



Intel[®] Server Board S2600TP Product Family and Intel[®] Compute Module HNS2600TP Product Family

Technical Product Specification

A document providing an overview of product features, functions, architecture, and support specifications

REVISION 1.48

APRIL 2017

INTEL[®] SERVER PRODUCTS AND SOLUTIONS

Revision History

Date	Revision Number	Modifications
November, 2014	1.00	1 st External Public Release
March, 2015	1.10	Updated BMC Sensor Table Removed Appendix Node Manager 2.0 IPMI Integrated Sensors Added POST Error Beep Codes Added maximum RPM for BMC fan control
August, 2015	1.20	Corrected the Mellanox* Connect-IB* InfiniBand* name Removed the 6G SAS bridge board (option1) from the support list
August, 2015	1.30	Added Intel® Compute Module HNS2600TP24 Family
September, 2015	1.31	Updated the Intel® Compute Module HNS2600TP24 family weight
November, 2015	1.32	Updated the product code Corrected pin-out tables references
March, 2016	1.40	Added support for Intel® Xeon® processor E5-2600 v4 product family
April, 2016	1.41	Added Intel® Server Chassis H2224XXLR2 CFM spec Added FXX2130PCRPS Power Supply
May, 2016	1.42	Edited 12G Bridge Boards description to include IT/IMR mode specification
June, 2016	1.43	Added H2312XXLR2 and H2216XXLR2 references
August, 2016	1.44	Updated embedded RAID description Updated 6Gbps Bridge Board description
October, 2016	1.45	Updated the E5-2600 v4 memory speed supported reference Added Intel® Server Board S2600TPTR and Intel® Compute Module HNS2600TP24STR TPM references Intel® ESRT2 SATA DOM support for RAID-0 and RAID-1 Typographical corrections
January, 2017	1.46	Updated the "PCIe* Clock Source by Slot" table Explicitly state Slot #3 has lane reversal on the CPU side
March, 2017	1.47	Added Intel® Server Board S2600TPTR and Intel® Compute Module HNS2600TP24STR
April, 2017	1.48	Added S2600TPFR Mellanox IB card has no driver support for Windows OS Errata: Removed "ED2 – 4: CATERR due to CPU 3-strike timeout" from CATERR Sensor section

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 product Intel® Server Board S2600TP Product Family and Intel® Compute Module HNS2600TP Product Family 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 © 2017 Intel Corporation. All rights reserved.

Table of Contents

1	Introduction.....	1
1.1	Chapter Outline.....	1
1.2	Server Board Use Disclaimer	2
2	Product Features Overview.....	3
2.1	Components and Features Identification	6
2.1.1	Components Identification	7
2.2	Rear Connectors and Back Panel Feature Identification	7
2.3	Intel® Light Guided Diagnostic LED	9
2.4	Jumper Identification	9
2.5	Mechanical Dimension.....	10
2.6	Product Architecture Overview	11
2.7	Power Docking Board.....	15
2.7.1	Standard Power Docking Board.....	15
2.7.2	SAS/PCIe* SFF Combo Power Docking Board	15
2.8	Bridge Board.....	16
2.8.1	6G SATA Bridge Board	16
2.8.2	12G SAS Bridge Board	17
2.8.3	12G SAS Bridge Board with RAID 5	17
2.8.4	12G SAS/PCIe* SFF Combo Bridge Board	18
2.9	Riser Card	18
2.9.1	Riser Slot 1 Riser Card.....	19
2.9.2	Riser Slot 2 Riser Card.....	19
2.10	I/O Module Carrier	19
2.11	Compute Module Fans.....	21
2.12	Air Duct.....	23
2.13	Intel® RAID C600 Upgrade Key.....	23
2.14	Intel® Remote Management Module 4 (Intel® RMM4) Lite.....	24
2.15	Breakout Board	24
2.16	System Software Overview	26
2.16.1	System BIOS.....	26
2.16.2	Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data.....	30
2.16.3	Baseboard Management Controller (BMC) Firmware	31
3	Processor Support.....	32
3.1	Processor Socket Assembly	32
3.2	Processor Thermal Design Power (TDP) Support	33

3.3	Processor Population Rules.....	33
3.4	Processor Initialization Error Summary.....	34
3.5	Processor Function Overview	37
3.5.1	Processor Core Features	37
3.5.2	Supported Technologies	37
3.6	Processor Heat Sink.....	40
4	Memory Support.....	41
4.1	Memory Subsystem Architecture.....	41
4.1.1	IMC Modes of Operation	42
4.1.2	Memory RASM Features.....	43
4.2	Supported DDR4-2400 memory for Intel® Xeon processor v4 Product Family....	44
4.3	Supported DDR4-2133 memory for Intel® Xeon processor v4 Product Family....	45
4.4	Memory Slot Identification and Population Rules	45
4.5	System Memory Sizing and Publishing.....	48
4.5.1	Effects of Memory Configuration on Memory Sizing	48
4.5.2	Publishing System Memory	49
4.5.3	Memory Initialization	49
5	Server Board I/O	53
5.1	PCI Express* Support.....	53
5.1.1	PCIe Enumeration and Allocation	53
5.1.2	PCIe Non-Transparent Bridge (NTB).....	54
5.2	Add-in Card Support	55
5.2.1	Riser Card Support	55
5.3	Serial ATA (SATA) Support.....	58
5.3.1	Staggered Disk Spin-Up	59
5.4	Embedded SATA RAID Support.....	60
5.4.1	Intel® Rapid Storage Technology (RSTe) 4.0	60
5.4.2	Intel® Embedded Server RAID Technology 2 (ESRT2).....	61
5.5	Network Interface.....	63
5.5.1	MAC Address Definition	64
5.5.2	LAN Manageability.....	64
5.6	Video Support	64
5.7	Universal Serial Bus (USB) Ports	65
5.8	Serial Port.....	66
5.9	InfiniBand* Controller	66
5.9.1	Device Interfaces	66
5.9.2	Quad Small Form-factor Pluggable (QSFP+) Connector.....	68

6	Connector and Header	69
6.1	Power Connectors	69
6.1.1	Main Power Connector	69
6.1.2	Backup Power Control Connector.....	69
6.2	System Management Headers.....	70
6.2.1	Intel® Remote Management Module 4 (Intel® RMM4) Lite Connector.....	70
6.2.2	IPMB Header.....	70
6.2.3	Control Panel Connector	70
6.3	Bridge Board Connector.....	71
6.3.1	Power Button	72
6.4	I/O Connectors.....	73
6.4.1	PCI Express* Connectors.....	73
6.4.2	VGA Connector.....	83
6.4.3	NIC Connectors	83
6.4.4	SATA Connectors	84
6.4.5	SATA SGPIO Connectors.....	84
6.4.6	Hard Drive Activity (Input) LED Header	85
6.4.7	Intel® RAID C600 Upgrade Key Connector.....	85
6.4.8	Serial Port Connectors.....	85
6.4.9	USB Connectors.....	86
6.4.10	3QSFP+ for InfiniBand*	86
6.4.11	UART Header.....	87
6.5	Fan Headers.....	87
6.5.1	FAN Control Cable Connector	87
6.5.2	Discrete System FAN Connector.....	88
6.6	Power Docking Board Connectors.....	88
7	Configuration Jumpers	91
7.1	BMC Force Update (J7A2)	92
7.2	ME Force Update (J5D2).....	93
7.3	Password Clear (J7A6).....	93
7.4	BIOS Recovery Mode (J7A7).....	94
7.5	BIOS Default (J7A3).....	96
8	Intel® Light-Guided Diagnostics.....	97
8.1	Status LED	97
8.2	ID LED.....	100
8.3	BMC Boot/Reset Status LED Indicators	100
8.4	InfiniBand* Link/Activity LED	101

8.5	POST Code Diagnostic LEDs	101
9	Platform Management.....	103
9.1	Management Feature Set Overview	103
9.1.1	IPMI 2.0 Features Overview	103
9.1.2	Non IPMI Features Overview	104
9.2	Platform Management Features and Functions.....	106
9.2.1	Power Subsystem	106
9.2.2	Advanced Configuration and Power Interface (ACPI)	106
9.2.3	System Initialization.....	107
9.2.4	System Event Log (SEL)	108
9.3	Sensor Monitoring	108
9.3.1	Sensor Scanning.....	108
9.3.2	Sensor Rearm Behavior	108
9.3.3	BIOS Event-Only Sensors	109
9.3.4	Margin Sensors.....	110
9.3.5	IPMI Watchdog Sensor	110
9.3.6	BMC Watchdog Sensor	110
9.3.7	BMC System Management Health Monitoring.....	110
9.3.8	VR Watchdog Timer	110
9.3.9	System Airflow Monitoring.....	110
9.3.10	Thermal Monitoring	111
9.3.11	Processor Sensors	114
9.3.12	Voltage Monitoring.....	117
9.3.13	Fan Monitoring	117
9.3.14	Standard Fan Management.....	119
9.3.15	Power Management Bus (PMBus*).....	126
9.3.16	Power Supply Dynamic Redundancy Sensor	126
9.3.17	Component Fault LED Control	127
9.3.18	CMOS Battery Monitoring.....	128
9.4	Intel® Intelligent Power Node Manager (NM).....	128
9.4.1	Hardware Requirements	129
9.4.2	Features.....	129
9.4.3	ME System Management Bus (SMBus*) Interface.....	129
9.4.4	PECI 3.0	129
9.4.5	NM “Discovery” OEM SDR.....	129
9.4.6	SmaRT/CLST	130
9.5	Basic and Advanced Server Management Features	131

9.5.1	Dedicated Management Port	132
9.5.2	Embedded Web Server.....	132
9.5.3	Advanced Management Feature Support (RMM4 Lite).....	134
10	Thermal Management	139
11	System Security	141
11.1	Password Setup	141
11.1.1	System Administrator Password Rights	142
11.1.2	Authorized System User Password Rights and Restrictions.....	142
11.2	Front Panel Lockout.....	143
11.3	Trusted Platform Module (TPM) support.....	143
11.3.1	TPM security BIOS	144
11.3.2	Physical presence	144
11.3.3	TPM security setup options	144
11.4	Intel® Trusted eXecution Technology (TXT)	146
12	Environmental Limits Specification	147
13	Power Supply Specification Guidelines	148
13.1	Power Supply DC Output Connector.....	148
13.2	Power Supply DC Output Specification.....	148
13.2.1	Output Power/Currents.....	148
	Standby Output.....	149
13.2.2	Voltage Regulation	149
13.2.3	Dynamic Loading.....	149
13.2.4	Capacitive Loading	149
13.2.5	Grounding.....	150
13.2.6	Closed-loop Stability	150
13.2.7	Residual Voltage Immunity in Standby Mode.....	150
13.2.8	Common Mode Noise.....	150
13.2.9	Soft Starting	150
13.2.10	Zero Load Stability Requirements	151
13.2.11	Hot Swap Requirements.....	151
13.2.12	Forced Load Sharing.....	151
13.2.13	Ripple/Noise.....	151
13.2.14	Timing Requirement	151
Appendix A:	Integration and Usage Tips	155
Appendix B:	Integrated BMC Sensor Tables.....	156
Appendix C:	BIOS Sensors and SEL Data.....	173
Appendix D:	POST Code Diagnostic LED Decoder	180

Appendix E: POST Code Errors 187
 POST Error Beep Codes 190
Appendix F: Statement of Volatility 191
Glossary 193
Reference Documents..... 195



List of Figures

Figure 1. Intel® Server Board S2600TPFR (demo picture)	3
Figure 2. Intel® Compute Module HNS2600TPFR (demo picture).....	4
Figure 3. Server Board Components (S2600TPFR).....	7
Figure 4. Compute Module Components.....	7
Figure 5. Server Board Rear Connectors	8
Figure 6. HNS2600TPR Compute Module Back Panel	8
Figure 7. HNS2600TP24R Compute Module Back Panel.....	8
Figure 8. Intel® Light Guided Diagnostic LED	9
Figure 9. Jumper Identification	9
Figure 10. Server Board Dimension.....	10
Figure 11. Compute Module Dimension.....	10
Figure 12. Intel® Server Board S2600TPR Block Diagram	12
Figure 13. Intel® Server Board S2600TPFR Block Diagram.....	13
Figure 14. Intel® Server Board S2600TPTR Block Diagram.....	14
Figure 15. SAS/PCIe* SFF Combo Power Docking Board Top View	16
Figure 16. 6G SATA Bridge Board Overview	17
Figure 17. 12G SAS Bridge Board Overview	17
Figure 18. 12G SAS Bridge Board with RAID 5 Overview	18
Figure 19. SAS/PCIe* SFF Combo Bridge Board Overview.....	18
Figure 20. Riser Card for Riser Slot #1	19
Figure 21. Riser Card for Riser Slot #2	19
Figure 22. I/O Module Carrier Installation	19
Figure 23. Installing the M.2 Device	20
Figure 24. AXXKPTPM2IOM I/O Module Carrier Connectors	21
Figure 25. Connecting the M.2 SATA Cable	21
Figure 26. Compute Module Fans	22
Figure 27. Air Duct	23
Figure 28. Intel® RAID C600 Upgrade Key	23
Figure 29. Intel® RMM4 Lite	24
Figure 30. Breakout Board Front and Rear View.....	25
Figure 31. Breakout Board Mechanical Drawing (Unit: mm).....	25
Figure 32. Processor Socket Assembly.....	32
Figure 33. Processor Socket ILM	32
Figure 34. Processor Heat Sink Overview.....	40
Figure 35. Integrated Memory Controller Functional Block Diagram	41

Figure 36. Intel® Server Board S2600TPR Product Family DIMM Slot Layout	47
Figure 37. Add-in Card Support Block Diagram (S2600TPR)	55
Figure 38. Server Board Riser Slots (S2600TPFR)	55
Figure 39. SATA Support	58
Figure 40. SATA RAID 5 Upgrade Key	62
Figure 41. Network Interface Connectors	63
Figure 42. RJ45 NIC Port LED	64
Figure 43. USB Ports Block Diagram	66
Figure 44. Serial Port A Location	66
Figure 45. Jumper Location	91
Figure 46. Status LED (E) and ID LED (D)	97
Figure 47. InfiniBand* Link LED (K) and InfiniBand* Activity LED (J)	101
Figure 48. Rear Panel Diagnostic LEDs	102
Figure 49. High-level Fan Speed Control Process	123
Figure 50. Air Flow and Fan Identification	139
Figure 51. Turn On/Off Timing (Power Supply Signals – 5VSB)	153
Figure 52. Turn On/Off Timing (Power Supply Signals – 12VSB)	154
Figure 53. Diagnostic LED Placement Diagram	180

List of Tables

Table 1. Intel® Server Board S2600TPR Product Family Feature Set.....	4
Table 2. Intel® Compute Module HNS2600TPR Product Family Feature Set.....	6
Table 3. Product Weight and Packaging	11
Table 5. POST Hot-Keys.....	28
Table 6. Mixed Processor Configurations Error Summary.....	35
Table 9. DIMM Nomenclature	46
Table 10. Supported DIMM Populations.....	47
Table 11. PCIe* Port Routing – CPU 1	56
Table 12. PCIe* Port Routing – CPU 2	56
Table 13. SATA and sSATA Controller BIOS Utility Setup Options	58
Table 14. SATA and sSATA Controller Feature Support	59
Table 15. Onboard Video Resolution and Refresh Rate (Hz).....	65
Table 16. Network Port Configuration	67
Table 17. Main Power Supply Connector 6-pin 2x3 Connector.....	69
Table 18. Backup Power Control Connector	69
Table 19. Intel® RMM4 Lite Connector Pin-out	70
Table 20. IPMB Header 4-pin	70
Table 21. Control Panel Connector.....	70
Table 22. Bridge Board Connector	71
Table 23. SATA DOM Connector Pin-out.....	72
Table 24. USB 2.0 Type-A Connector Pin-out.....	73
Table 25. 5V_AUX Power Connector Pin-out.....	73
Table 26. CPU1 and CPU2 PCIe* Bus Connectivity.....	73
Table 27. PCIe* x16 Riser Slot 1 Connector	74
Table 28. PCIe* x24 Riser Slot 2 Connector.....	76
Table 29. PCIe* x24 Riser Slot 3 Connector.....	78
Table 30. PCIe* x16 Riser Slot 4 Connector.....	81
Table 31. PCIe* Riser ID Assignment.....	82
Table 32. PCIe* Clock Source by Slot	83
Table 33. VGA External Video Connector	83
Table 34. RJ-45 10/100/1000 NIC Connector Pin-out.....	84
Table 35. SATA Connector	84
Table 36. SATA SGPIO Connector	85
Table 37. SATA HDD Activity (Input) LED Header	85
Table 38. Storage Upgrade Key Connector	85

Table 39. Internal 9-pin Serial A	85
Table 40. External USB port Connector.....	86
Table 41. Internal USB Connector	86
Table 42. QSFP+ Connector	86
Table 43. UART Header	87
Table 44. Baseboard Fan Connector.....	87
Table 45. Baseboard Fan Connector.....	88
Table 46. Main Power Input Connector	88
Table 47. Fan Control Signal Connector	89
Table 48. Compute Module Fan Connector	89
Table 49. Main Power Output Connector	89
Table 50. 40 pin Misc. Signal Connector (HNS2600TP24R/HNS2600TP24SR only).....	89
Table 51. Jumper Modes Selection.....	91
Table 52. Force Integrated BMC Update Jumper (J7A2)	92
Table 53. Force ME Update Jumper (J5D2).....	93
Table 54. Password Clear Jumper (J7A6).....	94
Table 55. BIOS Recovery Mode Jumper (J7A7)	95
Table 56. BIOS Default Jumper.....	96
Table 57. Status LED State Definitions	98
Table 58. ID LED	100
Table 59. BMC Boot/Reset Status LED Indicators	100
Table 60. InfiniBand* Link/Activity LED	101
Table 61. ACPI Power States	106
Table 62. Processor Sensors.....	114
Table 63. Processor Status Sensor Implementation.....	115
Table 64. Component Fault LEDs.....	127
Table 65. Intel® Remote Management Module 4 (RMM4) Options.....	131
Table 66. Basic and Advanced Server Management Features Overview.....	131
Table 67. Air Flow	139
Table 68. TPM Setup Utility – Security Configuration Screen Fields	145
Table 69. Server Board Design Specifications	147
Table 70. Power Supply DC Power Input Connector Pin-out.....	148
Table 71. Minimum 1200W/1600W Load Ratings.....	148
Table 72. Minimum 2130W Load Ratings.....	148
Table 73. Voltage Regulation Limits	149
Table 74. Transient Load Requirements	149
Table 75. Capacitive Loading Conditions.....	150

Table 76. Ripples and Noise	151
Table 77. Timing Requirements (5VSB)	152
Table 78. Timing Requirements (12VSB)	153
Table 79. BMC Sensor Table	158
Table 80. BIOS Sensor and SEL Data	173
Table 81. POST Code LED Example	181
Table 82. MRC Fatal Error Codes	181
Table 83. MRC Progress Codes	182
Table 84. POST Progress Codes	183
Table 85. POST Error Codes and Messages	187
Table 86. POST Error Beep Codes	190
Table 87. Glossary	193

<This page is intentionally left blank.>

1 Introduction

This *Technical Product Specification (TPS)* provides specific information detailing the features, functionality, and high-level architecture of the Intel® Server Board S2600TP product family and the Intel® Compute Module HNS2600TP product family.

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 – Memory Support
- Chapter 5 – Server Board I/O
- Chapter 6 – Connector and Header
- Chapter 7 – Configuration Jumpers
- Chapter 8 – Intel® Light-Guided Diagnostics
- Chapter 9 – Platform Management
- Chapter 10 – Thermal Management
- Chapter 11 – System Security
- Chapter 12 – Environmental Limits Specification
- Chapter 13 – Power Supply Specification Guidelines
- Appendix A – Integration and Usage Tips
- Appendix B – Integrated BMC Sensor Tables
- Appendix C – BIOS Sensors and SEL Data
- Appendix D – POST Code Diagnostic LED Decoder
- Appendix E – POST Code Errors
- Appendix F – Statement of Volatility
- Glossary
- Reference Documents

1.2 Server Board Use Disclaimer

Intel Corporation server boards contain a number of high-density VLSI (Very Large Scale Integration) 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 air flow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of their published operating or non-operating limits.

2 Product Features Overview

The Intel® Server Board S2600TP product family is a monolithic printed circuit board (PCB) assembly with features designed to support the high performance and high density computing markets. This server board is designed to support the Intel® Xeon® processor E5-2600 v3 and v4 product family. Previous generation Intel® Xeon® processors are not supported.

The Intel® Server Board S2600TP product family contains two server board options. Many of the features and functions of the server board family are common. A board will be identified by its name which has described features or functions unique to it.

- S2600TPFR – With onboard InfiniBand® controller providing one external rear QSFP+ port



Figure 1. Intel® Server Board S2600TPFR (demo picture)

- S2600TPR – Without onboard InfiniBand® controller, no external rear QSFP+ port
- S2600TPTR – S2600TPR board with on-board TPM 2.0 chip
- S2600TPNR – S2600TPR board with ADR-on-Reset cold-reset feature support when using an NVDIMM module

The Intel® Compute Module HNS2600TP product family provides two compute module options, each integrated with either of the server board from the Intel® Server Board S2600TP product family.



Figure 2. Intel® Compute Module HNS2600TPFR (demo picture)

The following table provides a high-level product feature list.

Table 1. Intel® Server Board S2600TPR Product Family Feature Set

Feature	Description
Processor Support	<ul style="list-style-type: none">Two LGA2011-3 (Socket R3) processor socketsSupport for one or two Intel® Xeon® processors E5-2600 v3 and v4 product familyMaximum supported Thermal Design Power (TDP) of up to 160 W
Memory Support	<ul style="list-style-type: none">Sixteen DIMM slots in total across eight memory channelsRegistered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM)Memory DDR4 data transfer rates of 1600/1866/2133/2400 MT/s
Chipset	Intel® C612chipset
External I/O Connections	<ul style="list-style-type: none">DB-15 video connectorTwo RJ-45 1GbE Network Interface Controller (NIC) portsOne dedicated RJ-45 port for remote server managementOne stacked two port USB 2.0 (port 0/1) connectorOne InfiniBand* FDR QSFP+ port (S2600TPFR only)

Feature	Description
Internal I/O connectors/headers	<ul style="list-style-type: none"> ▪ Bridge slot to extend board I/O <ul style="list-style-type: none"> ○ Four SATA 6Gb/s signals to backplane ○ Front control panel signals ○ One SATA 6Gb/s port for SATA DOM ○ One USB 2.0 connector (port 10) ▪ One internal USB 2.0 connector (port 6/7) ▪ One 2x7 pin header for system fan module ▪ One 1x12 pin control panel header ▪ One DH-10 serial Port A connector ▪ One SATA 6Gb/s port for SATA DOM ▪ Four SATA 6Gb/s connectors (port 0/1/2/3) ▪ One 2x4 pin header for Intel® RMM4 Lite ▪ One 1x4 pin header for Storage Upgrade Key ▪ One 1x8 pin backup power control connector
PCIe Support	PCIe* 3.0 (2.5, 5, 8 GT/s)
Power Connections	Two sets of 2x3 pin connectors (main power 1/2)
System Fan Support	<ul style="list-style-type: none"> ▪ One 2x7 pin fan control connector for Intel compute module and chassis ▪ Three 1x8 pin fan connectors for third-party chassis
Video	<ul style="list-style-type: none"> ▪ Integrated 2D video graphics controller ▪ 16MB DDR3 memory
Riser Support	<ul style="list-style-type: none"> ▪ Four riser slots <ul style="list-style-type: none"> ○ Riser slot 1 provides x16 PCIe* 3.0 lanes ○ Riser slot 2 provides <ul style="list-style-type: none"> ▪ x24 PCIe* 3.0 lanes for S2600TPR ▪ x16 PCIe* 3.0 lanes for S2600TPFR ○ Riser slot 3 provides x24 PCIe* 3.0 lanes ○ Riser slot 4 provides x16 PCIe* 3.0 lanes ▪ One bridge board slot for board I/O expansion
On-board storage controllers and options	<ul style="list-style-type: none"> ▪ 5 on-board SATA 6Gb/s ports, one of them is SATA DOM ▪ 5 ports SATA 6Gb/s signal is integrated in the bridge board slot and connect to backplane via bridge slot.
RAID Support	<ul style="list-style-type: none"> ▪ Intel® Rapid Storage RAID Technology (RSTe) 4.0 ▪ Intel® Embedded Server RAID Technology 2 (ESRT2) with optional Intel® RAID C600 Upgrade Key to enable SATA RAID 5
Server Management	<ul style="list-style-type: none"> ▪ Onboard Emulex* Pilot III* Controller ▪ Support for Intel® Remote Management Module 4 Lite solutions ▪ Intel® Light-Guided Diagnostics on field replaceable units ▪ Support for Intel® System Management Software ▪ Support for Intel® Intelligent Power Node Manager (Need PMBus*-compliant power supply)

Warning! The riser slot 1 on the server board is designed for plugging in ONLY the riser card. Plugging in the PCIe* card may cause permanent server board and PCIe* card damage.

