# Intel® Server Board S2400GP

## *Technical Product Specification*

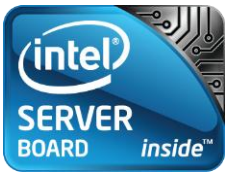


*Intel order number G50295-002*

**Revision 1.01**

**May, 2012**

**Enterprise Platforms and Services Division - Marketing**

# Revision History

| Date | Revision Number | Modifications |
|---|---|---|
| May 2012 | 1.0 | Initial release. |
| May 2012 | 1.01 | Add NIC Port MAC address. |

# Disclaimers

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel®'s Terms and Conditions of Sale for such products, Intel® assumes no liability whatsoever, and Intel® disclaims any express or implied warranty, relating to sale and/or use of Intel® products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel® products are not intended for use in medical, life-saving, or life sustaining applications. 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 Intel® Server Board S2400GP may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

Intel, Pentium, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation.

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

Copyright © Intel Corporation 2011 – 2012.

# Table of Contents

# List of Figures

# List of Tables

**<This page is intentionally left blank.>**

# 1. Introduction

This Technical Product Specification (TPS) provides board specific information detailing the features, functionality, and high-level architecture of the Intel® Server Board S2400GP.

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

## 1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Overview
- Chapter 3 – Functional Architecture
- Chapter 4 – System Security
- Chapter 5 – BIOS Setup Interface
- Chapter 6 – Platform Management Functional Overview
- Chapter 7 – Advanced Management Feature Support (RMM4)
- Chapter 8 – On-board Connector/Header Overview
- Chapter 9 – Jumper Blocks
- Chapter 10 – Intel® Light Guided Diagnostics
- Chapter 11 – Environmental Limits Specifications
- Appendix A – Integration and Usage Tips
- Appendix B – Integrated BMC Sensor Tables
- Appendix C – POST Code Diagnostic LED Decoder
- Appendix D – POST Code Errors
- Appendix E – Supported Intel® Server Chassis
- 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. Overview

The Intel® Server Board S2400GP is monolithic printed circuit boards (PCBs) with features designed to support the pedestal and rack server markets.

## 2.1 Intel® Server Boards S2400GP Feature Set

**Table 1. Intel® Server Boards S2400GP Feature Set**

| Feature | Description |
|---|---|
| Processors | ▪ Support for one or two Intel® Xeon® E5-2400 series Processor(s) in an FC-LGA 1356 Socket B2 package with Thermal Design Power up to 95W<br>▪ 6.4 GT/s, 7.2 GT/s and 8.0 GT/s Intel® QuickPath Interconnect (Intel® QPI)<br>▪ EVRD (Enterprise Voltage Regulator-Down) 12 |
| Memory | ▪ Six memory channels, Twelve memory DIMMs (two DIMMs per channel, three channels per processor socket)<br>▪ Support for 1066/1333 MT/s Unbuffered (UDIMM) LVDDR3 or DDR3 memory<br>▪ Support for 800/1066/1333/1600 MT/s ECC Registered (RDIMM) DDR3 memory<br>▪ Support for 800/1066/1333 ECC Registered (RDIMM) LVDDR3 memory<br>▪ Support for 1066/1333 Load Reduced (LRDIMM) DDR3 memory<br>▪ Support for 1066 Load Reduced (LRDIMM) LVDDR3 memory<br>▪ No support for mixing of RDIMMs and UDIMMs<br>▪ No support for Quad Rank DIMMs |
| Chipset | Intel® C602 (-A) chipset with support for storage option upgrade keys |
| Cooling Fan Support | Support for<br>▪ Two processor fans (4-pin headers)<br>▪ Six front system fans (6-pin headers)<br>▪ One rear system fan (4-pin headers)<br>▪ 3-pin fans are compatible with all fan headers |
| Add-in Card Slots | Six expansion slots:<br>▪ Slot6: PCI Express* Gen3 x16 slot, from the first processor<br>▪ Slot5: PCI Express* Gen3 x4 slot (supports x8 mechanically, x4 electrically), from the second processor<br>▪ Slot4: PCI Express* Gen3 x8 slot, from the second processor<br>▪ Slot3: PCI Express* Gen3 x16 slot, from the second processor<br>▪ Slot2: PCI Express* Gen3 x4 slot (x8 mechanically, x4 electrically) , from the first processor<br>▪ Slot1: 32-bit/33 MHz PCI slot, from PCH |
| Hard Drive and Optical Drive Support | ▪ Optical devices are supported.<br>▪ Two SATA connectors at 6Gbps and four SATA connectors at 3Gbps.<br>▪ Up to eight SAS2.0 connectors at 3Gbps with optional Intel® C600 RAID Upgrade Keys. |
| RAID Support | ▪ Intel® RSTe SW RAID 0/1/10/5<br>▪ LSI* SW RAID 0/1/10/5 |

| Feature | Description |
|---|---|
| I/O control support | External connections:<br>▪ DB9 serial port A connection<br>▪ Two RJ-45 NIC connectors for 10/100/1000 Mb connections: Dual GbE through the Intel® Ethernet Controller I350 (for S2400GP2 Sku)<br>▪ Four RJ-45 NIC connectors for 10/100/1000 Mb connections: Dual GbE through the Intel® Ethernet Controller I350 (for S2400GP4 sku)<br>▪ Four USB 2.0 ports at the back of the board<br>▪ One DB-15 video connector<br>Internal connections:<br>▪ Two 2x5pin USB headers, each supports two USB 2.0 ports<br>▪ One internal Type-A USB 2.0 port<br>▪ One 9pin USB header for eUSB SSD<br>▪ One DH10 serial port B header<br>▪ One SSI-compliant 32-pin front control panel header<br>▪ One 1x7 pin header for optional Intel® Local Control Panel support. |
| Video Support | ▪ Integrated 2D video controller<br>▪ Dual monitor video mode is supported<br>▪ 16 MB DDR3 Memory |
| LAN | Two Gigabit Ethernet Ports through the Intel® Ethernet Controller I350 PHYs (for S2400GP2 sku)<br>Four Gigabit Ethernet Ports through the Intel® Ethernet Controller I350 PHY (for S2400GP4 sku) |
| Security | Trusted Platform Module (Optional) |
| Server Management | ▪ Onboard ServerEngines* LLC Pilot III* Controller<br>▪ Support for Intel® Remote Management Module 4 solutions (Optional).<br>▪ Support for Intel® Remote Management Module 4 Lite solutions (Optional).<br>▪ Intel® Light-Guided Diagnostics on field replaceable units<br>▪ Support for Intel® System Management Software<br>▪ Support for Intel® Intelligent Power Node Manager (Need PMBus*-compliant power supply) |
| BIOS Flash | Winbond* 64MB Flash |
| Form Factor | SSI EEB 12”x13” compliant form factor |
| Compatible Intel® Server Chassis | Intel® Server Chassis P4000M Chassis |

## 2.2 Server Board Layout

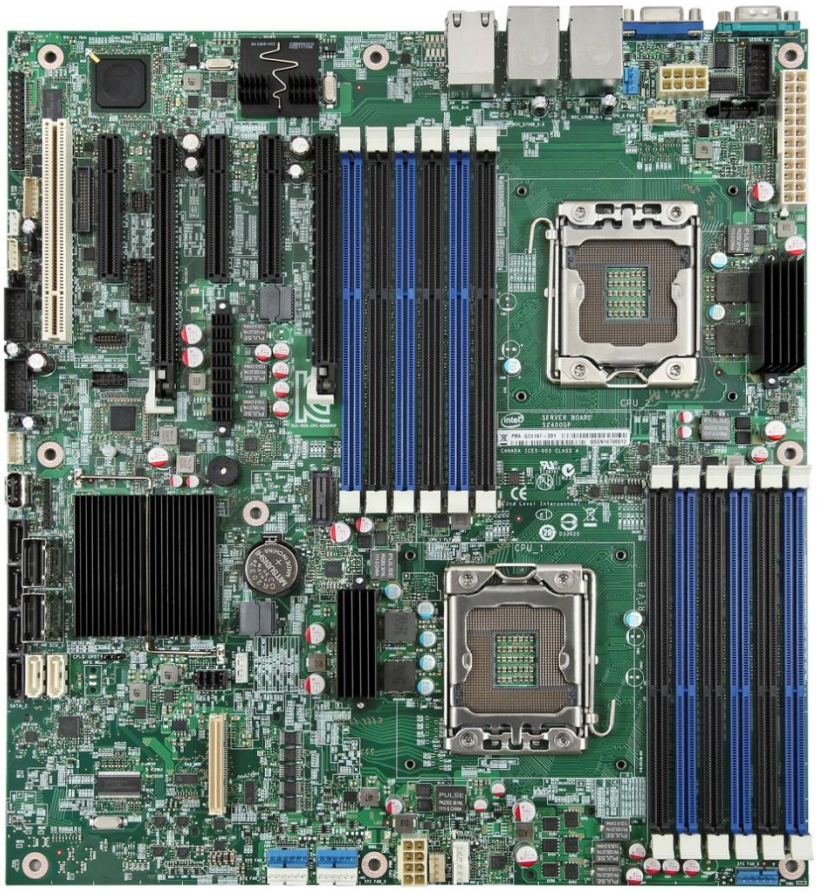

**Figure 1. Intel® Server Board S2400GP Layout (S2400GP4 as example)**

## 2.3 Server Board Connector and Component Layout

The following figure shows the layout of the server board. Each connector and major component is identified by a number or letter, and a description is given below the figure.

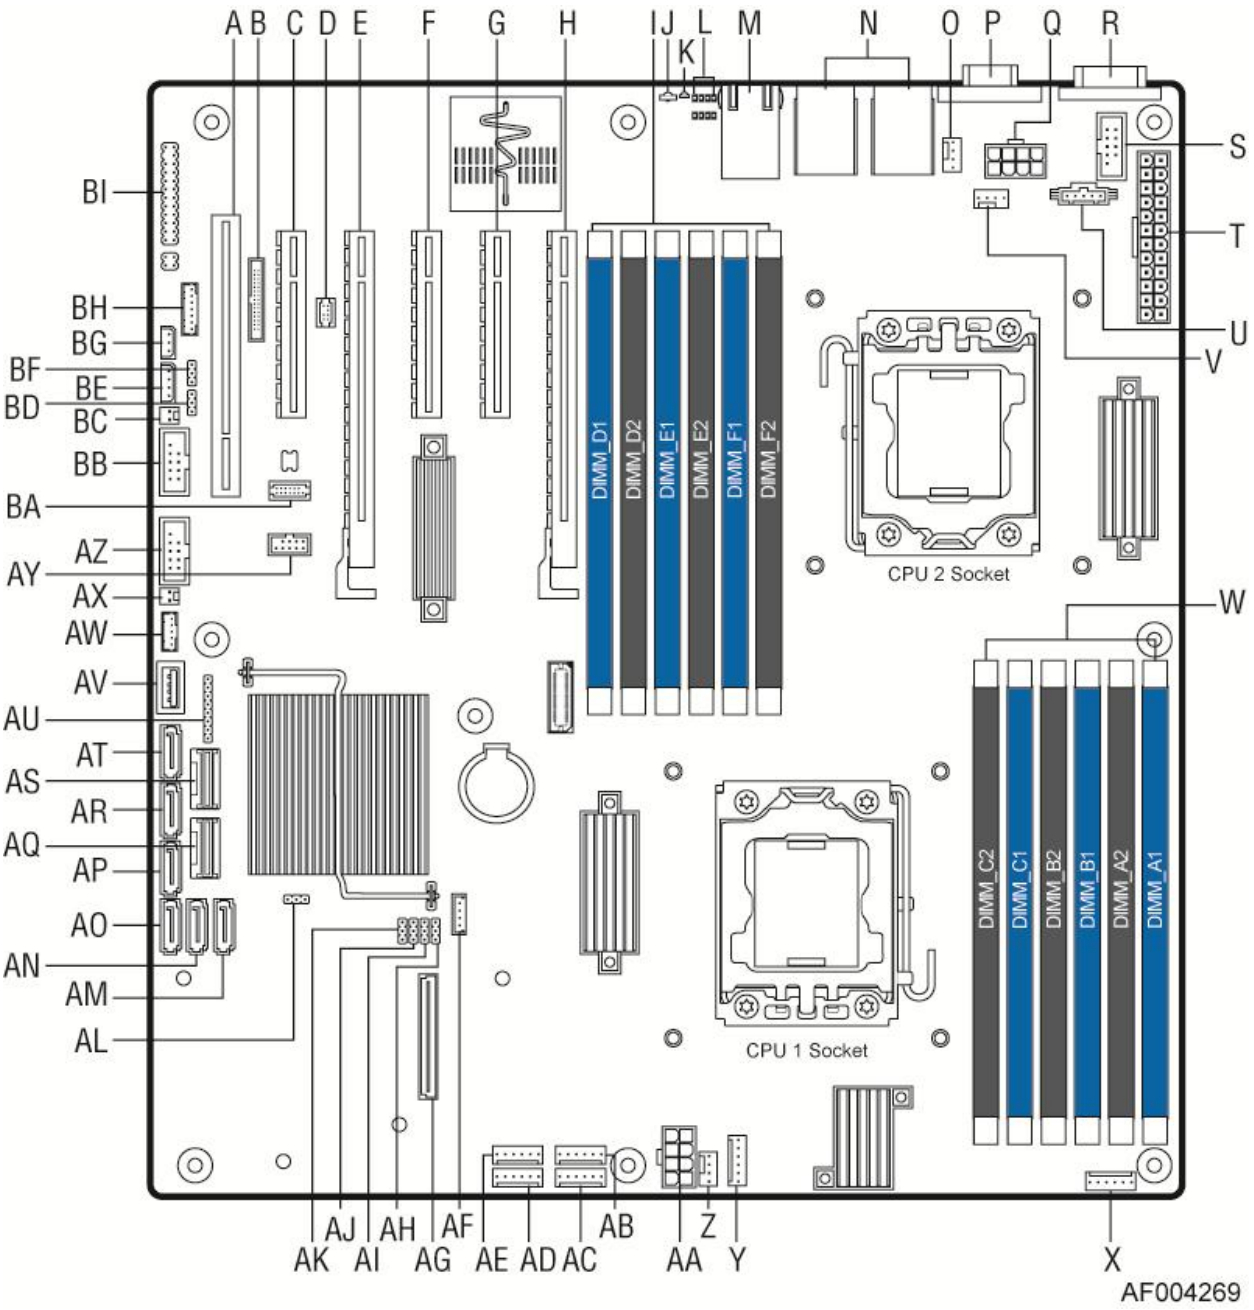


**Figure 2. Intel® Server Board S2400GP Components**

**Table 2. Intel® Server Boards S2400GP Components**

| Item | Description | Item | Description |
|---|---|---|---|
| A | Slot 1, 32bit/33MHz PCI | AA | Processor 1 power |
| B | RMM4 header | AB | System fan 4 header |
| C | Slot 2, PCI Express* Gen3 x4 (x8 mechanical connector) | AC | System fan 3 header |
| D | RMM4 Lite header | AD | System fan 1 header |
| E | Slot 3, PCI Express* Gen3 x16 (work only when CPU2 installed) | AE | System fan 2 header |
| F | Slot 4, PCI Express* Gen3 x8 (work only when CPU2 installed) | AF | Storage upgrade key |
| G | Slot 5, PCI Express* Gen3 x4 (work only when CPU2 installed and x8 mechanical connector) | AG | SAS/ROC Module connector |
| H | Slot 6, PCI Express* Gen3 x16 | AH | BIOS Recovery jumper |
| I | DIMM sockets from Processor 2 socket (Channel D, E, F) | AI | Password Clear jumper |
| J | System Status LED<br>Rear I/O connectors | AJ | BIOS Default jumper |
| K | ID (identify) LED | AK | ME Force Update jumper |
| L | Diagnostic LEDs | AL | CPLD Online Update jumper |
| M | NIC 3/4 (only on S2400GP4) | AM | SATA 6G port 0 |
| N | USB 0/1/2/3, NIC 1/2 | AN | SATA 6G port 1 |
| O | System fan 7 header | AO | SATA 3G port 2 |
| P | VGA header | AP | SATA 3G port 3 |
| Q | Processor 2 power | AQ | SCU1 Mini SAS port (4-7) |
| R | Serial Port A header | AR | SATA 3G port 4 |
| S | Serial Port B header | AS | SCU0 Mini SAS port (0-3) |
| T | Main power connector | AT | SATA 3G port 5 |
| U | Power supply auxiliary header | AU | CPLD JTAG |
| V | Processor 2 Fan header | AV | USB Type A port 8 |
| W | DIMM sockets from Processor 1 socket (Channel A, B, C) | AW | SATA SGPIO header |
| X | System fan 6 header | AX | HDD LED header |
| Y | System fan 5 header | AY | eUSB SSD |
| Z | Processor 1 Fan header | AZ | USB header port 6/7 to front panel |
| BA | TPM header | BB | USB header port 4/5 |
| BC | Chassis Intrusion header | BD | PECI header |
| BE | IPMB header | BF | BMC Force Update jumper |
| BG | HSBP_I2C header | BH | LCP header |
| BI | Front panel header | | |

