LGA2011-3 (Socket R3) processor sockets • Support for one or two Intel® Xeon® processors E5-2600 v3 product family • Maximum supported Thermal Design Power (TDP) of up to 145 W.
Memory	• 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor • Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) • Memory data transfer rates: o DDR4 RDIMM: 1600 MT/s (3DPC), 1866 MT/s (2DPC) and 2133 MT/s (1DPC) o DDR4 LRDIMM: 1600 MT/s (3DPC), 2133 MT/s (2DPC & 1DPC) • DDR4 standard I/O voltage of 1.2V
Chipset	Intel® C612 chipset

Feature	Description
External I/O connections	• DB-15 Video connectors o Front and Back on non-storage systems o Back only on storage systems (12 x 3.5" and 24 x 2.5" drive support) • RJ-45 Serial Port A connector • Dual RJ-45 Network Interface connectors supporting either : o 10 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WTT) or o 1 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WT2) • Dedicated RJ-45 server management port • Three USB 2.0 / 3.0 connectors on back panel • Two USB 2.0 / 3.0 ports on front panel (non-storage models only)
Internal I/O connectors / headers	• One Type-A USB 2.0 connector • One 2x5 pin connector providing front panel support for two USB 2.0 ports • One 2x10 pin connector providing front panel support for two USB 2.0 / 3.0 ports • One 2x15 pin SSI-EEB compliant front panel header • One 2x7 pin Front Panel Video connector • One 1x7 pin header for optional Intel® Local Control Panel (LCP) support • One DH-10 Serial Port B connector
I/O Module Accessory Options	The server board includes a proprietary on-board connector allowing for the installation of a variety of available I/O modules. An installed I/O module can be supported in addition to standard on-board features and add-in PCIe* cards. • AXX4P1GBPWLIOM – Quad port RJ45 1 GbE based on Intel® Ethernet Controller I350 • AXX10GBTWLIOM – Dual port RJ-45 10GBase-T I/O Module based on Intel® Ethernet Controller x540 • AXX10GBNIAIOM – Dual port SFP+ 10 GbE module based on Intel® 82599 10 GbE controller • AXX1FDRIBIOM – Single port QSFP FDR 56 GT/S speed InfiniBand* module • AXX2FDRIBIOM – Dual port QSFP FDR 56 GT/S speed infiniband* module • AXX1P40FRTIOM – Single port QSFP+ 40 GbE module • AXX2P40FRTIOM – Dual port QSFP+ 40 GbE module
System Fans	• Six dual rotor managed system fans • One power supply fan for each installed power supply module
Riser Card Support	Concurrent support for two PCIe* riser cards. Each riser card slot has support for the following riser card options: • Single add-in card slot – PCIe* x16, x16 mechanical
Video	• Integrated 2D Video Controller • 16 MB DDR3 Memory
On-board storage controllers and options	• 10 x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) • Two single port SATA connectors capable of supporting up to 6 Gb/sec • Two 4-port mini-SAS HD (SFF-8643) connectors capable of supporting up to 6 Gb/sec SATA • One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices • Optional SAS IOC/ROC support via on-board Intel® Integrated RAID module connector • Embedded Software SATA RAID o Intel® Rapid Storage RAID Technology (RSTe) 4.0 o Intel® Embedded Server RAID Technology 2 (ESRT2) with optional RAID 5 key support
Security	• Intel® Trusted Platform Module (TPM) - AXXTPME5 (Accessory Option)
Server Management	• Integrated Baseboard Management Controller, IPMI 2.0 compliant • Support for Intel® Server Management Software • On-board RJ45 management port • Advanced Server Management via an Intel® Remote Management Module 4 Lite (Accessory Option)

Feature	Description
Power Supply Options	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 Three power supply options: • AC 750W Platinum • DC 750W Gold • AC 1100W Platinum
Storage Options	12Gb/sec Hot Swap Backplane Options: • 8x – 2.5” SATA/SAS • 4x - 2.5” SATA/SAS + 4x - 2.5” PCIe NVM Express* (Not Hot Swappable) – subject to change • 4x – 3.5” SATA/SAS Storage Bay Options: • 4x – 3.5” SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support • 8x – 2.5” SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support (capable) • 4x - 2.5” SATA/SAS + 4x - 2.5” PCIe* SSD
Supported Rack Mount Kit Accessory Options	• AXXPRAIL – Tool-less rack mount rail kit – 800mm max travel length • AXXELVRAIL – Enhanced value rack mount rail kit - 424mm max travel length • AXX1U2UCMA – Cable Management Arm – (*supported with AXXPRAIL only) • AXX2POSTBRCKT – 2-post fixed mount bracket kit