Table 2. Intel® Compute Module HNS2600TPR Product Family Feature Set

Feature ¹	Description
Server Board	Intel® Server Board S2600TPR product family <ul style="list-style-type: none"> ▪ HNS2600TPR – include Intel® Server Board S2600TPR ▪ HNS2600TPFR – include Intel® Server Board S2600TPFR ▪ HNS2600TPNR – include Intel® Server Board S2600TPNR ▪ HNS2600TP24R – include Intel® Server Board S2600TPR and Dual Port Intel® X540 10GbE I/O Module (RJ45) ▪ HNS2600TP24SR – include Intel® Server Board S2600TPR and Dual Port Intel® 82599 10GbE I/O Module (SFP+) ▪ HNS2600TP24STR – include Intel® Server Board S2600TPTR and Dual Port Intel® 82599 10GbE I/O Module (SFP+)
Processor Support	Maximum supported Thermal Design Power (TDP) of up to 145 W
Heat Sink	<ul style="list-style-type: none"> ▪ One Cu/Al 84x106mm heat sink for CPU 1 ▪ One Ex-Al 84x106mm heat sink for CPU 2
Fan	Three sets of 40x56mm dual rotor system fans
Riser Support	<ul style="list-style-type: none"> ▪ One riser card with bracket in riser slot 1 to support one PCIe* 3.0 x16 low profile card (default)² ▪ One I/O module riser and carrier kit in riser slot 2 to support an Intel® I/O Expansion Module (optional) <i>Note: Riser slot 3 and 4 cannot be used with the bridge board installed.</i>
Compute Module Board	<ul style="list-style-type: none"> ▪ Four types of bridge boards: <ul style="list-style-type: none"> ○ 6G SATA Bridge Board (Default in HNS2600TPR/HNS2600TPFR) ○ 12G SAS Bridge Board (Optional for HNS2600TPR/HNS2600TPFR) ○ 12G SAS Bridge Board with RAID 5 (Optional for HNS2600TPR/HNS2600TPFR) ○ 12G SAS/NVMe Combo Bridge Board (Default in HNS2600TP24R/HNS2600TP24SR/ HNS2600TP24STR) ▪ One compute module power docking board
Air Duct	One transparent air duct

Notes:

1. The table only lists features that are unique to the compute module or different from the server board.
2. ONLY low profile PCIe* card can be installed on riser slot 1 riser card of the compute module.

2.1 Components and Features Identification

This section provides a general overview of the server board and compute module, identifying key features and component locations. The majority of the items identified are common in the product family.

2.1.1 Components Identification




Figure 3. Server Board Components (S2600TPFR)





Figure 4. Compute Module Components

2.2 Rear Connectors and Back Panel Feature Identification

The Intel® Server Board S2600TPR product family has the following board rear connector placement.



	Description		Description
A	NIC port 1 (RJ45)	G	Dedicated Management Port (RJ45)
B	NIC port 2 (RJ45)	H	InfiniBand* Port (QSFP+, S2600TPFR only)
C	Video out (DB-15)	I	POST Code LEDs (8 LEDs)
D	ID LED	J	InfiniBand* Activity LED (S2600TPFR only)
E	Status LED	K	InfiniBand* Link LED (S2600TPFR only)
F	Dual port USB		

Figure 5. Server Board Rear Connectors

The Intel® Compute Module HNS2600TP product family has the following back panel features.




Figure 6. HNS2600TPR Compute Module Back Panel




Figure 7. HNS2600TP24R Compute Module Back Panel

2.3 Intel® Light Guided Diagnostic LED




Figure 8. Intel® Light Guided Diagnostic LED

2.4 Jumper Identification




Figure 9. Jumper Identification

2.5 Mechanical Dimension




Figure 10. Server Board Dimension

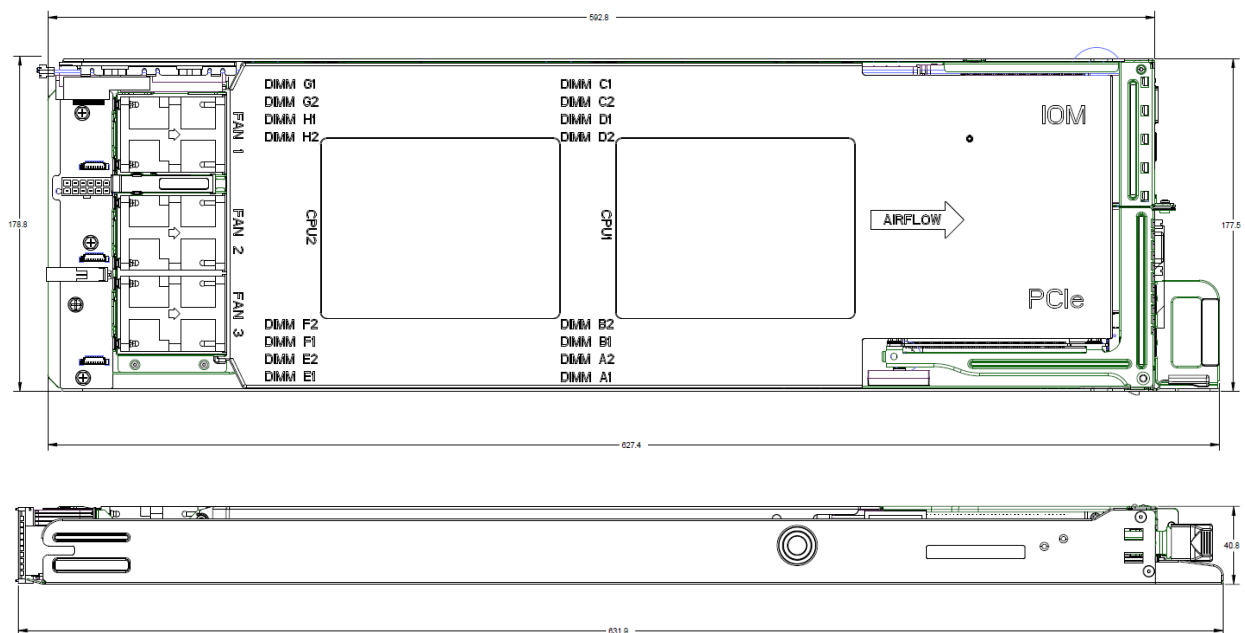


Figure 11. Compute Module Dimension

Approximate product weight is listed in the following table for reference. Variations are expected with real shipping products.

Table 3. Product Weight and Packaging

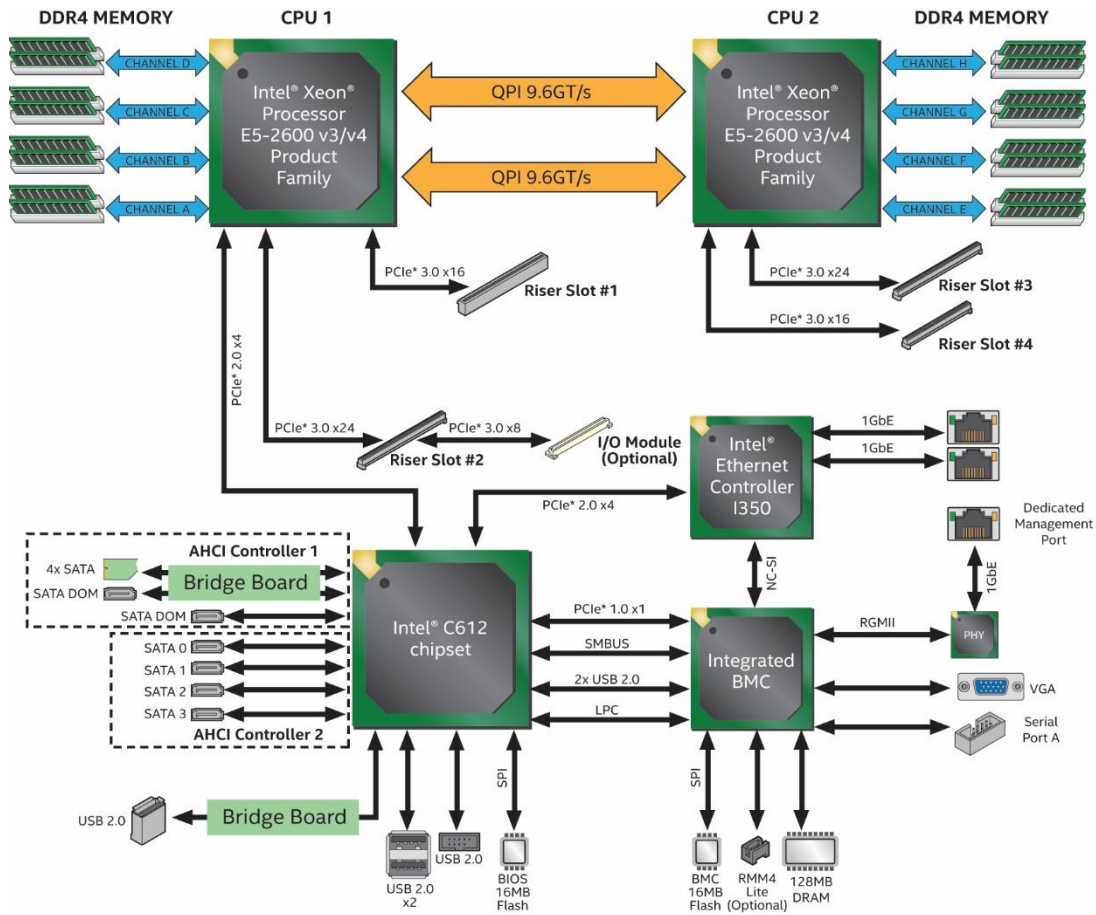
Product Code	Quantity per Box	Box Dimension (mm)	Net Weight (kg)	Package Weight (kg)
BBS2600TPR	10	578X438X301	15.2	20.7
BBS2600TPFR	10	578X438X301	15.9	21.4
BBS2600TPTR	10	578X438X301	15.2	20.7
BBS2600TPNR	10	578X438X301	15.2	20.7
HNS2600TPR	1	716X269X158	3.5	4.7
HNS2600TPFR	1	716X269X158	3.6	4.8
HNS2600TPNR	1	716X269X158	3.5	4.7
HNS2600TP24R	1	716X269X158	3.56	4.78
HNS2600TP24SR	1	716X269X158	3.56	4.78
HNS2600TP24STR	1	716X269X158	3.56	4.78

2.6 Product Architecture Overview

The Intel® Server Board S2600TP product family is a purpose built, rack-optimized, liquid cooling friendly server board used in a high-density rack system. It is designed around the integrated features and functions of the Intel® Xeon® processor E5-2600 v3 and v4 product family, the Intel® C612 chipset, and other supporting components including the Integrated BMC, the Intel® I350 network interface controller, and the Mellanox® Connect-IB® InfiniBand® (S2600TPFR only).

The half-width board size allows four boards to reside in a standard multi-compute module 2U Intel® Server Chassis H2000G product family, for high-performance and high-density computing platforms.

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



AF006716

Figure 12. Intel® Server Board S2600TPR Block Diagram

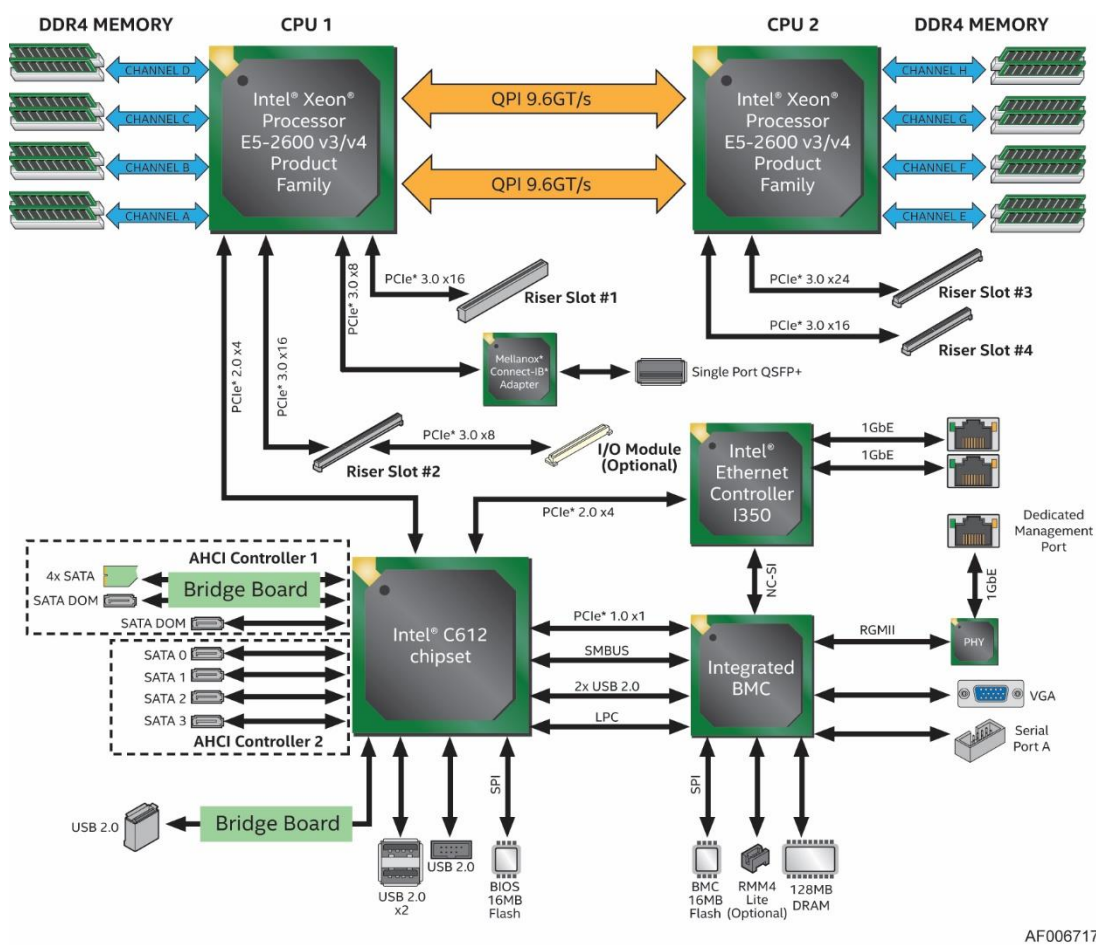


Figure 13. Intel® Server Board S2600TPFR Block Diagram

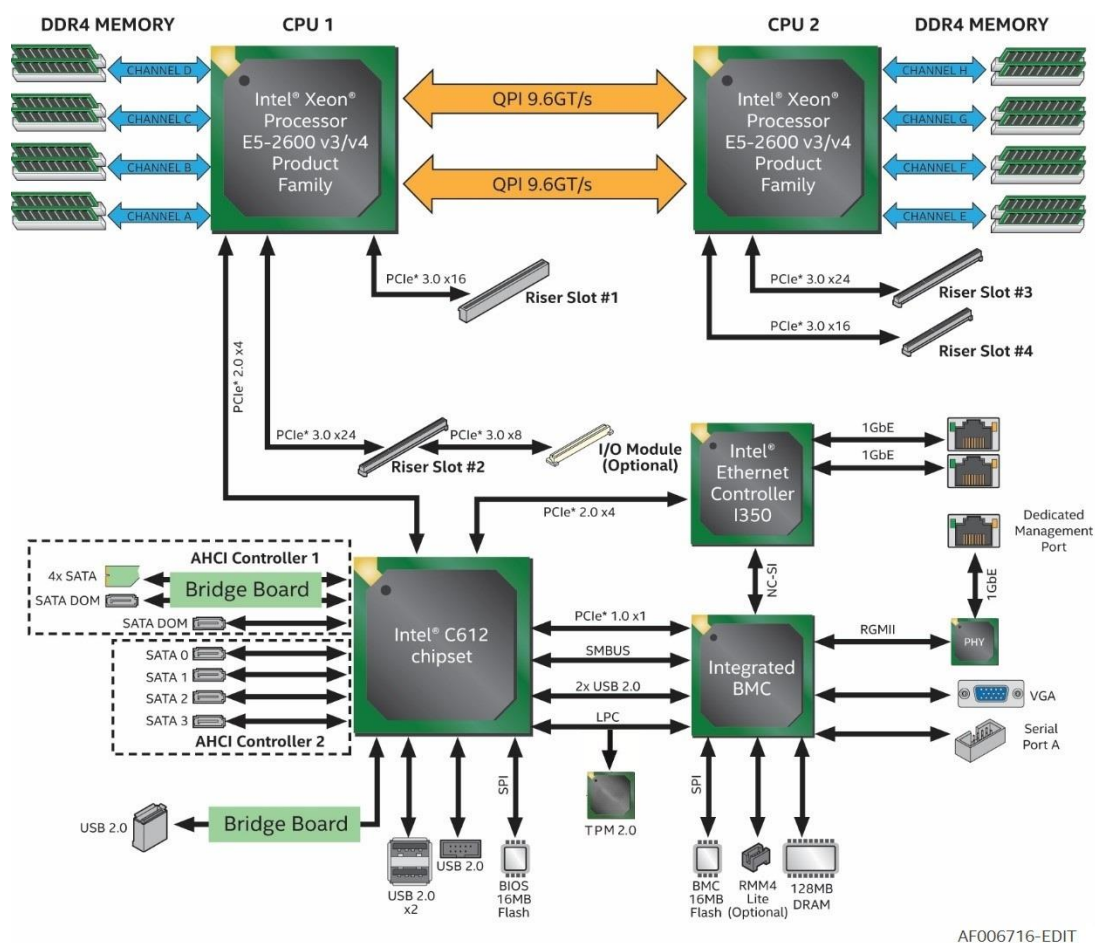



Figure 14. Intel® Server Board S2600TPTR Block Diagram

The Intel® Compute Module HNS2600TP product family provides a series of features including the power docking board, bridge boards, riser cards, fans, and the air duct.

2.7 Power Docking Board

2.7.1 Standard Power Docking Board


The power docking board provides hot swap docking of 12V main power between the compute module and the server. It supports three dual rotor fan connections, 12V main power hot swap controller, and current sensing. The standard power docking board is intended to support the usage of Intel® Compute Module HNS2600TPR or HNS2600TPFR.



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

2.7.2 SAS/PCIe* SFF Combo Power Docking Board

The SAS/PCIe* SFF Combo Power Docking Board is only used in Intel® Compute Module HNS2600TP24R, HNS2600TP24SR or HNS2600TP24STR.



A	40 pin Misc. Signal Connector (to bridge board)
B	2x7 pin Fan Control Connector
C	8 pin Connector for Fan 1

D	2x6 pin Main Power Output Connector
E	8 pin Connector for Fan 2
F	8 pin Connector for Fan 3
G	Power Blade Card Edge Connector (to BIB)
H	40 pin Misc. Signal Card Edge Connector (to BIB)

Figure 15. SAS/PCIe* SFF Combo Power Docking Board Top View

2.8 Bridge Board

There are five types of bridge boards that implement different features and functions.

- 6G SATA bridge board (default)
- 12G SAS bridge board (optional)
- 12G SAS bridge board with RAID 5 (optional)
- 12G SAS/PCIe* SFF Combo Bridge Board (default)


Computer Module	6G SATA bridge board	12G SAS bridge board	12G SAS bridge board with RAID 5	12G SAS/PCIe* SFF Combo Bridge Board
HNS2600TPR	Default	Optional	Optional	Not Support
HNS2600TPFR	Default	Optional	Optional	Not Support
HNS2600TP24R	Not Support	Not Support	Not Support	Default
HNS2600TP24SR	Not Support	Not Support	Not Support	Default
HNS2600TP24STR	Not Support	Not Support	Not Support	Default

Table 4. Computer Module and Bridge board support matrix

Note: All 12G SAS bridge boards require two processors installed to be functional.

2.8.1 6G SATA Bridge Board

The 6G SATA bridge board provides hot swap interconnect of all electrical signals to the backplane of the server chassis (except for main 12V power). It supports up to 4x lanes of SATA, a 7-pin SATA connector for SATA DOM devices, and a type-A USB connector for USB flash device. One bridge board is used per one compute module. The bridge board is secured with three loose screws to the compute module. The bridge board support [embedded SATA RAID Support](#).




Label	Description
A	2x40-pin card edge connector (to the backplane)
B	USB 2.0 Type-A connector
C	2-pin 5V power
D	SATA DOM port connector
E	2x40-pin card edge connector (to the bridge board connector on the server board)

Figure 16. 6G SATA Bridge Board Overview

2.8.2 12G SAS Bridge Board

The optional 12G SAS bridge board has one embedded LSI* SAS 3008 (IMR mode) controller to support up to four SAS/SATA ports with RAID 0, 1, and 10 support, a 7-pin SATA connector for SATA DOM devices, and a UART (Universal Asynchronous Receiver/Transmitter) header. One bridge board is used per one compute module, connecting to the bridge board slot and Riser Slot 3.




Label	Description
A	2x40-pin card edge connector (to the backplane)
B	UART header
C	2-pin 5V power
D	SATA DOM port connector
E	2x40-pin card edge connector (to the bridge board connector on the server board)
F	200-pin connector (to Riser Slot 3 on the server board)

Figure 17. 12G SAS Bridge Board Overview

2.8.3 12G SAS Bridge Board with RAID 5

The optional 12G SAS bridge board with RAID 5 has one embedded LSI* SAS 3008 (IMR mode) controller to support up to four SAS/SATA ports with RAID 0, 1, 10, and RAID 5 support, a 7-pin SATA connector for SATA DOM devices, and a UART (Universal Asynchronous Receiver/Transmitter) header. One bridge board is used per one compute module, connecting to the bridge board slot and Riser Slot 3.




Label	Description
A	2x40-pin card edge connector (to the backplane)
B	UART header
C	2-pin 5V power
D	SATA DOM port connector
E	2x40-pin card edge connector (to the bridge board connector on the server board)
F	200-pin connector (to Riser Slot 3 on the server board)

Figure 18. 12G SAS Bridge Board with RAID 5 Overview

2.8.4 12G SAS/PCIe* SFF Combo Bridge Board

The 12G SAS/PCIe* SFF combo bridge board has one embedded LSI* SAS 3008 (IT mode) controller to support up to six 12Gb/s SAS ports, two x4 PCIe* 3.0 lanes to support up to two PCIe* SFF devices, one 7-pin SATA connector for SATA DOM devices, and one UART header. One bridge board is pre-installed in each compute module for 24 x 2.5" drive solution, connecting to the bridge board slot and Riser Slot 3.



A	100 pin card edge connector (to backplane)
B	UART header
C	SATA DOM connector
D	80 pin card edge connector (to server board)
E	200 pin card edge connector (to server board)
F	40 pin Misc. Signal Card Edge Connector (to power docking board)

Figure 19. SAS/PCIe* SFF Combo Bridge Board Overview

Note: All 12G SAS Bridge Board requires CPU2 installed to be functional.

2.9 Riser Card

There are two types of riser cards:

- Riser slot 1 riser card (for Riser slot 1 only)
- Riser slot 2 riser card (for Riser slot 2 only)

2.9.1 Riser Slot 1 Riser Card

The riser card for Riser Slot 1 has one PCIe* 3.0 x16 slot.




Figure 20. Riser Card for Riser Slot #1

2.9.2 Riser Slot 2 Riser Card

The riser card for Riser Slot 2 has one PCIe* 3.0 x16 slot (x8 lanes are for I/O module carrier) which can only support Intel® I/O Module carrier.




Figure 21. Riser Card for Riser Slot #2

2.10 I/O Module Carrier

To broaden the standard on-board feature set, the server board supports the option of adding a single I/O module providing external ports for a variety of networking interfaces. The I/O module attaches to a high density 80-pin connector of the I/O module carrier on the riser slot 2 riser card.




Figure 22. I/O Module Carrier Installation

The I/O module carrier board is included in the optional accessory kit. It is horizontally installed to the riser slot 2 riser card. The board provides electrical connectivity for installing an Intel® I/O Expansion Module and a SATA based M.2 form factor (NGFF, Next Generation

Form Factor) storage device. It supports up to x8 lanes of PCIe* 3.0 for the I/O module, and a 7 pin SATA header for the M.2 device. The I/O module carrier has two types AXXKPTPM2IOM and AXXKPTPIOM, only AXXKPTPM2IOM can support SATA based M.2. But due to mechanical limitation, the AXXKPTPM2IOM cannot support the computer module with onboard IB* module.

