Intel® Server Board S2600CW Family

Technical Product Specification

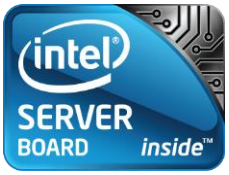

Revision1.3

September, 2015

Intel® Server Boards and Systems

Revision History

Date	Revision Number	Modifications
September, 2014	1.0	Initial release
November, 2014	1.10	Include information for S2600CW2S, S2600CWT and S2600CWTS
February, 2015	1.11	Updated chapter 3.4.1, 3.4.7 and 7.5
July, 2015	1.12	Updated chapter 5.3 Sensor Monitoring.
August, 2015	1.2	Updated Post Progress Codes and MRC Progress Codes
September, 2015	1.3	Updated Memory support Table3~4

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.

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.

Intel may make changes to specifications and product descriptions at any time, without notice. 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 information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document 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.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

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 is a registered trademark of Intel Corporation.

*Other names and brands are the property of their respective owners.

Table of Contents

List of Figures

List of Tables

< This page intentionally left blank. >

1. Introduction

This Technical Product Specification (TPS) provides information on the Intel® Server Board S2600CW including architecture, features, and functionality.

In addition, you can obtain design-level information for a given subsystem by ordering the External Product Specifications (EPS) for the specific subsystem. EPS documents are not publicly available and you must order them through your local Intel representative.

1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Intel® Server Board S2600CW Overview
- Chapter 3 – Intel® Server Board S2600CW Functional Architecture
- Chapter 4 – System Security
- Chapter 5 – Intel® Server Board S2600CW Platform Management
- Chapter 6 – Intel® Intelligent Power Node Manager (NM) Support Overview
- Chapter 7 – Intel® Server Board S2600CW Connector/Header Locations and Pin-outs
- Chapter 8 – Intel® Server Board S2600CW Jumper Blocks
- Chapter 9 – Intel® Light Guided Diagnostics
- Chapter 10 – Power Supply Specification Guidelines
- Chapter 11 – Design and Environmental Specifications
- Appendix A: Integration and Usage Tips
- Appendix B: Compatible Intel® Server Chassis
- Appendix C: BMC Sensor Tables
- Appendix D: Platform Specific BMC Appendix
- Appendix E: POST Code Diagnostic LED Decoder
- Appendix F: POST Error Code
- Glossary
- Reference Documents

1.2 Server Board Use Disclaimer

Intel® Server Boards contain a number of high-density VLSI (Very Large Scale Integration) and power delivery components that require adequate airflow for cooling. Intel ensures through its own chassis development and testing that when Intel server building blocks are used together, the fully integrated system meets 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 the published operating or non-operating limits.

2. Intel® Server Board S2600CW Overview

The Intel® Server Board S2600CW is a monolithic printed circuit board (PCB) with features designed to support the pedestal server markets. This server board is designed to support the Intel® Xeon® processor E5-2600 v3 product family. Previous generation Intel® Xeon® processors are not supported.

The Intel® Server Board S2600CW family includes different board configurations:

- Intel® Server Board S2600CW2: dual 1GbE NIC ports (I350)
- Intel® Server Board S2600CW2S: dual 1GbE NIC ports (I350) and onboard SAS controller
- Intel® Server Board S2600CWT: dual 10GbE NIC ports (X540)
- Intel® Server Board S2600CWTS: dual 10GbE NIC ports (X540) and onboard SAS controller

2.1 Intel® Server Board S2600CW Feature Set

Table 1. Intel® Server Board S2600CW Feature Set

Feature	Description
Processors	▪ Two LGA2011-3 (Socket R3) processor sockets ▪ Support for one or two Intel® Xeon® processors E5-2600 v3 product family ▪ Maximum supported Thermal Design Power (TDP) of up to 145 W
Memory	▪ Eight memory channels (four channels for each processor socket) ▪ Two DIMM slots for each channel ▪ Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) ▪ DDR4 Memory data transfer rates: 1333, 1600, 1866, and 2133 MT/s
Chipset	Intel® C612 chipset
Cooling Fan Support	▪ Two processor fans (4-pin headers) ▪ Six front system fans (6-pin headers) ▪ One rear system fan (4-pin header)
Add-in Card Slots	▪ Support up to six expansion slots ▪ From PCH: – Slot 1: PCIe Gen II x4 ▪ From the first processor: – Slot 5: PCIe Gen III x16 connector. Electrical x16 for S2600CW2 or S2600CWT, electrical x8 for S2600CW2S or S2600CWTS – Slot 6 PCIe Gen III x16 electrical with x16 physical connector ▪ From the second processor: – Slot 2: PCIe Gen III x16 electrical with x16 physical connector – Slot 3: PCIe Gen III x8 electrical with x8 physical connector – Slot 4: PCIe Gen III x16 electrical with x16 physical connector

Feature	Description
RAID Support	▪ PCH SATA 6G – ESRT2 RAID 0/1/10 – Optional RAID 5 is supported through the ESRT2 RAID5 upgrade key ▪ LSI* SAS3008 SAS 12G – Integrated RAID 0/1/1E/10 – Integrated MegaRaid 0/1/10/5/50 upgrade is supported through the upgrade key AXXRPFKHY5
External back panel I/O Connectors	▪ One DB-15 video connector ▪ Two NIC ports (I350 for 1GbE or X540 for 10GbE) ▪ One NIC port dedicated to server management (DMN) ▪ Two USB 3.0 ports and two USB 2.0 ports
Internal I/O Connectors/Headers	▪ Two 7-pin SATA 6G ports ▪ Two mini-SAS HD connectors supporting eight SATA 6Gbps transfer rate ▪ Two mini-SAS HD connectors supporting eight SAS 12Gbps transfer rate (S2600CW2S and S2600CWTS only) ▪ One 2x10-pin connector providing front panel support for two USB 3.0 ports ▪ One internal Type-A USB 2.0 port ▪ One internal USB port to support low profile eUSB SSD ▪ One DH-10 Serial Port B connector ▪ One 24-pin SSI-EEB compliant front panel header ▪ One TPM connector ▪ One M.2/NGFF connector ▪ One RMM4 LITE connector
Video Support	▪ Integrated 2D Video Controller ▪ 16 MB DDR3 Memory
Server Management	▪ Support for Intel® Remote Management Module 4 solutions ▪ Intel® Light-Guided Diagnostics on field replaceable units ▪ Support for Intel® System Management Software ▪ Support for Intel® Intelligent Power Node Manager (PMBus*-compliant power supply needed)
Security	Intel® TPM – AXXTPME5 (Accessory Option)
Form Factor	SSI EEB (12”x13”)
Compatible Intel® Server Chassis	Intel® Server Chassis P4304XXMFEN2 and P4304XXMUXX

2.2 Server Board Layout

The following diagram shows the board layout for S2600CW2S and S2600CWTS with on-board SAS controller.

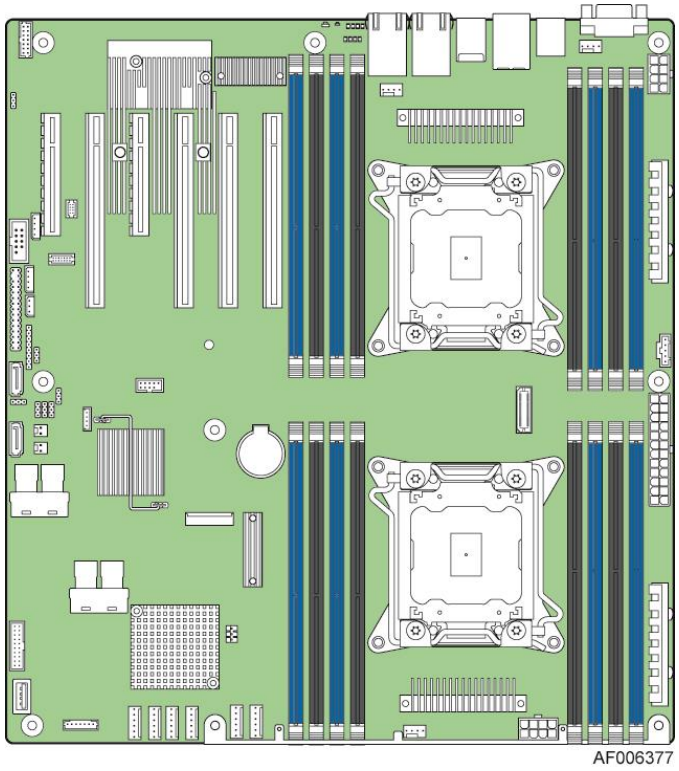

Figure 1. Intel® Server Board S2600CW2S and S2600CWTS

The following diagram shows the board layout for S2600CW2 and S2600CWT without on-board SAS controller.

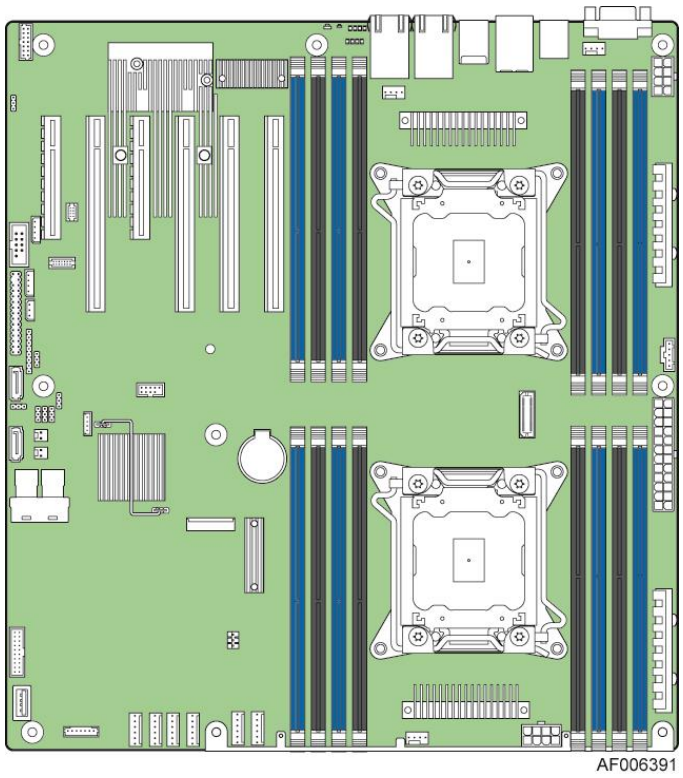

Figure 2. Intel® Server Board S2600CW2 and S2600CWT

2.2.1 Server Board Connector and Component Layout

The following figure shows the layout of the server board and the location of each connector and major component except jumpers. The locations of jumpers can be found in Chapter 8.

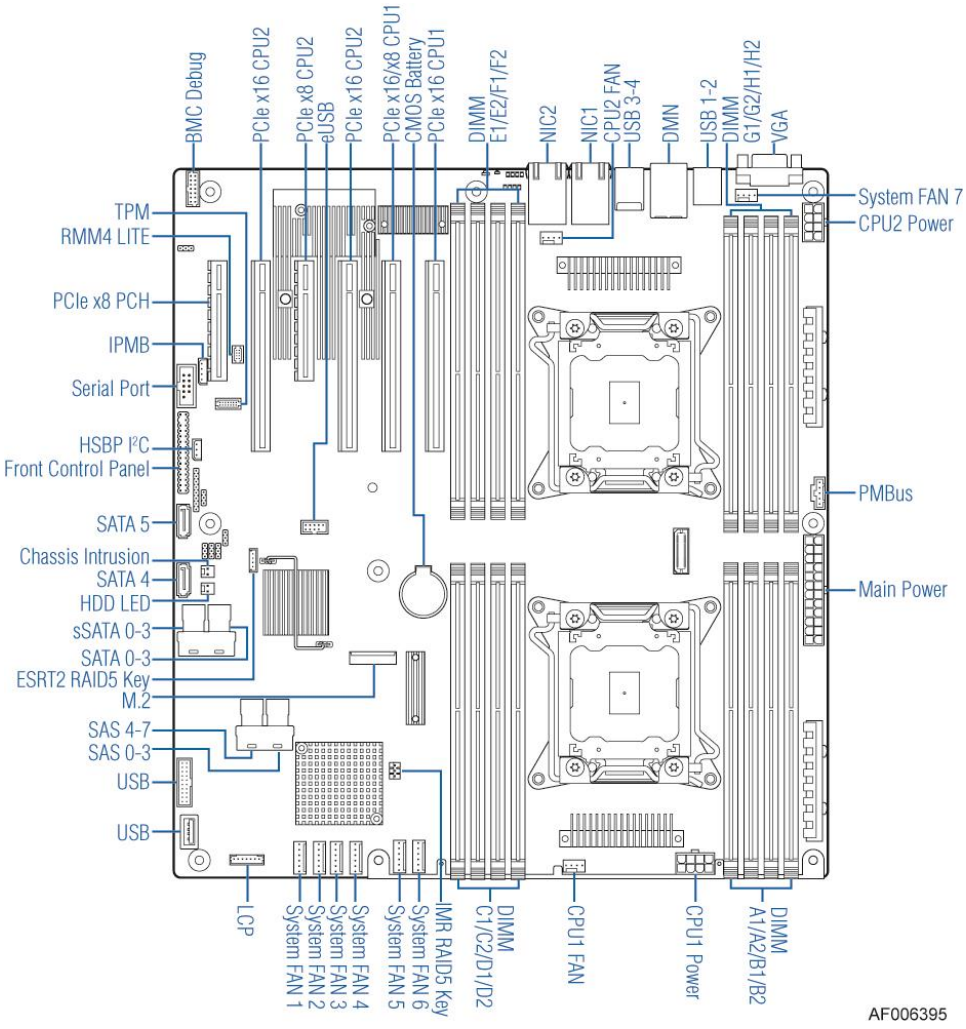

Figure 3. Connector and Component Layout

2.2.2 Server Board Mechanical Drawings

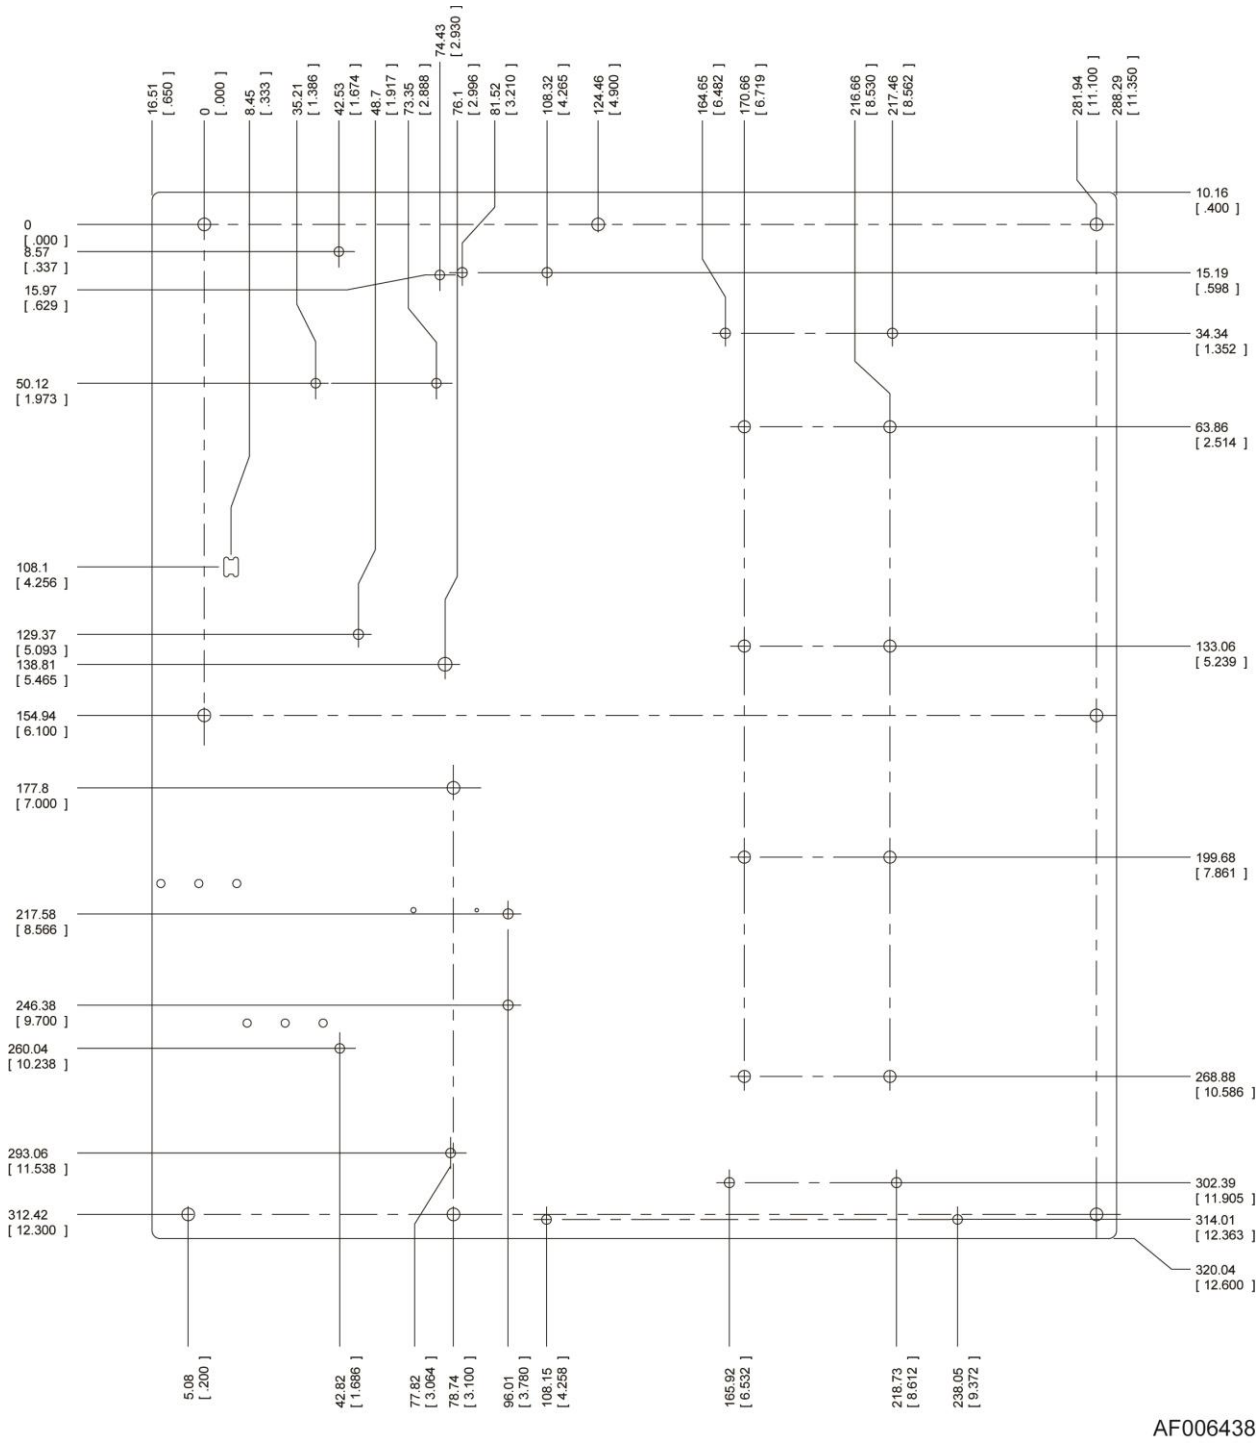

Figure 4. Mounting Hole Locations (1 of 2)

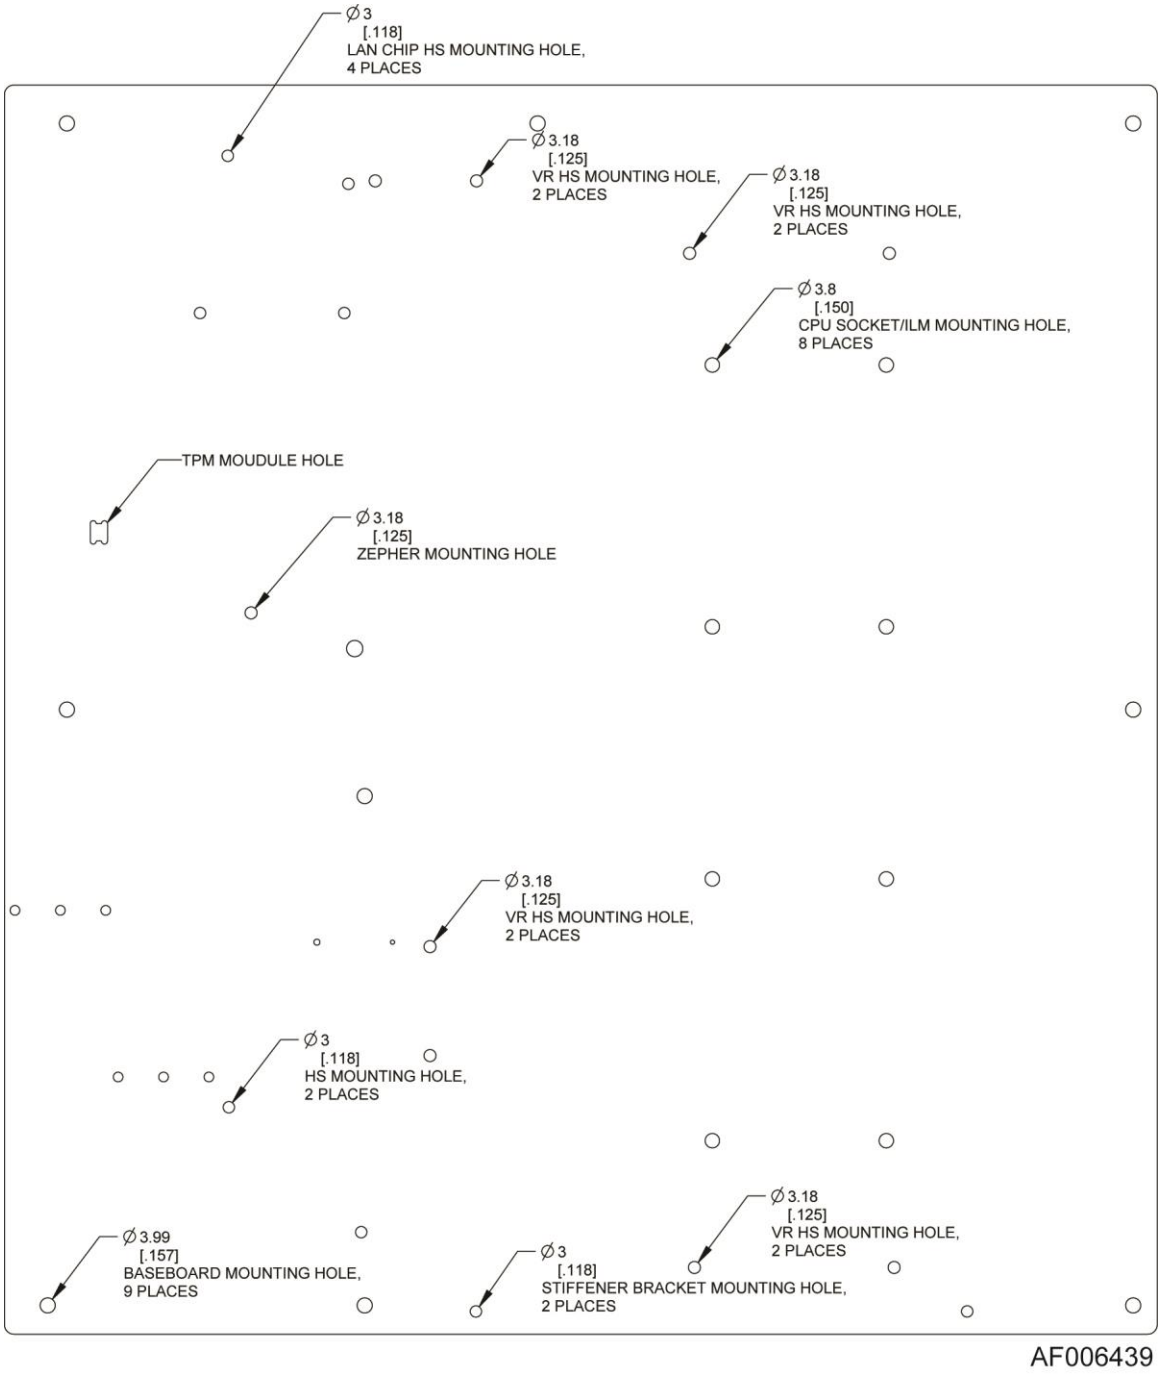

Figure 5. Mounting Hole Locations (2 of 2)

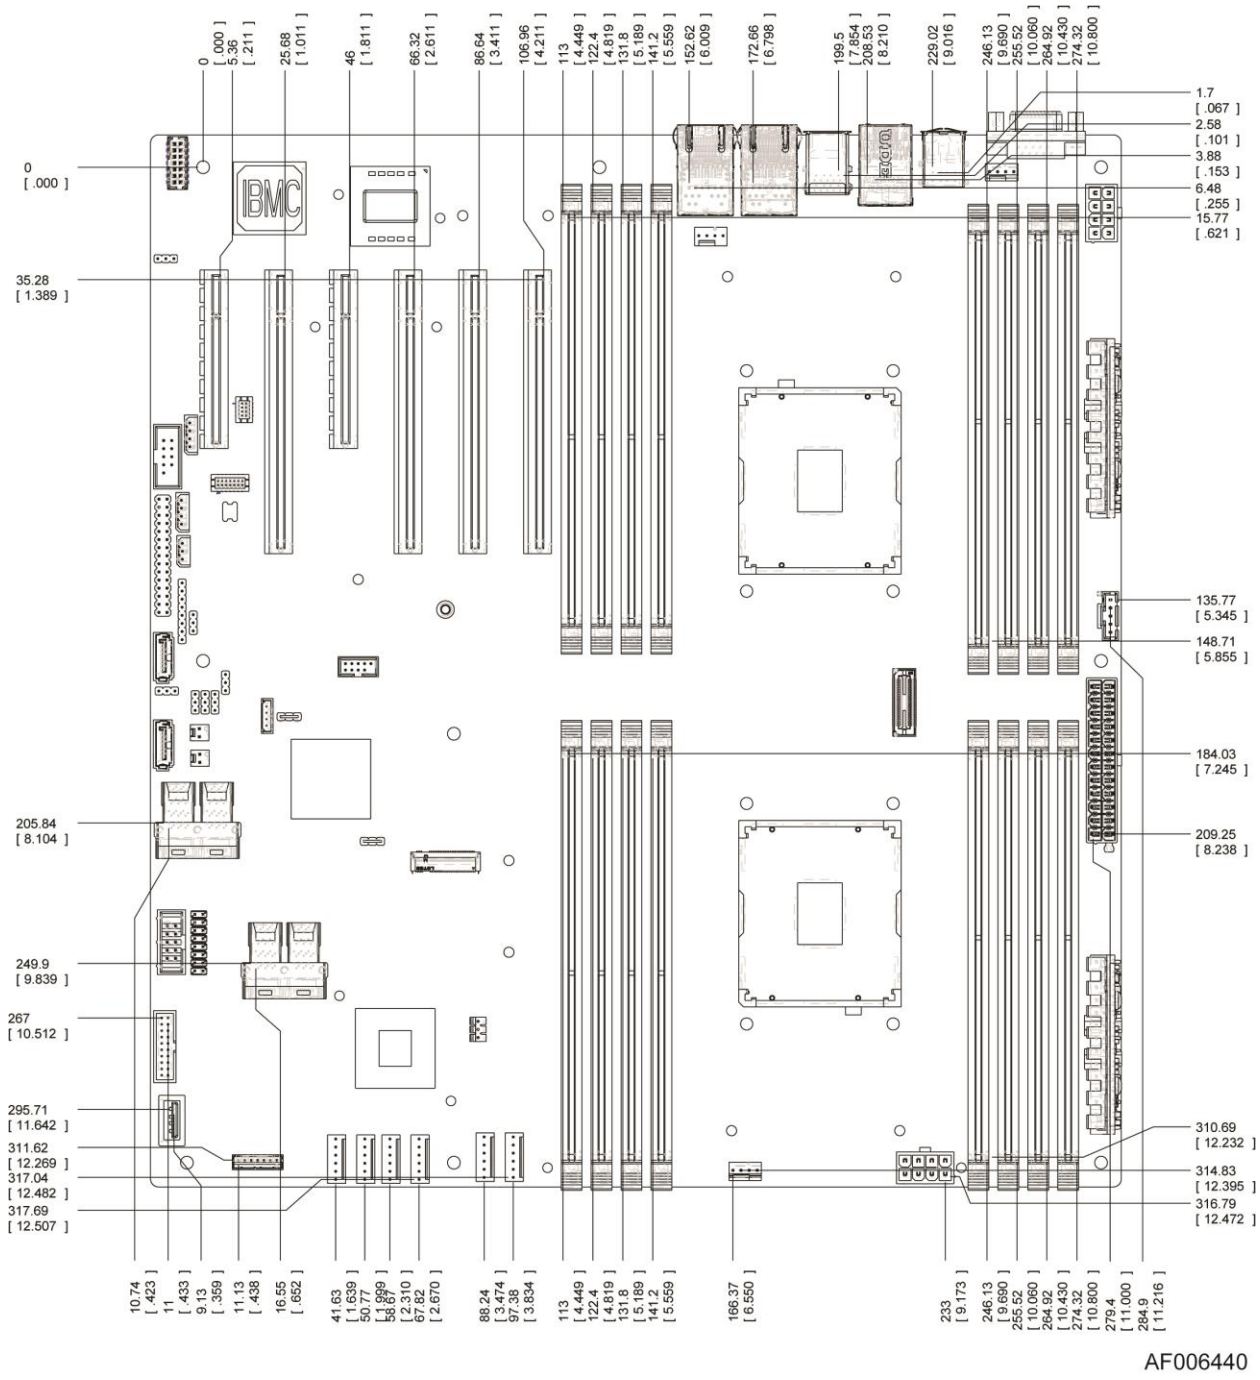

Figure 6. Major Connector Pin-1 Locations (1 of 2)

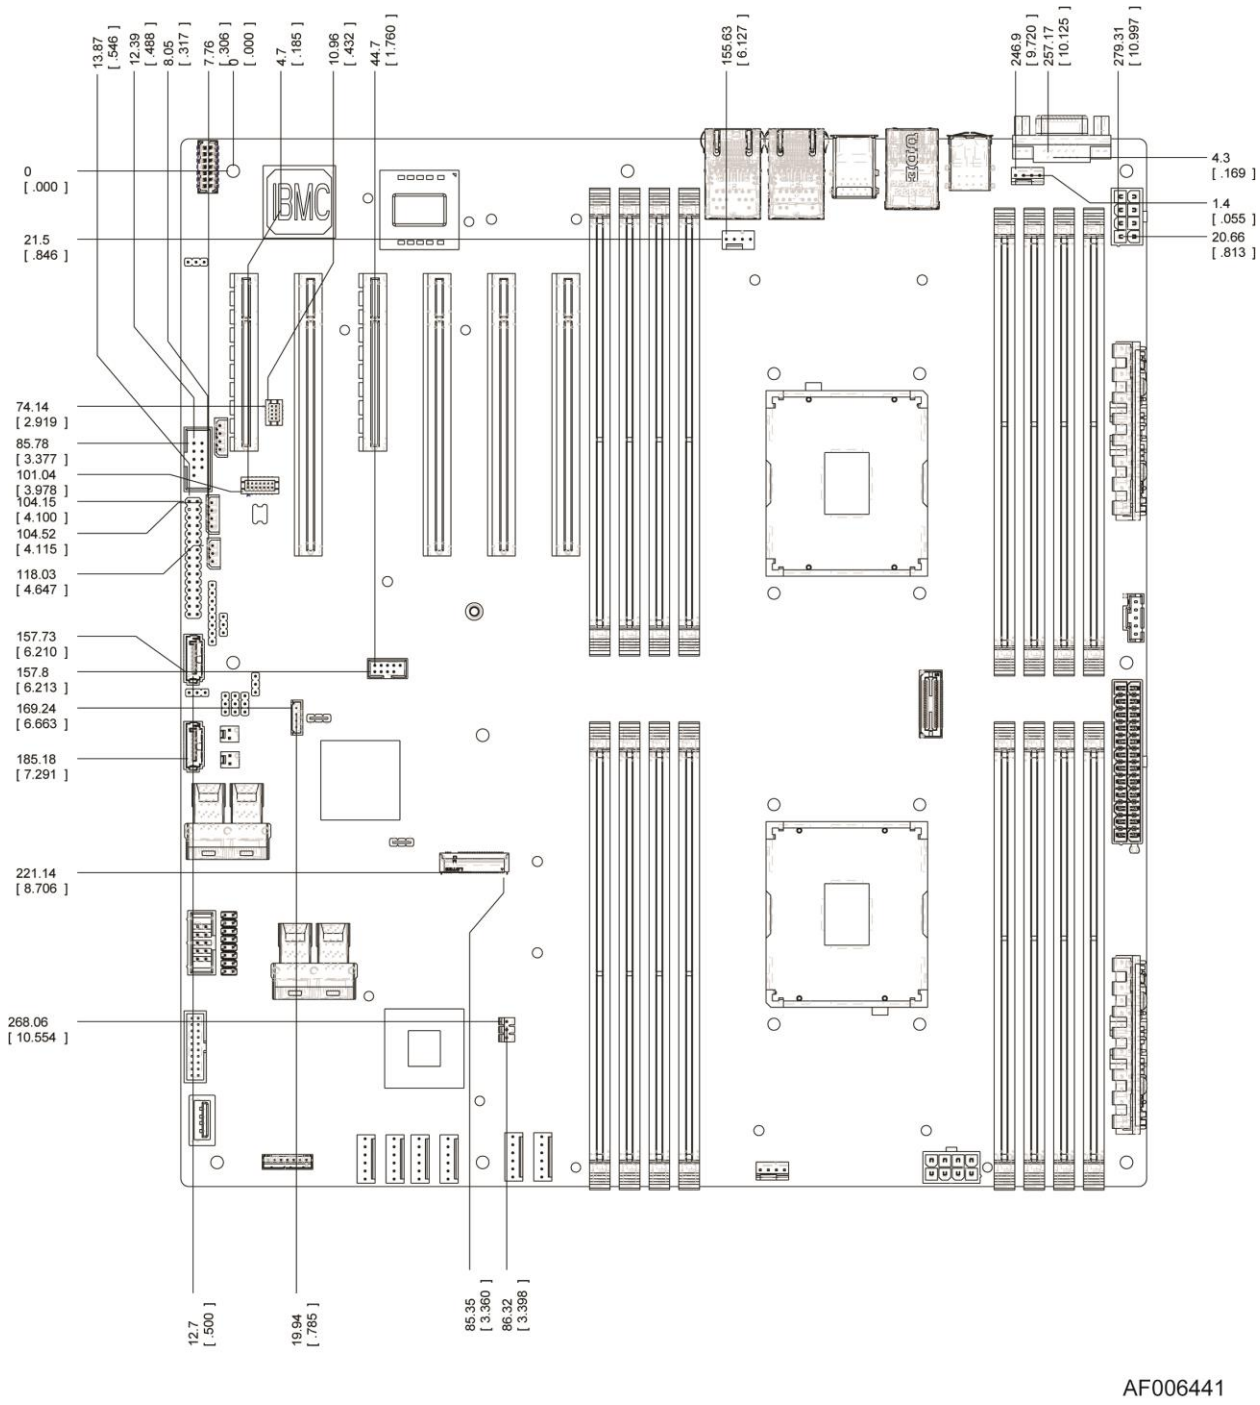

Figure 7. Major Connector Pin-1 Locations (2 of 2)