Intel® Server System R2000WT Product Family

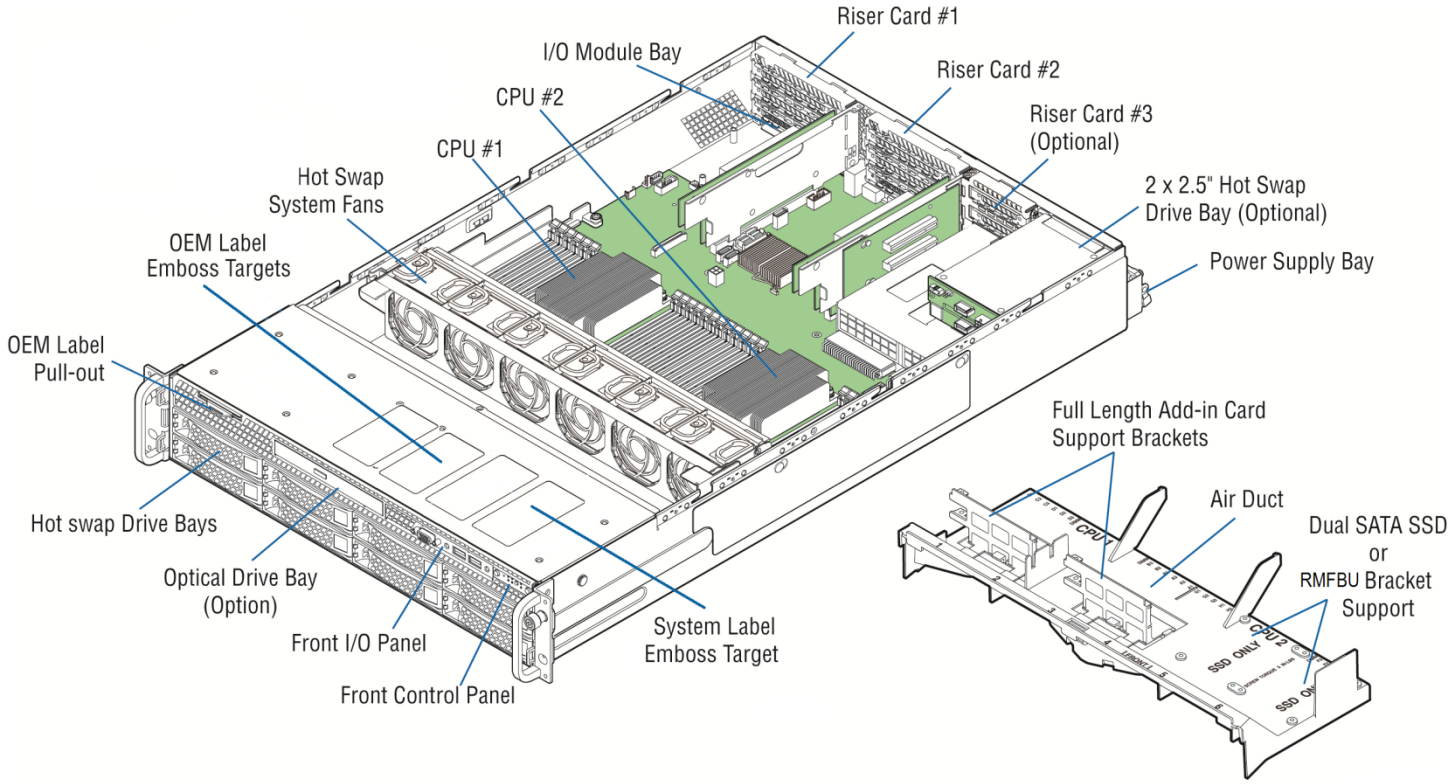


Figure 38. Intel® Server System R2000WT

Table 68. Intel® Server System R2000WT Product Family Feature Set

Feature	Description
Chassis Type	2U Rack Mount Chassis
Server Board Options	• Intel® Server Board S2600WT w/Dual 1GbE ports – S2600WT2 • Intel® Server Board S2600WT w/Dual 10GbE ports – S2600WTT
Processor Support	• Two LGA2011-3 (Socket R3) processor sockets • Support for one or two Intel® Xeon® processors E5-2600 v3 product family • Maximum supported Thermal Design Power (TDP) of up to 145 W.
Memory	• 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor • Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) • Memory data transfer rates: o DDR4 RDIMM: 1600 MT/s (3DPC), 1866 MT/s (2DPC) and 2133 MT/s (1DPC) o DDR4 LRDIMM: 1600 MT/s (3DPC), 2133 MT/s (2DPC & 1DPC) • DDR4 standard I/O voltage of 1.2V
Chipset	Intel® C612 chipset

Feature	Description
External I/O connections	• DB-15 Video connectors o Front and Back on non-storage systems o Back only on storage systems (12 x 3.5" and 24 x 2.5" drive support) • RJ-45 Serial Port A connector • Dual RJ-45 Network Interface connectors supporting either : o 10 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WTT) or o 1 GbE RJ-45 connectors (Intel Server Board Product Code – S2600WT2) • Dedicated RJ-45 server management port • Three USB 2.0 / 3.0 connectors on back panel • Two USB 2.0 / 3.0 ports on front panel (non-storage models only) • One USB 2.0 port on rack handle (storage models only)