I/O Module Carrier	M.2 Support	Supported Computer Modules
AXXKPTPM2IOM	Yes	<ul style="list-style-type: none"> • HNS2600TPR • HNS2600TP24R • HNS2600TP24SR • HNS2600TP24STR
AXXKPTPIOM	No	<ul style="list-style-type: none"> • HNS2600TPR • HNS2600TPFR • HNS2600TP24R • HNS2600TP24SR • HNS2600TP24STR

The M.2 slot is on the backside of the AXXKPTPM2IOM, it can support M.2 2280 SSD which size is 80.0 mm X 22.0 mm X 3.8 mm. User can install the M.2 device to the M.2 slot (see the letter A on Figure 22) and fix it with the screw (see the letter B on the Figure 22)




Figure 23. Installing the M.2 Device

User still needs to connect the SATA connector (see B on Figure 23) on the AXXKPTPM2IOM to the STAT connector on the server with SATA cable

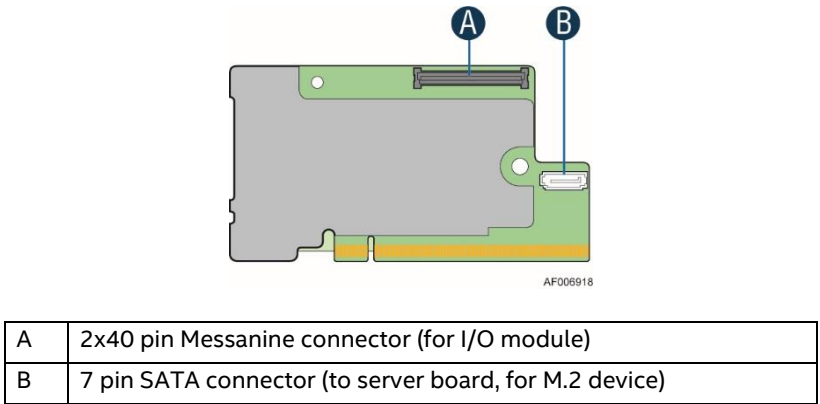


Figure 24. AXXKPTPM2IOM I/O Module Carrier Connectors




Figure 25. Connecting the M.2 SATA Cable

2.11 Compute Module Fans

The cooling subsystem for the compute module consists of three 40 x 40 x 56 dual rotor fans and one air duct. These components provide the necessary cooling and airflow.




Figure 26. Compute Module Fans

Note: The Intel® Compute Module HNS2600TP product family does not support redundant cooling. If one of the compute module fans fails, it is recommended to replace the failed fan as soon as possible.

Each fan within the compute module can support multiple speeds. Fan speed may change automatically when any temperature sensor reading changes. The fan speed control algorithm is programmed into the server board's BMC.

Each fan connector within the module supplies a tachometer signal that allows the BMC to monitor the status of each fan. If one of the fans fails, the status LED on the server board will light up.

The fan control signal is from the BMC on the mother board to the power docking board and then is distributed to three sets of dual rotor fans. The expected maximum RPM is 25,000.

2.12 Air Duct

Each compute module requires the use of a transparent plastic air duct to direct airflow over critical areas within the compute module. To maintain the necessary airflow, the air duct must be properly installed. Before sliding the compute module into the chassis, make sure the air duct is installed properly.




Figure 27. Air Duct

There are two different types of air duct in different compute modules.

Air Duct Type A (Intel part number: H44809-xxx) is only for Intel® Compute Module HNS2600TPR/HNS2600TPFR. Air Duct Type B (Intel part number: H70127-xxx) is only for Intel® Compute Module HNS2600TP24.

Warning: The air duct is pre-installed in compute module. It is required to use correct type of air duct in compute module for proper system cooling.

2.13 Intel® RAID C600 Upgrade Key

The Intel® RAID C600 Upgrade Key RKSATA4R5 is supported. With the optional key installed on the server board, software SATA RAID 5 is enabled.




Figure 28. Intel® RAID C600 Upgrade Key

2.14 Intel® Remote Management Module 4 (Intel® RMM4) Lite

The optional Intel® RMM4 Lite is a small board that unlocks the advanced management features when installed on the server board.




Figure 29. Intel® RMM4 Lite

2.15 Breakout Board

Intel provides a breakout board which is designed for the server board only I/O peripherals in a third-party chassis. It is not a standard accessory of the Intel® Compute Module HNS2600TPR product family or Intel® Server Chassis H2000G product family.

The breakout board provides:

- One 7 pin SATA connector for 6Gb/s SATA DOM
- One mini-SAS HD SFF-8643 connector for 4x lanes of 6Gb/s SATA
- One 7 pin connector for miscellaneous signals:
 - Status LED
 - NMI switch
 - SMBus
 - 3.3V auxiliary power (maximum current 50mA)




Figure 30. Breakout Board Front and Rear View

The breakout board has reserved holes for users to design their own bracket to fix the board into the server system. See the following mechanical drawing for details.




Figure 31. Breakout Board Mechanical Drawing (Unit: mm)

2.16 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® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 and v4 product family*
- *Intel® Server System BMC Firmware External Product Specification for Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 and v4 product family*

2.16.1 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.

2.16.1.1 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 product families and the Intel® C610 chipset family.
- **OEMID** = Three-character OEM BIOS Identifier, to identify the board BIOS “owner”.
 - “**86B**” is used for Intel 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

2.16.1.2 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 5. 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

2.16.1.3 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.

2.16.1.4 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.

2.16.1.5 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.

2.16.1.6 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 given operating system. Full BIOS update instructions are provided when update packages are downloaded from the Intel web site.

2.16.1.7 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 introduces 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* (see Figure 42) is set, 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.

2.16.2 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.*

2.16.2.1 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 given 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.

2.16.3 Baseboard Management Controller (BMC) Firmware

See Platform Management.

3 Processor Support

The server board includes two Socket-R3 (LGA 2011-3) processor sockets and can support one or two of the Intel® Xeon® processor E5-2600 v3/v4 product family, with a Thermal Design Power (TDP) of up to 160W.

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 to the server board.

The following illustration identifies each sub-assembly component.




Figure 32. Processor Socket Assembly



Figure 33. Processor Socket ILM

The narrow ILM has a 56 x 94mm heat sink mounting hole pattern.

Note: *The pins inside the CPU socket are extremely sensitive. Other than the CPU, no object should make contact with the pins inside the CPU socket. A damaged CPU Socket pin may render the socket inoperable, and will produce erroneous CPU or other system errors if used.*

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 and v4 product family TDP guidelines up to and including 160W. The compute module described in this document is designed to support the Intel® Xeon® Processor E5-2600 v3 and v4 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 their 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 the processors that 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 Product Architecture Overview for details.*

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 that 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.

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/v4 product family and Intel® C612 chipset product family 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 Code may be displayed to the screen, will be logged to the BIOS Setup Error Manager, an error event will be logged to the System Event Log (SEL), and the system will continue to boot.

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

Table 6. Mixed Processor Configurations Error Summary

Error	Severity	System Action
Processor family not Identical	Fatal	The BIOS detects the error condition and responds as follows: <ul style="list-style-type: none"> ▪ Halts at POST Code 0xE6. ▪ Halts with 3 long beeps and 1 short beep. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor model not Identical	Fatal	The BIOS detects the error condition and responds as follows: <ul style="list-style-type: none"> ▪ Logs the POST Error Code into the System Event Log (SEL). ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays "0196: Processor model mismatch detected" message in the Error Manager. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor cores/threads not identical	Fatal	The BIOS detects the error condition and responds as follows: <ul style="list-style-type: none"> ▪ Halts at POST Code 0xE5. ▪ Halts with 3 long beeps and 1 short beep. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor cache not identical	Fatal	The BIOS detects the error condition and responds as follows: <ul style="list-style-type: none"> ▪ Halts at POST Code 0xE5. ▪ Halts with 3 long beeps and 1 short beep. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.

Error	Severity	System Action
Processor frequency (speed) not identical	Fatal	<p>The BIOS detects the processor frequency difference, and responds as follows:</p> <ul style="list-style-type: none"> ▪ Adjusts all processor frequencies to the highest common frequency. ▪ No error is generated – this is not an error condition. ▪ Continues to boot the system successfully. <p>If the frequencies for all processors cannot be adjusted to be the same, then this <u>is</u> an error, and the BIOS responds as follows:</p> <ul style="list-style-type: none"> ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Does not disable the processor. ▪ Displays “0197: Processor speeds unable to synchronize” message in the Error Manager. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor Intel® QuickPath Interconnect link frequencies not identical	Fatal	<p>The BIOS detects the QPI link frequencies and responds as follows:</p> <ul style="list-style-type: none"> ▪ Adjusts all QPI interconnect link frequencies to the highest common frequency. ▪ No error is generated – this is not an error condition. ▪ Continues to boot the system successfully. <p>If the link frequencies for all QPI links cannot be adjusted to be the same, then this <u>is</u> an error, and the BIOS responds as follows:</p> <ul style="list-style-type: none"> ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays “0195: Processor Intel(R) QPI link frequencies unable to synchronize” message in the Error Manager. ▪ Does not disable the processor. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor microcode update missing	Minor	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> ▪ Logs the POST Error Code into the SEL. ▪ Displays “818x: Processor 0x microcode update not found” message in the Error Manager or on the screen. ▪ The system continues to boot in a degraded state, regardless of the setting of POST Error Pause in the Setup.
Processor microcode update failed	Major	<p>The BIOS detects the error condition and responds as follows:</p> <ul style="list-style-type: none"> ▪ Logs the POST Error Code into the SEL. ▪ Displays “816x: Processor 0x unable to apply microcode update” message in the Error Manager or on the screen. ▪ Takes Major Error action. The system may continue to boot in a degraded state, depending on the setting of POST Error Pause in Setup, or may halt with the POST Error Code in the Error Manager waiting for operator intervention.

3.5 Processor Function Overview

The Intel® Xeon® processor E5-2600 v3/v4 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 3.0 Express* links capable of 8.0 GT/s, and four 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. For more comprehensive processor specific information, refer to the Intel® Xeon® processor E5-2600 v3/v4 product family documents listed in the Reference Documents list.

3.5.1 Processor Core Features

- Up to 12 execution cores (Intel® Xeon® processor E5-2600 v3/v4 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 SpeedStep® Technology
- Intel® Advanced Vector Extensions 2 (Intel® AVX2)
- Intel® Node Manager 3.0
- Intel® Secure Key
- Intel® OS Guard
- Intel® Quick Data Technology

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 the virtualization performance and robustness. 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

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 a malicious worm attempts to insert code in the buffer, the processor disables code execution, preventing damage and worm 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 firmware. This is in essence 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 PTASCUPS 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. Among other things, the Digital Random Number Generator (DRNG) using the RDRAND instruction is useful for generating high-quality keys for cryptographic protocols.

3.5.2.12 Intel® OS Guard

Protects the 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.

3.6 Processor Heat Sink

Two types of heat sinks are included in the compute module package.

- On CPU 1 – 1U Cu/Al 84mm x 106mm Heat Sink (Rear Heat Sink)
- On CPU 2 – 1U Ex-Al 84mm x 106mm Heat Sink (Front Heat Sink)

Warning: *The two heat sinks are NOT interchangeable.*

This heat sink is designed for optimal cooling and performance. To achieve better cooling performance, you must properly attach the heat sink bottom base with TIM (thermal interface material). ShinEtsu* G-751 or 7783D or Honeywell* PCM45F TIM is recommended. The mechanical performance of the heat sink must satisfy mechanical requirement of the processor.




Figure 34. Processor Heat Sink Overview

Note: *The passive heat sink is Intel standard thermal solution for 1U/2U rack chassis.*

4 Memory Support

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

4.1 Memory Subsystem Architecture




Figure 35. Integrated Memory Controller Functional 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 using x4 or x8 DRAM technology
- DIMMs organized as Single Rank (SR), Dual Rank (DR), or Quad Rank (QR)
- Maximum of 8 ranks per channel
- DIMM sizes of 4 GB, 8 GB, 16 GB, 32 or 64¹ GB depending on ranks and technology
- DIMM speeds of 1600, 1866, 2133 or 2400¹ 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

¹ Intel® Xeon® processor E5-2600 v4 product family only

4.1.1 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 an 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).

4.1.2 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 multi-bit 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.

4.2 Supported DDR4-2400 memory for Intel® Xeon processor v4 Product Family

Table 7. DDR4-2400 DIMM Support Guidelines for Intel® Xeon processor v4 Product Family

Type	Ranks Per DIMM and Data Width	DIMM Capacity (GB)		Speed (MT/s); Voltage (V); Slot Per Channel (SPC) and DIMM Per Channel (DPC)		
				1SPC	2SPC	
				1DPC	1DPC	2DPC
		4Gb	8Gb	1.2V	1.2V	1.2V
RDIMM	SRx8	4GB	8GB	2400	2400	2133
RDIMM	SRx4	8GB	16GB	2400	2400	2133
RDIMM	DRx8	8GB	16GB	2400	2400	2133
RDIMM	DRx4	16GB	32GB	2400	2400	2133
LRDIMM	QRx4	32GB	64GB	2400	2400	2400
LRDIMM 3DS	8Rx4	64GB	128GB	2400	2400	2400

4.3 Supported DDR4-2133 memory for Intel® Xeon processor v4 Product Family

Table 8. DDR4-2133 DIMM Support Guidelines for Intel® Xeon processor v4 Product Family

Type	Ranks Per DIMM and Data Width	DIMM Capacity (GB)		Speed (MT/s); Voltage (V); Slot Per Channel (SPC) and DIMM Per Channel (DPC)		
				1SPC	2SPC	
				1DPC	1DPC	2DPC
		4Gb	8Gb	1.2V	1.2V	1.2V
RDIMM	SRx8	4GB	8GB	2133	2133	1866
RDIMM	SRx4	8GB	16GB	2133	2133	1866
RDIMM	DRx8	8GB	16GB	2133	2133	1866
RDIMM	DRx4	16GB	32GB	2133	2133	1866
LRDIMM	QRx4	32GB	64GB	2133	2133	2133
LRDIMM 3DS	8Rx4	64GB	128GB	2133	2133	2133

4.4 Memory Slot Identification and Population Rules

Note: Although mixed DIMM configurations are supported, Intel only 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 S2600TP product family each memory channel supports two memory slots, for a total possible 16 DIMMs installed.
- 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.
- The silk screened DIMM slot identifiers on the board provide information about the channel, and therefore the processor to which they belong. 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 as long as 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 and Error Management) in the BIOS setup is applied commonly across processor sockets.
- All DIMMs must be DDR4 DIMMs.
- Mixing of LRDIMM with any other DIMM type is not allowed per platform.
- Mixing of DDR4 operating frequencies is not validated within a socket (Intra-socket) or across sockets by Intel. If DIMMs with different frequencies are mixed, all DIMMs run at the common lowest frequency.
- A maximum of eight logical ranks (ranks seen by the host) per channel is allowed.
- 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.

On the Intel® Server Board S2600TPR product family, a total of sixteen DIMM slots are provided (two CPUs – four channels per CPU and two DIMMs per channel). The nomenclature for DIMM sockets is detailed in the following table.

Table 9. DIMM Nomenclature

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	B1	C1	D1	E1	F1	G1	H1
A2	B2	C2	D2	E2	F2	G2	H2




Figure 36. Intel® Server Board S2600TPR Product Family DIMM Slot Layout

The following are the DIMM population requirements.

Table 10. Supported DIMM Populations

Total DIMM#	Processor Socket 1 = Populated								Processor Socket 2 = Populated								Mirror Mode Support
	A1	A2	B1	B2	C1	C2	D1	D2	E1	E2	F1	F2	G1	G2	H1	H2	
1 DIMM	X																No
2 DIMMs	X	X															No
	X		X														Yes
3 DIMMs	X								X								No
	X		X		X												No
	X								X		X						No
4 DIMMs	X	X	X	X													Yes
	X	X	X		X												No
	X		X		X		X										Yes
	X		X						X		X						Yes
	X	X							X	X							No
									X	X	X	X					Yes
5 DIMMs	X	X	X	X	X												No
	X		X		X		X		X								No
6 DIMMs	X	X	X	X	X	X											No
	X		X		X		X		X		X						No
8 DIMMs	X	X	X	X	X	X	X	X									Yes
	X	X	X	X					X	X	X	X					Yes
	X		X		X		X		X		X		X		X		Yes
12 DIMMs	X	X	X	X	X	X	X	X	X	X	X	X					No
16 DIMMs	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	Yes

4.5 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 PCI/PCIe configuration. RAS configurations reduce the memory space available in return for the RAS features. PCI/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.

4.5.1 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 $((2 * 16GB) * 4 \text{ channels}) * 2 \text{ CPU sockets} = 256GB$.

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 $((2 * 16GB) * 4 \text{ channels}) * 2 \text{ CPU sockets} = 256GB$.

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

4.5.2 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)
 - ACPI NVS table
 - Processor microcode
 - Memory Mapped I/O (MMIO)
 - Manageability Engine (ME)
 - BIOS flash
-

4.5.3 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).

4.5.3.1 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 EEPROM 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.

4.5.3.2 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 RDIMM, or 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**.

Note: 3 QR LRDIMMs on a channel is an acceptable configuration if operating in Rank Multiplication mode with RM = 2 or 4. In this case each QR LRDIMM appears to be a DR or SR DIMM.

- 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 efforts 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**.

4.5.3.3 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**.

4.5.3.4 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.

4.5.3.5 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.

4.5.3.6 RAS Mode Initialization

If configured, the DIMM configuration is validated for specified RAS mode. If the enabled DIMM configuration is compliant for the RAS mode selected, then the necessary register settings are done and the RAS mode is started into operation.

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 Server Board I/O

The server board input/output features are provided via the embedded features and functions of several onboard components including: the Integrated I/O Module (IIO) of the Intel® Xeon® processor E5-2600 v3/v4 product family, the Intel® C612 chipset, the Intel® Ethernet controller I350, and the I/O controllers embedded within the Emulex® Pilot-III Management Controller.

See the block diagram for an overview of the features and interconnects of each of the major subsystem components.

5.1 PCI Express* Support

The Integrated I/O (IIO) module of the Intel® Xeon® processor E5-2600 v3/v4 product family provides the PCI express interface for general purpose PCI Express* (PCIe) devices at up to Gen 3 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 an x4 DMI2 lane which can be connected to a PCH, or operate as an x4 PCIe 2.0 port.

5.1.1 PCIe Enumeration and Allocation

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

5.1.2 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/v4 product family 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.

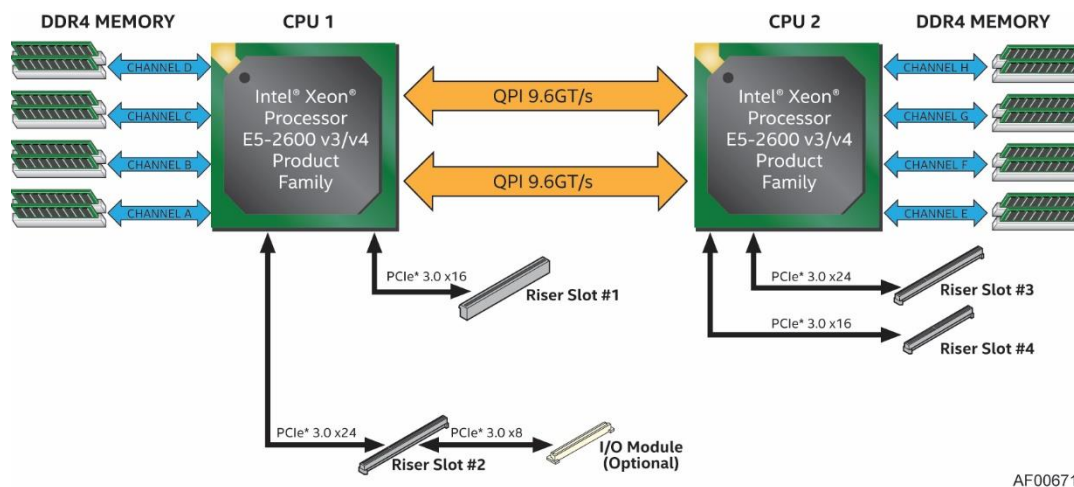
There are 3 NTB supported configurations:

- 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/v4 product family system.

Note: When NTB is enabled in BIOS Setup, Spread Spectrum Clocking (SSC) will be automatically disabled.

5.2 Add-in Card Support

The following sub-sections describe the server board features that are directly supported by the processor I/O module. These include the Riser Card Slots, Network Interface, and connectors for the optional I/O modules and SAS Module. Features and functions of the Intel® C612 Series chipset will be described in its own dedicated section.




AF006719

Figure 37. Add-in Card Support Block Diagram (S2600TPR)

5.2.1 Riser Card Support

The server board includes features for concurrent support of several add-in card types including: PCIe add-in cards via 3 riser card slots (RISER_SLOT_1, RISER_SLOT_3, and RISER_SLOT_4), and Intel® I/O module options via 1 riser card slot (RISER_SLOT_2). The following illustration identifies the location of the onboard connector features and general board placement for add-in modules and riser cards.



AF006446

Figure 38. Server Board Riser Slots (S2600TPFR)

Following is the scope of PCIe connection from processors.

Table 11. PCIe* Port Routing – CPU 1

CPU 1			
PCI Ports	Device (D)	Function (F)	On-board Device
Port DMI 2/PCIe* x4	D0	F0	Chipset
Port 1A - x4	D1	F0	InfiniBand* on S2600TPFR Riser Slot 2 on S2600TPR
Port 1B - x4	D1	F1	InfiniBand* on S2600TPFR Riser Slot 2 on S2600TPR
Port 2A - x4	D2	F0	Riser Slot 1
Port 2B - x4	D2	F1	Riser Slot 1
Port 2C - x4	D2	F2	Riser Slot 1
Port 2D - x4	D2	F3	Riser Slot 1
Port 3A - x4	D3	F0	Riser Slot 2
Port 3B - x4	D3	F1	Riser Slot 2
Port 3C - x4	D3	F2	Riser Slot 2
Port 3D - x4	D3	F3	Riser Slot 2

Table 12. PCIe* Port Routing – CPU 2

CPU 2			
PCI Ports	Device (D)	Function (F)	On-board Device
Port DMI 2/PCIe* x4	D0	F0	Not connect
Port 1A - x4	D1	F0	Riser Slot 3
Port 1B - x4	D1	F1	Riser Slot 3
Port 2A - x4	D2	F0	Riser Slot 4
Port 2B - x4	D2	F1	Riser Slot 4
Port 2C - x4	D2	F2	Riser Slot 4
Port 2D - x4	D2	F3	Riser Slot 4
Port 3A - x4	D3	F0	Riser Slot 3
Port 3B - x4	D3	F1	Riser Slot 3
Port 3C - x4	D3	F2	Riser Slot 3
Port 3D - x4	D3	F3	Riser Slot 3

Notes:

1. All riser slots are defined specially for dedicated risers only. Plugging in a normal PCIe riser or PCIe add-in card directly causes danger and may burn out the add-in riser or card.
-

-
2. *Riser slot 3 and 4 can be used only in dual processor configurations. Any graphic add-in card in riser slot 3 and 4 cannot output video, meaning that the default video out is still from on-board integrated BMC.*
-

5.3 Serial ATA (SATA) Support

The server board utilizes two chipset embedded AHCI SATA controllers, identified as **SATA** and **sSATA** ("s" is for secondary), 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 to the bridge board connector and then to the backplane through the bridge board
- One SATA port to the bridge board connector for the one SATA DOM connector on the bridge board
- One SATA port for the SATA DOM connector on the bridge board

The AHCI **sSATA** controller provides four SATA ports on the server board.




Figure 39. SATA Support

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

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

SATA Controller	sSATA Controller	Supported
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
Command Completion Coalescing	Reduces interrupt and completion overhead by allowing a specified number of commands to complete and then generating an interrupt to process the commands		N/A

5.3.1 Staggered Disk Spin-Up

Because of the high density of disk drives that can be attached to the onboard AHCI SATA controller and the sSATA controller, the combined startup power demand surge for all drives

at once can be much higher than the normal running power requirements and could require a much larger power supply for startup than for normal operations.

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

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

5.4 Embedded SATA RAID Support

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

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

Using the <F2> BIOS Setup Utility, accessed during system POST, options are available to enable/disable 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.

5.4.1 Intel® Rapid Storage Technology (RSTe) 4.0

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

- **RAID Level 0** 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.
- **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 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 specified operating system
- At least two SATA drives needed to support RAID levels 0 or 1
- At least three SATA drives needed to support RAID level 5
- At least four SATA drives needed to support RAID level 10

With Intel® RSTe 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 hitting 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


5.4.2 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 six drives, enabling higher throughput for data intensive applications such as video editing. Supports two SATA DOM devices.
- Data security is offered through **RAID Level 1**, which performs mirroring. Supports two SATA DOM devices.
- **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 four 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 three or more drives. By striping parity, and rotating it across all disks, fault tolerance of any single drive is achieved while only consuming one drive worth of capacity. That is, a 3-drive RAID 5 has the capacity of two drives, or a 4-drive RAID 5 has the capacity of three 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.



AF006450

Figure 40. 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.

5.5 Network Interface

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




Figure 41. Network Interface Connectors

Network interface support is provided from the onboard Intel® i350 NIC, which is a dual-port, compact component with two fully integrated GbE Media Access Control (MAC) and Physical Layer (PHY) ports. The Intel® i350 NIC provides the server board with support for dual LAN ports designed for 10/100/1000 Mbps operation. Refer to the *Intel® i350 Gigabit Ethernet Controller Datasheet* for full details of the NIC feature set.


The NIC device provides a standard IEEE 802.3 Ethernet interface for 1000BASE-T, 100BASE-TX, and 10BASE-T applications (802.3, 802.3u, and 802.3ab) and is capable of transmitting and receiving data at rates of 1000 Mbps, 100 Mbps, or 10 Mbps.

Intel® i350 will be used in conjunction with the Emulex* PILOT III BMC for in band Management traffic. The BMC will communicate with Intel® i350 over an NC-SI interface (RMII physical). The NIC will be on standby power so that the BMC can send management traffic over the NC-SI interface to the network during sleep states S4 and S5.