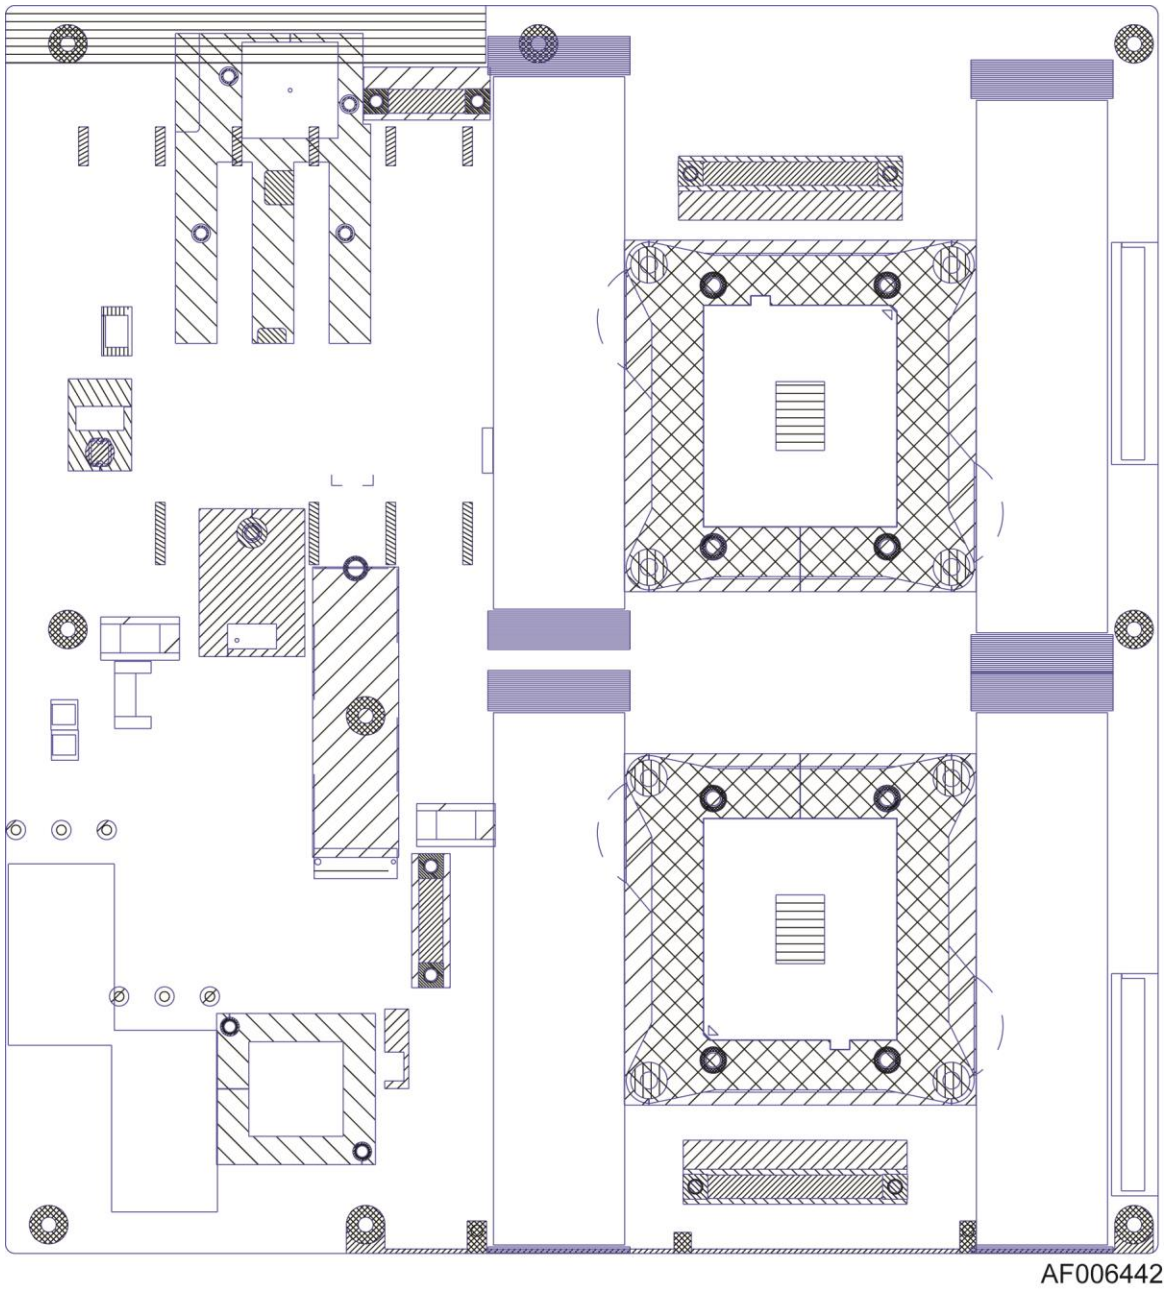

Figure 8. Primary Side Card-Side Keep-out Zone

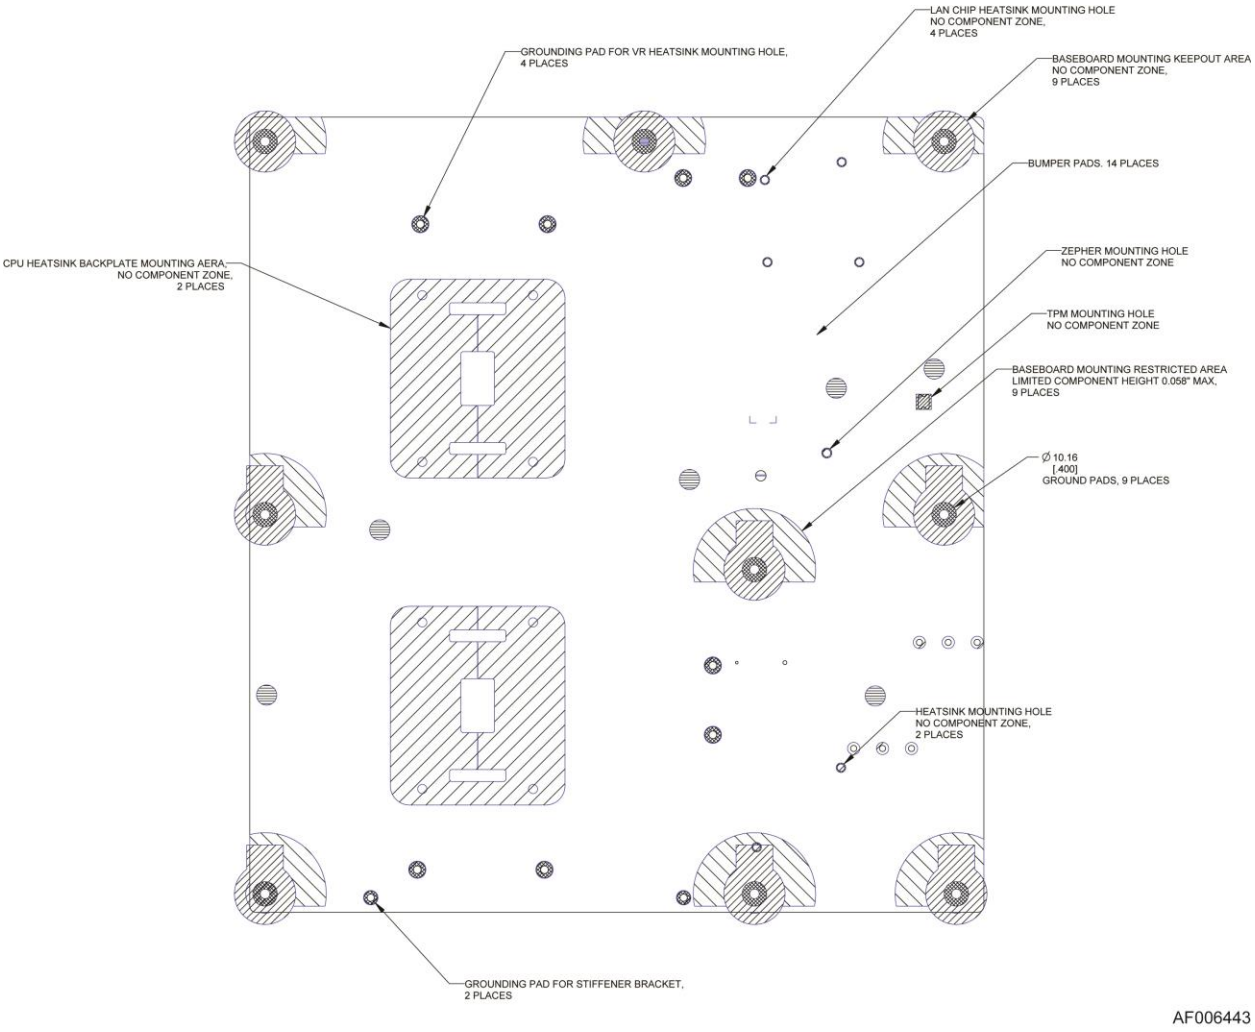

Figure 9. Second Side Keep-out Zone

2.2.3 Server Board Rear I/O Layout

The following drawing shows the layout of the rear I/O components for the server boards.

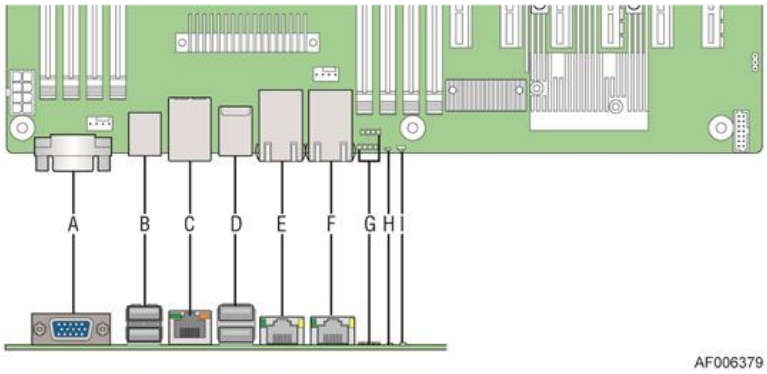

Callout	Description	Callout	Description
A	Video	B	USB 2.0
C	Dedicated Management NIC (DMN)	D	USB 3.0
E	NIC1	F	NIC2
G	Diagnostic LEDs	H	ID LED
I	System Status LED		

Figure 10. Rear I/O Layout of Intel® Server Board S2600CW

3. Intel® Server Board S2600CW Functional Architecture

The architecture and design of the Intel® Server Board S2600CW is based on the Intel® Xeon® E5-2600 v3 processors and the Intel® C612 chipset. This chapter provides a high-level description of the functionality associated with each chipset component and the architectural blocks that make up the server boards.

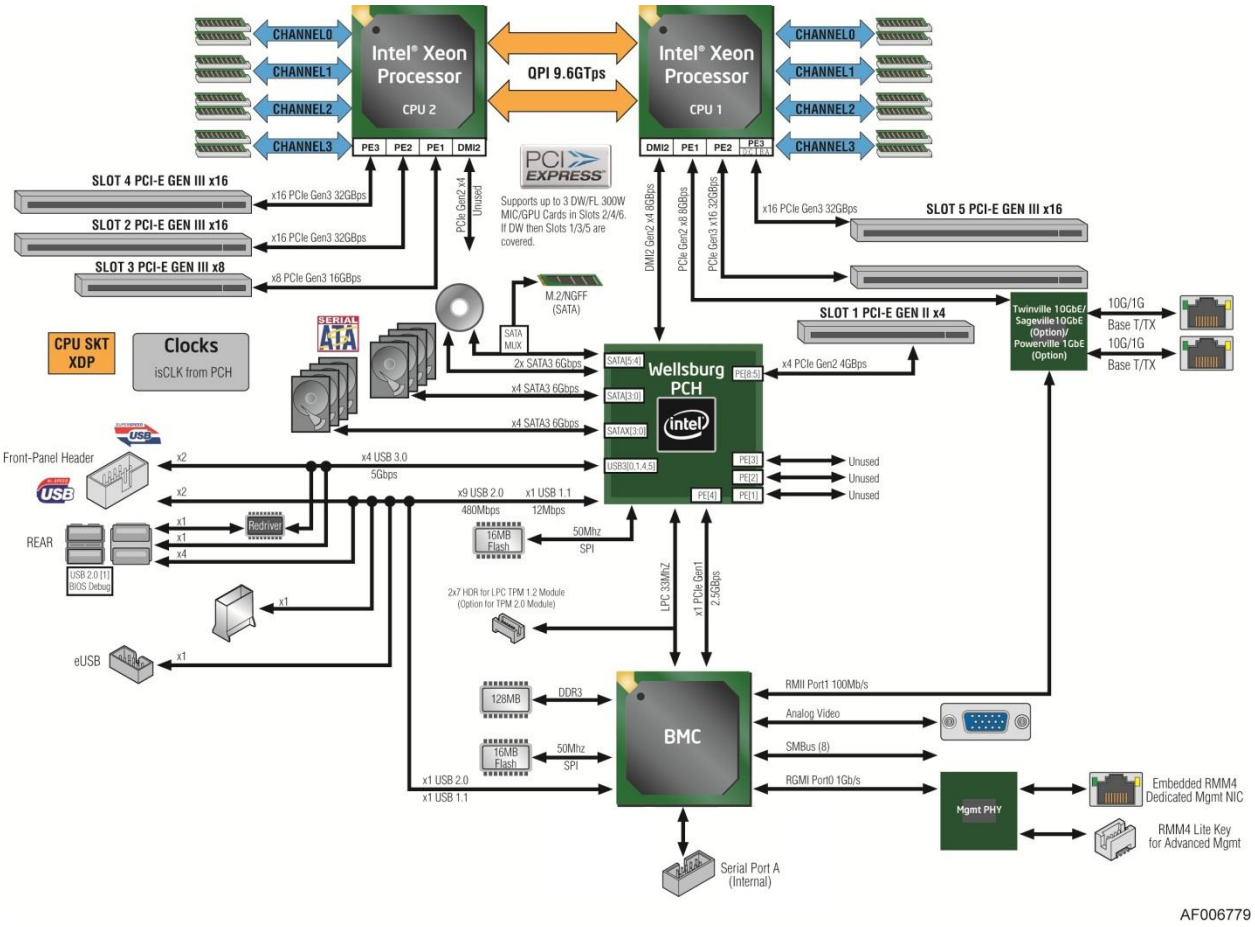

Figure 11. Intel® Server Board S2600CW2/S2600CWT Functional Block Diagram

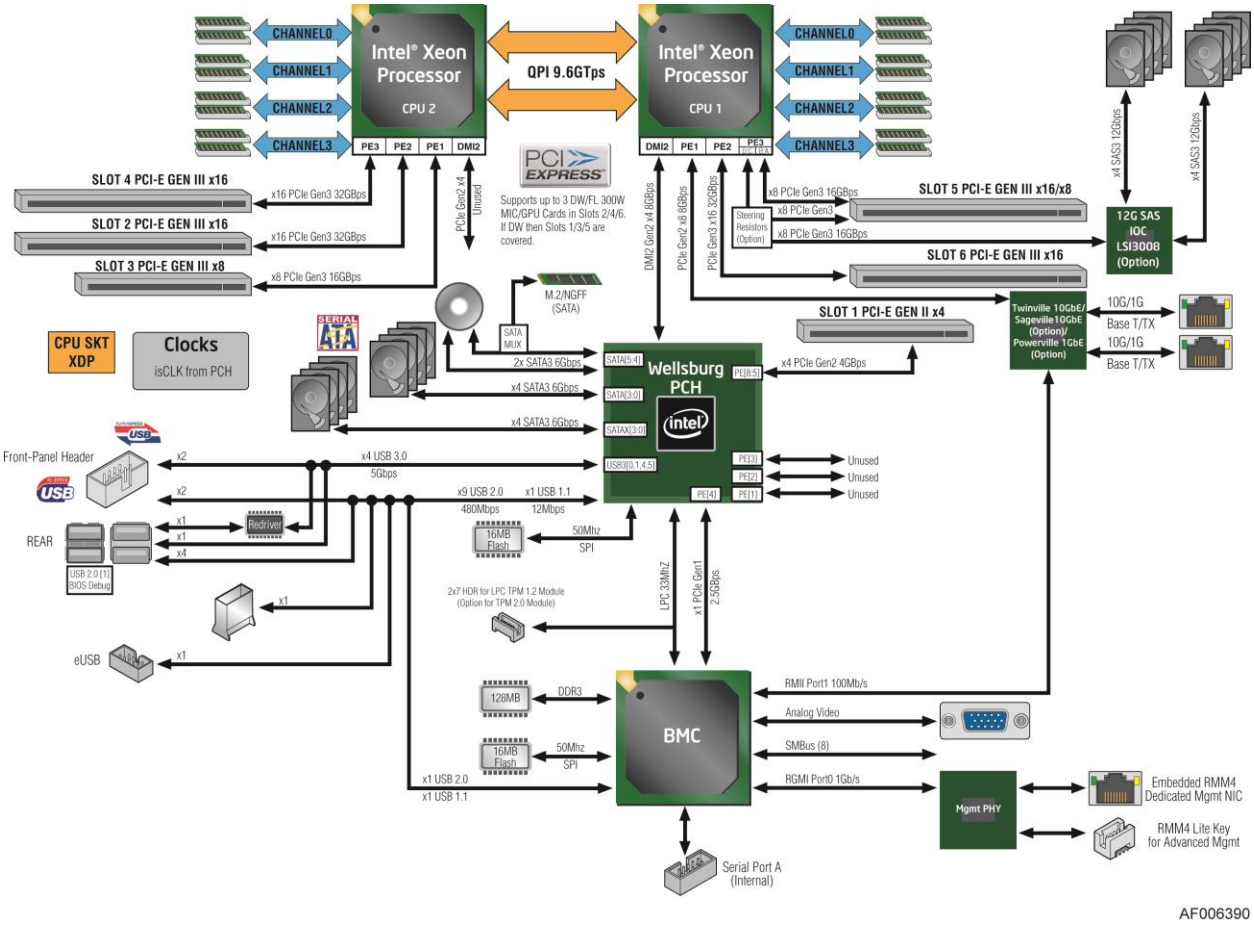

Figure 12. Intel® Server Board S2600CW2S/S2600CWTS Functional Block Diagram

3.1 Processor Support

The server board includes two Socket-R3 (LGA2011-3) processor sockets and can support two processors from the Intel® Xeon® processor E5-2600 v3 product family with a Thermal Design Power (TDP) of up to 145W.

Previous generation Intel® Xeon® processors are not supported on the Intel® Server Boards described in this document.

Visit the Intel website for a complete list of supported processors.

3.1.1 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 illustration below identifies each sub-assembly component.

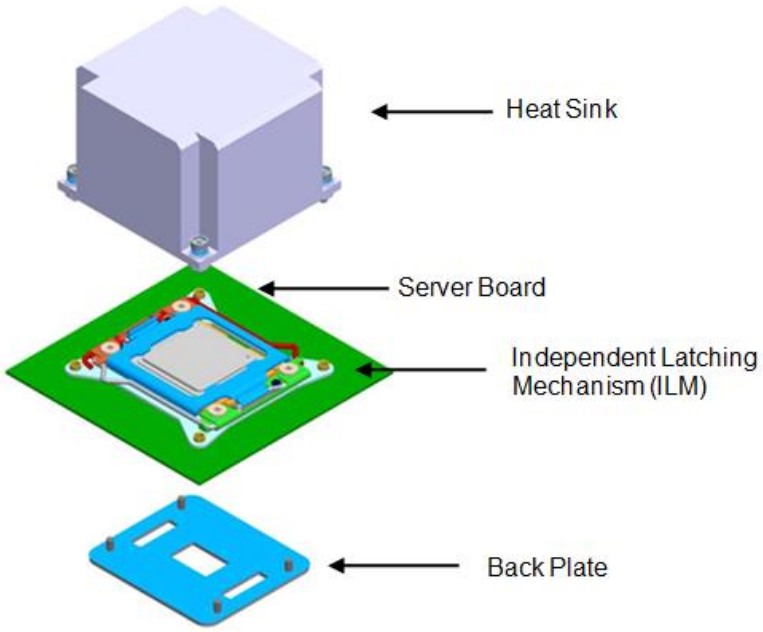

Figure 13. Processor Socket Assembly

3.1.2 Processor Population Rules

***Note**: Although the server board does support dual-processor configurations consisting of different processors that meet the defined criteria below, Intel does not perform validation testing of these configurations. 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 CPU1.

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 QPI link frequencies may operate together if they are otherwise compatible and if a common link frequency can be selected. The common link frequency will 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.

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

- **Fatal**: If the system can boot, it pauses 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", regardless of whether the "Post Error Pause" setup option is enabled or disabled.

  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 System Event Log (SEL) with the POST Error Code.

  The system cannot boot unless the error is resolved. The user needs to replace the faulty part and restart the system.

  For Fatal Errors during processor initialization, the System Status LED will be set to a steady Amber color, indicating an unrecoverable system failure condition.
- **Major**: If the "Post Error Pause" setup option is enabled, the system goes directly to the Error Manager to display the error, and logs the POST Error Code to the SEL. Operator intervention is required to continue booting the system.

  Otherwise, if "POST Error Pause" is disabled, the system continues to boot and no prompt is given for the error, although the Post Error Code is logged to the Error Manager and in a SEL message.
- **Minor**: The message is displayed on the screen or on the Error Manager screen, 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 2. Mixed Processor Configurations

Error	Severity	System Action
Processor family not Identical	Fatal	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the System Event Log (SEL). ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays **0194: Processor family mismatch detected** message in the Error Manager. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.

Error	Severity	System Action
Processor model not Identical	Fatal	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays **0196: Processor model mismatch detected** message in the Error Manager. ▪ Takes Fatal Error action (see above) and 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: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays **0191: Processor core/thread count mismatch detected** message in the Error Manager. ▪ 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: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays **0192: Processor cache size mismatch detected** message in the Error Manager. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor frequency (speed) not identical	Fatal	The BIOS detects the processor frequency difference, and responds as follows: ▪ Adjusts all processor frequencies to the highest common frequency. ▪ No error is generated – **this is not an error condition.** ▪ Continues to boot the system successfully. If the frequencies for all processors **cannot be adjusted to be the same,** then this **is** an error, and the BIOS responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Does not disable the processor. ▪ Displays **0197: Processor speeds unable to synchronize** message in the Error Manager. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.

Error	Severity	System Action
Processor Intel® QuickPath Interconnect link frequencies not identical	Fatal	The BIOS detects the QPI link frequencies and responds as follows: ▪ Adjusts all QPI interconnect link frequencies to the highest common frequency. ▪ No error is generated – **this is not an error condition.** ▪ Continues to boot the system successfully. If the link frequencies for all QPI links **cannot be adjusted to be the same,** then this **is** an error, and the BIOS responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ 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	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Displays **818x: Processor 0x microcode update not found** message in the Error Manager or on the screen. ▪ The system continues to boot in a degraded state, regardless of the setting of POST Error Pause in the Setup.
Processor microcode update failed	Major	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Displays **816x: Processor 0x unable to apply microcode update** message in the Error Manager or on the screen. ▪ Takes Major Error action. The system may continue to boot in a degraded state, depending on the setting of POST Error Pause in Setup, or may halt with the POST Error Code in the Error Manager waiting for operator intervention.

3.2 Processor Functions Overview

The Intel® Xeon® processor E5-2600 v3 product family combines several key system components into a single processor package, including the CPU cores, Integrated Memory Controller (IMC), and Integrated IO Module (IIO). In addition, each processor package includes two Intel® QuickPath Interconnect point-to-point links capable of up to 9.6 GT/s, up to 40 lanes of Gen 3 PCI Express* links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* Gen 2 interface with a peak transfer rate of 5.0 GT/s. The processor supports up to 46 bits of physical address space and 48 bits of virtual address space.

The following sections 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 product family documents listed in the Reference Documents list.

Processor Feature Details:

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

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
- Trusted Platform Module (TPM) 1.2

3.2.1 Intel® Virtualization Technology (Intel® VT) for Intel® 64 and IA-32 Intel® Architecture (Intel® VT-x)

Intel® VT provides hardware assist to the virtualization software, reducing its size, cost, and complexity. Special attention is also given to reduce the virtualization overheads occurring in cache, I/O, and memory. Intel® VT-x specifications and functional descriptions are included in the Intel® 64 and IA-32 Architectures Software Developer's Manual.

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

A general requirement for I/O virtualization models is the ability to isolate and restrict device accesses to the resources owned by the partition managing the device. Intel® VT for Directed I/O provides VMM software with the following capabilities:

- I/O device assignment: for flexibly assigning I/O devices to VMs and extending the protection and isolation properties of VMs for I/O operations.
- DMA remapping: for supporting address translations for Direct Memory Accesses (DMA) from devices.
- Interrupt remapping: for supporting isolation and routing of interrupts from devices and external interrupt controllers to appropriate VMs.
- Interrupt posting: for supporting direct delivery of virtual interrupts from devices and external interrupt controllers to virtual processors.
- Reliability: for recording and reporting of DMA and interrupt errors to system software that may otherwise corrupt memory or impact VM isolation.

Intel® VT-d specifications and functional descriptions are included in the Intel® Virtualization Technology for Directed I/O application documents.

3.2.3 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 so that those wishing to rely on the platform can make an appropriate trust decision. The Intel® TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

3.2.4 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.

3.2.5 Advanced Encryption Standard (AES)

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

3.2.6 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.

3.2.7 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.

3.2.8 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.

3.2.9 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.

3.2.10 Intel® Node Manager 3.0

Intel® Node Manager 3.0 enables the PTAS-CUPS (Power Thermal Aware Scheduling – Compute Usage Per Second) feature of the Intel® Server Platform Services 3.0 Intel® ME FW. This is 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 PTAS-CUPS data can then be used in conjunction with the Intel® Server Platform Services 3.0 Intel® Node Manager power monitoring/controls and a remote management application (such as the Intel® Data Center Manager [Intel® DCM]) to create a dynamic, automated, closed-loop data center management and monitoring system.

3.2.11 Intel® Secure Key

It is 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. Please see more details from Intel® Digital Random Number Generator Software Implementation Guide. It is intended to provide a complete source of technical information on the RDRAND Instruction usage, including code examples.

3.2.12 Intel® OS Guard

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.

3.2.13 Trusted Platform Module (TPM)

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

3.3 Integrated Memory Controller (IMC) and Memory Subsystem

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

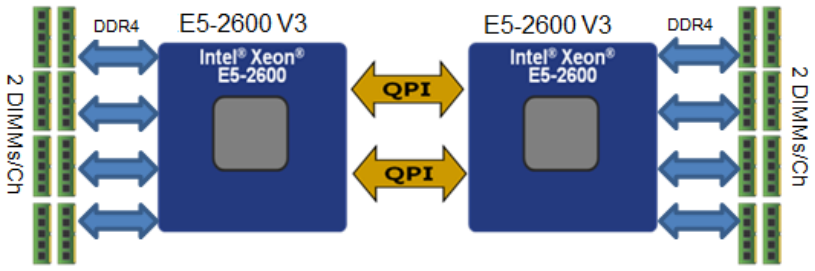

Figure 14. Memory Subsystem for Intel® Server Board S2600CW

Each installed processor includes an integrated memory controller (IMC). Each processor supports 4 memory channels capable of supporting up to 2 DIMMs per channel. The processor IMC supports the following:

- DDR4 ECC RDIMM
- DDR4 ECC LRDIMM
- Support for 4 Gb and 8 Gb DRAM Technologies
- Max Ranks per DDR channel – DDR4 LRDIMM: 16 (SR, DR, QR, 8R)
- Max Ranks per DDR channel – DDR4 RDIMM: 4 (SR, DR)
- IMCs can operate in either Independent mode (Maximum Performance mode) or Lockstep mode (RASM)
- Memory RASM support:
  - DRAM Single Device Data Correction (SDDC)
  - 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

3.3.1 Supported Memory

Table 3. RDIMM Support

Ranks Per DIMM and Data Width	Memory Capacity Per DIMM		Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)	
			1DPC	2DPC
			1.2V	1.2V
SRx4	8GB	16GB	2133MT/s	1866 MT/s
SRx8	4GB	8GB	2133 MT/s	1866 MT/s
DRx8	8GB	16GB	2133 MT/s	1866 MT/s
DRx4	16GB	32GB	2133 MT/s	1866 MT/s

Table 4. LRDIMM Support

Ranks Per DIMM and Data Width	Memory Capacity Per DIMM		Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)	
			1DPC	2DPC
			1.2V	1.2V
QRx4	32GB	64GB	2133 MT/s	2133 MT/s

3.3.2 Memory Population Rules

- Each installed processor provides four channels of memory. On the Intel® Server Board S2600CW each memory channel supports two memory slots, for a total possible 16 DIMMs installed.
- System memory is organized into physical slots on DDR4 memory channels that belong to processor sockets.
- The memory channels from processor socket 1 are identified as Channel A, B, C and D. The memory channels from processor socket 2 are identified as Channel E, F, G and H.
- Each memory slot on the server board is identified by channel, and slot number within that channel. For example, DIMM_A1 is the first slot on Channel A on processor 1; DIMM_E1 is the first DIMM socket on Channel E on processor 2.
- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- A processor may be installed without populating the associated memory slots provided a second processor is installed with associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS and Error Management) in the BIOS setup is applied commonly across processor sockets.
- The BLUE memory slots on the server board identify the first memory slot for a given memory channel.

DIMM population rules require that DIMMs within a channel be populated starting with the BLUE DIMM slot or DIMM farthest from the processor in a “fill-farthest” approach. In addition, when populating a quad-rank DIMM with a single- or dual-rank DIMM in the same channel, the quad rank DIMM must be populated farthest from the processor. Note that quad-rank DIMMs and UDIMMs are not allowed in three slots populated configurations. Intel® MRC will check for correct DIMM placement.

The nomenclature for DIMM sockets on the Intel® Server Board S2600CW is detailed in the following table.