## 2.4 Server Board Mechanical Drawings

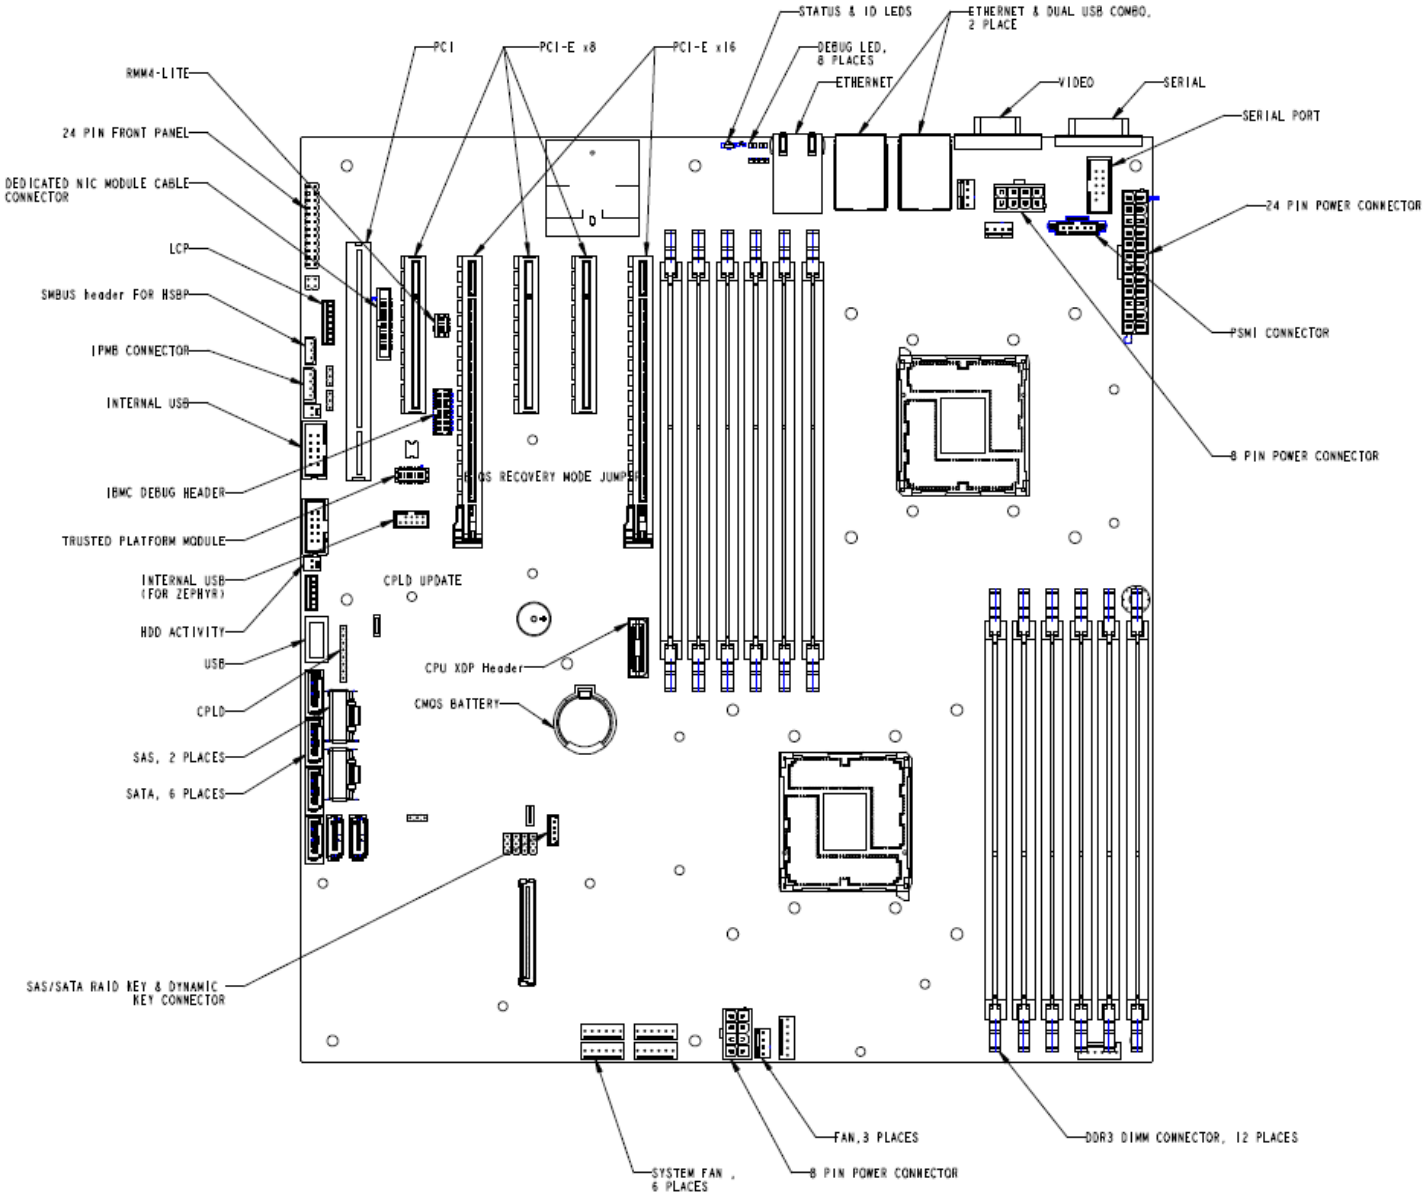


**Figure 3. Major Board Components**

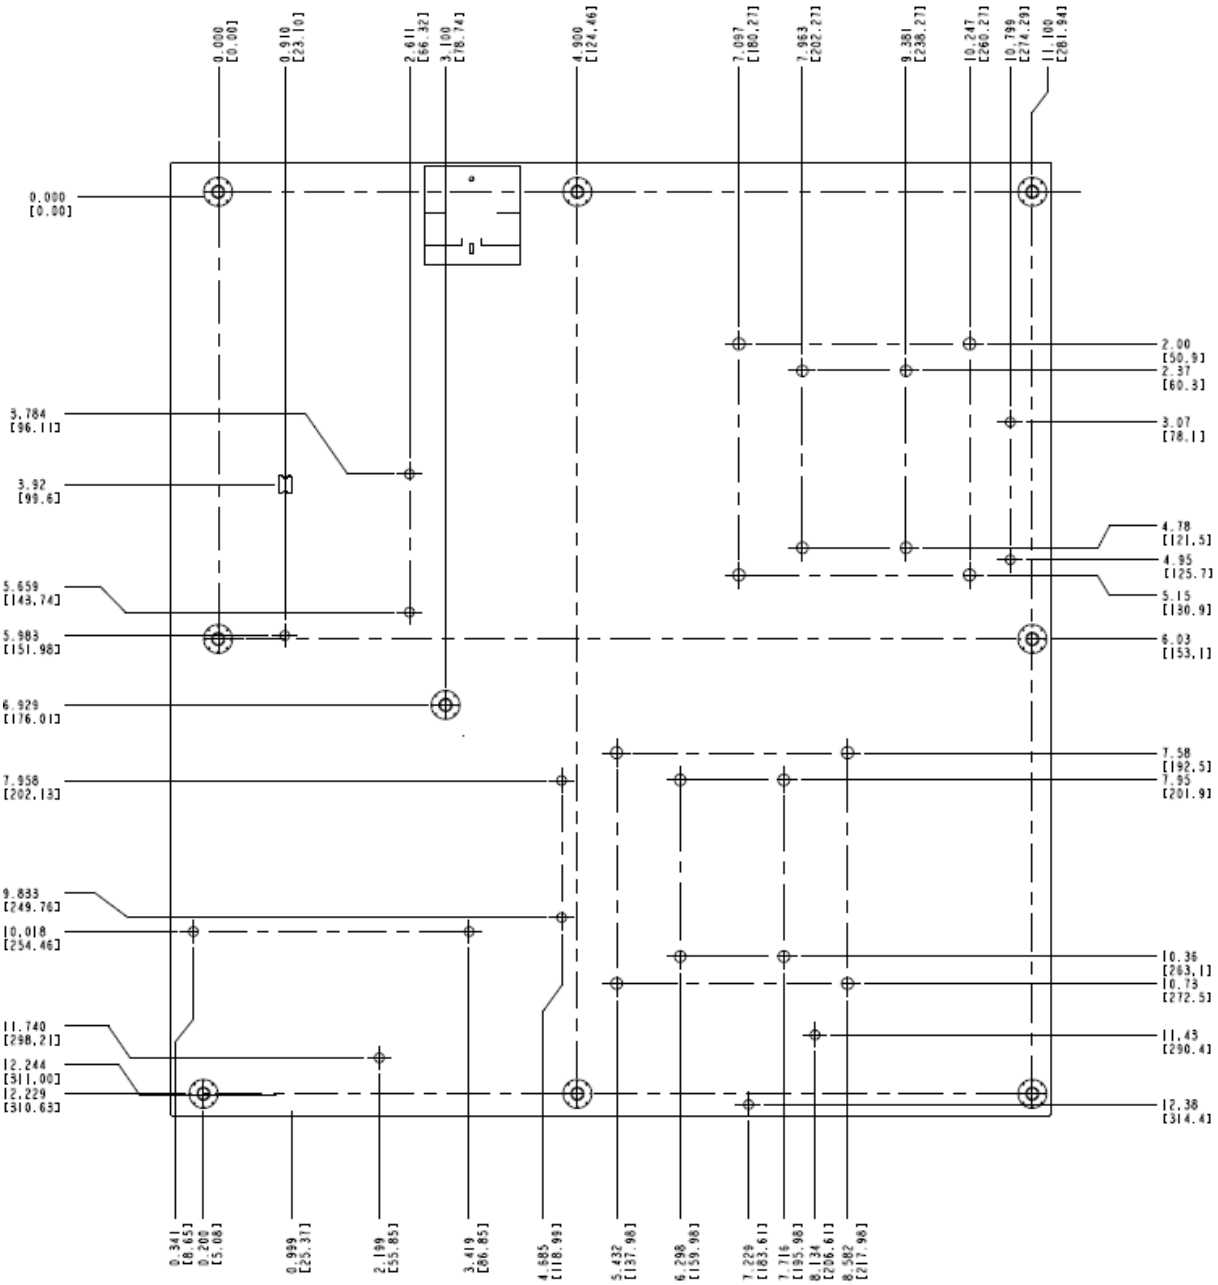


**Figure 4. Intel® Server Board S2400GP – Mounting Hole Locations (1 of 2)**

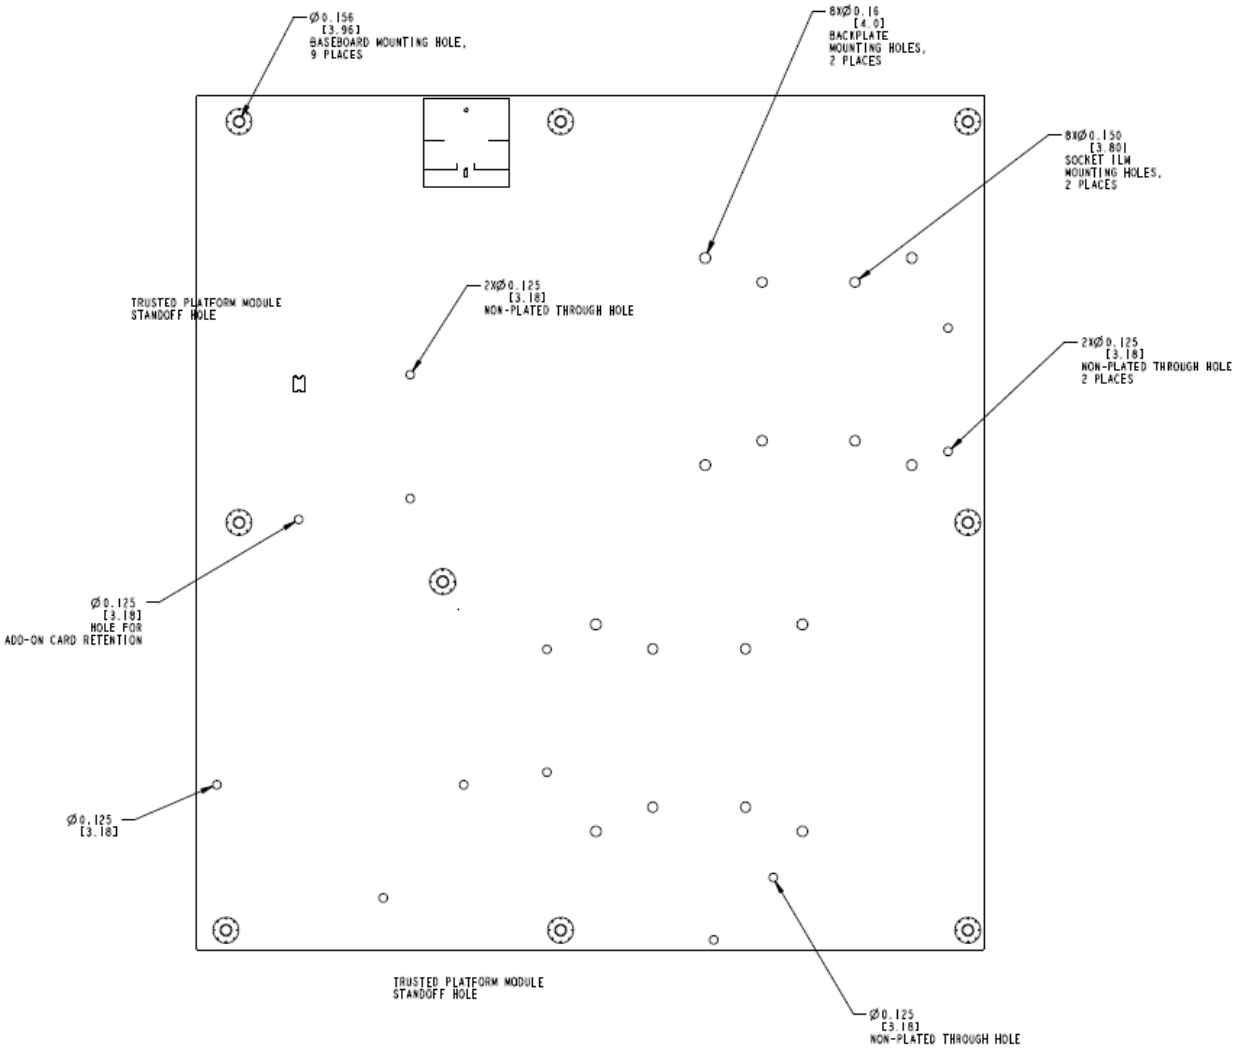


**Figure 5. Intel® Server Board S2400GP – Mounting Hole Locations (2 of 2)**

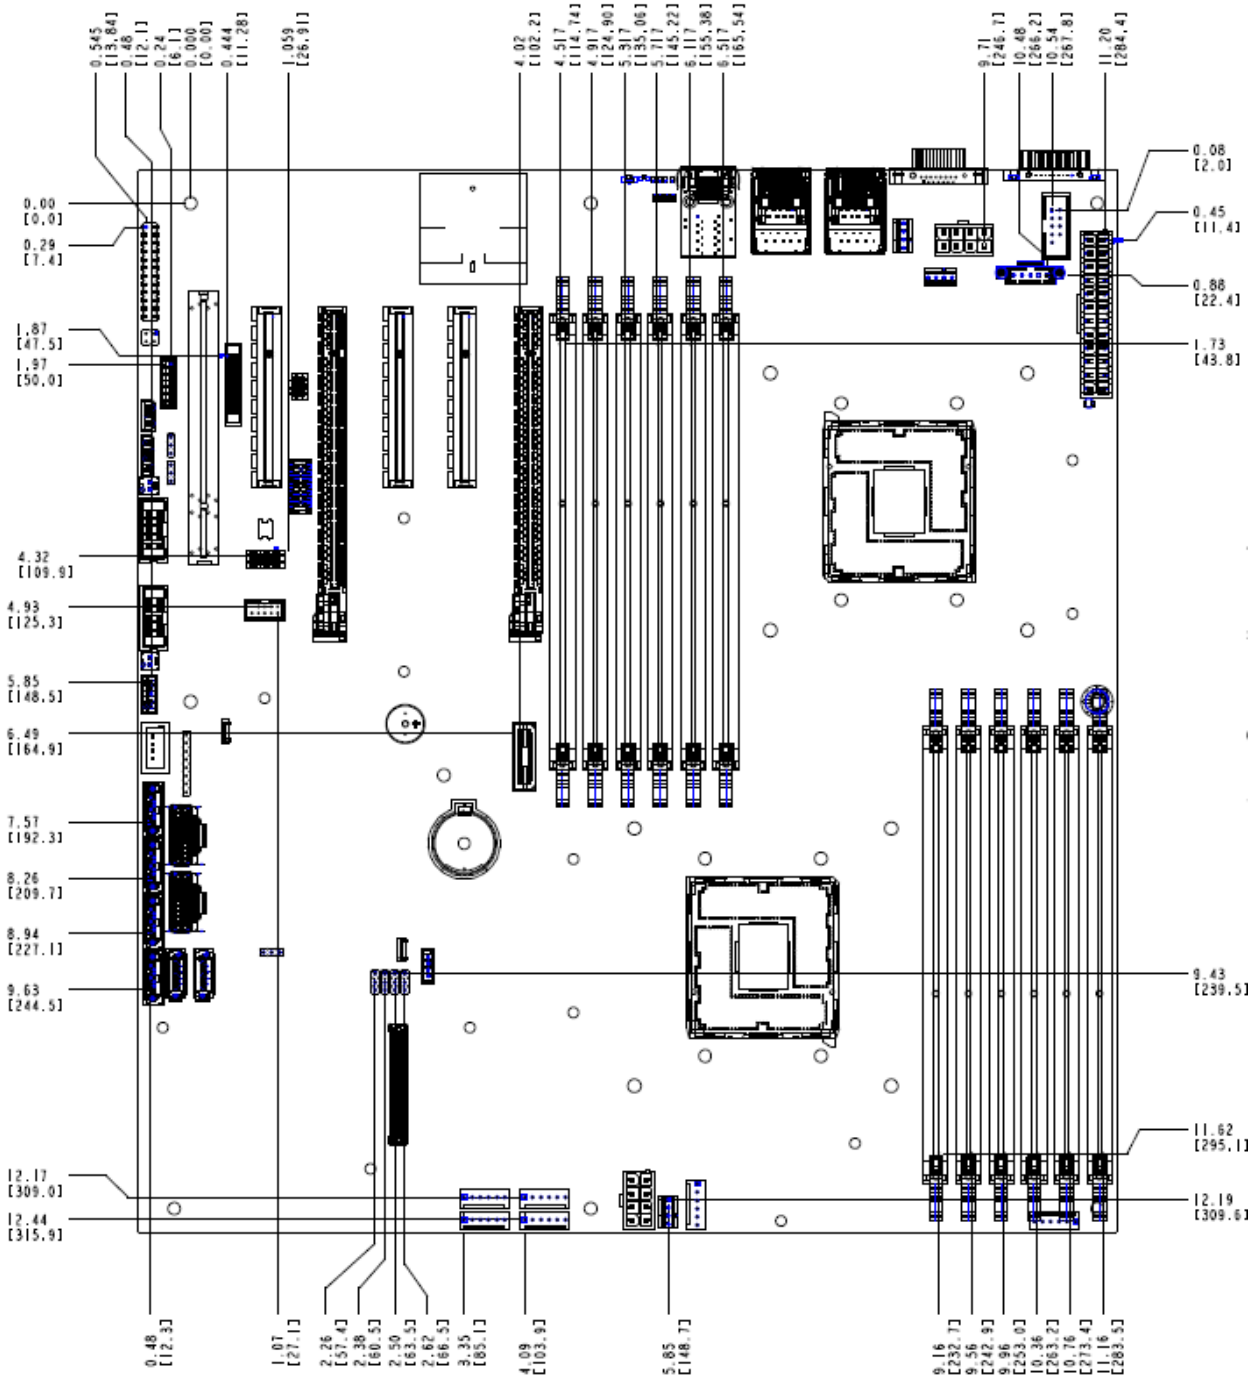


**Figure 6. Intel® Server Boards S2400GP – Major Connector Pin-1 Locations (1 of 3)**

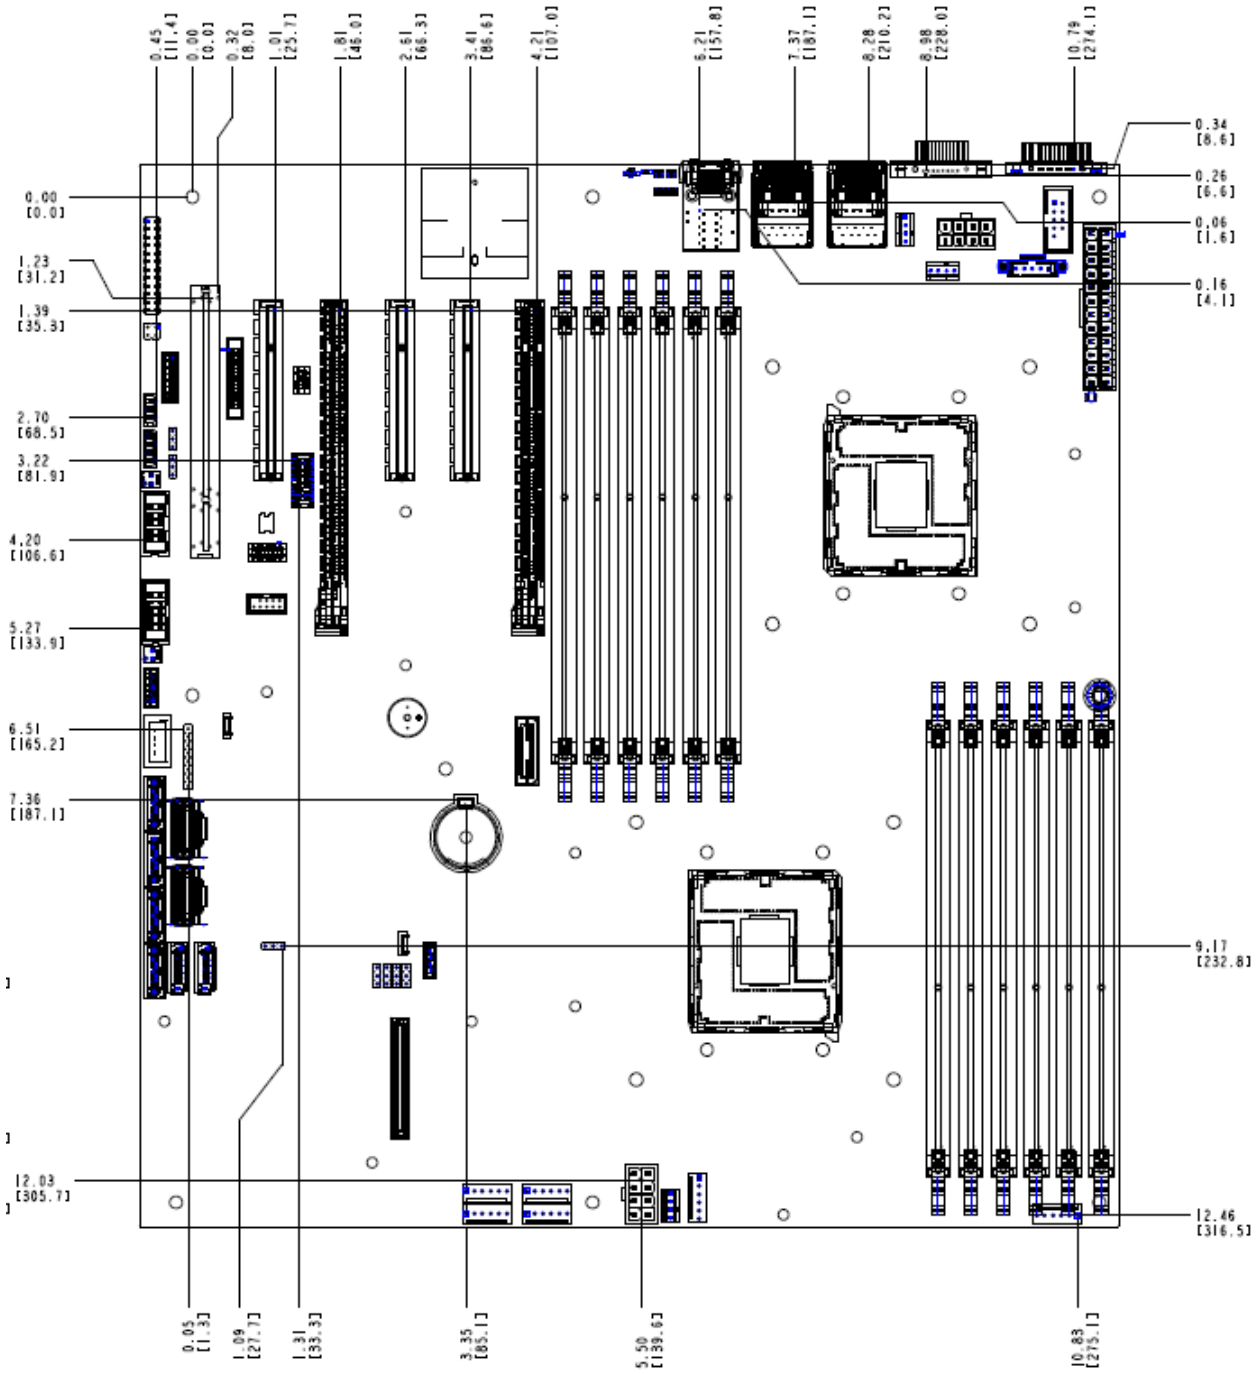


**Figure 7. Intel® Server Boards S2400GP – Major Connector Pin-1 Locations (2 of 3)**

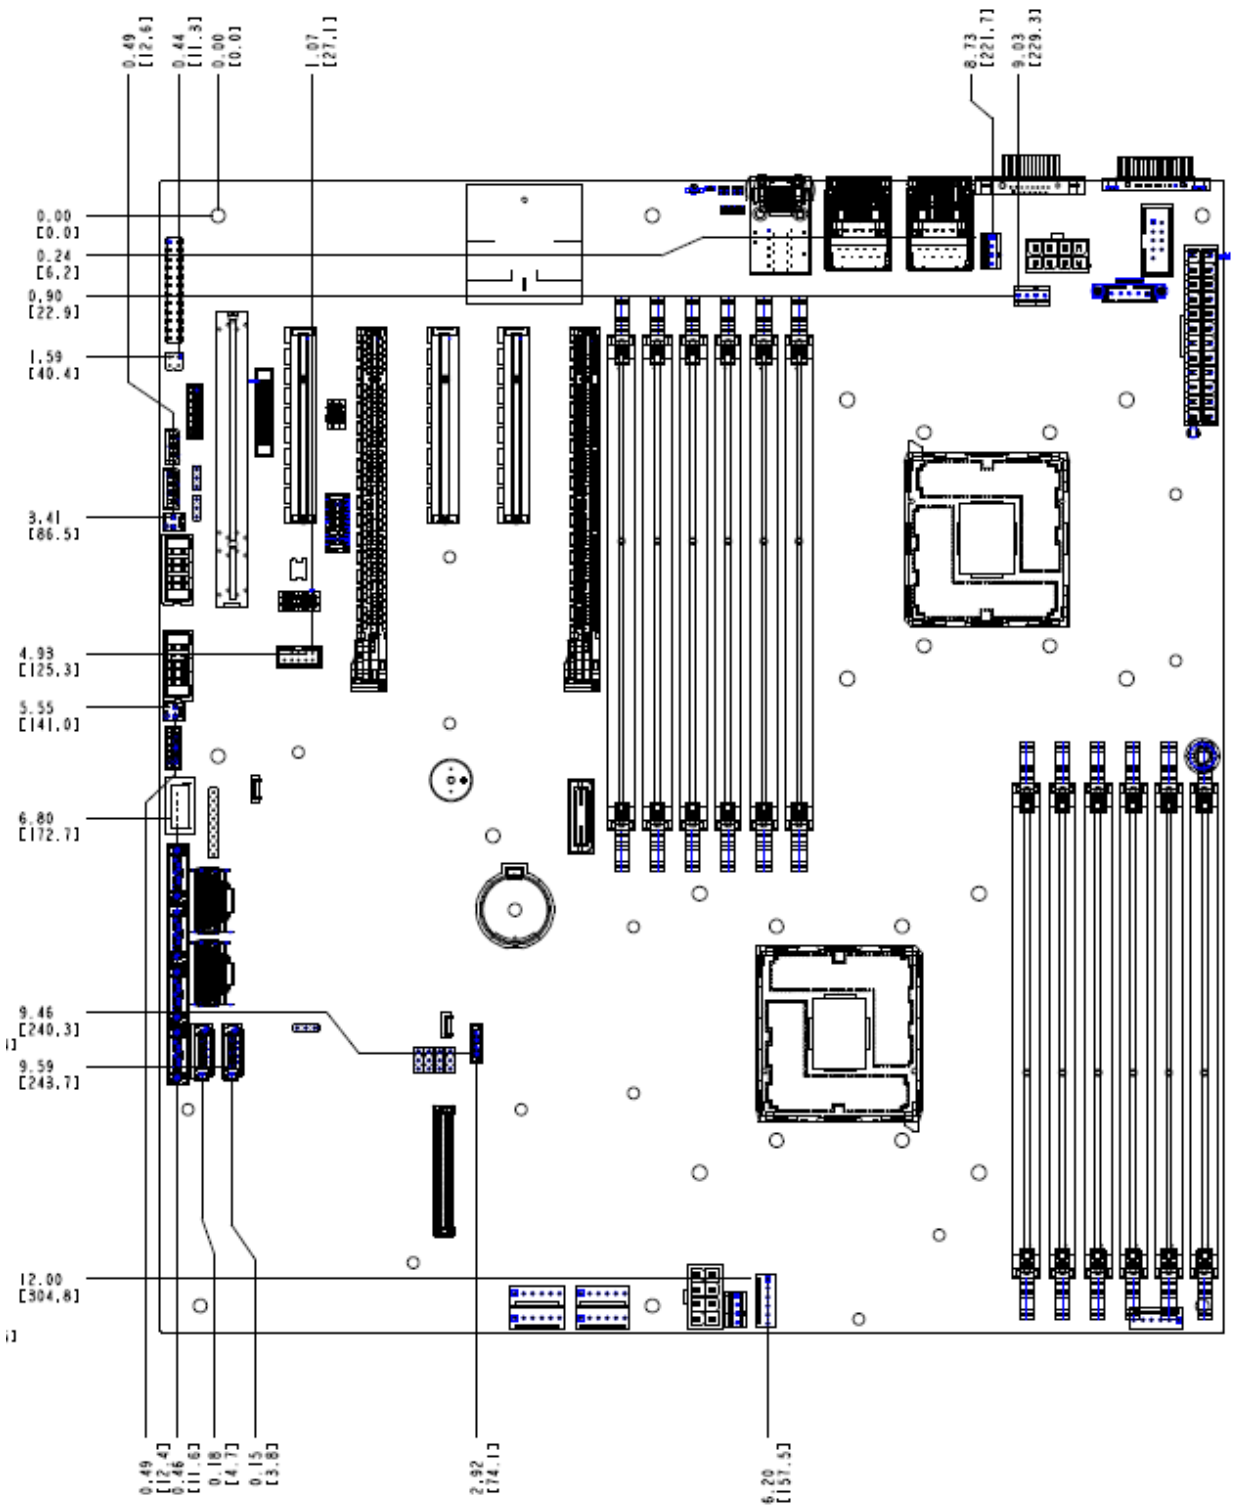


**Figure 8. Intel® Server Boards S2400GP – Major Connector Pin-1 Locations (2 of 3)**

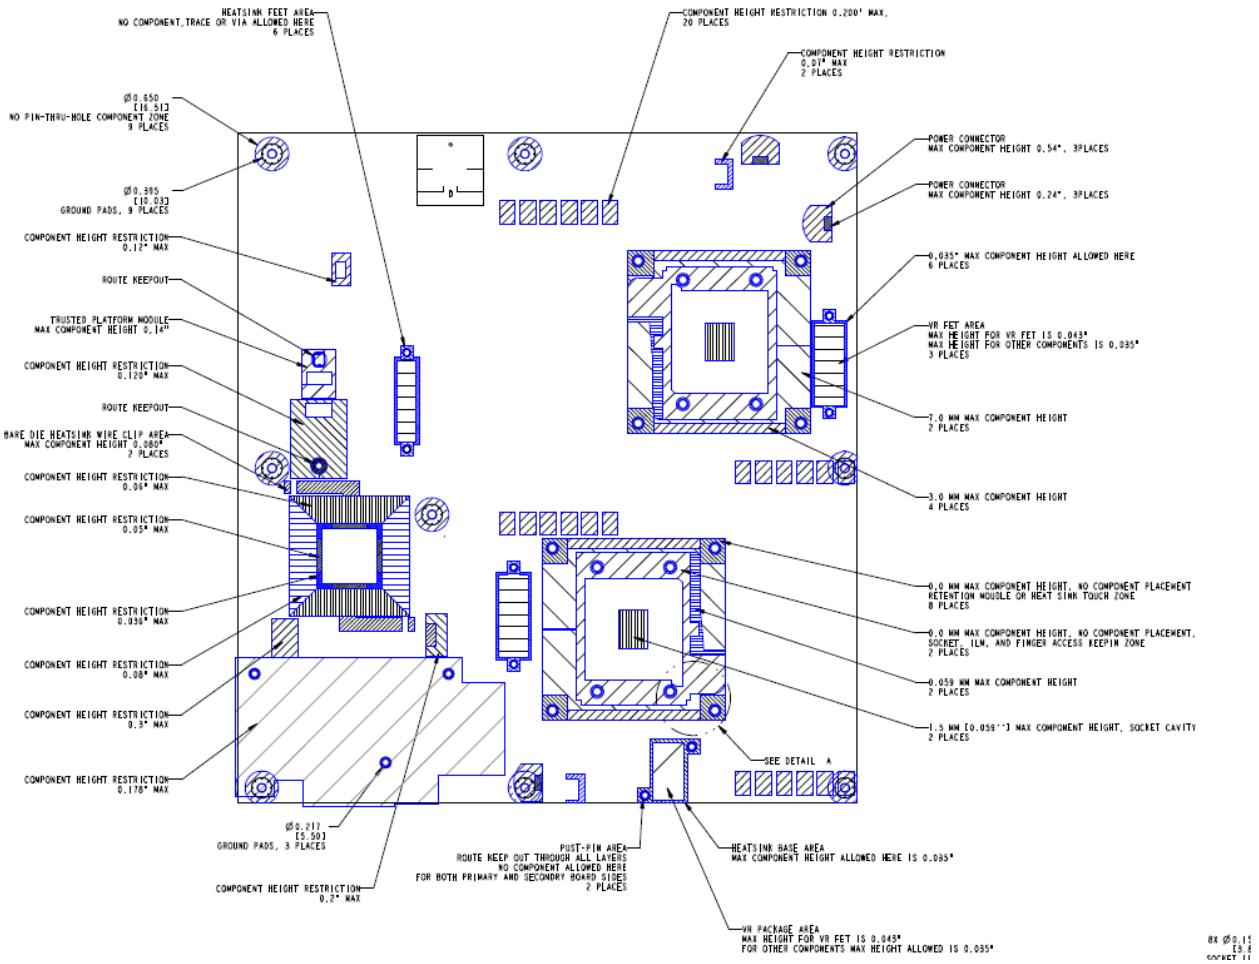


**Figure 9. Intel® Server Boards S2400GP – Primary Side Keep-out Zone**

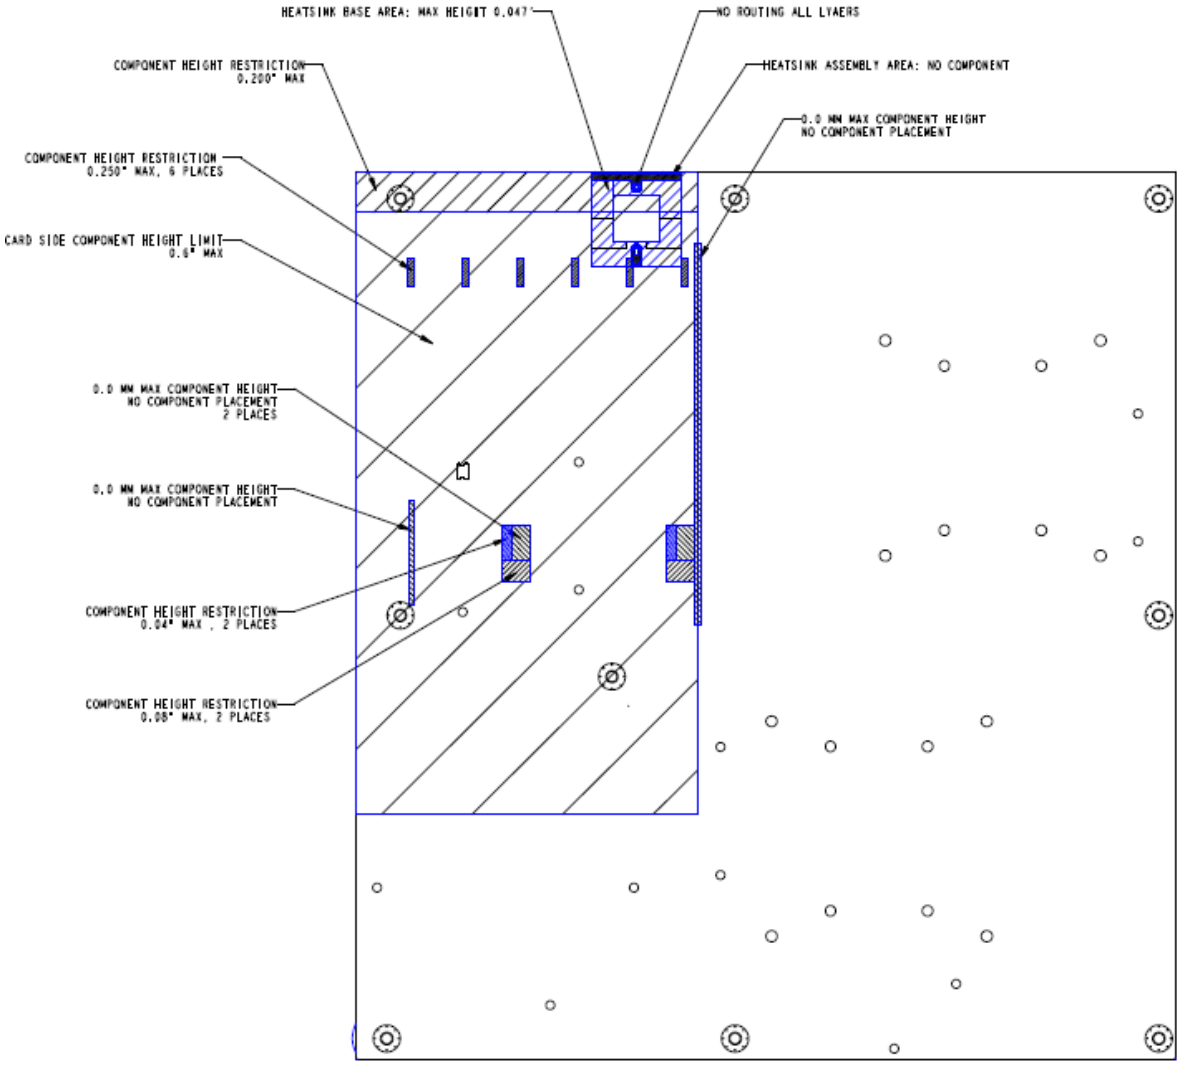


**Figure 10. Intel® Server Boards S2400GP – Primary Side Card Side Keep-out Zone**

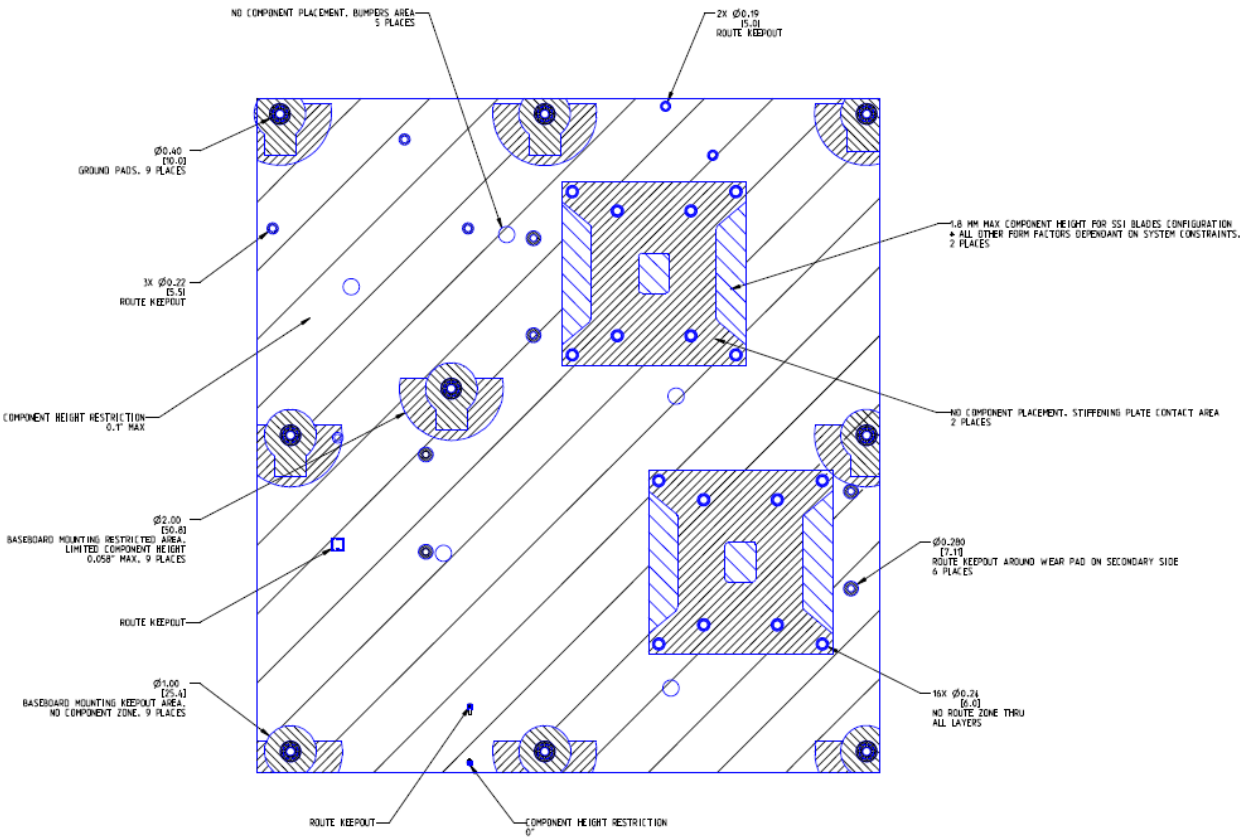


**Figure 11. Intel® Server Boards S2400GP – Second Side Keep-out Zone**

## 2.5 Server Board Rear I/O Layout

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

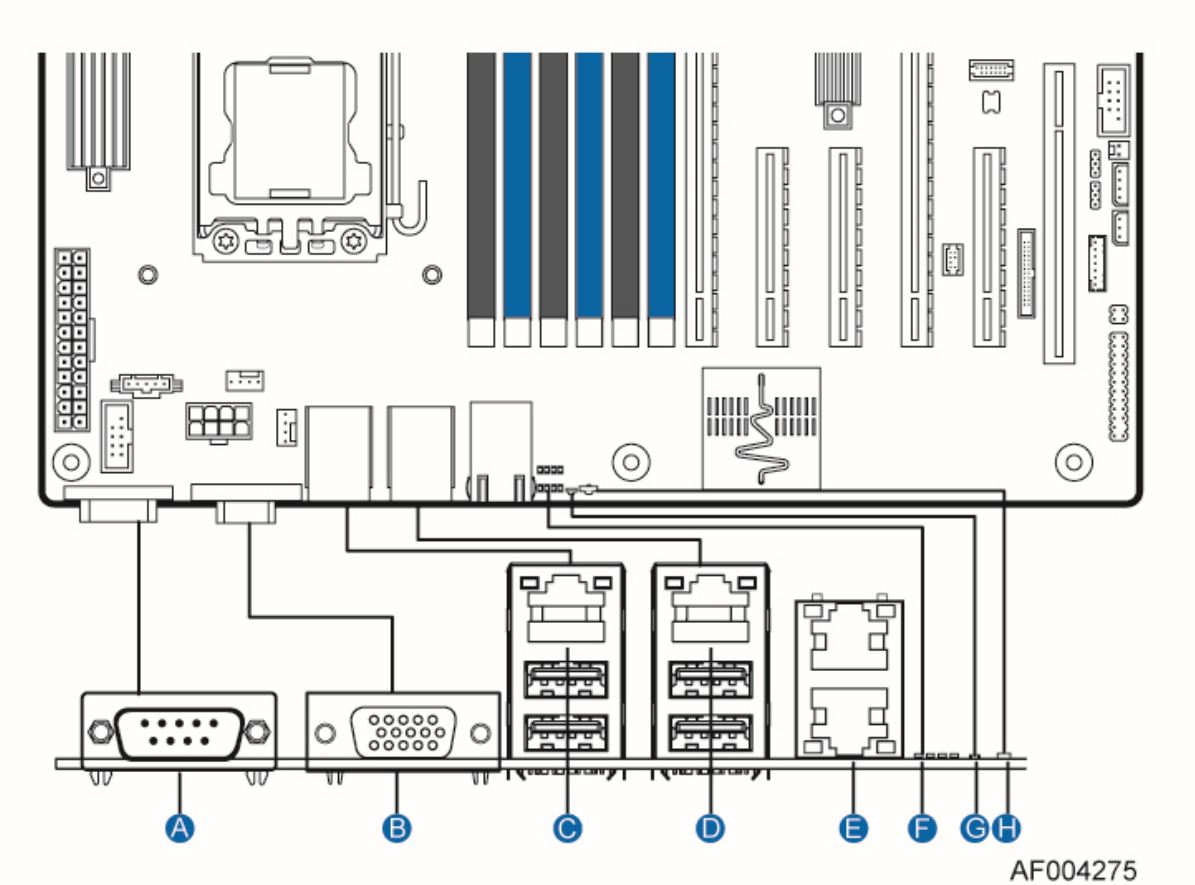


| | | | |
|---|---|---|---|
| A | Serial Port A | E | NIC Port 3 (top), 4 (bottom) for S2400GP4 sku only |
| B | Video | F | Diagnostic LEDs |
| C | NIC Port 1 (1 Gb) Management Port USB_0(top), 1(bottom) | G | ID LED |
| D | NIC Port 2 _USB Port 2 (top), 3 (bottom) | H | System Status LED |

**Figure 12. Intel® Server Boards S2400GP Rear I/O Layout**

# 3. Functional Architecture

The architecture and design of the Intel® Server Board S2400GP is based on the Intel® C600 chipset. The chipset is designed for systems based on the Intel® Xeon® processor in an FC-LGA 1356 Socket B2 package with Intel® QuickPath Interconnect (Intel® QPI).