Internal I/O connectors / headers	• One Type-A USB 2.0 connector • One 2x5 pin connector providing front panel support for two USB 2.0 ports • One 2x10 pin connector providing front panel support for two USB 2.0 / 3.0 ports • One 2x15 pin SSI-EEB compliant front panel header • One 2x7pin Front Panel Video connector • One 1x7pin header for optional Intel® Local Control Panel (LCP) support • One DH-10 Serial Port B connector
I/O Module Accessory Options	The server board includes a proprietary on-board connector allowing for the installation of a variety of available I/O modules. An installed I/O module can be supported in addition to standard on-board features and add-in PCIe* cards. • AXX4P1GBPWLIOM – Quad port RJ45 1 GbE based on Intel® Ethernet Controller I350 • AXX10GBTWLIOM – Dual port RJ-45 10GBase-T I/O Module based on Intel® Ethernet Controller x540 • AXX10GBNIAIOM – Dual port SFP+ 10 GbE module based on Intel® 82599 10 GbE controller • AXX1FDRIBIOM – Single port QSFP FDR 56 GT/S speed InfiniBand* module • AXX2FDRIBIOM – Dual port QSFP FDR 56 GT/S speed infiniband* module • AXX1P40FRTIOM – Single port QSFP+ 40 GbE module • AXX2P40FRTIOM – Dual port QSFP+ 40 GbE module
System Fans	• Six managed hot swap system fans • One power supply fan for each installed power supply module
Riser Card Support	Support for three riser cards. • Riser #1 – PCIe* Gen3 x24 – up to 3 PCIe* slots • Riser #2 – PCIe* Gen3 x24 – up to 3 PCIe* slots • Riser #3 – PCIe* Gen3 x8 + DMI x4 (operating in PCIe* mode) – up to 2 PCIe* slots (Optional) With three riser cards installed, up to 8 possible add-in cards can be supported: • 4 Full Height / Full Length + 2 Full Height / Half Length add-in cards via Risers #1 and #2 • 2 low profile add-in cards via Riser #3 (option) • See Chapter 10 for available riser card options.
Video	• Integrated 2D Video Controller • 16 MB DDR3 Memory