Table 5. Intel® Server Board S2600CW DIMM Nomenclature

Processor Socket 1								Processor Socket 2							
(0) Channel A		(1) Channel B		(2) Channel C		(3) Channel D		(0) Channel A		(1) Channel B		(2) Channel C		(3) Channel D	
A1	**A2**	**B1**	**B2**	**C1**	**C2**	**D1**	**D2**	**E1**	**E2**	**F1**	**F2**	**G1**	**G2**	**H1**	**H2**

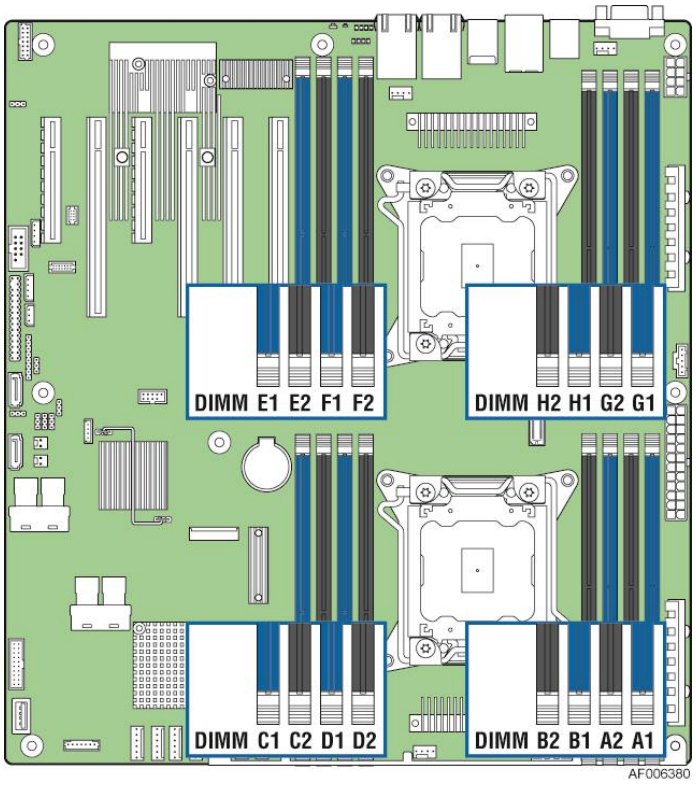

Figure 15. Intel® Server Board S2600CW DIMM Slot Layout

The following are generic DIMM population requirements that generally apply to the Intel® Server Board S2600CW.

- 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 or across sockets by Intel. If DIMMs with different frequencies are mixed, all DIMMs will run at the common lowest frequency.
- A maximum of 8 logical ranks (ranks seen by the host) per channel is allowed.
- DIMMs with different timing parameters can be installed on different slots within the same channel, but only timings that support the slowest DIMM will be applied to all. As a consequence, faster DIMMs will be operated at timings supported by the slowest DIMM populated.
- When one DIMM is used, it must be populated in the BLUE DIMM slot (farthest away from the CPU) of a given channel.
- When single-, dual-, and quad-rank DIMMs are populated for 2DPC, always populate the higher number rank DIMM first (starting from the farthest slot), for example, first quad-rank, then dual-rank, and last single-rank DIMM.

3.3.3 Effects of Memory Configuration on Memory Sizing

The system BIOS supports 4 memory configurations – Independent Channel Mode (Maximum Performance 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. Independent Channel mode is also known as Maximum Performance mode.
- **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 two 16GB Quad Rank DIMMs on each of 4 channels on each of 2 processor sockets, the total installed memory will be (((2 * 16GB) * 4 channels) * 2 CPU sockets) = 256GB.

  For a 16GB QR DIMM, each rank would be 4GB. With one rank reserved on each channel, that would 32GB reserved. So the available effective memory size would be 256GB - 32GB, or 224GB.
- **Mirroring Mode:** Mirroring creates a duplicate image of the memory that is in use, which uses half of the available memory to mirror the other half. This reduces the available memory size to half of the installed physical memory.

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

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

3.3.4 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 as well 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*

3.3.5 RAS 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 the OS even when 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". Addresses reliability by improving signal integrity at the physical layer. Additionally, assists 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 thus enhancing the containment of the fault and improving the reliability of the system.
- **Rank Level / Multi Rank Level Memory Sparing**: Dynamic failover 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.

3.3.6 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).

3.3.6.1 DIMM Discovery

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

Potential Error Cases:

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

3.3.6.2 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 equal to 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**.

3.3.6.3 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**.

3.3.6.4 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**.

3.3.6.5 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.

3.3.6.6 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.

3.4 System IO

The server board Input/Output features are provided via the embedded features and functions of several onboard components including the Integrated I/O Module (IIO) of the Intel® Xeon® E5-2600 v3 processor, the Intel® C612 series chipset, the Intel® Ethernet controller I350 or X540, and the I/O controllers embedded within the Emulex* Pilot-III Management Controller.

3.4.1 PCI Express* Support

The integrated I/O module incorporates the PCI Express* interface and supports up to 40 lanes of PCI Express. The Intel® Server Board S2600CW supports six PCIe slots from two processors and PCH:

- From PCH:
  - Slot 1: PCIe Gen II x4
- From the first processor:
  - Slot 5: PCIe Gen III x16/x8 electrical with x16 physical connector. Electrical x16 for S2600CW2 or S2600CWT, electrical x8 for S2600CW2S or S2600CWTS
  - Slot 6 PCIe Gen III x16 electrical with x16 physical connector
- From the second processor:
  - Slot 2: PCIe Gen III x16 electrical with x16 physical connector
  - Slot 3: PCIe Gen III x8 electrical with x8 physical connector
  - Slot 4: PCIe Gen III x16 electrical with x16 physical connector

The server board supports up to three full length full height double-width PCIe cards on slot 2, slot 4, and slot 6.

Standard PCIe slots can provide up to 25W power. PCIe slot 2, 4, and 6 can provide up to 75W power. The higher power requirement needs direct power cables from power supplies.

Note:

1. *Both CPU power connectors need to be connected in order for all the PCIe slots to work even only one CPU is installed.*

2. *While PCIe slot 5 is also an x16 connector, only 3 x16 PCIe slots on the board can support up to 75W slot power at a time. This is a configuration limitation based on the System Power budget. When used as riser slot, slot 6 can each provide up to 75W power to the riser.*

Below is a list of possible power configurations supporting different high power PCIe cards:

- x16 150W card
  - 75W from PCIe slot and 75W from PSU direct cable attach: 75W (2x3 conn)

- x16 225W card
  - 75W from PCIe slot and 150W from PSU direct cable attach: 75W (2x3 conn) + 75W (2x3 conn) or
  - 75W from PCIe slot and 150W from PSU direct cable attach: 75W (2x3 conn) + 75W (2x4 conn) or
  - 75W from PCIe slot and 150W from PSU direct cable attach: 150W (2x4 conn)
- x16 300W card
  - 75W from PCIe slot and 225W from PSU direct cable attach: 75W (2x3 conn) + 150W (2x4 conn)

PCIe slot 6 supports risers. The riser slot supports standard x16 PCIe connector pin-outs. The riser slots can support the following PCIe slot configurations:

- Riser with two x4 PCIe slots
- Riser with one x4 PCIe slot & one x8 PCIe slot
- Riser with two x8 PCIe slots
- Riser with one x16 PCIe slot

3.4.2 PCIe Enumeration and Allocation

The BIOS assigns PCI bus numbers in a depth-first hierarchy, in accordance with the PCI Local Bus Specification, Revision 2.2. The bus number is incremented when the BIOS encounters a PCI-PCI bridge device.

Scanning continues on the secondary side of the bridge until all subordinate buses are assigned numbers. PCI bus number assignments may vary from boot to boot with varying presence of PCI devices with PCI-PCI bridges.

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

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

The BIOS automatically assigns IRQs to devices in the system for legacy compatibility. A method is not provided to manually configure the IRQs for devices.

The following table shows the PCI layout for the Intel® Server Board S2600CW.

CPU 0	S2600CW2 S2600CWT	S2600CW2S S2600CWTS

PCI Ports	Bus (B)	Device (D)	Function (F)	On-board Device	On-board Device
PE1A	0	1	0	I350/X540	I350/X540
PE1B	0	1	1	I350/X540	I350/X540
PE2A	0	2	0	PCIE Slot6	PCIE Slot6
PE2B	0	2	1	PCIE Slot6	PCIE Slot6
PE2C	0	2	2	PCIE Slot6	PCIE Slot6
PE2D	0	2	3	PCIE Slot6	PCIE Slot6
PE3A	0	3	0	PCIE Slot5	PCIE Slot5
PE3B	0	3	1	PCIE Slot5	PCIE Slot5
PE3C	0	3	2	PCIE Slot5	LSISAS3008
PE3D	0	3	3	PCIE Slot5	LSISAS3008
DMI	0	0	0	PCIE Slot5	PCH

CPU 1				S2600CW2 S2600CWT	S2600CW2S S2600CWTS
PCI Ports	Bus (B)	Device (D)	Function (F)	On-board Device	On-board Device
PE1A	0x80	1	0	PCIE Slot3	PCIE Slot3
PE1B	0x80	1	1	PCIE Slot3	PCIE Slot3
PE2A	0x80	2	0	PCIE Slot2	PCIE Slot2
PE2B	0x80	2	1	PCIE Slot2	PCIE Slot2
PE2C	0x80	2	2	PCIE Slot2	PCIE Slot2
PE2D	0x80	2	3	PCIE Slot2	PCIE Slot2
PE3A	0x80	3	0	PCIE Slot4	PCIE Slot4
PE3B	0x80	3	1	PCIE Slot4	PCIE Slot4
PE3C	0x80	3	2	PCIE Slot4	PCIE Slot4
PE3D	0x80	3	3	PCIE Slot4	PCIE Slot4
DMI	N/A	N/A	N/A	Not connected	Not connected

3.4.3 PCIe Non-Transparent Bridge (NTB)

PCI Express Non-Transparent Bridge (NTB) acts as a gateway that enables high performance, low overhead communication between two intelligent subsystems, the local and the remote subsystems. The NTB allows a local processor to independently configure and control the local subsystem, provides isolation of the local host memory domain from the remote host memory domain while enabling status and data exchange between the two domains.

The PCI Express Port 3A of Intel® Xeon® Processor E5-2600 Product Families can be configured to be a transparent bridge or an NTB with x4/x8/x16 link width and Gen1/Gen2/Gen3 link speed. This NTB port could be attached to another NTB port or PCI Express Root Port on another subsystem. NTB supports three 64bit BARs as configuration

space or prefetchable memory windows that can access both 32bit and 64bit address space through 64bit BARs.

There are 3 NTB supported configurations:

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

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

3.4.4 PCIe SSD Support

The board supports PCIe SSD as cabled add-in card on existing PCIe slots. The x16 PCIe add-in card can support up to four x4 PCIe SSD 2.5" drives with mini-SAS HD cables connected to the hot-swap backplane.

The ingredients for PCIe SSD:

- PCIe SSD cabled add-in card: cable PCIe link through add-in card in existing PCIe slots, support four x4 PCIe SSD 2.5" drives.
- Combo hot-swap backplane supporting SAS/SATA drives and PCIe SSDs.
- Four x4 mini-SAS HD cables connecting add-in card and combo HSBP.
- 8 x 2.5" Drives: 4 SAS/SATA + 4 PCIe SSD 2.5"

Hot-plug and enclosure management features are not supported with PCIe SSD devices.

The PCIe SSD add-in-card needs to be used on electrically x16 PCIe slot. On S2600CW2 and S2600CWT, the add-in-card can be used on PCIe slot 2/4/5/6. On S2600CW2S and S2600CWTS, the add-in-card can be used on PCIe slot 2/4/6 (PCIe slot 5 is electrically x8 bandwidth).

Note*:*

Intel® Server Chassis P4304XXMFEN2/P4304XXMUXX support up to airflow 200LFM/55°C. Some PCIe SSD devices may require airflow 300LFM (or above), performance mode in BIOS setup menu will need to be enabled.

When add-in card form factor PCIe SSD device requiring 300LFM (or above) used in Intel® Server Chassis P4304XXMUXX, the device, the card should be installed on PCIe slot 3/4/5/6, in addition to enabling the performance mode in BIOS setup menu.

Please also see the note in Section "HTA support for Intel® Server Chassis P4304XXMFEN2 / P4304XXMUXX with Intel® Server Board S2600CW" in Intel® Server Chassis P4304XXMFEN2 / P4304XXMUXX TPS.

3.4.5 Serial ATA (SATA) Support

The Intel® C612 Series chipset provides the server board with support for up to ten Serial ATA (SATA) ports from two integrated controllers identified as SATA and sSATA.

The Intel® Server Board S2600CW family on-board SATA connectors include:

- Two 7-pin single-port SATA connectors labeled "SATA-4" and "SATA-5"; each port is capable of supporting up to 6 Gbps. These connectors are intended for use with optical drives or SATADOM devices.
- Two 4-port mini-SAS HD connectors labeled "SATA_0-3" and "sSATA_0-3". All eight ports are capable of supporting up to 6 Gbps.

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

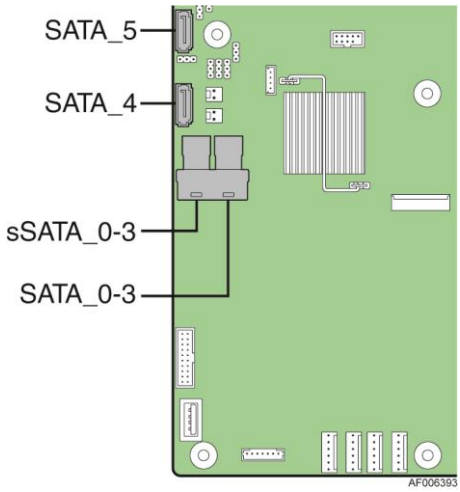

3.4.6 SATADOM Support

SATADOM devices can be used on the 7-pin SATA connectors (SATA-4 and SATA-5). There are configuration limitations when using SATADOM due to SATADOM clearance. Two slim type SATADOMs can be used at a time, or one slim type SATADOM on SATA-4 connector and one low profile SATADOM on SATA-5 connector.

The board design supports the Apacer* power delivery option. The Apacer* SDM 7+2 connector has the power and ground pins on the side of the 7-pin SATA connector but is compatible with a standard 7-pin SATA cable.

Note:

Once M.2/NGFF device is used on the server board, SATA-4 port cannot be used. SATA-5 can be used at the same time when M.2/NGFF device is used.

SATADOM may not be used under these conditions:

- *A single-width half-length or full-length add-in card is installed on PCIe slot 1, or*
- *A double-width 3/4-length or full-length add-in card is installed on PCIe slot 2.*

3.4.7 M.2/NGFF Support

M.2, formerly known as the Next Generation Form Factor (NGFF), is a specification for computer expansion cards and associated connectors. It is a small form factor module supporting SSD/Memory-offloading technology using SATA or PCIe x4 links. The server board uses a SATA mux to select between M.2/NGFF connector and 7-pin SATA connector.

Note*:*

- *Although M.2/NGFF allows support for many different interfaces, the server board supports only the SATA interface (all other interfaces such as PCIe and USB are not supported).*
- *Once M.2/NGFF device is used on the server board, SATA-4 port cannot be used. SATA-5 can be used at the same time when M.2/NGFF device is used.*
- *In order to support full performance from Intel S3500 M.2 device, please use "Performance Mode" in Intel server BIOS setup options to meet thermal requirement. Using "Acoustic Mode" may result in Intel S3500 M.2 device running with reduced performance.*

The following diagram identifies the location of the M.2/NGFF connector on the board.

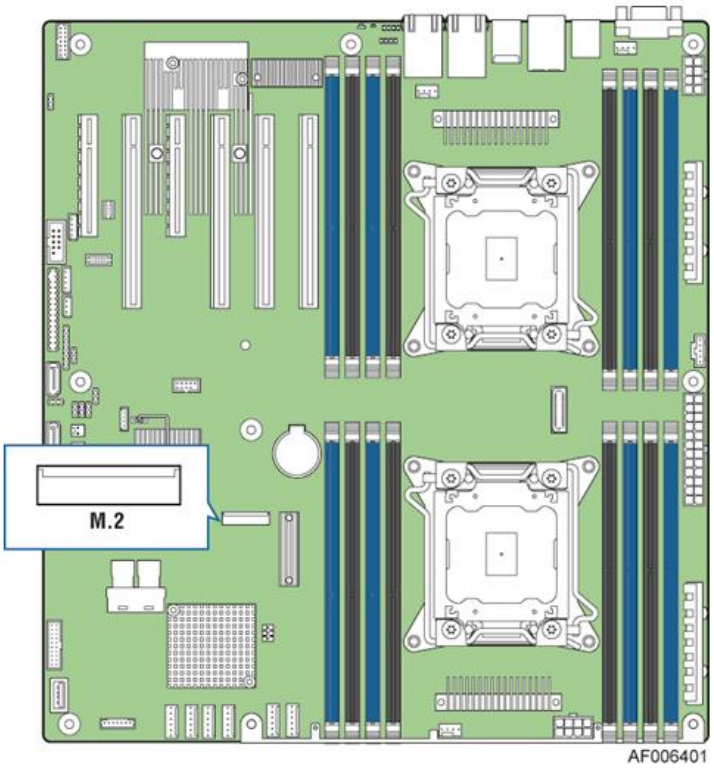

The following diagram shows how the M.2/NGFF modules are installed on the connector.

Note: The insertion angle may vary.

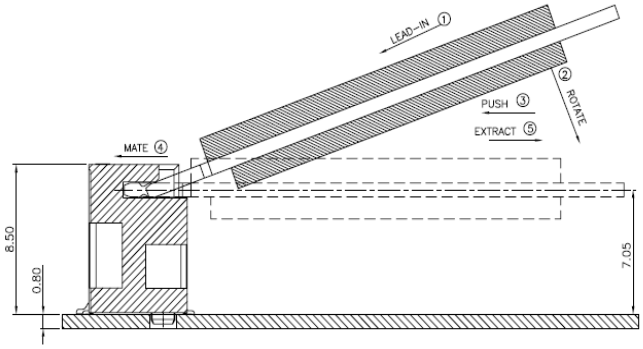

3.4.8 Embedded SATA RAID Support

The Intel® Server Board S2600CW2/S2600CWT has embedded support for two SATA SW RAID options:

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

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

In addition to the SATA SW RAID options, the Intel® Server Board S2600CW2S/S2600CWTS supports Integrated MegaRaid RAID (IMR).

3.4.8.1 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 SW-RAID option must be enabled in <F2> BIOS Setup
- Intel® RSTe option must be selected in <F2> BIOS Setup
- Intel® RSTe drivers must be loaded for the specified operating system
- At least two SATA drives are needed to support RAID levels 0 or 1
- At least three SATA drives are needed to support RAID levels 5 or 10

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

- A boot-time, pre-operating system environment, text mode user interface that allows the user to manage the RAID configuration on the system. Its feature set is kept simple to keep size to a minimum, but allows the user to create and delete RAID volumes and

select recovery options when problems occur. The user interface can be accessed by hitting the <CTRL-I > keys during system POST.

- Providing 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.
- Providing the user with a status of the RAID volumes at each boot up.

3.4.8.2 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
- Native support for RAID Levels – 0, 1, 10
- Optional support for RAID Level 5
  - Enabled with the addition of an optionally installed SATA RAID 5 Upgrade Key
- Maximum drive support = 8
- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*
- OS Support = Windows 7*, Windows 2008*, Windows 2003*, RHEL*, SLES, other Linux* variants using partial source builds
- Utilities = Windows* GUI and CLI, Linux GUI and CLI, DOS CLI, and EFI CLI

The following diagram identifies the location of the key to enable ESRT2 RAID level 5.

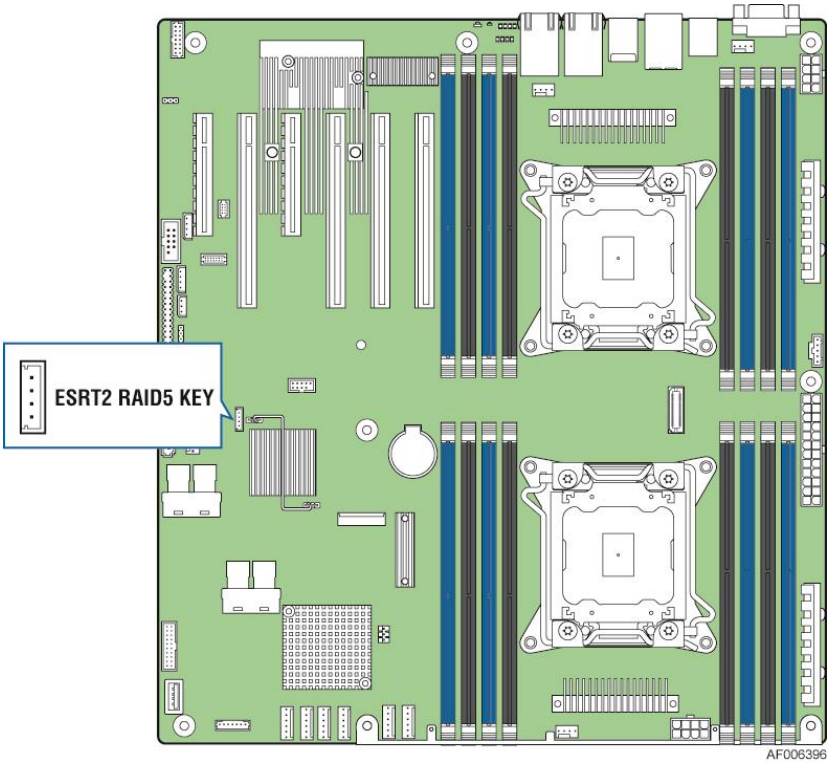

Using the <F2> BIOS Setup Utility, accessed during system POST, system setup options are available to enable/disable the Software RAID feature. By default, the embedded SATA Software RAID feature is disabled.

3.4.9 Serial Attached SCSI (SAS) Support

The Intel® Server Board S2600CW2S/S2600CWTS supports a Gen3 12G SAS IO Controller with LSISAS3008 (S2600CW2/S2600CWT does not support this). The 8-port SAS (12G) or SATA (6G) will connect to a 1x2 Right Angle (RA) mini-SAS HD connector and be used for Hot Swap Backplane (HSBP) connectivity.

The following list summarizes the features of the LSISAS3008 controller:

- Providing an eight-lane 8Gbps PCIe 3.0 host interface
- Providing an eight-port 12Gbps SAS and 6Gbps SATA interface
- Providing a full-featured hardware-based RAID solution that supports RAID levels 0, 1, 1E, and 10
- 16Mbytes Flash ROM memory
- 256kbit MRAM memory for write journaling support
- Up to two 36-pin RA mini-SAS HD connectors

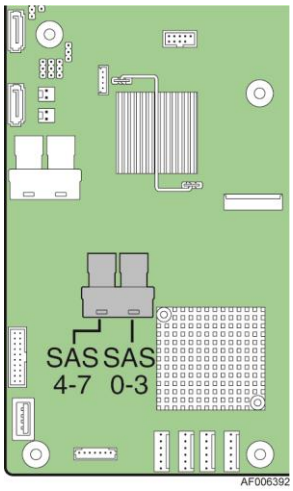

3.4.10 Integrated MegaRAID Support

Integrated MegaRaid RAID (IMR) (RAID 0/1/10) support is included with based LSISAS3008 functionality. Integrated MegaRaid (IMR) (RAID 5) upgrade is supported through installation of an activation key.

The following diagram identifies the location of the activation key for the IMR RAID 5.

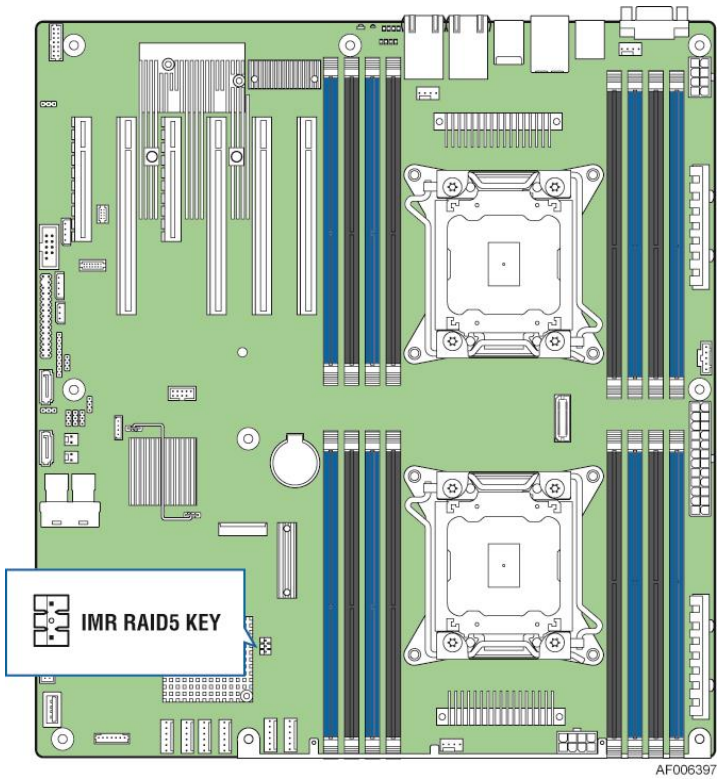

3.4.11 USB Support

The Intel® C612 Series chipset has up to two Enhanced Host Controller Interface (EHCI) host controllers that support USB high-speed signaling. The Intel® C612 Series chipset supports up to 14 USB 2.0 ports of which up to six can be used as USB3.0 ports.

3.4.11.1 USB connectors

The server board provides the following USB ports:

- Two USB 2.0 ports on the board rear end, next to the VGA connector.
- Two USB 3.0 ports on the board rear end, next to the NIC1 connector.
- One 2x10-pin USB 3.0 header on the board, providing USB connectivity to the front panel.
- One Type-A USB header on the board.

3.4.11.2 eUSB Module

Smart Modular Z-U130 Value Solid State Drive (SSD) is an embedded USB2.0 (eUSB2.0) storage solution built around high performance Intel® NAND flash memory. This module uses single-level cell Intel® NAND flash memory with cache programming and dual plane feature set designed to improve overall module performance. The Intel® Z-U130 Value SSD supports the Universal Serial Bus (USB) specification v2.0 and is backward compatible with v1.1. The module uses industry standard connectors which are available in two sizes. This device can be used with operating systems compatible with the USB Mass Storage Class specification v1.0.

The following diagram identifies the location of the eUSB connector on the board.

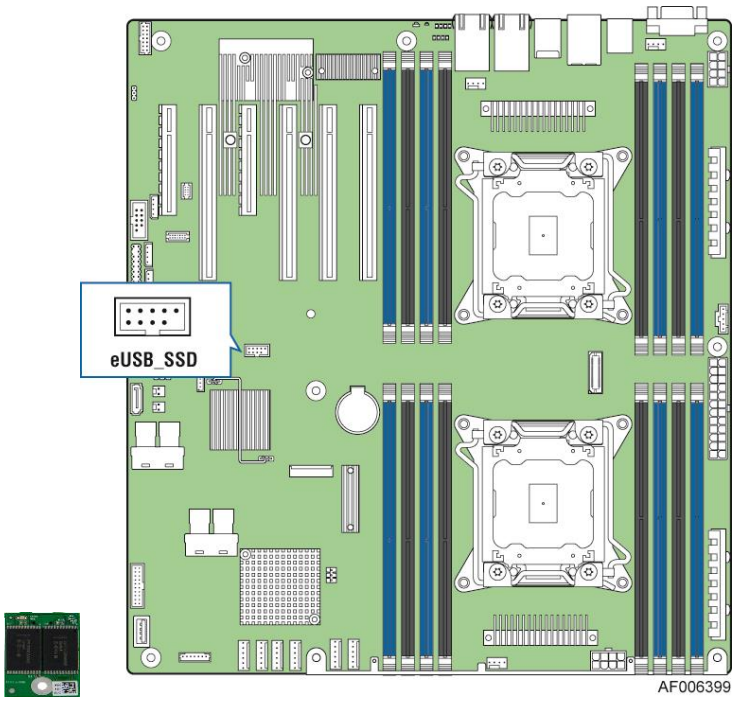

3.4.12 Graphics Controller and Video Support

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

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

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

Table 6. Video Modes

2D Mode	2D Video Mode Support			
	8 bpp	16 bpp	24 bpp	32 bpp
640x480	X	X	X	X
800x600	X	X	X	X
1024x768	X	X	X	X
1152x864	X	X	X	X
1280x1024	X	X	X	X
1600x1200 [1]	X	X		

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

3.4.12.1 Dual Video and Add-In Video Adapters

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

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

Configurations with add-in video cards can get more complicated on server boards that have two or more CPU sockets. Some multi-socket boards have PCIe slots capable of hosting an add-in video card which are attached to the IIOs of CPU sockets other than CPU Socket 1. However, only one CPU Socket can be designated as “Legacy VGA Socket” as required in POST.

To provide for this, there is another PCI Configuration option to control “Legacy VGA Socket”. The rules for this are:

- This option appears only on boards which have the possibility of an add-in video adapter in a PCIe slot on a CPU socket other than socket 1.
- When present, the option is grayed out and unavailable unless an add-in video card is actually installed in a PCIe slot connected to the other socket.
- Because the Onboard Video is “hardwired” to CPU Socket 1, whenever Legacy VGA Socket is set to a CPU Socket other than Socket 1, that disables both Onboard Video.

3.4.12.1.1 Dual Monitor Video

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

In the single video mode, the onboard video controller or the add-in video adapter is detected during the POST. In the dual video mode, the onboard video controller is enabled and is the primary video device while the add-in video adapter is allocated resources and is considered as the secondary video device.

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

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

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

  Case 1: Onboard Video only active display.

  Onboard Video = Enabled (grayout, can't change)

  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)

  Add-in Video Adapter = Disabled (grayout, can't change)

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

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

  Onboard Video = Enabled

  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)

  Add-in Video Adapter = Disabled

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

  Onboard Video = Disabled

  Legacy VGA Socket = CPU Socket 1 (grayout, can't change)

  Add-in Video Adapter = Enabled

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

Onboard Video = Enabled

Legacy VGA Socket = CPU Socket 1 (grayout, can't change)

Add-in Video Adapter = Enabled

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

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

Onboard Video = Enabled

Legacy VGA Socket = CPU Socket 1

Add-in Video Adapter = Disabled (grayout, can't change)

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

Onboard Video = Disabled (grayout, can't change)

Legacy VGA Socket = CPU Socket 2

Add-in Video Adapter = Enabled (grayout, can't change)

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

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

Onboard Video = Enabled

Legacy VGA Socket = CPU Socket 1

Add-in Video Adapter = Disabled

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

Onboard Video = Disabled

Legacy VGA Socket = CPU Socket 1

Add-in Video Adapter = Enabled

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

Onboard Video = Enabled

Legacy VGA Socket = CPU Socket 1

Add-in Video Adapter = Enabled

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

Onboard Video = Disabled (grayout, can't change)

Legacy VGA Socket = CPU Socket 2

Add-in Video Adapte = Enabled (grayout, can't change)

3.4.12.2 Setting Video Configuration Options using the BIOS Setup Utility

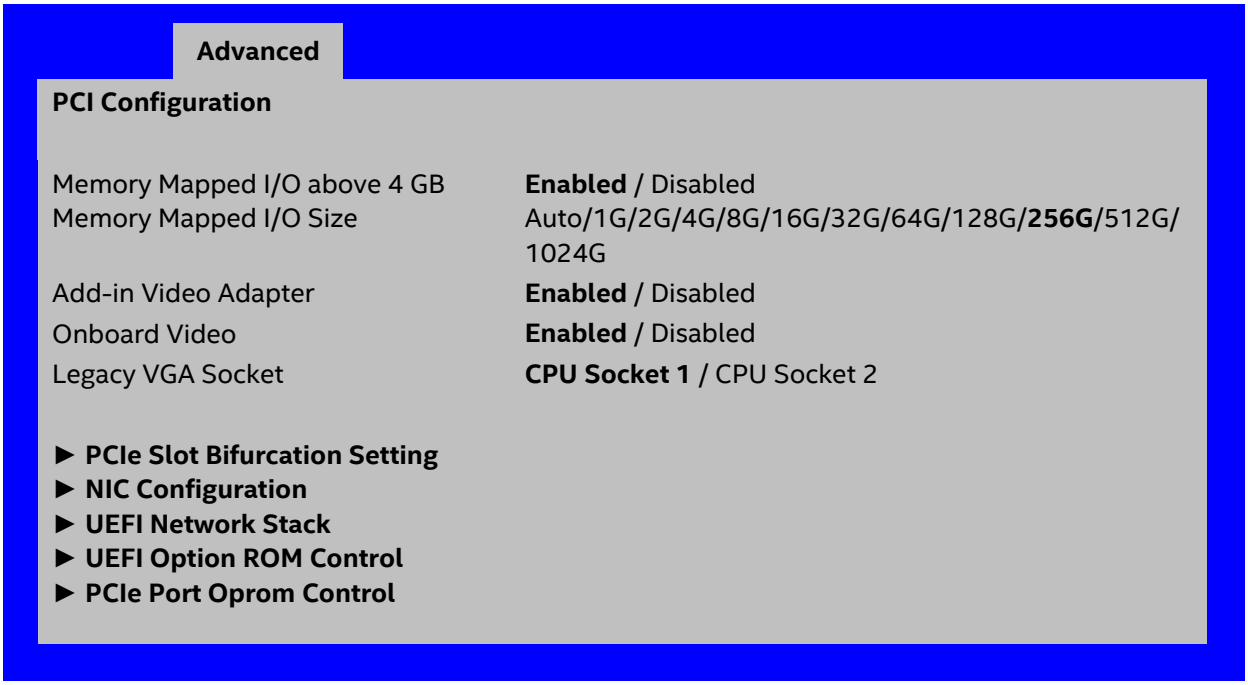

Figure 16. BIOS Setup Utility – Video Configuration Options

1. **Add-in Video Adapter**

   Option Values: **Enabled**
   Disabled

   Help Text:

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

   Comments: This option must be enabled to use an add-in card as a primary POST Legacy Video device.