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.

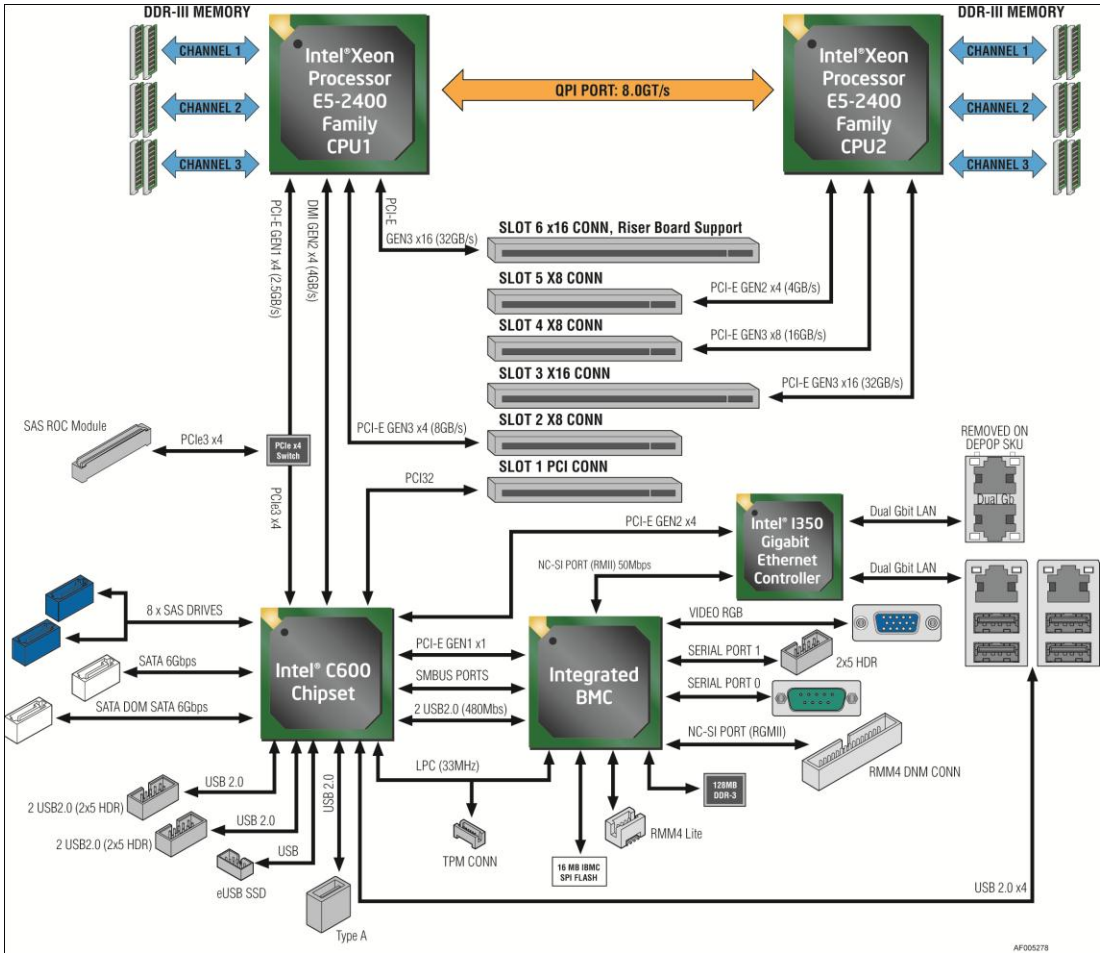


**Figure 13. Intel® Server Board S2400GP Functional Block Diagram**

## 3.1 Processor Support

The Intel® Server Board S2400GP includes two Socket-B2 (LGA-1356) processor sockets and can support one or two of the following processor: Intel® Xeon® processor E5-2400 product family, with a Thermal Design Power (TDP) of up to 95W.

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

Visit the Intel web site for a complete list of supported processors.

## 3.2 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.

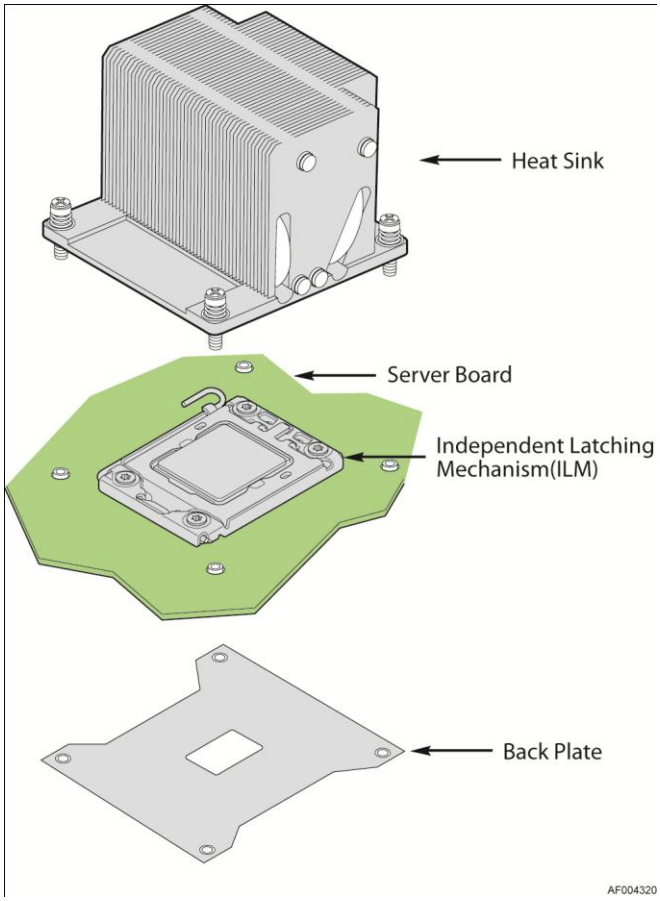


**Figure 14. Processor Socket Assembly**

## 3.3 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 this configuration. For optimal system performance in dual-processor configurations, Intel® recommends that identical processors be installed.

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

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 size for all levels of processor cache memory.
- Processors with different speeds can be mixed in a system, given the prior rules are met. If this condition is detected, all processor speeds are set to the lowest common denominator (highest common speed) and an error is reported.
- Processors which have different Intel® Quickpath (QPI) Link Frequencies may operate together if they are otherwise compatible and if a common link frequency can be selected. The common link frequency would be the highest link frequency that all installed processors can achieve.
- Processor stepping within a common processor family can be mixed as long as it is listed in the processor specification updates published by Intel Corporation.

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 E5-2400 product family and Intel® C600 chipset product family architecture. The errors fall into one of the following two 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.
- **Major:** If the "POST Error Pause" option in BIOS Setup is disabled, the system will log the error to the BIOS Setup Utility Error Manager and then continue to boot. No POST error message is given. If the "POST Error Pause" option in BIOS Setup is enabled, the error is logged and the system goes directly to the Error Manager in BIOS Setup.
- **Minor:** The message is displayed on the screen or on the Error manager screen, and the POST Error code is logged to 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 does not have any effect on this error.

**Table 3. Mixed Processor Configurations**

| Error | Severity | System Action |
|---|---|---|
| Processor family not Identical | Fatal | The BIOS detects the error condition and responds as follows:<br>Logs the POST Error Code into the System Event Log (SEL).<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Displays "0194: Processor family mismatch detected" message in the Error Manager.<br>Takes Fatal Error action (see above) and will not boot until the fault condition is remedied. |
| Processor model not Identical | Fatal | The BIOS detects the error condition and responds as follows:<br>Logs the POST Error Code into the System Event Log (SEL).<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Displays "0196: Processor model mismatch detected" message in the Error Manager.<br>Takes Fatal Error action (see above) and will not boot until the fault condition is remedied. |

| Error | Severity | System Action |
|---|---|---|
| Processor cores/threads not identical | Fatal | The BIOS detects the error condition and responds as follows:<br>Logs the POST Error Code into the SEL.<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Displays “0191: Processor core/thread count mismatch detected” message in the Error Manager.<br>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:<br>Logs the POST Error Code into the SEL.<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Displays “0192: Processor cache size mismatch detected message in the Error Manager.<br>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:<br>Adjusts all processor frequencies to the highest common frequency.<br>No error is generated – this is not an error condition.<br>Continues to boot the system successfully.<br>If the frequencies for all processors cannot be adjusted to be the same, then this is an error, and the BIOS responds as follows:<br>Logs the POST Error Code into the SEL.<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Does not disable the processor.<br>Displays “0197: Processor speeds unable to synchronize” message in the Error Manager.<br>Takes Fatal Error action (see above) and will not boot until the fault condition is remedied. |
| Processor Intel® QuickPath Interconnect link frequencies not identical | Fatal | The BIOS detects the QPI link frequencies and responds as follows:<br>Adjusts all QPI interconnect link frequencies to highest common frequency.<br>No error is generated – this is not an error condition.<br>Continues to boot the system successfully.<br>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:<br>Logs the POST Error Code into the SEL.<br>Alerts the BMC to set the System Status LED to steady Amber.<br>Displays “0195: Processor Intel® QPI link frequencies unable to synchronize” message in the Error Manager.<br>Does not disable the processor.<br>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:<br>Logs the POST Error Code into the SEL.<br>Displays “818x: Processor 0x microcode update not found” message in the Error Manager or on the screen.<br>The system continues to boot in a degraded state, regardless of the setting of POST Error Pause in the Setup. |

| Error | Severity | System Action |
|---|---|---|
| Processor microcode update failed | Major | The BIOS detects the error condition and responds as follows:<br>Logs the POST Error Code into the SEL.<br>Displays “816x: Processor 0x unable to apply microcode update” message in the Error Manager or on the screen.<br>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. |

## 3.4 Processor Function Overview

With the release of the Intel® Xeon® processor E5-2400 product family, several key system components, including the CPU, Integrated Memory Controller (IMC), and Integrated IO Module (IIO), have been combined into a single processor package and feature per socket; One Intel® QuickPath Interconnect point-to-point links capable of up to 8.0 GT/s, up to 24 lanes of Gen 3 PCI Express* links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* Gen 1 interface with a peak transfer rate of 2.5 GT/s. The processor supports up to 46 bits of physical address space and 48-bit of virtual address space.

The following sections will provide an overview of the key processor features and functions that help to define the performance and architecture of the server board. For more comprehensive processor specific information, refer to the Intel® Xeon® processor E5-2400 product family documents listed in the Reference Document list.

Processor Feature Details:

- Up to eight execution cores
- Each core supports two threads (Intel® Hyper-Threading Technology), up to 16 threads per socket
- 46-bit physical addressing and 48-bit virtual addressing
- 1GB 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 20 MB last level cache (LLC): up to 2.5 MB per core instruction/data last level cache (LLC), shared among all cores

Supported Technologies:

- Intel® Virtualization Technology (Intel® VT)
- Intel® Virtualization Technology for Directed I/O (Intel® VT-d)
- Intel® Virtualization Technology “Sandy Bridge” Processor Extensions
- Intel® Trusted Execution Technology (Intel® TXT)
- Intel® 64 Architecture
- Intel® Streaming SIMD Extensions 4.1 (Intel® SSE4.1)
- Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2)
- Intel® Advanced Vector Extensions (Intel® AVX)
- Intel® Hyper-Threading Technology
- Execute Disable Bit
- Intel® Turbo Boost Technology
- Intel® Intelligent Power Technology
- Enhanced Intel® SpeedStep Technology
- Intel® Data Direct I/O (DDIO) Technology

## 3.5 Intel® QuickPath Interconnect

The Intel® QuickPath Interconnect is a high speed, packetized, point-to-point interconnect used in the processor. The narrow high-speed links stitch together processors in distributed shared memory and integrated I/O platform architecture. It offers much higher bandwidth with low latency. The Intel® QuickPath Interconnect has an efficient architecture allowing more interconnect performance to be achieved in real systems. It has a snoop protocol optimized for low latency and high scalability, as well as packet and lane structures enabling quick completions of transactions. Reliability, availability, and serviceability features (RAS) are built into the architecture.

The physical connectivity of each interconnect link is made up of twenty differential signal pairs plus a differential forwarded clock. Each port supports a link pair consisting of two uni-directional links to complete the connection between two components. This supports traffic in both directions simultaneously. To facilitate flexibility and longevity, the interconnect is defined as having five layers: Physical, Link, Routing, Transport, and Protocol.

The Intel® QuickPath Interconnect includes a cache coherency protocol to keep the distributed memory and caching structures coherent during system operation. It supports both low-latency source snooping and a scalable home snoop behavior. The coherency protocol provides for direct cache-to-cache transfers for optimal latency.

## 3.6 Integrated Memory Controller (IMC) and Memory Subsystem

Integrated into the processor is a memory controller. Each processor provides three DDR3 channels that support the following:

- Unbuffered DDR3 and registered DDR3 DIMMs
- LR DIMM (Load Reduced DIMM) for buffered memory solutions demanding higher capacity memory subsystems
- Independent channel mode or lockstep mode
- Data burst length of eight cycles for all memory organization modes
- Memory DDR3 data transfer rates of 800, 1066, 1333, and 1600 MT/s
- 64-bit wide channels plus 8-bits of ECC support for each channel
- DDR3 standard I/O Voltage of 1.5 V and DDR3 Low Voltage of 1.35 V
- 1Gb, 2Gb, and 4Gb DDR3 DRAM technologies supported for these devices:
    - UDIMM DDR3 – SR x8 and x16 data widths, DR – x8 data width
    - RDIMM DDR3 – SR,DR, and QR – x4 and x8 data widths
    - LRDIMM DDR3 – QR – x4 and x8 data widths with direct map or with rank multiplication
- Up to 8 ranks supported per memory channel, 1, 2 or 4 ranks per DIMM
- Open with adaptive idle page close timer or closed page policy
- Per channel memory test and initialization engine can initialize DRAM to all logical zeros with valid ECC (with or without data scrambler) or a predefined test pattern
- Isochronous access support for Quality of Service (QoS)
- Minimum memory configuration: independent channel support with one DIMM populated
- Integrated dual SMBus* master controllers

- Command launch modes of 1n/2n
- RAS Support:
    - Rank Level Sparing and Device Tagging
    - Demand and Patrol Scrubbing
    - DRAM Single Device Data Correction (SDDC) for any single x4 or x8 DRAM device. Independent channel mode supports x4 SDDC. x8 SDDC requires lockstep mode
    - Lockstep mode where channels 0 and 1 and channels 2 and 3 are operated in lockstep mode
    - Data scrambling with address to ease detection of write errors to an incorrect address.
    - Error reporting through Machine Check Architecture
    - Read Retry during CRC error handling checks by iMC
    - Channel mirroring within a socket
        - CPU1 Channel Mirror Pairs B and C
        - CPU2 Channel Mirror Pairs E and F
    - Error Containment Recovery
- Improved Thermal Throttling with dynamic Closed Loop Thermal Throttling (CLTT)
- Memory thermal monitoring support for DIMM temperature

### 3.6.1 Supported Memory

**Table 4. UDIMM Support Guidelines**

| Ranks Per DIMM and Data Width | Memory Capacity Per DIMM1 | | | Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)[2,3] | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | | | 1 Slot per Channel | | 2 Slots per Channel | | | |
| | | | | 1DPC | | 1DPC | | 2DPC | |
| | | | | 1.35V | 1.5V | 1.35V | 1.5V | 1.35V | 1.5V |
| SRx8 Non-ECC | 1GB | 2GB | 4GB | n/a | 1066, 1333, | n/a | 1066, 1333 | n/a | 1066 |
| DRx8 Non-ECC | 2GB | 4GB | 8GB | n/a | 1066, 1333, | n/a | 1066, 1333 | n/a | 1066 |
| SRx16 Non-ECC | 512MB | 1GB | 2GB | n/a | 1066, 1333, | n/a | 1066, 1333 | n/a | 1066 |
| SRx8 ECC | 1GB | 2GB | 4GB | 1066, 1333 | 1066, 1333, | 1066, 1333 | 1066, 1333 | 1066 | 1066 |
| DRx8 ECC | 2GB | 4GB | 8GB | 1066, 1333 | 1066, 1333, | 1066, 1333 | 1066, 1333 | 1066 | 1066 |

**Notes:**

1. Supported DRAM Densities are 1Gb, 2Gb and 4Gb. Only 2Gb and 4Gb are validated by Intel®.
2. Command Address Timing is 1N for 1DPC and 2N for 2DPC.
3. For Memory Population Rules, please refer to the Romley Platform Design Guide.

| | |
|---|---|
| | Supported and Validated |
| | Supported but not Validated |

**Table 5. RDIMM Support Guidelines**

| Ranks Per DIMM and Data Width | Memory Capacity Per DIMM[1] | | | Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)[2,3,4] | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | | | 1 Slot per Channel | | 2 Slots per Channel | | | |
| | | | | 1DPC | | 1DPC | | 2DPC | |
| | | | | 1.35V | 1.5V | 1.35V | 1.5V | 1.35V | 1.5V |
| SRx8 | 1GB | 2GB | 4GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| DRx8 | 2GB | 4GB | 8GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| SRx4 | 2GB | 4GB | 8GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| DRx4 | 4GB | 8GB | 16GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| QRx4 | 8GB | 16GB | 32GB | 800 | 800 | 800 | 800 | 800 | 800 |
| QRx8 | 4GB | 8GB | 16GB | 800 | 800 | 800 | 800 | 800 | 800 |

**Notes:**

1. Supported DRAM Densities are 1Gb, 2Gb and 4Gb. Only 2Gb and 4Gb are validated by Intel®.
2. Command Address Timing is 1N.
3. For Memory Population Rules, please refer to the Romley Platform Design Guide.
4. QR RDIMMs are supported but not validated by Intel®, QR LRDIMMs are supported and validated by Intel®.

| | |
|---|---|
| | Supported and Validated |
| | Supported but not Validated |

**Table 6. LRDIMM Support Guidelines**

| Ranks Per DIMM and Data Width | Memory Capacity Per DIMM[1] | | LR-DIMM Speed (MT/s) and Voltage Validated by Slot Per Channel (SPC) and DIMM Per Channel (DPC) [2,3,4] | | | | | |
|---|---|---|---|---|---|---|---|---|
| | | | 1 Slot per Channel | | 2 Slots per Channel | | | |
| | | | 1DPC | | 1DPC | | 2DPC | |
| | | | 1.35V | 1.5V | 1.35V | 1.5V | 1.35V | 1.5V |
| QRx8(P)[6] | 8GB | 16GB | 1066 | 1066 1333 | 1066 | 1066 1333 | 1066 | 1066 |
| QRx4(DDP)[6] | 16GB | 32GB | 1066 | 1066 1333 | 1066 | 1066 1333 | 1066 | 1066 |

**Notes:**

1. Physical Rank is used to calculate DIMM Capacity.
2. Supported and validated DRAM Densities are 2Gb and 4Gb.
3. Command Address Timing is 1N.
4. The speeds are estimated targets and will be verified through simulation.
5. For Memory Population Rules, please refer to the Romley Platform Design Guide.
6. DDP - Dual Die Package DRAM stacking. P – Planer monolithic DRAM Die.

| | |
|---|---|
| | Supported and Validated |
| | Supported but not Validated |

### 3.6.2 Memory population rules

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

Each processor provides four banks of memory, each capable of supporting up to 2 DIMMs.

- DIMMs are organized into physical slots on DDR3 memory channels that belong to processor sockets.
- The memory channels from processor socket 1 are identified as Channel A, B and C. The memory channels from processor socket 2 are identified as Channel D, E and F.
- The silk screened DIMM slot identifiers on the board provide information about the channel, and therefore the processor to which they belong. For example, DIMM_A1 is the first slot on Channel A on processor 1; DIMM_D1 is the first DIMM socket on Channel D on processor 2.
- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- A processor may be installed without populating the associated memory slots provided a second processor is installed with associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.

  Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS, Error Management,) in the BIOS setup are applied commonly across processor sockets.

On the Intel® Server Board S2400GP, a total of 12 DIMM slots is provided (2 CPUs and 3 Channels/CPU). The nomenclature for DIMM sockets is detailed in the following table:

**Table 7. Intel® Server Board S2400GP DIMM Nomenclature**

| Processor Socket 1 | | | Processor Socket 2 | | |
|---|---|---|---|---|---|
| (0) Channel A | (1) Channel B | (2) Channel C | (0) Channel D | (1) Channel E | (2) Channel F |

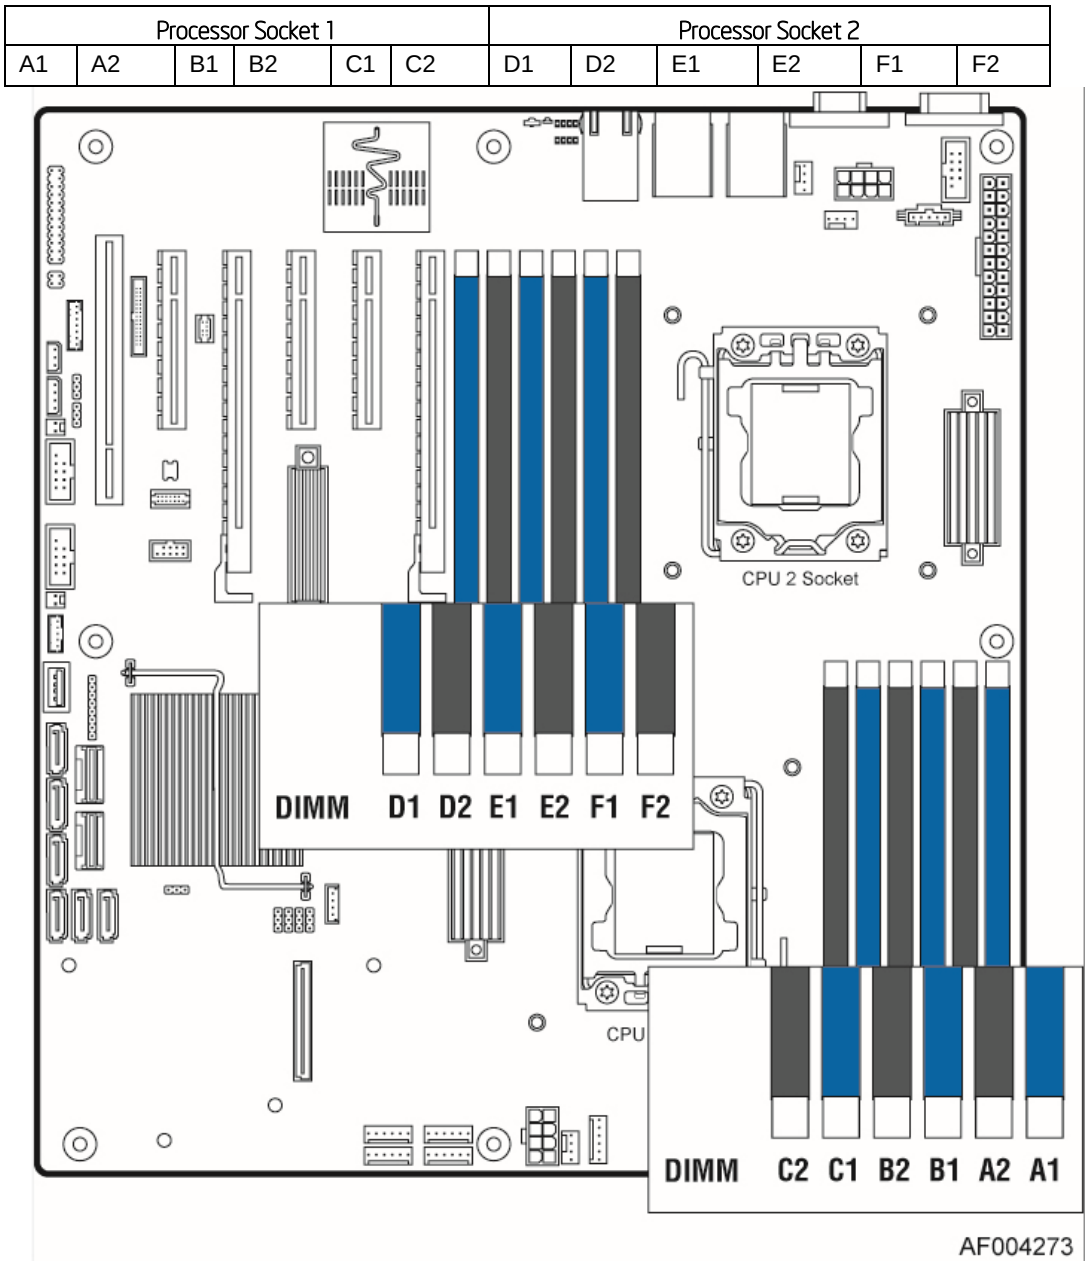


**Figure 15. Intel® Server Board S2400GP DIMM Slot Layout**

The following are generic DIMM population requirements that generally apply to the Intel® Server Board S2400GP.

- All DIMMs must be DDR3 DIMMs
- Registered DIMMs must be ECC only; Unbuffered DIMMs can be ECC or non-ECC. However, Intel® only validates and supports ECC memory for its server products.
- Mixing of Registered and Unbuffered DIMMs is not allowed per platform.
- Mixing of LRDIMM with any other DIMM type is not allowed per platform.

- Mixing of DDR3 voltages is not validated within a socket or across sockets by Intel®. If 1.35V (DDR3L) and 1.50V (DDR3) DIMMs are mixed, the DIMMs will run at 1.50V.
- Mixing of DDR3 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 eight logical ranks (ranks seen by the host) per channel is allowed.
- Mixing of ECC and non-ECC DIMMs is not allowed per platform.
- 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.

DIMM population rules require that DIMMs within a channel be populated starting with the BLUE DIMM slot or DIMM farthest from the processor in a “fill-farthest” approach. In addition, when populating a Quad-rank DIMM with a Single- or Dual-rank DIMM in the same channel, the Quad-rank DIMM must be populated farthest from the processor. Intel® MRC will check for correct DIMM placement.

### 3.6.3 Publishing System Memory

- The BIOS displays the “Total Memory” of the system during POST if Display Logo is disabled in the BIOS setup. This is the total size of memory discovered by the BIOS during POST, and is the sum of the individual sizes of installed DDR3 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 DDR3 DIMMs that are active (not disabled) and not used as redundant units.
- The BIOS provides the total memory of the system in the main page of the BIOS setup. This total is the same as the amount described by the first bullet above.
- If Display Logo 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.

**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

### 3.6.4 RAS Features

The server board supports the following memory RAS modes:

- Independent Channel Mode
- Rank Sparing Mode
- Mirrored Channel Mode
- Lockstep Channel Mode
- Single Device Data Correction (SDDC)
- Error Correction Code (ECC) Memory
- Demand Scrubbing for ECC Memory
- Patrol Scrubbing for ECC Memory

Regardless of RAS mode, the requirements for populating within a channel given in the section 3.6.2 must be met at all times. Note that support of RAS modes that require matching DIMM population between channels (Mirrored and Lockstep) require that ECC DIMMs be populated. Independent Channel Mode is the only mode that supports non-ECC DIMMs in addition to ECC DIMMs.

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

#### 3.6.4.1 Independent Channel Mode

In non-ECC and x4 SDDC configurations, each channel is running independently (nonlock-step), that is, each cache-line from memory is provided by a channel. To deliver the 64-byte cache-line of data, each channel is bursting eight 8-byte chunks. Back to back data transfer in the same direction and within the same rank can be sent back-to-back without any dead-cycle. The independent channel mode is the recommended method to deliver most efficient power and bandwidth as long as the x8 SDDC is not required.

#### 3.6.4.2 Rank Sparing Mode

In Rank Sparing Mode, one rank is a spare of the other ranks on the same channel. The spare rank is held in reserve and is not available as system memory. The spare rank must have identical or larger memory capacity than all the other ranks (sparing source ranks) on the same channel. After sparing, the sparing source rank will be lost.

Rank Sparing Mode enhances the system's RAS capability by "swapping out" failing ranks of DIMMs. Rank Sparing is strictly channel and rank oriented. Each memory channel is a Sparing Domain.

For Rank Sparing to be available as a RAS option, there must be 2 or more single rank or dual rank DIMMs, or at least one quad rank DIMM installed on each memory channel.

Rank Sparing Mode is enabled or disabled in the Memory RAS and Performance Configuration screen in the <F2> Bios Setup Utility.

When Sparing Mode is operational, for each channel, the largest size memory rank is reserved as a “spare” and is not used during normal operations. The impact on Effective Memory Size is to subtract the sum of the reserved ranks from the total amount of installed memory.

Hardware registers count the number of Correctable ECC Errors for each rank of memory on each channel during operations and compare the count against a Correctable Error Threshold. When the correctable error count for a given rank hits the threshold value, that rank is deemed to be “failing”, and it triggers a Sparing Fail Over (SFO) event for the channel in which that rank resides. The data in the failing rank is copied to the Spare Rank for that channel, and the Spare Rank replaces the failing rank in the IMC’s address translation registers.

An SFO Event is logged to the BMC SEL. The failing rank is then disabled, and any further Correctable Errors on that now non-redundant channel will be disregarded.

The correctable error that triggered the SFO may be logged to the BMC SEL, if it was the first one to occur in the system. That first correctable error event will be the only one logged for the system. However, since each channel is a Sparing Domain, the correctable error counting continues for other channels which are still in a redundant state. There can be as many SFO Events as there are memory channels with DIMMs installed.

#### 3.6.4.3 Mirrored Channel Mode

Channel Mirroring Mode gives the best memory RAS capability by maintaining two copies of the data in main memory. If there is an Uncorrectable ECC Error, the channel with the error is disabled and the system continues with the “good” channel, but in a non-redundant configuration.

For Mirroring mode to be to be available as a RAS option, the DIMM population must be identical between each pair of memory channels that participate. Not all channel pairs need to have memory installed, but for each pair, the configuration must match. If the configuration is not matched up properly, the memory operating mode falls back to Independent Channel Mode.

Mirroring Mode is enabled/disabled in the Memory RAS and Performance Configuration screen in the <F2> BIOS Setup Utility.

When Mirroring Mode is operational, each channel in a pair is “mirrored” by the other channel. The impact on Effective Memory size is to reduce by half the total amount of installed memory available for use.

When Mirroring Mode is operational, the system treats Correctable Errors the same way as it would in Independent channel mode. There is a correctable error threshold. Correctable error counts accumulate by rank, and the first event is logged.

What Mirroring primarily protects against is the possibility of an Uncorrectable ECC Error occurring with critical data “in process”. Without Mirroring, the system would be expected to “Blue Screen” and halt, possibly with serious impact to operations. But with Mirroring Mode in operation, an Uncorrectable ECC Error from one channel becomes a Mirroring Fail Over (MFO) event instead, in which the IMC retrieves the correct data from the “mirror image” channel and disables the failed channel. Since the ECC Error was corrected in the process of the MFO Event, the ECC Error is demoted to a Correctable ECC Error. The channel pair becomes a single non-redundant channel, but without impacting operations, and the Mirroring Fail Over Event is

logged to SEL to alert the user that there is memory hardware that has failed and needs to be replaced.

In Mirrored Channel Mode, the memory contents are mirrored between Channel B and Channel C and also between Channel E and Channel F. As a result of the mirroring, the total physical memory available to the system is half of what is populated. Mirrored Channel Mode requires that Channel B and Channel C, and Channel E and Channel F must be populated identically with regards to size and organization. DIMM slot populations within a channel do not have to be identical but the same DIMM slot location across Channel B and Channel C and across Channel E and Channel F must be populated the same.