The NIC supports the normal RJ-45 LINK/Activity speed LEDs as well as the Preset ID function. These LEDs are powered from a Standby voltage rail.

The link/activity LED (at the right of the connector) indicates network connection when on, and transmit/receive activity when blinking. The speed LED (at the left of the connector)

indicates 1000-Mbps operation when green, 100-Mbps operation when amber, and 10-Mbps when off. The following table provides an overview of the LEDs.



LED Color	LED State	NIC State
Green/Amber (B)	Off	10 Mbps
	Amber	100 Mbps
	Green	1000 Mbps
Green (A)	On	Active Connection
	Blinking	Transmit/Receive activity

Figure 42. RJ45 NIC Port LED

5.5.1 MAC Address Definition

The Intel® Server Board S2600TP product family has the following four MAC addresses assigned to it at the Intel factory:

- NIC1 Port1 MAC address (for OS usage)
- NIC1 Port2 MAC address = NIC1 Port1 MAC address + 1 (for OS usage)
- BMC LAN channel 1 MAC address = NIC1 Port1 MAC address + 2
- BMC LAN channel 2 MAC address = NIC1 Port1 MAC address + 3
- BMC LAN channel 3 (Dedicated Server Management NIC) MAC address = NIC1 MAC address + 4

The Intel® Server Board S2600TPR has a white MAC address sticker included with the board. The sticker displays the NIC1 Port1 MAC address in both bar code and alphanumeric formats.

5.5.2 LAN Manageability

Port 1 of the Intel® i350 NIC will be used by the BMC firmware to send management traffic.

5.6 Video Support

There is a video controller which is actually a functional block included in the Baseboard Management Controller integrated on the server board.

There is a PCIe x1 Gen1 link between the BMC video controller and the Integrated IO of the processor in CPU Socket 1, which is the Legacy Processor Socket for boot purposes. During the PCI enumeration and initialization, 16 MB of memory is reserved for video use.

The Onboard Video Controller can support the 2D video resolutions shown in the following table.

Table 15. Onboard Video Resolution and Refresh Rate (Hz)

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

The user can use an add-in PCIe* video adapter to either replace or complement the Onboard Video Controller.

There are enable/disable options in BIOS Setup 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.

5.7 Universal Serial Bus (USB) Ports

There are eight USB 2.0 ports and six USB 3.0 ports available from Intel® C612 chipset. All ports are high-speed, full-speed and low-speed capable. A total of five USB 2.0 dedicated ports are used. The USB port distribution is as follows:

- Emulex* BMC PILOT III consumes two USB 2.0 ports (one USB 1.1 and one USB 2.0).
- Two external USB 2.0 ports on the rear side of server board.
- One internal USB 2.0 port for extension of front-panel USB port on server board.
- One internal USB 2.0 port on bridge board of the compute module.




Figure 43. USB Ports Block Diagram

5.8 Serial Port

The server board has support for one serial port - Serial Port A.




Figure 44. Serial Port A Location

Serial Port A is an internal 10-pin DH-10 connector labeled “Serial_A”.

5.9 InfiniBand* Controller

Intel® Server Board S2600TPFR is populated with a new generation InfiniBand* adapter device. Mellanox* Connect-IB* providing single port 10/20/40/56 Gb/s InfiniBand* interfaces.

Major features and functions include:

- Single InfiniBand* Port: SDR/DDR/QDR/FDR10/FDR with port remapping in firmware.
- Performance optimization: Achieving single port line-rate bandwidth.
- PCI Express 3.0* x8 to achieve 2.5GT/s, 5GT/s, 8GT/s link rate per lane.
- Low power consumption: 7.9 Watt typical.

5.9.1 Device Interfaces

Following is a list of major interfaces of Mellanox* Connect-IB* chip:

- **Clock and Reset signals:** Include core clock input and chip reset signals.

- **Uplink Bus:** The PCI Express* bus is a high-speed uplink interface used to connect Connect-IB* to the host processor. The Connect-IB* supports a PCI Express 3.0* x8 uplink connection with transfer rates of 2.5GT/s, 5GT/s, and 8GT/s per lane. Throughout this document, the PCI Express* interface may also be referred to as the “uplink” interface.
- **Network Interface:** Single network port connecting the device to a network fabric in one of the configurations described in the following table.

Table 16. Network Port Configuration

Port Configured as
10/20/40/56 Gb/s InfiniBand*

- **Flash interface:** Chip initialization and host boot.
- **I²C Compatible Interfaces:** For chip, QSFP+ connector, and chassis configure and monitor.
- **Management Link:** Connect to BMC through SMBus* and NC-SI.
- **Others including:** MDIO, GPIO, and JTAG.

Notes:

The following features are unsupported in the current Connect-IB* firmware

- Service types not supported:
 - SyncUMR
 - Mellanox transport
 - PTP
 - RAW IPv6
 - PTP (ieee 1588)
- Windows Operating System drivers
- SR-IOV
- Connect-IB® currently supports only a single physical function model
- INT-A not supported for EQs only MSI-X
- PCI VPD write flow (RO flow supported)

- Streaming receive queue (STRQ) and collapsed CQ
- Precise clock synchronization over the network (IEEE 1588)
- Data integrity validation of control structures
- NC-SI interface is not enabled in Connect-IB* firmware
- PCIe Function Level Reset (FLR)

5.9.2 Quad Small Form-factor Pluggable (QSFP+) Connector

Port of the Mellanox* Connect-IB* is connected to a single QSFP+ connector on Intel® Server Board S2600TPR (available on SKU: S2600TPFR).

The QSFP+ module and all pins shall withstand 500V electrostatic discharge based on the Human Body Model per JEDEC JESD22-A114-B.

The module shall meet ESD requirements given in *EN61000-4-2, criterion B test specification* such that when installed in a properly grounded cage and chassis the units are subjected to 12KV air discharges during operation and 8KV direct contact discharges to the case.

6 Connector and Header

6.1 Power Connectors

6.1.1 Main Power Connector

To facilitate customers who want to cable to this board from a power supply, the power connector is implemented through two 6pin Minifit Jr* connectors, which can be used to deliver 12amps per pin or 60+Amps total. Note that no over-voltage protective circuits will exist on the board.

Table 17. Main Power Supply Connector 6-pin 2x3 Connector

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

6.1.2 Backup Power Control Connector

The backup (or auxiliary) power connector is used for power control when the baseboard is used in a 3rd party chassis.

It is expected that the customer will use only pin 7 or pin 8 for delivering standby power. The connector is capable of delivering up to 3 amps. Connector type is the AMPMODU MTE Interconnection System or equivalent.

Table 18. Backup Power Control Connector

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

6.2 System Management Headers

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

A 7-pin Intel® RMM4 Lite connector (J1A2) is included on the server board to support the optional Intel® Remote Management Module 4. There is no support for third-party management cards on this server board.

Note: This connector is not compatible with the Intel® Remote Management Module 3 (Intel® RMM3).

Table 19. Intel® RMM4 Lite Connector Pin-out

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

6.2.2 IPMB Header

Table 20. IPMB Header 4-pin

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

6.2.3 Control Panel Connector

This connector is used for front panel control when the baseboard is used in a 3rd party chassis.

This connector is designed to use either discrete 2 pin connectors (similar to ATX chassis) or a ribbon cable. The R470-5VSB represents a resistor to 5VSB to limit the LED current, and the actual value of the resistor will be determined by engineering based on a reasonable brightness at around 8ma. The reset button and the ID button may be mutually exclusive if the KEY pin is required; else the Key can be another GND.

Table 21. Control Panel Connector

Pin	Signal	Pin	Signal
1	FP_ID_LED_N	2	FP_P5V_AUX_0
3	FP_HD ACT_LED_N	4	FP_P5V_AUX_1
5	FP_PWR_LED_N	6	FP_P5V_AUX_2

Pin	Signal	Pin	Signal
7	GND	8	FP_PWR_BTN_N
9	GND	10	FP_ID_BTN_N
11	FP_RST_BTN_N	12	Key

6.3 Bridge Board Connector

The bridge board delivers SATA/SAS signals, disk backplane management signals, BMC SMBus*es as well as control panel and miscellaneous compute module specific signals.

Table 22. Bridge Board Connector

Pin	Signal	Pin	Signal
80	SATA_SAS_SEL	79	GND
78	GND	77	GND
76	SATA6G_P0_RX_DP	75	SATA6G_P0_TX_DN
74	SATA6G_P0_RX_DN	73	SATA6G_P0_TX_DP
72	GND	71	GND
70	SATA6G_P1_TX_DP	69	SATA6G_P1_RX_DN
68	SATA6G_P1_TX_DN	67	SATA6G_P1_RX_DP
66	GND	65	GND
64	SATA6G_P2_RX_DP	63	SATA6G_P2_TX_DN
62	SATA6G_P2_RX_DN	61	SATA6G_P2_TX_DP
60	GND	59	GND
58	SATA6G_P3_TX_DP	57	SATA6G_P3_RX_DN
56	SATA6G_P3_TX_DN	55	SATA6G_P3_RX_DP
54	GND	53	GND
52	SGPIO SATA_CLOCK	51	PWRGD_PSU
50	BMC_NODE_ID1	49	SGPIO_SATA_LOAD
48	BMC_NODE_ID2	47	SGPIO_SATA_DATAOUT0
46	BMC_NODE_ID3	45	SGPIO_SATA_DATAOUT1
KEY			
44	BMC_NODE_ID4	43	PS_EN_PSU_N
42	SPA_SIN_N	41	IRQ_SML1_PMBUS_ALERT_N
40	SPA_SOUT_N	39	GND
38	FP_NMI_BTN_N	37	SMB_PMBUS_CLK
36	FP_PWR_BTN_N	35	SMB_PMBUS_DATA
34	FP_RST_BTN_N	33	GND
32	FP_ID_BTN_N	31	SMB_HSBP_STBY_LVC3_CLK
30	FP_ID_LED_N	29	SMB_HSBP_STBY_LVC3_DATA
28	FP_PWR_LED_N	27	GND
26	FP_LED_STATUS_GREEN_N	25	SMB_CHAS_SENSOR_STBY_LVC3_CLK
24	FP_LED_STATUS_AMBER_N	23	SMB_CHAS_SENSOR_STBY_LVC3_DATA
22	FP_Activity_LED_N	21	GND
20	FP_HDD_ACT_LED_N	19	SMB_IPMB_5VSTBY_CLK

Pin	Signal	Pin	Signal
18	GND	17	SMB_IPMB_5VSTBY_DATA
16	USB2_FP_DN	15	GND
14	USB2_FP_DP	13	Spare
12	GND	11	FM_PS_ALL_NODE_OFF
10	SATA6G_P4_RX_DP	9	FM_NODE_PRESENT_N (GND)
8	SATA6G_P4_RX_DN	7	GND
6	GND	5	SATA6G_P4_TX_DP
4	FM_USB_OC_FP_N	3	SATA6G_P4_TX_DN
2	P5V Aux	1	P5V Aux

Combined system BIOS and the Integrated BMC support provide the functionality of the various supported control panel buttons and LEDs. The following sections describe the supported functionality of each control panel feature.

6.3.1 Power Button

The BIOS supports a front control panel power button. Pressing the power button initiates a request that the Integrated BMC forwards to the ACPI power state machines in the chipset. It is monitored by the Integrated BMC and does not directly control power on the power supply.

▪ Power Button — Off to On

The Integrated BMC monitors the power button and the wake-up event signals from the chipset. A transition from either source results in the Integrated BMC starting the power-up sequence. Since the processors are not executing, the BIOS does not participate in this sequence. The hardware receives the power good and reset signals from the Integrated BMC and then transitions to an ON state.

The System Control Interrupt (SCI) is masked. The BIOS sets up the power button event to generate an SMI and checks the power button status bit in the ACPI hardware registers when an SMI occurs. If the status bit is set, the BIOS sets the ACPI power state of the machine in the chipset to the OFF state. The Integrated BMC monitors power state signals from the chipset and de-asserts PS_PWR_ON to the power supply. As a safety mechanism, if the BIOS fails to service the request, the Integrated BMC automatically powers off the system in four to five seconds.

If an ACPI operating system is running, pressing the power button switch generates a request through SCI to the operating system to shut down the system. The operating system retains control of the system and the operating system policy determines the sleep state into which the system transitions, if any. Otherwise, the BIOS turns off the system.

The tables below list the connector pin definitions on the bridge board.

Table 23. SATA DOM Connector Pin-out

Pin	Signal Description
1	GND

2	SATA0_TXP
3	SATA0_TXN
4	GND
5	SATA0_RXN
6	SATA0_RXP
7	P5V_SATA for SATA DOM

Table 24. USB 2.0 Type-A Connector Pin-out

Pin	Signal Description
1	P5V_USB
2	USB2_P0N
3	USB2_P0P
4	GND

Table 25. 5V_AUX Power Connector Pin-out

Pin	Signal Description
1	GND
2	P5V

6.4 I/O Connectors

6.4.1 PCI Express* Connectors

The Intel® Server Board S2600TPR uses four PCI Express* slots physically with different pin-out definition. Each riser slot has dedicated usage and cannot be used for normal PCIe based add-in card.

- Riser Slot 1: Provide PCIe x16 to Riser. (Using standard 164-pin connector)
- Riser Slot 2: Provide PCIe x24 to Riser to support an x8 Intel® I/O Expansion Module. (Using 200-pin HSEC8 Connector)
- Riser Slot 3: Provide PCIe x24 to Riser. (Using 200-pin HSEC8 Connector)
- Riser Slot 4: Provide PCIe x16 to Riser. (Using 120-pin HSEC8 Connector)

The PCIe bus segment usage from the processors is listed as below.

Table 26. CPU1 and CPU2 PCIe* Bus Connectivity

CPU	Port	IOU	Width	Connection
CPU1	DMI2	IOU2	x4	Reserved for DMI2
CPU1	PE1	IOU2	X8	FDR IB (if depop, combine with riser 2)
CPU1	PE2	IOU0	x16	Riser 1

CPU1	PE3	IOU1	X24	Riser 2 (Depoped IB x8 to IOM on Riser)
CPU2	DMI2	IOU2	x4	Reserved for DMI2
CPU2	PE1	IOU2	x8	Riser 3
CPU2	PE2	IOU0	x16	Riser 4
CPU2	PE3	IOU1	X24	Riser 3 +IOU2

The pin-outs for the slots are shown in the following tables.

Table 27. PCIe* x16 Riser Slot 1 Connector

Pin	Pin Name	Description	Pin	Pin Name	Description
B1	12V	20W 3.3V generated on riser	A1	12V	20W 3.3V generated on riser
B2	12V	66W for GPU	A2	12V	66W for GPU
B3	12V	66W for GPU	A3	12V	66W for GPU
B4	12V	66W for GPU	A4	SMDATA	SMB_3V3STBY_DATA
B5	SMCLK	SMB_3V3STBY_CLK	A5	3.3VAUX	For wake on LAN
B6	3.3VAUX	For wake on LAN	A6	Spare	
B7	GND		A7	Spare	
B8	Spare		A8	Spare	
B9	Spare		A9	Spare	
B10	Spare		A10	Spare	
B11	Spare		A11	Spare	
KEY					
B12	THRTL_N	THROTTLE_RISER1_N	A12	Spare	
B13	Spare		A13	GND	
B14	GND		A14	PERST#	
B15	Spare		A15	WAKE#	
B16	Spare		A16	GND	
B17	GND		A17	REFCLK+	Clock pair 1
B18	PETxP0	Tx Lane 0+	A18	REFCLK-	Clock pair 1
B19	PETxN0	Tx Lane 0-	A19	GND	
B20	GND		A20	PERxP0	Rx Lane 0+
B21	GND		A21	PERxN0	Rx Lane 0-
B22	PETxP1	Tx Lane 1+	A22	GND	
B23	PETxN1	Tx Lane 1-	A23	GND	
B24	GND		A24	PERxP1	Rx Lane 1+
B25	GND		A25	PERxN1	Rx Lane 1-
B26	PETxP2	Tx Lane 2+	A26	GND	
B27	PETxN2	Tx Lane 2-	A27	GND	
B28	GND		A28	PERxP2	Rx Lane 2+
B29	GND		A29	PERxN2	Rx Lane 2-
B30	PETxP3	Tx Lane 3+	A30	GND	

Pin	Pin Name	Description	Pin	Pin Name	Description
B31	PETxN3	Tx Lane 3-	A31	GND	
B32	GND		A32	PERxP3	Rx Lane 3+
B33	GND		A33	PERxN3	Rx Lane 3-
B34	PETxP4	Tx Lane 4+	A34	GND	
B35	PETxN4	Tx Lane 4-	A35	GND	
B36	GND		A36	PERxP4	Rx Lane 4+
B37	GND		A37	PERxN4	Rx Lane 4-
B38	PETxP5	Tx Lane 5+	A38	GND	
B39	PETxN5	Tx Lane 5-	A39	GND	
B40	GND		A40	PERxP5	Rx Lane 5+
B41	GND		A41	PERxN5	Rx Lane 5-
B42	PETxP6	Tx Lane 6+	A42	GND	
B43	PETxN6	Tx Lane 6-	A43	GND	
B44	GND		A44	PERxP6	Rx Lane 6+
B45	GND		A45	PERxN6	Rx Lane 6-
B46	PETxP7	Tx Lane 7+	A46	GND	
B47	PETxN7	Tx Lane 7-	A47	GND	
B48	GND		A48	PERxP7	Rx Lane 7+
B49	GND		A49	PERxN7	Rx Lane 7-
B50	PETxP8	Tx Lane 8+	A50	GND	
B51	PETxN8	Tx Lane 8-	A51	GND	
B52	GND		A52	PERxP8	Rx Lane 8+
B53	GND		A53	PERxN8	Rx Lane 8-
B54	PETxP9	Tx Lane 9+	A54	GND	
B55	PETxN9	Tx Lane 9-	A55	GND	
B56	GND		A56	PERxP9	Rx Lane 9+
B57	GND		A57	PERxN9	Rx Lane 9-
B58	PETxP10	Tx Lane 10+	A58	GND	
B59	PETxN10	Tx Lane 10-	A59	GND	
B60	GND		A60	PERxP10	Rx Lane 10+
B61	GND		A61	PERxN10	Rx Lane 10-
B62	PETxP11	Tx Lane 11+	A62	GND	
B63	PETxN11	Tx Lane 11-	A63	GND	
B64	GND		A64	PERxP11	Rx Lane 11+
B65	GND		A65	PERxN11	Rx Lane 11-
B66	PETxP12	Tx Lane 12+	A66	GND	
B67	PETxN12	Tx Lane 12-	A67	GND	
B68	GND		A68	PERxP12	Rx Lane 12+
B69	GND		A69	PERxN12	Rx Lane 12-
B70	PETxP13	Tx Lane 13+	A70	GND	
B71	PETxN13	Tx Lane 13-	A71	GND	
B72	GND		A72	PERxP13	Rx Lane 13+

Pin	Pin Name	Description	Pin	Pin Name	Description
B73	GND		A73	PERxN13	Rx Lane 13-
B74	PETxP14	Tx Lane 14+	A74	GND	
B75	PETxN14	Tx Lane 14-	A75	GND	
B76	GND		A76	PERxP14	Rx Lane 14+
B77	REFCLK+	Clock pair 2	A77	PERxN14	Rx Lane 14-
B78	REFCLK-	Clock pair 2	A78	GND	
B79	GND		A79	PERxP15	Rx Lane 15+
B80	PETxP15	Tx Lane 15+	A80	PERxN15	Rx Lane 15-
B81	PETxN15	Tx Lane 15-	A81	GND	
B82	GND		A82	Riser_ID0	(See Table 30)

Table 28. PCIe* x24 Riser Slot 2 Connector

Pin	Pin Name	Description	Pin	Pin Name	Description
1	FM_RISER_ID1	(See Table 30)	2	GND	
3	GND		4	PETxP0	Tx Lane 0+
5	PERxP0	Rx Lane 0+	6	PETxN0	Tx Lane 0-
7	PERxN0	Rx Lane 0-	8	GND	
9	GND		10	PETxP1	Tx Lane 1+
11	PERxP1	Rx Lane 1+	12	PETxN1	Tx Lane 1-
13	PERxN1	Rx Lane 1-	14	GND	
15	GND		16	PETxP2	Tx Lane 2+
17	PERxP2	Rx Lane 2+	18	PETxN2	Tx Lane 2-
19	PERxN2	Rx Lane 2-	20	GND	
21	GND		22	PETxP3	Tx Lane 3+
23	PERxP3	Rx Lane 3+	24	PETxN3	Tx Lane 3-
25	PERxN3	Rx Lane 3-	26	GND	
27	GND		28	PETxP4	Tx Lane 4+
29	PERxP4	Rx Lane 4+	30	PETxN4	Tx Lane 4-
31	PERxN4	Rx Lane 4-	32	GND	
33	GND		34	PETxP5	Tx Lane 5+
35	PERxP5	Rx Lane 5+	36	PETxN5	Tx Lane 5-
37	PERxN5	Rx Lane 5-	38	GND	
39	GND		40	PETxP6	Tx Lane 6+
41	PERxP6	Rx Lane 6+	42	PETxN6	Tx Lane 6-
43	PERxN6	Rx Lane 6-	44	GND	
45	GND		46	PETxP7	Tx Lane 7+
47	PERxP7	Rx Lane 7+	48	PETxN7	Tx Lane 7-
49	PERxN7	Rx Lane 7-	50	GND	
51	GND		52	PETxP7 (IB)	Tx Lane 7+ (IB)
53	PERxP7 (IB)	Rx Lane 7+ (IB)	54	PETxN7 (IB)	Tx Lane 7- (IB)

Pin	Pin Name	Description	Pin	Pin Name	Description
55	PERxN7 (IB)	Rx Lane 7- (IB)	56	GND	
57	GND		58	PETxP6 (IB)	Tx Lane 6+ (IB)
59	PERxP6 (IB)	Rx Lane 6+ (IB)	60	PETxN6 (IB)	Tx Lane 6- (IB)
61	PERxN6 (IB)	Rx Lane 6- (IB)	62	GND	
63	GND		64	GND	
KEY					
65	GND		66	GND	
67	GND		68	PETxP5 (IB)	Tx Lane 5+ (IB)
69	PERxP5 (IB)	Rx Lane 5+ (IB)	70	PETxN5 (IB)	Tx Lane 5- (IB)
71	PERxN5 (IB)	Rx Lane 5- (IB)	72	GND	
73	GND		74	PETxP4 (IB)	Tx Lane 4+ (IB)
75	PERxP4 (IB)	Rx Lane 4+ (IB)	76	PETxN4 (IB)	Tx Lane 4- (IB)
77	PERxN4 (IB)	Rx Lane 4- (IB)	78	GND	
79	GND		80	PETxP3 (IB)	Tx Lane 3+ (IB)
81	PERxP3 (IB)	Rx Lane 3+ (IB)	82	PETxN3 (IB)	Tx Lane 3- (IB)
83	PERxN3 (IB)	Rx Lane 3- (IB)	84	GND	
85	GND		86	PETxP2 (IB)	Tx Lane 2+ (IB)
87	PERxP2 (IB)	Rx Lane 2+ (IB)	88	PETxN2 (IB)	Tx Lane 2- (IB)
89	PERxN2 (IB)	Rx Lane 2- (IB)	90	GND	
91	GND		92	PETxP1 (IB)	Tx Lane 1+ (IB)
93	PERxP1 (IB)	Rx Lane 1+ (IB)	94	PETxN1 (IB)	Tx Lane 1- (IB)
95	PERxN1 (IB)	Rx Lane 1- (IB)	96	GND	
97	GND		98	PETxP0 (IB)	Tx Lane 0+ (IB)
99	PERxP0 (IB)	Rx Lane 0+ (IB)	100	PETxN0 (IB)	Tx Lane 0- (IB)
101	PERxN0 (IB)	Rx Lane 0- (IB)	102	GND	
103	GND		104	PETxP8	Tx Lane 8+
105	PERxP8	Rx Lane 8+	106	PETxN8	Tx Lane 8-
107	PERxN8	Rx Lane 8-	108	GND	
109	GND		110	PETxP9	Tx Lane 9+
111	PERxP9	Rx Lane 9+	112	PETxN9	Tx Lane 9-
113	PERxN9	Rx Lane 9-	114	GND	
115	GND		116	PETxP10	Tx Lane 10+
117	PERxP10	Rx Lane 10+	118	PETxN10	Tx Lane 10-
119	PERxN10	Rx Lane 10-	120	GND	
121	GND		122	PETxP11	Tx Lane 11+
123	PERxP11	Rx Lane 11+	124	PETxN11	Tx Lane 11-
125	PERxN11	Rx Lane 11-	126	GND	
127	GND		128	PETxP12	Tx Lane 12+
129	PERxP12	Rx Lane 12+	130	PETxN12	Tx Lane 12-
131	PERxN12	Rx Lane 12-	132	GND	
133	GND		134	PETxP13	Tx Lane 13+
135	PERxP13	Rx Lane 13+	136	PETxN13	Tx Lane 13-