   If there is no add-in video card in any PCIe slot connected to CPU Socket 1 with the Legacy VGA Socket option set to *CPU Socket 1*, this option is set to *Disabled* and grayed out and unavailable.

   If there is no add-in video card in any PCIe slot connected to CPU Socket 2 with the Legacy VGA Socket option set to *CPU Socket 2*, this option is set to *Disabled* and grayed out and unavailable.

   If the Legacy VGA Socket option is set to CPU Socket 1 with both Add-in Video Adapter and Onboard Video Enabled, the onboard video device works as primary video device while add-in video adapter as secondary.

2. **Onboard Video**

Option Values: **Enabled**
Disabled

Help Text:

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

Comments: When disabled, the system requires an add-in video card for the video to be seen. When there is no add-in video card installed, Onboard Video is set to *Enabled* and grayed out so it cannot be changed.

If there is an add-in video card installed in a PCIe slot connected to CPU Socket 1, and the Legacy VGA Socket option is set to *CPU Socket 1*, then this Onboard Video option is available to be set and default as Disabled.

If there is an add-in video card installed on a PCIe slot connected to CPU Socket 2, and the Legacy VGA Socket option is set to *CPU Socket 2*, this option is grayed out and unavailable, with a value set to *Disabled*. This is because the Onboard Video is connected to CPU Socket 1, and is not functional when CPU Socket 2 is the active path for video. When Legacy VGA Socket is set back to *CPU Socket 1*, this option becomes available again and is set to its default value of *Enabled*.

Note: *This option does not appear on some models.*

3. **Legacy VGA Socket**

Option Values: **CPU Socket 1**
CPU Socket 2

Help Text:

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

Comments: This option is necessary when using an add-in video card on a PCIe slot attached to CPU Socket 2, due to a limitation of the processor IIO. The Legacy video device can be connected through either socket but there is a setting that must be set on only one of the two. This option allows the switch to using a video card in a slot connected to CPU Socket 2.

This option does not appear unless the BIOS is running on a board which has one processor installed on CPU Socket 2 and can potentially have a video card installed in a PCIe slot connected to CPU Socket 2.

This option is grayed out as unavailable and set to *CPU Socket 1* unless there is a processor installed on CPU Socket 2 and a video card installed in a PCIe slot connected to CPU Socket 2. When this option is active and is set to *CPU Socket 2*, then both Onboard Video and Dual Monitor Video are set to *Disabled* and grayed out as unavailable. This is because the Onboard Video is a PCIe device connected to CPU Socket 1, and is unavailable when the Legacy VGA Socket is set to Socket 2.

3.4.13 Trusted Platform Module

The Trusted Platform Module (TPM) provides platform security functions such as hash, encryption, and secure storage, and works in conjunction with the processor's TXT functionality.

The TPM is a small board that provides hardware level security for the server and resides on the LPC bus.

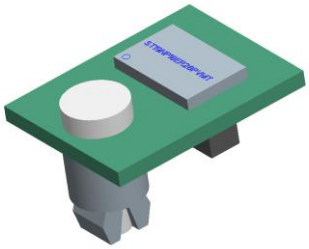

Figure 17. TPM Module

3.4.14 Network Support

The Intel® Server Board S2600CW2/S2600CW2S provides 1Gb network connectivity with the Intel® I350 dual port controller, and the Intel® Server Board S2600CWT/S2600CWTS provides 10Gb network connectivity with the Intel® X540 dual port controller. The controllers are fully integrated MAC/PHY in a single low power package that supports dual-port 1G/10G Ethernet designs.

The board also provides a 1G RJ45 Dedicated Management NIC port (DMN) for the Intel® Remote Management Module (RMM4 Lite). The DMN is active with or without the RMM4 Lite key installed.

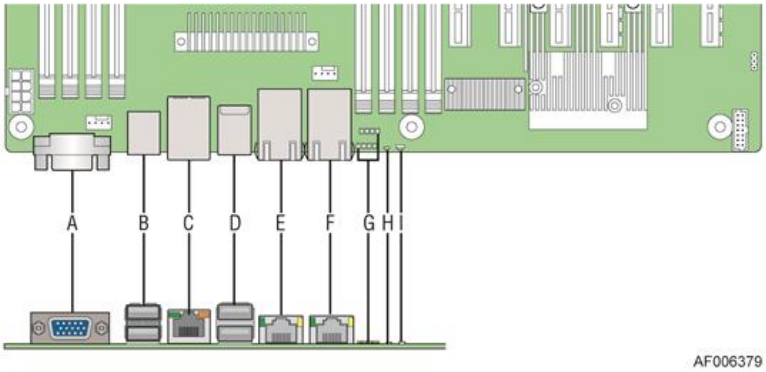

Callout	Description	Callout	Description
A	Video	B	USB 2.0
C	Dedicated Management NIC (DMN)	D	USB 3.0
E	NIC1	F	NIC2
G	System Diagnostic LED	H	ID LED
I	System Status LED		

The server board S2600CW2/S2600CW2S (I350) has the following MAC addresses assigned:

- NIC 1 MAC address
- NIC 2 MAC address – Assigned the NIC 1 MAC address +1
- Integrated BMC LAN Channel MAC0 address – Assigned the NIC 1 MAC address +2
- Integrated BMC LAN Channel MAC1 address – Assigned the NIC 1 MAC address +3
- Intel® Remote Management Module (Intel® RMM) MAC address – Assigned the NIC 1 MAC address +4

The server board S2600CWT/S2600CWTS (X540) has the following MAC addresses assigned:

- NIC 1 MAC address
- NIC 2 MAC address – Assigned the NIC 1 MAC address +1
- Integrated BMC LAN Channel MAC0 address – Assigned the NIC 1 MAC address +2
- Integrated BMC LAN Channel MAC1 address – Assigned the NIC 1 MAC address +3
- Intel® Remote Management Module (Intel® RMM) MAC address – Assigned the NIC 1 MAC address +4
- NIC 1 SAN MAC address - Assigned the NIC 1 MAC address +5
- NIC 2 SAN MAC address - Assigned the NIC 1 MAC address +6

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

and transmit/receive activity when blinking. The LED at the right of the connector indicates link speed as defined in the following table.

Table 7. External RJ45 NIC Port LED Definition

LED Color	LED State	NIC State
Green/Amber (Right)	Off	3rd Fastest (100 Mbps for X540)
	Amber/Yellow	2nd Fastest (1 Gbps for X540)
	Green	Fastest (10 Gbps for X540)
Green (Left)	On	Active Connection
	Blinking	Transmit/Receive activity

4. System Security

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

- Password Protection
- Front Panel Lockout
- Trusted Platform Module (TPM) support
- Intel® Trusted Execution Technology

4.1 BIOS Setup Utility Security Option Configuration

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

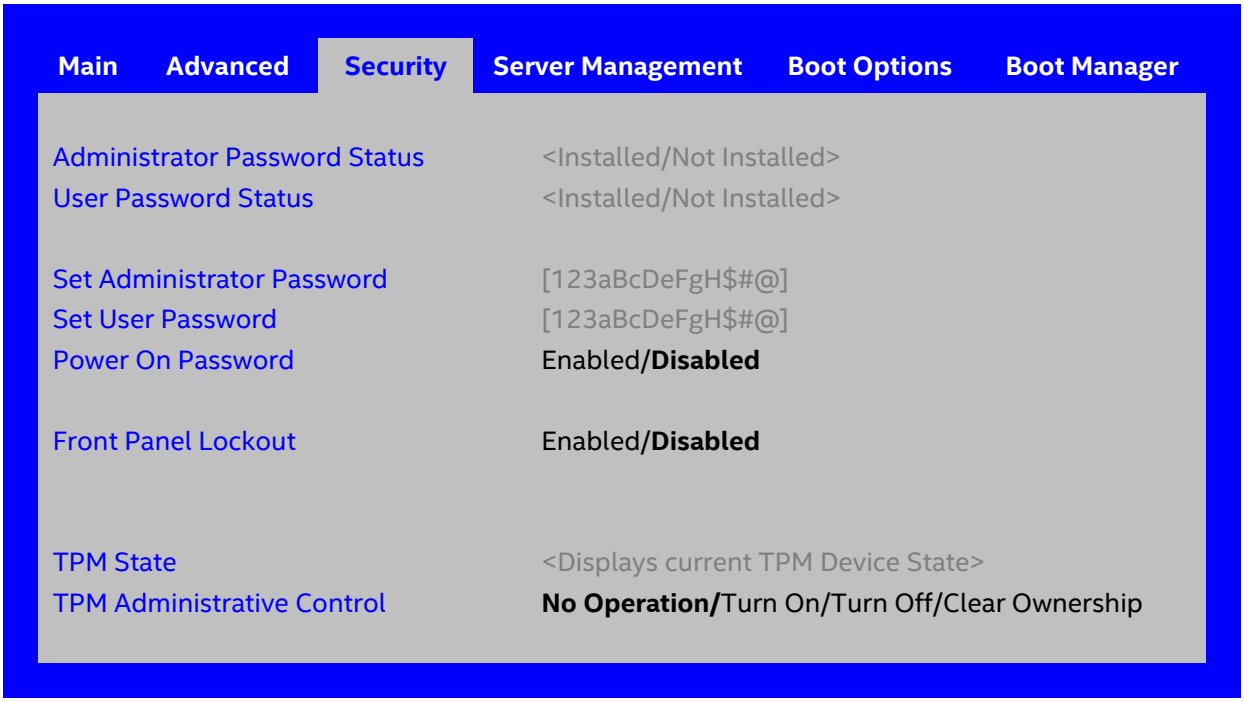

Table 8. Setup Utility – Security Configuration Screen Fields

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 the 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.	

4.2 BIOS Password Protection

The BIOS uses passwords to prevent unauthorized tampering with the server setup. Passwords can restrict entry to the BIOS Setup, restrict use of the Boot Popup menu, and suppress automatic USB device reordering.

There is also an option to require a Power On password entry in order to boot the system. If the Power On Password function is enabled in Setup, the BIOS will halt early in POST to request a password before continuing POST.

Both Administrator and User passwords are supported by the BIOS. An Administrator password must be installed in order to set the User password. The maximum length of a password is 14 characters. A password can have alphanumeric (a-z, A-Z, 0-9) characters and it is case sensitive. Certain special characters are also allowed, from the following set:

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

The Administrator and User passwords must be different from each other. An error message will be displayed if there is an attempt to enter the same password for one as for the other.

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 eight characters in length, and must include at least one each of alphabetic, numeric, and special characters. If a “weak” password is entered, a popup warning message will be displayed, although 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 the BIOS Setup or other explicit means of changing the passwords. Clearing the Administrator password will also clear the User password.

Alternatively, the passwords can be cleared by using the Password Clear jumper if necessary. Resetting the BIOS configuration settings to the default values (by any method) has no effect on the Administrator and User passwords.

Entering the User password allows the user to modify only the System Time and System Date in the Setup Main screen. Other setup fields can be modified only if the Administrator password has been entered. If any password is set, a password is required to enter the BIOS setup.

The Administrator has control over all fields in the BIOS setup, including the ability to clear the User password and the Administrator password.

It is strongly recommended that at least an Administrator Password be set, because not having set a password gives everyone who boots the system the equivalent of Administrative access. Unless an Administrator password is installed, any User can go into the Setup and change the BIOS settings at will.

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 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.

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.

4.3 Trusted Platform Module (TPM) Support

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

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

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

4.3.1 TPM Security BIOS

The BIOS TPM support conforms to the *TPM PC Client Implementation Specification for Conventional BIOS* and to 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.

4.3.2 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. A 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).

4.3.3 TPM Security Setup Options

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

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

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

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

4.4 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5-4600/2600/2400/1600 v3 product families support Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment. Designed to help protect against software-based attacks, Intel® Trusted Execution Technology integrates new security features and capabilities into the processor, chipset, and other platform components. When used in conjunction with Intel® Virtualization Technology, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications.

This hardware-rooted security provides a general-purpose, safer computing environment capable of running a wide variety of operating systems and applications to increase the

confidentiality and integrity of sensitive information without compromising the usability of the platform.

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

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

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

5. Intel® Server Board S2600CW 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 product families should be referenced for more in-depth and design level platform management information.

5.1 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.

5.1.1 IPMI 2.0 Features Overview

- Baseboard management controller (BMC)
- IPMI Watchdog timer
- Messaging support, including command bridging and user/session support
- Chassis device functionality, including power/reset control and BIOS boot flags support
- Event receiver device: The BMC receives and processes events from other platform subsystems.
- Field Replaceable Unit (FRU) inventory device functionality: The BMC supports access to system FRU devices using IPMI FRU commands.
- System Event Log (SEL) device functionality: The BMC supports and provides access to a SEL including SEL Severity Tracking and the Extended SEL.

- Sensor Data Record (SDR) repository device functionality: The BMC supports storage and access of system SDRs.
- Sensor device and sensor scanning/monitoring: The BMC provides IPMI management of sensors. It polls sensors to monitor and report system health.
- IPMI interfaces
  - Host interfaces include system management software (SMS) with receive message queue support, and server management mode (SMM)
  - IPMB interface
  - LAN interface that supports the IPMI-over-LAN protocol (RMCP, RMCP+)
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS.
- BMC self-test: The BMC performs initialization and runtime self-tests and makes results available to external entities.

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

5.1.2 Non-IPMI Features Overview

The BMC supports the following non-IPMI features.

- In-circuit BMC firmware update.
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Chassis intrusion detection (dependent on platform support).
- Fan speed control with SDR. Fan redundancy monitoring and support.
- Enhancements to fan speed control.
- Power supply redundancy monitoring and support.
- Hot-swap fan support.
- Acoustic management: Support for multiple fan profiles.
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval.
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.
- Power state retention.
- Power fault analysis.
- Intel® Light-Guided Diagnostics.

- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support.
- Email alerting.
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server.
  - Human-readable SEL
  - Additional system configurability
  - Additional system monitoring capability
  - Enhanced online help
- Integrated KVM.
- Enhancements to KVM redirection.
  - Support for higher resolution
- Integrated Remote Media Redirection.
- Lightweight Directory Access Protocol (LDAP) support.
- Intel® Intelligent Power Node Manager support.
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
  - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.5 compliance (product-specific).
- Management support for PMBus* rev1.2 compliant power supplies.
- BMC Data Repository (Managed Data Region Feature).
- Support for an Intel® Local Control Display Panel.
- System Airflow Monitoring.
- Exit Air Temperature Monitoring.
- Ethernet Controller Thermal Monitoring.
- Global Aggregate Temperature Margin Sensor.
- Memory Thermal Management.
- Power Supply Fan Sensors.

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

5.2 Platform Management Features and Functions

5.2.1 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 *POWER_ON*)	Turns power on or off
CPU Thermal	Processor Thermtrip	Turns power off
PCH Thermal	PCH Thermtrip	Turns power off
WOL (Wake On LAN)	LAN	Turns power on

5.2.2 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states.

Table 9. ACPI Power States

State	Supported	Description
S0	Yes	Working. ▪ The front panel power LED is on (not controlled by the BMC). ▪ The fans spin at the normal speed, as determined by sensor inputs. ▪ Front panel buttons work normally.
S1	No	Not supported.
S2	No	Not supported.

State	Supported	Description
S3	No	Supported only on Workstation platforms. See appropriate Platform Specific Information for more information.
S4	No	Not supported.
S5	Yes	Soft off. ▪ The front panel buttons are not locked. ▪ The fans are stopped. ▪ The power-up process goes through the normal boot process. ▪ The power, reset, front panel NMI, and ID buttons are unlocked.

5.2.3 System Initialization

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

5.2.3.1 Processor Tcontrol Setting

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

5.2.3.2 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.

5.2.3.3 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 will deactivate POST LEDs after POST completes.

5.2.4 Watchdog Timer

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

5.2.5 System Event Log (SEL)

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

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

5.3 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.

5.3.1 Sensor Scanning

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

system and the availability of the physical sensor when in that state. For example, non-standby voltages are not monitored when the system is in S4 or S5 power state.

5.3.2 Sensor Rearm Behavior

5.3.2.1 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.

5.3.2.2 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.

5.3.3 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.

5.3.4 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.

5.3.5 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.

5.3.6 BMC Watchdog Sensor

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

5.3.7 BMC System Management Health Monitoring

The BMC tracks the health of each of its IPMI sensors and reports 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.

5.3.8 VR Watchdog Timer

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

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

5.3.9 System Airflow Monitoring

The sensor is valid only for Intel server chassis. BMC provides an IPMI sensor to report the volumetric system airflow in CFM (cubic feet per minute). The airflow 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.

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

5.3.10.1 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 a reference value.

- Thermal fault indication sensors – These are discrete sensors that indicate a specific thermal fault condition.

5.3.10.2 Processor DTS-Spec Margin Sensor(s)

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family incorporate a DTS based thermal spec. This allows a much more accurate control of the thermal solution and enables 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. DTS 2.0 is implemented on new Intel board generation DTS 2.0 incorporated platform-visible thermal data interfaces and internal algorithms for calculating the relevant thermal data. As the major difference between the DTS 1.0 and DTS 2.0 is that allows the CPUs to automatically calculate thermal gap/margin to DTS profile as input for Fan Speed Control. DTS 2.0 helps to further optimize system acoustics. Please refer to iBL #455822 Platform Digital Thermal Sensor (DTS) Based Thermal Specifications and Overview – Rev 1.5 for more details about DTS 2.0

5.3.10.3 Processor Thermal Margin Sensor(s)

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

5.3.10.4 Processor Thermal Control Monitoring (Prochot)

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

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

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

5.3.10.5 Processor Voltage Regulator (VRD) Over-Temperature Sensor

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

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).

5.3.10.6 Inlet Temperature Sensor

Each platform supports a thermal sensor for monitoring the inlet temperature. For Intel server chassis, the inlet temperature sensor is on front panel with address 21h. For 3rd chassis, sensor 20h which is on the front edge of baseboard can be used as inlet temperature sensor with several degrees offset from actual inlet temperature.

5.3.10.7 Baseboard Ambient Temperature Sensor(s)

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

5.3.10.8 Server South Bridge (SSB) Thermal Monitoring

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

5.3.10.9 Exit Air Temperature Monitoring

The sensor is only valid for Intel server chassis. 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 airflow 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 airflow (cubic feet per meter or CFM), and altitude range.

The Exit Air temp sensor is only available when PMBus* power supplies are installed.

5.3.10.10 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 monitors the sensors and uses this data as an input to the fan speed control. The BMC instantiates an IPMI temperature sensor for each device on the baseboard.

5.3.10.11 Memory VRD-Hot Sensor(s)

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

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

5.3.10.12 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 instantiates 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.

5.3.10.13 Processor ThermTrip

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

5.3.10.14 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.

5.3.10.15 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).

5.3.11 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 10. 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

5.3.11.1 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 11.

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

Table 11. 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.

5.3.11.2 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 HW to allow power-on of the system.

At BMC startup, the BMC checks for the fault condition and sets 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.

5.3.11.3 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.

5.3.11.4 CATERR Sensor

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

The CATERR signal is defined as having three states:

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

All processors in a system have their CATERR pins tied together. The pin is used as a communication path to signal a catastrophic system event to all CPUs. The BMC has direct access to this aggregate CATERR signal.

The BMC only monitors for the "CATERR held low" condition. A pulsed low condition is ignored by the BMC. If a CATERR-low condition is detected, the BMC logs an error message to

the SEL against the CATERR sensor and the default action after logging the SEL entry is to reset the system. The BIOS setup utility provides an option to disable or enable system reset by the BMC for detection of this condition.

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

When the BMC detects that this aggregate CATERR signal has asserted, it can then go through PECI to query each CPU to determine which one was the source of the error and write an OEM code identifying the CPU slot into an event data byte in the SEL entry. If PECI is non-functional (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 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.

5.3.11.5 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).

5.3.12 Voltage Monitoring

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

5.3.12.1 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 FW 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.

5.3.13 Fan Monitoring

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

5.3.13.1 Fan Tach Sensors

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

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

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

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

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.

5.3.13.2 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).

5.3.13.3 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 transitions 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.

5.3.13.4 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.

5.3.13.5 Monitoring for “Fans Off” Scenario

On Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family, it is likely that there will be situations where specific fans are turned off based on current system conditions. BMC Fan monitoring will comprehend this scenario and not log false failure events. The recommended method is for the BMC FW to halt updates to the value of the associated fan tach sensor and set that sensor’s IPMI sensor state to “reading-state-unavailable” when this mode is active. Management software must comprehend this state for fan tach sensors and not report these as failure conditions.

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

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

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

5.3.14 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 the system comes out of “Fans-off” mode.
  - Any associated temperature sensor is in a critical state. The SDR describes which temperature-threshold violations cause fan boost for each fan domain.
  - The BMC is in firmware update mode, or the operational firmware is corrupted.
  - If any of the above conditions apply, the fans are set to a fixed boost state speed.
- Nominal

- A fan domain's nominal fan speed can be configured as static (fixed value) or controlled by the state of one or more associated temperature sensors.

5.3.14.1 Hot-Swap Fans

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

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

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

5.3.14.2 Fan Redundancy Detection

The BMC supports redundant fan monitoring and implements a fan redundancy sensor. A fan redundancy sensor generates events when its associated set of fans transitions between redundant and non-redundant states, as determined by the number and health of the fans. The definition of fan redundancy is configuration dependent. The BMC allows redundancy to be configured on a per fan redundancy sensor basis through OEM SDR records.

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

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

5.3.14.3 Fan Domains

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

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

The same device may drive multiple PWM signals.

5.3.14.4 Nominal Fan Speed

A fan domain's nominal fan speed can be configured as static (fixed value) or controlled by the state of one or more associated temperature sensors.

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

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

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

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

5.3.14.5 Thermal and Acoustic Management

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

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

5.3.14.6 Thermal Sensor Input to Fan Speed Control

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

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

- Baseboard temperature sensors
- CPU DTS-Spec margin sensors
- DIMM thermal margin sensors

- Exit air temperature sensor
- PCH Temperature sensor
- Global aggregate thermal margin sensors
- SSB (Intel® C612 Series Chipset) temperature sensor
- On-board Ethernet controller temperature sensors (support for this is specific to the Ethernet controller being used)
- On-board SAS controller temperature sensors
- CPU VR Temperature sensor
- DIMM VR Temperature sensor
- BMC Temperature sensor
- DIMM VRM Temperature sensor

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

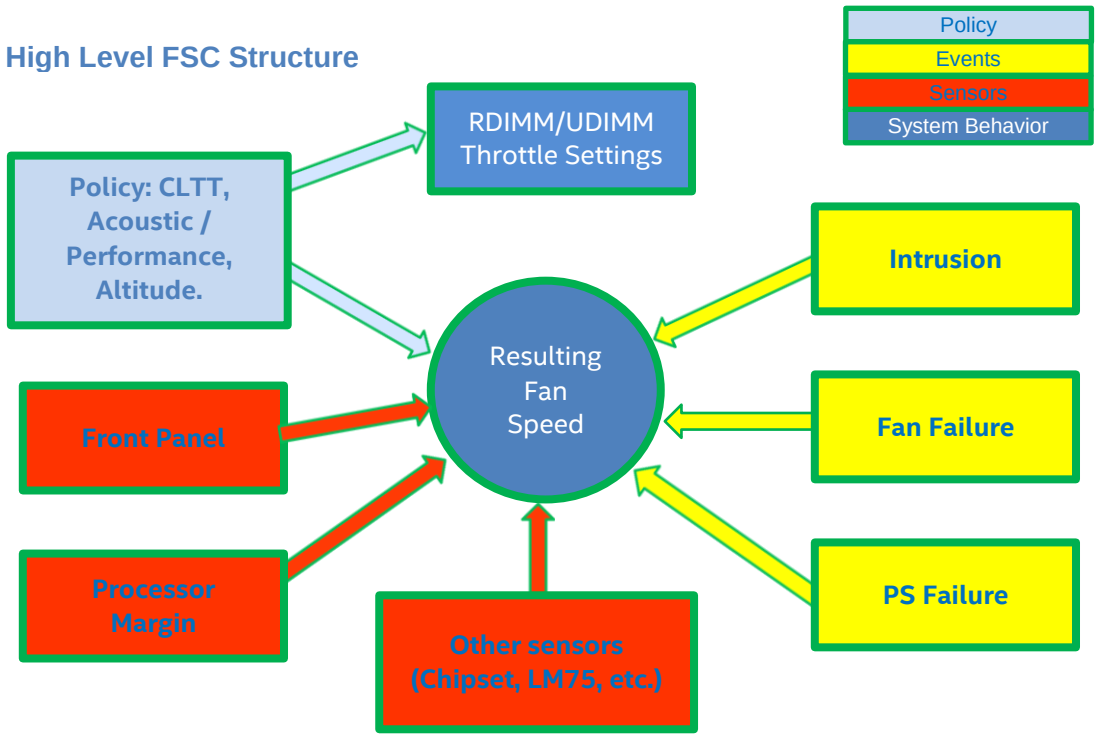

Figure 18. High-level Fan Speed Control Process

5.3.14.6.1 Processor Thermal Management

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

5.3.14.6.2 Memory Thermal Management

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

5.3.14.6.2.1 Memory Thermal Throttling

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

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

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

Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family introduce a new type of CLTT which is referred to as Hybrid CTLL for which the integrated Memory controller estimates the DRAM temperature in between actual reads of the TSODs. Hybrid CLTT shall be used on all Intel® Server Systems supporting the Intel® Xeon® processor E5-2600 v3 product family that have DIMMs with thermal sensors. Therefore, the terms Dynamic-CLTT and Static-CLTT are really referring to this “hybrid” mode. Note that if the IMC’s polling of the TSODs is interrupted, the temperature readings that the BMC gets from the IMC shall be these estimated values.

5.3.14.6.3 DIMM Temperature Sensor Input to Fan Speed Control

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

5.3.14.6.4 Dynamic (Hybrid) CLTT

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

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

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

5.3.14.6.5 Autoprofile

Intel server board implements auto-profile feature to improve upon previous platform configuration-dependent FSC and maintain competitive acoustics within the market. This feature is not available for third party customization.

BIOS and BMC will handshake to automatically understand configuration details and automatically select the optimal fan speed control profile in the BMC.
Customers will only select a performance or an acoustic profile selection from the BIOS menu for EPSD system 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 recommend if customer installed MICs or any other high power add-in cards (higher than 75W) or PCIe add-in cards which requires excessive cooling.

5.3.14.6.6 ASHRAE Compliance

Auto-profile algorithm will be implemented for PCSD products from this generation. There will be no manual selection of profiles at different altitudes, but altitude impact will be well covered by auto-profile.

5.3.14.7 Power Supply Fan Speed Control

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

5.3.14.7.1 System Control of Power Supply Fans

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

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

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

Some products require that the power supply internal thermal sensors be 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.

5.3.14.8 Fan Boosting due to Fan Failures

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

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

5.3.14.9 Programmable Fan PWM Offset

The system provides a BIOS Setup option to boost the system fan speed by a programmable positive offset or a “Max” setting. Setting the programmable offset causes the BMC to add the offset to the fan speeds 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 35°C.

5.3.15 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 Website http://www.powersig.org/.

5.3.16 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.

5.3.17 Component Fault LED Control

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

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

- **Fan fault LEDs** – A fan fault LED is associated with each fan. The BMC lights a fan fault LED if the associated fan tach sensor has a lower critical threshold event status

asserted. Fan tach sensors are manual re-arm sensors. Once the lower critical threshold is crossed, the LED remains lit until the sensor is rearmed. These sensors are rearmed at system DC power-on and system reset.

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

Table 12. Component Fault LEDs

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

5.3.18 NMI (Diagnostic Interrupt) Sensor

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

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

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

5.3.19 LAN Leash Event Monitoring

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

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

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

5.3.20 Add-in Module Presence Sensor

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

5.3.21 CMOS Battery Monitoring

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

5.4 Embedded Web Server

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

5.5 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.

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

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

The following table identifies both Basic and Advanced server management features.

Table 13. 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
Chassis Intrusion Detection	X	X
Fan Redundancy Monitoring	X	X
Hot-Swap Fan Support	X	X

Feature	Basic	Advanced
Acoustic Management	X	X
Diagnostic Beep Code Support	X	X
Power State Retention	X	X
ARP/DHCP Support	X	X
PECI Thermal Management Support	X	X
E-mail Alerting	X	X
Embedded Web Server	X	X
SSH Support	X	X
Integrated KVM		X
Integrated Remote Media Redirection		X
Lightweight Directory Access Protocol (LDAP)	X	X
Intel® Intelligent Power Node Manager Support	X	X
SMASH CLP	X	X

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

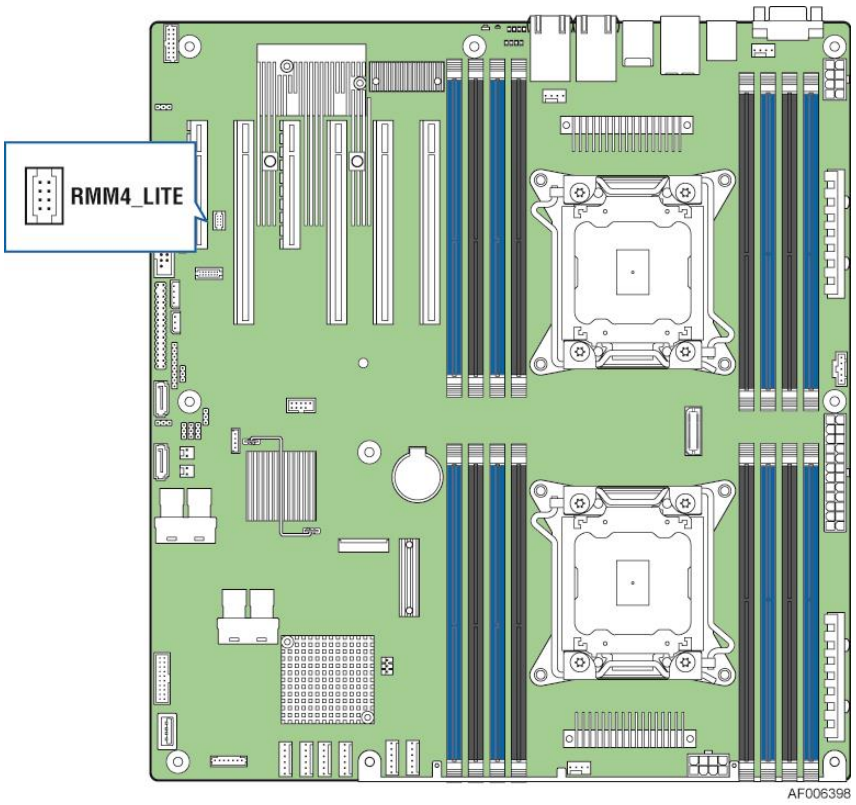

Figure 19. Intel® RMM4 Lite Activation Key Location

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

Key Features of the RMM4 add-on are:

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

5.5.1 Keyboard, Video, Mouse (KVM) Redirection

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

The BMC supports an embedded KVM application (*Remote Console*) that can be launched from the embedded web server from a remote console. USB1.1 or USB 2.0 based mouse and keyboard redirection are supported. It is also possible to use the KVM-redirection (KVM-r) session concurrently with media-redirection (media-r). This feature allows a user to interactively use the keyboard, video, and mouse (KVM) functions of the remote server as if the user were physically at the managed server. KVM redirection console supports the following keyboard layouts: English, Dutch, French, German, Italian, Russian, and Spanish.

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

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

Other attributes of this feature include:

- Encryption of the redirected screen, keyboard, and mouse
- Compression of the redirected screen
- Ability to select a mouse configuration based on the OS type
- 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+)

5.5.2 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.

5.5.3 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 2Mbps 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.

5.5.4 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.

5.5.5 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 FW update) will require the session to be re-established.

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

5.5.6 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.

5.5.7 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.

5.5.8 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 a local Floppy drive, a local USB device, or 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. A BMC reset (for example, due to an BMC reset after BMC FW update) will require the session to be re-established.
- The mounted device is visible to (and 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 install.

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

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

5.5.8.1 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.

5.5.8.2 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*.

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

Power management deals with requirements to manage processor power consumption and manage power at the platform level to meet critical business needs. Node Manager (NM) is a platform resident technology that enforces power capping and thermal-triggered power capping policies for the platform. These policies are applied by exploiting subsystem 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.

6.1 Hardware Requirements

NM is supported only on platforms that have the NM FW functionality loaded and enabled on the Management Engine (ME) in the SSB and that have a BMC present to support the external LAN interface to the ME. NM power limiting features require 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.

6.2 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.

6.3 ME System Management Bus (SMBus*) Interface

- The ME uses the SMLink0 on the SSB in multi-master mode as a dedicated bus for communication with the BMC using the IPMB protocol. The BMC FW considers this a secondary IPMB bus and runs at 400 kHz.
- The ME uses the SMLink1 on the SSB in multi-master mode bus for communication with PMBus* devices in the power supplies for support of various NM-related features. This bus is shared with the BMC, which polls these PMBus* power supplies for sensor monitoring purposes (for example, power supply status, input power, and so on). This bus runs at 100 KHz.
- The Management Engine has access to the “Host SMBus*”.