#### 3.6.4.4 Lockstep Channel Mode

In lockstep channel mode the cache-line is split across channels. This is done to support Single Device Data Correction (SDDC) for DRAM devices with 8-bit wide data ports. Also, the same address is used on both channels, such that an address error on any channel is detectable by bad ECC. The iMC module always accumulates 32-bytes before forwarding data so there is no latency benefit for disabling ECC.

Lockstep channels must be populated identically. That is, each DIMM in one channel must have a corresponding DIMM of identical organization (number ranks, number banks, number rows, number columns). DIMMs may be of different speed grades, but the iMC module will be configured to operate all DIMMs according to the slowest parameters present by the Memory Reference Code (MRC).

Performance in lockstep mode cannot be as high as with independent channels. The burst length for DDR3 DIMMs is eight which is shared between two channels that are in lockstep mode. Each channel of the pair provides 32 bytes to produce the 64-byte cache-line. DRAMs on independent channels are configured to deliver a burst length of eight. The maximum read bandwidth for a given Rank is half of peak. There is another
draw back in using lockstep mode, that is, higher power consumption since the total activation power is about twice of the independent channel operation if comparing to same type of DIMMs.

In Lockstep Channel Mode, each memory access is a 128-bit data access that spans Channel B and Channel C, and Channel E and Channel F. Lockstep Channel mode is the only RAS mode that allows SDDC for x8 devices. Lockstep Channel Mode requires that Channel B and Channel C, and Channel E and Channel F must be populated identically with regards to size and organization. DIMM slot populations within a channel do not have to be identical but the same DIMM slot location across Channel B and Channel C and across Channel E and Channel F must be populated the same.

#### 3.6.4.5 Single Device Data Correction (SDDC)

SDDC – Single Device Data Correction is a technique by which data can be replaced by the IMC from an entire x4 DRAM device which is failing, using a combination of CRC plus parity. This is an automatic IMC driven hardware. It can be extended to x8 DRAM technology by placing the system in Channel Lockstep Mode.

#### 3.6.4.6 Error Correction Code (ECC) Memory

ECC uses “extra bits” – 64-bit data in a 72-bit DRAM array – to add an 8-bit calculated “Hamming Code” to each 64 bits of data. This additional encoding enables the memory

controller to detect and report single or multiple bit errors when data is read, and to correct single-bit errors.

#### *3.1.1.1.1 Correctable Memory ECC Error Handling*

A “Correctable ECC Error” is one in which a single-bit error in memory contents is detected and corrected by use of the ECC Hamming Code included in the memory data. For a correctable error, data integrity is preserved, but it may be a warning sign of a true failure to come. Note that some correctable errors are expected to occur.

The system BIOS has logic to cope with the random factor in correctable ECC errors. Rather than reporting every correctable error that occurs, the BIOS has a threshold and only logs a correctable error when a threshold value is reached. Additional correctable errors that occur after the threshold has been reached are disregarded. In addition, on the expectation the server system may have extremely long operational runs without being rebooted, there is a “Leaky Bucket” algorithm incorporated into the correctable error counting and comparing mechanism. The “Leaky Bucket” algorithm reduces the correctable error count as a function of time – as the system remains running for a certain amount of time, the correctable error count will “leak out” of the counting registers. This prevents correctable error counts from building up over an extended runtime.

The correctable memory error threshold value is a configurable option in the <F2> BIOS Setup Utility, where you can configure it for 20/10/5/ALL/None

Once a correctable memory error threshold is reached, the event is logged to the System Event Log (SEL) and the appropriate memory slot fault LED is lit to indicate on which DIMM the correctable error threshold crossing occurred.

#### *3.1.1.1.2 Uncorrectable Memory ECC Error Handling*

All multi-bit “detectable but not correctable“ memory errors are classified as Uncorrectable Memory ECC Errors. This is generally a fatal error.

However, before returning control to the OS drivers through Machine Check Exception (MCE) or Non-Maskable Interrupt (NMI), the Uncorrectable Memory ECC Error is logged to the SEL, the appropriate memory slot fault LED is lit, and the System Status LED state is changed to solid Amber.

### 3.6.4.7 Demand Scrubbing for ECC Memory

Demand scrubbing is the ability to write corrected data back to the memory once a correctable error is detected on a read transaction. This allows for correction of data in memory at detect, and decrease the chances of a second error on the same address accumulating to cause a multi-bit error (MBE) condition.

Demand Scrubbing is enabled/disabled (default is enabled) in the Memory Configuration screen in Setup.

### 3.6.4.8 Patrol Scrubbing for ECC Memory

Patrol scrubs are intended to ensure that data with a correctable error does not remain in DRAM long enough to stand a significant chance of further corruption to an uncorrectable stage.

## 3.7 Processor Integrated I/O Module (IIO)

The processor's integrated I/O module provides features traditionally supported through chipset components. The integrated I/O module provides the following features:

- **PCI Express* Interfaces**: The integrated I/O module incorporates the PCI Express* interface and supports up to 24 lanes of PCI Express*. Following are key attributes of the PCI Express* interface:
  - Gen3 speeds at 8 GT/s (no 8b/10b encoding)
    - The Intel® Server Board S2400GP supports PCIe slots from two processors:
  - From the first processor:
    - Slot 2: PCIe Gen3 x4 electrical with x8 physical connector
    - Slot 6: PCIe Gen3 x16 electrical with x16 physical connector
  - From the second processor:
    - Slot 3: PCIe Gen3 x16 electrical with x16 physical connector
    - Slot 4: PCIe Gen3 x8 electrical with x8 physical connector
    - Slot 5: PCIe Gen3 x4 electrical with x8 physical connector
- **DMI2 Interface to the PCH**: The platform requires an interface to the legacy Southbridge (PCH) which provides basic, legacy functions required for the server platform and operating systems. Since only one PCH is required and allowed for the system, any sockets which do not connect to PCH would use this port as a standard x4 PCI Express* 2.0 interface.
- **Integrated IOAPIC**: Provides support for PCI Express* devices implementing legacy interrupt messages without interrupt sharing.
- **Non Transparent Bridge:** 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.
- **Intel® QuickData Technology**: Used for efficient, high bandwidth data movement between two locations in memory or from memory to I/O.

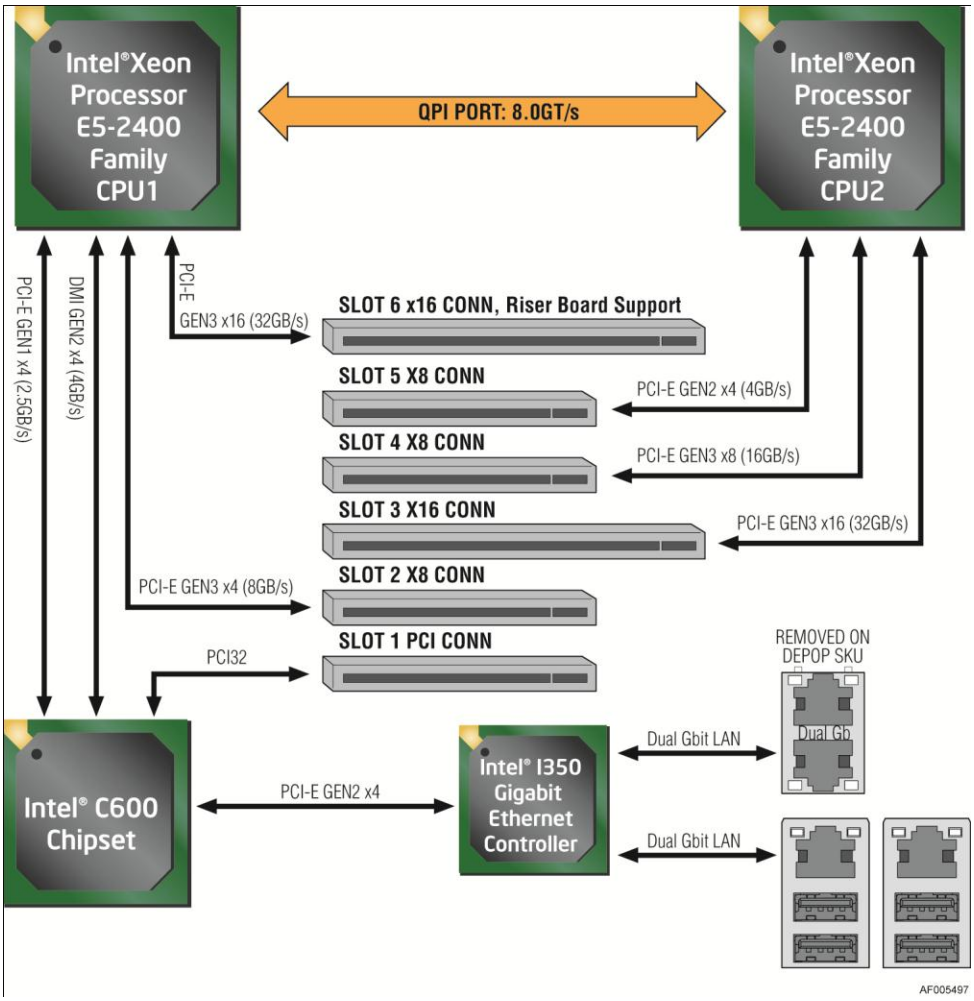


**Figure 16. Functional Block Diagram of Processor IIO Sub-system**

The following sub-sections will describe the server board features that are directly supported by the processor IIO module. These include the Riser Card Slots, Network Interface, and connectors for the optional I/O modules and SAS Module. Features and functions of the Intel® C600 Series chipset will be described in its own dedicated section.

### 3.7.1 Network Interface

Network connectivity is provided by means of two onboard Intel® Ethernet Controller I350 providing up to two 10/100/1000 Mb Ethernet ports. The NIC chip is supported by implementing x4 PCIe Gen2 signals from the Intel® C600 PCH.

On the Intel® Server Board S2400GP, two for S2400GP2 and four for S2400GP4 external 10/100/1000 Mb RJ45 Ethernet ports are provided. Each Ethernet port drives two or four LEDs located on each network interface connector. The LED at the right of the connector is the link/activity LED and indicates network connection when on, and transmit/receive activity when blinking. The LED at the left of the connector indicates link speed as defined in the following table:

**Table 8. External RJ45 NIC Port LED Definition**

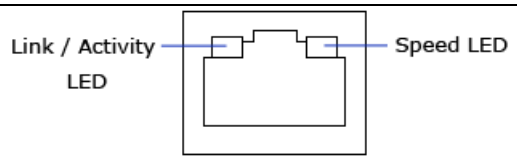


| LED Color | LED State | NIC State |
|---|---|---|
| Green/Amber (Right) | Off | 10 Mbps |
| | Amber | 100 Mbps |
| | Green | 1000 Mbps |
| Green (Left) | On | Active Connection |
| | Blinking | Transmit/Receive activity |

The server board has seven MAC addresses programmed at the factory for S2400GP4. MAC addresses are assigned as follows:

- NIC 1 MAC address (for OS usage)
- NIC 2 MAC address = NIC 1 MAC address + 1 (for OS usage)
- NIC 3 MAC address = NIC 1 MAC address + 2 (for OS usage)
- NIC 4 MAC address = NIC 1 MAC address + 3 (for OS usage)
- BMC LAN channel 1 MAC address = NIC1 MAC address + 4
- BMC LAN channel 2 MAC address = NIC1 MAC address + 5
- BMC LAN channel 3 (RMM) MAC address = NIC1 MAC address + 6

The server board has five MAC addresses programmed at the factory for S2400GP2. MAC addresses are assigned as follows:

- NIC 1 MAC address (for OS usage)
- NIC 2 MAC address = NIC 1 MAC address + 1 (for OS usage)
- BMC LAN channel 1 MAC address = NIC1 MAC address + 2
- BMC LAN channel 2 MAC address = NIC1 MAC address + 3
- BMC LAN channel 3 (RMM) MAC address = NIC1 MAC address + 4

The printed MAC address on the server board and/or server system is assigned to NIC1 on the server board.

## 3.8 Intel® C602-A Chipset Functional Overview

The following sub-sections will provide an overview of the key features and functions of the Intel® C602-A chipset used on the server board. For more comprehensive chipset specific information, refer to the Intel® C600 Series chipset documents listed in the Reference Document

list in Chapter 1.

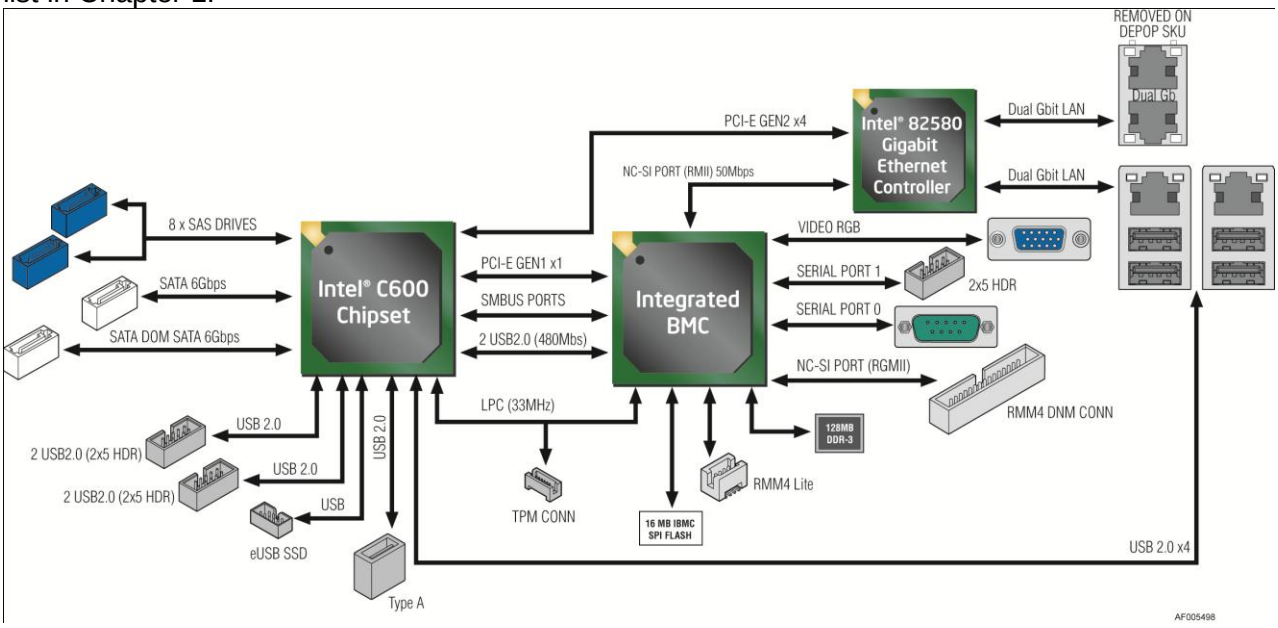


**Figure 17. Functional Block Diagram – Chipset Supported Features and Functions**

On the Intel® Server Boards S2400GP, the chipset provides support for the following on-board functions:

- Digital Media Interface (DMI)
- PCI Express* Interface
- Serial ATA (SATA) Controller
- Serial Attached SCSI (SAS)/SATA Controller
- AHCI
- Rapid Storage Technology
- PCI Interface
- Low Pin Count (LPC) interface
- Serial Peripheral Interface (SPI)
- Compatibility Modules (DMA Controller, Timer/Counters, Interrupt Controller)
- Advanced Programmable Interrupt Controller (APIC)
- Universal Serial Bus (USB) Controller
- Gigabit Ethernet Controller
- RTC
- GPIO
- Enhanced Power Management
- Manageability
- System Management Bus (SMBus* 2.0)
- Intel® Active Management Technology (Intel® AMT)

- Integrated NVSRAM controller
- Virtualization Technology for Direct I/O (Intel® VT-d)
- JTAG Boundary-Scan
- KVM/Serial Over LAN (SOL) Function

### 3.8.1 Digital Media Interface (DMI)

Digital Media Interface (DMI) is the chip-to-chip connection between the processor and C600 chipset. This high-speed interface integrates advanced priority-based servicing allowing for concurrent traffic and true isochronous transfer capabilities. Base functionality is completely software-transparent, permitting current and legacy software to operate normally.

### 3.8.2 PCI Express* Interface

The C600 chipset provides up to 8 PCI Express* Root Ports, supporting the *PCI Express* Base Specification*, Revision 2.0. Each Root Port x1 lane supports up to 5 Gb/s bandwidth in each direction (10 Gb/s concurrent). PCI Express* Root Ports 1-4 or Ports 5-8 can independently be configured to support four x1s, two x2s, one x2 and two x1s,or one x4 port widths.

### 3.8.3 Serial ATA (SATA) Controller

The C600 chipset has two integrated SATA host controllers that support independent DMA operation on up to six ports and supports data transfer rates of up to 6.0 Gb/s (600 MB/s) on up to two ports (Port 0 and 1 Only) while all ports support rates up to 3.0 Gb/s (300 MB/s) and up to 1.5 Gb/s (150 MB/s). The SATA controller contains two modes of operation – a legacy mode using I/O space, and an AHCI mode using memory space. Software that uses legacy mode will not have AHCI capabilities. The C600 chipset supports the *Serial ATA Specification*, Revision 3.0. The C600 also supports several optional sections of the *Serial ATA II: Extensions to Serial ATA 1.0 Specificatio*n, Revision 1.0 (AHCI support is required for some elements).

### 3.8.4 AHCI

The C600 chipset provides hardware support for Advanced Host Controller Interface (AHCI), a standardized programming interface for SATA host controllers. Platforms supporting AHCI may take advantage of performance features such as no master/slave designation for SATA devices—each device is treated as a master—and hardware assisted native command queuing. AHCI also provides usability enhancements such as Hot-Plug. AHCI requires appropriate software support (for example, an AHCI driver) and for some features, hardware support in the SATA device or additional platform hardware.

### 3.8.5 Rapid Storage Technology

The C600 chipset provides support for Intel® Rapid Storage Technology, providing both AHCI (see above for details on AHCI) and integrated RAID functionality. The industry-leading RAID capability provides high-performance RAID 0, 1, 5, and 10 functionality on up to 6 SATA ports of the C600 chipset. Matrix RAID support is provided to allow multiple RAID levels to be combined on a single set of hard drives, such as RAID 0 and RAID 1 on two disks. Other RAID features include hot-spare support, SMART alerting, and RAID 0 auto replace. Software components include an Option ROM for pre-boot configuration and boot functionality, a Microsoft Windows* compatible driver, and a user interface for configuration and management of the RAID capability of the C600 chipset.

### 3.8.6 PCI Interface

The C600 chipset PCI interface provides a 33 MHz, Revision 2.3 implementation. The C600 chipset integrates a PCI arbiter that supports up to four external PCI bus masters in addition to the internal C600 chipset requests. This allows for combinations of up to four PCI down devices and PCI slots.

### 3.8.7 Low Pin Count (LPC) Interface

The C600 chipset implements an LPC Interface as described in the *LPC 1.1 Specification*. The Low Pin Count (LPC) bridge function of the C600 resides in PCI Device 31: Function 0. In addition to the LPC bridge interface function, D31:F0 contains other functional units including DMA, interrupt controllers, timers, power management, system management, GPIO, and RTC.

### 3.8.8 Serial Peripheral Interface (SPI)

The C600 chipset implements an SPI Interface as an alternative interface for the BIOS flash device. An SPI flash device can be used as a replacement for the FWH, and is required to support Gigabit Ethernet and Intel® Active Management Technology. The C600 chipset supports up to two SPI flash devices with speeds up to 50 MHz, utilizing two chip select pins.

### 3.8.9 Compatibility Modules (DMA Controller, Timer/Counters, Interrupt Controller)

The DMA controller incorporates the logic of two 82C37 DMA controllers, with seven independently programmable channels. Channels 0–3 are hardwired to 8-bit, count-by-byte transfers, and channels 5–7 are hardwired to 16-bit, count-by-word transfers. Any two of the seven DMA channels can be programmed to support fast Type-F transfers. Channel 4 is reserved as a generic bus master request.

The C600 chipset supports LPC DMA, which is similar to ISA DMA, through the C600 chipset's DMA controller. LPC DMA is handled through the use of the LDRQ# lines from peripherals and special encoding on LAD[3:0] from the host. Single, Demand, Verify, and Increment modes are supported on the LPC interface.

The timer/counter block contains three counters that are equivalent in function to those found in one 82C54 programmable interval timer. These three counters are combined to provide the system timer function, and speaker tone. The 14.31818 MHz oscillator input provides the clock source for these three counters.

The C600 chipset provides an ISA-Compatible Programmable Interrupt Controller (PIC) that incorporates the functionality of two, 82C59 interrupt controllers. The two interrupt controllers are cascaded so that 14 external and two internal interrupts are possible. In addition, the C600 chipset supports a serial interrupt scheme.

### 3.8.10 Advanced Programmable Interrupt Controller (APIC)

In addition to the standard ISA compatible Programmable Interrupt controller (PIC) described in the previous section, the C600 incorporates the Advanced Programmable Interrupt Controller (APIC).

### 3.8.11 Universal Serial Bus (USB) Controller

The C600 chipset has up to two Enhanced Host Controller Interface (EHCI) host controllers that support USB high-speed signaling. High-speed USB 2.0 allows data transfers up to 480 Mb/s which is 40 times faster than full-speed USB. The C600 chipset supports up to fourteen USB 2.0 ports. All fourteen ports are high-speed, full-speed, and low-speed capable.

### 3.8.12 Gigabit Ethernet Controller

The Gigabit Ethernet Controller provides a system interface using a PCI function. The controller provides a full memory-mapped or IO mapped interface along with a 64 bit address master support for systems using more than 4 GB of physical memory and DMA (Direct Memory Addressing) mechanisms for high performance data transfers. Its bus master capabilities enable the component to process high-level commands and perform multiple operations; this lowers processor utilization by off-loading communication tasks from the processor. Two large configurable transmit and receive FIFOs (up to 20 KB each) help prevent data under-runs and overruns while waiting for bus accesses. This enables the integrated LAN controller to transmit data with minimum interframe spacing (IFS).

The LAN controller can operate at multiple speeds (10/100/1000 MB/s) and in either full duplex or half duplex mode. In full duplex mode the LAN controller adheres with the *IEEE 802.3x Flow Control Specification*. Half duplex performance is enhanced by a proprietary collision reduction mechanism.

### 3.8.13 RTC

The C600 chipset contains a Motorola MC146818B-compatible real-time clock with 256 bytes of battery-backed RAM. The real-time clock performs two key functions: keeping track of the time of day and storing system data, even when the system is powered down. The RTC operates on a 32.768 KHz crystal and a 3 V battery. The RTC also supports two lockable memory ranges. By setting bits in the configuration space, two 8-byte ranges can be locked to read and write accesses. This prevents unauthorized reading of passwords or other system security information. The RTC also supports a date alarm that allows for scheduling a wake up event up to 30 days in advance, rather than just 24 hours in advance.

### 3.8.14 GPIO

Various general purpose inputs and outputs are provided for custom system design. The number of inputs and outputs varies depending on the C600 chipset configuration.

### 3.8.15 Enhanced Power Management

The C600 chipset's power management functions include enhanced clock control and various low-power (suspend) states (for example, Suspend-to-RAM and Suspend-to-Disk). A hardware-based thermal management circuit permits software-independent entrance to low-power states. The C600 chipset contains full support for the *Advanced Configuration and Power Interface (ACPI) Specification*, Revision 4.0a.

### 3.8.16 Manageability

The chipset integrates several functions designed to manage the system and lower the total cost of ownership (TCO) of the system. These system management functions are designed to report errors, diagnose the system, and recover from system lockups without the aid of an external microcontroller.

- **TCO Timer.** The chipset's integrated programmable TCO timer is used to detect system locks. The first expiration of the timer generates an SMI# that the system can use to recover from a software lock. The second expiration of the timer causes a system reset to recover from a hardware lock.
- **Processor Present Indicator.** The chipset looks for the processor to fetch the first instruction after reset. If the processor does not fetch the first instruction, the chipset will reboot the system.
- **ECC Error Reporting.** When detecting an ECC error, the host controller has the ability to send one of several messages to the chipset. The host controller can instruct the chipset to generate SMI#, NMI, SERR#, or TCO interrupt.
- **Function Disable.** The chipset provides the ability to disable the following integrated functions: LAN, USB, LPC, SATA, PCI Express* or SMBus*. Once disabled, these functions no longer decode I/O, memory, or PCI configuration space. Also, no interrupts or power management events are generated from the disabled functions.
- **Intruder Detect.** The chipset provides an input signal (INTRUDER#) that can be attached to a switch that is activated by the system case being opened. The chipset can be programmed to generate an SMI# or TCO interrupt due to an active INTRUDER# signal.

### 3.8.17 System Management Bus (SMBus* 2.0)

The C600 chipset contains a SMBus* Host interface that allows the processor to communicate with SMBus* slaves. This interface is compatible with most I2C devices. Special I2C commands are implemented. The C600 chipset's SMBus* host controller provides a mechanism for the processor to initiate communications with SMBus* peripherals (slaves). Also, the C600 chipset supports slave functionality, including the Host Notify protocol. Hence, the host controller supports eight command protocols of the SMBus* interface (see *System Management Bus (SMBus*) Specification*, Version 2.0): Quick Command, Send Byte, Receive Byte, Write Byte/Word, Read Byte/Word, Process Call, Block Read/Write, and Host Notify.

The C600 chipset's SMBus* also implements hardware-based Packet Error Checking for data robustness and the Address Resolution Protocol (ARP) to dynamically provide address to all SMBus* devices.

### 3.8.18 Intel® Active Management Technology (Intel® AMT)

Intel® Active Management Technology (Intel® AMT) is the next generation of client manageability using the wired network. Intel® AMT is a set of advanced manageability features developed as a direct result of IT customer feedback gained through Intel® market research. With the new implementation of System Defense in C600 chipset, the advanced manageability feature set of Intel® AMT is further enhanced.

### 3.8.19 Integrated NVSRAM Controller

The C600 chipset has an integrated NVSRAM controller that supports up to 32KB external device. The host processor can read and write data to the NVSRAM component.

### 3.8.20 Intel® Virtualization Technology for Direct I/O (Intel® VT-d)

The C600 chipset provides hardware support for implementation of Intel® Virtualization Technology with Directed I/O (Intel® VT-d). Intel® VT-d consists of technology components that support the virtualization of platforms based on Intel® Architecture Processors. Intel® VT-d Technology enables multiple operating systems and applications to run in independent partitions. A partition behaves like a virtual machine (VM) and provides isolation and protection across partitions. Each partition is allocated its own subset of host physical memory.

### 3.8.21 JTAG Boundary-Scan

The C600 chipset adds the industry standard JTAG interface and enables Boundary-Scan in place of the XOR chains used in previous generations of chipsets. Boundary-Scan can be used to ensure device connectivity during the board manufacturing process. The JTAG interface allows system manufacturers to improve efficiency by using industry available tools to test the C600 chipset on an assembled board. Since JTAG is a serial interface, it eliminates the need to create probe points for every pin in an XOR chain. This eases pin breakout and trace routing and simplifies the interface between the system and a bed-of-nails tester.

### 3.8.22 KVM/Serial Over LAN (SOL) Function

These functions support redirection of keyboard, mouse, and text screen to a terminal window on a remote console. The keyboard, mouse, and text redirection enables the control of the client machine through the network without the need to be physically near that machine. Text, mouse, and keyboard redirection allows the remote machine to control and configure the client by entering BIOS setup. The KVM/SOL function emulates a standard PCI serial port and redirects the data from the serial port to the management console using LAN. KVM has additional requirements of internal graphics and SOL may be used when KVM is not supported.

### 3.8.23 On-board Serial Attached SCSI (SAS)/Serial ATA (SATA) Support and Options

The Intel® C602-A chipset provides storage support through two integrated controllers: AHCI and SCU. By default the server board will support up to 10 SATA ports: Two single 6Gb/sec SATA ports routed from the AHCI controller to the two white SATA connectors labeled “SATA_0” and “SATA_1”, four 3Gb/sec SATA ports routed from the AHCI controller to the four black SATA connectors labeled “SATA_2” to “SATA_5”, and four 3Gb/sec SATA ports routed from the SCU to the SFF8087 miniSAS port labeled “SCU_0”.

**Note:** The miniSAS connector labeled “SCU 1” is NOT functional by default and is only enabled with the addition of an Intel® RAID C600 Upgrade Key option supporting eight AS/SATA ports.

Standard are two embedded software RAID options using the storage ports configured from the SCU only:

- Intel® Embedded Server RAID Technology 2 (ESRT2) based on LSI* MegaRAID SW RAID technology supporting SATA RAID levels 0,1,10

- Intel® Rapid Storage Technology (RSTe) supporting SATA RAID levels 0,1,5,10

The server board is capable of supporting additional chipset embedded SAS and RAID options from the SCU controller when configured with one of several available Intel® RAID C600 Upgrade Keys. Upgrade keys install onto a 4-pin connector on the server board labeled "STOR_UPG_KEY". The following table identifies available upgrade key options and their supported features.

**Table 9. Intel® RAID C600 Upgrade Key Options**

| Product Code | Color | On-Server Board SATA/SAS Capable Controller | On-Server Board AHCI Capable SATA Controller |
|---|---|---|---|
| No Key | N/A | Intel® RSTE 4 ports SATA R0,1,10,5 or Intel® ESRT2 4 ports SATA R0,1,10 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10 |
| RKSATA4R5 | Black | Intel® RSTE 4 ports SATA R0,1,10,5 or Intel® ESRT2 4 ports SATA R0,1,10,5 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10,5 |
| RKSATA8 | Blue | Intel® RSTE 8 ports SATA R0,1,10,5 or Intel® ESRT2 8 ports SATA R0,1,10 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10 |
| RKSATA8R5 | White | Intel® RSTE 8 ports SATA R0,1,10,5 or Intel® ESRT2 8 ports SATA R0,1,10,5 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10,5 |
| RKSAS4 | Green | Intel® RSTE 4 ports SAS R0,1,10 or Intel® ESRT2 4 ports SAS R0,1,10 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10 |
| RKSAS4R5 | Yellow | Intel® RSTE 4 ports SAS R0,1,10 or Intel® ESRT2 4 ports SAS R0,1,10,5 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10,5 |
| RKSAS8 | Orange | Intel® RSTE 8 ports SAS R0,1,10 or Intel® ESRT2 8 ports SAS R0,1,10 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10 |
| RKSAS8R5 | Purple | Intel® RSTE 8 ports SAS R0,1,10 or Intel® ESRT2 8 ports SAS R0,1,10,5 | Intel® RSTE SATA R0,1,10,5 or Intel® ESRT2 SATA R0,1,10,5 |

Additional information for the on-board RAID features and functionality can be found in the *Intel® RAID Software Users Guide* (Intel Document Number D29305-018).

The system includes support for two embedded software RAID options:

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

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.

### 3.8.24 Intel® Embedded Server RAID Technology 2 (ESRT2)

Features of the embedded software RAID option Intel® Embedded Server RAID Technology 2 (ESRT2) include the following:

- Based on LSI* MegaRAID Software Stack
- Software RAID with system providing memory and CPU utilization
- Supported RAID Levels – 0,1,5,10
  - 4 and 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key
  - 4 and 8 Port SAS RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key
- Maximum drive support = Eight (with or without SAS expander option installed)
- 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

### 3.8.25 Intel® Rapid Storage Technology (RSTe)

Features of the embedded software RAID option Intel® Rapid Storage Technology (RSTe) include the following:

- Software RAID with system providing memory and CPU utilization
- Supported RAID Levels – 0,1,5,10
  - 4 Port SATA RAID 5 available standard (no option key required)
  - 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key
  - No SAS RAID 5 support
- Maximum drive support = 32 (in arrays with 8 port SAS), 16 (in arrays with 4 port SAS), 128 (JBOD)
- Open Source Compliance = Yes (uses MDRAID)
- OS Support = Windows 7*, Windows 2008*, Windows 2003*, RHEL* 6.2 and later, SLES* 11 w/SP2 and later, VMWare* 5.x.
- Utilities = Windows* GUI and CLI, Linux* CLI, DOS CLI, and EFI CLI
- Uses Matrix Storage Manager for Windows*
- MDRAID supported in Linux* (Does not require a driver)

**Note:** No boot drive support to targets attached through SAS expander card.

## 3.9 Integrated Baseboard Management Controller (INTEGRATED BMC) Overview

The server board utilizes the I/O controller, Graphics Controller, and Baseboard Management features of the Server Engines* Pilot-III Server Management Controller. The following is an overview of the features as implemented on the server board from each embedded controller.