Feature	Description
On-board storage controllers and options	• 10 x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) • Two single port SATA connectors capable of supporting up to 6 Gb/sec • Two 4-port mini-SAS HD (SFF-8643) connectors capable of supporting up to 6 Gb/sec /SATA • One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices • Optional SAS IOC/ROC support via on-board Intel® Integrated RAID module connector • Embedded Software SATA RAID o Intel® Rapid Storage RAID Technology (RSTe) 4.0 o Intel® Embedded Server RAID Technology 2 (ESRT2) with optional RAID 5 key support
Security	• Intel® Trusted Platform Module (TPM) - AXXTPME5 (Accessory Option)
Server Management	• Integrated Baseboard Management Controller, IPMI 2.0 compliant • Support for Intel® Server Management Software • On-board RJ45 management port • Advanced Server Management via an Intel® Remote Management Module 4 Lite (Accessory Option)
Power Supply Options	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 Three power supply options: • AC 750W Platinum • DC 750W Gold • AC 1100W Platinum
Storage Options	**12Gb/sec Hot Swap Backplane Options:** • 8 x 2.5” SATA/SAS • 8 x 2.5” Combo Backplane - SATA/SAS + up to 4 x PCIe NVM Express* (Not Hot Swappable) • 8 x 2.5” Dual Port SATA/SAS • 8 x 3.5” SATA/SAS • 12 x 3.5” SATA/SAS **12 Gb/sec 24 port SAS Expander Support (Accessory Options)** • Internal mount • PCIe* add-in **Storage Bay Options:** • 8 x 3.5” SATA/SAS Hot Swap Drive Bays + Optical Drive support + front panel I/O • 12 x 3.5” SATA/SAS Hot Swap Drive Bays (Storage model) • 8 x 2.5” SATA/SAS Hot Swap Drive Bays + Optical Drive support + front panel I/O • 16 x 2.5” SATA/SAS Hot Swap Drive Bays + Optical Drive support + front panel I/O • 24 x 2.5” SATA/SAS Hot Swap Drive Bays (Storage model) • 2 x 2.5” SATA SSD Back of Chassis Hot Swap Drive Bays (Accessory Option) • 2 x internal fixed mount 2.5” SSDs (All SKUs)
Supported Rack Mount Kit Accessory Options	• AXXPRAIL – Tool-less rack mount rail kit – 800mm max travel length • AXXELVRAIL – Enhanced value rack mount rail kit - 424mm max travel length • AXX1U2UCMA – Cable Management Arm – (*supported with AXXPRAIL only) • AXX2POSTBRCKT – 2-post fixed mount bracket kit (not supported with 12 and 24 drive storage SKUs)

Refer to the Technical Product Specification for each Intel® Server System product family for more information.

Glossary

This appendix contains important terms used in this document. For ease of use, numeric entries are listed first (for example, “82460GX”) followed by alpha entries (for example, “AGP 4x”). Acronyms are followed by non-acronyms.

Term	Definition
ACPI	Advanced Configuration and Power Interface
AP	Application Processor
APIC	Advanced Programmable Interrupt Control
ARP	Address Resolution Protocal
ASIC	Application Specific Integrated Circuit
ASMI	Advanced Server Management Interface
BIOS	Basic Input/Output System
BIST	Built-In Self Test
BMC	Baseboard Management Controller
BPP	Bits per pixel
Bridge	Circuitry connecting one computer bus to another, allowing an agent on one to access the other
BSP	Bootstrap Processor
Byte	8-bit quantity
CBC	Chassis Bridge Controller (A microcontroller connected to one or more other CBCs, together they bridge the IPMB buses of multiple chassis.
CEK	Common Enabling Kit
CHAP	Challenge Handshake Authentication Protocol
CMOS	Complementary Metal-oxide-semiconductor In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory, which normally resides on the server board.
DHCP	Dynamic Host Configuration Protocal
DPC	Direct Platform Control
EEPROM	Electrically Erasable Programmable Read-Only Memory
EHCI	Enhanced Host Controller Interface
EMP	Emergency Management Port
EPS	External Product Specification
ESB2	Enterprise South Bridge 2
FBD	Fully Buffered DIMM
F MB	Flexible Mother Board
FRB	Fault Resilient Booting
FRU	Field Replaceable Unit
FSB	Front Side Bus
GB	1024 MB
GPA	Guest Physical Address
GPIO	General Purpose I/O
GTL	Gunning Transceiver Logic
HPA	Host Physical Address
HSC	Hot-swap Controller
Hz	Hertz (1 cycle/second)
I^2C	Inter-Integrated Circuit Bus