6.4 PECI 3.0

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

6.5 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 whether 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 the installed power supplies do in fact support PMBus*. If the installed power supplies are not PMBus*-compliant, the NM “discovery” OEM SDR should not be loaded.

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

6.6 SmaRT/CLST

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

System power control modes include:

- SmaRT: Low AC input voltage event; results in a one-time momentary throttle for each event to the maximum throttle state.
- 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.
- 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 HW for a short period (~100ms) to reduce overall power consumption. It is expected that the interruption to the fans will be of short enough duration to avoid false lower threshold crossings for the fan tach sensors; however, this may need to be comprehended by the fan monitoring FW if it does have this side-effect.

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

6.6.1 Dependencies on PMBus*-compliant Power Supply Support

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

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

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

---

7. Intel® Server Board S2600CW Connector/Header Locations and Pin-outs

7.1 Power Connectors

7.1.1 Main Power Connector

Main server board power is supplied by one 12-pin power connector. The connector is labeled as “MAIN PWR” on the left bottom of the server board. The following table provides the pin-out for “MAIN PWR” connector.

Table 14. Main Power Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	P3V3	13	P3V3
2	P3V3	14	N12V
3	GND	15	GND
4	P5V	16	FM_PS_EN_PSU_ON
5	GND	17	GND
6	P5V	18	GND
7	GND	19	GND
8	PWRGD_PS_PWROK_PSU_R1	20	NC_PS_RES_TP
9	P5V_STBY_PSU	21	P5V
10	P12V	22	P5V
11	P12V	23	P5V
12	P3V3	24	GND

7.1.2 CPU Power Connectors

On the server board are two white 8-pin CPU power connectors labeled “CPU_1 PWR” and “CPU_2 PWR”. The following table provides the pin-out for both connectors.

Table 15. CPU_1 Power Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	GND	5	P12V1
2	GND	6	P12V1
3	GND	7	P12V3A
4	GND	8	P12V3A

Table 16. CPU_2 Power Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	GND	5	P12V2
2	GND	6	P12V2

Pin	Signal Name	Pin	Signal Name
3	GND	7	P12V3B
4	GND	8	P12V3B

7.2 Front Panel Header and Connectors

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

7.2.1 Front Panel Header

Included on the left edge of the server board is a 24-pin SSI-compatible front panel header which provides various front panel features including:

- Power/Sleep Button
- System ID Button
- NMI Button
- NIC Activity LEDs
- Hard Drive Activity LEDs
- System Status LED
- System ID LED

The following table provides the pin-out for this 24-pin header.

Table 17. Front Panel Header Pin-out

Pin	Signal Name	Pin	Signal Name
1	P3V3_AUX	2	P3V3_AUX
3	Key	4	P5V_STBY
5	FP_PWR_LED_BUF_N	6	FP_ID_LED_BUF_N
7	P3V3	8	FP_LED_STATUS_GREEN_BUF_N
9	LED_HDD_ACTIVITY_N	10	FP_LED_STATUS_AMBER_BUF_N
11	FP_PWR_BTN_N	12	LED_NIC_LINK0_ACT_BUF_N
13	GND	14	LED_NIC_LINK0_LNKUP_BUF_N
15	FP_RST_BTN_N	16	SMB_SENSOR_3V3STBY_DATA
17	GND	18	SMB_SENSOR_3V3STBY_CLK
19	FP_ID_BTN_N	20	FP_CHASSIS_INTRUSION
21	PU_FM_SIO_TEMP_SENSOR	22	LED_NIC_LINK1_ACT_BUF_N
23	FP_NMI_BTN_N	24	LED_NIC_LINK1_LNKUP_BUF_N

7.2.2 Front Panel USB Connector

The server board includes a 20-pin connector, which when cabled, can provide up to two USB 3.0 ports to a front panel. The following table provides the connector pin-out.

Table 18. Front Panel USB 3.0 Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	P5V_AUX_USB_FP_USB3	key	KEY
2	USB3_01_FB_RX_DN	19	P5V_AUX_USB_FP_USB3
3	USB3_01_FB_RX_DP	18	USB3_00_FB_RX_DN
4	GND	17	USB3_00_FB_RX_DP
5	USB3_01_FB_TX_DN	16	GND
6	USB3_01_FB_TX_DP	15	USB3_00_FB_TX_DN
7	GND	14	USB3_00_FB_TX_DP
8	USB2_13_FB_DN	13	GND
9	USB2_13_FB_DP	12	USB2_8_FB_DN
10	TP_FM_OC5_FP_R_N	11	USB2_8_FB_DP

7.3 On-board Storage Connectors

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

7.3.1 SATA 6Gbps Connectors

The server board includes two 7-pin SATA connectors capable of transfer rates of up to 6Gbps. The following table provides the pin-out for both connectors.

Table 19. SATA 6Gbps Connector Pin-out

Pin	Signal Name
1	GND
2	SATA_TX_P
3	SATA_TX_N
4	GND
5	SATA_RX_N
6	SATA_RX_P
7	GND

The server board also includes two mini-SAS HD ports, each supporting four SATA 6Gbps transfer rates. The following table provides the pin-out for both connectors.

Table 20. Mini-SAS HD Connectors for SATA 6Gbps Pin-out

Pin	Signal Name	Pin	Signal Name
1A1	TP_SAS1_BACKPLANE_TYPE	2A1	TP_SAS0_BACKPLANE_TYPE
1B1	GND	2B1	GND
1C1	SGPIO_SSATA_DATAOUT0_R1	2C1	SGPIO_SATA_DATAOUT0_R1

Pin	Signal Name	Pin	Signal Name
1D1	PU_DATAIN1_SAS1	2D1	PU_DATAIN1_SAS0
1A2	SGPIO_SSATA_CLOCK_R1	2A2	SGPIO_SATA_CLOCK_R1
1B2	SGPIO_SSATA_LOAD_R1	2B2	SGPIO_SATA_LOAD_R1
1C2	GND	2C2	GND
1D2	PD_SAS1_CONTROLLER_TYPE	2D2	PD_SAS0_CONTROLLER_TYPE

7.3.2 SAS Connectors

The server boards S2600CW2S and S2600CWTS include two mini-SAS HD connectors supporting up to SAS 12Gbps transfer rates. The following table provides the pin-out for each connector.

Table 21. Mini-SAS HD Connectors for SAS 12Gbps Pin-out

Pin	Signal Name	Pin	Signal Name
1A1	TP_SAS1_BACKPLANE_TYPE	2A1	TP_SAS0_BACKPLANE_TYPE
1B1	GND	2B1	GND
1C1	SGPIO_SSATA_DATAOUT0_R1	2C1	SGPIO_SATA_DATAOUT0_R1
1D1	SGPIO_DATAIN	2D1	SGPIO_DATAIN
1A2	SGPIO_SSATA_CLOCK_R1	2A2	SGPIO_SATA_CLOCK_R1
1B2	SGPIO_SSATA_LOAD_R1	2B2	SGPIO_SATA_LOAD_R1
1C2	GND	2C2	GND
1D2	PD_SAS1_CONTROLLER_TYPE	2D2	PD_SAS0_CONTROLLER_TYPE

7.3.3 HSBP I²C Header

Table 22. HSBP I²C Header Pin-out

Pin	Signal Name
1	SMB_HSBP_3V3STBY_DATA
2	GND
3	SMB_HSBP_3V3STBY_CLK

7.3.4 HDD LED Header

The server board includes a 2-pin hard drive activity LED header used with some SAS/SATA controller add-in cards. The header has the following pin-out.

Table 23. HDD LED Header Pin-out

Pin	Signal Name	Pin	Signal Name
1	LED_HDD_ACT_N	2	NA

7.3.5 Internal Type-A USB Connector

The server board includes one internal Type-A USB connector. The following table provides the pin-out for this connector.

Table 24. Type-A USB Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	P5V	2	USB2_P2_F_DN
3	USB2_P2_F_DP	4	GND

7.3.6 Internal eUSB SSD Header

The server board includes one 10-pin internal eUSB header with an intended usage of supporting USB SSD devices. The following table provides the pin-out for this connector.

Table 25. eUSB SSD Header Pin-out

Pin	Signal Name	Pin	Signal Name
1	5V	2	NC
3	USB2_PCH_P12_DN	4	NC
5	USB2_PCH_P12_DP	6	NC
7	GND	8	NC
9	Key	10	LED_HDD_ACT_ZEPHER_N

7.3.7 M.2/NGFF Header

The server board includes one M.2/NGFF header. The following table provides the pin-out for this connector.

Table 26. M.2/NGFF Header Pin-out

Pin	Signal Name	Pin	Signal Name
		75	GND
74	3.3v	73	GND
72	3.3v	71	GND
70	3.3v	69	PEDET(GND-SATA)
68	SUSCKL(32kHZ)(1)(0/3.3v)	67	N/C
66	Key	65	Key
64	Key	63	Key

Pin	Signal Name	Pin	Signal Name
62	Key	61	Key
60	Key	59	Key
58	Reserved	57	GND
56	Reserved	55	N/C
54	N/C	53	N/C
52	N/C	51	GND
50	N/C	49	SATA-A+
48	N/C	47	SATA-A-
46	N/C	45	GND
44	N/C	43	SATA-B-
42	N/C	41	SATA-B+
40	N/C	39	GND
38	DEVSLP(1)(0/3.3v)	37	N/C
36	N/C	35	N/C
34	N/C	33	GND
32	N/C	31	N/C
30	N/C	29	N/C
28	N/C	27	GND
26	N/C	25	N/C
24	N/C	23	N/C
22	N/C	21	GND
20	N/C	19	N/C
18	3.3v	17	N/C
16	3.3v	15	GND
14	3.3v	13	N/C
12	3.3v	11	N/C
10	DAS/DSS#(O)(OD)	9	GND
8	N/C	7	N/C
6	N/C	5	N/C
4	3.3v	3	GND
2	3.3v	1	GND

7.4 Management and Security Connectors

7.4.1 RMM4 Lite Connector

A 7-pin Intel® RMM4 Lite connector 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.

Table 27. RMM4 Lite Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	P3V3_AUX	2	DI
3	KEY	4	CLK
5	DO	6	GND
7	CS_N	8	GND

7.4.2 TPM Connector

Table 28. TPM Connector Pin-out

Pin	Signal Name	Pin	Signal Name
1	Key	2	LPC_LAD<1>
3	LPC_LAD<0>	4	GND
5	IRQ_SERIAL	6	LPC_FRAME_N
7	P3V3	8	GND
9	RST_IBMC_NIC_N_R2	10	CLK_33M_TPM
11	LPC_LAD<3>	12	GND
13	GND	14	LPC_LAD<2>

7.4.3 PMBus* Connector

Table 29. PMBus* Connector Pin-out

Pin	Signal Name
1	SMB_PMBUS_CLK_R
2	SMB_PMBUS_DATA_R
3	IRQ_SML1_PMBUS_ALERT_RC_N
4	GND
5	P3V3

7.4.4 Chassis Intrusion Header

The server board includes a 2-pin chassis intrusion header which can be used when the chassis is configured with a chassis intrusion switch. The header has the following pin-out.

Table 30. Chassis Intrusion Header Pin-out

Header State	Description
Pins 1 and 2 closed	FM_INTRUDER_HDR_N is pulled HIGH. Chassis cover is closed.
Pins 1 and 2 open	FM_INTRUDER_HDR_N is pulled LOW. Chassis cover is removed.

7.4.5 IPMB Connector

Table 31. IPMB Connector Pin-out

Pin	Signal Name
1	SMB_IPMB_5VSTBY_DATA
2	GND
3	SMB_IPMB_5VSTBY_CLK
4	P5V_STBY

7.5 FAN Connectors

The server board provides support for nine fans. Seven of them are system cooling fans, and two of them are CPU fans. The expected maximum RPM is 25000.

7.5.1 System FAN Connectors

The six system cooling fan connectors near the front edge of the board are 6-pin connectors; the one system cooling fan connector near the edge of the board is a 4-pin connector. The following table provides the pin-out for all system fan connectors.

Table 32. 6-pin System FAN Connector Pin-out

Pin	Signal Name
1	GND
2	12V
3	TACH
4	PWM
5	PRSNT
6	FAULT

Table 33. 4-pin System FAN Connector Pin-out

Pin	Signal Name
1	GND
2	12V
3	TACH
4	PWM

7.5.2 CPU FAN Connector

The two CPU fan connectors are 4-pin fan connectors. The following table provides the pin-out for CPU fan connectors.

Table 34. CPU FAN Connector Pin-out

Pin	Signal Name
1	GND
2	12V
3	TACH
4	PWM

7.6 Serial Port and Video Connectors

7.6.1 Serial Port Connector

The server board includes one internal DH-10 serial port connector.

Table 35. Serial Port B Connector Pin-out

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

7.6.2 Video Connector

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

Table 36. Video Connector Pin-out

Pin	Signal Name
1	CRT_RED
2	CRT_GREEN
3	CRT_BLUE
4	N/C
5	GND
6	GND
7	GND
8	GND
9	P5V
10	GND
11	NC
12	CRT_DDCDATA
13	CRT_HSYNC
14	CRT_VSYNC
15	CRT_DDCCLK

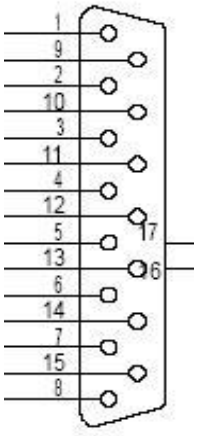

Figure 20. Video Connector Pin-out

***Note**: Intel Corporation server boards support peripheral components and can contain a number of high-density VLSI and power delivery components that need adequate airflow to cool. Intel's own chassis are designed and tested to meet the intended thermal requirements of these components when the fully integrated system is used together. It is the responsibility of the system integrator that chooses not to use Intel developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of its published operating or non-operating limits.*

7.7 PCIe Riser Slot

The following table provides the pin-out for PCIe slot 6 as a riser slot.

PIN	IO	SIGNAL DESCRIPTION Riser Define (PCIe Spec)	PIN	IO	SIGNAL DESCRIPTION Riser Define (PCIe Spec)
A1	NC	TP_PRSNT1_N	B1	PWR	P12V
A2	PWR	P12V	B2	PWR	P12V
A3	PWR	P12V	B3	PWR	P12V (RSVD)
A4	GND	GND	B4	GND	GND
A5	PWR	P3V3 (TCK)	B5	I	SMB_CLK_3VSB_PCI
A6	I	TDI	B6	IO	SMB_DAT_3VSB_PCI
A7	I	SAS_MODULE_ENABLE (TDO)	B7	GND	GND
A8	PWR	P3V3 (TMS)	B8	PWR	P3V3
A9	PWR	P3V3	B9	I	TRST#
A10	PWR	P3V3	B10	PWR	P3V3_AUX
A11	I	PERST_N	B11	O	WAKE_N
	KEY			KEY	

	KEY			KEY	
A12	GND	GND	B12	PWR	FM_THROTLE_SLOT6_N (RSVD)
A13	I	REFCLKP_1	B13	GND	GND
A14	I	REFCLKN_1	B14	I	PETP0
A15	GND	GND	B15	I	PETN0
A16	O	PERP0	B16	GND	GND
A17	O	PERN0	B17	I	MUX_RST_N (PRSNT2_N)
A18	GND	GND	B18	GND	GND
A19	PWR	P3V3 (RSVD)	B19	I	PETP1
A20	GND	GND	B20	I	PETN1
A21	O	PERP1	B21	GND	GND
A22	O	PERN1	B22	GND	GND
A23	GND	GND	B23	I	PETP2
A24	GND	GND	B24	I	PETN2
A25	O	PERP2	B25	GND	GND
A26	O	PERN2	B26	GND	GND
A27	GND	GND	B27	I	PETP3
A28	GND	GND	B28	I	PETN3
A29	O	PERP3	B29	GND	GND
A30	O	PERN3	B30	PWR	P3V3 (RSVD)
A31	GND	GND	B31	O	PD via 4.75Kohm (PRNST2_N)
A32	I	REFCLKP_2 (RSVD)	B32	GND	GND
A33	I	REFCLKN_2 (RSVD)	B33	I	PETP4
A34	GND	GND	B34	I	PETN4
A35	O	PERP4	B35	GND	GND
A36	O	PERN4	B36	GND	GND
A37	GND	GND	B37	I	PETP5
A38	GND	GND	B38	I	PETN5
A39	O	PERP5	B39	GND	GND
A40	O	PERN5	B40	GND	GND
A41	GND	GND	B41	I	PETP6
A42	GND	GND	B42	I	PETN6
A43	O	PERP6	B43	GND	GND
A44	O	PERN6	B44	GND	GND
A45	GND	GND	B45	I	PETP7
A46	GND	GND	B46	I	PETN7
A47	O	PERP7	B47	GND	GND
A48	O	PERN7	B48	O	LINK_WIDTH_ID0 (PRSNT2_N)
A49	GND	GND	B49	GND	GND
A50	O	LINK_WIDTH_ID1 (RSVD)	B50	I	PETP8

A51	GND	GND	B51	I	PETN8
A52	O	PERP8	B52	GND	GND
A53	O	PERN8	B53	GND	GND
A54	GND	GND	B54	I	PETP9
A55	GND	GND	B55	I	PETN9
A56	O	PERP9	B56	GND	GND
A57	O	PERN9	B57	GND	GND
A58	GND	GND	B58	I	PETP10
A59	GND	GND	B59	I	PETN10
A60	O	PERP10	B60	GND	GND
A61	O	PERN10	B61	GND	GND
A62	GND	GND	B62	I	PETP11
A63	GND	GND	B63	I	PETN11
A64	O	PERP11	B64	GND	GND
A65	O	PERN11	B65	GND	GND
A66	GND	GND	B66	I	PETP12
A67	GND	GND	B67	I	PETN12
A68	O	PERP12	B68	GND	GND
A69	O	PERN12	B69	GND	GND
A70	GND	GND	B70	I	PETP13
A71	GND	GND	B71	I	PETN13
A72	O	PERP13	B72	GND	GND
A73	O	PERN13	B73	GND	GND
A74	GND	GND	B74	I	PETP14
A75	GND	GND	B75	I	PETN14
A76	O	PERP14	B76	GND	GND
A77	O	PERN14	B77	GND	GND
A78	GND	GND	B78	I	PETP15
A79	GND	GND	B79	I	PETN15
A80	O	PERP15	B80	GND	GND
A81	O	PERN15	B81	I	NC (PRSNT2_N)
A82	GND	GND	B82	I	NC (RSVD)

PCIe riser slot reuses the same mechanical slot as the standard PCIe x16 slot and only redefines several RSVD, PRSNT and JTAG pins for dedicated riser usage. Standard PCIe X16 card is also supported in the riser slot.

A summary of all changes made on the PCIe x16 riser slot compared with industry PCIe Spec:

- Reuse Pin B3 as the P12V Power.
- Reuse Pin A5/A8/A19/B30 as P3V3 Power.

- Reuse Pin A7 as SAS_MODULE_ENABLE. This change applies to standard PCIe slot as well
- Reuse Pin B17 as MUX_RST_N.
- Reuse Pin B48/A50/B31 as LINK_WIDTH_ID.
- Reuse A32/A33 as additional PCIe Clocks to the 2nd slot on the risers.

NOTES:

- The 3rd REFCLK is not used.
- LINK_WIDTH_ID2 is not used on slot 6, which is pulled to GND via a 4.75Kohm resistor.
- For pin B12, it is connected to system throttling signal to throttle input for Xeon Phi.

8. Intel® Server Board S2600CW Jumper Blocks

The server boards have several 3-pin jumper blocks that you can use to configure, protect, or recover specific features of the server boards.

The following symbol identifies Pin 1 on each jumper block on the silkscreen: ▼

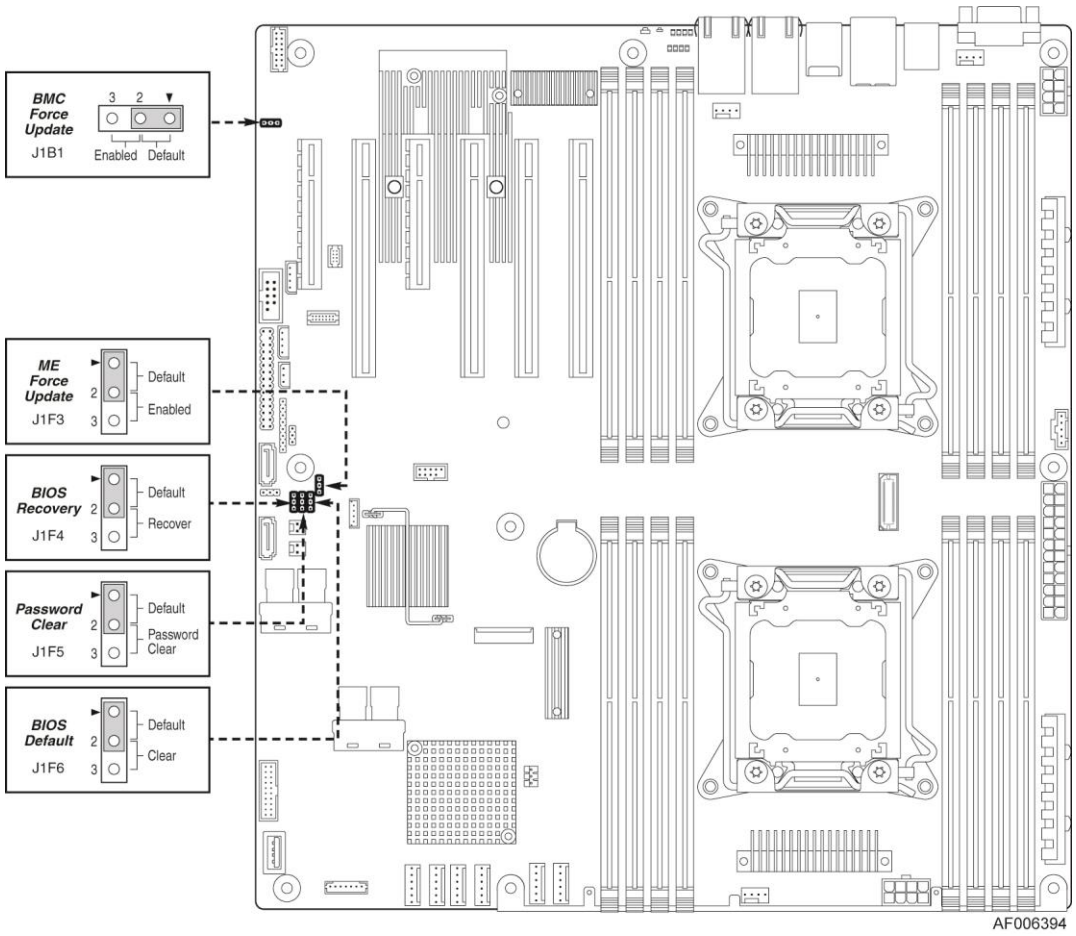

Figure 21. Jumper Blocks

Table 37. Server Board Jumpers

Jumper Name	Pins	System Results
BIOS Recovery	1-2	Pins 1-2 should be connected for normal system operation. **(Default)**
	2-3	The main system BIOS does not boot with pins 2-3 connected. The system only boots from EFI-bootable recovery media with a recovery BIOS image present.
BIOS Default (that is, CMOS Clear)	1-2	These pins should have a jumper in place for normal system operation. **(Default)**
	2-3	If pins 2-3 are connected when AC power unplugged, the CMOS settings clear in 5 seconds. Pins 2-3 should not be connected for normal system operation.
ME Force Update	1-2	ME Firmware Force Update Mode – Disabled **(Default)**
	2-3	ME Firmware Force Update Mode – Enabled

Jumper Name	Pins	System Results
BMC Force Update	1-2	BMC Firmware Force Update Mode – Disabled **(Default)**
	2-3	BMC Firmware Force Update Mode – Enabled
Password Clear	1-2	These pins should have a jumper in place for normal system operation. **(Default)**
	2-3	To clear administrator and user passwords, power on the system with pins 2-3 connected. The administrator and user passwords clear in 5-10 seconds after power on. Pins 2-3 should not be connected for normal system operation.

8.1 BIOS Default and Password Reset Usage Procedure

The BIOS Default and Password Reset recovery features are designed such that the desired operation can be achieved with minimal system downtime. The usage procedure for these two features has changed from previous generation Intel® Server Boards. The following procedure outlines the new usage model.

8.1.1 Set BIOS to Default (Clearing the CMOS)

To clear the CMOS, perform the following steps:

1. Power down the server. Unplug the power cord.
2. Open the server chassis. For instructions, see your server chassis documentation.
3. Move jumper from the default operating position (covering pins 1 and 2) to the reset/clear position (covering pins 2 and 3).
4. Wait five seconds.
5. Move the jumper back to the default position (covering pins 1 and 2).
6. Close the server chassis.
7. Install AC power cord.
8. Power up the server and access the BIOS setup utility by <F2>
9. Verify the BIOS default operation was successful by view the Error Manager screen. Two errors should be logged:

   5220 BIOS Settings reset to default settings
   0012 System RTC date/time not set

The CMOS is now cleared and can be reset by going into the BIOS setup.

***Note**: This jumper does not reset Administrator or User passwords. In order to reset passwords, the Password Clear jumper must be used. The system will automatically power on after AC is applied to the system.*

8.1.2 Clearing the Password

This jumper causes both the User password and the Administrator password to be cleared if they were set. The operator should be aware that this creates a security gap until passwords have been installed again through the <F2> BIOS Setup utility. This is the only method by which the Administrator and User passwords can be cleared unconditionally. Other than this

jumper, passwords can only be set or cleared by changing them explicitly in BIOS Setup or by similar means. No method of resetting BIOS configuration settings to default values will affect either the Administrator or User passwords.

To clear the password, perform the following steps:

1. Power down the server. For safety, unplug the power cords.
2. Remove the system top cover.
3. Move the "*Password Clear*" jumper from the default operating position (covering pins 1 and 2) to the password clear position (covering pins 2 and 3).
4. Reinstall the system top cover and reattach the power cords.
5. Power up the server and access the <F2> BIOS Setup utility.
6. Verify the password clear operation was successful by viewing the Error Manager screen. Two errors should be logged:

```
5221  Passwords cleared by jumper
5224  Password clear jumper is set
```

7. Exit the BIOS Setup utility and power down the server. For safety, remove the AC power cords.
8. Remove the system top cover and move the "*Password Clear*" jumper back to the default operating position (covering pins 1 and 2).
9. Reinstall the system top cover and reattach the AC power cords.
10. Power up the server.

The password is now cleared and can be reset by going into the BIOS setup.

8.2 Integrated BMC Force Update Procedure

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

1. Power down and remove the AC power cord.
2. Open the server chassis. For instructions, see your server chassis documentation.
3. Move jumper from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Close the server chassis.
5. Reconnect the AC cord and power up the server.

6. Perform the Integrated BMC firmware update procedure as documented in the README.TXT file that is included in the given Integrated BMC firmware update package. After successful completion of the firmware update process, the firmware update utility may generate an error stating that the Integrated BMC is still in update mode.
7. Power down and remove the AC power cord.
8. Open the server chassis.
9. Move jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
10. Close the server chassis.
11. Reconnect the AC cord and power up the server.

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

8.3 ME Force Update Jumper

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

1. Power down and remove the AC power cord.
2. Open the server chassis. For instructions, see your server chassis documentation.
3. Move jumper from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).
4. Close the server chassis.
5. Reconnect the AC cord and power up the server.
6. Boot to the EFI shell and update the ME firmware using the “MEComplete.cap” file using the following command:

   ```
   iflash32 /u /ni MEComplete.cap
   ```
7. Power down and remove the AC power cord.
8. Open the server chassis.
9. Move jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).
10. Close the server chassis.
11. Reconnect the AC cord and power up the server.

8.4 BIOS Recovery Jumper

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

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

9. Intel® Light Guided Diagnostics

Both server boards have 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 boards.

9.1 5-volt Stand-by LED

Several server management features of these server boards require a 5-V stand-by voltage supplied from the power supply. The features and components that require this voltage must be present when the system is powered-down. The LED is illuminated when AC power is applied to the platform and 5-V stand-by voltage is supplied to the server board by the power supply.

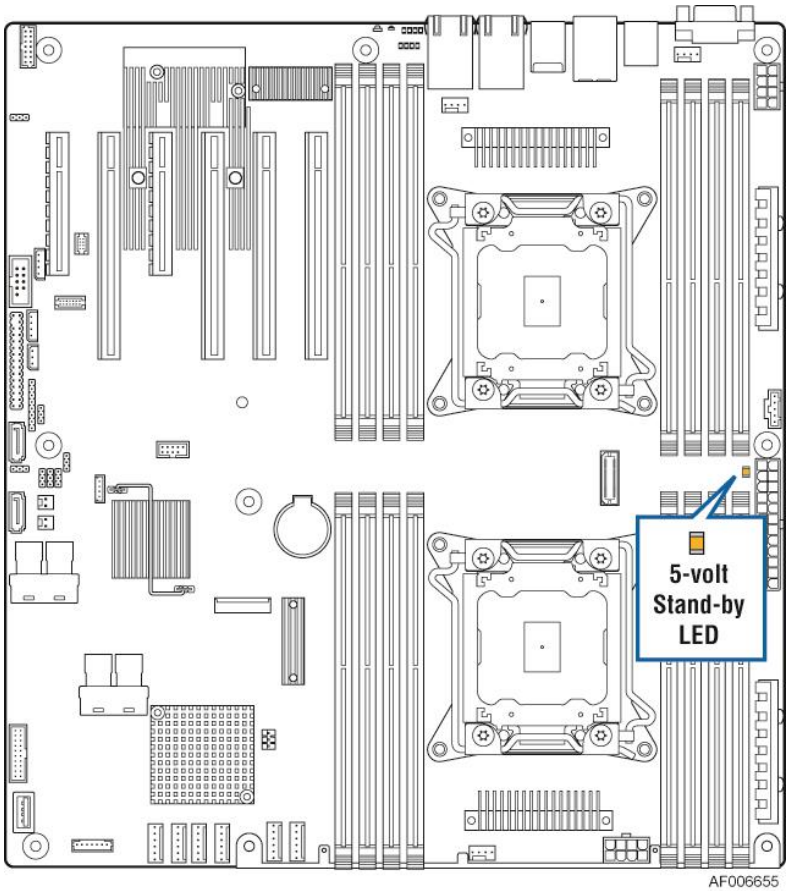

Figure 22. 5-volt Stand-by Status LED Location

9.2 Fan Fault LEDs

Fan fault LEDs are present for the two CPU fans. The fan fault LEDs illuminate when the corresponding fan has fault.

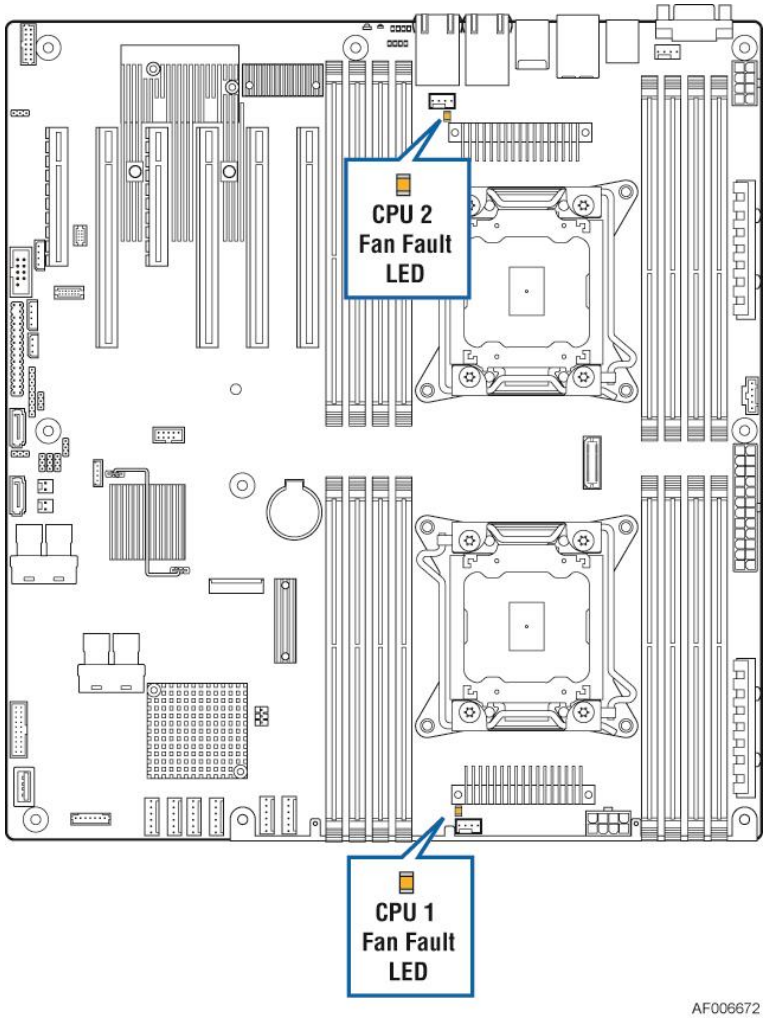

Figure 23. Fan Fault LED's Location

9.3 DIMM Fault LEDs

The server board provides memory fault LED for each DIMM socket. These LEDs are located as shown in the following figure. The DIMM fault LED illuminates when the corresponding DIMM slot has memory installed and a memory error occurs.