Pin	Pin Name	Description	Pin	Pin Name	Description
137	PERxN13	Rx Lane 13-	138	GND	
139	GND		140	PETxP14	Tx Lane 14+
141	PERxP14	Rx Lane 14+	142	PETxN14	Tx Lane 14-
143	PERxN14	Rx Lane 14-	144	GND	
145	GND		146	PETxP15	Tx Lane 15+
147	PERxP15	Rx Lane 15+	148	PETxN15	Tx Lane 15-
149	PERxN15	Rx Lane 15-	150	GND	
151	GND		152	REFCLK1-	Clock pair 1
153	REFCLK0-	Clock pair 0	154	REFCLK1+	Clock pair 1
155	REFCLK0+	Clock pair 0	156	GND	
157	GND		158	PCIE_WAKE#	
159	REFCLK2-	Clock pair 2	160	PCIE_RST#	
161	REFCLK2+	Clock pair 2	162	GND	
163	GND		164	GPU_NODE_ON	
165	SMB_DATA_HOST	SMB_HOST_3V3_DATA	166	RSVD	
167	SMB_CLK_HOST	SMB_HOST_3V3_CLK	168	RSVD	
169	GND		170	THROTTLE_N	THROTTLE_RISER2_N
171	FAN_BMC_TACH11	FAN_BMC_TACH11_R	172	GND	
173	FAN_BMC_TACH10	FAN_BMC_TACH10_R	174	BMC_MIC_MUX_RST_N	for SMB mux
175	FAN_BMC_TACH9	FAN_BMC_TACH9_R	176	RSVD	
177	FAN_BMC_TACH8	FAN_BMC_TACH8_R	178	RSVD	
179	FAN_BMC_TACH7	FAN_BMC_TACH7_R	180	FAN_PWM1_OUT	
181	FAN_BMC_TACH6	FAN_BMC_TACH6_R	182	GND	
183	LED_RIOM_ACT_N		184	P3V3_AUX	
185	IOM_PRESENT_N	IOM_PRESENT_N	186	SMB_CLK	
187	P5V_AUX		188	GND	
189	SMB_DATA	PCI_SLOT2_DATA	190	RSVD	
191	GND		192	RSVD	
193	PWRGD_GPU		194	GND	
195	GND		196	P12V	
197	P12V		198	P12V	
199	P12V		200	P12V	

Table 29. PCIe* x24 Riser Slot 3 Connector

Pin	Pin Name	Description	Pin	Pin Name	Description
1	RISER_ID2	(See Table 30)	2	GND	
3	GND		4	PETxP7	Tx Lane 7+
5	PERxP7	Rx Lane 7+	6	PETxN7	Tx Lane 7-
7	PERxN7	Rx Lane 7-	8	GND	
9	GND		10	PETxP6	Tx Lane 6+

Pin	Pin Name	Description	Pin	Pin Name	Description
11	PERxP6	Rx Lane 6+	12	PETxN6	Tx Lane 6-
13	PERxN6	Rx Lane 6-	14	GND	
15	GND		16	PETxP5	Tx Lane 5+
17	PERxP5	Rx Lane 5+	18	PETxN5	Tx Lane 5-
19	PERxN5	Rx Lane 5-	20	GND	
21	GND		22	PETxP4	Tx Lane 4+
23	PERxP4	Rx Lane 4+	24	PETxN4	Tx Lane 4-
25	PERxN4	Rx Lane 4-	26	GND	
27	GND		28	PETxP3	Tx Lane 3+
29	PERxP3	Rx Lane 3+	30	PETxN3	Tx Lane 3-
31	PERxN3	Rx Lane 3-	32	GND	
33	GND		34	PETxP2	Tx Lane 2+
35	PERxP2	Rx Lane 2+	36	PETxN2	Tx Lane 2-
37	PERxN2	Rx Lane 2-	38	GND	
39	GND		40	PETxP1	Tx Lane 1+
41	PERxP1	Rx Lane 1+	42	PETxN1	Tx Lane 1-
43	PERxN1	Rx Lane 1-	44	GND	
45	GND		46	PETxP0	Tx Lane 0+
47	PERxP0	Rx Lane 0+	48	PETxN0	Tx Lane 0-
49	PERxN0	Rx Lane 0-	50	GND	
51	GND		52	PETxP15	Tx Lane 15+
53	PERxP15	Rx Lane 15+	54	PETxN15	Tx Lane 15-
55	PERxN15	Rx Lane 15-	56	GND	
57	GND		58	PETxP14	Tx Lane 14+
59	PERxP14	Rx Lane 14+	60	PETxN14	Tx Lane 14-
61	PERxN14	Rx Lane 14-	62	GND	
63	GND		64	RSVD	
KEY					
65	GND		66	GND	
67	PERxP13	Rx Lane 13+	68	PETxP13	Tx Lane 13+
69	PERxN13	Rx Lane 13-	70	PETxN13	Tx Lane 13-
71	GND		72	GND	
73	GND		74	PETxP12	Tx Lane 12+
75	PERxP12	Rx Lane 12+	76	PETxN12	Tx Lane 12-
77	PERxN12	Rx Lane 12-	78	GND	
79	GND		80	PETxP11	Tx Lane 11+
81	PERxP11	Rx Lane 11+	82	PETxN11	Tx Lane 11-
83	PERxN11	Rx Lane 11-	84	GND	
85	GND		86	PETxP10	Tx Lane 10+
87	PERxP10	Rx Lane 10+	88	PETxN10	Tx Lane 10-
89	PERxN10	Rx Lane 10-	90	GND	
91	GND		92	PETxP9	Tx Lane 9+
93	PERxP9	Rx Lane 9+	94	PETxN9	Tx Lane 9-
95	PERxN9	Rx Lane 9-	96	GND	
97	GND		98	PETxP8	Tx Lane 8+
99	PERxP8	Rx Lane 8+	100	PETxN8	Tx Lane 8-
101	PERxN8	Rx Lane 8-	102	GND	

Pin	Pin Name	Description	Pin	Pin Name	Description
103	GND		104	PETxP7	Tx Lane 7+
105	PERxP7	Rx Lane 7+	106	PETxN7	Tx Lane 7-
107	PERxN7	Rx Lane 7-	108	GND	
109	GND		110	PETxP6	Tx Lane 6+
111	PERxP6	Rx Lane 6+	112	PETxN6	Tx Lane 6-
113	PERxN6	Rx Lane 6-	114	GND	
115	GND		116	PETxP5	Tx Lane 5+
117	PERxP5	Rx Lane 5+	118	PETxN5	Tx Lane 5-
119	PERxN5	Rx Lane 5-	120	GND	
121	GND		122	PETxP4	Tx Lane 4+
123	PERxP4	Rx Lane 4+	124	PETxN4	Tx Lane 4-
125	PERxN4	Rx Lane 4-	126	GND	
127	GND		128	PETxP3	Tx Lane 3+
129	PERxP3	Rx Lane 3+	130	PETxN3	Tx Lane 3-
131	PERxN3	Rx Lane 3-	132	GND	
133	GND		134	PETxP2	Tx Lane 2+
135	PERxP2	Rx Lane 2+	136	PETxN2	Tx Lane 2-
137	PERxN2	Rx Lane 2-	138	GND	
139	GND		140	PETxP1	Tx Lane 1+
141	PERxP1	Rx Lane 1+	142	PETxN1	Tx Lane 1-
143	PERxN1	Rx Lane 1-	144	GND	
145	GND		146	PETxP0	Tx Lane 0+
147	PERxP0	Rx Lane 0+	148	PETxN0	Tx Lane 0-
149	PERxN0	Rx Lane 0-	150	GND	
151	GND		152	REFCLK1-	Clock pair 1
153	REFCLK0-	Clock pair 0	154	REFCLK1+	Clock pair 1
155	REFCLK0+	Clock pair 0	156	GND	
157	GND		158	PCIE_WAKE#	
159	REFCLK2-	Clock pair 2	160	PCIE_RST#	
161	REFCLK2+	Clock pair 2	162	GND	
163	GND		164	RSVD	
165	RSVD		166	RSVD	
167	RSVD		168	RSVD	
169	GND		170	THROTTLE_N	THROTTLE_RISER3_N
171	RSVD		172	GND	
173	GND		174	RSVD	
175	RSVD		176	RSVD	
177	RSVD		178	RSVD	
179	RSVD		180	RSVD	
181	RSVD		182	GND	
183	RSVD		184	P3V3_AUX	
185	SAS_PRESENT_N	SAS_PRESENT_N	186	SMB_CLK	SMB_PCI_SLOT3_CLK
187	RSVD		188	GND	
189	SMB_DATA	SMB_PCI_SLOT3_DAT A	190	RSVD	
191	GND		192	RSVD	
193	RSVD		194	GND	

Pin	Pin Name	Description	Pin	Pin Name	Description
195	GND		196	P12V	
197	P12V		198	P12V	
199	P12V		200	P12V	

Table 30. PCIe* x16 Riser Slot 4 Connector

Pin	PCIE Riser	Description	Pin	PCIE Riser	Description
1	FM_RISER_ID3	for PCIe configuration	2	GND	
3	GND		4	PETxP15	Tx Lane 15+
5	PERxP15	Rx Lane 15+	6	PETxN15	Tx Lane 15-
7	PERxN15	Rx Lane 15-	8	GND	
9	GND		10	PETxP14	Tx Lane 14+
11	PERxP14	Rx Lane 14+	12	PETxN14	Tx Lane 14-
13	PERxN14	Rx Lane 14-	14	GND	
15	GND		16	PETxP13	Tx Lane 13+
17	PERxP13	Rx Lane 13+	18	PETxN13	Tx Lane 13-
19	PERxN13	Rx Lane 13-	20	GND	
21	GND		22	PETxP12	Tx Lane 12+
23	PERxP12	Rx Lane 12+	24	PETxN12	Tx Lane 12-
25	PERxN12	Rx Lane 12-	26	GND	
27	GND		28	PETxP11	Tx Lane 11+
29	PERxP11	Rx Lane 11+	30	PETxN11	Tx Lane 11-
31	PERxN11	Rx Lane 11-	32	GND	
33	GND		34	PETxP10	Tx Lane 10+
35	PERxP10	Rx Lane 10+	36	PETxN10	Tx Lane 10-
37	PERxN10	Rx Lane 10-	38	GND	
39	GND		40	PETxP9	Tx Lane 9+
41	PERxP9	Rx Lane 9+	42	PETxN9	Tx Lane 9-
43	PERxN9	Rx Lane 9-	44	GND	
45	GND		46	PETxP8	Tx Lane 8+
47	PERxP8	Rx Lane 8+	48	PETxN8	Tx Lane 8-
49	PERxN8	Rx Lane 8-	50	GND	
51	GND		52	PETxP7	Tx Lane 7+
53	PERxP7	Rx Lane 7+	54	PETxN7	Tx Lane 7-
55	PERxN7	Rx Lane 7-	56	GND	
57	GND		58	PETxP6	Tx Lane 6+
59	PERxP6	Rx Lane 6+	60	PETxN6	Tx Lane 6-
61	PERxN6	Rx Lane 6-	62	GND	
63	GND		64	RSVD	

Pin	PCIE Riser	Description	Pin	PCIE Riser	Description
	Key			Key	
65	RSVD		66	GND	
67	GND		68	PETxP5	Tx Lane 5+
69	PERxP5	Rx Lane 5+	70	PETxN5	Tx Lane 5-
71	PERxN5	Rx Lane 5-	72	GND	
73	GND		74	PETxP4	Tx Lane 4+
75	PERxP4	Rx Lane 4+	76	PETxN4	Tx Lane 4-
77	PERxN4	Rx Lane 4-	78	GND	
79	GND		80	PETxP3	Tx Lane 3+
81	PERxP3	Rx Lane 3+	82	PETxN3	Tx Lane 3-
83	PERxN3	Rx Lane 3-	84	GND	
85	GND		86	PETxP2	Tx Lane 2+
87	PERxP2	Rx Lane 2+	88	PETxN2	Tx Lane 2-
89	PERxN2	Rx Lane 2-	90	GND	
91	GND		92	PETxP1	Tx Lane 1+
93	PERxP1	Rx Lane 1+	94	PETxN1	Tx Lane 1-
95	PERxN1	Rx Lane 1-	96	GND	
97	GND		98	PETxP0	Tx Lane 0+
99	PERxP0	Rx Lane 0+	100	PETxN0	Tx Lane 0-
101	PERxN0	Rx Lane 0-	102	GND	
103	GND		104	FM_THROTTLE_N	
105	REFCLK0-	Clock pair 0	106	RSVD	
107	REFCLK0+	Clock pair 0	108	PCIE_WAKE#	
109	GND		110	PCIE_RST#	
111	REFCLK1-	Clock pair 1	112	GND	
113	REFCLK1+	Clock pair 1	114	P3V3_AUX	
115	GND		116	P12V	
117	SMB_DATA		118	P12V	
119	SMB_CLK		120	P12V	

Table 31. PCIe* Riser ID Assignment

Description	CPU1		CPU2	
	Riser ID (0)	Riser ID (1)	Riser ID (2)	Riser ID (3)
Riser 1 1x16	1			
Riser 1 2x8	0			
Riser 2 1x16 + 1x8 (S2600TPR only)		1		
Riser 2 2x8 + 1x8 (S2600TPR only)		0		
Riser 3 1x16 + 1x8			1	
Riser 3 2x8 + 1x8			0	
Riser4 1x16				1

Description	CPU1		CPU2	
	Riser ID (0)	Riser ID (1)	Riser ID (2)	Riser ID (3)
Riser4 2x8				0

Note: PCIe bifurcation setting in the BIOS will override riser slot pin setting.

Table 32. PCIe* Clock Source by Slot

PCI Express Clocks Source according Riser's lane				
PCIE SLOT RISER	x16 PCIE PIN	x8 (lane 0~7) PCIE PIN	x8 (lane 8~15) PCIE PIN	x8 (lane 0~7) PCIE PIN ³
SLOT 1	PIN A17/A18	PIN B77/B78	PIN A17/A18	NA
SLOT 2 ¹	PIN 153/155	PIN 152/154	PIN 153/155	PIN 159/161
SLOT 3 ²	PIN 153/155	PIN 153/155	PIN 152/154	PIN 159/161
SLOT 4	PIN 105/107	PIN 111/113	PIN 105/107	NA

Notes:

1. PCIe* Slot #2 riser design using x8 (Lanes 0~7) on CPU 1A-1B is only for S2600TP board
2. PCIe* Slot #3 has lane reversal on the CPU side
3. Part of x16 PCIE (Can configure as 2x8 PCIE)

Please review the S2600TP MSU January 2017 publication for more details.

6.4.2 VGA Connector

The following table details the pin-out definition of the external VGA connector.

Table 33. VGA External Video Connector

Pin	Signal Name	Description
1	V_IO_R_CONN	Red (analog color signal R)
2	V_IO_G_CONN	Green (analog color signal G)
3	V_IO_B_CONN	Blue (analog color signal B)
4	NC	No connection
5	GND	Ground
6	GND	Ground
7	GND	Ground
8	GND	Ground
9	P5V	No fuse protection
10	GND	Ground
11	NC	No connection
12	V_IO_DDCDAT	DDCDAT
13	V_IO_HSYNC_CONN	HSYNC (horizontal sync)
14	V_IO_VSYNC_CONN	VSYNC (vertical sync)
15	V_IO_DDCLK	DDCLK

6.4.3 NIC Connectors

The server board provides three independent RJ-45 connectors on the back edge of the board (JA6A1, JA5A1, and JA2A1). The pin-outs for NIC connectors are identical and are defined in the following table.

Table 34. RJ-45 10/100/1000 NIC Connector Pin-out

Pin	Signal Name
1	LED_NIC_LINK0_100_N
2	LED_NIC_LINK0_1000_R_N
3	NIC_0_0_DP
4	NIC_0_0_DN
5	NIC_0_1_DP
6	NIC_0_1_DN
7	NIC_CT1
8	NIC_CT2
9	NIC_0_2_DP
10	NIC_0_2_DN
11	NIC_0_3_DP
12	NIC_0_3_DN
13	LED_NIC_LINK0_LNKUP_N
14	LED_NIC_LINK0_ACT_R_N

6.4.4 SATA Connectors

The server board provides five SATA port connectors from SATA_0 to SATA_3 ports and SATA DOM support on board. Additional four SATA ports are provided through the bridge board.

The pin configuration for each connector is identical and defined in the following table.

Table 35. SATA Connector

Pin	Signal Name	Description
1	GND	Ground
2	SATA_TX_P	Positive side of transmit differential pair
3	SATA_TX_N	Negative side of transmit differential pair
4	GND	Ground
5	SATA_RX_N	Negative side of receive differential pair
6	SATA_RX_P	Positive side of receive differential pair
7	P5V_SATA/GND	+5V for DOM (J7C3) or Ground for SATA signals

Note: SATA DOM requiring external power cannot be used with SATA DOM port on board.

6.4.5 SATA SGPIO Connectors

The server board provides one SATA SGPIO connector from PCH.

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

Table 36. SATA SGPIO Connector

Pin	Description
1	SGPIO_SSATA_CLOCK
2	SGPIO_SSATA_LOAD
3	GND
4	SGPIO_SSATA_DATAOUT0
5	SGPIO_SSATA_DATAOUT1

6.4.6 Hard Drive Activity (Input) LED Header

Table 37. SATA HDD Activity (Input) LED Header

Pin	Description
1	LED_HD_ACTIVE_N
2	NC

6.4.7 Intel® RAID C600 Upgrade Key Connector

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

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

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

Table 38. Storage Upgrade Key Connector

Pin	Signal Description
1	GND
2	PU_KEY
3	GND
4	PCH_SATA_RAID_KEY

6.4.8 Serial Port Connectors

The server board provides one internal 9-pin serial type A header. The following tables define the pin-outs.

Table 39. Internal 9-pin Serial A

Pin	Signal Name	Pin	Signal Name
1	SPA_DCD	2	SPA_DSR
3	SPA_SIN_N	4	SPA_RTS
5	SPA_SOUT_N	6	SPA_CTS

Pin	Signal Name	Pin	Signal Name
7	SPA_DTR	8	SPA_RI
9	GND	10	KEY

6.4.9 USB Connectors

The following table details the pin-out of the external stack USB port connectors found on the back edge of the server board.

Table 40. External USB port Connector

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

One 2x5 connector on the server board provides an option to support two additional internal USB ports. The pin-out is detailed in the following table.

Table 41. Internal USB Connector

Pin	Signal Name	Pin	Signal Name
1	+5V	2	+5V
3	USB_N	4	USB_N
5	USB_P	6	USB_P
7	GND	8	GND
9	Key	10	NC

6.4.10 3QSFP+ for InfiniBand*

The following table details the pin-out of the QSFP+ connector found on the back edge of the server board. This port is available only on board SKU S2600TPFR.

Table 42. QSFP+ Connector

Pin	Signal	Pin	Signal
19	GND	38	GND
18	IB_RX0_DN0	37	IB_RX0_DN1
17	IB_RX0_DP0	36	IB_RX0_DP1
16	GND	35	GND
15	IB_RX0_DN2	34	IB_RX0_DN3
14	IB_RX0_DP2	33	IB_RX0_DP3
13	GND	32	GND

Pin	Signal	Pin	Signal
12	SMB_IB_QSFP0_DATA	31	IB_MODPRSL_PORT1_N
11	SMB_IB_QSFP0_CLK	30	IB_INT_PORT1
10	P3V3_RX_PORT0	29	P3V3_TX_PORT0
9	IB_RESET_PORT1	28	P3V3_PORT0
8	IB_MODSEIL_PORT1_N	27	PD_QSFP0_LPMODE
7	GND	26	GND
6	IB_TX0_DP3	25	IB_TX0_DP2
5	IB_TX0_DN3	24	IB_TX0_DN2
4	GND	23	GND
3	IB_TX0_DP1	22	IB_TX0_DP0
2	IB_TX0_DN1	21	IB_TX0_DN0
1	GND	20	GND

6.4.11 UART Header

The following table details the pin-out of the UART header. This header is only available on 12G SAS bridge board with RAID 5.

Table 43. UART Header

Pin	Signal
1	SP0_SAS_TX_P1V8
2	GND
3	SP0_SAS_RX_P1V8
4	P1V8

6.5 Fan Headers

6.5.1 FAN Control Cable Connector

To facilitate the connection of 3 x40mm double rotor fans, a 14 pin header is provided, and all fans will share a PWM. Both rotor Tachs can be monitored.

Table 44. Baseboard Fan Connector

Pin	Signal Name	Pin	Signal Name
1	FAN_PWM_OUT	2	Key
3	FAN_BMC_TACH0	4	FAN_BMC_TACH1
5	FAN_BMC_TACH2	6	FAN_BMC_TACH3
7	FAN_BMC_TACH4	8	FAN_BMC_TACH5
9	PS_HOTSWAP_EN	10	GND
11	SMB_HOST_3V3_CLK	12	SMB_HOST_3V3_DATA
13	NODE_ADR0(GND)	14	PWRGD_PS_PWROK

The SMBus* is used to connect to the hot swap controller that provides inrush current protection and can measure the power being used by the compute module. The NODE_ON signal is used to turn on the hot swap controller. Note that the polarity is correct as the ADI1276 controller uses a high true enable signal. When the compute module is turned off, the fans will continue to rotate at a preset rate; this rate is selected by Intel and preset by the Fan manufacturer. This is done to stop air recirculation between compute modules. When docking the board to a live 12V rail, the fans could spin up immediately; they may be required to phase their connection to power to minimize the inrush current. Bench testing of the fans should determine if this is necessary.

6.5.2 Discrete System FAN Connector

To support the 3rd party chassis, three discrete system FAN connectors are provided on baseboard. They are used for connecting FANs with tach meters directly.

Table 45. Baseboard Fan Connector

J1K1		J5K2		J1K2	
Pin	Signal Description	Pin	Signal Description	Pin	Signal Description
1	GND	1	GND	1	GND
2	P12V	2	P12V	2	P12V
3	BMC_TACH0_R	3	BMC_TACH2_R	3	BMC_TACH4_R
4	PWM0	4	PWM0	4	PWM0
5	GND	5	GND	5	GND
6	P12V	6	P12V	6	P12V
7	BMC_TACH1_R	7	BMC_TACH3_R	7	BMC_TACH5_R
8	PWM0	8	PWM0	8	PWM0

6.6 Power Docking Board Connectors

The table below lists the connector type and pin definition on the power docking board.

Table 46. Main Power Input Connector

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

Table 47. Fan Control Signal Connector

Pin	Signal Description	Pin	Signal Description
1	PWM1	2	Reserved
3	Tach0	4	Tach1
5	Tach2	6	Tach3
7	Tach4	8	Tach5
9	NODE_ON	10	GND
11	SMBUS_R4 CLK	12	SMBUS_R4 DAT
13	NODE_ADR0	14	NODE_PWRGD

Table 48. Compute Module Fan Connector

Pin	Signal Description
1	GND
2	P12V
3	TACH0
4	PWM1
5	GND
6	P12V
7	TACH1
8	PWM1

Table 49. Main Power Output Connector

Pin	Signal Description	Pin	Signal Description
1	GND	7	P12V_HS
2	GND	8	P12V_HS
3	GND	9	P12V_HS
4	GND	10	P12V_HS
5	GND	11	P12V_HS
6	GND	12	P12V_HS

Table 50. 40 pin Misc. Signal Connector (HNS2600TP24R/HNS2600TP24SR only)

Pin	Signal Description	Pin	Signal Description
1	PE_SMB_CLK	2	PE_SMB_DATA
3	GND	4	FM_ALL_NODE_OFF
5	RST_RTMR_1_PRTRST_N_R	6	GND
7	FP Activity LED_N	8	SMB_IPMB_5VSTBY_BP_DATA
9	FP Health LEDA_N	10	SMB_IPMB_5VSTBY_BP_CLK
11	FP Health LEDG_N	12	GND

Pin	Signal Description	Pin	Signal Description
13	FP_PWR_LED_BUF_R_N	14	SMB_SENSOR_3V3STBY_BP_DATA
15	FP_ID_LED_BUF_R_N	16	SMB_SENSOR_3V3STBY_BP_CLK
17	FP_ID_BTN_R_N	18	GND
19	FP_RST_BTN_R_N	20	SMB_HSBP_3V3_BP_DATA
21	FP_PWR_BTN_R_N	22	SMB_HSBP_3V3_BP_CLK
23	FM_IBMC_NODEID_1	24	GND
25	FM_IBMC_NODEID_0	26	SMB_PMBUS_BP_DATA
27	GND	28	SMB_PMBUS_BP_CLK
29	SGPIO_SAS12G_1_CLOCK_R1	30	GND
31	GND	32	IRQ_SML1_PMBUS_ALERT_N
33	SGPIO_SAS12G_0_CLK	34	FM_NODE_ON_N
35	SGPIO_SAS12G_0_LD	36	SGPIO_SAS12G_1_DATAIN1_R1
37	SGPIO_SAS12G_0_Data_Out	38	SGPIO_SAS12G_1_DATAOUT0_R1
39	PWROK	40	SGPIO_SAS12G_1_LOAD_R1

7 Configuration Jumpers

The following table provides a summary and description of configuration, test, and debug jumpers. The server board has several 3-pin jumper blocks that can be used. Pin 1 on each jumper block can be identified by the following symbol on the silkscreen: ▼




Figure 45. Jumper Location

Table 51. Jumper Modes Selection

Jumper Name	Jumper Position	Mode of Operation	Note
J7A2: BMC Force Update jumper	1-2	Normal	Normal mode
	2-3	Update	BMC in force update mode
J7A3: BIOS Default	1-2	Normal	Normal mode
	2-3	Clear BIOS Settings	BIOS settings are reset to factory default
J7A6: Password Clear	1-2	Normal	Normal mode, password in protection
	2-3	Clear Password	BIOS password is cleared
J7A7: BIOS Recovery Mode	1-2	Normal	Normal mode
	2-3	Recovery	BIOS in recovery mode
J5D2: ME Force Update	1-2	Normal	Normal mode
	2-3	Update	ME in force update mode

7.1 BMC Force Update (J7A2)

When performing a standard BMC firmware update procedure, the update utility places the BMC into an update mode, allowing the firmware to load safely onto the flash device. In the unlikely event the BMC firmware update process fails due to the BMC not being in the proper update state, the server board provides a BMC Force Update jumper (J7A2) which will force the BMC into the proper update state. The following procedure should be followed in the event the standard BMC firmware update process fails.