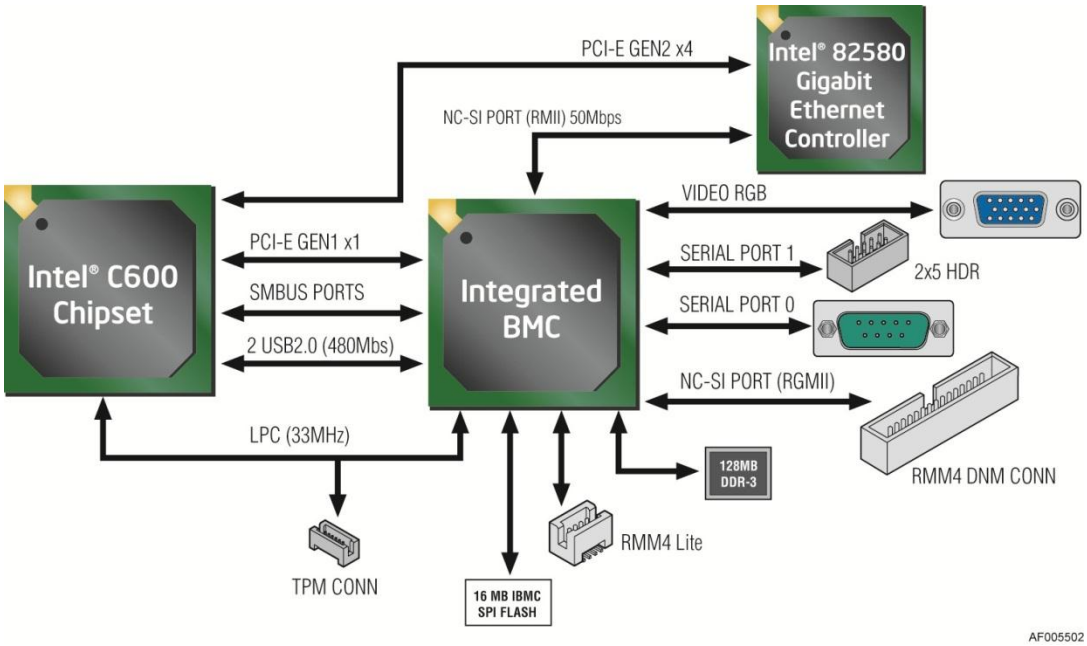


**Figure 18. Integrated Baseboard Management Controller (BMC) Overview**

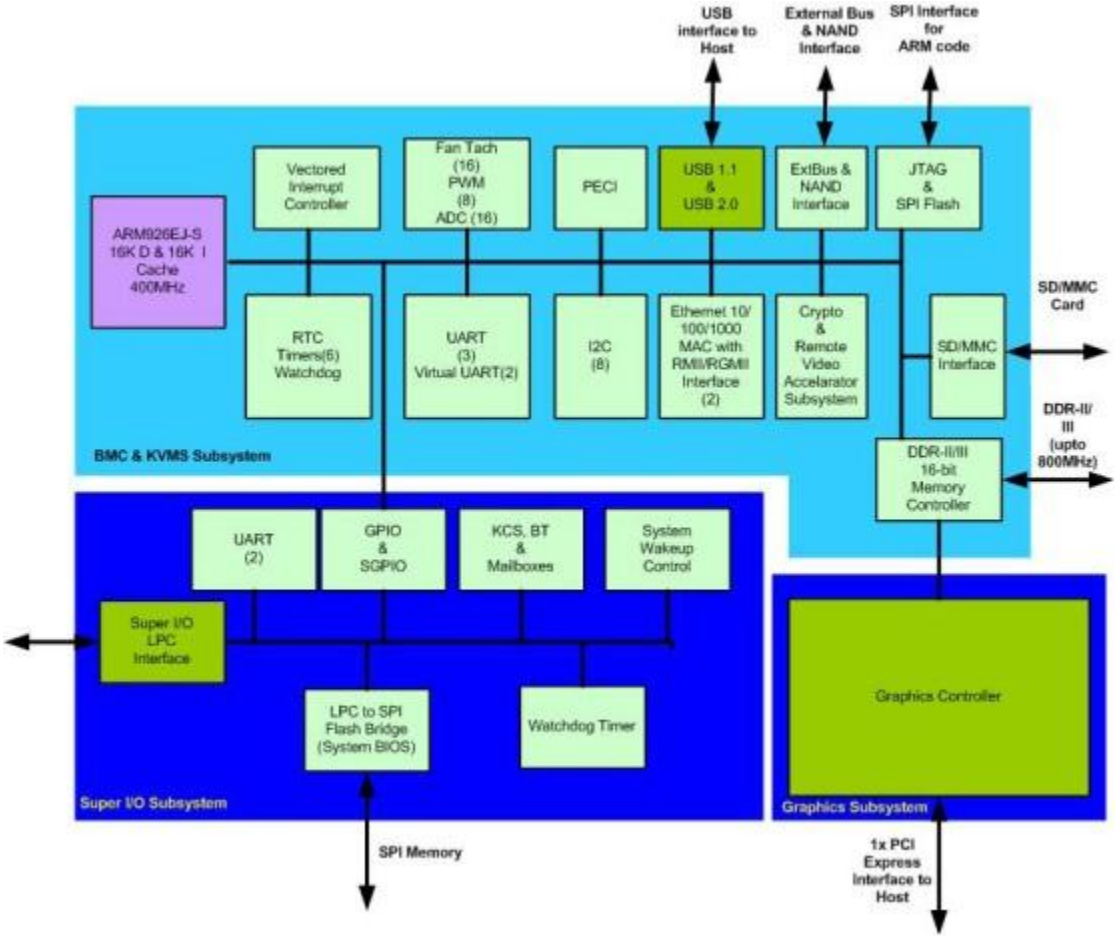

**Figure 19. Integrated BMC Hardware**

## 3.10 Super I/O Controller

The integrated super I/O controller provides support for the following features as implemented on the server board:

- Two Fully Functional Serial Ports, compatible with the 16C550
- Serial IRQ Support
- Up to 16 Shared direct GPIO's
- Serial GPIO support for 80 general purpose inputs and 80 general purpose outputs available for host processor
- Programmable Wake-up Event Support
- Plug and Play Register Set
- Power Supply Control
- Host SPI bridge for system BIOS support

### 3.10.1 Keyboard and Mouse Support

The server board does not support PS/2 interface keyboards and mice. However, the system BIOS recognizes USB specification-compliant keyboard and mice.

### 3.10.2 Wake-up Control

The super I/O contains functionality that allows various events to power on and power off the system.

## 3.11 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 supporting 128MB of memory
- Supports display resolutions up to 1600 x 1200 16bpp @ 60Hz
- 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 10. 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 |

| 2D Mode | 2D Video Mode Support | | | |
|---|---|---|---|---|
| | 8 bpp | 16 bpp | 24 bpp | 32 bpp |
| 1152x864 | X | X | X | X |
| 1280x1024 | X | X | X | X |
| 1600x1200** | X | X | | |

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

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

The BIOS supports dual-video mode when an add-in video card is installed.

- In the single mode (dual monitor video = disabled), the on-board video controller is disabled when an add-in video card is detected.
- In the dual mode (on-board video = enabled, dual monitor video = enabled), the on-board video controller is enabled and is the primary video device. The add-in video card is allocated resources and is considered the secondary video device. The BIOS Setup utility provides options to configure the feature as follows:

**Table 11. Video mode**

| | | |
|---|---|---|
| On-board Video | Enabled<br>Disabled | |
| Dual Monitor Video | Enabled<br>Disabled | Shaded if on-board video is set to "Disabled" |

## 3.12 Baseboard Management Controller

The server board utilizes the following features of the embedded baseboard management controller.

- IPMI 2.0 Compliant
- 400MHz 32-bit ARM9 processor with memory management unit (MMU)
- Two independent10/100/1000 Ethernet Controllers with RMII/RGMII support
- DDR2/3 16-bit interface with up to 800 MHz operation
- 12 10-bit ADCs
- Fourteen fan tachometers
- Eight Pulse Width Modulators (PWM)
- Chassis intrusion logic
- JTAG Master

- Eight I2C interfaces with master-slave and SMBus timeout support. All interfaces are SMBus 2.0 compliant.
- Parallel general-purpose I/O Ports (16 direct, 32 shared)
- Serial general-purpose I/O Ports (80 in and 80 out)
- Three UARTs
- Platform Environmental Control Interface (PECI)
- Six general-purpose timers
- Interrupt controller
- Multiple SPI flash interfaces
- NAND/Memory interface
- Sixteen mailbox registers for communication between the BMC and host
- LPC ROM interface
- BMC watchdog timer capability
- SD/MMC card controller with DMA support
- LED support with programmable blink rate controls on GPIOs
- Port 80h snooping capability
- Secondary Service Processor (SSP), which provides the HW capability of off-loading time critical processing tasks from the main ARM core.

### 3.12.1 Remote KVMS Support

- USB 2.0 interface for Keyboard, Mouse and Remote storage such as CD/DVD ROM and floppy
- USB 1.1/USB 2.0 interface for PS2 to USB bridging, remote Keyboard and Mouse
- Hardware Based Video Compression and Redirection Logic
- Supports both text and Graphics redirection
- Hardware assisted Video redirection using the Frame Processing Engine
- Direct interface to the Integrated Graphics Controller registers and Frame buffer
- Hardware-based encryption engine

### 3.12.2 Integrated BMC Embedded LAN Channel

The Integrated BMC hardware includes two dedicated 10/100 network interfaces. These interfaces are not shared with the host system. At any time, only one dedicated interface may be enabled for management traffic. The default active interface is the NIC 1 port.

For these channels, support can be enabled for IPMI-over-LAN and DHCP. For security reasons, embedded LAN channels have the following default settings:

- IP Address: Static.
- All users disabled.

# 4 System Security

## 4.1 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 8 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 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 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, since 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 Setup and change 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

## 4.2 Trusted Platform Module (TPM) Support

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” 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 TPM security BIOS

The BIOS TPM support conforms to the TPM PC Client Specific – Implementation Specification for Conventional BIOS, version 1.2, and to the TPM Interface specification, version 1.2. The BIOS adheres to the Microsoft Vista* BitLocker requirement. 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.

### 4.3.1 Physical Presence

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

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

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

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

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

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

### 4.3.3 Security Screen

To enter the BIOS Setup, press the F2 function key during boot time when the OEM or Intel® logo displays. The following message displays on the diagnostics screen and under the Quiet Boot logo screen:

Press <F2> to enter setup

When the Setup is entered, the Main screen displays. The BIOS Setup utility provides the Security screen to enable and set the user and administrative passwords and to lock out the front panel buttons so they cannot be used. The Intel® Server Board S2400GP provides TPM settings through the security screen.

To access this screen from the Main screen, select the **Security** option.

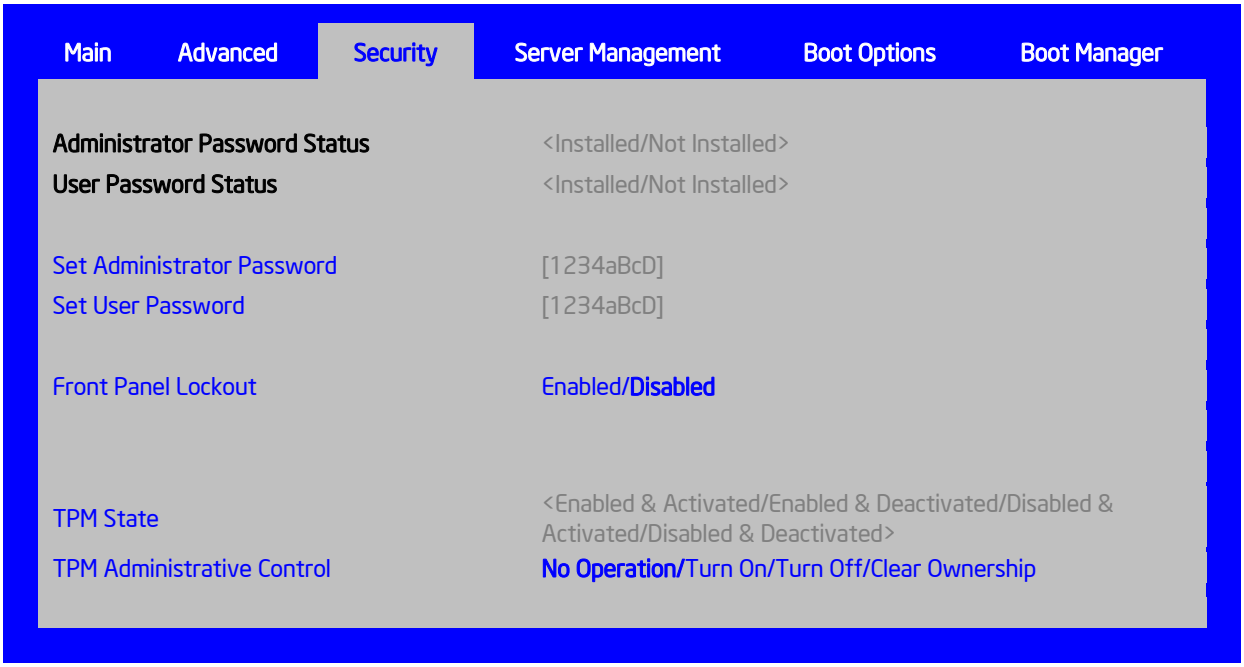


**Figure 20. Setup Utility – TPM Configuration Screen**

**Table 12. TPM Setup Utility – Security Configuration Screen Fields**

| Setup Item | Options | Help Text | Comments |
|---|---|---|---|
| TPM State* | Enabled and Activated<br>Enabled and Deactivated<br>Disabled and Activated<br>Disabled and Deactivated | | Information only.<br>Shows the current TPM device state.<br><br>A disabled TPM device will not execute commands that use TPM functions and TPM security operations will not be available.<br><br>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.<br><br>An enabled and activated TPM executes all commands that use TPM functions and TPM security operations will be available. |
| TPM Administrative Control** | No Operation<br>Turn On<br>Turn Off<br>Clear Ownership | [No Operation] - No changes to current state.<br>[Turn On] - Enables and activates TPM.<br>[Turn Off] - Disables and deactivates TPM.<br>[Clear Ownership] - Removes the TPM ownership authentication and returns the TPM to a factory default state.<br>**Note:** The BIOS setting returns to [No Operation] on every boot cycle by default. | |

## 4.4 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5-4600/2600/2400/1600 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 Specification*, Revision 1.2.

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

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

# 5 Technology Support

## 5.1 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5 4600/2600/2400/1600 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 and Intel® VT for Directed IO, with an active TPM, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications.

## 5.2 Intel® Virtualization Technology – Intel® VT-x/VT-d/VT-c

Intel® Virtualization Technology consists of three components which are integrated and interrelated, but which address different areas of Virtualization.

- **Intel® Virtualization Technology (VT-x)** is processor-related and provides capabilities needed to provide hardware assist to a Virtual Machine Monitor (VMM).
- **Intel® Virtualization Technology for Directed I/O (VT-d*)*** is primarily concerned with virtualizing I/O efficiently in a VMM environment. This would generally be a chipset I/O feature, but in the Second Generation Intel® Core™ Processor Family there is an Integrated I/O unit embedded in the processor, and the IIO is also enabled for VT-d.
- **Intel® Virtualization Technology for Connectivity (VT-c)** is primarily concerned I/O hardware assist features, complementary to but independent of VT-d.

Intel® VT-x is designed to support multiple software environments sharing same hardware resources. Each software environment may consist of OS and applications. The Intel® Virtualization Technology features can be enabled or disabled in the BIOS setup. The default behavior is disabled.

Intel® VT-d is supported jointly by the Intel® Xeon® Processor E5 4600/2600/2400/1600 Product Families and the C600 chipset. Both support DMA remapping from inbound PCI Express* memory Guest Physical Address (GPA) to Host Physical Address (HPA). PCI devices are directly assigned to a virtual machine leading to a robust and efficient virtualization.

The Intel® S4600/S2600/S2400/S1600/S1400 Server Board Family BIOS publishes the DMAR table in the ACPI Tables. For each DMA Remapping Engine in the platform, one exact entry of DRHD (DMA Remapping Hardware Unit Definition) structure is added to the DMAR. The DRHD structure in turn contains a Device Scope structure that describes the PCI endpoints and/or sub-hierarchies handled by the particular DMA Remapping Engine.

Similarly, there are reserved memory regions typically allocated by the BIOS at boot time. The BIOS marks these regions as either reserved or unavailable in the system address memory map reported to the OS. Some of these regions can be a target of DMA requests from one or more devices in the system, while the OS or executive is active. The BIOS reports each such memory region using exactly one RMRR (Reserved Memory Region Reporting) structure in the DMAR. Each RMRR has a Device Scope listing the devices in the system that can cause a DMA request to the region.

For more information on the DMAR table and the DRHD entry format, refer to the *Intel® Virtualization Technology for Directed I/O Architecture Specification*. For more general information about VT-x, VT-d, and VT-c, a good reference is *Enabling Intel® Virtualization Technology Features and Benefits White Paper*.

## 5.3 Intel® Intelligent Power Node Manager

Data centers are faced with power and cooling challenges that are driven by increasing numbers of servers deployed and server density in the face of several data center power and cooling constraints. In this type of environment, Information Technology (IT) needs the ability to monitor actual platform power consumption and control power allocation to servers and racks in order to solve specific data center problems including the following issues.

**Table 13. Intel® Intelligent Power Node Manager**

| IT Challenge | Requirement |
|---|---|
| Over-allocation of power | ▪ Ability to monitor actual power consumption<br>▪ Control capability that can maintain a power budget to enable dynamic power allocation to each server |
| Under-population of rack space | Control capability that can maintain a power budget to enable increased rack population. |
| High energy costs | Control capability that can maintain a power budget to ensure that a set energy cost can be achieved |
| Capacity planning | ▪ Ability to monitor actual power consumption to enable power usage modeling over time and a given planning period<br>▪ Ability to understand cooling demand from a temperature and airflow perspective |
| Detection and correction of hot spots | ▪ Control capability that reduces platform power consumption to protect a server in a hot-spot<br>▪ Ability to monitor server inlet temperatures to enable greater rack utilization in areas with adequate cooling. |

The requirements listed above are those that are addressed by the C600 chipset Management Engine (ME) and Intel® Intelligent Power Node Manager (NM) technology. The ME/NM combination is a power and thermal control capability on the platform, which exposes external interfaces that allow IT (through external management software) to query the ME about platform power capability and consumption, thermal characteristics, and specify policy directives (for example, set a platform power budget).

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

The NM feature is implemented by a complementary architecture utilizing the ME, BMC, BIOS, and an ACPI-compliant OS. The ME provides the NM policy engine and power control/limiting functions (referred to as Node Manager or NM) while the BMC provides the external LAN link by which external management software can interact with the feature. The BIOS provides system power information utilized by the NM algorithms and also exports ACPI Source Language (ASL)

code used by OS-Directed Power Management (OSPM) for negotiating processor P and T state changes for power limiting. PMBus*-compliant power supplies provide the capability to monitoring input power consumption, which is necessary to support NM.

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

- **Platform Power Monitoring and Limiting:** The ME/NM monitors platform power consumption and hold average power over duration. It can be queried to return actual power at any given instance. The power limiting capability is to allow external management software to address key IT issues by setting a power budget for each server. For example, if there is a physical limit on the power available in a room, then IT can decide to allocate power to different servers based on their usage – servers running critical systems can be allowed more power than servers that are running less critical workload.
- **Inlet Air Temperature Monitoring:** The ME/NM monitors server inlet air temperatures periodically. If there is an alert threshold in effect, then ME/NM issues an alert when the inlet (room) temperature exceeds the specified value. The threshold value can be set by policy.
- **Memory Subsystem Power Limiting:** The ME/NM monitors memory power consumption. Memory power consumption is estimated using average bandwidth utilization information
- **Processor Power monitoring and limiting:** The ME/NM monitors processor or socket power consumption and holds average power over duration. It can be queried to return actual power at any given instant. The monitoring process of the ME will be used to limit the processor power consumption through processor P-states and dynamic core allocation
- **Core allocation at boot time:** Restrict the number of cores for OS/VMM use by limiting how many cores are active at boot time. After the cores are turned off, the CPU will limit how many working cores are visible to BIOS and OS/VMM. The cores that are turned off cannot be turned on dynamically after the OS has started. It can be changed only at the next system reboot.
- **Core allocation at run-time:** This particular use case provides a higher level processor power control mechanism to a user at run-time, after booting. An external agent can dynamically use or not use cores in the processor subsystem by requesting ME/NM to control them, specifying the number of cores to use or not use.

### 5.3.1 Hardware Requirements

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

# 6 Platform Management Functional Overview

Platform management functionality is supported by several hardware and software components integrated on the server board that work together to control system functions, monitor and report system health, and control various thermal and performance features in order to maintain (when possible) server functionality in the event of component failure and/or environmentally stressed conditions.

This chapter provides a high level overview of the platform management features and functionality implemented on the server board. For more in depth and design level Platform Management information, please reference the *BMC Core Firmware External Product Specification (EPS)* and *BIOS Core External Product Specification (EPS)* for Intel® Server products based on the Intel® Xeon® processor E5-2400 product families.

## 6.1 Baseboard Management Controller (BMC) Firmware Feature Support

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.

### 6.1.1 IPMI 2.0 Features

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

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

### 6.1.2 Non IPMI Features

The BMC supports the following non-IPMI features.

- In-circuit BMC firmware update
- 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
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Enable/Disable of System Reset Due CPU Errors
- Chassis intrusion detection
- Fan speed control
- Fan redundancy monitoring and support
- Hot-swap fan support
- Power Supply Fan Sensors
- System Airflow Monitoring
- Exit Air Temperature Monitoring
- Acoustic management: Support for multiple fan profiles
- Ethernet Controller Thermal Monitoring
- Global Aggregate Temperature Margin Sensor
- Platform environment control interface (PECI) thermal management support
- Memory Thermal Management
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Power supply redundancy monitoring and support
- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.
- Intel® Intelligent Power Node Manager support
- 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.
- Local Control Display Panel support
- Power state retention
- Power fault analysis
- Intel® Light-Guided Diagnostics

- 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).
- E-mail alerting
- Embedded web server
    - Support for embedded web server UI in Basic Manageability feature set.
    - Human-readable SEL
    - Additional system configurability
    - Additional system monitoring capability
    - Enhanced on-line help
- Integrated KVM
- Integrated Remote Media Redirection
- Local Directory Access Protocol (LDAP) support
- Sensor and SEL logging additions/enhancements (for example, additional thermal monitoring capability)
- SEL Severity Tracking and the Extended SEL
- 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 1.1 compliance (product-specific).
- Management support for PMBus* rev1.2 compliant power supplies
- 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

### 6.1.3 New Manageability Features

Intel® S1400/S1600/S2400/S2600 Server Platforms offer a number of changes and additions to the manageability features that are supported on the previous generation of servers. The following is a list of the more significant changes that are common to this generation Integrated BMC based Intel® Server boards:

- Sensor and SEL logging additions/enhancements (for example, additional thermal monitoring capability)
- SEL Severity Tracking and the Extended SEL
- 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)

- Enhancements to fan speed control.
- DCMI 1.1 compliance (product-specific).
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server
  - Human-readable SEL
  - Additional system configurability
  - Additional system monitoring capability
  - Enhanced on-line help
- Enhancements to KVM redirection
  - Support for higher resolution
- Support for EU Lot6 compliance
- Management support for PMBus* rev1.2 compliant power supplies
- BMC Data Repository (Managed Data Region Feature)
- 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

## 6.2 Basic and Advanced Features

The bellowing table lists basic and advanced feature support. Individual features may vary by platform. See the appropriate Platform Specific EPS addendum for more information.

**Table 14. Basic and Advanced Features**

| 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 |

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

## 6.3 Integrated BMC Hardware: Emulex* Pilot III

### 6.3.1 Emulex* Pilot III Baseboard Management Controller Functionality

The Integrated BMC is provided by an embedded ARM9 controller and associated peripheral functionality that is required for IPMI-based server management. Firmware usage of these hardware features is platform dependent.

The following is a summary of the Integrated BMC management hardware features that comprise the BMC:

- 400MHz 32-bit ARM9 processor with memory management unit (MMU)
- Two independent10/100/1000 Ethernet Controllers with Reduced Media Independent Interface (RMII)/Reduced Gigabit Media Independent Interface (RGMII) support
- DDR2/3 16-bit interface with up to 800 MHz operation
- 16 10-bit ADCs
- Sixteen fan tachometers
- Eight Pulse Width Modulators (PWM)
- Chassis intrusion logic
- JTAG Master
- Eight $I^2C$ interfaces with master-slave and SMBus* timeout support. All interfaces are SMBus* 2.0 compliant.
- Parallel general-purpose I/O Ports (16 direct, 32 shared)
- Serial general-purpose I/O Ports (80 in and 80 out)
- Three UARTs
- Platform Environmental Control Interface (PECI)
- Six general-purpose timers

- Interrupt controller
- Multiple Serial Peripheral Interface (SPI) flash interfaces
- NAND/Memory interface
- Sixteen mailbox registers for communication between the BMC and host
- LPC ROM interface
- BMC watchdog timer capability
- SD/MMC card controller with DMA support
- LED support with programmable blink rate controls on GPIOs
- Port 80h snooping capability
- Secondary Service Processor (SSP), which provides the HW capability of offloading time critical processing tasks from the main ARM core.

Emulex* Pilot III contains an integrated SIO, KVMS subsystem and graphics controller with the following features:

## 6.4 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states:

**Table 15. ACPI Power States**

| State | Supported | Description |
|---|---|---|
| S0 | Yes | Working.<br>▪ The front panel power LED is on (not controlled by the BMC).<br>▪ The fans spin at the normal speed, as determined by sensor inputs.<br>▪ Front panel buttons work normally. |
| S1 | Yes | Sleeping. Hardware context is maintained; equates to processor and chipset clocks being stopped.<br>▪ The front panel power LED blinks at a rate of 1 Hz with a 50% duty cycle (not controlled by the BMC).<br>▪ The watchdog timer is stopped.<br>▪ The power, reset, front panel NMI, and ID buttons are unprotected.<br>▪ Fan speed control is determined by available SDRs. Fans may be set to a fixed state, or basic fan management can be applied.<br>The BMC detects that the system has exited the ACPI S1 sleep state when the BIOS SMI handler notifies it. |
| S2 | No | Not supported. |
| S3 | No | Supported only on Workstation platforms. See appropriate Platform Specific Information for more information. |
| S4 | No | Not supported. |
| S5 | Yes | Soft off.<br>▪ The front panel buttons are not locked.<br>▪ The fans are stopped.<br>▪ The power-up process goes through the normal boot process.<br>▪ The power, reset, front panel NMI, and ID buttons are unlocked. |

## 6.5 Power Control Sources

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

**Table 16. Power Control Initiators**

| 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 |
| Command | 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 | CPU Thermtrip | Turns power off |
| WOL(Wake On LAN) | LAN | Turns power on |

## 6.6 BMC Watchdog

The BMC FW is increasingly called upon to perform system functions that are time-critical in that failure to provide these functions in a timely manner can result in system or component damage. Intel® S1400/S1600/S2400/S2600/S4600 Server Platforms introduce a BMC watchdog feature to provide a safe-guard against this scenario by providing an automatic recovery mechanism. It also can provide automatic recovery of functionality that has failed due to a fatal FW defect triggered by a rare sequence of events or a BMC hang due to some type of HW glitch (for example, power).

This feature is comprised of a set of capabilities whose purpose is to detect misbehaving subsections of BMC firmware, the BMC CPU itself, or HW subsystems of the BMC component, and to take appropriate action to restore proper operation. The action taken is dependent on the nature of the detected failure and may result in a restart of the BMC CPU, one or more BMC HW subsystems, or a restart of malfunctioning FW subsystems.

The BMC watchdog feature will only allow up to three resets of the BMC CPU (such as HW reset) or entire FW stack (such as a SW reset) before giving up and remaining in the uBOOT code. This count is cleared upon cycling of power to the BMC or upon continuous operation of the BMC without a watchdog-generated reset occurring for a period of > 30 minutes. The BMC FW logs a SEL event indicating that a watchdog-generated BMC reset (either soft or hard reset) has occurred. This event may be logged after the actual reset has occurred. Refer sensor section for details for the related sensor definition. The BMC will also indicate a degraded system status on the Front Panel Status LED after an BMC HW reset or FW stack reset. This state (which follows the state of the associated sensor) will be cleared upon system reset or (AC or DC) power cycle.

**Note:** A reset of the BMC may result in the following system degradations that will require a system reset or power cycle to correct:

1. Timeout value for the rotation period can be set using this parameter. Potentially, there will be incorrect ACPI Power State reported by the BMC.
2. Reversion of temporary test modes for the BMC back to normal operational modes.
3. FP status LED and DIMM fault LEDs may not reflect BIOS detected errors.

## 6.7 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.

## 6.8 Sensor Monitoring

The BMC monitors system hardware and reports system health. Some of the sensors include those for monitoring.

- Component, board, and platform temperatures
- Board and platform voltages
- System fan presence and tach
- Chassis intrusion
- Front Panel NMI
- Front Panel Power and System Reset Buttons
- SMI timeout

- Processor errors

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.

See *Appendix B – Integrated BMC Sensor Tables* for additional sensor information.

## 6.9 Field Replaceable Unit (FRU) Inventory Device

The BMC implements the interface for logical FRU inventory devices as specified in the *Intelligent Platform Management Interface Specification, Version 2.0.* This functionality provides commands used for accessing and managing the FRU inventory information. These commands can be delivered through all interfaces.

The BMC provides FRU device command access to its own FRU device and to the FRU devices throughout the server. The FRU device ID mapping is defined in the Platform Specific Information. The BMC controls the mapping of the FRU device ID to the physical device

## 6.10 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 65,502 bytes (approximately 64 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. Any command that results in an overflow of the SEL beyond the allocated space is rejected with an "Out of Space" IPMI completion code (C4h).

Events logged to the SEL can be viewed using Intel®'s SELVIEW utility, Embedded Web Server, and Active System Console.

## 6.11 System 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: sleep, nominal, and boost. 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.

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

### 6.11.1 Thermal and Acoustic Management

The S2400GP offers multiple thermal and acoustic management features to maintain comprehensive thermal protection as well as intelligent fan speed control. The features can be adjusted in BIOS interface with path BIOS > Advanced > System Acoustic and Performance Configuration.
This feature refers to enhanced fan management to keep the system optimally cooled while reducing the amount of noise generated by the system fans. Aggressive acoustics standards might require a trade-off between fan speed and system performance parameters that contribute to the cooling requirements, primarily memory bandwidth. The BIOS, BMC, and SDRs work together to provide control over how this trade-off is determined.
This capability requires the BMC to access temperature sensors on the individual memory DIMMs. Additionally, closed-loop thermal throttling is only supported with buffered DIMMs.

### 6.11.2 Setting Throttling Mode

Select the most appropriate memory thermal throttling mechanism for memory sub-system from [Auto], [DCLTT], [SCLTT] and [SOLTT].

[Auto] – BIOS automatically detect and identify the appropriate thermal throttling mechanism based on DIMM type, airflow input, DIMM sensor availability.
[DCLTT] – Dynamic Closed Loop Thermal Throttling: for the SOD DIMM with system airflow input
[SCLTT] – Static Close Loop Thermal Throttling: for the SOD DIMM without system airflow input
[SOLTT] – Static Open Loop Thermal Throttling: for the DIMMs without sensor on DIMM (SOD)

The default setting is [Auto]

### 6.11.3 Altitude

Select the proper altitude that the system is distributed from [300m or less], [301m-900m], [901m-1500m], [Above 1500m] options. Lower altitude selection can lead to potential thermal risk. And higher altitude selection provides better cooling but with undesired acoustic and fan power consumption. If the altitude is known, higher altitude is recommended in order to provide sufficient cooling. The default setting is [301m – 900m].