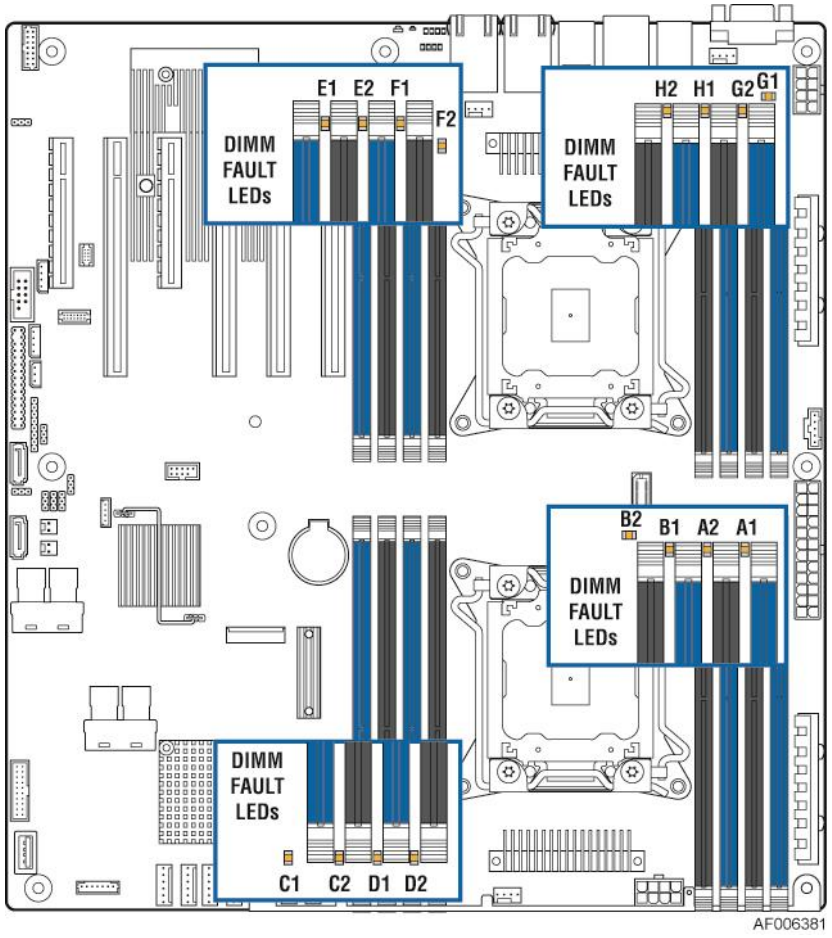

Figure 24. DIMM Fault LED's Location

9.4 System ID LED, System Status LED, and POST Code Diagnostic LEDs

The server board provides LEDs for system ID, system status, and POST code. These LEDs are located in the rear I/O area of the server board as shown in the following figure.

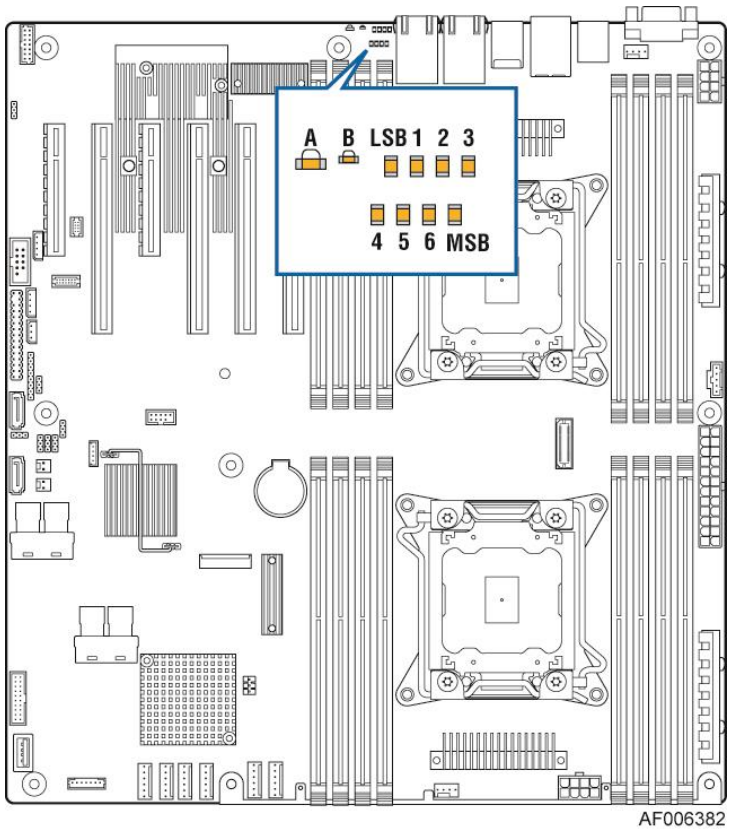

Callout	Description
A	System Status LED
B	System ID LED
LSB 1 2 3 4 5 6 MSB	POST Code Diagnostic LEDs

Figure 25. Location of System Status, System ID, and POST Code Diagnostic LEDs

9.4.1 System ID LED

You can illuminate the blue System ID LED using either of the following two mechanisms:

- By pressing the System ID Button on the system front control panel, the ID LED displays a solid blue color until the button is pressed again.
- By issuing the appropriate hex IPMI "Chassis Identify" value, the ID LED will either blink blue for 15 seconds and turn off or will blink indefinitely until the appropriate hex IPMI Chassis Identify value is issue to turn it off.

9.4.2 System Status LED

The bi-color (green/amber) System Status LED operates as follows.

Table 38. System Status LED

Color	State	Criticality	Description
Green	Solid on	Ok	Indicates that the System Status is "Healthy". The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running.
Green	~1 Hz blink	Degraded	**System degraded:** 1. Redundancy loss such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. 2. Fan warning or failure when the number of fully operational fans is less than minimum number needed to cool the system. 3. 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. 4. Power supply predictive failure occurred while redundant power supply configuration was present. 5. Unable to use all of the installed memory (more than 1 DIMM installed). 6. Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the user no longer has spared DIMMs indicating a redundancy lost condition. Corresponding DIMM LED lit. 7. In mirrored configuration, when memory mirroring takes place and system loses memory redundancy. 8. Battery failure. 9. BMC executing in uBoot. (Indicated by Chassis ID blinking at 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. 10. BMC booting Linux*. (Indicated by Chassis ID solid ON). System in degraded state (no manageability). Control has been passed from BMC uBoot to BMC Linux* itself. It will be in this state for ~10-~20 seconds. 11. BMC Watchdog has reset the BMC. 12. Power Unit sensor offset for configuration error is asserted. 13. HDD HSC is off-line or degraded.
Amber	~1 Hz blink	Non-critical	**Non-fatal alarm – system is likely to fail:** 1. Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply and PROCHOT (Therm Ctrl) sensors. 2. VRD Hot asserted. 3. Minimum number of fans to cool the system not present or failed.

Color	State	Criticality	Description
			4. Hard drive fault. 5. Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present). 6. In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window.
Amber	Solid on	Critical, non-recoverable	**Fatal alarm – system has failed or shutdown:** 1. CPU CATERR signal asserted. 2. MSID mismatch detected (CATERR also asserts for this case). 3. CPU 1 is missing. 4. CPU ThermalTrip. 5. No power good – power fault. 6. DIMM failure when there is only 1 DIMM present and hence no good memory present. 7. Runtime memory uncorrectable error in non-redundant mode[1]. 8. DIMM Thermal Trip or equivalent. 9. SSB Thermal Trip or equivalent. 10. CPU ERR2 signal asserted. 11. BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition). 12. Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition). 13. 240VA fault
Off	N/A	Not ready	AC power off

Note:

* When the server is powered down (transitions to the DC-off state or S5), the BMC is still on standby power and retains the sensor and front panel status LED state established before the power-down event. If the system status is normal when the system is powered down (the LED is in a solid green state), the system status LED is off.

9.4.3 POST Code Diagnostic LEDs

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 boards. To assist in troubleshooting a system hang during the POST process, you can use the diagnostic LEDs to identify the last POST process executed.

Table 39. POST Code Diagnostic LEDs

A. Diagnostic LED #7 (MSB LED)	E. Diagnostic LED #3
B. Diagnostic LED #6	F. Diagnostic LED #2
C. Diagnostic LED #5	G. Diagnostic LED #1
D. Diagnostic LED #4	H. Diagnostic LED #0 (LSB LED)

10. Power Supply Specification Guidelines

10.1 Power System Options Overview

The Intel® Server Board S2600CW can work with 550-W fixed power supply, or 750-W/1600-W redundant power supplies shipped with the Intel® Server Chassis.

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 from a 750W power supply and the Power Distribution Board as an option used in an Intel® Server Chassis. The intent of this section is to provide customers with a guide to assist in defining and/or selecting a power supply for custom server platform designs that utilize the server board detailed in this document.*

10.2 750-W Power Supply

This specification defines a 750W redundant power supply that supports server systems. This power supply has 2 outputs; 12V and 12V standby. The AC input is auto ranging and power factor corrected.

10.2.1 Mechanical Overview

The physical size of the power supply enclosure is 39/40mm x 74mm x 185mm. The power supply contains a single 40mm fan. The power supply has a card edge output that interfaces with a 2x25 card edge connector in the system. The AC plugs directly into the external face of the power supply. Refer to the following figure. All dimensions are nominal.

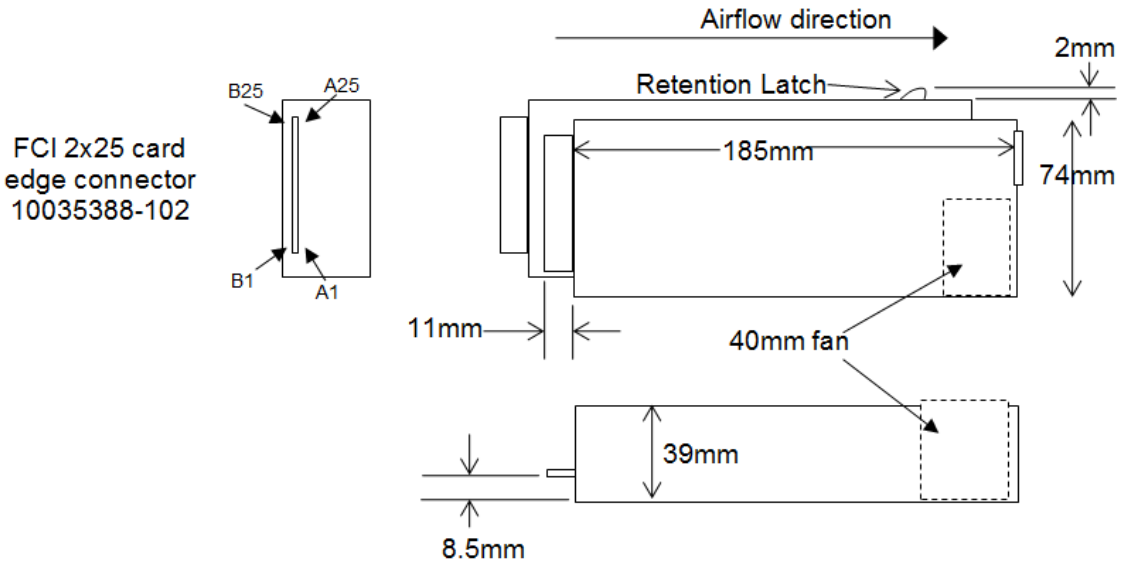

Figure 26. 750-W Power Supply Outline Drawing

10.2.1.1 DC Output Connector

The power supply uses a card edge output connection for power and signal that is compatible with a 2x25 Power Card Edge connector (equivalent to 2x25 pin configuration of the FCI power card connector 10035388-102LF).

Table 40. DC Output Connector

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

10.2.1.2 Handle Retention

The power supply has a handle to assist extraction. The module is able to be inserted and extracted without the assistance of tools. The power supply has a latch which retains the power supply into the system and prevents the power supply from being inserted or extracted from the system when the AC power cord is pulled into the power supply.

The handle protects the operator from any burn hazard.

10.2.1.3 LED Marking and Identification

The power supply uses a bi-color LED: Amber and Green. Below are table showing the LED states for each power supply operating state and the LED's wavelength characteristics.

Refer to the *Intel® LED Wavelength and Intensity Specification* for more details.

Table 41. LED Characteristics

	Min λd Wavelength	Nominal λd Wavelength	Max λd Wavelength	Units
Green	562	565	568	nm
Amber	607	610	613	nm

Table 42. Power Supply LED Functionality

Power Supply Condition	LED State
Output ON and OK.	GREEN
No AC power to all power supplies.	OFF
AC present/Only 12VSB on (PS off) or PS in Cold redundant state.	1Hz Blink GREEN
AC cord unplugged or AC power lost; with a second power supply in parallel still with AC input power.	AMBER
Power supply warning events where the power supply continues to operate; high temp, high power, high current, slow fan.	1Hz Blink Amber
Power supply critical event causing a shutdown; failure, OCP, OVP, Fan Fail.	AMBER
Power supply FW updating.	2Hz Blink GREEN

10.2.1.4 Temperature Requirements

The power supply operates within all specified limits over the T_{op} temperature range. All airflow passes through the power supply and not over the exterior surfaces of the power supply.

Table 43. Environmental Requirements

Item	Description	Min	Max	Units
$T_{op\_sc\_red}$	Operating temperature range; spreadcore redundant (60% load, 3000m, spreadcore system flow impedance[1])	0	60	°C
$T_{op\_sc\_nr}$	Operating temperature range; spreadcore non-redundant (100% load, 3000m, spreadcore system flow impedance[1])	0	50	°C
$T_{op\_rackped\_900}$	Operating temperature range; rack/pedestal 900m (100% load, 900m, rack/pedestal system flow impedance[1])	0	45	°C
$T_{op\_rackped\_3000}$	Operating temperature range; rack/pedestal 3000m (100% load, 3000m, rack/pedestal system flow impedance[1])	0	40	°C

Texit	Maximum exit air temperature		68	°C
T_{non-op}	Non-operating temperature range	-40	70	°C
Altitude	Maximum operating altitude[2]		3050	m

Notes:

1. Under normal conditions, the exit air temperature shall be less than 65°C. 68°C is provided for absolute worst case conditions and is expected only to exist when the inlet ambient reaches 60°C.
2. $T_{op\_rackped\_900}$ condition only requires maximum altitude of 900m.

The power supply meets UL enclosure requirements for temperature rise limits. All sides of the power supply, with exception to the air exhaust side, are classified as "Handle, knobs, grips, and so on", and held for short periods of time only.

10.2.2 AC Input Requirements

10.2.2.1 Power Factor

The power supply meets the power factor requirements stated in the *Energy Star* Program Requirements for Computer Servers*. These requirements are stated below.

Table 44. Power Factor Requirements for Computer Servers

Output Power	10% Load	20% Load	50% Load	100% Load
Power factor	> 0.65	> 0.80	> 0.90	> 0.95

Tested at 230VAC, 50Hz and 60Hz and 115VAC, 60Hz

Tested according to *Generalized Internal Power Supply Efficiency Testing Protocol Rev 6.4.3*. This is posted at http://efficientpowersupplies.epri.com/methods.asp.

10.2.2.2 AC Inlet Connector

The AC input connector is an *IEC 320 C-14* power inlet. This inlet is rated for 10A/250VAC.

10.2.2.3 AC Input Voltage Specification

The power supply operates within all specified limits over the following input voltage range. Harmonic distortion of up to 10% of the rated line voltage does not cause the power supply to go out of specified limits. Application of an input voltage below 85VAC does not cause damage to the power supply, including a blown fuse.

Table 45. AC Input Voltage Range

Parameter	MIN	Rated	V_{MAX}	Start up VAC	Power Off VAC
Voltage (110)	90 V_{rms}	100-127 V_{rms}	140 V_{rms}	85VAC +/-4VAC	70VAC +/-5VAC
Voltage (220)	**180 V_{rms}**	200-240 V_{rms}	264 V_{rms}		
Frequency	47 Hz	50/60	63 Hz		

Notes:

1. Maximum input current at low input voltage range shall be measured at 90VAC, at maximum load.
2. Maximum input current at high input voltage range shall be measured at 180VAC, at maximum load.
3. This requirement is not to be used for determining agency input current markings.

10.2.2.4 AC Line Dropout/Holdup

An AC line dropout is defined as that when the AC input drops to 0VAC at any phase of the AC line for any length of time. During an AC dropout, the power supply meets dynamic voltage regulation requirements. An AC line dropout of any duration does not cause tripping of control signals or protection circuits. If the AC dropout lasts longer than the holdup time, the power supply recovers and meets all turn on requirements. The power supply meets the AC dropout requirement over rated AC voltages and frequencies. A dropout of the AC line for any duration does not cause damage to the power supply.

Table 46. AC Line Holdup Time

Loading	Holdup Time
70%	12msec

10.2.2.4.1 AC Line 12VSB Holdup

The 12VSB output voltage stays in regulation under its full load (static or dynamic) during an AC dropout of **70ms min** (=12VSB holdup time) whether the power supply is in ON or OFF state (PSON asserted or de-asserted).

10.2.2.5 AC Line Fuse

The power supply has one line fused in the **single line fuse** on the line (Hot) wire of the AC input. The line fusing is acceptable for all safety agency requirements. The input is a slow blow type. AC inrush current does not cause the AC line fuse to blow under any conditions. All protection circuits in the power supply does not cause the AC fuse to blow unless a component in the power supply has failed. This includes DC output load short conditions.

10.2.2.6 AC Line Transient Specification

AC line transient conditions are defined as “sag” and “surge” conditions. “Sag” conditions are also commonly referred to as “brownout”; these conditions are defined as the AC line voltage drops below nominal voltage conditions. “Surge” is defined to refer to conditions when the AC line voltage rises above nominal voltage.

The power supply meets the requirements under the following AC line sag and surge conditions.

Table 47. AC Line Sag Transient Performance

AC Line Sag (10sec interval between each sagging)				
Duration	Sag	Operating AC Voltage	Line Frequency	Performance Criteria
0 to ½ AC cycle	95%	Nominal AC Voltage ranges	50/60Hz	No loss of function or performance
> 1 AC cycle	>30%	Nominal AC Voltage ranges	50/60Hz	Loss of function acceptable, self recoverable

Table 48. AC Line Surge Transient Performance

AC Line Surge				
Duration	Surge	Operating AC Voltage	Line Frequency	Performance Criteria
Continuous	10%	Nominal AC Voltages	50/60Hz	No loss of function or performance
0 to ½ AC cycle	30%	Mid-point of nominal AC Voltages	50/60Hz	No loss of function or performance

10.2.2.7 Power Recovery

The power supply shall recover automatically after an AC power failure. AC power failure is defined to be any loss of AC power that exceeds the dropout criteria.

10.2.3 Efficiency

The following table provides the required minimum efficiency level at various loading conditions. These are provided at three different load levels; 100%, 50%, 20%, and 10%. Output shall be loaded according to the proportional loading method defined by 80 Plus in *Generalized Internal Power Supply Efficiency Testing Protocol Rev. 6.4.3*. This is posted at http://efficientpowersupplies.epri.com/methods.asp.

Table 49. Silver Efficiency Requirement

Loading	100% of Maximum	50% of Maximum	20% of Maximum	10% of Maximum
Minimum Efficiency	91%	94%	90%	82%

The power supply passes with enough margins to make sure that all power supplies meet these efficiency requirements in production.

10.2.4 DC Output Specification

10.2.4.1 Output Power/Currents

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

Table 50. Minimum Load Ratings

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

10.2.4.2 Standby Output

The 12VSB output is present when an AC input greater than the power supply turn on voltage is applied.

10.2.4.3 Voltage Regulation

The power supply output voltages 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 51. Voltage Regulation Limits

Parameter	Tolerance	Min	Nom	Max	Units
+12V	-5%/+5%	+11.40	+12.00	+12.60	V_{rms}
+12V stby	-5%/+5%	+11.40	+12.00	+12.60	V_{rms}

10.2.4.4 Dynamic Loading

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

Table 52. Transient Load Requirements

Output	Δ Step Load Size (See note 2)	Load Slew Rate	Test Capacitive Load
+12VSB	1.0A	0.25 A/μsec	20 μF
+12V	60% of max load	0.25 A/μsec	2000 μF

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

10.2.4.5 Capacitive Loading

The power supply is stable and meets all requirements with the following capacitive loading ranges.

Table 53. Capacitive Loading Conditions

Output	Min	Max	Units
+12VSB	20	3100	μF
+12V	500	25000	μF

10.2.4.6 Grounding

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

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

10.2.4.7 Residual Voltage Immunity in Standby Mode

The power supply is immune to any residual voltage placed on its outputs (Typically a leakage voltage through the system from standby output) up to 500mV. There is neither additional heat generated, nor stressing of any internal components with this voltage applied to any individual or all outputs simultaneously. It also does not trip the protection circuits during turn on.

The residual voltage at the power supply outputs for no load condition does not exceed 100mV when AC voltage is applied and the PSON# signal is de-asserted.

10.2.4.8 Common Mode Noise

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

The measurement is made across a 100Ω resistor between each of DC outputs, including ground at the DC power connector and chassis ground (power subsystem enclosure).

The test setup shall use a FET probe such as Tektronix* model P6046 or equivalent.

10.2.4.9 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 remains within the limits with the capacitive load specified. The hot swap test is conducted when the system is operating under static, dynamic, and zero loading conditions. The power supply uses a latching mechanism to prevent insertion and extraction of the power supply when the AC power cord is inserted into the power supply.

10.2.4.10 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 does not affect the load sharing or output voltages of the other supplies still operating. The supplies are 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.

10.2.4.11 Ripple/Noise

The maximum allowed ripple/noise output of the power supply is defined in the table below. This is measured over a bandwidth of 10Hz to 20MHz at the power supply output connectors.

A 10μF tantalum capacitor in parallel with a 0.1μF ceramic capacitor is placed at the point of measurement.

Table 54. Ripples and Noise

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

The test setup shall be as shown below.

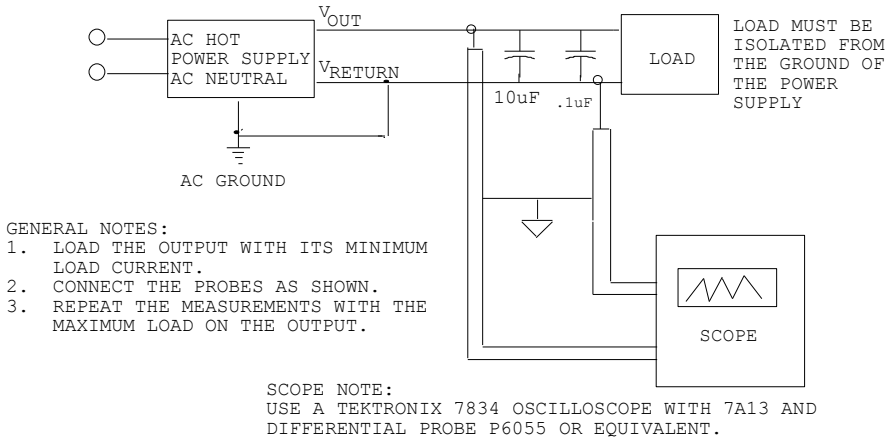

Figure 27. Differential Noise Test Setup

***Note**: When performing this test, the probe clips and capacitors should be located close to the load.*

10.2.4.12 Timing Requirements

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

Table 55. Timing Requirements

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 12VSB 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

Item	Description	Min	Max	Units
$T_{pson\_on\_delay}$	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 12VSB being in regulation to O/Ps being in regulation at AC turn on.	50	1000	ms
T_{12VSB_holdup}	Time the 12VSB output voltage stays within regulation after loss of AC.	70		ms

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

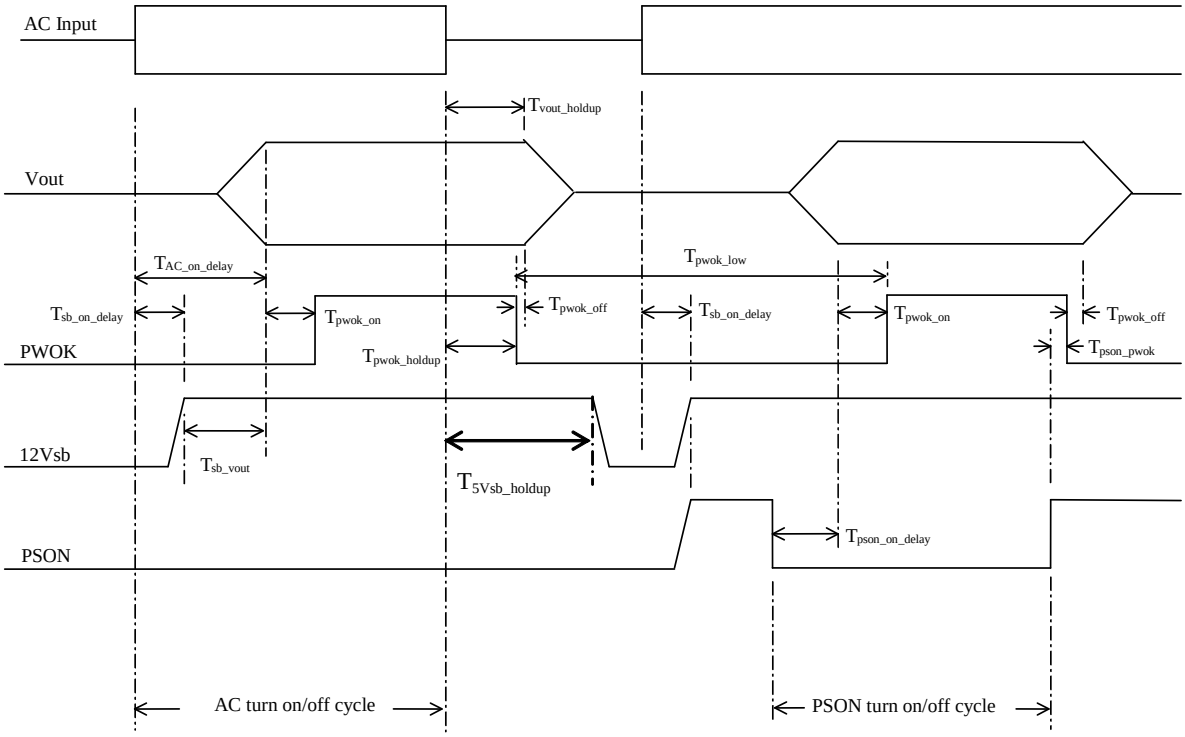

Figure 28. Turn On/Off Timing (Power Supply Signals)

10.2.5 Protection Circuits

Protection circuits inside the power supply causes only the power supply's main outputs to shut down. If the power supply latches-off due to a protection circuit tripping, an AC cycle OFF for 15sec and a PSON# cycle HIGH for one second are able to reset the power supply.

10.2.5.1 Current Limit (OCP)

The power supply has current limit to prevent the outputs from exceeding the values shown in table below. If the current limits are exceeded the power supply shuts down and latches-off. The latch will be cleared by toggling the PSON# signal or by an AC power interruption. The power supply is not damaged from repeated power cycling in this condition. 12VSB will be auto-recovered after removing OCP limit.

Table 56. Over Current Protection

Output Voltage	Input Voltage Range	Over Current Limits
+12V	90 – 264VAC	72A min; 78A max
12VSB	90 – 264VAC	2.5A min; 3.5A max

10.2.5.2 Over Voltage Protection (OVP)

The power supply over voltage protection is locally sensed. The power supply shuts down and latches off after an over voltage condition occurs. This latch is cleared by toggling the PSON# signal or by an AC power interruption. The values are measured at the output of the power supply's connectors. The voltage does not exceed the maximum levels when measured at the power connectors of the power supply connector during any single point of fail. The voltage does not trip any lower than the minimum levels when measured at the power connector. 12VSB will be auto-recovered after removing OVP limit.

Table 57. Over Voltage Protection (OVP) Limits

Output voltage	Min (v)	Max (v)
+12V	13.3	14.5
+12VSB	13.3	14.5

10.2.5.3 Over Temperature Protection (OTP)

The power supply will be protected against over temperature conditions caused by loss of fan cooling or excessive ambient temperature. In an OTP condition, the PSU will shut down. When the power supply temperature drops to within specified limits, the power supply shall restore power automatically, while the 12VSB remains always on. The OTP circuit must have built in margin such that the power supply will not oscillate on and off due to temperature recovering condition. The OTP trip level shall have a minimum of 4°C of ambient temperature margin.

10.2.6 Control and Indicator Functions

The following sections define the input and output signals from the power supply. Signals that can be defined as low true use the following convention: Signal# = low true.

10.2.6.1 PSON# Input Signal

The PSON# signal is required to remotely turn on/off the power supply. PSON# is an active low signal that turns on the +12V power rail. When this signal is not pulled low by the system, or left open, the outputs (except the +12VSB) turn off. This signal is pulled to a standby voltage by a pull-up resistor internal to the power supply. Refer to Table 55 for the timing diagram.

Table 58. PSON# Signal Characteristic

Signal Type	Accepts an open collector/drain input from the system. Pull-up to VSB located in power supply.	
PSON# = Low	ON	
PSON# = High or Open	OFF	
	MIN	MAX
Logic level low (power supply ON)	0V	1.0V
Logic level high (power supply OFF)	2.0V	3.46V
Source current, Vpson = low		4mA
Power up delay: $T_{pson\_on\_delay}$	5msec	400msec
PWOK delay: T_{pson_pwok}		50msec

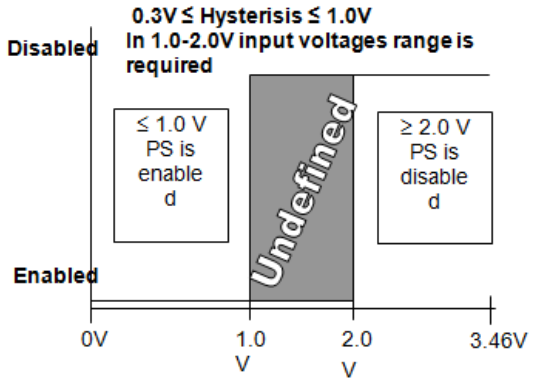

Figure 29. PSON# Required Signal Characteristic

10.2.6.2 PWOK (Power OK) Output Signal

PWOK is a power OK signal and will be pulled HIGH by the power supply to indicate that all the outputs are within the regulation limits of the power supply. When any output voltage falls below regulation limits or when AC power has been removed for a time sufficiently long so that power supply operation is no longer guaranteed, PWOK will be de-asserted to a LOW state. See the following table for a representation of the timing characteristics of PWOK. The start of the PWOK delay time shall be inhibited as long as any power supply output is in current limit.

Table 59. PWOK Signal Characteristics

Signal Type	Open collector/drain output from power supply. Pull-up to VSB located in the power supply.	
PWOK = High	Power OK	
PWOK = Low	Power Not OK	
	MIN	MAX
Logic level low voltage, Isink=400uA	0V	0.4V
Logic level high voltage, Isource=200μA	2.4V	3.46V
Sink current, PWOK = low		400uA
Source current, PWOK = high		2mA
PWOK delay: T_{pwok_on}	100ms	1000ms
PWOK rise and fall time		100μsec
Power down delay: T_{pwok_off}	1ms	200msec

A recommended implementation of the Power Ok circuits is shown below.

***Note**: the Power Ok circuits should be compatible with 5V pull up resistor (>10k) and 3.3V pull up resistor (>6.8k).*

10.2.6.3 SMBAlert# Signal

This signal indicates that the power supply is experiencing a problem that the user should investigate. This shall be asserted due to Critical events or Warning events. The signal shall be activated in case the critical component temperature reaches a warning threshold, general failure, over-current, over-voltage, under-voltage, or fan failure. This signal may also indicate the power supply is reaching its end of life or is operating in an environment exceeding the specified limits.

This signal is to be asserted in parallel with LED turning solid Amber or blink Amber.

Table 60. SMBAlert# Signal Characteristics

Signal Type (Active Low)	Open collector/drain output from power supply. Pull-up to VSB located in system.	
Alert# = High	OK	
Alert# = Low	Power Alert to system	
	MIN	MAX
Logic level low voltage, Isink=4 mA	0 V	0.4 V
Logic level high voltage, Isink=50 μA		3.46 V
Sink current, Alert# = low		4 mA
Sink current, Alert# = high		50 μA
Alert# rise and fall time		100 μs

10.2.7 Thermal CLST

The power supply shall assert the SMBAlert signal when a temperature sensor crosses a warning threshold. Refer to the *Intel® Common Hardware and Firmware Requirements for CRPS Power Supplier* for detailed requirements.

10.2.8 Power Supply Diagnostic "Black Box"

The power supply saves the latest PMBus* data and other pertinent data into nonvolatile memory when a critical event shuts down the power supply. This data is accessible by the SMBus* interface with an external source providing power to the 12Vstby output.