Table 52. Force Integrated BMC Update Jumper (J7A2)

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

Steps to perform Force BMC Update:

1. Plug out the compute module.
2. Remove the air duct. See the *Service Guide* for instructions.
3. Move the jumper (J7A2) from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Restore the air duct to the compute module.
5. Insert the compute module back to the chassis.
6. Power on the compute module by pressing the power button on the front panel.
7. Perform the BMC firmware update procedure as documented in the *Release Notes* included in the given BMC firmware update package. After successful completion of the firmware update process, the firmware update utility may generate an error stating the BMC is still in update mode.
8. Power down and plug out the compute module.
9. Remove the air duct.
10. Move the jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
11. Restore the air duct to the compute module.
12. Plug in the compute module back to the chassis and power up the server.

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

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

7.2 ME Force Update (J5D2)

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

Table 53. Force ME Update Jumper (J5D2)

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

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

Steps to perform the Force ME Update:

1. Plug out the compute module from the chassis.
2. Remove the air duct. See the *Service Guide* for instructions.
3. Move the jumper (J5D2) from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Restore the air duct back to the compute module.
5. Plug in the compute module back to the chassis.
6. Perform the ME firmware update procedure as documented in the Release Notes file that is included in the given system update package.
7. After update process is done, plug out the compute module out of the chassis.
8. Remove the air duct.
9. Move the jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
10. Restore the compute module back to the chassis.

7.3 Password Clear (J7A6)

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

Table 54. Password Clear Jumper (J7A6)

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

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

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

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

The recommended steps to clear the User and Administrator passwords are:

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

7.4 BIOS Recovery Mode (J7A7)

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

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

- Recovery Mode Jumper – This jumper causes the BIOS to boot in Recovery Mode.

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

Table 55. BIOS Recovery Mode Jumper (J7A7)

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

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

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

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

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

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

The following steps demonstrate this recovery process:

1. Plug out the compute module and remove the air duct.
2. Move the jumper (J7A7) from the default operating position (covering pins 1 and 2) to the BIOS Recovery position (covering pins 2 and 3).
3. Restore the air duct back to the compute module.
4. Plug in the compute module back to the chassis.
5. Power on the compute module.
6. The BIOS will load and boot with the backup BIOS image without any video or display.

7. When the compute module boots into the EFI shell directly, the BIOS recovery is successful.
8. Power off the compute module.
9. Plug out the compute module from the chassis.
10. Remove the air duct and put the jumper (J7A7) back to the normal position (covering pins 1 and 2).
11. Restore the air duct and put the compute module back to the chassis.
12. A normal BIOS update can be performed if needed.

7.5 BIOS Default (J7A3)

Table 56. BIOS Default Jumper

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

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

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

The recommended steps to reset to the BIOS defaults are:

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

8 Intel® Light-Guided Diagnostics

Intel® Server Board S2600TPR has several onboard diagnostic LEDs to assist in troubleshooting board-level issues. This section provides a description of the location and function of each LED on the server board.

8.1 Status LED

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

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

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

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

The following table maps the server state to the LED state.




Figure 46. Status LED (E) and ID LED (D)

Table 57. Status LED State Definitions

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

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

8.2 ID LED

The ID LED provides a visual indication of the server board or compute module being serviced. The state of the ID LED is affected by the following:

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

Table 58. ID LED

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

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

8.3 BMC Boot/Reset Status LED Indicators

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

Table 59. 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.

BMC Boot/Reset State	Chassis ID LED	Status LED	Comment
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.

8.4 InfiniBand® Link/Activity LED

The server board provides dedicated LEDs for InfiniBand® Link/Activity. They are located on the baseboard rear, near diagnostic LED set. This set of LEDs works only on the Intel® Server Board S2600TPFR.

The following table maps the system state to the LED state.




Figure 47. InfiniBand® Link LED (K) and InfiniBand® Activity LED (J)

Table 60. InfiniBand® Link/Activity LED

LED Color	LED State	NIC State
Amber (Right)	Off	No Logical Link
	Blinking	Logical Link established
Green (Left)	Off	No Physical Link
	On	Physical Link established

8.5 POST Code Diagnostic LEDs

Eight amber POST code diagnostic LEDs are located on the back left edge of the server board, in the rear I/O area near the QSFP+ connector.

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




Figure 48. Rear Panel Diagnostic LEDs

9 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 system 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/v4 product family should be referenced for more in-depth and design level platform management information.

9.1 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.

9.1.1 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*.

9.1.2 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.
- Fan speed control with SDR
- Fan redundancy monitoring and support
- Enhancements to fan speed control.
- Power supply redundancy monitoring and support
- Acoustic management: Support for multiple fan profiles
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.
- Power state retention
- Power fault analysis
- Intel® Light-Guided Diagnostics
- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.

- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support
- 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.
- Enhancements to KVM redirection
 - Support for higher resolution
- Integrated Remote Media Redirection
- Lightweight Directory Access Protocol (LDAP) support
- Intel® Intelligent Power Node Manager support
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
 - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.5 compliance (product-specific).
- Management support for PMBus* rev1.2 compliant power supplies
- BMC Data Repository (Managed Data Region Feature)
- 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 Firmware 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.

9.2 Platform Management Features and Functions

9.2.1 Power Subsystem

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

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

9.2.2 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states.

Table 61. ACPI Power States

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

State	Supported	Description
S5	Yes	Soft off. <ul style="list-style-type: none"> ▪ The front panel buttons are not locked. ▪ The fans are stopped. ▪ The power-up process goes through the normal boot process. ▪ The power, reset, and ID buttons are unlocked.

9.2.3 System Initialization

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

9.2.3.1 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 PECL Proxy mechanism provided by Manageability Engine (ME). The BMC uses these values as part of the fan-speed-control algorithm.

9.2.3.2 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.

9.2.3.3 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.

9.2.4 System Event Log (SEL)

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

The BMC allocates 95231 bytes (approx 93 KB) of non-volatile storage space to store system events. The SEL timestamps may not be in order. Up to 3,639 SEL records can be stored at a time. Because the SEL is circular, any command that results in an overflow of the SEL beyond the allocated space will overwrite the oldest entries in the SEL, while setting the overflow flag.

9.3 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.

9.3.1 Sensor Scanning

The value of many of the BMC's sensors is derived by the BMC firmware 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 firmware 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 S4 or S5 power state.

9.3.2 Sensor Rearm Behavior

9.3.2.1 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.

9.3.2.2 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.

9.3.3 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.

9.3.4 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.

9.3.5 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.

9.3.6 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 firmware stack is reset or the BMC CPU itself is reset.

9.3.7 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. If a successful sensor read is completed, the counter resets to zero.

9.3.8 VR Watchdog Timer

The BMC firmware 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 firmware powers down the system.

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

9.3.9 System Airflow Monitoring

This sensor is only valid in Intel chassis. 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 PWM values. The specific Pulse Width Modulation (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.

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

Additionally, the BMC firmware 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.

9.3.10.1 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.

9.3.10.2 Processor DTS-Spec Margin Sensor(s)

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3/v4 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.

9.3.10.3 Processor Thermal Margin Sensor(s)

Each processor supports a physical thermal margin sensor per core that is readable through the PECL interface. This provides a relative value representing a thermal margin from the core's throttling thermal trip point. Assuming that temp controlled throttling is enabled; the physical core temp 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.

9.3.10.4 Processor Thermal Control Monitoring (Prochot)

The BMC firmware 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 PECL-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.

9.3.10.5 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 initiate throttling to reduce processor power draw, therefore indirectly lowering the VRD temperature.

There is one processor VRD_HOT# signal per CPU slot. The BMC instantiates one discrete IPMI sensor for each VRD_HOT# signal. This sensor monitors a digital signal that indicates whether a processor VRD is running in an over-temperature condition. When the BMC detects that this signal is asserted it will cause a sensor assertion which will result in an event being written into the sensor event log (SEL).

9.3.10.6 Inlet Temperature Sensor

Each platform supports a thermal sensor for monitoring the inlet temperature. There are four potential sources for inlet temperature reading.

9.3.10.7 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.

9.3.10.8 Server South Bridge (SSB) Thermal Monitoring

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

9.3.10.9 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, the volumetric system air flow (cubic feet per meter or CFM), and altitude range.

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.

9.3.10.10 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 an input to the fan speed control. The BMC will instantiate an IPMI temperature sensor for each device on the baseboard.

9.3.10.11 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/v4 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.

9.3.10.12 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.

9.3.10.13 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.

9.3.10.14 Server South Bridge (SSB) ThermTrip Monitoring

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

9.3.10.15 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).

9.3.11 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 62. 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

9.3.11.1 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. An 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 shown in Table 64.

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

Table 63. 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.

9.3.11.2 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 slots are not populated as required by the platform hardware 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 the missing processor into the correct slot.
3. AC power on the server.

9.3.11.3 ERR2 Timeout Monitoring

The BMC supports an ERR2 Timeout Sensor (1 per CPU) that asserts if a CPU's ERR2 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 be always do so reliably. A continuously asserted ERR2 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.

9.3.11.4 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 PECL 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 PECL is non-functional (it isn't 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/v4 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.

9.3.11.5 MSID Mismatch Sensor

The BMC supports an *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).

9.3.12 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.

9.3.12.1 DIMM Voltage Sensors

Some systems support either LVDDR (Low Voltage DDR) memory or regular (non-LVDDR) memory. During POST, the system BIOS detects which type of memory is installed and configures the hardware to deliver the correct voltage.

Since the nominal voltage range is different, this necessitates the ability to set different thresholds for any associated IPMI voltage sensors. The BMC firmware supports this by implementing separate sensors (that is, separate IPMI sensor numbers) for each nominal voltage range supported for a single physical sensor and it enables/disables the correct IPMI sensor based on which type memory is installed. The sensor data records for both these DIMM voltage sensor types have scanning disabled by default. Once the BIOS has completed its POST routine, it is responsible for communicating the DIMM voltage type to the BMC which will then enable sensor scanning of the correct DIMM voltage sensor.

9.3.13 Fan Monitoring

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

9.3.13.1 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 not to 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:

1. The system is reset or power-cycled.
2. 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.

9.3.13.2 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).

9.3.13.3 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.

9.3.13.4 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.

9.3.13.5 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 firmware 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 hardware for only 100ms, which should not be long enough to result in triggering a fan fault SEL event.

9.3.14 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.
 - See section 9.3.14.3 for more details.

9.3.14.1 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.

9.3.14.2 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.

9.3.14.3 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.

9.3.14.4 Thermal and Acoustic Management

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

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

9.3.14.5 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 the fan speed control:

- Front panel temperature sensor
- Baseboard temperature sensors
- CPU DTS-Spec margin sensors
- DIMM thermal margin sensors
- Exit air temperature sensor
- Global aggregate thermal margin sensors
- SSB (Intel® C610 Series Chipset) temperature sensor
- On-board Ethernet controller temperature sensors (support for this is specific to the Ethernet controller being used)
- Add-in Intel® SAS/IO module temperature sensor(s) (if present)
- Power supply thermal sensors (only available on PMBus*-compliant power supplies)

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

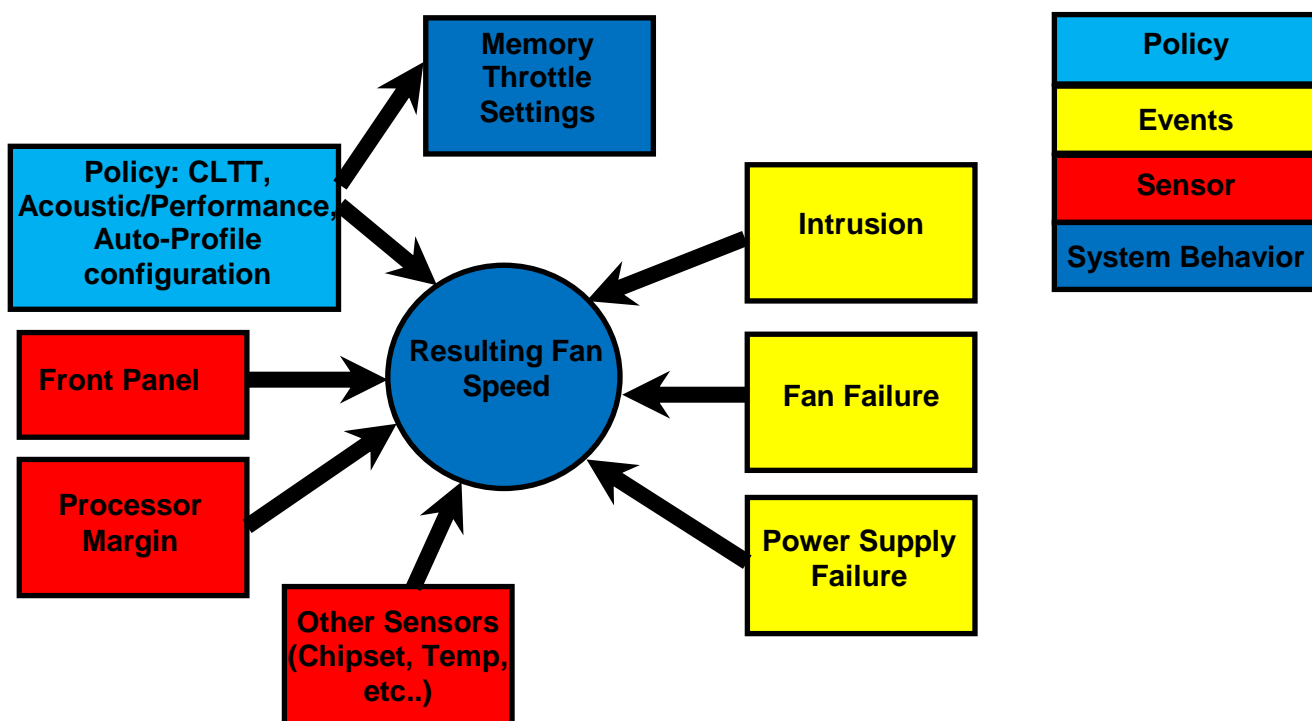


Figure 49. High-level Fan Speed Control Process

9.3.14.5.1 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.

9.3.14.5.2 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. This section provides an overview of this management capability from a BMC perspective.

9.3.14.5.2.1 Memory Thermal Throttling

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

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

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

9.3.14.5.3 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.

9.3.14.5.4 Dynamic (Hybrid) CLTT

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

In order to support this feature, the BMC firmware 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, altitude setting, 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 firmware 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.

9.3.14.5.5 Auto-profile

The auto-profile feature is to improve upon previous platform configuration-dependent FSC and maintain competitive acoustics. This feature is not available for third-party customization.

The BIOS and BMC will handshake to automatically understand configuration details and automatically select the optimal fan speed control profile in the BMC.

Users will only select a performance or an acoustic profile selection from the BIOS menu for use with Intel® Server Chassis and the fan speed control will be optimal for the configuration loaded.

Users can still choose performance or acoustic profile in BIOS setting. Default is acoustic. Performance option is recommended if any other high power add-in cards (higher than 75W) are installed.

9.3.14.5.6ASHRAE Compliance

There is no manual selection of profiles at different altitudes. Altitude impact is covered by auto-profile.

9.3.14.6 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.

9.3.14.6.1System 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 firmware 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.

9.3.14.6.2Use 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.

9.3.14.7 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.

9.3.14.8 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 that it would otherwise be driving the fans to. 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 fans speeds in instances where the fan speed optimized for a given platform may not be sufficient when a high end add-in 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 35C.

9.3.15 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*, see the System Management Interface Forum Web site <http://www.powersig.org/>.

9.3.16 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. See the appropriate Platform Specific Information for power unit redundancy support.

This sensor is configured as 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.

9.3.17 Component Fault LED Control

Several sets of component fault LEDs are supported on the server board. See Figure 8. Intel® Light Guided Diagnostic LED. Some LEDs are owned by the BMC and some by the BIOS.

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

- **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 hardware 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)

Table 64. Component Fault LEDs

Component	Owner	Color	State	Description
Fan Fault LED	BMC	Amber	Solid On	Fan failed
		Amber	Off	Fan working correctly

Component	Owner	Color	State	Description
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.

9.3.18 CMOS Battery Monitoring

The BMC monitors the voltage level from the CMOS battery, which provides battery backup to the chipset RTC. 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 software enable bit to allow the BMC firmware to poll the battery voltage at a relatively slow rate in order to conserve battery power.

9.4 Intel® Intelligent Power Node Manager (NM)

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 settings (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 monitor 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.

9.4.1 Hardware Requirements

NM is supported only on platforms that have the NM firmware 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 feature 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 server board.

9.4.2 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.

9.4.3 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 firmware 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*”.

9.4.4 PECI 3.0

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

9.4.5 NM “Discovery” OEM SDR

An 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.

9.4.6 SmaRT/CLST

The power supply optimization provided by SmaRT/CLST relies on a platform hardware capability as well as ME firmware 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 firmware 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 a fixed maximum throttle time and a fixed throttle release ramp time.
3. Thermal Protection CLST: High power supply thermal event; results in a throttling hiccup mode with a fixed maximum throttle time and a fixed throttle release ramp time.

When the SMBAlert# signal is asserted, the fans will be gated by hardware 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 firmware if it does have this side-effect.

ME firmware 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 firmware support for this sensor. Please refer to the ME firmware EPS for SEL log details.

9.4.6.1.1 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.

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>

9.5 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 65. Intel® Remote Management Module 4 (RMM4) Options

Intel Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE	Intel® Remote Management Module 4 Lite (ROHS Compliant). EOL by June 2017 (Last Order Date)	RMM4 Lite Activation Key	Enables KVM & media redirection
AXXRMM4LITE2	Intel® Remote Management Module 4 Lite v2 (ROHS Free)	RMM4 Lite Activation Key	Enables KVM & media redirection

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

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

Table 66. Basic and Advanced Server Management Features Overview

Feature	Basic	Advanced
IPMI 2.0 Feature Support	X	X
In-circuit BMC Firmware Update	X	X
FRB 2	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

9.5.1 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.

9.5.2 Embedded Web Server

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

- Microsoft Internet Explorer 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 grey-out font in that user's UI display. The web GUI also provides a launch point for some of the advanced features, such as KVM and media redirection. These features are grayed out in the GUI unless the system has been updated to support these advanced features. The embedded web server only displays US English or Chinese language output.

Additional features supported by the web GUI includes:

- Presents all the Basic features to the users
- Power on/off/reset the server and view current power state
- Displays BIOS, BMC, ME and SDR version information
- Display overall system health.
- Configuration of various IPMI over LAN parameters for both IPV4 and IPV6
- Configuration of alerting (SNMP and SMTP)
- Display system asset information for the product, board, and chassis.
- Display of BMC-owned sensors (name, status, current reading, enabled thresholds), including color-code status of sensors.

- Provides ability to filter sensors based on sensor type (Voltage, Temperature, Fan and Power supply related)
- Automatic refresh of sensor data with a configurable refresh rate
- On-line help
- Display/clear SEL (display is in easily understandable human readable format)
- Supports 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 as is available over IPMI over LAN.
- Ability to get and set Node Manager (NM) power policies
- Display of power consumed by the server
- Ability to view and configure VLAN settings
- Warn user the reconfiguration of IP address will cause disconnect.
- Capability to block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.
- Server Power Control – Ability to force into Setup on a reset
- System POST results – The web server provides the system's Power-On Self Test (POST) sequence for the previous two boot cycles, including timestamps. The timestamps may be viewed in relative to the start of POST or the previous POST code.
- Customizable ports – The web server provides the ability to customize the port numbers used for SMASH, http, https, KVM, secure KVM, remote media, and secure remote media.

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

9.5.3 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; upto 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.

9.5.3.1 Keyboard, Video, Mouse (KVM) Redirection

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

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

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

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

Other attributes of this feature include:

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

KVM redirection feature supports the following resolutions and refresh rates:

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

9.5.3.2 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.

9.5.3.3 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.

9.5.3.4 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.

9.5.3.5 Availability

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

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

9.5.3.6 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.

9.5.3.7 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.

9.5.3.8 Media Redirection

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

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

The following capabilities are supported:

- The operation of remotely mounted devices is independent of the local devices on the server. Both remote and local devices are useable in parallel.
- Either IDE (CD-ROM, floppy) or USB devices can be mounted as a remote device to the server.
- It is possible to boot all supported operating systems from the remotely mounted device and to boot from disk IMAGE (*.IMG) and CD-ROM or DVD-ROM ISO files. See the Tested/supported Operating System List for more information.
- Media redirection supports redirection for both a virtual CD device and a virtual Floppy/USB device concurrently. The CD device may be either a local CD drive or else an ISO image file; the Floppy/USB device may be either a local Floppy drive, a local USB device, or else a disk image file.
- The media redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.
- A remote media session is maintained even when the server is powered-off (in standby mode). No restart of the remote media session is required during a server reset or power on/off. An BMC reset (for example, due to an BMC reset after BMC firmware update) will require the session to be re-established.
- The mounted device is visible to (and useable by) managed system's OS and BIOS in both pre-boot and post-boot states.
- The mounted device shows up in the BIOS boot order and it is possible to change the BIOS boot order to boot from this remote device.
- It is possible to install an operating system on a bare metal server (no OS present) using the remotely mounted device. This may also require the use of KVM-r to configure the OS during installation.

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

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

9.5.3.8.1 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.

9.5.3.8.2 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 Thermal Management

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



AF006862

Figure 50. Air Flow and Fan Identification

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

Table 67. Air Flow

	Single Compute Module Airflow
With Intel® Server Chassis H2312XXKR2/LR2	6 ~ 28CFM
With Intel® Server Chassis H2216XXKR2/LR2	5 ~ 38CFM
With Intel® Server Chassis H2224XXKR2/LR2	7 ~ 35CFM

The compute module supports short-term, excursion-based, operation up to 45°C (ASHRAE A4) with limited performance impact. The configuration requirements and limitations are described in the configuration matrix found in the *Power Budget and Thermal Configuration Tool*, available as a download online at <http://www.intel.com/support>.

The installation and functionality of several components are used to maintain compute module thermals. They include three compute module fans, air duct, and installed CPU heat sinks.

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

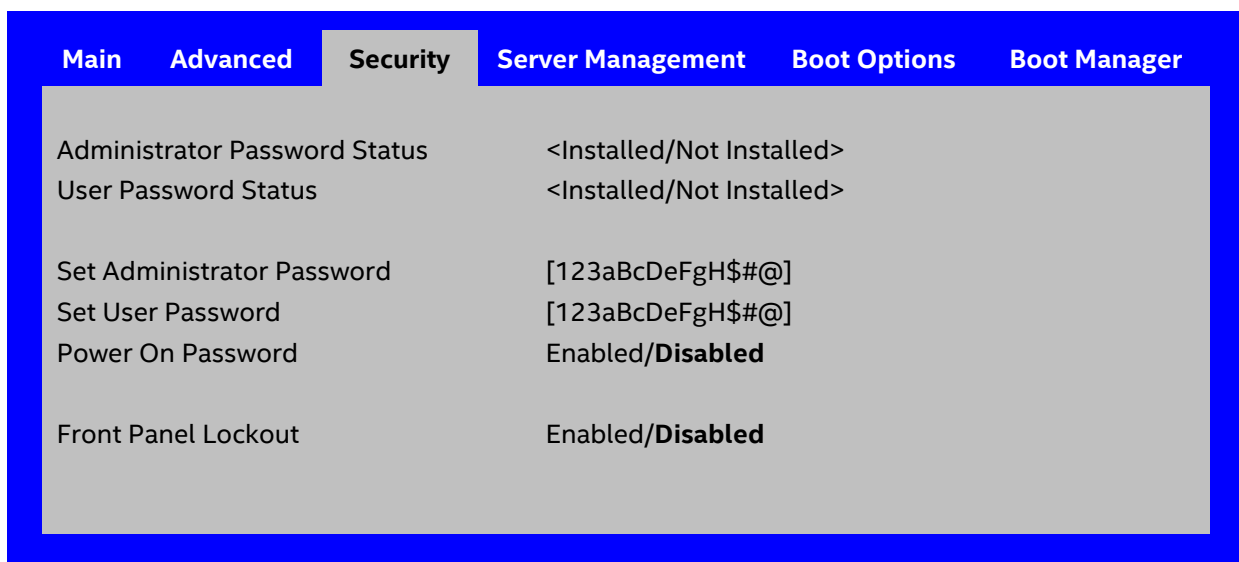
- The compute module operating ambient is designed for sustained operation up to 35°C (ASHRAE Class A2) with short-term excursion-based operation up to 45°C (ASHRAE Class A4).
 - The compute module can operate up to 40°C (ASHRAE Class A3) for up to 900 hours per year.
 - The compute module can operate up to 45°C (ASHRAE Class A4) for up to 90 hours per year.
 - The compute module performance may be impacted when operating within the extended operating temperature range.
 - There is no long-term system reliability impact when operating at the extended temperature range within the approved limits.
- Specific configuration requirements and limitations are documented in the configuration matrix found in the *Power Budget and Thermal Configuration Tool*, available as a download online at <http://www.intel.com/support>.
- The CPU-1 processor + CPU heat sink must be installed first. The CPU-2 heat sink must be installed at all times, with or without a processor installed.
- Memory Slot population requirements –
 - **DIMM Population Rules on CPU-1** – Install DIMMs in order; Channels A, B, C, and D.
 - **DIMM Population Rules on CPU-2** – Install DIMMs in order; Channels E, F, G, and H.
- With the compute module operating, the air duct must be installed at all times.