Term	Definition
IA	Intel® Architecture
IBF	Input Buffer
ICH	I/O Controller Hub
ICMB	Intelligent Chassis Management Bus
IERR	Internal Error
IFB	I/O and Firmware Bridge
ILM	Independent Loading Mechanism
IMC	Integrated Memory Controller
INTR	Interrupt
I/OAT	I/O Acceleration Technology
IOH	I/O Hub
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IR	Infrared
ITP	In-Target Probe
KB	1024 bytes
KCS	Keyboard Controller Style
KVM	Keyboard, Video, Mouse
LAN	Local Area Network
LCD	Liquid Crystal Display
LDAP	Local Directory Authentication Protocol
LED	Light Emitting Diode
LPC	Low Pin Count
LUN	Logical Unit Number
MAC	Media Access Control
MB	1024 KB
MCH	Memory Controller Hub
MD2	Message Digest 2 – Hashing Algorithm
MD5	Message Digest 5 – Hashing Algorithm – Higher Security
ME	Management Engine
MMU	Memory Management Unit
ms	Milliseconds
MTTR	Memory Type Range Register
Mux	Multiplexor
NIC	Network Interface Controller
NMI	Nonmaskable Interrupt
OBF	Output Buffer
OEM	Original Equipment Manufacturer
Ohm	Unit of electrical resistance
OVP	Over-voltage Protection
PECI	Platform Environment Control Interface
PEF	Platform Event Filtering
PEP	Platform Event Paging
PIA	Platform Information Area (This feature configures the firmware for the platform hardware)

Term	Definition
PLD	Programmable Logic Device
PMI	Platform Management Interrupt
POST	Power-On Self Test
PSMI	Power Supply Management Interface
PWM	Pulse-Width Modulation
QPI	QuickPath Interconnect
RAM	Random Access Memory
RASUM	Reliability, Availability, Serviceability, Usability, and Manageability
RISC	Reduced Instruction Set Computing
RMII	Reduced Media-Independent Interface
ROM	Read Only Memory
RTC	Real-Time Clock (Component of ICH peripheral chip on the server board)
SDR	Sensor Data Record
SECC	Single Edge Connector Cartridge
SEEPROM	Serial Electrically Erasable Programmable Read-Only Memory
SEL	System Event Log
SIO	Server Input/Output
SMBUS*	System Management BUS
SMI	Server Management Interrupt (SMI is the highest priority non-maskable interrupt)
SMM	Server Management Mode
SMS	Server Management Software
SNMP	Simple Network Management Protocol
SPS	Server Platform Services
SSE2	Streaming SIMD Extensions 2
SSE3	Streaming SIMD Extensions 3
SSE4	Streaming SIMD Extensions 4
TBD	To Be Determined
TDP	Thermal Design Power
TIM	Thermal Interface Material
UART	Universal Asynchronous Receiver/Transmitter
UDP	User Datagram Protocol
UHCI	Universal Host Controller Interface
URS	Unified Retention System
UTC	Universal time coordinare
VID	Voltage Identification
VRD	Voltage Regulator Down
VT	Virtualization Technology
Word	16-bit quantity
WS-MAN	Web Services for Management
ZIF	Zero Insertion Force

Reference Documents

- *Advanced Configuration and Power Interface Specification*, Revision 3.0, http://www.acpi.info/.
- *Intelligent Platform Management Bus Communications Protocol Specification*, Version 1.0. 1998. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- *Intelligent Platform Management Interface Specification*, Version 2.0. 2004. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification, Version 1.1, 02/01/02, Intel Corporation.
- Intel® Remote Management Module User's Guide, Intel Corporation.
- Alert Standard Format (ASF) Specification, Version 2.0, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., http://www.dmtf.org.
- *Intel® Server System BIOS External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 V3 product family – (Intel NDA Required)*
- *Intel® Server System BIOS Setup Utility Guide for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 V3 product family*
- *Intel® Server System BMC Firmware External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 V3 product family – (Intel NDA Required)*
- *SmaRT & CLST Architecture on Intel Systems and Power Supplies Specification (Doc Reference # 461024)*
- *Intel Integrated RAID Module RMS25PB080, RMS25PB040, RMS25CB080, and RMS25CB040 Hardware Users Guide*
- *Intel® Remote Management Module 4 Technical Product Specification*
- *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*
- *Intel® Server System R1000WT Technical Product Specification*
- *Intel® Server System R2000WT Technical Product Specification*
- *Intel® Ethernet Controller I350 Family Product Brief*
- *Intel® Ethernet Controller X540 Family Product Brief*
- *Intel® Chipset C610 product family ("Wellsburg") External Design Specification – (Intel NDA Required)*
- *Intel® Xeon® Processor E5-4600/2600/2400/1600 v3 Product Families ("Haswell") External Design Specification – (Intel NDA Required)*

NOTES