Refer to the *Intel® Common Hardware and Firmware Requirements for CRPS Power Supplier* for detailed requirements.

10.2.9 Firmware Uploader

The power supply has the capability to update its firmware by the PMBus* interface while it is in standby mode. This FW can be updated when in the system and in standby mode and outside the system with power applied to the 12Vstby pins.

Refer to the *Intel® Common Hardware and Firmware Requirements for CRPS Power Supplier* for detailed requirements.

10.3 Higer Power Common Redundant Power Distribution Board (PDB)

The Power Distribution Board (PDB) for the Intel® Server Chassis P4000M supports the Common Redundant power supply in a 1+1 redundant configuration. The PDB is designed to plug directly to the output connector of the PS and it contains 3 DC/DC power converters to produce other required voltages: +3.3VDC, +5VDC, and 5V standby along with additional over current protection circuit for the 12V rails.

The Intel® Server Chassis P4304XXMUXX family includes this PDB.

10.3.1 Mechanical Overview

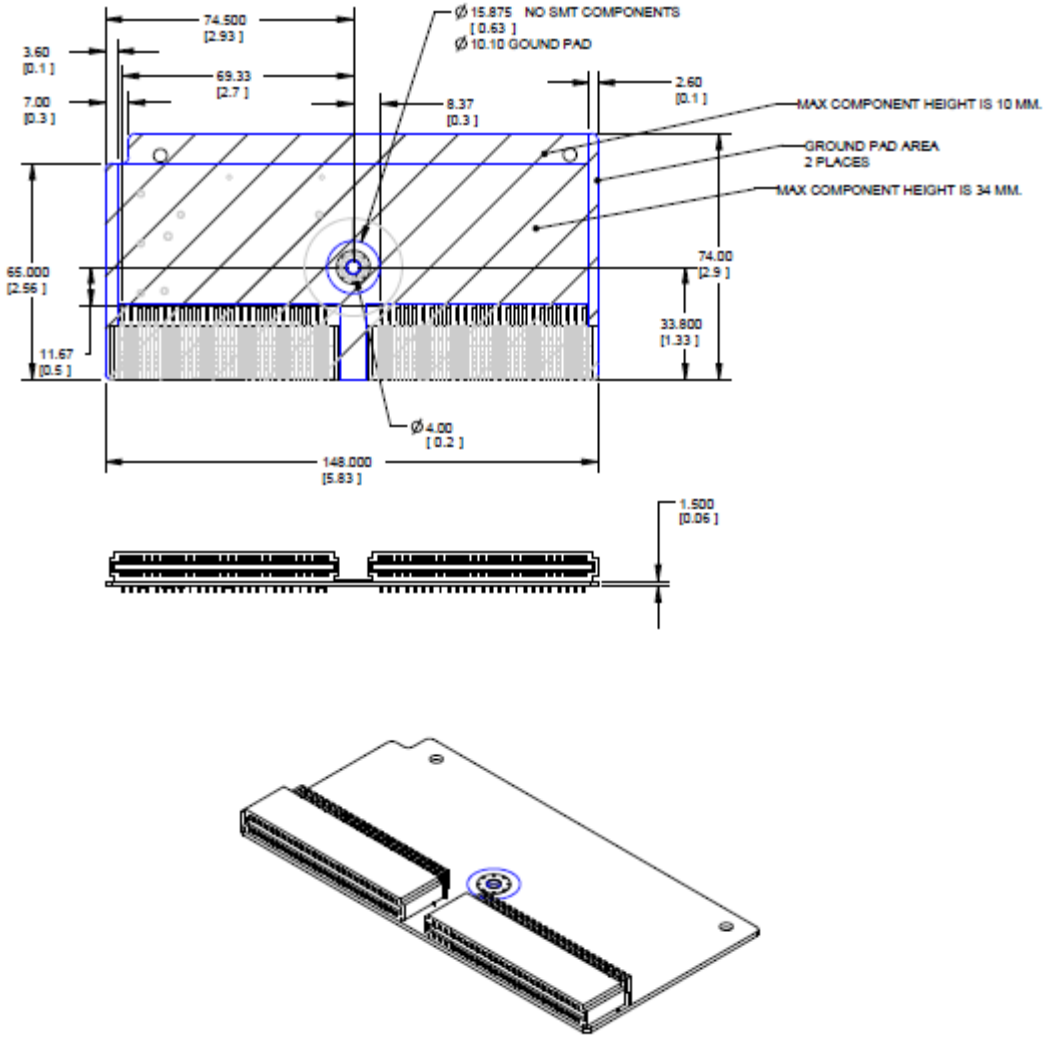

Figure 30. Outline Drawing

10.3.1.1 Airflow Requirements

The power distribution board shall get enough airflow for cooling DC/DC converters from the fans located in the Power Supply modules. Below is a basic drawing showing airflow direction.

The amount of cooling airflow that will be available to the DC/DC converters is to be no less than 1.2M/s.

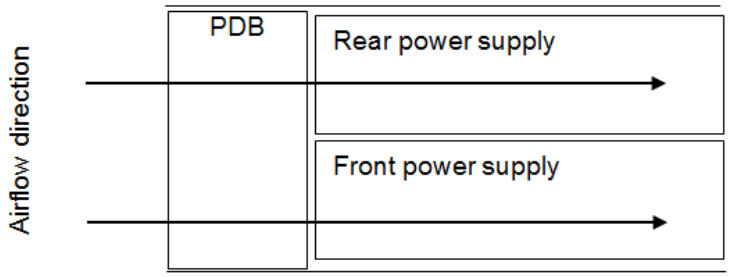

Figure 31. Airflow Diagram

10.3.1.2 DC/DC Converter Cooling

The DC/DC converters on the power distribution board are in series with the airflow path with the power supplies.

10.3.1.3 Temperature Requirements

The PDB operates within all specified limits over the Top temperature range. Some amount of airflow shall pass over the PDB.

Table 61. Thermal Requirements

Item	Description	Min	Max	Units
T_{op}	Operating temperature range.	0	50	°C
T_{non-op}	Non-operating temperature range.	-40	70	°C

10.3.1.4 Efficiency

Each DC/DC converter shall have a **minimum** efficiency of **85%** at 50% ~ 100% loads and over +12V line voltage range and over temperature and humidity range.

10.3.2 DC Output Specification

10.3.2.1 Input Connector (Power Distribution Mating Connector)

The power distribution provides two power pin, a card edge output connection for power and signal that is compatible with a 2x25 Power Card Edge connector (equivalent to 2x25 pin configuration of the FCI power card connector 10035388-102LF). The FCI power card edge connector is a new version of the PCE from FCI used to raise the card edge by 0.031" to allow for future 0.093" PCBs in the system. The card edge connector has no keying features; the keying method is accomplished by the system sheet metal.

Table 62. Input Connector and Pin Assignment

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

*The compatibility Pin is used for soft compatibility check. The two compatibility pins are connected directly.

10.3.2.2 Output Wire Harness

The power distribution board has a wire harness output with the following connectors.

Listed or recognized component appliance wiring material (AVLV2), CN, rated min 85°C shall be used for all output wiring.

Table 63. PDB Cable Length

From	Length, mm	To connector #	No of pins	Description
Power Supply cover exit hole	280	P1	24	Baseboard Power Connector
Power Supply cover exit hole	300	P2	8	Processor 0 connector
Power Supply cover exit hole	500	P3	8	Processor 1 connector
Power Supply cover exit hole	900	P4	5	Power FRU/PMBus* connector

From	Length, mm	To connector #	No of pins	Description
Power Supply cover exit hole	500	P5	5	SATA peripheral power connector for 5.25”
Extension from P5	100	P6	5	SATA peripheral power connector for 5.25”
Extension from P6	100	P7	4	Peripheral Power Connector for 5.25”/HSBP Power
Power Supply cover exit hole	600	P8	4	1x4 legacy HSBP Power Connector
Extension from P8	75	P9	4	1x4 legacy HSBP Power Connector
Power supply cover exit hole	700	P10	4	1x4 legacy HSBP Power/Fixed HDD adaptor Connection
Extension from P10	75	P11	4	1x4 legacy HSBP Power/Fixed HDD adaptor Connection
Connector only (no cable)	N/a	P12	4	Aux baseboard power connector for PCIe slots
Connector only (no cable)	N/a	P13	4	GFX card aux connectors
Connector only (no cable)	N/a	P14	4	
Connector only (no cable)	N/a	P15	4	
Connector only (no cable)	N/a	P16	4	

10.3.2.2.1 Baseboard Power Connector (P1)

- Connector housing: 24-pin Molex* Mini-Fit Jr. 39-01-2245 or equivalent
- Contact: Molex* Mini-Fit, HCS Plus, Female, Crimp 44476 or equivalent

Table 64. P1 Baseboard Power Connector

Pin	Signal	18 AWG Color	Pin	Signal	18 AWG Color
1	+3.3VDC	Orange	13	+3.3VDC	Orange
	3.3V RS	Orange (24AWG)			
2	+3.3VDC	Orange	14	-12VDC	Blue
3	COM	Black	15	COM	Black
4	+5VDC	Red	16	PSON#	Green (24AWG)
5	COM	Black	17	COM	Black
6	+5VDC	Red	18	COM	Black
7	COM	Black	19	COM	Black
8	PWR OK	Gray (24AWG)	20	Reserved	N.C.
9	5 VSB	Purple	21	+5VDC	Red
10	+12V1	Yellow	22	+5VDC	Red
11	+12V1	Yellow	23	+5VDC	Red
12	+3.3VDC	Orange	24	COM	Black

10.3.2.2.2 Processor#0 Power Connector (P2)

- Connector housing: 8-pin Molex* 39-01-2080 or equivalent
- Contact: Molex* Mini-Fit, HCS Plus, Female, Crimp 44476 or equivalent

Table 65. P0 Processor Power Connector

Pin	Signal	18 AWG color	Pin	Signal	18 AWG Color
1	COM	Black	5*	+12V1	White
2	COM	Black	6	+12V1	White
3	COM	Black	7	+12V1	White
4	COM	Black	8	+12V1	White

10.3.2.2.3 Processor#1 Power Connector (P3)

- Connector housing: 8-pin Molex* 39-01-2080 or equivalent
- Contact: Molex* Mini-Fit, HCS Plus, Female, Crimp 44476 or equivalent

Table 66. P1 Processor Power Connector

Pin	Signal	18 AWG color	Pin	Signal	18 AWG Color
1	COM	Black	5	+12V1	Brown
2	COM	Black	6	+12V1	Brown
3	COM	Black	7	+12V1	Brown
4	COM	Black	8	+12V1	Brown

10.3.2.2.4 Power Signal Connector (P4)

- Connector housing: 5-pin Molex* 50-57-9405 or equivalent
- Contacts: Molex* 16-02-0087 or equivalent

Table 67. Power Signal Connector

Pin	Signal	24 AWG Color
1	I^2C Clock	White
2	I^2C Data	Yellow
3	SMBAlert#	Red
4	COM	Black
5	3.3RS	Orange

10.3.2.2.5 2x2 12V Connector (P12-P16)

Connector header: Foxconn* p/n HM3502E-P1 or equivalent

Table 68. P12 12V Connectors

Pin	Signal	18 AWG color	Pin	Signal	18 AWG Color
1	COM	Black	5	+12V1	Yellow
2	COM	Black	6	+12V1	Yellow

Table 69. P13-P16 12V Connectors

Pin	Signal	18 AWG color	Pin	Signal	18 AWG Color
1	COM	Black	5	+12V2	Yellow
2	COM	Black	6	+12V2	Yellow

10.3.2.2.6 Legacy 1x4 Peripheral Power Connectors (P7, P8, P9, P10)

- Connector housing: Molex* 0015-24-4048 or equivalent
- Contact: Molex* 0002-08-1201 or equivalent

Table 70. P8, P9 Legacy Peripheral Power Connectors

Pin	Signal	18 AWG Color
1	+12V3	Green
2	COM	Black
3	COM	Black
4	+5 VDC	Red

Table 71. P7, P10, P11 Legacy Peripheral Power Connectors

Pin	Signal	18 AWG Color
1	+12V3	Green
2	COM	Black
3	COM	Black
4	+5 VDC	Red

10.3.2.2.7 SATA 1x5 Peripheral Power Connectors (P5, P6)

- Connector housing: Molex* 0675-82-0000 or equivalent
- Contact: Molex* 0675-81-0000 or equivalent

Table 72. SATA Peripheral Power Connectors

Pin	Signal	18 AWG Color
1	+3.3VDC	Orange
2	COM	Black

Pin	Signal	18 AWG Color
3	+5VDC	Red
4	COM	Black
5	+12V2	Yellow

10.3.2.3 Grounding

The ground of the pins of the PDB output connectors provides the power return path. The output connector ground pins is connected to safety ground (PDB enclosure). This grounding is well designed to ensure passing the maximum allowed Common Mode Noise levels.

10.3.2.4 Remote Sense

Below is listed the remote sense requirements and connection points for all the converters on the PDB and the main 12V output of the power supply.

Table 73. Remote Sense Connection Points

Converter	+ Sense Location	- Sense Location
Power supply main 12V	On PDB	On PDB
12V/3.3V	P20 (1x5 signal connector)	P20 (1x5 signal connector)
12V/5V	On PDB	On PDB
12V/-12V	none	none
12Vstby/5Vstby	none	none

Table 74. Remote Sense Requirements

Characteristic	Requirement
+3.3V remote sense input impedance	200Ω (measure from +3.3V on P1 2x12 connector to +3.3V sense on P20 1x5 signal connector)
+3.3V remote sense drop	200mV (remote sense must be able to regulate out 200mV drop on the +3.3V and return path; from the 2x12 connector to the remote sense points)
Max remote sense current draw	< 5mA

10.3.2.5 12V Rail Distribution

The following table shows the configuration of the 12V rails and what connectors and components in the system they are powering.

Table 75. 12V Rail Distribution

	P2	P3	P12	P1	P8	P9	P10	P11	P5,6,7	P13	P14	P15	P16	P17	P18	P19	P20				
	2x4	2x4	2x2	2x12	1x4	1x4	1x4	1x4	(2) 1x5, 1x4	GPU1		GPU2		GPU3		GPU4			OCP		

	CPU 1	Memory1	CPU 2	Memory2	PCIe	Fans	Misc	HDD and peripherals					2x3	2x4	2x3	2x4	2x3	2x4	2x3	2x4	Total Current	Min	Nominal	Max
12V1	17.8 A	10.5 A	17.8 A	10.5 A	21.7 A	10.0 A	3.0 A														91 A	91	95.5	100
12V2													6.3 A	12.5 A	6.3 A	12.5 A	6.3 A	12.5 A	6.3 A	12.5 A	76 A	76	88	100
12V3								18.0 A													18 A	18	19	20

Note:

+12V current to PCIe slots may be supplied from four different connectors. 12V1 on P2, 12V2 on P3, 12V3 on P1, and 12V3 on P12. P12 is reserved for board that needs 4 x GPU cards powered. P1 is the main 12V power for PCIe slot; but additional 12V power can be connected to P2 and/or P3. The motherboard MUST NOT short any of the 12V rails or connectors together.

10.3.2.6 Hard Drive 12V Rail Configuration Options

The following table shows the hard drive configuration options using the defined power connectors. In some cases additional converter or "Y" cables are needed.

Table 76. Hard Drive 12V Rail Configuration Options

	P8	P9	P10	P11	P5	P6	P7
	1x4	1x4	1x4	1x4	1x5	1x5	1x4
	18						
3 x 2.5" 8xHDD BP	HDD1 8 x 2.5	HDD2 8 x 2.5	N/a	N/a	N/a	N/a	HDD3 8 x 2.5
2 x 3.5" 4xHDD BP	HDD1 4x3.5		HDD1 4x3.5		peripheral bay		
1 x 3.5" 8xHDD BP	HDD1 8x3.5		N/a	N/a	peripheral bay		
8 x 3.5" fixed SATA	2xfixed	2xfixed	2xfixed	2xfixed	peripheral bay		
8 x 3.5" fixed SAS	2xfixed	2xfixed	2xfixed	2xfixed	peripheral bay		

10.3.2.7 DC/DC Converters Loading

The following table defines power and current ratings of three DC/DC converters located on the PDB, each powered from +12V rail. The three converters meet both static and dynamic voltage regulation requirements for the minimum and maximum loading conditions.

Table 77. DC/DC Converters Load Ratings

	+12VDC Input DC/DC Converters		
	+3.3V Converter	+5V Converter	-12V Converter
MAX Load	25A	25A	0.5A
MIN Static/Dynamic Load	0A	0A	0A

	+12VDC Input DC/DC Converters		
	+3.3V Converter	+5V Converter	-12V Converter
Max Output Power	3.3V x25A =82.5W	5V x25A =125W	12V x0.5A =6W

10.3.2.8 5VSB Loading

There is also one DC/DC converter that converts the 12V standby into 5V standby.

Table 78. 5VSB Loading

	12V stby/5V stby DC/DC Converters
MAX Load	8A
MIN Static/Dynamic Load	0.1
Max Output Power	5V x8A =40W

10.3.2.9 DC/DC Converters Voltage Regulation

The DC/DC converters' output voltages stay within the following voltage limits when operating at steady state and dynamic loading conditions. These limits include the peak-peak ripple/noise specified in Table 82. The 3.3V and 5V outputs are measured at the remote sense point, all other voltages measured at the output harness connectors.

Table 79. Voltage Regulation Limits

Converter Output	Tolerance	Min	Nom	Max	Units
+ 3.3VDC	-4%/+5%	+3.20	+3.30	+3.46	VDC
+ 5VDC	-4%/+5%	+4.80	+5.00	+5.25	VDC
5Vstby	-4%/+5%	+4.80	+5.00	+5.25	VDC

10.3.2.10 DC/DC Converters Dynamic Loading

The output voltages remains within limits specified in table above for the step loading and capacitive loading specified in the table below. The load transient repetition rate is only a test specification. The Δ step load may occur anywhere within the minimum load to the maximum load shown in Table 77 and Table 78.

Table 80. Transient Load Requirements

Output	Max Δ Step Load Size	Max Load Slew Rate	Test Capacitive Load
+ 3.3VDC	5A	0.25 A/μs	250 μF
+ 5VDC	5A	0.25 A/μs	400 μF
+5Vsb	0.5A	0.25A/μs	20 μF

10.3.2.11 DC/DC Converter Capacitive Loading

The DC/DC converters are stable and meet all requirements with the following capacitive loading ranges. Minimum capacitive loading applies to static load only.

Table 81. Capacitive Loading Conditions

Converter Output	Min	Max	Units
+3.3VDC	250	6800	μF
+5VDC	400	4700	μF
5Vstby	20	350	μF

10.3.2.12 DC/DC Converters Closed Loop Stability

Each DC/DC converter is unconditionally stable under all line/load/transient load conditions including capacitive load ranges specified in Section 10.3.2.11. A minimum of: **45 degrees phase margin** and **-10dB-gain margin** is required. The PDB provides proof of the unit's closed-loop stability with local sensing through the submission of Bode plots. Closed-loop stability must be ensured at the maximum and minimum loads as applicable.

10.3.2.13 Common Mode Noise

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

- The measurement shall be made across a 100Ω resistor between each of DC outputs, including ground, at the DC power connector and chassis ground (power subsystem enclosure).
- The test set-up shall use a FET probe such as Tektronix* model P6046 or equivalent.

10.3.2.14 Ripple/Noise

The maximum allowed ripple/noise output of each DC/DC Converter is defined in the table below. This is measured over a bandwidth of 0Hz to 20MHz at the PDB output connectors. A 10μF tantalum capacitor in parallel with a 0.1μF ceramic capacitor are placed at the point of measurement.

Table 82. Ripple and Noise

+3.3V	+5V	-12V	+5VSB
50mVp-p	50mVp-p	120mVp-p	50mVp-p

The test setup shall be as shown below.

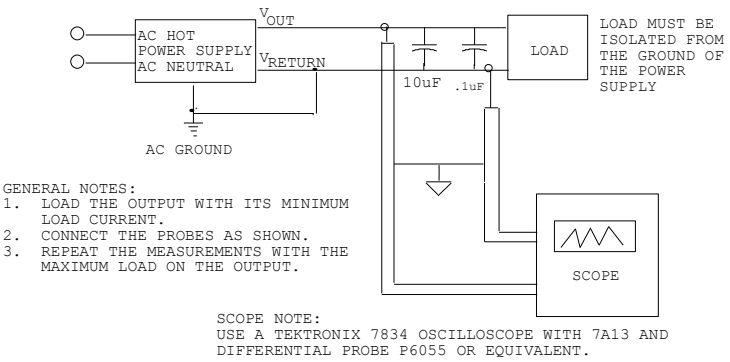

Note: When performing this test, the probe clips and capacitors should be located close to the load.

Figure 32. Differential Noise Test Setup

10.3.2.15 Timing Requirements

Below are timing requirements for the power on/off of the PDB DC/DC converters. The +3.3V, +5V and +12V output voltages should start to rise approximately at the same time. All outputs must rise monotonically.

Table 83. Output Voltage Timing

Description	Min	Max	Units
Output voltage rise time for each main output; 3.3V, 5V, -12V, and 5Vstby.	1.0	20	msec
The main DC/DC converters (3.3V, 5V, -12V) shall be in regulation limits within this time after the 12V input has reached 11.4V.		20	msec
The main DC/DC converters (3.3V, 5V, -12V) must drop below regulation limits within this time after the 12V input has dropped below 11.4V.		20	msec
The 5Vstby converter shall be in regulation limits within this time after the 12Vstby has reach 11.4V.		20	msec
The 5Vstby converter must power off within this time after the 12Vstby input has dropped below 11.4V.		20	msec

10.3.2.16 Residual Voltage Immunity in Standby Mode

Each DC/DC converter is immune to any residual voltage placed on its respective output (typically a leakage voltage through the system from standby output) up to 500mV. This residual voltage does not have any adverse effect on each DC/DC converter, such as: no additional power dissipation or over-stressing/over-heating any internal components or adversely affecting the turn-on performance (no protection circuits tripping during turn on).

While in Stand-by mode, at no load condition, the residual voltage on each DC/DC converter output does not exceed 100mV.

10.3.3 Protection Circuits

The PDB shall shut down all the DC/DC converters on the PDB and the power supply (by PSON) if there is a fault condition on the PDB (OVP or OCP). If the PDB DC/DC converter latches-off due to a protection circuit tripping, an AC cycle OFF for 15sec min or a PSON# cycle HIGH for 1sec shall be able to reset the power supply and the PDB.

10.3.3.1 Over Current Protection (OCP)/240VA Protection

Each DC/DC converter output on PDB has individual OCP protection circuits. The PS+PDB combo shall shutdown and latch off after an over current condition occurs. This latch shall be cleared by toggling the PSON# signal or by an AC power interruption. The values are measured at the PDB harness connectors. The DC/DC converters shall not be damaged from repeated power cycling in this condition. Also, the +12V output from the power supply is divided on the PDB into 4 channels and +12V4 is limited to 240VA of power. There are current sensors and limit circuits to shut down the entire PS+PDB combo if the limit is exceeded. The limits are listed in below table. -12V and 5VSB is protected under over current or shorted conditions so that no damage can occur to the power supply. Auto-recovery feature is a requirement on 5VSB rail.

Table 84. PDB Over Current Protection Limits/240VA Protection

Output Voltage	Min OCP Trip Limits	Max OCP Trip Limits	Usage	Connectors
+3.3V	27A	32A	PCIe, Misc	P1, P5, P6
+5V	27A	32A	PCIe, HDD, Misc	P1, P5, P6
+12V1	91A	100A	CPU1 + memory Fans, Misc	P1-P3, P12
+12V2	76A	100A	HDD and peripherals	P13-P16
+12V3	18A	20A		P5-P11

10.3.3.2 Over Voltage Protection (OVP)

Each DC/DC converter output on PDB have individual OVP protection circuits built in and it shall be locally sensed. The PS+PDB combo shall shutdown and latch off after an over voltage condition occurs. This latch shall be cleared by toggling the PSON# signal or by an AC power interruption. The table below displays the over voltage limits. The values are measured at the PDB harness connectors. The voltage shall never exceed the maximum levels when measured at the power pins of the output harness connector during any single point of fail. The voltage shall never trip any lower than the minimum levels when measured at the power pins of the PDB connector.

Table 85. Over Voltage Protection (OVP) Limits

Output Voltage	OVP Min (v)	OVP Max (v)
+3.3V	3.9	4.8
+5V	5.7	6.5

Output Voltage	OVP Min (v)	OVP Max (v)
-12V	-13.3	-15.5
+5VSB	5.7	6.5

10.3.4 PWOK (Power OK) Signal

The PDB connects the PWOK signals from the power supply modules and the DC/DC converters to a common PWOK signal. This common PWOK signal connects to the PWOK pin on P1. The DC/DC convert PWOK signals have open collector outputs.

10.3.4.1 System PWOK Requirements

The system will connect the PWOK signal to 3.3V or 5V by a pull-up resistor. The maximum sink current of the power supplies are 0.5mA. The minimum resistance of the pull-up resistor is stated below depending upon the motherboard's pull-up voltage. Refer to the *CRPS Power Supply Specification* for signal details.

Table 86. System PWOK Requirements

Motherboard Pull-up Voltage	MIN Resistance Value (ohms)
5V	10K
3.3V	6.8K

10.3.5 PSON Signal

The PDB connects the power supplies PSON signals together and connect them to the PSON signal on P1. Refer to the *CRPS Power Supply Specification* for signal details.

10.3.6 PMBus*

The PDB has no components on it to support PMBus*. It only needs to connect the power supply PMBus* signals (clock, data, SMBAlert#) and pass them to the 1x5 signal connector.

10.3.6.1 Addressing

The PDB addresses the power supply as follows on the PDB: 0 = open, 1 = grounded.

Table 87. PDB Addressing

	Power Supply Position 1	Power Supply Position 2
PDB addressing Address0/Address1	0/0	0/1
Power supply PMBus* device address	B0h	B2h

11. Design and Environmental Specifications

11.1 Intel® Server Board S2600CW Design Specifications

The following table defines the Intel® Server Board S2600CW operating and non-operating environmental limits. Operation of the Intel® Server Board S2600CW 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 system reliability.

Table 88. Server Board Design Specifications

Parameter	Limits
Operating Temperature	Board: 0°C to +55°C System: +10°C to +35°C
Non-Operating Temperature	-40°C to +70°C
Non-Operating Humidity	50% to 90%, non-condensing with a maximum wet bulb of 28°C (at temperatures from 25°C to 35°C)
Acoustic noise	Sound power: 7.0BA with hard disk drive stress only at room ambient temperature (23 +/-2C)
Shock, operating	Half sine, 2g peak, 11 mSec
Shock, unpackaged	System: 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

Note:

1. Chassis design must provide proper airflow to avoid exceeding the processor maximum case temperature.

***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.*

***Disclaimer Note**: Intel Corporation server boards contain a number of high-density VLSI and power delivery components that need adequate airflow to cool. Intel ensures through its own chassis development and testing that when Intel server building blocks are used together, the fully integrated system will meet the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel developed server building blocks to consult vendor datasheets and operating parameters to determine the*

amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible, if components fail or the server board does not operate correctly when used outside any of its published operating or non-operating limits.

11.2 MTBF

The following is the calculated Mean Time Between Failures (MTBF) 40°C (ambient air). These values are derived using a historical failure rate and multiplied by factors for application, electrical and/or thermal stress and for device maturity. You should view MTBF estimates as "reference numbers" only.

- Calculation Model: Telcordia* Issue 2, method I case 3
- Operating Temperature: Server in 40°C ambient air
- Operating Environment: Ground Benign, Controlled
- Duty Cycle: 100%
- Quality Level: II

Table 89. MTBF Estimate

Assembly	Failure Rate	MTBF
Motherboard	4261	234,708

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, 5-V standby is still present even though the server board is powered off.
- This server board supports the Intel® Xeon® Processor E5-2600 v3 product family with a Thermal Design Power (TDP) of up to and including 145 Watts. Previous generations of the Intel® Xeon® processors are not supported.
- 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.
- This server board only supports DDR4 DIMMs (RDIMM and LRDIMM). Mixing of RDIMMs and LRDIMMs is not supported.
- For the best performance, the number of DDR4 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.
- The Intel® Remote Management Module 4 (Intel® RMM4) connector is not compatible with any previous versions of the Intel® Remote Management Module (Product Order Code – AXXRMM, AXXRMM2, and AXXRMM3).
- Clear the CMOS with AC power cord plugged. Removing the AC power before performing the CMOS clear operation causes the system to automatically power up and immediately power down after the CMOS clear procedure is followed and AC power is re-applied. If this happens, remove the AC power cord, wait 30 seconds, and then re-connect the AC power cord. Power up the system and proceed to the <F2> BIOS Setup utility to reset the desired settings.
- 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: Compatible Intel® Server Chassis

The Intel® Server Board S2600CW can be used inside the Intel® Server Chassis P4000M family.

Table 90. Compatible Intel® Server Chassis

Chassis Name	System Fans	Storage Drives	Power Supply
P4304XXMFEN2	Two fixed Fans	Fixed HDD trays	One 550-W Non-redundant
P4304XXMUXX	Five redundant Fans	Fixed HDD trays	N/A. Compatible with 750-W or 1600-W Redundant PSUs.

Appendix C: 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 Codes

Sensor table given below lists the sensor identification numbers and information regarding the sensor type, name, supported thresholds, assertion and de-assertion information, and a brief description of the sensor purpose. Refer to 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, Version 2.0*. 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, Version 2.0*. 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, Version 2.0*, 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. Check platform EPS section for platform applicability and platform chassis section for chassis specific sensors. Redundancy sensors are present only on systems with appropriate hardware to support redundancy (for instance, fan or power supply).