11 System Security

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

- Password Protection
- Front Panel Lockout

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



11.1 Password Setup

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

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

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

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

Passwords are case sensitive.

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

Once set, a password can be cleared by changing it to a null string. This requires the Administrator password, and must be done through BIOS Setup or other explicit means of changing the passwords. Clearing the Administrator password will also clear the User password. Passwords can also be cleared by using the Password Clear jumper on the server board. See Chapter 7 – Configuration 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.

11.1.1 System Administrator Password Rights

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

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

11.1.2 Authorized System User Password Rights and Restrictions

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

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

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

11.2 Front Panel Lockout

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

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

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

11.3 Trusted Platform Module (TPM) support

The Intel® Server Board S2600TPTR and the Intel® Compute Module HNS2600TP24STR has the option to support a Trusted Platform Module (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 2.0 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.

11.3.1 TPM security BIOS

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

Measures and stores the boot process in the TPM microcontroller to allow a TPM enabled operating system to verify system boot integrity.

- Produces 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.

11.3.2 Physical presence

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

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

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).

11.3.3 TPM security setup options

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

Table 68. 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.

11.4 Intel® Trusted eXecution Technology (TXT)

The Intel® Xeon® Processor E5-2600 v3 and v4 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 or v2.0, as defined by the Trusted Computing Group TPM PC Client Specifications, Revision 1.2 or 2.0.

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/>

12 Environmental Limits Specification

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

Note: *The Energy Star compliance is at system level, but not board level. Use of Intel boards alone does not guarantee Energy Star compliance.*

Table 69. Server Board Design Specifications

Parameter	Limits
Operating Temperature	+10°C to +35°C with the maximum rate of change not to exceed 10°C per hour.
Non-Operating Temperature	-40°C to +70°C
Non-Operating Humidity	90%, non-condensing at 35°C
Acoustic noise	Sound power: 7.0BA with hard disk drive stress only at room ambient temperature (23 +/-2°C)
Shock, operating	Half sine, 2g peak, 11 mSec
Shock, unpackaged	Trapezoidal, 25g, velocity change 205 inches/second (80 lbs. to < 100 lbs.)
Vibration, unpackaged	5 Hz to 500 Hz, 2.20 g RMS random
Shock and vibration, packaged	ISTA (International Safe Transit Association) Test Procedure 3A
ESD	+/-12 KV except I/O port +/- 8 KV per Intel® Environmental Test Specification
System Cooling Requirement in BTU/Hr.	1600 Watt Max – 5459 BTU/hour
System Cooling Requirement in BTU/Hr.	2130 Watt Max – 7268 BTU/hour

Disclaimer Note: *Intel ensures the unpackaged server board and system meet the shock requirement mentioned above through its own chassis development and system configuration. It is the responsibility of the system integrator to determine the proper shock level of the board and system if the system integrator chooses different system configuration or different chassis. 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.*

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 1600W and 2130W 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 boards detailed in this document.

13.1 Power Supply DC Output Connector

The server board includes two main power Minifit Jr* connectors allowing for power supplies to attach directly to the server board. The connectors are two sets of 2x3 pin and can be used to deliver 12amps per pin or 60+Amps total. Note that no over-voltage protective circuits will exist on the board.

Table 70. Power Supply DC Power Input Connector Pin-out

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

13.2 Power Supply DC Output Specification

13.2.1 Output Power/Currents

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

Table 71. Minimum 1200W/1600W Load Ratings

Parameter	Min	Max.	Peak ^{1,2}	Unit
12V main	0.0	60.0	72.0	A
5Vstby	0.0	2.0	2.4	A

Table 72. Minimum 2130W Load Ratings

Parameter	Min	Max.	Peak	Unit
12V main (240 – 264 VAC)	0.0	178.0	210.0	A
12V stby	0.0	3.5	4.0	A

Notes:

1. Peak combined power for all outputs shall not exceed 800W.

2. 12Vstby must be able to provide 4.0A peak load with single power supply. The power supply fan is allowed to run in standby mode for loads > 1.5A.
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. The peak load requirement should apply to full operating temperature range.
4. The setting of $I_{Peak} < I_{OCW} < I_{OCP}$ needs to be followed to make the CLST work reasonably.
5. Power supply must protect itself in case system doesn't take any action to reduce load based on SMBAlert# signal asserting.
6. The power supply shall support 25msec peak power at 20% duty cycle step loading for an average current at the current rating.

Standby Output

The 5VSB (12VSB for the 2130W PSU) 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.

13.2.2 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 73. Voltage Regulation Limits

Parameter	Tolerance	Min	Nom	Max	Units
+12V	- 5%/+5%	+11.40	+12.00	+12.60	V _{rms}
+5V stby	- 5%/+5%	+4.75	+5.00	+5.25	V _{rms}
+12V stby (2130W)	- 5%/+5%	+11.40	+12.00	+12.60	V _{rms}

13.2.3 Dynamic Loading

The output voltages shall remain within limits specified for the step loading and capacitive loading specified in the following table. The load transient repetition rate shall be tested between 50 Hz and 5 KHz 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 74. Transient Load Requirements

Output	Δ Step Load Size	Load Slew Rate	Test capacitive Load
+5VSB	1.0A	0.25 A/ μ sec	20 μ F
+12VSB	1.0A	0.25 A/ μ sec	20 μ F
+12V	60% of max load	0.25 A/ μ sec	2000 μ F

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

13.2.4 Capacitive Loading

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

Table 75. Capacitive Loading Conditions

Output	MIN	MAX	Units
+5VSB	20	3100	μ F
+12VSB	20	3100	μ F
+12V	500	25000	μ F

13.2.5 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 m Ω . This path may be used to carry DC current.

13.2.6 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.

13.2.7 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.

13.2.8 Common Mode Noise

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

13.2.9 Soft Starting

The Power Supply shall contain 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.

13.2.10 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.

13.2.11 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. The power supply shall use a latching mechanism to prevent insertion and extraction of the power supply when the AC power cord is inserted into the power supply.

13.2.12 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 12VSB output is not required to actively share current between power supplies (passive sharing). The 12VSB output 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.

13.2.13 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 10 μ F tantalum capacitor in parallel with a 0.1 μ F ceramic capacitor is placed at the point of measurement.

Table 76. Ripples and Noise

+12V main	+5VSB	+12VSB
120mVp-p	50mVp-p	120mVp-p

13.2.14 Timing Requirement

These are the timing requirements for the power supply operation. The output voltages must rise from 10% to within regulation limits ($T_{\text{vout_rise}}$) within 5 to 70ms. For 5VSB, 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 through the AC input, with PS_ON held low and the PS_ON signal, with the AC input applied.

Table 77. Timing Requirements (5VSB)

Item	Description	Min	Max	Units
T _{vout_rise}	Output voltage rise time	5.0 *	70 *	ms
T _{sb_on_delay}	Delay from AC being applied to 5VSB being within regulation.		1500	ms
T _{ac_on_delay}	Delay from AC being applied to all output voltages being within regulation.		3000	ms
T _{vout_holdup}	Time 12Vl output voltage stay within regulation after loss of AC.	13		ms
T _{pwok_holdup}	Delay from loss of AC to de-assertion of PWOK	12		ms
T _{pson_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
T _{pwok_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
T _{pwok_low}	Duration of PWOK being in the de-asserted state during an off/on cycle using AC or the PSON signal.	100		ms
T _{sb_vout}	Delay from 5VSB being in regulation to O/Ps being in regulation at AC turn on.	50	1000	ms
T _{5VSB_holdup}	Time the 5VSB output voltage stays within regulation after loss of AC.	70		ms

Note:

* The 5VSB output voltage rise time shall be from 1.0ms to 25ms.

Table 78. Timing Requirements (12VSB)

Item	Description	Min	Max	Units
Tvout_rise	Output voltage rise time	5.0 *	70 *	ms
Tsb_on_delay	Delay from AC being applied to 12VSB being within regulation.		1500	ms
Tac_on_delay	Delay from AC being applied to all output voltages being within regulation.		3000	ms
Tvout_holdup	Time 12V output voltage stay within regulation after loss of AC.	11		ms
Tpwok_holdup	Delay from loss of AC to de-assertion of PWOK	10		ms
Tpson_on_delay	Delay from PSON# active to output voltages within regulation limits.	5	400	ms
Tpson_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
Tpwok_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 12VSB being in regulation to O/Ps being in regulation at AC turn on.	50	1000	ms
T12VSB_holdup	Time the 12VSB output voltage stays within regulation after loss of AC.	70		ms

Note:

* The 12VSB output voltage rise time shall be from 5.0ms to 10ms.

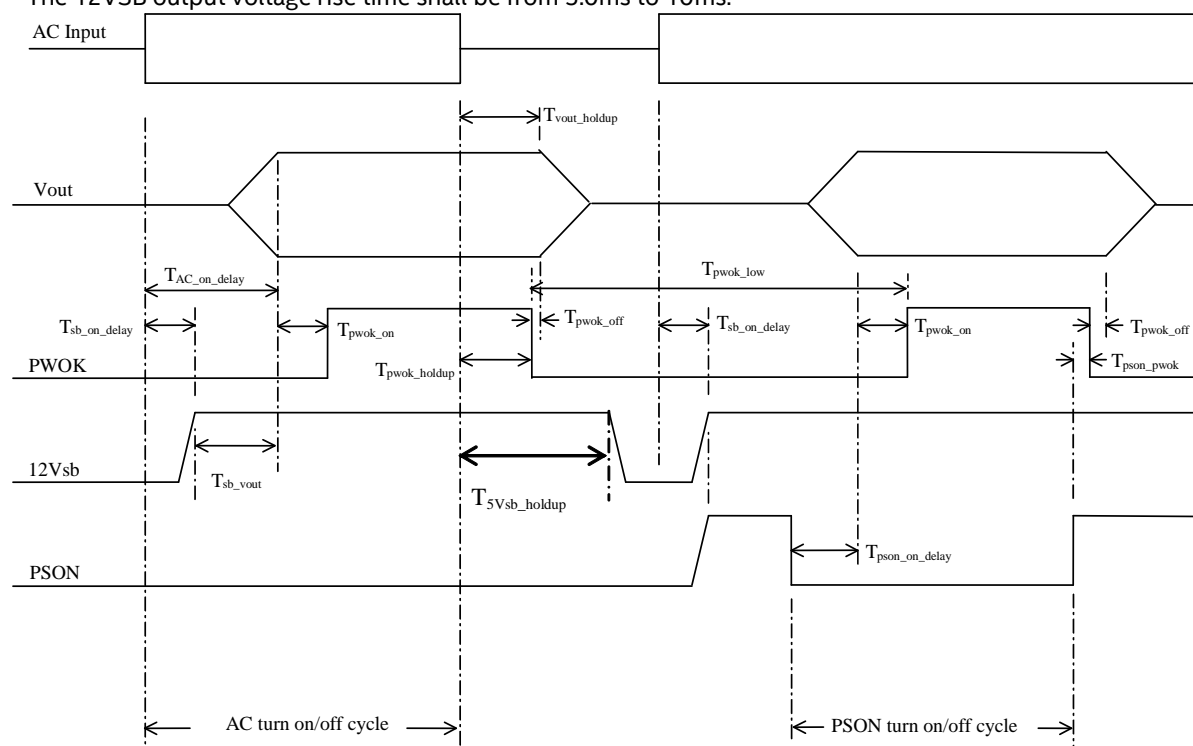


Figure 51. Turn On/Off Timing (Power Supply Signals – 5VSB)

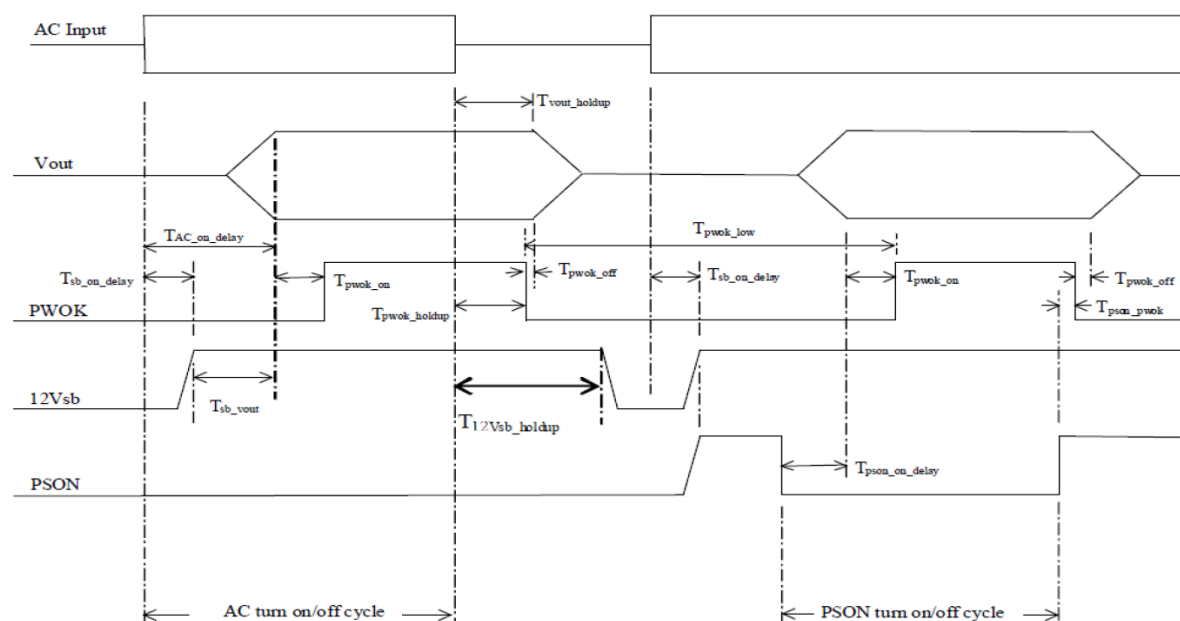


Figure 52. Turn On/Off Timing (Power Supply Signals – 12VSB)

Appendix A: Integration and Usage Tips

- When adding or removing components or peripherals from the server board, AC power must be removed. With AC power plugged into the server board, 5V standby is still present even though the server board is powered off.
- This server board supports the Intel® Xeon® Processor E5-2600 v3/v4 product family with a Thermal Design Power (TDP) of up to and including 160 Watts. Previous generations of the Intel® Xeon® processors are not supported.
- Processors must be installed in order. CPU 1 must be populated for the server board to operate.
- 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.
- For the best performance, the number of DIMMs installed should be balanced across both processor sockets and memory channels. For example, a two-DIMM configuration performs better than a one-DIMM configuration. In a two-DIMM configuration, DIMMs should be installed in DIMM sockets A1 and D1. A six-DIMM configuration (DIMM sockets A1, B1, C1, D1, E1, and F1) performs better than a three-DIMM configuration (DIMM sockets A1, B1, and C1).
- Normal Integrated BMC functionality is disabled with the BMC Force Update jumper set to the “enabled” position (pins 2-3). The server should never be run with the BMC Force Update jumper set in this position and should only be used when the standard firmware update process fails. This jumper should remain in the default (disabled) position (pins 1-2) when the server is running normally.
- When performing a normal BIOS update procedure, the BIOS recovery jumper must be set to its default position (pins 1-2).

Appendix B: Integrated BMC Sensor Tables

This appendix lists the sensor identification numbers and information about the sensor type, name, supported thresholds, assertion and de-assertion information, and a brief description of the sensor purpose. See the *Intelligent Platform Management Interface Specification, Version 2.0*, for sensor and event/reading-type table information.

- **Sensor Type**

The sensor type references the values in the Sensor Type Codes table in the *Intelligent Platform Management Interface Specification Second Generation v2.0*. It provides a context to interpret the sensor.

- **Event/Reading Type**

The event/reading type references values from the Event/Reading Type Code Ranges and the Generic Event/Reading Type Code tables in the *Intelligent Platform Management Interface Specification Second Generation v2.0*. Digital sensors are specific type of discrete sensors that only have two states.

- **Event Thresholds/Triggers**

The following event thresholds are supported for threshold type 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 Code* or *Sensor Type Code* tables in the *Intelligent Platform Management Interface Specification Second Generation v2.0*, depending on whether the sensor event/reading type is generic or a sensor-specific response.

- **Assertion/Deassertion**

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

As: Assertion

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 by means of the *Get Sensor Reading* command. Unless otherwise indicated, 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 associated sensor. For threshold-based sensors, these 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 in the comment column to describe a sensor:

A: Auto-rearm

M: Manual rearm

I: Rearm by init agent

- **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 front 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 79. BMC Sensor Table

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 ¹ (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					

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
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						
					03 - Power cycle						
					08 - Timer interrupt						
Physical Security (Physical Scrtty)	04h	Chassis Intrusion is chassis-specific	Physical Security 05h	Sensor Specific 6Fh	00 - Chassis intrusion	Degraded OK	As and De	–	Trig Offset	A	X
					04 - LAN leash lost						
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	-

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
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 ¹ (Fan Redundancy)	0Ch	Chassis-specific	Fan 04h	Generic 0Bh	00 – Fully redundant	OK	As and De	–	Trig Offset	A	–
					01 - Redundancy lost	Degraded					
					02 - Redundancy degraded	Degraded					
					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

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
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	-
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 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

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
Front Panel Temperature (Front Panel Temp)	21h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc] UNR	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
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

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
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
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 ² (Chassis specific sensor names)	30h–3Fh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal	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 ³ (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 ³	51h		Power Supply		00 - Presence	OK		-		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
(PS2 Status)		Chassis-specific	08h	Sensor Specific 6Fh	01 - Failure	Degraded	As and De		Trig Offset		
					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
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/ FIVR	Fatal	As and De	-	Trig Offset	M	X
					07 - Presence	OK					

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 2 Status (P2 Status)	71h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip/ FIVR	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	–
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	–

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 ERR2 Timeout (CPU ERR2)	7Ch	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	A	–
IERR recovery dump info (IERR Rec Info)	7Dh	All	OEM sensor type D1h	OEM defined 70h	00h – Dump successfully 01h – Dump failure	OK	As	–	Trig Offset	A	–
Internal Error (IERR)	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	–
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	-

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
VRD Over 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 Fail 1 ³ (PS1 Fan Fail 1)	A0h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	X
Power Supply 1 Fan Fail 2 ³ (PS1 Fan Fail 2)	A1h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	X
PHI 1 Status (GPGPU1 Status)	A2h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
PHI 2 Status (GPGPU2 Status)	A3h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
Power Supply 2 Fan Fail 1 ³ (PS2 Fan Fail 1)	A4h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	X
Power Supply 2 Fan Fail 2 ³ (PS2 Fan Fail 2)	A5h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	X
PHI 3 Status (GPGPU3 Status)	A6h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
PHI 4 Status (GPGPU4 Status)	A7h	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
PHI 1 Avg Pwr	AAh	Platform Specific	Power 03h	Threshold 01h	-	-	-	Analog	-	-	-
PHI 2 Avg Pwr	ABh	Platform Specific	Power 03h	Threshold 01h	-	-	-	Analog	-	-	-
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	-

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
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	[I] [c,nc]	nc = Degraded c = Non-fatal ²	As and De	Analog	R, T	M	-
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 over temperature	Fatal	As and De	–	Trig Offset	M	X
PHI 1 Temp (GPGPU1 Core Temp)	C4h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
PHI 2 Temp (GPGPU2 Core Temp)	C5h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
PHI 3 Temp (GPGPU3 Core Temp)	C6h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
PHI 4 Temp (GPGPU4 Core Temp)	C7h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-

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 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	-
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	Degraded	-	-	-	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
Baseboard Temperature 5 (Platform Specific)	D5h	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)	D6h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
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	–
Hot-swap Backplane 4 Temperature (HSBP 4 Temp)	E0h	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Rear Hard Disk Drive 0 -1 Status (Rear HDD 0 - 1 Stat)	E2h - E3h	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					
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					

Notes:

1. Redundancy sensors will be only present on systems with appropriate hardware to support redundancy (for instance, fan or power supply). Note that power supply redundancy may be lost even when both supplies are operational if the system is loaded beyond the capacity of a single power supply.
2. This is only applicable when the system doesn't support redundant fans. When fan redundancy is supported, then the contribution to system state is driven by the fan redundancy sensor, not individual sensors. On a system with fan redundancy, the individual sensor severities will read the same as the fan redundancy sensor's severity.

3. This is only applicable when the system doesn't support redundant power supplies. When redundancy is supported, then the contribution to system state is driven by the power unit redundancy sensor. On a system with power supply redundancy, the individual sensor severities will read the same as the power unit redundancy sensor's severity.

Appendix C: BIOS Sensors and SEL Data

BIOS owns a set of IPMI-compliant Sensors. These are actually divided in ownership between BIOS POST (GID = 01) and BIOS SMI Handler (GID = 33). The SMI Handler Sensors are typically for logging runtime error events, but they are active during POST and may log errors such as Correctable Memory ECC Errors if they occur.

It is important to remember that a Sensor is uniquely identified by the combination of Sensor Owner plus Sensor Number. There are cases where the same Sensor Number is used with different Sensor Owners – this is not a conflict. For example, in the BIOS Sensors list there is a Sensor Number 83h for Sensor Owner 01h (BIOS POST) as well as for Sensor Owner 33h (SMI Handler), but these are two distinct sensors reporting the same type of event from different sources (Generator IDs 01h and 33h).

On the other hand, each distinct Sensor (GID + Sensor Number) is defined by one specific Sensor Type describing the kind of data being reported, and one specific Event Type describing the type of event and the format of the data being reported.