### 6.11.4 Set Fan Profile

[Performance] and [Acoustic] fan profiles are available to select. The Acoustic mode offers best acoustic experience and appropriate cooling capability covering mainstream and majority of the add-in cards. Performance mode is designed to provide sufficient cooling capability covering all kinds of add-in cards on the market. The default setting is [Performance]

### 6.11.5 Fan PWM Offset

This feature is reserved for manual adjustment to the minimum fan speed curves. The valid range is from [0 to 100] which stands for 0% to 100% PWM adding to the minimum fan speed. This feature is valid when Quiet Fan Idle Mode is at Enabled state. The default setting is [0]

### 6.11.6 Quiet Fan Idle Mode

This feature can be [Enabled] or [Disabled]. If enabled, the fan will either stopped or shift to a lower speed when the aggregate sensor temperatures are satisfied indicating the system is at ideal thermal/light loading conditions. When the aggregate sensor temperatures not satisfied, the fan will shift back to normal control curves. If disabled, the fan will never stopped or shift into lower fan speed whatever the aggregate sensor temperatures are satisfied or not. The default setting is [Disabled]

**Note**:
1. The above features may or may not be in effective depends on the actual thermal characters of a specific system.
2. Refer to the Intel server system TPS for the board in Intel chassis thermal and acoustic management
3. Refer to Fan Control Whitepaper for the board in 3rd party chassis fan speed control customization

### 6.11.7 Fan Profiles

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

The BMC supports eight fan profiles, numbered from 0 to 7.

**Table 17. Fan Profiles**

| Type | Profile | Details |
|---|---|---|
| OLTT | 0 | Acoustic, 300M altitude |
| OLTT | 1 | Performance, 300M altitude |
| OLTT | 2 | Acoustic, 900M altitude |
| OLTT | 3 | Performance, 900M altitude |
| OLTT | 4 | Acoustic, 1500M altitude |
| OLTT | 5 | Performance, 1500M altitude |
| OLTT | 6 | Acoustic, 3000M altitude |
| OLTT | 7 | Performance, 3000M altitude |
| CLTT | 0 | 300M altitude |
| CLTT | 2 | 900M altitude |
| CLTT | 4 | 1500M altitude |
| CLTT | 6 | 3000M altitude |

Each group of profiles allows for varying fan control policies based on the altitude. For a given altitude, the Tcontrol SDRs associated with an acoustics-optimized profile generate less noise than the equivalent performance-optimized profile by driving lower fan speeds, and the BIOS reduces thermal management requirements by configuring more aggressive memory throttling.

The BMC only supports enabling a fan profile through the command if that profile is supported on all fan domains defined for the given system. **It is important to configure platform Sensor Data Records (SDRs) so that all desired fan profiles are supported on each fan domain**. If no single profile is supported across all domains, the BMC, by default, uses profile 0 and does not allow it to be changed.

### 6.11.8 Thermal Sensor Input to Fan Speed Control

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

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

- 
- Front Panel Temperature Sensor1
- Baseboard Temperature Sensor2
- CPU Margin Sensors3,5,6
- DIMM Thermal Margin Sensors3,5
- Exit Air Temperature Sensor1, 4, 8
- PCH Temperature Sensor4,6
- On-board Ethernet Controller Temperature Sensors4, 6
- Add-In Intel SAS/IO Module Temperature Sensors4, 6
- PSU Thermal Sensor4, 9
- CPU VR Temperature Sensors4, 7
- DIMM VR Temperature Sensors4, 7
- iBMC Temperature Sensor4, 7
- Global Aggregate Thermal Margin Sensors3, 8

**Note**:

1. For fan speed control in Intel chassis
2. For fan speed control in 3rd party chassis
3. Temperature margin from throttling threshold
4. Absolute temperature
5. PECI value
6. On-die sensor
7. On-board sensor
8. Virtual sensor
9. Available only when PSU has PMBus

The following illustration provides a simple model showing the fan speed control structure that implements the resulting fan speeds.

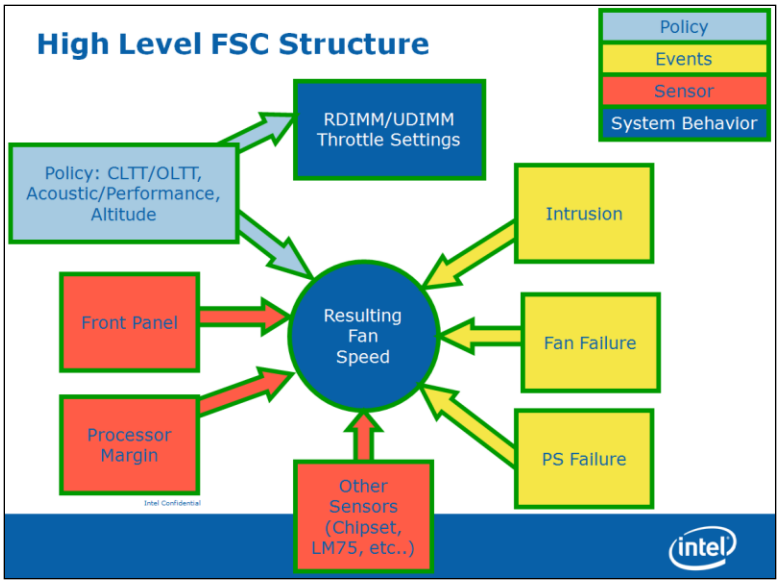


**Figure 21. Fan Speed Control Process**

### 6.11.9 Memory Thermal Throttling

The server board provides support for system thermal management through open loop throttling (OLTT) and closed loop throttling (CLTT) of system memory. Normal system operation uses closed-loop thermal throttling (CLTT) and DIMM temperature monitoring as major factors in overall thermal and acoustics management. In the event that BIOS is unable to configure the system for CLTT, it defaults to open-loop thermal throttling (OLTT). In the OLTT mode, it is assumed that the DIMM temperature sensors are not available for fan speed control. 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's fan speed control functionality is linked to the memory throttling mechanism used.

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

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

Both Static and Dynamic CLTT modes implement a Hybrid Closed Loop Thermal Throttling mechanism whereby the Integrated Memory Controller estimates the DRAM temperature in between actual reads of the memory thermal sensors.

## 6.12 Messaging Interfaces

The BMC supports the following communications interfaces:

- Host SMS interface by means of low pin count (LPC)/keyboard controller style (KCS) interface
- Host SMM interface by means of low pin count (LPC)/keyboard controller style (KCS) interface
- Intelligent Platform Management Bus (IPMB) I2C interface
- LAN interface using the IPMI-over-LAN protocols

Every messaging interface is assigned an IPMI channel ID by IPMI 2.0.

**Table 18. Messaging Interfaces**

| Channel ID | Interface | Supports Sessions |
|---|---|---|
| 0 | Primary IPMB | No |
| 1 | LAN 1 | Yes |
| 2 | LAN 2 | Yes |
| 3 | LAN3 [1] (Provided by the Intel® Dedicated Server Management NIC) | Yes |
| 4 | Reserved | Yes |
| 5 | USB | No |
| 6 | Secondary IPMB | No |
| 7 | SMM | No |
| 8– 0Dh | Reserved | – |
| 0Eh | Self [2] | – |
| 0Fh | SMS/Receive Message Queue | No |

**Notes:**

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

### 6.12.1 User Model

The BMC supports the IPMI 2.0 user model. 15 user IDs are supported. These 15 users can be assigned to any channel. The following restrictions are placed on user-related operations:

1. User names for User IDs 1 and 2 cannot be changed. These are always “” (Null/blank) and “root” respectively.
2. User 2 (“root”) always has the administrator privilege level.
3. All user passwords (including passwords for 1 and 2) may be modified.

User IDs 3-15 may be used freely, with the condition that user names are unique. Therefore, no other users can be named “” (Null), “root,” or any other existing user name.

### 6.12.2 IPMB Communication Interface

The IPMB communication interface uses the 100 KB/s version of an $I^2C$ bus as its physical medium. For more information on $I^2C$ specifications, see *The $I^2C$ Bus and How to Use It.* The IPMB implementation in the BMC is compliant with the *IPMB v1.0, revision 1.0.*

The BMC IPMB slave address is 20h.

The BMC both sends and receives IPMB messages over the IPMB interface. Non-IPMB messages received by means of the IPMB interface are discarded.

Messages sent by the BMC can either be originated by the BMC, such as initialization agent operation, or by another source. One example is KCS-IPMB bridging.

## 6.12.3 LAN Interface

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

See the *Intelligent Platform Management Interface Specification Second Generation v2.0* for details about the IPMI-over-LAN protocol.

Run-time determination of LAN channel capabilities can be determined by both standard IPMI defined mechanisms.

### 6.12.3.1 RMCP/ASF Messaging

The BMC supports RMCP ping discovery in which the BMC responds with a pong message to an RMCP/ASF ping request. This is implemented per the *Intelligent Platform Management Interface Specification Second Generation v2.0*.

### 6.12.3.2 BMC LAN Channels

The BMC supports three RMII/RGMII ports that can be used for communicating with Ethernet devices. Two ports are used for communication with the on-board NICs and one is used for communication with an Ethernet PHY located on an optional RMM4 add-in module.

#### *6.12.3.2.1 Baseboard NICs*

The on-board Ethernet controller provides support for a Network Controller Sideband Interface (NC-SI) manageability interface. This provides a sideband high-speed connection for manageability traffic to the BMC while still allowing for a simultaneous host access to the OS if desired.

The NC-SI is a DMTF industry standard protocol for the side band management LAN interface. This protocol provides a fast multi-drop interface for management traffic.

The baseboard NIC(s) are connected to a single BMC RMII/RGMII port that is configured for RMII operation. The NC-SI protocol is used for this connection and provides a 100 Mb/s full-duplex multi-drop interface which allows multiple NICs to be connected to the BMC. The physical layer is based upon RMII, however RMII is a point-to-point bus whereas NC-SI allows 1 master and up to 4 slaves. The logical layer (configuration commands) is incompatible with RMII.

The server board will provide support for a dedicated management channel that can be configured to be hidden from the host and only used by the BMC. This mode of operation is configured through a BIOS setup option.

#### 6.12.3.2.2 Dedicated Management Channel

An additional LAN channel dedicated to BMC usage and not available to host SW is supported through an optional RMM4 add-in card. There is only a PHY device present on the RMM4 add-in card. The BMC has a built-in MAC module that uses the RGMII interface to link with the card's PHY. Therefore, for this dedicated management interface, the PHY and MAC are located in different devices.

The PHY on the RMM4 connects to the BMC's other RMII/RGMII interface (that is, the one that is not connected to the baseboard NICs). This BMC port is configured for RGMII usage.
In addition to the use of an RMM4 add-in card for a dedicated management channel, on systems that support multiple Ethernet ports on the baseboard, the system BIOS provides a setup option to allow one of these baseboard ports to be dedicated to the BMC for manageability purposes. When this is enabled, that port is hidden from the OS.

#### 6.12.3.2.3 Concurrent Server Management Use of Multiple Ethernet Controllers

The BMC FW supports concurrent OOB LAN management sessions for the following combination:

- Two on-board NIC ports
- One on-board NIC and the optional dedicated RMM4 add-in management NIC.
- Two on-board NICs and optional dedicated RMM4 add-in management NIC.

All NIC ports must be on different subnets for the above concurrent usage models.
MAC addresses are assigned for management NICs from a pool of up to three MAC addresses allocated specifically for manageability. The total number of MAC addresses in the pool is dependent on the product HW constraints (for example, a board with two NIC ports available for manageability would have a MAC allocation pool of 2 addresses). For these channels, support can be enabled for IPMI-over-LAN and DHCP.

For security reasons, embedded LAN channels have the following default settings:

- IP Address: Static
- All users disabled

IPMI-enabled network interfaces may not be placed on the same subnet. This includes the Intel® Dedicated Server Management NIC and either of the BMC's embedded network interfaces.

Host-BMC communication over the same physical LAN connection – also known as "loopback" – is not supported. This includes "ping" operations.

On server boards with more than two onboard NIC ports, only the first two ports can be used as BMC LAN channels. The remaining ports have no BMC connectivity.

Maximum bandwidth supported by BMC LAN channels are as follows:

- BMC LAN1 (Baseboard NIC port) ----- 100Mb (10Mb in DC off state)
- BMC LAN 2 (Baseboard NIC port) ----- 100Mb (10Mb in DC off state)

- BMC LAN 3 (Dedicated NIC) ----- 100Mb

#### 6.12.3.3 IPV6 Support

In addition to IPv4, the server board has support for IPv6 for manageability channels. Configuration of IPv6 is provided by extensions to the IPMI Set and Get LAN Configuration Parameters commands as well as through a Web Console IPv6 configuration web page.

The BMC supports IPv4 and IPv6 simultaneously so they are both configured separately and completely independently. For example, IPv4 can be DHCP configured while IPv6 is statically configured or vice versa.

The parameters for IPv6 are similar to the parameters for IPv4 with the following differences:

- An IPv6 address is 16 bytes vs. 4 bytes for IPv4.
- An IPv6 prefix is 0 to 128 bits whereas IPv4 has a 4 byte subnet mask.
- The IPv6 Enable parameter must be set before any IPv6 packets will be sent or received on that channel.
- There are two variants of automatic IP Address Source configuration vs. just DHCP for IPv4.

The three possible IPv6 IP Address Sources for configuring the BMC are:

**Static (Manual)**: The IP, Prefix, and Gateway parameters are manually configured by the user. The BMC ignores any Router Advertisement messages received over the network.

**DHCPv6**: The IP comes from running a DHCPv6 client on the BMC and receiving the IP from a DHCPv6 server somewhere on the network. The Prefix and Gateway are configured by Router Advertisements from the local router. The IP, Prefix, and Gateway are read-only parameters to the BMC user in this mode.

**Stateless auto-config**: The Prefix and Gateway are configured by the router through Router Advertisements. The BMC derives its IP in two parts: the upper network portion comes from the router and the lower unique portion comes from the BMC's channel MAC address. The 6-byte MAC address is converted into an 8-byte value per the EUI-64* standard. For example, a MAC value of 00:15:17:FE:2F:62 converts into a EUI-64 value of 215:17ff:fefe:2f62. If the BMC receives a Router Advertisement from a router at IP 1:2:3:4::1 with a prefix of 64, it would then generate for itself an IP of 1:2:3:4:215:17ff:fefe:2f62. The IP, Prefix, and Gateway are read-only parameters to the BMC user in this mode.

IPv6 can be used with the BMC's Web Console, JViewer (remote KVM and Media), and Systems Management Architecture for Server Hardware – Command Line Protocol (SMASH-CLP) interface (ssh). There is no standard yet on how IPMI RMCP or RMCP+ should operate over IPv6 so that is not currently supported.

#### 6.12.3.4 LAN Failover

The BMC FW provides a LAN failover capability such that the failure of the system HW associated with one LAN link will result in traffic being rerouted to an alternate link. This

functionality is configurable through IPMI methods as well as through the BMC's Embedded UI, allowing for user to specify the physical LAN links constitute the redundant network paths or physical LAN links constitute different network paths. BMC will support only a all or nothing" approach – that is, all interfaces bonded together, or none are bonded together.

The LAN Failover feature applies only to BMC LAN traffic. It bonds all available Ethernet devices but only one is active at a time. When enabled, If the active connection's leash is lost, one of the secondary connections is automatically configured so that it has the same IP address. Traffic immediately resumes on the new active connection.

The LAN Failover enable/disable command may be sent at any time. After it has been enabled, standard IPMI commands for setting channel configuration that specify a LAN channel other than the first will return an error code.

#### 6.12.3.5 BMC IP Address Configuration

Enabling the BMC's network interfaces requires using the *Set LAN Configuration Parameter* command to configure LAN configuration parameter 4, *IP Address Source*. The BMC supports this parameter as follows:

- 1h, static address (manually configured): Supported on all management NICs. This is the BMC's default value.
- 2h, address obtained by BMC running DHCP: Supported only on embedded management NICs.

IP Address Source value 4h, address obtained by BMC running other address assignment protocol, is not supported on any management NIC.

Attempting to set an unsupported IP address source value has no effect, and the BMC returns error code 0xCC, Invalid data field-in request. Note that values 0h and 3h are no longer supported, and will return a 0xCC error completion code.

##### *6.12.3.5.1 Static IP Address (IP Address Source Values 0h, 1h, and 3h)*

The BMC supports static IP address assignment on all of its management NICs. The IP address source parameter must be set to "static" before the IP address; the subnet mask or gateway address can be manually set.

The BMC takes no special action when the following IP address source is specified as the IP address source for any management NIC:1h – Static address (manually configured)

The *Set LAN Configuration Parameter* command must be used to configure LAN configuration parameter 3, *IP Address*, with an appropriate value.

The BIOS does not monitor the value of this parameter, and it does not execute DHCP for the BMC under any circumstances, regardless of the BMC configuration.

#### 6.12.3.5.2 Static LAN Configuration Parameters

When the IP Address Configuration parameter is set to 01h (static), the following parameters may be changed by the user:

- LAN configuration parameter 3 (IP Address)
- LAN configuration parameter 6 (Subnet Mask)
- LAN configuration parameter 12 (Default Gateway Address)

When changing from DHCP to Static configuration, the initial values of these three parameters will be equivalent to the existing DHCP-set parameters. Additionally, the BMC observes the following network safety precautions:

1. The user may only set a subnet mask that is valid, per IPv4 and RFC 950 (*Internet Standard Subnetting Procedure*). Invalid subnet values return a 0xCC (Invalid Data Field in Request) completion code, and the subnet mask is not set. If no valid mask has been previously set, default subnet mask is 0.0.0.0.
2. The user may only set a default gateway address that can potentially exist within the subnet specified above. Default gateway addresses outside the BMC's subnet are technically unreachable and the BMC will not set the default gateway address to an unreachable value. The BMC returns a 0xCC (Invalid Data Field in Request) completion code for default gateway addresses outside its subnet.
3. If a command is issued to set the default gateway IP address before the BMC's IP address and subnet mask are set, the default gateway IP address is not updated and the BMC returns 0xCC.

If the BMC's IP address on a LAN channel changes while a LAN session is in progress over that channel, the BMC does not take action to close the session except through a normal session timeout. The remote client must re-sync with the new IP address. The BMC's new IP address is only available in-band through the "Get LAN Configuration Parameters" command.

#### 6.12.3.5.3 Enabling/Disabling Dynamic Host Configuration (DHCP) Protocol

The BMC DHCP feature is activated by using the *Set LAN Configuration Parameter* command to set LAN configuration parameter 4, *IP Address Source*, to 2h: "address obtained by BMC running DHCP". Once this parameter is set, the BMC initiates the DHCP process within approximately 100 ms.

If the BMC has previously been assigned an IP address through DHCP or the *Set LAN Configuration Parameter* command, it requests that same IP address to be reassigned. If the BMC does not receive the same IP address, system management software must be reconfigured to use the new IP address. The new address is only available in-band, through the IPMI *Get LAN Configuration Parameters* command.

Changing the *IP Address Source* parameter from 2h to any other supported value will cause the BMC to stop the DHCP process. The BMC uses the most recently obtained IP address until it is reconfigured.

If the physical LAN connection is lost (that is, the cable is unplugged), the BMC will not re-initiate the DHCP process when the connection is re-established.

#### *6.12.3.5.4 DHCP-related LAN Configuration Parameters*

Users may not change the following LAN parameters while the DHCP is enabled:

- LAN configuration parameter 3 (IP Address)
- LAN configuration parameter 6 (Subnet Mask)
- LAN configuration parameter 12 (Default Gateway Address)

To prevent users from disrupting the BMC's LAN configuration, the BMC treats these parameters as read-only while DHCP is enabled for the associated LAN channel. Using the *Set LAN Configuration Parameter* command to attempt to change one of these parameters under such circumstances has no effect, and the BMC returns error code 0xD5, "Cannot Execute Command. Command, or request parameter(s) are not supported in present state."

### 6.12.3.6 DHCP BMC Hostname

The BMC allows setting a DHCP Hostname using the *Set/Get LAN Configuration Parameters* command.

- DHCP Hostname can be set regardless of the IP Address source configured on the BMC. But this parameter is only used if the IP Address source is set to DHCP.
- When Byte 2 is set to "Update in progress", all the 16 Block Data Bytes (Bytes 3 – 18) must be present in the request.
- When Block Size < 16, it must be the last Block request in this series. In other words Byte 2 is equal to "Update is complete" on that request.
- Whenever Block Size < 16, the Block data bytes must end with a NULL Character or Byte (=0).
- All Block write requests are updated into a local Memory byte array. When Byte 2 is set to "Update is Complete", the Local Memory is committed to the NV Storage. Local Memory is reset to NULL after changes are committed.
- When Byte 1 (Block Selector = 1), firmware resets all the 64 bytes local memory. This can be used to undo any changes after the last "Update in Progress".
- User should always set the hostname starting from block selector 1 after the last "Update is complete". If the user skips block selector 1 while setting the hostname, the BMC will record the hostname as "NULL," because the first block contains NULL data.
- This scheme effectively does not allow a user to make a partial Hostname change. Any Hostname change needs to start from Block 1.
- Byte 64 ( Block Selector 04h byte 16) is always ignored and set to NULL by BMC which effectively means we can set only 63 bytes.
- User is responsible for keeping track of the Set series of commands and Local Memory contents.

While BMC firmware is in "Set Hostname in Progress" (Update not complete), the firmware continues using the Previous Hostname for DHCP purposes.

### 6.12.4 Address Resolution Protocol (ARP)

The BMC can receive and respond to ARP requests on BMC NICs. Gratuitous ARPs are supported, and disabled by default.

### 6.12.5 Internet Control Message Protocol (ICMP)

The BMC supports the following ICMP message types targeting the BMC over integrated NICs:

- Echo request (ping): The BMC sends an Echo Reply.
- Destination unreachable: If message is associated with an active socket connection within the BMC, the BMC closes the socket.

### 6.12.6 Virtual Local Area Network (VLAN)

The BMC supports VLAN as defined by IPMI 2.0 specifications. VLAN is supported internally by the BMC, not through switches. VLAN provides a way of grouping a set of systems together so that they form a logical network. This feature can be used to set up a management VLAN where only devices which are members of the VLAN will receive packets related to management and members of the VLAN will be isolated from any other network traffic. Please note that VLAN does not change the behavior of the host network setting, it only affects the BMC LAN communication.

LAN configuration options are now supported (by means of the Set LAN Config Parameters command, parameters 20 and 21) that allow support for 802.1Q VLAN (Layer 2). This allows VLAN headers/packets to be used for IPMI LAN sessions. VLAN ID's are entered and enabled by means of parameter 20 of the Set LAN Config Parameters IPMI command. When a VLAN ID is configured and enabled, the BMC only accepts packets with that VLAN tag/ID. Conversely, all BMC generated LAN packets on the channel include the given VLAN tag/ID. Valid VLAN ID's are 1 through 4094, VLAN ID's of 0 and 4095 are reserved, per the 802.1Q VLAN specification. Only one VLAN can be enabled at any point in time on a LAN channel. If an existing VLAN is enabled, it must first be disabled prior to configuring a new VLAN on the same LAN channel.

Parameter 21 (VLAN Priority) of the Set LAN Config Parameters IPMI command is now implemented and a range from 0-7 will be allowed for VLAN Priorities. Please note that bits 3 and 4 of Parameter 21 are considered reserved bits.

Parameter 25 (VLAN Destination Address) of the Set LAN Config Parameters IPMI command is not supported and returns a completion code of 0x80 (parameter not supported) for any read/write of parameter 25.

If the BMC IP address source is DHCP, then the following behavior is seen:

- If the BMC is first configured for DHCP (prior to enabling VLAN), when VLAN is enabled, the BMC performs a discovery on the new VLAN in order to obtain a new BMC IP address.
- If the BMC is configured for DHCP (before disabling VLAN), when VLAN is disabled, the BMC performs a discovery on the LAN in order to obtain a new BMC IP address.

If the BMC IP address source is Static, then the following behavior is seen:

- If the BMC is first configured for static (prior to enabling VLAN), when VLAN is enabled, the BMC has the same IP address that was configured before. It is left to the management application to configure a different IP address if that is not suitable for VLAN.

- If the BMC is configure for static (prior to disabling VLAN), when VLAN is disabled, the BMC has the same IP address that was configured before. It is left to the management application to configure a different IP address if that is not suitable for LAN.

### 6.12.7 Secure Shell (SSH)

Secure Shell (SSH) connections are supported for SMASH-CLP sessions to the BMC.

### 6.12.8 Serial-over-LAN (SOL 2.0)

The BMC supports IPMI 2.0 SOL.

IPMI 2.0 introduced a standard serial-over-LAN feature. This is implemented as a standard payload type (01h) over RMCP+.

Three commands are implemented for SOL 2.0 configuration.

- "Get SOL 2.0 Configuration Parameters" and "Set SOL 2.0 Configuration Parameters": These commands are used to get and set the values of the SOL configuration parameters. The parameters are implemented on a per-channel basis.
- "Activating SOL": This command is not accepted by the BMC. It is sent by the BMC when SOL is activated to notify a remote client of the switch to SOL.
- Activating a SOL session requires an existing IPMI-over-LAN session. If encryption is used, it should be negotiated when the IPMI-over LAN session is established.

### 6.12.9 Platform Event Filter (PEF)

The BMC includes the ability to generate a selectable action, such as a system power-off or reset, when a match occurs to one of a configurable set of events. This capability is called *Platform Event Filtering,* or PEF. One of the available PEF actions is to trigger the BMC to send a LAN alert to one or more destinations.

The BMC supports 20 PEF filters. The first twelve entries in the PEF filter table are pre-configured (but may be changed by the user). The remaining entries are left blank, and may be configured by the user.

**Table 19. Factory Configured PEF Table Entries**

| Event Filter Number | Offset Mask | Events |
|---|---|---|
| 1 | Non-critical, critical and non-recoverable | Temperature sensor out of range |
| 2 | Non-critical, critical and non-recoverable | Voltage sensor out of range |
| 3 | Non-critical, critical and non-recoverable | Fan failure |
| 4 | General chassis intrusion | Chassis intrusion (security violation) |
| 5 | Failure and predictive failure | Power supply failure |
| 6 | Uncorrectable ECC | BIOS |
| 7 | POST error | BIOS: POST code error |
| 8 | FRB2 | Watchdog Timer expiration for FRB2 |
| 9 | Policy Correction Time | Node Manager |
| 10 | Power down, power cycle, and reset | Watchdog timer |
| 11 | OEM system boot event | System restart (reboot) |
| 12 | Drive Failure, Predicted Failure | Hot Swap Controller |

Additionally, the BMC supports the following PEF actions:

- Power off
- Power cycle
- Reset
- OEM action
- Alerts

The “Diagnostic interrupt” action is not supported.

### 6.12.10 LAN Alerting

The BMC supports sending embedded LAN alerts, called SNMP PET (Platform Event traps), and SMTP email alerts.

The BMC supports a minimum of four LAN alert destinations.

#### 6.12.10.1 SNMP Platform Event Traps (PETs)

This feature enables a target system to send SNMP traps to a designated IP address by means of LAN. These alerts are formatted per the *Intelligent Platform Management Interface Specification Second Generation v2.0*. A Modular Information Block (MIB) file associated with the traps is provided with the BMC firmware to facilitate interpretation of the traps by external software. The format of the MIB file is covered under RFC 2578.

### 6.12.11 Alert Policy Table

Associated with each PEF entry is an alert policy that determines which IPMI channel the alert is to be sent. There is a maximum of 20 alert policy entries. There are no pre-configured entries in the alert policy table because the destination types and alerts may vary by user. Each entry in the alert policy table contains four bytes for a maximum table size of 80 bytes.

#### 6.12.11.1 E-mail Alerting

The Embedded Email Alerting feature allows the user to receive e-mails alerts indicating issues with the server. This allows e-mail alerting in an OS-absent (for example, Pre-OS and OS-Hung) situation. This feature provides support for sending e-mail by means of SMTP, the Simple Mail Transport Protocol as defined in Internet RC 821. The e-mail alert provides a text string that describes a simple description of the event. SMTP alerting is configured using the embedded web server.

### 6.12.12 SM-CLP (SM-CLP Lite)

SMASH refers to Systems Management Architecture for Server Hardware. SMASH is defined by a suite of specifications, managed by the DMTF, that standardize the manageability interfaces for server hardware. CLP refers to Command Line Protocol. SM-CLP is defined by the *Server Management Command Line Protocol Specification (SM-CLP) ver1.0*, which is part of the SMASH suite of specifications. The specifications and further information on SMASH can be found at the DMTF website (http://www.dmtf.org/).

The BMC provides an embedded “lite” version of SM-CLP that is syntax-compatible but not considered fully compliant with the DMTF standards.

The SM-CLP utilized by a remote user by connecting a remote system through one of the system NICs. It is possible for third party management applications to create scripts using this CLP and execute them on server to retrieve information or perform management tasks such as reboot the server, configure events, and so on.

The BMC embedded SM-CLP feature includes the following capabilities:

- Power on/off/reset the server.
- Get the system power state.
- Clear the System Event Log (SEL).
- Get the interpreted SEL in a readable format.
- Initiate/terminate an Serial Over LAN session.
- Support “help” to provide helpful information
- Get/set the system ID LED.
- Get the system GUID
- Get/set configuration of user accounts.
- Get/set configuration of LAN parameters.
- Embedded CLP communication should support SSH connection.
- Provide current status of platform sensors including current values. Sensors include voltage, temperature, fans, power supplies, and redundancy (power unit and fan redundancy).

The embedded web server is supported over any system NIC port that is enabled for server management capabilities.

### 6.12.13 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 RMM4 dedicated add-in management NIC. At least two concurrent web sessions from up to two different users is supported. The embedded web user interface shall support the following client web browsers:

- Microsoft Internet Explorer 7.0*
- Microsoft Internet Explorer 8.0*
- Microsoft Internet Explorer 9.0*
- Mozilla Firefox 3.0*
- Mozilla Firefox 3.5*
- Mozilla Firefox 3.6*

The embedded web user interface supports strong security (authentication, encryption, and firewall support) since it enables remote server configuration and control. Embedded web server uses ports #80 and #443. 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.

Additional features supported by the web GUI includes:

- Presents all the Basic features to the users.
- Power on/off/reset the server and view current power state.
- Displays BIOS, BMC, ME and SDR version information.
- Display overall system health.
- Configuration of various IPMI over LAN parameters for both IPV4 and IPV6
- Configuration of alerting (SNMP and SMTP).
- Display system asset information for the product, board, and chassis.
- Display of BMC-owned sensors (name, status, current reading, enabled thresholds), including color-code status of sensors.
- Provides ability to filter sensors based on sensor type (Voltage, Temperature, Fan and Power supply related)
- Automatic refresh of sensor data with a configurable refresh rate.
- On-line help.
- Display/clear SEL (display is in easily understandable human readable format).
- Supports major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*).
- Automatically logs out after user-configurable inactivity period.
- 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 "diagnostic 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.
- Severity level indication of SEL events. The web server UI displays the severity level associated with each event in the SEL. The severity level correlates with the front panel system status LED ( "OK", "Degraded", "Non-Fatal", or "Fatal").

- 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.