Table 91. Integrated BMC Core Sensors

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Unit Status *(Pwr Unit Status)*	01h	All	Power Unit 09h	Sensor Specific 6Fh	00 - Power down	OK	As and De	–	Trig Offset	A	X
					02 - 240 VA power down	Fatal					
					04 - A/C lost	OK					
					05 - Soft power control failure	Fatal					
					06 - Power unit failure						
Power Unit Redundancy[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					

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
					07 – Redundant: Transition from non-redundant state.	Degraded					
IPMI Watchdog *(IPMI Watchdog)*	03h	All	Watchdog 2 23h	Sensor Specific 6Fh	00 - Timer expired, status only 01 - Hard reset 02 - Power down 03 - Power cycle 08 - Timer interrupt	OK	As	–	Trig Offset	A	X
Physical Security *(Physical Scrty)*	04h	Chassis Intrusion is chassis-specific	Physical Security 05h	Sensor Specific 6Fh	00 - Chassis intrusion 04 - LAN leash lost	Degraded OK	As and De	–	Trig Offset	A	X
FP Interrupt *(FP NMI Diag Int)*	05h	Chassis - specific	Critical Interrupt 13h	Sensor Specific 6Fh	00 - Front panel NMI/diagnostic interrupt	OK	As	–	Trig Offset	A	–
QPI Correctable Event (QPI Corr Sensor)	06h	All	Critical Event 13h	72h							
QPI Uncorrectable Event (QPI Fatl Sensor)	07h	All	Critical Event 13h	73h							
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

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
Button Sensor *(Button)*	09h	All	Button/Switch 14h	Sensor Specific 6Fh	00 – Power Button 02 – Reset Button	OK	AS	_	Trig Offset	A	X
BMC Watchdog	0Ah	All	Mgmt System Health 28h	Digital Discrete 03h	01 – State Asserted	Degraded	As	–	Trig Offset	A	-
Voltage Regulator Watchdog (VR Watchdog)	0Bh	All	Voltage 02h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Fan Redundancy[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					

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
SSB Thermal Trip *(SSB Therm Trip)*	0Dh	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
IO Module Presence (IO Mod Presence)	0Eh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	-
SAS Module Presence (SAS Mod Presence)	0Fh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	X
BMC Firmware Health (BMC FW Health)	10h	All	Mgmt Health 28h	Sensor Specific 6Fh	04 – Sensor Failure	Degraded	As	-	Trig Offset	A	X
System Airflow (System Airflow)	11h	All	Other Units 0Bh	Threshold 01h	–	–	–	Analog	–	–	–
FW Update Status	12h	All	Version Change 2Bh	OEM defined 70h	**00h – Update started** **01h – Update completed successfully.** 02h – Update failure	OK	As	_	Trig Offset	A	_
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

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
IO Module2 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
NM Health *(NM Health)*	19h	Platform-specific	OEM DCh	OEM defined 73h	-	-	-	-	-	-	–
NM Capabilities *(NM Capabilities)*	1Ah	Platform-specific	OEM DCh	OEM defined 74h	-	-	-	-	-	-	–
Baseboard Temperature 1 *(*Platform Specific*)*	20h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Front Panel Temperature *(Front Panel Temp)*	21h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SSB Temperature *(SSB Temp)*	22h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
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

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 4 *(*Platform Specific*)*	25h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Module Temperature (I/O Mod Temp)	26h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 1 Temperature (PCI Riser 1 Temp)	27h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Riser Temperature (IO Riser Temp)	28h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 1 Temperature (HSBP 1 Temp)	29h	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 2 Temperature (HSBP 2 Temp)	2Ah	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 3 Temperature (HSBP 3 Temp)	2Bh	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 2 Temperature (PCI Riser 2 Temp)	2Ch	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SAS Module Temperature (SAS Mod Temp)	2Dh	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Exit Air Temperature (Exit Air Temp)	2Eh	Chassis and Platform Specific	Temperature 01h	Threshold 01h	This sensor does not generate any events.	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Network Interface Controller Temperature (LAN NIC Temp)	2Fh	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Fan Tachometer Sensors (Chassis specific sensor names)	30h–3Fh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal[Note3]	As and De	Analog	R, T	M	-
Fan Present Sensors *(Fan x Present)*	40h–4Fh	Chassis and Platform Specific	Fan 04h	Generic 08h	01 - Device inserted	OK	As and De	-	Triggered Offset	Auto	-
Power Supply 1 Status[Note2] *(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[Note2] *(PS2 Status)*	51h	Chassis-specific	Power Supply 08h	Sensor Specific 6Fh	00 - Presence	OK	As and De	–	Trig Offset	A	X
					01 - Failure	Degraded					
					02 – Predictive Failure	Degraded					
					03 - A/C lost	Degraded					
					06 – Configuration error	OK					
Power Supply 1 AC Power Input *(PS1 Power In)*	54h	Chassis-specific	Other Units 0Bh	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Supply 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
					01- Drive Fault	Degraded					X
					07 - Rebuild/Remap in progress	Degraded					X
Processor 1 Status *(P1 Status)*	70h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 2 Status *(P2 Status)*	71h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 3 Status *(P3 Status)*	72h	Platform-specific	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	–	Trig Offset	M	X
					07 - Presence	OK					
Processor 4 Status	73h	Platform-	Processor	Sensor	01 - Thermal trip	Fatal	As	–	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
(P4 Status)		specific	07h	Specific 6Fh	07 - Presence	OK	and De				
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	–
Processor ERR2 Timeout *(CPU ERR2)*	7Ch	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	A	–
Catastrophic Error *(CATERR)*	80h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	–

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
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	-
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 Tachometer 1 *(PS1 Fan Tach 1)*	A0h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Supply 1 Fan Tachometer 2 *(PS1 Fan Tach 2)*	A1h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
MIC 1 Status *(GPGPU1 Status)*	A2h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
MIC 2 Status *(GPGPU2 Status)*	A3h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
Power Supply 2 Fan Tachometer 1 *(PS2 Fan Tach 1)*	A4h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 2 *(PS2 Fan Tach 2)*	A5h	Chassis-specific	Fan 04h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
MIC 3 Status *(GPGPU3 Status)*	A6h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
MIC 4 Status *(GPGPU4 Status)*	A7h	Platform Specific	Status C0h	OEM Defined 70h	-	-	-	-	-	-	-
Processor 1 DIMM Aggregate Thermal Margin 1 *(P1 DIMM Thrm Mrgn1)*	B0h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 1 DIMM Aggregate Thermal Margin 2 *(P1 DIMM Thrm Mrgn2)*	B1h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–

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 DIMM Aggregate Thermal Margin 1 *(P2 DIMM Thrm Mrgn1)*	B2h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 2 DIMM Aggregate Thermal Margin 2 *(P2 DIMM Thrm Mrgn2)*	B3h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 3 DIMM Aggregate Thermal Margin 1 *(P3 DIMM Thrm Mrgn1)*	B4h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 3 DIMM Aggregate Thermal Margin 2 *(P3 DIMM Thrm Mrgn2)*	B5h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 4 DIMM Aggregate Thermal Margin 1 *(P4 DIMM Thrm Mrgn1)*	B6h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 4 DIMM Aggregate Thermal Margin 2 *(P4 DIMM Thrm Mrgn2)*	B7h	Platform Specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Node Auto-Shutdown Sensor *(Auto Shutdown)*	B8h	Multi-Node Specific	Power Unit 09h	Generic – digital discrete 03h	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	A	-
Fan Tachometer Sensors (Chassis specific sensor names)	BAh–BFh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal[2]	As and De	Analog	R, T	M	-
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	-

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 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
MIC 1 Temp (GPGPU1 Core Temp)	C4h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 2 Temp (GPGPU2 Core Temp)	C5h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 3 Temp (GPGPU3 Core Temp)	C6h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
MIC 4 Temp (GPGPU4 Core Temp)	C7h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	-	-	-	-
Global Aggregate Temperature Margin 1 (Agg Therm Mrgn 1)	C8h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 2 (Agg Therm Mrgn 2)	C9h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 3 (Agg Therm Mrgn 3)	CAh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 4 (Agg Therm Mrgn 4)	CBh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Global Aggregate Temperature Margin 5 (Agg Therm Mrgn 5)	CCh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 6 (Agg Therm Mrgn 6)	CDh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 7 (Agg Therm Mrgn 7)	CEh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 8 (Agg Therm Mrgn 8)	CFh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Baseboard +12V *(BB +12.0V)*	D0h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Voltage Fault *(Voltage Fault)*	D1h	All	Voltage 02h	Discrete 03h	01 – Asserted	-	-	-	-	A	-
Baseboard CMOS Battery *(BB +3.3V Vbat)*	DEh	All	Voltage 02h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
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					

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
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

Notes:

1. Redundancy sensors are present only on systems with appropriate hardware to support redundancy (for instance, fan or power supply).
2. This is applicable only when the system does not support redundant fans. When fan redundancy is supported, then the contribution to system state is driven by the fan redundancy sensor.

Appendix D: Platform Specific BMC Appendix

This is an addendum document to BMC Core EPS. This document describes platform and chassis specific information.

Product ID

Bytes 11:12 (product ID) of *Get Device ID* command response: 71h 00h

IPMI Channel ID Assignments

Below table provides the information of BMC channels' assignments.

Channel ID	Interface	Supports Sessions
0	Primary IPMB	No
1	LAN1	Yes
2	LAN2	Yes
3	LAN3 [1] (Provided by the Intel® Remote Management Module 4)	Yes
4	Serial (COM2 terminal mode only)	Yes
5	USB	No
6	SMLink1 (IPMB connection to Node Manager. Bridged through BMC)	No
7	SMM	No
8-0Dh	Reserved	–
0Eh	Self [2]	–
0Fh	SMS / Receive Message Queue	No

Notes:

1. Optional HW supported by the server system.
2. Refers to the actual channel used to send the request.

Baseboard Specific Sensors

Sensor Number	Sensor Name
D5h	MEM VRM Temp
D6h	SAS IOC Temp

ACPI S3 Sleep State Support

Not supported.

Processor Support for Intel® Server Board S2600CW

- Intel® Xeon® E5-2600 v3 up to 145 Watt

Supported Chassis

- Intel® Server Chassis P4304XXMFEN2
- Intel® Server Chassis P4304XXMUXX

Chassis-specific Sensors

Fan Tachometer Sensors:

Intel® Server Chassis	Fan Tachometer Sensors (Sensor Number)	Fan Presence Sensors (Sensor Number)
P4304XXMFEN2	System Fan 1 (30h)	NA
	System Fan 2 (31h)	NA
P4304XXMUXX	System Fan 1 (30h)	Fan 1 Present (40h)
	System Fan 2 (31h)	Fan 2 Present (41h)
	System Fan 3 (32h)	Fan 3 Present (42h)
	System Fan 4 (33h)	Fan 4 Present (43h)
	System Fan 5 (34h)	Fan 5 Present (44h)

Hot-plug Fan Support

Supported on Intel® Server Chassis P4000 Redundant (Union Peak Medium) only

Fan Redundancy Support

Supported on Intel® Server Chassis P4000 Redundant (Union Peak Medium) only

Fan Domain Definition

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
P4304XXMUXX	0	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) LAN/BMC Temp (23h) HSBP 1 Temp (29h) SSB Temp (22h) LAN NIC Temp (2Fh) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h)	System Fan 1 (30h)

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
		P2 DTS Therm Mgn (84h) MIC 1 Margin (C4h) MIC 2 Margin (C5h) MIC 3 Margin (C6h)	
	1	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) LAN/BMC Temp (23h) HSBP 1 Temp (29h) SSB Temp (22h) LAN NIC Temp (2Fh) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h) MIC 1 Margin (C4h) MIC 2 Margin (C5h) MIC 3 Margin (C6h)	System Fan 2 (31h)
	2	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) SAS IOC Temp (D6h) MEM EFVRD Temp (24h) MEM VRM Temp (D5h) P1 VRD Temp (25h) HSBP 1 Temp (29h) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h) MIC 1 Margin (C4h) MIC 2 Margin (C5h) MIC 3 Margin (C6h)	System Fan 3 (32h)
	3	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) SAS IOC Temp (D6h) MEM EFVRD Temp (24h) MEM VRM Temp (D5h) P1 VRD Temp (25h) HSBP 1 Temp (29h)	System Fan 4 (33h)

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
		Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h) MIC 1 Margin (C4h) MIC 2 Margin (C5h) MIC 3 Margin (C6h)	
	4	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) SAS IOC Temp (D6h) MEM EFVRD Temp (24h) MEM VRM Temp (D5h) P1 VRD Temp (25h) HSBP 1 Temp (29h) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h) MIC 1 Margin (C4h) MIC 2 Margin (C5h) MIC 3 Margin (C6h)	System Fan 5 (34h)
	5	PS1 Temperature (5Ch) PS2 Temperature (5Dh)	Power supply fans
P4304XXMFEN2	0	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) LAN/BMC Temp (23h) HSBP 1 Temp (29h) SSB Temp (22h) LAN NIC Temp (2Fh) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h)	System Fan 1 (30h)
	1	DIMM Thrm Mrgn 1 (B0h) DIMM Thrm Mrgn 2 (B1h) DIMM Thrm Mrgn 3 (B2h) DIMM Thrm Mrgn 4 (B3h) SAS IOC Temp (D6h) MEM EFVRD Temp (24h) MEM VRM Temp (D5h) P1 VRD Temp (25h)	System Fan 2 (31h)

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
		HSBP 1 Temp (29h) Exit Air Temp (2Eh) P1 DTS Therm Mgn (83h) P2 DTS Therm Mgn (84h)	

HSC Availability

- Intel® Server Chassis P4304XXMUXX
  - 4-bay 3.5” HDD – FUP4X35S3HSBP
  - 8-bay 2.5” HDD Combo – FXX8X25PCIHSBP
  - 8-bay 2.5” HDD SAS Only – FXX8X25S3HSBP
- Intel® Server Chassis P4304XXMFEN2
  - 4-bay 3.5” HDD – FUP4X35S3HSBP
  - 8-bay 2.5” HDD Combo – FXX8X25PCIHSBP
  - 8-bay 2.5” HDD SAS Only – FXX8X25S3HSBP

Power Unit Redundancy Support

- Intel® Server Chassis P4304XXMUXX

Redundant Fans only for Intel® Server Chassis

- Intel® Server Chassis P4304XXMUXX

Fan Fault LED Support

Fan fault LEDs are available on the baseboard and on the hot-swap redundant fans available on the Intel® Server Chassis P4000 Redundant (Union Peak Medium).

Memory Throttling Support

The baseboard supports this feature.

Appendix E: POST Code Diagnostic LED Decoder

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 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, the Diagnostic LEDs can be used to identify the last POST process that was executed.

Each POST code is represented by a sequence of eight amber diagnostic LEDs. The POST codes are divided into two groups of LEDs as shown in the figure below.

The diagnostic LED #7 is labeled as “MSB”, and the diagnostic LED #0 is labeled as “LSB”.

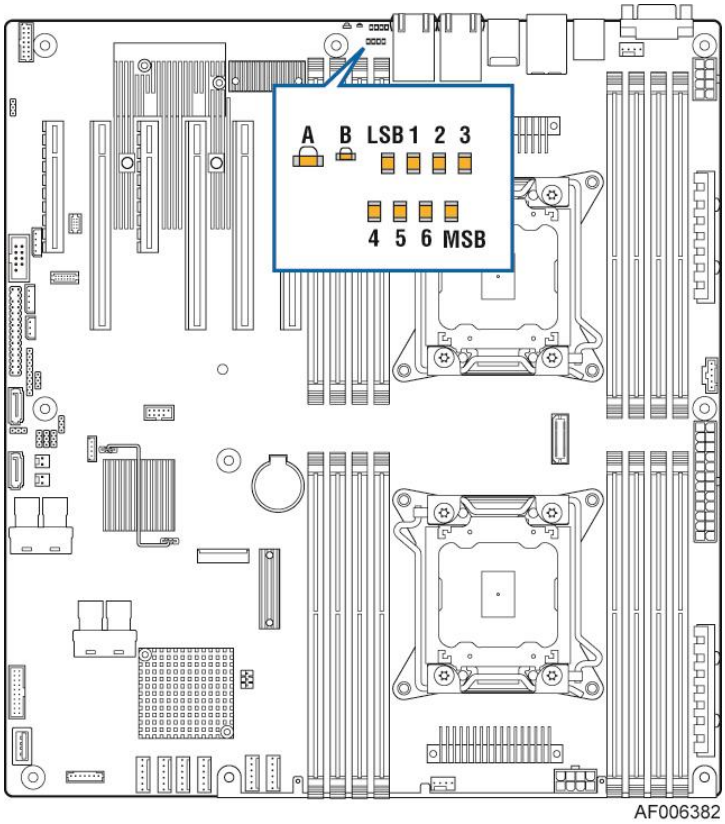

Figure 33. POST Code Diagnostic LED Decoder

A – System Status LED

B – System ID LED

LSB 1 2 3 4 5 6 MSB – Diagnostic LED

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

Table 92. POST Progress Code LED Example

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

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

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

Table 93. 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
QPI RC (Fully leverage without platform change)	
0xA1	Collect info such as SBSP, Boot Mode, Reset type, etc.
0xA3	Setup minimum path between SBSP and other sockets
0xA7	Topology discovery and route calculation
0xA8	Program final route
0xA9	Program final IO SAD setting
0xAA	Protocol layer and other uncore settings
0xAB	Transition links to full speed operation

0xAC	Phy layer settings
0xAD	Link Layer settings
0xAE	Coherency settings
0xAF	QPI initialization done
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
0x15	NB PEIM
0x19	SB PEIM
MRC Progress Codes *At this point the MRC Progress Code sequence ix executed.* *See Table 94.*	
0x31	Memory Installed
0x32	CPU PEIM (CPU Init)
0x33	CPU PEIM (Cache Init)
0x4F	DEX 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
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)
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)

POST Memory Initialization MRC Diagnostic Codes

There are two types of POST Diagnostic Codes displayed by the MRC during memory initialization: Progress Codes and Fatal Error Codes.

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

Table 94. MRC Progress Codes

Progress Code	Main Sequence	Subsequences / subfunctions
0xB0	Detect DIMM population	--n/a--
0x B1	Set DDR4 frequency	--n/a--
0x B2	Gather remaining SPD data	--n/a--
0x B3	Program registers on the memory controller level	--n/a--
0x B4	Evaluate RAS modes and save rank information	--n/a--
0x B5	Program registers on the channel level	--n/a--
0x B6	Perform the JEDEC defined initialization sequence	--n/a--
0x B7	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
0x B8	Initialize CLTT/OLTT	--n/a--
0x B9	Hardware memory test and init	--n/a--
0x BA	Execute software memory init	--n/a--
0x BB	Program memory map and interleaving	--n/a--
0x BC	Program RAS configuration	--n/a--
0x BF	MRC is done	--n/a--

Memory Initialization at the beginning of POST includes multiple functions, including: discovery, channel training, validation that the DIMM population is acceptable and functional,

initialization of the IMC and other hardware settings, and initialization of applicable RAS configurations.

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

Table 95. MRC Fatal Error Codes

Error Code	Fatal Error Code Explanation (With MRC Internal Minor Code)
0xE8	No Usable Memory Error • 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	Memory is locked by Intel® Trusted Execution Technology and is inaccessible.
0xEA	DDR4 Channel Training Error • 01h = Error on read DQ/DQS (Data/Data Strobe) Init • 02h = Error on Receive Enable • 03h = Error on Write Leveling • 04h = Error on write DQ/DQS (Data/Data Strobe)
0xEB	Memory Test Failure • 01h = Software memtest failure • 02h = Hardware on Receive Enable • 03h = Hardware Memtest failure in Lockstep Channel mode requiring a channel to be disabled. *This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.*
0xED	DIMM Configuration/Population Error • 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	Indicates a CLTT table structure error

Appendix F: POST Error Code

Most error conditions encountered during POST are reported using POST Error Codes. These codes represent specific failures, warnings, or informational messages that are identified with particular hardware units. These POST Error Codes may be displayed in the Error Manager display screen, and are always automatically logged to the System Event Log (SEL). Being logged to SEL means that the error information is available to System Management applications, including Remote and Out of Band (OOB) management. 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 column 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.
- **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.

Table 96. 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
8130	Processor 01 disabled	Major

Error Code	Error Message	Response
8131	Processor 02 disabled	Major
8132	Processor 03 disabled	Major
8133	Processor 04 disabled	Major
8160	Processor 01 unable to apply microcode update	Major
8161	Processor 02 unable to apply microcode update	Major
8162	Processor 03 unable to apply microcode update	Major
8163	Processor 04 unable to apply microcode update	Major
8170	Processor 01 failed Self Test (BIST)	Major
8171	Processor 02 failed Self Test (BIST)	Major
8172	Processor 03 failed Self Test (BIST)	Major
8173	Processor 04 failed Self Test (BIST)	Major
8180	Processor 01 microcode update not found	Minor
8181	Processor 02 microcode update not found	Minor
8182	Processor 03 microcode update not found	Minor
8183	Processor 04 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 Selftest	Major
83A1	Management Engine (ME) Failed to respond.	Major
84F2	Baseboard management controller failed to respond	Major
84F3	Baseboard management controller in update mode	Major
84F4	Sensor data record empty	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8501	DIMM Population Error	Major
8520	DIMM_A1 failed test/initialization	Major
8521	DIMM_A2 failed test/initialization	Major
8522	DIMM_A3 failed test/initialization	Major
8523	DIMM_B1 failed test/initialization	Major
8524	DIMM_B2 failed test/initialization	Major
8525	DIMM_B3 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8527	DIMM_C2 failed test/initialization	Major
8528	DIMM_C3 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major
852A	DIMM_D2 failed test/initialization	Major
852B	DIMM_D3 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852D	DIMM_E2 failed test/initialization	Major
852E	DIMM_E3 failed test/initialization	Major

Error Code	Error Message	Response
852F	DIMM_F1 failed test/initialization	Major
8530	DIMM_F2 failed test/initialization	Major
8531	DIMM_F3 failed test/initialization	Major
8532	DIMM_G1 failed test/initialization	Major
8533	DIMM_G2 failed test/initialization	Major
8534	DIMM_G3 failed test/initialization	Major
8535	DIMM_H1 failed test/initialization	Major
8536	DIMM_H2 failed test/initialization	Major
8537	DIMM_H3 failed test/initialization	Major
8538	DIMM_I1 failed test/initialization	Major
8539	DIMM_I2 failed test/initialization	Major
853A	DIMM_I3 failed test/initialization	Major
853B	DIMM_J1 failed test/initialization	Major
853C	DIMM_J2 failed test/initialization	Major
853D	DIMM_J3 failed test/initialization	Major
853E	DIMM_K1 failed test/initialization	Major
853F (Go to 85C0)	DIMM_K2 failed test/initialization	Major
8540	DIMM_A1 disabled	Major
8541	DIMM_A2 disabled	Major
8542	DIMM_A3 disabled	Major
8543	DIMM_B1 disabled	Major
8544	DIMM_B2 disabled	Major
8545	DIMM_B3 disabled	Major
8546	DIMM_C1 disabled	Major
8547	DIMM_C2 disabled	Major
8548	DIMM_C3 disabled	Major
8549	DIMM_D1 disabled	Major
854A	DIMM_D2 disabled	Major
854B	DIMM_D3 disabled	Major
854C	DIMM_E1 disabled	Major
854D	DIMM_E2 disabled	Major
854E	DIMM_E3 disabled	Major
854F	DIMM_F1 disabled	Major
8550	DIMM_F2 disabled	Major
8551	DIMM_F3 disabled	Major
8552	DIMM_G1 disabled	Major
8553	DIMM_G2 disabled	Major
8554	DIMM_G3 disabled	Major
8555	DIMM_H1 disabled	Major
8556	DIMM_H2 disabled	Major
8557	DIMM_H3 disabled	Major
8558	DIMM_I1 disabled	Major

Error Code	Error Message	Response
8559	DIMM_I2 disabled	Major
855A	DIMM_I3 disabled	Major
855B	DIMM_J1 disabled	Major
855C	DIMM_J2 disabled	Major
855D	DIMM_J3 disabled	Major
855E	DIMM_K1 disabled	Major
855F (Go to 85D0)	DIMM_K2 disabled	Major
8560	DIMM_A1 encountered a Serial Presence Detection (SPD) failure	Major
8561	DIMM_A2 encountered a Serial Presence Detection (SPD) failure	Major
8562	DIMM_A3 encountered a Serial Presence Detection (SPD) failure	Major
8563	DIMM_B1 encountered a Serial Presence Detection (SPD) failure	Major
8564	DIMM_B2 encountered a Serial Presence Detection (SPD) failure	Major
8565	DIMM_B3 encountered a Serial Presence Detection (SPD) failure	Major
8566	DIMM_C1 encountered a Serial Presence Detection (SPD) failure	Major
8567	DIMM_C2 encountered a Serial Presence Detection (SPD) failure	Major
8568	DIMM_C3 encountered a Serial Presence Detection (SPD) failure	Major
8569	DIMM_D1 encountered a Serial Presence Detection (SPD) failure	Major
856A	DIMM_D2 encountered a Serial Presence Detection (SPD) failure	Major
856B	DIMM_D3 encountered a Serial Presence Detection (SPD) failure	Major
856C	DIMM_E1 encountered a Serial Presence Detection (SPD) failure	Major
856D	DIMM_E2 encountered a Serial Presence Detection (SPD) failure	Major
856E	DIMM_E3 encountered a Serial Presence Detection (SPD) failure	Major
856F	DIMM_F1 encountered a Serial Presence Detection (SPD) failure	Major
8570	DIMM_F2 encountered a Serial Presence Detection (SPD) failure	Major
8571	DIMM_F3 encountered a Serial Presence Detection (SPD) failure	Major
8572	DIMM_G1 encountered a Serial Presence Detection (SPD) failure	Major
8573	DIMM_G2 encountered a Serial Presence Detection (SPD) failure	Major
8574	DIMM_G3 encountered a Serial Presence Detection (SPD) failure	Major
8575	DIMM_H1 encountered a Serial Presence Detection (SPD) failure	Major
8576	DIMM_H2 encountered a Serial Presence Detection (SPD) failure	Major
8577	DIMM_H3 encountered a Serial Presence Detection (SPD) failure	Major
8578	DIMM_I1 encountered a Serial Presence Detection (SPD) failure	Major
8579	DIMM_I2 encountered a Serial Presence Detection (SPD) failure	Major
857A	DIMM_I3 encountered a Serial Presence Detection (SPD) failure	Major
857B	DIMM_J1 encountered a Serial Presence Detection (SPD) failure	Major
857C	DIMM_J2 encountered a Serial Presence Detection (SPD) failure	Major
857D	DIMM_J3 encountered a Serial Presence Detection (SPD) failure	Major
857E	DIMM_K1 encountered a Serial Presence Detection (SPD) failure	Major
857F (Go to 85E0)	DIMM_K2 encountered a Serial Presence Detection (SPD) failure	Major
85C0	DIMM_K3 failed test/initialization	Major
85C1	DIMM_L1 failed test/initialization	Major

Error Code	Error Message	Response
85C2	DIMM_L2 failed test/initialization	Major
85C3	DIMM_L3 failed test/initialization	Major
85C4	DIMM_M1 failed test/initialization	Major
85C5	DIMM_M2 failed test/initialization	Major
85C6	DIMM_M3 failed test/initialization	Major
85C7	DIMM_N1 failed test/initialization	Major
85C8	DIMM_N2 failed test/initialization	Major
85C9	DIMM_N3 failed test/initialization	Major
85CA	DIMM_O1 failed test/initialization	Major
85CB	DIMM_O2 failed test/initialization	Major
85CC	DIMM_O3 failed test/initialization	Major
85CD	DIMM_P1 failed test/initialization	Major
85CE	DIMM_P2 failed test/initialization	Major
85CF	DIMM_P3 failed test/initialization	Major
85D0	DIMM_K3 disabled	Major
85D1	DIMM_L1 disabled	Major
85D2	DIMM_L2 disabled	Major
85D3	DIMM_L3 disabled	Major
85D4	DIMM_M1 disabled	Major
85D5	DIMM_M2 disabled	Major
85D6	DIMM_M3 disabled	Major
85D7	DIMM_N1 disabled	Major
85D8	DIMM_N2 disabled	Major
85D9	DIMM_N3 disabled	Major
85DA	DIMM_O1 disabled	Major
85DB	DIMM_O2 disabled	Major
85DC	DIMM_O3 disabled	Major
85DD	DIMM_P1 disabled	Major
85DE	DIMM_P2 disabled	Major
85DF	DIMM_P3 disabled	Major
85E0	DIMM_K3 encountered a Serial Presence Detection (SPD) failure	Major
85E1	DIMM_L1 encountered a Serial Presence Detection (SPD) failure	Major
85E2	DIMM_L2 encountered a Serial Presence Detection (SPD) failure	Major
85E3	DIMM_L3 encountered a Serial Presence Detection (SPD) failure	Major
85E4	DIMM_M1 encountered a Serial Presence Detection (SPD) failure	Major
85E5	DIMM_M2 encountered a Serial Presence Detection (SPD) failure	Major
85E6	DIMM_M3 encountered a Serial Presence Detection (SPD) failure	Major
85E7	DIMM_N1 encountered a Serial Presence Detection (SPD) failure	Major
85E8	DIMM_N2 encountered a Serial Presence Detection (SPD) failure	Major
85E9	DIMM_N3 encountered a Serial Presence Detection (SPD) failure	Major
85EA	DIMM_O1 encountered a Serial Presence Detection (SPD) failure	Major
85EB	DIMM_O2 encountered a Serial Presence Detection (SPD) failure	Major

Error Code	Error Message	Response
85EC	DIMM_O3 encountered a Serial Presence Detection (SPD) failure	Major
85ED	DIMM_P1 encountered a Serial Presence Detection (SPD) failure	Major
85EE	DIMM_P2 encountered a Serial Presence Detection (SPD) failure	Major
85EF	DIMM_P3 encountered a Serial Presence Detection (SPD) failure	Major
8604	POST Reclaim of non-critical NVRAM variables	Minor
8605	BIOS Settings are corrupted	Major
92A3	Serial port component was not detected	Major
92A9	Serial port component encountered a resource conflict error	Major
A000	TPM device not detected.	Minor
A001	TPM device missing or not responding.	Minor
A002	TPM device failure.	Minor
A003	TPM device failed self-test.	Minor
A100	BIOS ACM Error	Major
A421	PCI component encountered a SERR error	Fatal
A5A0	PCI Express* component encountered a PERR error	Minor
A5A1	PCI Express* component encountered an SERR error	Fatal

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

Table 97. POST Error Beep Codes

Beeps	Error Message	POST Progress Code	Description
3	Memory error	See Table 93.	System halted because a fatal error related to the memory was detected.
1 long	Intel® TXT security violation	0xAE, 0xAF	System halted because Intel® Trusted Execution Technology detected a potential violation of system security.

POST Error Beep Code

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

Table 98. Integrated BMC Beep Codes

Code	Reason for Beep	Associated Sensors
1-5-2-1	No CPUs installed or first CPU socket is empty.	CPU Missing Sensor

1-5-2-4	MSID Mismatch.	MSID Mismatch Sensor
1-5-4-2	Power fault: DC power is unexpectedly lost (power good dropout).	Power unit – power unit failure offset
1-5-4-4	Power control fault (power good assertion timeout).	Power unit – soft power control failure offset
1-5-1-2	VR Watchdog Timer sensor assertion.	VR Watchdog Timer
1-5-1-4	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.	PS Status

Appendix G: Statement of Volatility

This Appendix describes the volatile and non-volatile components on the Intel® Server Board S2600CW 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 S2600CW is included in the table below. The sections below the table provide additional information about the fields in this table.

Component Type	Size	Board Locaiton	User Data	Name
None-Volatile	16MB	U2F1	No	BMC FW flash ROM
None-Volatile	16MB	U4G2	No	BIOS flash ROM
None-Volatile	256K	U8L1	No	I350 EEROM
None-Volatile	16MB	U2A1	No	X540 flash ROM
Volatile	256MB	U1A1	No	BMC FW SDRAM
None-Volatile	16MB	U7V1	No	LSISAS3008 NOR flash
None-Volatile	64KB	U7W2	No	LSISAS3008 SBL EEPROM
None-Volatile	32KB	U3H2	No	LSISAS3008 nvSRAM (or mRAM)

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 appendix 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.

Term	Definition
ACPI	Advanced Configuration and Power Interface
AP	Application Processor
APIC	Advanced Programmable Interrupt Control
ASIC	Application Specific Integrated Circuit
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
CBC	Chassis Bridge Controller (A microcontroller connected to one or more other CBCs, together they bridge the IPMB buses of multiple chassis)
CEK	Common Enabling Kit
CHAP	Challenge Handshake Authentication Protocol
CMOS	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.
DPC	Direct Platform Control
EEPROM	Electrically Erasable Programmable Read-Only Memory
EHCI	Enhanced Host Controller Interface
EMP	Emergency Management Port
EPS	External Product Specification
FMB	Flexible MotherBoard
FMC	Flex Management Connector
FMM	Flex Management Module
FRB	Fault Resilient Booting
FRU	Field Replaceable Unit
FSB	Front Side Bus
GB	1024MB
GPIO	General Purpose I/O
GTL	Gunning Transceiver Logic
HSC	Hot-Swap Controller
Hz	Hertz (1 cycle/second)
I^2C	Inter-Integrated Circuit Bus
IA	Intel® Architecture

Term	Definition
IBF	Input Buffer
ICH	I/O Controller Hub
ICMB	Intelligent Chassis Management Bus
IERR	Internal Error
IFB	I/O and Firmware Bridge
INTR	Interrupt
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IR	Infrared
ITP	In-Target Probe
KB	1024 bytes
KCS	Keyboard Controller Style
LAN	Local Area Network
LCD	Liquid Crystal Display
LED	Light Emitting Diode
LPC	Low Pin Count
LUN	Logical Unit Number
MAC	Media Access Control
MB	1024KB
mBMC	National Semiconductor© PC87431x mini BMC
MCH	Memory Controller Hub
MD2	Message Digest 2 – Hashing Algorithm
MD5	Message Digest 5 – Hashing Algorithm – Higher Security
ms	milliseconds
MTTR	Memory Tpe Range Register
Mux	Multiplexor
NIC	Network Interface Controller
NMI	Nonmaskable Interrupt
NTB	Non-Transparent Bridge
OBF	Output Buffer
OEM	Original Equipment Manufacturer
Ohm	Unit of electrical resistance
PCH	Platform Controller Hub
PEF	Platform Event Filtering
PEP	Platform Event Paging
PIA	Platform Information Area (This feature configures the firmware for the platform hardware)
PLD	Programmable Logic Device
PMI	Platform Management Interrupt
POST	Power-On Self Test

Term	Definition
PSMI	Power Supply Management Interface
PWM	Pulse-Width Modulation
RAM	Random Access Memory
RASUM	Reliability, Availability, Serviceability, Usability, and Manageability
RISC	Reduced Instruction Set Computing
ROM	Read Only Memory
RTC	Real-Time Clock (Component of ICH peripheral chip on the server board)
SDR	Sensor Data Record
SECC	Single Edge Connector Cartridge
SEEPROM	Serial Electrically Erasable Programmable Read-Only Memory
SEL	System Event Log
SIO	Server Input/Output
SMI	Server Management Interrupt (SMI is the highest priority nonmaskable interrupt)
SMM	Server Management Mode
SMS	Server Management Software
SNMP	Simple Network Management Protocol
TBD	To Be Determined
TIM	Thermal Interface Material
UART	Universal Asynchronous Receiver/Transmitter
UDP	User Datagram Protocol
UHCI	Universal Host Controller Interface
UTC	Universal time coordinare
VID	Voltage Identification
VRD	Voltage Regulator Down
Word	16-bit quantity
ZIF	Zero Insertion Force

Reference Documents

- *Advanced Configuration and Power Interface Specification*, Revision 3.0, http://www.acpi.info/.
- *Intelligent Platform Management Bus Communications Protocol Specification*, Version 1.0. 1998. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- *Intelligent Platform Management Interface Specification*, Version 2.0. 2004. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- *Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification*, Version 1.1, 02/01/02, Intel Corporation.
- *Intel® Remote Management Module User's Guide*, Intel Corporation.
- *Alert Standard Format (ASF) Specification*, Version 2.0, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., http://www.dmtf.org.
- *Intel® Server System BIOS External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 V3 product family*
- *Intel® Server System BMC Firmware External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor E5-2600 V3 product family*
- *Intel® Remote Management Module 4 Technical Product Specification*
- *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*
- *Intel® Ethernet Controller I350 Family Product Brief*
- *Intel® Ethernet Controller X540 Family Product Brief*