Table 80. BIOS Sensor and SEL Data

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
Mirroring Redundancy State	01h	33h (SMI Handler)	0Ch (Memory)	0Bh (Discrete, Redundancy State) 0h = Fully Redundant 2h = Redundancy Degraded	ED2 = [7:4] = Mirroring Domain 0-1 = Channel Pair for Socket [3:2] = Reserved [1:0] = Rank on DIMM 0-3 = Rank Number ED3 = [7:5] = Socket ID 0-3 = CPU1-4 [4:3] = Channel 0-3 = Channel A-D for Socket [2:0] = DIMM 0-2 = DIMM 1-3 on Channel

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
Memory RAS Configuration Status	02h	01h (BIOS POST)	0Ch (Memory)	09h (Digital Discrete) 0h = RAS Configuration Disabled 1h = RAS Configuration Enabled	ED2 = [7:4] = Reserved [3:0] Config Err 0 = None 3 = Invalid DIMM Config for RAS Mode ED3 = [7:4] = Reserved [3:0] = RAS Mode 0 = None 1 = Mirroring 2 = Lockstep 4 = Rank Sparing
Memory ECC Error	02h	33h (SMI Handler)	0Ch (Memory)	6Fh (Sensor Specific Offset) 0h = Correctable Error 1h = Uncorrectable Error	ED2 = [7:2] = Reserved [1:0] = Rank on DIMM 0-3 = Rank Number ED3 = [7:5] = Socket ID 0-3 = CPU1-4 [4:3] = Channel 0-3 = Channel A-D for Socket [2:0] = DIMM 0-2 = DIMM 1-3 on Channel
Legacy PCI Error	03h	33h (SMI Handler)	13h (Critical Interrupt)	6Fh (Sensor Specific Offset) 4h = PCI PERR 5h = PCI SERR	ED2 = [7:0] = Bus Number ED3 = [7:3] = Device Number [2:0] = Function Number

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
PCIe Fatal Error (Standard AER Errors) (see Sensor 14h for continuation)	04h	33h (SMI Handler)	13h (Critical Interrupt)	70h (OEM Discrete) <hr/> 0h = Data Link Layer Protocol Error 1h = Surprise Link Down Error 2h = Completer Abort 3h = Unsupported Request 4h = Poisoned TLP 5h = Flow Control Protocol 6h = Completion Timeout 7h = Receiver Buffer Overflow 8h = ACS Violation 9h = Malformed TLP Ah = ECRC Error Bh = Received Fatal Message From Downstream Ch = Unexpected Completion Dh = Received ERR_NONFATAL Message Eh = Uncorrectable Internal Fh = MC Blocked TLP	ED2 = [7:0] = Bus Number <hr/> ED3 = [7:3] = Device Number [2:0] = Function Number
PCIe Correctable Error (Standard AER Errors)	05h	33h (SMI Handler)	13h (Critical Interrupt)	71h (OEM Discrete) <hr/> 0h = Receiver Error 1h = Bad DLLP 2h = Bad TLP 3h = Replay Num Rollover 4h = Replay Timer timeout 5h = Advisory Non-fatal 6h = Link BW Changed 7h = Correctable Internal 8h = Header Log Overflow	ED2 = [7:0] = Bus Number <hr/> ED3 = [7:3] = Device Number [2:0] = Function Number
BIOS POST Error	06h	01h (BIOS POST)	0Fh (System Firmware Progress)	6Fh (Sensor Specific Offset) <hr/> 0h = System Firmware Error (POST Error Code)	ED2 = [7:0] = LSB of POST Error Code <hr/> ED3 = [7:0] MSB of POST Error Code

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
QPI Correctable Errors (reserved for Validation)	06h	33h (SMI Handler)	13h (Critical Interrupt)	72h (OEM Discrete)	ED2 = Reserved
				Offset Reserved	ED3 = Reserved
QPI Fatal Error (see Sensor 17h for continuation)	07h	33h (SMI Handler)	13h (Critical Interrupt)	73h (OEM Discrete)	ED2 =
				0h = Link Layer Uncorrectable ECC Error 1h = Protocol Layer Poisoned Packet Reception Error 2h = Link/PHY Init Failure with resultant degradation in link width 3h = CSI PHY Layer detected drift buffer alarm 4h = CSI PHY detected latency buffer rollover 5h = CSI PHY Init Failure 6h = CSI Link Layer generic control error (buffer overflow/underflow, credit underflow and so on.) 7h = Parity error in link or PHY layer 8h = Protocol layer timeout detected 9h = Protocol layer failed response Ah = Protocol layer illegal packet field, target Node ID and so on. Bh = Protocol Layer Queue/table overflow/underflow Ch = Viral Error Dh = Protocol Layer parity error Eh = Routing Table Error Fh = (unused)	[7:0] = Node ID ED2 = [7:0] = Node ID 0-3 = CPU1-4 ED3 = No Data
Chipset Proprietary (reserved for Validation)	08h	33h (SMI Handler)	19h (Chipset)	75h (OEM Discrete)	ED2 = Reserved
				Offset Reserved	ED3 = Reserved

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
QPI Link Width Reduced	09h	01h (BIOS POST)	13h (Critical Interrupt)	77h (OEM Discrete) 1h = Reduced to ½ width 2h = Reduced to ¼ width	ED2 = [7:0] = Node ID 0-3 = CPU1-4 ED3 = No Data
Memory Error Extension (reserved for Validation)	10h	33h (SMI Handler)	0Ch (Memory)	7Fh (OEM Discrete) Offset Reserved	ED2 = Reserved ED3 = Reserved
Sparing Redundancy State	11h	33h (SMI Handler)	0Ch (Memory)	0Bh (Discrete, Redundancy State) 0h = Fully Redundant 2h = Redundancy Degraded	ED2 = [7:4] = Sparing Domain 0-3 = Channel A-D for Socket [3:2] = Reserved [1:0] = Rank on DIMM 0-3 = Rank Number ED3 = [7:5] = Socket ID 0-3 = CPU1-4 [4:3] = Channel 0-3 = Channel A-D for Socket [2:0] = DIMM 0-2 = DIMM 1-3 on Channel
Memory RAS Mode Select	12h	01h (BIOS POST)	0Ch (Memory)	09h (Digital Discrete) 0h = RAS Configuration Disabled 1h = RAS Configuration Enabled	ED2 = Prior Mode [7:4] = Reserved [3:0] = RAS Mode 0 = None 1 = Mirroring 2 = Lockstep 4 = Rank Sparing ED3 = Selected Mode [7:4] = Reserved [3:0] = RAS Mode 0 = None 1 = Mirroring 2 = Lockstep 4 = Rank Sparing

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type Offset Values	Event Data 2 Event Data 3
Memory Parity Error	13h	33h (SMI Handler)	0Ch (Memory)	6Fh (Sensor Specific Offset) 2h = Address Parity Error	ED2 = Validity [7:5] = Reserved [4] = Channel Validity Check 0 = ED3 Chan # Not Valid 1 = ED3 Chan # Is Valid [3] = DIMM Validity Check 0 = ED3 DIMM # Not Valid 1 = ED3 DIMM # Is Valid [2:0] = Error Type 0 = Not Known 2 = Address Parity Error ED3 = Location [7:5] = Socket ID 0-3 = CPU1-4 [4:2] = Channel 0-3 = Channel A-D for Socket [1:0] = DIMM 0-2 = DIMM 1-3 on Channel
PCIe Fatal Error#2 (Standard AER Errors) (continuation of <u>Sensor 04h</u>)	14h	33h (SMI Handler)	13h (Critical Interrupt)	76h (OEM Discrete) 0h = Atomic Egress Blocked 1h = TLP Prefix Blocked Fh = Unspecified Non-AER Fatal Error	ED2 = [7:0] = Bus Number ED3 = [7:3] = Device Number [2:0] = Function Number

Sensor Name	Sensor Number	Sensor Owner (GID)	Sensor Type	Event/Reading Type	Event Data 2
				Offset Values	Event Data 3
QPI Fatal Error (continuation of <u>Sensor 07h</u>)	17h	33h (SMI Handler)	13h (Critical Interrupt)	74h (OEM Discrete)	ED2 =
				0h = Illegal inbound request 1h = IIO Write Cache Uncorrectable Data ECC Error 2h = IIO CSR crossing 32-bit boundary Error 3h = IIO Received XPF physical/logical redirect interrupt inbound 4h = IIO Illegal SAD or Illegal or non-existent address or memory 5h = IIO Write Cache Coherency Violation 6Fh (Sensor Specific Offset)	[7:0] = Node ID 0-3 = CPU1-4
				1h = System Boot Event 5h = Time Synch	ED3 = No Data
System Event	83h	01h (BIOS POST)	12h (System Event))	6Fh (Sensor Specific Offset)	ED2 = (only for Time Synch)
				5h = Time Synch	[7:0] Synch # 00h = 1 st in pair 80h = 2 nd in pair
					ED3 = No Data
System Event	83h	33h (SMI Handler)	12h (System Event))	6Fh (Sensor Specific Offset)	ED2 = (only for Time Synch)
				5h = Time Synch	[7:0] Synch # 00h = 1 st in pair 80h = 2 nd in pair
					ED3 = No Data

Appendix D: POST Code Diagnostic LED Decoder

During the system POST process, Diagnostic LED Codes are used extensively as a mechanism to indicate progress and Fatal Halt conditions independently of the video display. If the system hangs or halts, the Diagnostic LED display can help determine the reason even when video is not available.

These Diagnostic LEDs are equivalent to the Legacy “Port 80 POST Codes”, and a Legacy I/O Port 80 output will be displayed as a Diagnostic LED code.

The Diagnostic LEDs are a set of LEDs found on the back edge of the server board. There are 8 Diagnostic LEDs which form a 2 hex digit (8 bit) code read left-to-right as facing the rear of the server.




Figure 53. Diagnostic LED Placement Diagram

An LED which is ON represents a 1 bit value, and an LED which is OFF represents a 0 bit value. The LED bit values are read as Most Significant Bit to the left, Least Significant Bit to the right.

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

Table 81. POST 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.

POST Memory Initialization MRC Diagnostic Codes

This is a brief list of the Diagnostic LED codes displayed during memory initialization by the Memory Reference Code (MRC), the BIOS component responsible for it. There are two types of POST Diagnostic Codes used by the MRC, Fatal Error Codes and Progress Codes.

MRC Fatal Error Codes are necessary because if the Memory Initialization fails badly for some reason – like no usable memory installed – the system would not have the resources to give any other error indication. So in the case of a major failure during Memory Initialization, the system outputs a Fatal Error Code to Port 80 (the Diagnostic LEDs) and executes a Halt. These Fatal Error Halts do not change the Status LED, and they do not get logged as SEL Events.

The MRC Progress Codes are displays to the Diagnostic LEDs that show the execution point in the MRC operational path at each step. The intent is that if the system hangs during execution of the MRC, the LED display will tell at what point in the code the system was executing.

Be aware that these are Diagnostic LED display codes used in early POST by the MRC. Later in POST these same Diagnostic LED display codes are used for other BIOS Progress Codes.

Also, be aware that the MRC Fatal Error Codes and MRC Progress Codes are **not controlled by the BIOS** and are subject to change at the discretion of the Memory Reference Code teams.

Table 82. MRC Fatal Error Codes

Error Code	Fatal Error Code Explanation (with MRC Internal Minor Code)
0xE8	<u>No Usable Memory Error:</u> 01h = No memory was detected via 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.
0xE9	<u>Memory is locked by Intel® Trusted Execution Technology and is inaccessible.</u>

Error Code	Fatal Error Code Explanation (with MRC Internal Minor Code)
0xEA	<u>DDR4 Channel Training Error:</u> 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)
0xEB	<u>Memory Test Failure:</u> 01h = Software memtest failure. 02h = Hardware memtest failed. 03h = Hardware memtest failure in Lockstep Channel mode requiring a channel to be disabled. <i>This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.</i>
0xED	<u>DIMM Configuration/Population Error:</u> 01h = Different DIMM types (RDIMM, LRDIMM) are detected installed in the system. 02h = Violation of DIMM population rules. 03h = The third DIMM slot cannot be populated when QR DIMMs are installed. 04h = UDIMMs are not supported. 05h = Unsupported DIMM Voltage.
0xEF	<u>Indicates a CLTT table structure error.</u>

Table 83. MRC Progress Codes

Progress Code	Main Sequence	Subsequences/Subfunctions
0xB0	Detect DIMM population	—n/a—
0xB1	Set DDR4 frequency	—n/a—
0xB2	Gather remaining SPD data	—n/a—
0xB3	Program registers on the memory controller level	—n/a—
0xB4	Evaluate RAS modes and save rank information	—n/a—
0xB5	Program registers on the channel level	—n/a—
0xB6	Perform the JEDEC defined initialization sequence	—n/a—
0xB7	Train DDR4 ranks	—n/a—
0x01	↓	Read DQ/DQS training
0x02	↓	Receive Enable training
0x03	↓	Write Leveling training
0x04	↓	Write DQ/DQS training
0x05	↓	DDR channel training done
0xB8	Initialize CLTT	—n/a—

Progress Code	Main Sequence	Subsequences/Subfunctions
0xB9	Hardware memory test and init	—n/a—
0xBA	Execute software memory init	—n/a—
0xBB	Program memory map and interleaving	—n/a—
0xBC	Program RAS configuration	—n/a—
0xBF	MRC is done	—n/a—

POST Progress Code Checkpoints

During the system boot process, the BIOS executes a number of platform configuration processes, each of which is assigned a specific POST Progress Code, a 2-digit hexadecimal number. As each configuration routine is started, the BIOS displays the POST Progress Code on the Diagnostic LEDs found on the back edge of the server board.

To assist in troubleshooting a system hang during the POST process, the POST Progress Code displayed in the Diagnostic LEDs can be used to identify the last POST process to begin execution.

Table 84. POST Progress Codes

Progress Code	Description
SEC Phase	
0x01	First POST code after CPU reset
0x02	Microcode load begin
0x03	CRAM initialization begin
0x04	Pei Cache When Disabled
0x05	SEC Core At Power On Begin.
0x06	Early CPU initialization during Sec Phase.
0x07	Early SB initialization during Sec Phase.
0x08	Early NB initialization during Sec Phase.
0x09	End Of Sec Phase.
0x0E	Microcode Not Found.
0x0F	Microcode Not Loaded.
PEI Phase	
0x10	PEI Core
0x11	CPU PEIM

Progress Code	Description
0x15	NB PEIM
0x19	SB PEIM
MRC Progress Codes <i>At this point the MRC Progress Code sequence is executed See Table 81.</i>	
0x31	Memory Installed
0x32	CPU PEIM (CPU Init)
0x33	CPU PEIM (Cache Init)
0x4F	Dxe IPL started
DXE Phase	
0x60	DXE Core started
0x61	DXE NVRAM Init
0x62	DXE Setup Init
0x63	DXE CPU Init
0x65	DXE CPU BSP Select
0x66	DXE CPU AP Init
0x68	DXE PCI Host Bridge Init
0x69	DXE NB Init
0x6A	DXE NB SMM Init
0x70	DXE SB Init
0x71	DXE SB SMM Init
0x72	DXE SB devices Init
0x78	DXE ACPI Init
0x79	DXE CSM Init
0x80	DXE BDS Started
0x81	DXE BDS connect drivers
0x82	DXE PCI Bus begin
0x83	DXE PCI Bus HPC Init
0x84	DXE PCI Bus enumeration
0x85	DXE PCI Bus resource requested
0x86	DXE PCI Bus assign resource

Progress Code	Description
0x87	DXE CON_OUT connect
0x88	DXE CON_IN connect
0x89	DXE SIO Init
0x8A	DXE USB start
0x8B	DXE USB reset
0x8C	DXE USB detect
0x8D	DXE USB enable
0x91	DXE IDE begin
0x92	DXE IDE reset
0x93	DXE IDE detect
0x94	DXE IDE enable
0x95	DXE SCSI begin
0x96	DXE SCSI reset
0x97	DXE SCSI detect
0x98	DXE SCSI enable
0x99	DXE verifying SETUP password
0x9B	DXE SETUP start
0x9C	DXE SETUP input wait
0x9D	DXE Ready to Boot
0x9E	DXE Legacy Boot
0x9F	DXE Exit Boot Services
0xC0	RT Set Virtual Address Map Begin
0xC2	DXE Legacy Option ROM init
0xC3	DXE Reset system
0xC4	DXE USB Hot plug
0xC5	DXE PCI BUS Hot plug
0xC6	DXE NVRAM cleanup
0xC7	DXE ACPI Enable
0x00	Clear POST Code
S3 Resume	
0x40	S3 Resume PEIM (S3 started)

Progress Code	Description
0x41	S3 Resume PEIM (S3 boot script)
0x42	S3 Resume PEIM (S3 Video Repost)
0x43	S3 Resume PEIM (S3 OS wake)
BIOS Recovery	
0x46	PEIM which detected forced Recovery condition
0x47	PEIM which detected User Recovery condition
0x48	Recovery PEIM (Recovery started)
0x49	Recovery PEIM (Capsule found)
0x4A	Recovery PEIM (Capsule loaded)

Appendix E: POST Code Errors

The table below lists the supported POST Error Codes, with a descriptive Error Message text for each. There is also a Response listed, which classifies the error as Minor, Major, or Fatal depending on how serious the error is and what action the system should take.

The Response section in the following table indicates one of these actions:

- **Minor:** The 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 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 on 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.

Table 85. 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
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

Error Code	Error Message	Response
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
8523	DIMM_B1 failed test/initialization	Major
8524	DIMM_B2 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8527	DIMM_C2 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major
852A	DIMM_D2 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852D	DIMM_E2 failed test/initialization	Major
852F	DIMM_F1 failed test/initialization	Major
8530	DIMM_F2 failed test/initialization	Major
8532	DIMM_G1 failed test/initialization	Major
8533	DIMM_G2 failed test/initialization	Major
8535	DIMM_H1 failed test/initialization	Major
8536	DIMM_H2 failed test/initialization	Major
8540	DIMM_A1 disabled	Major
8541	DIMM_A2 disabled	Major
8543	DIMM_B1 disabled	Major
8544	DIMM_B2 disabled	Major
8546	DIMM_C1 disabled	Major

Error Code	Error Message	Response
8547	DIMM_C2 disabled	Major
8549	DIMM_D1 disabled	Major
854A	DIMM_D2 disabled	Major
854C	DIMM_E1 disabled	Major
854D	DIMM_E2 disabled	Major
854F	DIMM_F1 disabled	Major
8550	DIMM_F2 disabled	Major
8552	DIMM_G1 disabled	Major
8553	DIMM_G2 disabled	Major
8555	DIMM_H1 disabled	Major
8556	DIMM_H2 disabled	Major
8560	DIMM_A1 encountered a Serial Presence Detection (SPD) failure	Major
8561	DIMM_A2 encountered a Serial Presence Detection (SPD) failure	Major
8563	DIMM_B1 encountered a Serial Presence Detection (SPD) failure	Major
8564	DIMM_B2 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
8569	DIMM_D1 encountered a Serial Presence Detection (SPD) failure	Major
856A	DIMM_D2 encountered a Serial Presence Detection (SPD) failure	Major
856C	DIMM_E1 encountered a Serial Presence Detection (SPD) failure	Major
856D	DIMM_E2 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
8572	DIMM_G1 encountered a Serial Presence Detection (SPD) failure	Major
8573	DIMM_G2 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
8604	POST Reclaim of non-critical NVRAM variables	Minor
8605	BIOS Settings are corrupted	Major
8606	NVRAM variable space was corrupted and has been reinitialized	Major
8607	Recovery boot has been initiated. Note: The Primary BIOS image may be corrupted or the system may hang during POST. A BIOS update is required.	Fatal
92A3	Serial port component was not detected	Major
92A9	Serial port component encountered a resource conflict error	Major
A000	TPM device not detected.	Minor
A001	TPM device missing or not responding.	Minor
A002	TPM device failure.	Minor
A003	TPM device failed self test.	Minor
A100	BIOS ACM Error	Major
A421	PCI component encountered a SERR error	Fatal

Error Code	Error Message	Response
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 POST Error Beep Codes. Prior to system video initialization, the BIOS uses these beep codes to inform users of error conditions. The beep code is followed by a user visible code displayed on the Diagnostic LEDs.

Table 86. 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 following Beep Codes are from the BMC.			
1-5-2-1	CPU socket population error	N/A	CPU1 socket is empty, or sockets are populated incorrectly – CPU1 must be populated before CPU2.
1-5-2-4	MSID Mismatch	N/A	MSID mismatch occurs if a processor is installed into a system board that has incompatible power capabilities.
1-5-4-2	Power fault	N/A	DC power unexpectedly lost (power good dropout) – Power unit sensors report power unit failure offset.
1-5-4-4	Power control fault	N/A	Power good assertion timeout – Power unit sensors report soft power control failure offset.
1-5-1-2	VR Watchdog Timer	N/A	VR controller DC power on sequence not completed in time.
1-5-1-4	Power Supply Status	N/A	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: Statement of Volatility

This Appendix describes the volatile and non-volatile components on the Intel® Server Board S2600TPR Product Family. It is not the intention of this document to include any components not directly on the listed Intel® Server Boards, such as the chassis components, processors, memory, hard drives, or add-in cards.

Server Board Components

Intel® servers contain several components that can be used to store data. A list of components for the Intel® Server Board S2600TPR is included in the table below. The sections below the table provide additional information about the fields in this table.

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16 MB	U5B2	No (firmware)	Firmware Flash
Non-Volatile	16 MB	U2D2	No (BIOS)	BIOS Flash
Non-Volatile	4 MB	U5M1	No	Connect-IB Flash
Non-Volatile	32 KB	U5A2	No	NIC EEPROM
Volatile	125 MB	U4B1	No	Firmware SDRAM

Component Type

Three types of components are on an Intel® Server Board. These types are:

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

Size

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

Board Location

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

User Data

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

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

- BIOS: The server board BIOS provides the capability to prevent unauthorized users from configuring BIOS settings when a BIOS password is set. This password is stored in BIOS flash, and is only used to set BIOS configuration access restrictions.
- BMC: The server boards support an Intelligent Platform Management Interface (IPMI) 2.0 conformant baseboard management controller (BMC). The BMC provides health monitoring, alerting and remote power control capabilities for the Intel® Server Board. The BMC does not have access to operating system level data.

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

Glossary

This glossary contains important terms used in the preceding chapters. For ease of use, numeric entries are listed first (for example, 82460GX) with alpha entries following (for example, AGP 4x). Acronyms are then entered in their respective place, with non-acronyms following.

Table 87. Glossary

Term	Definition
ACPI	Advanced Configuration and Power Interface
AP	Application Processor
ARP	Address Resolution Protocol
BIOS	Basic Input/Output System
BIST	Built-In Self Test
BMC	Baseboard Management Controller
Bridge	Circuitry connecting one computer bus to another, allowing an agent on one to access the other
BSP	Bootstrap Processor
Byte	8-bit quantity.
CATERR	On a catastrophic hardware event the core signals CATERR to the uncore. The core enters a halted state that can only be exited by a reset.
CMOS	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.
DCMI	Data Center Management Interface
DHCP	Dynamic Host Configuration Protocol
DIMM	Dual In-line Memory Module
EEPROM	Electrically Erasable Programmable Read-Only Memory
EPS	External Product Specification
FRB	Fault Resilient Booting
FRU	Field Replaceable Unit
GB	1024 MB
GPIO	General Purpose I/O
HSC	Hot-Swap Controller
Hz	Hertz (1 cycle/second)
I ² C	Inter-Integrated Circuit Bus
IA	Intel® Architecture
ILM	Independent Loading Mechanism
IMC	Integrated Memory Controller
IMR	Intel MegaRAID®
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IT	Initiator Target

Term	Definition
KB	1024 bytes
LAN	Local Area Network
LED	Light Emitting Diode
LSB	Least Significant Bit
LUN	Logical Unit Number
MAC	Media Access Control
MB	1024KB
ME	Management Engine
ms	Milliseconds
MSB	Most Significant Bit
NIC	Network Interface Controller
NMI	Nonmaskable Interrupt
NTB	Non-Transparent Bridge
OEM	Original Equipment Manufacturer
PECI	Platform Environment Control Interface
PEF	Platform Event Filtering
POST	Power-On Self Test
PWM	Pulse-Width Modulation
QPI	QuickPath Interconnect
QSFP+	Quad Small Form-factor Pluggable Plus
RAM	Random Access Memory
ROM	Read Only Memory
RTC	Real-Time Clock (Component of ICH peripheral chip on the server board)
RMM4	Remote Management Module 4
SDR	Sensor Data Record
EEPROM	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 nonmaskable interrupt)
SMM	Server Management Mode
SMS	Server Management Software
SNMP	Simple Network Management Protocol
TDP	Thermal Design Power
TIM	Thermal Interface Material
UART	Universal Asynchronous Receiver/Transmitter
VLSI	Very Large Scale Integration
VRD	Voltage Regulator Down
VT	Virtualization Technology

Reference Documents

- *Intel® Server Board S2600TPR Product Family and Intel® Compute Module HNS2600TPR Product Family Service Guide*
- *Intel® Server System BIOS External Product Specification for Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3/v4 product family*
- *Intel® Server System BMC Firmware External Product Specification for Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3/v4 product family*
- *Intel® Remote Management Module 4 Technical Product Specification*
- *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*
- *Intel® Xeon® Processor E5-4600/2600/2400/1600 v3/v4 Product Families External Design Specification*
- *Intel® Chipset C610 product family External Design Specification*
- *Intel® Ethernet Controller I350 Family Product Brief*
- *SmaRT & CLST Architecture on Intel Systems and Power Supplies Specification*
- *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.*
- *Alert Standard Format (ASF) Specification, Version 2.0, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., <http://www.dmtf.org>.*

Mouser Electronics

Authorized Distributor

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

[Intel:](#)

[BBS2600TPFR](#) [HNS2600TPR](#) [BBS2600TPR](#) [HNS2600TPFR](#)



**Стандарт
Электрон
Связь**

Мы молодая и активно развивающаяся компания в области поставок электронных компонентов. Мы поставляем электронные компоненты отечественного и импортного производства напрямую от производителей и с крупнейших складов мира.

Благодаря сотрудничеству с мировыми поставщиками мы осуществляем комплексные и плановые поставки широчайшего спектра электронных компонентов.

Собственная эффективная логистика и склад в обеспечивает надежную поставку продукции в точно указанные сроки по всей России.

Мы осуществляем техническую поддержку нашим клиентам и предпродажную проверку качества продукции. На все поставляемые продукты мы предоставляем гарантию .

Осуществляем поставки продукции под контролем ВП МО РФ на предприятия военно-промышленного комплекса России , а также работаем в рамках 275 ФЗ с открытием отдельных счетов в уполномоченном банке. Система менеджмента качества компании соответствует требованиям ГОСТ ISO 9001.

Минимальные сроки поставки, гибкие цены, неограниченный ассортимент и индивидуальный подход к клиентам являются основой для выстраивания долгосрочного и эффективного сотрудничества с предприятиями радиоэлектронной промышленности, предприятиями ВПК и научно-исследовательскими институтами России.

С нами вы становитесь еще успешнее!

Наши контакты:

Телефон: +7 812 627 14 35

Электронная почта: sales@st-electron.ru

Адрес: 198099, Санкт-Петербург,
Промышленная ул, дом № 19, литера Н,
помещение 100-Н Офис 331