### 6.12.14 Virtual Front Panel

- Virtual Front Panel is the module present as “Virtual Front Panel” on the left side in the embedded web server when "remote Control" tab is clicked.
- Main Purpose of the Virtual Front Panel is to provide the front panel functionality virtually.
- Virutal Front Panel (VFP) will mimic the status LED and Power LED status and Chassis ID alone. It is automatically in sync with BMC every 40 seconds.
- For any abnormal status LED state, Virtual Front Panel will get the reason behind the abnormal or status LED changes and displayed in VFP side.
- As Virtual Front Panel uses the chassis control command for power actions. It won’t log the Front button press event since Logging the front panel press event for Virtual Front Panel press will mislead the administrator.
- For Reset through Virtual Front Panel, the reset will be done by a “Chassis control” command.
- For Reset through Virtual Front Panel, the restart cause will be because of “Chassis control” command.
- During Power action, Power button/Reset button should not accept the next action until current Power action is complete and the acknowledgment from BMC is received.
- EWS will provide a valid message during Power action until it completes the current Power action.
- The VFP does not have any effect on whether the front panel is locked by “Set Front Panel Enables” command.
- The chassis ID LED provides a visual indication of a system being serviced. The state of the chassis ID LED is affected by the following actions:
- Toggled by turning the chassis ID button on or off.

- There is no precedence or lock-out mechanism for the control sources. When a new request arrives, previous requests are terminated. For example, if the chassis ID button is pressed, then the chassis ID LED changes to solid on. If the button is pressed again, then the chassis ID LED turns off.
- Note that the chassis ID will turn on because of the original chassis ID button press and will reflect in the Virtual Front Panel after VFP sync with BMC. Virtual Front Panel won't reflect the chassis LED software blinking through software command as there is no mechanism to get the chassis ID Led status.
- Only Infinite chassis ID ON/OFF through software command will reflect in EWS during automatic/manual EWS sync up with BMC.
- Virtual Front Panel help should available for virtual panel module.
- At present, NMI button in VFP is disabled in Romley. It can be used in future.

### 6.12.15 Embedded Platform Debug

The Embedded Platform Debug feature supports capturing low-level diagnostic data (applicable MSRs, PCI config-space registers, and so on.). This feature allows a user to export this data into a file that is retrievable through the embedded web GUI, as well as through host and remote IPMI methods, for the purpose of sending to an Intel® engineer for an enhanced debugging capability. The files are compressed, encrypted, and password protected. The file is not meant to be viewable by the end user but rather to provide additional debugging capability to an Intel® support engineer.

A list of data that may be captured using this feature includes but is not limited to:

- Platform sensor readings – This includes all "readable" sensors that can be accessed by the BMC FW and have associated SDRs populated in the SDR repository. This does not include any "event-only" sensors. (All BIOS sensors and some BMC and ME sensors are "event-only"; meaning that they are not readable using an IPMI Get Sensor Reading command but rather are used just for event logging purposes).
- SEL – The current SEL contents are saved in both hexadecimal and text format.
- CPU/memory register data – useful for diagnosing the cause of the following system errors: CATERR, ERR[2], SMI timeout, PERR, and SERR. The debug data is saved and time-stamped for the last 3 occurrences of the error conditions.
    - PCI error registers
    - MSR registers
    - MCH registers
- BMC configuration data
    - BMC FW debug log (that is, SysLog) – Captures FW debug messages.
    - *Non-volatile storage of captured data.* Some of the captured data will be stored persistently in the BMC's non-volatile flash memory and preserved across AC power cycles. Due to size limitations of the BMC's flash memory, it is not feasible to store all of the data persistently.
- *SMBIOS table data*. The entire SMBIOS table is captured from the last boot.

- *PCI configuration data for on-board devices and add-in cards*. The first 256 bytes of PCI configuration data is captured for each device for each boot.
- *System memory map.* The system memory map is provided by BIOS on the current boot. This includes the EFI memory map and the Legacy (E820) memory map depending on the current boot.
- *Power supplies debug capability.*
    - *Capture of power supply "black box" data and power supply asset information.* Power supply vendors are adding the capability to store debug data within the power supply itself. The platform debug feature provides a means to capture this data for each installed power supply. The data can be analyzed by Intel® for failure analysis and possibly provided to the power supply vendor as well. The BMC gets this data from the power supplies through PMBus* manufacturer-specific commands.
    - *Storage of system identification in power supply.* The BMC copies board and system serial numbers and part numbers into the power supply whenever a new power supply is installed in the system or when the system is first powered on. This information is included as part of the power supply black box data for each installed power supply.
- *Accessibility through IPMI interfaces.* The platform debug file can be accessed through an external IPMI interface (KCS or LAN).
- *POST code sequence for the two most recent boots*. This is a best-effort data collection by the BMC as the BMC real-time response cannot guarantee that all POST codes are captured.
- *Support for multiple debug files.* The platform debug feature provides the ability to save data to 2 separate files that are encrypted with different passwords.
    - File #1 is strictly for viewing by Intel® engineering and may contain BMC log messages (that is, syslog) and other debug data that Intel® FW developers deem useful in addition to the data specified in this document.
    - File #2 can be viewed by Intel® partners who have signed an NDA with Intel® and its contents are restricted to specific data items specified in this with the exception of the BMC syslog messages and power supply "black box" data.

#### 6.12.15.1 Output Data Format

The diagnostic feature shall output a password-protected compressed HTML file containing specific BMC and system information. This file is not intended for end-customer usage, this file is for customer support and engineering only.

#### 6.12.15.2 Output Data Availability

The diagnostic data shall be available on-demand through the embedded web server, KCS, or IPMI over LAN commands.

#### 6.12.15.3 Output Data Categories

The following tables list the data to be provided in the diagnostic output.

**Table 20. Diagnostic Data**

| Category | Data |
|---|---|
| Internal BMC Data | BMC uptime/load |
| | Process list |
| | Free Memory |
| | Detailed Memory List |
| | Filesystem List/Info |
| | BMC Network Info |
| | BMC Syslog |
| | BMC Configuration Data |
| External BMC Data | Hex SEL listing |
| | Human-readable SEL listing |
| | Human-readable sensor listing |
| External BIOS Data | BIOS configuration settings |
| | POST codes for the two most recent boots |
| System Data | SMBIOS table for the current boot |
| | 256 bytes of PCI config data for each PCI device |
| | Memory Map (EFI and Legacy) for current boot |

**Table 21. Additional Diagnostics on Error**

| Category | Data |
|---|---|
| System Data | First 256 bytes of PCI config data for each PCI device |
| | PCI error registers |
| | MSR registers |
| | MCH registers |

### 6.12.16 Data Center Management Interface (DCMI)

The *DCMI Specification* is an emerging standard that is targeted to provide a simplified management interface for Internet Portal Data Center (IPDC) customers. It is expected to become a requirement for server platforms which are targeted for IPDCs. DCMI is an IPMI-based standard that builds upon a set of required IPMI standard commands by adding a set of DCMI-specific IPMI OEM commands. Intel® S1400/S1600/S2400/S2600 Server Platforms will be implementing the mandatory DCMI features in the BMC firmware (DCMI 1.1 Errata 1 compliance). Please refer to DCMI 1.1 errata 1 specification for details. Only mandatory commands will be supported. No support for optional DCMI commands. Optional power management and SEL roll over feature is not supported. DCMI Asset tag will be independent of baseboard FRU asset Tag. Please refer table DCMI Group Extension Commands for more details on DCMI commands.

### 6.12.17 Lightweight Directory Authentication Protocol (LDAP)

The Lightweight Directory Access Protocol (LDAP) is an application protocol supported by the BMC for the purpose of authentication and authorization. The BMC user connects with an LDAP server for login authentication. This is only supported for non-IPMI logins including the embedded web UI and SM-CLP. IPMI users/passwords and sessions are not supported over LDAP.

LDAP can be configured (IP address of LDAP server, port, and so on.) through the BMC's Embedded Web UI. LDAP authentication and authorization is supported over the any NIC

configured for system management. The BMC uses a standard Open LDAP implementation for Linux*.

Only open LDAP is supported by BMC. Windows* and Novel* LDAP are not supported.

# 7 Advanced Management Feature Support (RMM4)

The integrated baseboard management controller has support for advanced management features which are enabled when an optional Intel® Remote Management Module 4 (RMM4) is installed.

RMM4 is comprised of two boards – RMM4 lite and the optional Dedicated Server Management NIC (DMN).

**Table 22. RMM4 Option Kits**

| Intel® Product Code | Description | Kit Contents | Benefits |
|---|---|---|---|
| AXXRMM4LITE | Intel® Remote Management Module 4 Lite | RMM4 Lite Activation Key | Enables KVM and media redirection through onboard NIC |
| AXXRMM4 | Intel® Remote Management Module 4 | RMM4 Lite Activation Key<br>Dedicated NIC Port Module | Dedicated NIC for management traffic. Higher bandwidth connectivity for KVM and media Redirection with 100Mbe NIC. |

On the server board each Intel® RMM4 component is installed at the following locations.

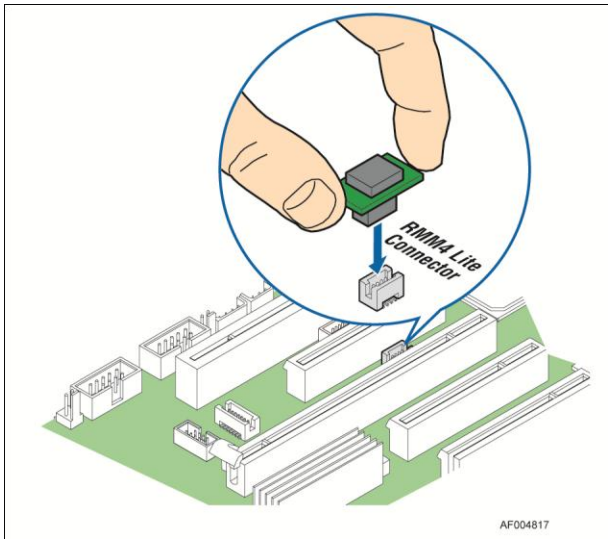


**Figure 22. Intel® RMM4 Lite Activation Key Installation**

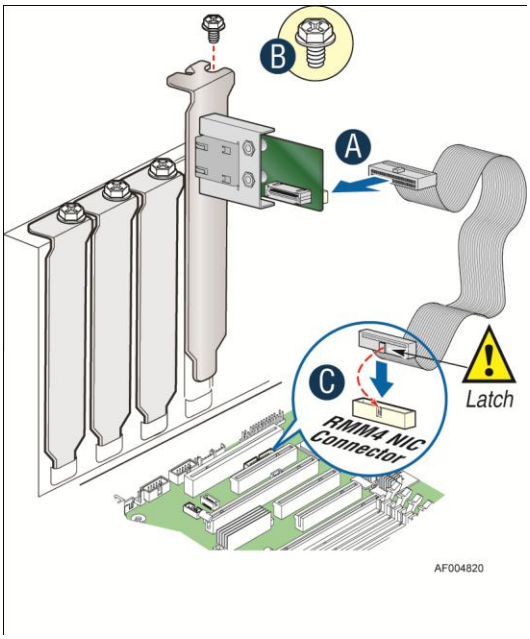


**Figure 23. Intel® RMM4 Dedicated Management NIC Installation**

If the optional Dedicated Server Management NIC is not used then the traffic can only go through the onboard Integrated BMC-shared NIC and will share network bandwidth with the host system. *Advanced* manageability features are supported over all NIC ports enabled for server manageability.

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

### 7.1.1 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, from the firewall and, in case of a private internal network, the NAT (Network Address Translation) settings have to be configured accordingly.

### 7.1.2 Performance

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

Enabling KVM and/or media encryption will degrade performance. Enabling video compression provides the fastest response while disabling compression provides better video quality.
For the best possible KVM performance, a 2Mb/sec link or higher is recommended.
The redirection of KVM over IP is performed in parallel with the local KVM without affecting the local KVM operation.

### 7.1.3 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.

### 7.1.4 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. An BMC reset (for example, due to an BMC Watchdog initiated reset or BMC reset after BMC FW update) will require the session to be re-established.

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

### 7.1.5 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 same server and start remote KVM sessions

### 7.1.6 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.

## 7.2 Media Redirection

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

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

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

- The media redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.
- A remote media session is maintained even when the server is powered-off (in standby mode). No restart of the remote media session is required during a server reset or power on/off. An BMC reset (for example, due to an BMC reset after BMC 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.

### 7.2.1 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.

### 7.2.2 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)

# 8 On-board Connector/Header Overview

## 8.1 Board Connector Information

The following section provides detailed information regarding all connectors, headers, and jumpers on the server boards.

The following table lists all connector types available on the board, as well as the corresponding preference designators printed on the silkscreen.

**Table 23. Board Connector Matrix**

| Connector | Quantity | Connector Type | Pin Count |
|---|---|---|---|
| Power supply | 4 | Main power<br>CPU 1 power<br>CPU 2 Power<br>P/S aux/IPMB | 24<br>8<br>8<br>5 |
| CPU | 2 | CPU sockets | 1356 |
| Main memory | 8 | DIMM sockets | 240 |
| PCI Express* x8 | 3 | Card edge | 98 |
| PCI Express* x16 | 2 | Card edge | 164 |
| 32-bit PCI | 1 | Card edge | 124 |
| Intel® RMM4 | 1 | Connector | 30 |
| Intel® RMM4 Lite | 1 | Connector | 7 |
| Storage Upgrade Key | 1 | Header | 4 |
| System fans | 6 | Header | 6 |
| System fans | 1 | Header | 4 |
| CPU fans | 2 | Header | 4 |
| Battery | 1 | Battery holder | 2 |
| Stacked RJ45/2xUSB | 2 | External LAN built-in magnetic and dual USB | 22 |
| Video | 1 | External DSub | 15 |
| Serial port A | 1 | Connector | 9 |
| Serial port B | 1 | Header | 9 |
| Front panel | 1 | Header | 24 |

| Connector | Quantity | Connector Type | Pin Count |
|---|---|---|---|
| Internal USB | 2 | Header | 10 |
| USB Solid State Drive | 1 | Header | 9 |
| Internal USB | 1 | Type-A USB | 4 |
| Chassis Intrusion | 1 | Header | 2 |
| Serial ATA | 6 | Header(white) | 7 |
| SAS | 2 | SFF8087 miniSAS | 36 |
| HSBP_I2C | 1 | Header | 3 |
| SATA SGPIO | 1 | Header | 4 |
| LCP | 1 | Header | 7 |
| IPMB | 1 | Header | 4 |
| Configuration jumpers | 6 | Jumpers | 3 |
| TPM | 1 | Connector | 14 |

## 8.2 Power Connectors

The main power supply connection uses an SSI-compliant 2x12 pin connector.

Three additional power-related connectors also exist:

- Two SSI-compliant 2x4 pin power connectors to provide 12-V power to the CPU voltage regulators and memory.
- One SSI-compliant 1x5 pin connector to provide $I^2C$ monitoring of the power supply.

The following tables define these connector pin-outs:

**Table 24. Main Power Connector Pin-out**

| Pin | Signal | Color | Pin | Signal | Color |
|---|---|---|---|---|---|
| 1 | +3.3 Vdc | Orange | 13 | +3.3 Vdc | Orange |
| 2 | +3.3 Vdc | Orange | 14 | -12 Vdc | Blue |
| 3 | GND | Black | 15 | GND | Black |
| 4 | +5 Vdc | Red | 16 | PS_ON# | Green |
| 5 | GND | Black | 17 | GND | Black |
| 6 | +5 Vdc | Red | 18 | GND | Black |
| 7 | GND | Black | 19 | GND | Black |
| 8 | PWR_OK | Gray | 20 | NC | White |
| 9 | 5 VSB | Purple | 21 | +5 Vdc | Red |
| 10 | +12 Vdc | Yellow | 22 | +5 Vdc | Red |
| 11 | +12 Vdc | Yellow | 23 | +5 Vdc | Red |
| 12 | +3.3 Vdc | Orange | 24 | GND | Black |

**Table 25. CPU 1/CPU 2 Power Connector Pin-out**

| Pin | Signal | Color |
|---|---|---|
| 1 | GND of Pin 5 | Black |
| 2 | GND of Pin 6 | Black |
| 3 | GND of Pin 7 | Black |
| 4 | GND of Pin 8 | Black |
| 5 | +12 Vdc CPU1 | Yellow/black |
| 6 | +12 Vdc CPU1 | Yellow/black |
| 7 | +12 Vdc DDR3_CPU1 | Yellow/black |
| 8 | +12 Vdc DDR3_CPU1 | Yellow/black |

**Table 26. Power Supply Auxiliary Signal Connector Pin-out**

| Pin | Signal | Color |
|---|---|---|
| 1 | SMB_CLK_FP_PWR_R | Orange |
| 2 | SMB_DAT_FP_PWR_R | Black |
| 3 | SMB_ALRT_3_ESB_R | Red |
| 4 | 3.3 V SENSE- | Yellow |
| 5 | 3.3 V SENSE+ | Green |

## 8.3 System Management Headers

### 8.3.1 Intel® Remote Management Module 4 Connector

A 30-pin Intel® RMM4 connector and a 7-pin Intel® RMM4 Lite connector are included on the server board to support the optional Intel® Remote Management Module 4 or Intel® Remote Management Module 4 Lite. This server board does not support third-party management cards.

**Note:** This connector is not compatible with the previous generation Intel® Remote Management Modules (Intel® RMM/RMM2/RMM3)

**Table 27. Intel® RMM4 Connector Pin-out**

| Pin | Signal Name | Pin | Signal Name |
|---|---|---|---|
| 1 | 3V3_AUX | 2 | MDIO |
| 3 | 3V3_AUX | 4 | MDC |
| 5 | GND | 6 | TXD_0 |
| 7 | GND | 8 | TXD_1 |
| 9 | GND | 10 | TXD_2 |
| 11 | GND | 12 | TXD_3 |
| 13 | GND | 14 | TX_CTL |
| 15 | GND | 16 | RX_CTL |
| 17 | GND | 18 | RXD_0 |
| 19 | GND | 20 | RXD_1 |
| 21 | GND | 22 | RXD_2 |
| 23 | 3V3_AUX | 24 | RXD_3 |
| 25 | 3V3_AUX | 26 | TX_CLK |
| 27 | 3V3_AUX | 28 | RX_CLK |
| 29 | GND | 30 | PRESENT# |

**Table 28. Intel® RMM4 – Lite Connector Pin-out**

| Pin | Signal Name | Pin | Signal Name |
|---|---|---|---|
| 1 | 3V3_AUX | 2 | SPI_RMM4_LITE_DI |
| 3 | N/A | 4 | SPI_RMM4_LITE_CLK |
| 5 | SPI_RMM4_LITE_DO | 6 | GND |
| 7 | SPI_RMM4_LITE_CS_N | 8 | GND |

### 8.3.2 TPM connector

**Table 29. TPM connector Pin-out**

| Pin | Signal Name | Pin | Signal Name |
|---|---|---|---|
| 1 | No pin | 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 | 10 | CLK_33M_TPM_CONN |
| 11 | LPC_LAD<3> | 12 | GND |
| 13 | GND | 14 | LPC_LAD<2> |

### 8.3.3 LCP Header

**Table 30. LCP Header Pin-out**

| Pin | Signal Name |
|---|---|
| 1 | SMB_SENSOR_3V3STBY_DATA |
| 2 | GND |
| 3 | SMB_SENSOR_3V3STBY_CLK |
| 4 | P3V3_AUX |
| 5 | FM_LCP_ENTER_N |
| 6 | FM_LCP_LEFT_N |
| 7 | FM_LCP_RIGHT_N |

### 8.3.4 HSBP Header

**Table 31. HSBP_I2C Header Pin-out**

| Pin | Signal Name |
|---|---|
| 1 | SMB_HSBP_3V3STBY_DATA |
| 2 | GND |
| 3 | SMB_HSBP_3V3STBY_CLK |

### 8.3.5 SATA SGPIO Header

**Table 32. SGPIO Header Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | SGPIO_CLOCK | SGPIO Clock Signal |
| 2 | SGPIO_LOAD | SGPIO Load Signal |
| 3 | SGPIO_DATAOUT0 | SGPIO Data Out |
| 4 | SGPIO_DATAOUT1 | SGPIO Data In |

## 8.4 Front Panel Connector

The server board provides a 30-pin SSI front panel connector for use with Intel® and third-party chassis. The following table provides the pin-out for this connector:

**Table 33. Front Panel SSI Standard 30-pin Connector Pin-out**

| Signal Description | Pin# | Pin# | Signal Description |
|---|---|---|---|
| P3V3_AUX | 1 | 2 | P3V3_AUX |
| KEY | | 4 | P5V_STBY |
| FP_PWR_LED_BUF_R_N | 5 | 6 | FP_ID_LED_BUF_R_N |
| P3V3 | 7 | 8 | FP_LED_STATUS_GREEN_R_N |
| LED_HDD_ACTIVITY_R_N | 9 | 10 | FP_LED_STATUS_AMBER_R_N |
| FP_PWR_BTN_N | 11 | 12 | LED _NIC_LINK0_ACT_FP_N |
| GROUND | 13 | 14 | LED _NIC_LINK0_LNKUP_FP_N |
| FP_RST_BTN_R_N | 15 | 16 | SMB_SENSOR_3V3STBY_DATA_R0 |
| GROUND | 17 | 18 | SMB_SENSOR_3V3STBY_CLK |
| FP_ID_BTN_R_N | 19 | 20 | FP_CHASSIS_INTRUSION |
| PU_FM_SIO_TEMP_SENSOR | 21 | 22 | LED_NIC_LINK1_ACT_FP_N |
| FP_NMI_BTN_R_N | 23 | 24 | LED_NIC_LINK1_LNKUP_FP_N |
| KEY | | | KEY |
| LED_NIC_LINK2_ACT_FP_N | 27 | 28 | LED_NIC_LINK3_ACT_FP_N |
| LED_NIC_LINK2_LNKUP_FP_N | 29 | 30 | LED_NIC_LINK3_LNKUP_FP_N |

## 8.5 I/O Connectors

### 8.5.1 VGA Connector

The following table details the pin-out definition of the VGA connector that is part of the stacked video/serial port A connector.

**Table 34. VGA Connector Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | V_IO_R_CONN | Red (analog color signal R) |
| 2 | V_IO_G_CONN | Green (analog color signal G) |
| 3 | V_IO_B_CONN | Blue (analog color signal B) |
| 4 | TP_VID_CONN_B4 | No connection |
| 5 | GND | Ground |
| 6 | GND | Ground |
| 7 | GND | Ground |
| 8 | GND | Ground |
| 9 | P5V | +5V DC |
| 10 | GND | Ground |
| 11 | TP_VID_CONN_B11 | No connection |
| 12 | V_IO_DDCDAT | DDCDAT |
| 13 | V_IO_HSYNC_CONN | HSYNC (horizontal sync) |
| 14 | V_IO_VSYNC_CONN | VSYNC (vertical sync) |
| 15 | V_IO_DDCCLK | DDCCLK |

### 8.5.2 NIC Connectors

The server board provides two stacked RJ-45/2xUSB connectors side-by-side on the back edge of the board. The pin-out for NIC connectors is identical and defined in the following table.

**Table 35. RJ-45 10/100/1000 NIC Connector Pin-out**

| Pin | Signal Name |
|---|---|
| 1 | GND |
| 2 | P1V8_NIC |
| 3 | NIC_A_MDI3P |
| 4 | NIC_A_MDI3N |
| 5 | NIC_A_MDI2P |
| 6 | NIC_A_MDI2N |
| 7 | NIC_A_MDI1P |
| 8 | NIC_A_MDI1N |
| 9 | NIC_A_MDI0P |
| 10 | NIC_A_MDI0N |
| 11 | NIC_LINKA_1000_N (LED |
| 12 | NIC_LINKA_100_N (LED) |
| 13 | NIC_ACT_LED_N |
| 14 | NIC_LINK_LED_N |
| 15 | GND |
| 16 | GND |

### 8.5.3 SATA Connectors

The server board provides up to 6 SATA connectors: SATA-0 to SATA-5, and 2 MiniSAS connectors: SCU-0 and SCU-1.

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

**Table 36. SATA, MiniSAS SCU-0 and SCU-1 Connector Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | GND | Ground |
| 2 | SATA TX_P_C | Positive side of transmit differential pair |
| 3 | SATA TX_N_C | Negative side of transmit differential pair |
| 4 | GND | Ground |
| 5 | SATA _RX_N_C | Negative side of receive differential pair |
| 6 | SATA _RX_P_C | Positive side of receive differential pair |
| 7 | GND | Ground |

| Signal Description | Pin# | Pin# | Signal Description |
|---|---|---|---|
| GROUND | A1 | B1 | GROUND |
| SAS0_RX_C_DP | A2 | B2 | SAS0_TX_C_DP |
| SAS0_RX_C_DN | A3 | B3 | SAS0_TX_C_DN |
| GROUND | A4 | B4 | GROUND |
| SAS1_RX_C_DP | A5 | B5 | SAS1_TX_C_DP |
| SAS1_RX_C_DN | A6 | B6 | SAS1_TX_C_DN |
| GROUND | A7 | B7 | GROUND |
| TP_SAS1_BACKPLANE_TYPE | A8 | B8 | SGPIO_SAS1_CLOCK |
| GROUND | A9 | B9 | SGPIO_SAS1_LOAD |
| SGPIO_SAS1_DATAOUT | A10 | B10 | GROUND |
| SGPIO_SAS1_DATAIN | A11 | B11 | PD_SAS1_CONTROLLER_TYPE |
| GROUND | A12 | B12 | GROUND |
| SAS2_RX_C_DP | A13 | B13 | SAS2_TX_C_DP |
| SAS2_RX_C_DN | A14 | B14 | SAS2_TX_C_DN |
| GROUND | A15 | B15 | GROUND |
| SAS3_RX_C_DP | A16 | B16 | SAS3_TX_C_DP |
| SAS3_RX_C_DN | A17 | B17 | SAS3_TX_C_DN |
| GROUND | A18 | B18 | GROUND |

| Signal Description | Pin# | Pin# | Signal Description |
|---|---|---|---|
| GROUND | MTH1 | MTH5 | GROUND |
| GROUND | MTH2 | MTH6 | GROUND |
| GROUND | MTH3 | MTH7 | GROUND |
| GROUND | MTH4 | MTH8 | GROUND |

| Signal Description | Pin# | Pin# | Signal Description |
|---|---|---|---|
| GROUND | A1 | B1 | GROUND |
| SAS4_RX_C_DP | A2 | B2 | SAS4_TX_C_DP |
| SAS4_RX_C_DN | A3 | B3 | SAS4_TX_C_DN |
| GROUND | A4 | B4 | GROUND |
| SAS5_RX_C_DP | A5 | B5 | SAS5_TX_C_DP |
| SAS5_RX_C_DN | A6 | B6 | SAS5_TX_C_DN |
| GROUND | A7 | B7 | GROUND |
| TP_SAS2_BACKPLANE_TYPE | A8 | B8 | SGPIO_SAS2_CLOCK |
| GROUND | A9 | B9 | SGPIO_SAS2_LOAD |
| SGPIO_SAS2_DATAOUT | A10 | B10 | GROUND |
| SGPIO_SAS2_DATAIN | A11 | B11 | PD_SAS2_CONTROLLER_TYPE |
| GROUND | A12 | B12 | GROUND |
| SAS6_RX_C_DP | A13 | B13 | SAS6_TX_C_DP |
| SAS6_RX_C_DN | A14 | B14 | SAS6_TX_C_DN |
| GROUND | A15 | B15 | GROUND |
| SAS7_RX_C_DP | A16 | B16 | SAS7_TX_C_DP |
| SAS7_RX_C_DN | A17 | B17 | SAS7_TX_C_DN |
| GROUND | A18 | B18 | GROUND |
| GROUND | MTH1 | MTH5 | GROUND |
| GROUND | MTH2 | MTH6 | GROUND |
| GROUND | MTH3 | MTH7 | GROUND |
| GROUND | MTH4 | MTH8 | GROUND |

### 8.5.4 Serial Port Connectors

The server board provides one external DB9 Serial A port and one internal 9-pin Serial B header. The following tables define the pin-outs.

**Table 37. External DB9 Serial A Port Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | SPA_DCD | DCD (carrier detect) |
| 2 | SPA_SIN_L | RXD (receive data) |
| 3 | SPA_SOUT_N | TXD (Transmit data) |
| 4 | SPA_DTR | DTR (Data terminal ready) |
| 5 | GND | Ground |
| 6 | SPA_DSR | DSR (data set ready) |
| 7 | SPA_RTS | RTS (request to send) |
| 8 | SPA_CTS | CTS (clear to send) |
| 9 | SPA_RI | RI (Ring Indicate) |

**Table 38. Internal 9-pin Serial B Header Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | SPB_DCD | DCD (carrier detect) |
| 2 | SPB_DSR | DSR (data set ready) |
| 3 | SPB_SIN_L | RXD (receive data) |
| 4 | SPB_RTS | RTS (request to send) |
| 5 | SPB_SOUT_N | TXD (Transmit data) |
| 6 | SPB_CTS | CTS (clear to send) |

| Pin | Signal Name | Description |
|---|---|---|
| 7 | SPB_DTR | DTR (Data terminal ready) |
| 8 | SPB_RI | RI (Ring indicate) |
| 9 | GND | Ground |

### 8.5.5 USB Connector

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

**Table 39. External USB Connector Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | USB_OC_5VSB | USB_PWR |
| 2 | USB_PN | DATAL0 (Differential data line paired with DATAH0) |
| 3 | USB_PP | DATAH0 (Differential data line paired with DATAL0) |
| 4 | GND | Ground |

Two 2x5 connectors on the server board provide support for four additional USB ports. J1G2 is recommended for front panel USB ports.

**Table 40. Internal USB Connector Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | USB_PWR_5V | USB power |
| 2 | USB_PWR_5V | USB power |
| 3 | USB _PN_CONN | USB port negative signal |
| 4 | USB _PN_CONN | USB port negative signal |
| 5 | USB _PP_CONN | USB port positive signal |
| 6 | USB _PP_CONN | USB port positive signal |
| 7 | Ground | |
| 8 | Ground | |
| 9 | Key | No pin |
| 10 | TP_USB _NC | Test point |

One low-profile 2x5 connector on the server boards provides an option to support a low-profile USB Solid State Drive

**Table 41. Pin-out of Internal Low-Profile USB Connector for Solid State Drive**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | +5V | USB power |
| 2 | USB_N | USB port negative signal |
| 3 | USB_P | USB port positive signal |
| 4 | GND | Ground |
| 5 | Key | No pin |
| 6 | NC | Not Connected |
| 7 | NC | Not Connected |
| 8 | NC | Not Connected |
| 9 | NC | Not Connected |
| 10 | LED# | Activity LED |

The server board provides one additional Type A USB port to support the installation of a USB device inside the server chassis.

**Table 42. Internal Type A USB Port Pin-out**

| Pin | Signal Name | Description |
|---|---|---|
| 1 | USB_PWR7_5V | USB_PWR |
| 2 | USB_PN | USB port negative signal |
| 3 | USB _PP | USB port positive signal |
| 4 | GND | Ground |

## 8.6 Fan Headers

The server board provides three SSI-compliant 4-pin and six SSI-compliant 6-pin fan headers to use as CPU and I/O cooling fans. 3-pin fans are supported on all fan headers. 6-pin fans are supported on headers SYS FAN_1 to SYS FAN_6. 4-pin fans are supported on headers SYS FAN_7. The pin configuration for each of the 4-pin and 6-pin fan headers is identical and defined in the following tables.

- Two 4-pin fan headers are designated as processor cooling fans:
  - CPU1 fan
  - CPU2 fan
- Four 6-pin fan headers are designated as hot-swap system fans:
  - Hot-swap system fan 1
  - Hot-swap system fan 2
  - Hot-swap system fan 3
  - Hot-swap system fan 4
  - Hot-swap system fan 5
  - Hot-swap system fan 6
- One 4-pin fan header is designated as a rear system fan:
  - System fan 7

**Table 43. SSI 4-pin Fan Header Pin-out**

| Pin | Signal Name | Type | Description |
|---|---|---|---|
| 1 | Ground | GND | Ground is the power supply ground |
| 2 | 12V | Power | Power supply 12 V |
| 3 | Fan Tach | In | FAN_TACH signal is connected to the BMC to monitor the fan speed |
| 4 | Fan PWM | Out | FAN_PWM signal to control fan speed |

**Table 44. SSI 6-pin Fan Header Pin-out**

| Pin | Signal Name | Type | Description |
|---|---|---|---|
| 1 | Ground | GND | Ground is the power supply ground |
| 2 | 12V | Power | Power supply 12 V |
| 3 | Fan Tach | In | FAN_TACH signal is connected to the BMC to monitor the fan speed |
| 4 | Fan PWM | Out | FAN_PWM signal to control fan speed |
| 5 | Fan Presence | In | Indicates the fan is present |
| 6 | Fan Fault LED | Out | Lights the fan fault LED |

**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.

# 9 Jumper Blocks

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

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

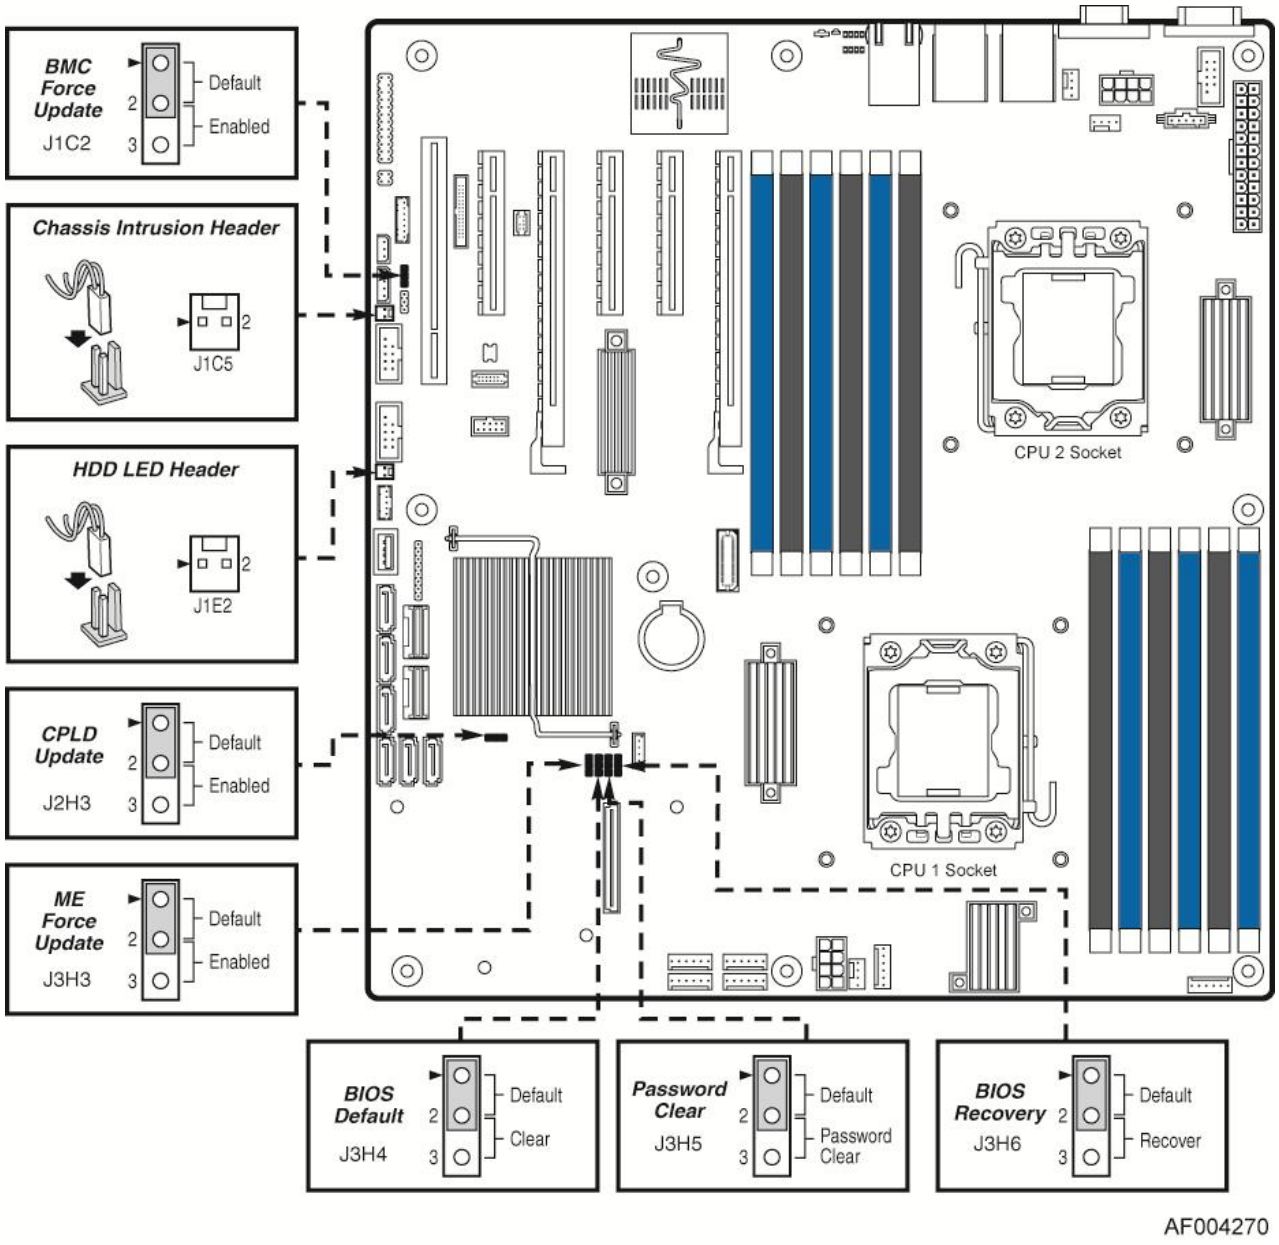


**Figure 24. Jumper Blocks (J1C2, J1C5, J1E2, J2H3, J3H3, J3H4, J3H5, J3H6)**

**Table 45. Server Board Jumpers (J1C2, J1E2, J1E4, J1E5, J1H1)**

| Jumper Name | Pins | System Results |
|---|---|---|
| J1C2: BMC | 1-2 | BMC Firmware Force Update Mode – Disabled (Default) |

| Jumper Name | Pins | System Results |
|---|---|---|
| Force Update | 2-3 | BMC Firmware Force Update Mode – Enabled |
| J3H6: BIOS Recovery | 1-2 | Pins 1-2 should be jumpered for normal system operation. (Default) |
| | 2-3 | The main system BIOS does not boot with pins 2-3 jumpered. The system only boots from EFI-bootable recovery media with a recovery BIOS image present. |
| J3H4: BIOS Default | 1-2 | These pins should have a jumper in place for normal system operation. (Default) |
| | 2-3 | If pins 2-3 are jumpered with AC power plugged in, the CMOS settings clear in 5 seconds. Pins 2-3 should not be jumpered for normal system operation. |
| J3H3: ME Force Update | 1-2 | ME Firmware Force Update Mode – Disabled (Default) |
| | 2-3 | ME Firmware Force Update Mode – Enabled |
| J3H5: Password Clear | 1-2 | These pins should have a jumper in place for normal system operation. |
| | 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. |
| J2H3: CPLD Update | 1-2 | CPLD online update disabled(Default) |
| | 2-3 | CPLD online update enabled |

**Note**:

1. For safety purposes, the power cord should be disconnected from a system before removing any system components or moving any of the on-board jumper blocks.
2. Access to jumper blocks near Riser Slot #1 is limited, removal of the riser card may be necessary to move these jumpers.
3. System Update and Recovery files are included in the System Update Packages (SUP) posted to Intel's web site.

## 9.1 BIOS Default (CMOS Clear) and Password Clear Usage Procedure

The BIOS default (J3H4) and Password Clear (J3H5) recovery features are designed to achieve the desired operation with minimum system down time. The usage procedure for these two features has changed from previous generation Intel® server boards. The following procedure outlines the new usage model.

### 9.1.1 Clearing CMOS (BIOS default)

1. Power down the server and unplug the AC power cord.
2. Open the server chassis. For instructions, see your server chassis documentation.
3. Move the jumper (J3H4) from the default operating position (covering pins 1 and 2) to the reset/clear position (covering pins 2 and 3).
4. Wait 5 seconds.
5. Move the jumper back to default position, covering pins 1 and 2.
6. Close the server chassis and reconnect the AC power cord.
7. Power up the server.

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

**Note:** Removing AC power before performing the CMOS Clear operation causes the system to automatically power up and immediately power down after the procedure is followed and AC

power is re-applied. If this happens, remove the AC power cord again, wait 30 seconds, and re-install the AC power cord. Power up the system and proceed to the <F2> BIOS Setup Utility to reset the desired settings.

### 9.1.2 Clearing the Password

1. Power down the server. Do not unplug the power cord.
2. Open the chassis. For instructions, see your server chassis documentation.
3. Move the jumper (J3H5) from the default operating position, covering pins 1 and 2, to the password clear position, covering pins 2 and 3.
4. Close the server chassis.
5. Power up the server and then press <F2> to enter the BIOS menu to check if the password is cleared.
6. Power down the server.
7. Open the chassis and move the jumper back to its default position, covering pins 1 and 2.
8. Close the server chassis.
9. Power up the server.

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

## 9.2 Force BMC Update Procedure

When performing a standard BMC firmware update procedure, the update utility places the BMC into an update mode, allowing the firmware to load safely onto the flash device. In the unlikely event the BMC firmware update process fails due to the BMC not being in the proper update state, the server board provides a Force BMC Update jumper (J1C2) that forces the BMC into the proper update state. In the event the standard BMC firmware update process fails, complete the following procedure.

1. Power down and remove the AC power cord.
2. Open the server chassis. See your server chassis documentation for instructions.
3. Move the jumper (J1C2) 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 BMC firmware update procedure as documented in the README.TXT file included in the given BMC firmware update package. After successful completion of the firmware update process, the firmware update utility may generate an error stating the BMC is still in update mode.
7. Power down and remove the AC power cord.
8. Open the server chassis.
9. Move the jumper (J1C2) 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 power cord and power up the server.

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

## 9.3 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.

**Note**: System Update and Recovery files are included in the System Update Packages (SUP) posted to Intel®'s web site.

1. Turn off the system and remove power cords.
2. Remove Riser Card Assembly #2.
3. Move the ME FRC UPD Jumper from the default (pins 1 and 2) operating position to the Force Update position (pins 2 and 3).
4. Re-attach system power cords.
5. Power on the system.
   **Note**: System Fans will boost and the BIOS Error Manager should report an 83A0 error code (ME in recovery mode).
6. Boot to the EFI shell and update the ME firmware using the "MEComplete####.cap" file (where #### = ME revision number) using the following command: iflash32 /u /ni MEComplete####.cap.
7. When update has successfully completed, power off system.
8. Remove AC power cords.
9. Move ME FRC UPD jumper back to the default position.
   **Note**: If the ME FRC UPD jumper is moved with AC power applied, the ME will not operate properly. The system will need have the AC power cords removed, wait for at least 10 seconds and then reinstalled to ensure proper operation.
10. Install the PCI Riser.
11. Install the AC power cords.
12. Power on the system.

## 9.4 BIOS Recovery Jumper

When the BIOS Recovery jumper block is moved from its default pin position, the system will boot into a BIOS Recovery Mode. It is used when the system BIOS has become corrupted and is non-functional, requiring a new BIOS image to be loaded on to the server board.
**Note**: The BIOS Recovery jumper is ONLY used to re-install a BIOS image in the event the BIOS has become corrupted. This jumper is NOT used when the BIOS is operating normally and you need to update the BIOS from one version to another.

The following procedure boots the recovery BIOS and flashes the normal BIOS:

1. Turn off the system power.
2. Move the BIOS recovery jumper to the recovery state.

3. Insert a bootable BIOS recovery media containing the new BIOS image files.
4. Turn on the system power.

The BIOS POST screen will appear displaying the progress, and the system will boot to the EFI shell. The EFI shell then executes the Startup.nsh batch file to start the flash update process.

The user should then switch off the power and return the recovery jumper to its normal position.

The user should not interrupt the BIOS POST on the first boot after recovery.

When the flash update completes:

1. Remove the recovery media.
2. Turn off the system power.
3. Restore the jumper to its original position.
4. Turn on the system power.
5. Re-flash any custom blocks, such as user binary or language blocks.

The system should now boot using the updated system BIOS.

# 10 Intel® Light Guided Diagnostics

Both server boards have several on-board 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.

## 10.1 5 V Stand-by LED

Several server management features of these server boards require a 5V stand-by voltage supplied from the power supply. The features and components that require this voltage that must be present when the system is "Off" include the Integrated BMC, on-board NICs, and optional Intel® RMM4/RMM4 Lite connector with Intel® RMM4/RMM4 Lite installed.

The LED is located near the main power connector on the lower left corner of the server board and is labeled "5VSB_LED". It is illuminated when AC power is applied to the platform and 5V stand-by voltage is supplied to the server board by the power supply.

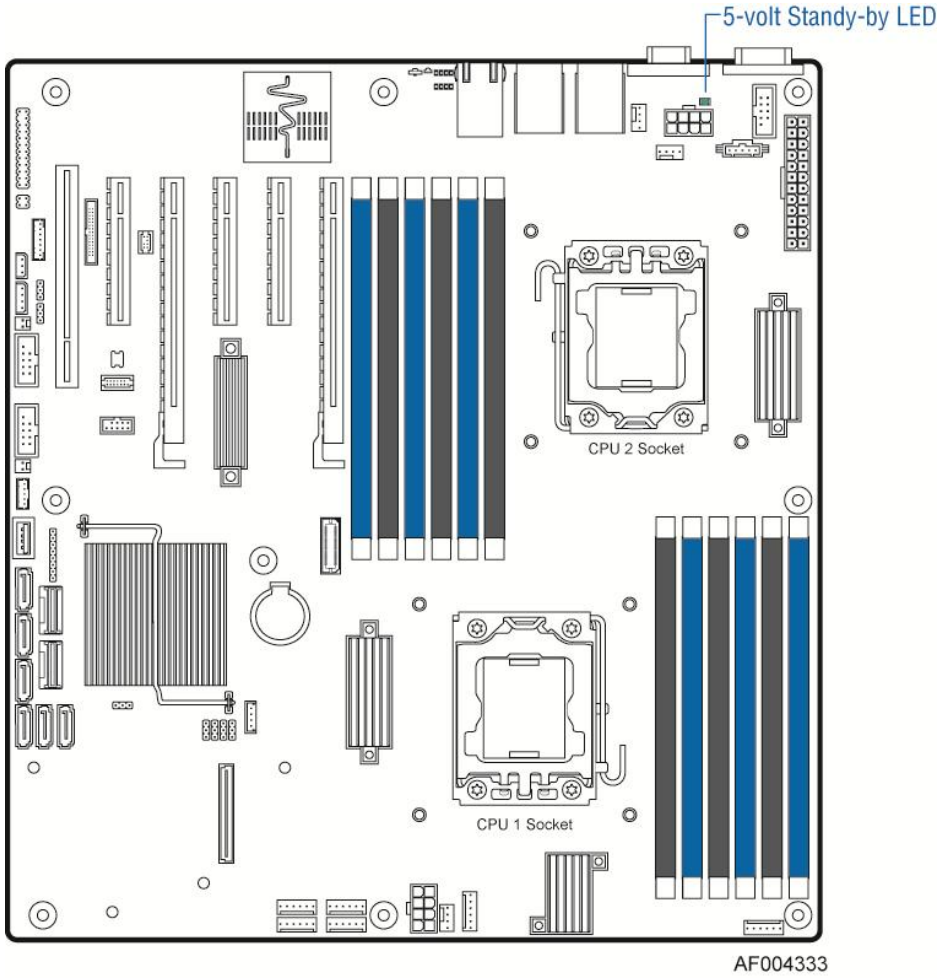


**Figure 25. 5 V Stand-by Status LED Location**

## 10.2 Fan Fault LEDs

Fan fault LEDs are present for the two CPU fans and are located near each CPU fan header.

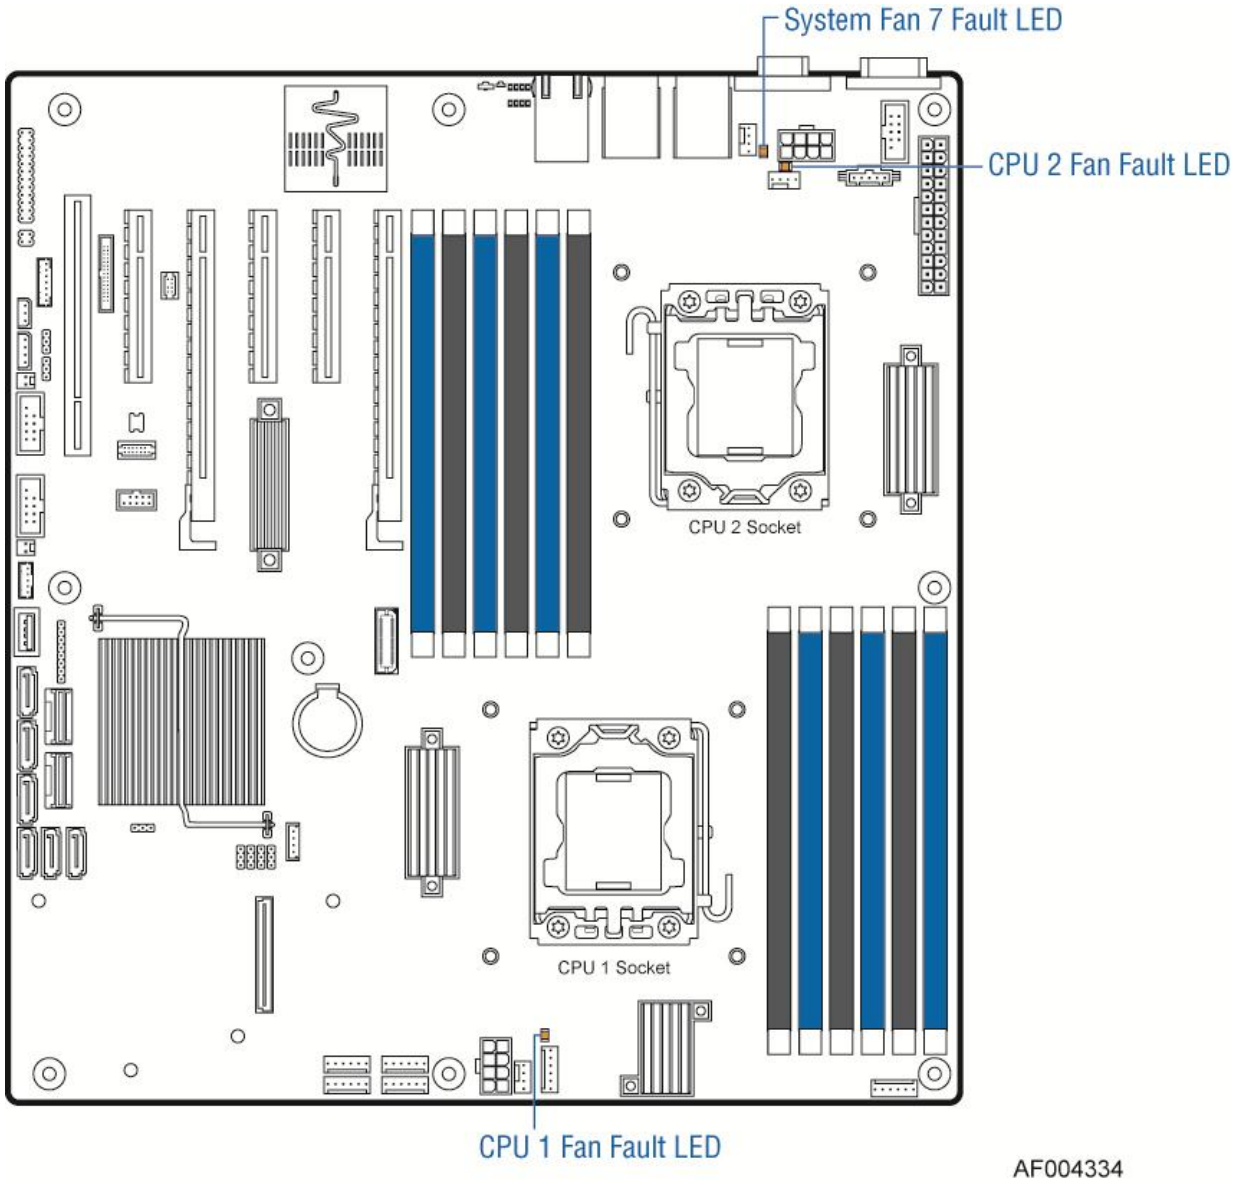


**Figure 26. Fan Fault LED Locations**

## 10.3 CPU Fault LEDs

CPU fault LEDs are present for the two CPUs and are located near each CPU fan header.

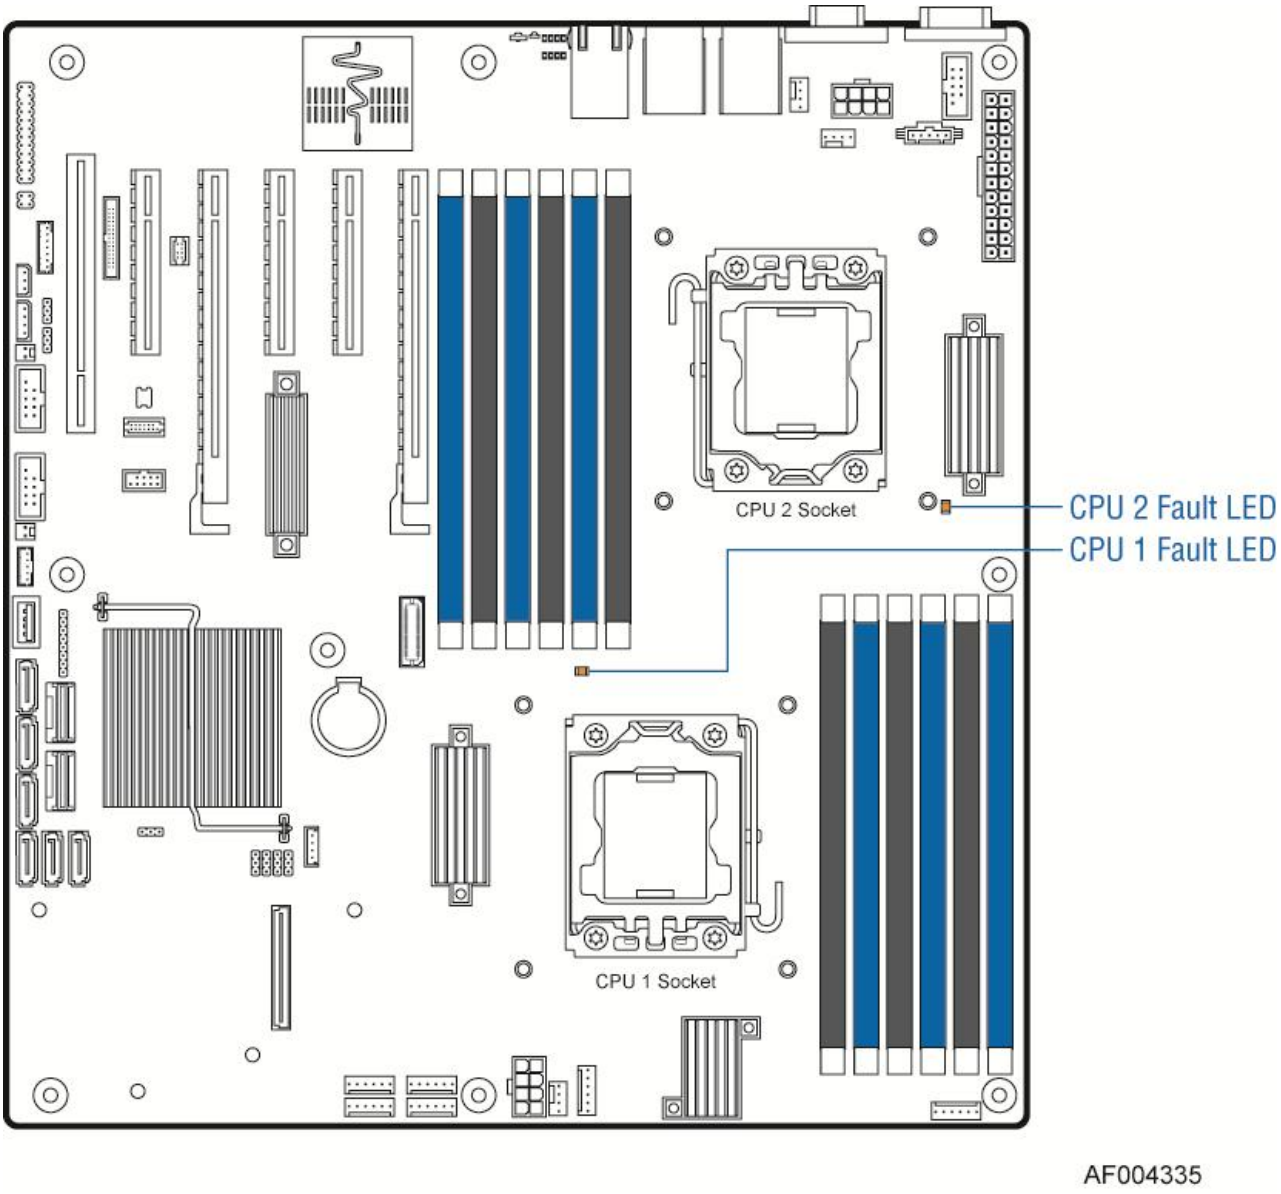


**Figure 27. CPU Fault LED Locations**

## 10.4 System Status LED

The server boards provide a System Status LED. The following figures show the location of the LED.

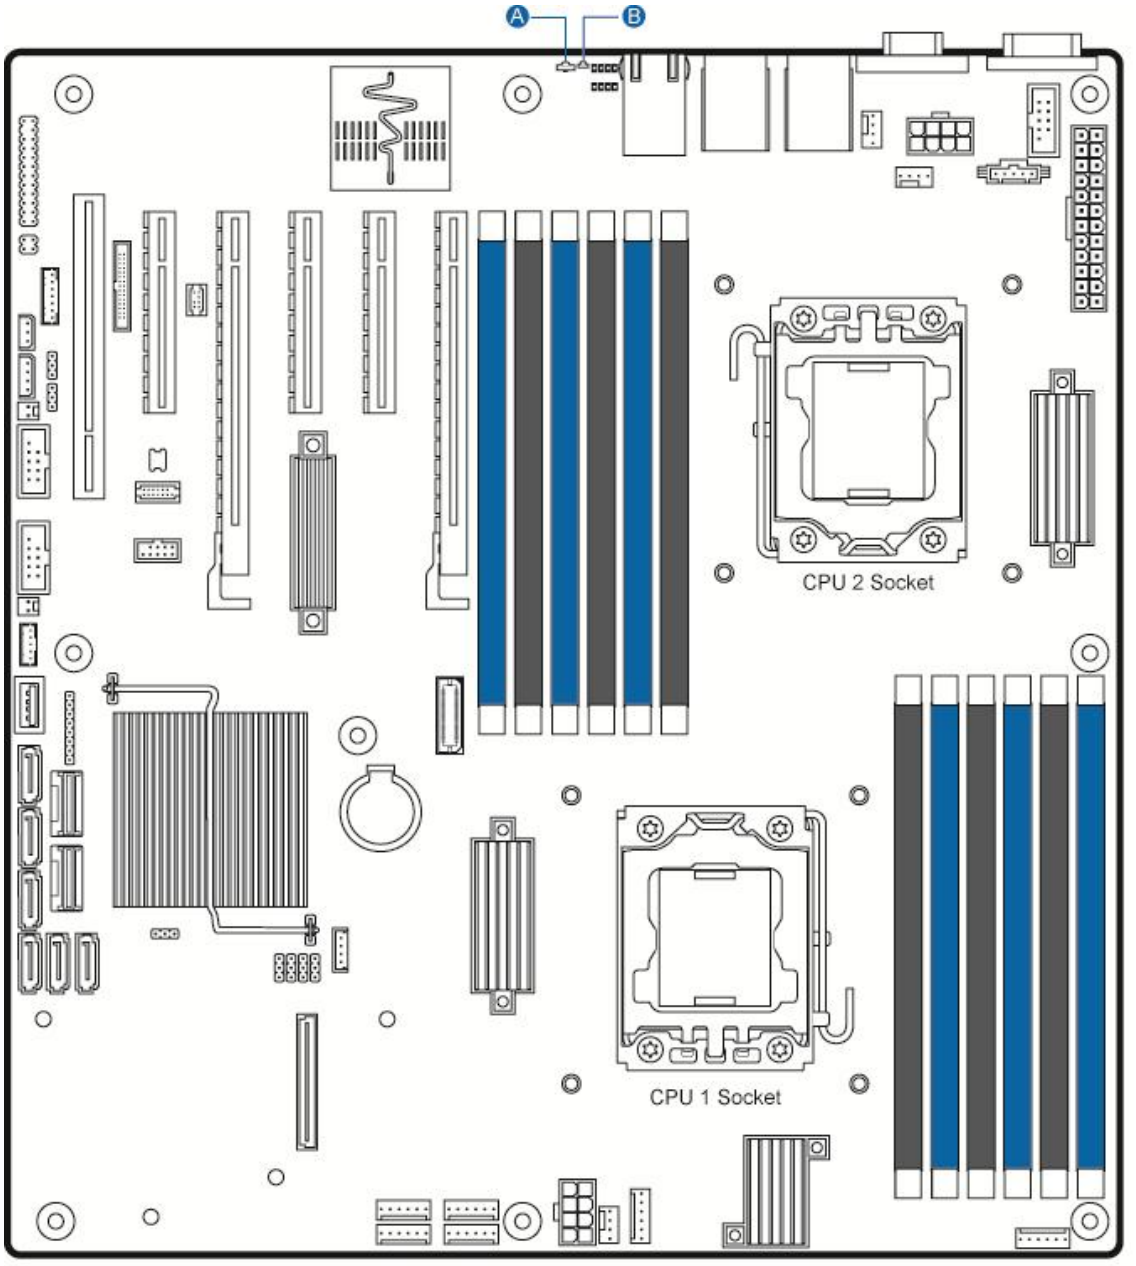


**Figure 28. System Status LED Location**

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

**Table 46. System Status LED**

| Color | State | Criticality | Description |
|---|---|---|---|
| Off | System is not operating | Not ready | 1. System is powered off (AC and/or DC).<br>2. System is in EuP Lot6 Off Mode.<br>3. System is in S5 Soft-Off State.<br>4. System is in S4 Hibernate Sleep State. |
| Green | Solid on | Ok | Indicates that the System is running (in S0 State) and its status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running. |
| Green | ~1 Hz blink | Degraded - system is operating in a degraded state although still functional, or system is operating in a redundant state but with an impending failure warning | System degraded:<br>▪ Redundancy loss, such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities.<br>▪ Fan warning or failure when the number of fully operational fans is more than minimum number needed to cool the system.<br>▪ 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.<br>▪ Power supply predictive failure occurred while redundant power supply configuration was present.<br>▪ Unable to use all of the installed memory (one or more DIMMs failed/disabled but functional memory remains available)<br>▪ 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.<br>▪ Uncorrectable memory error has occurred in memory Mirroring Mode, causing Loss of Redundancy.<br>▪ Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in fully redundant RAS Mirroring Mode.<br>▪ Battery failure.<br>▪ 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.<br>▪ 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.<br>▪ BMC Watchdog has reset the BMC.<br>▪ Power Unit sensor offset for configuration error is asserted.<br>▪ HDD HSC is off-line or degraded. |
| Amber | ~1 Hz blink | Non-critical - System is operating in a degraded state with an impending failure warning, although still functioning | Non-fatal alarm – system is likely to fail:<br>▪ 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.<br>▪ VRD Hot asserted.<br>▪ Minimum number of fans to cool the system not present or failed<br>▪ Hard drive fault |

| Color | State | Criticality | Description |
|---|---|---|---|
| | | | ▪ Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present)<br>▪ In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window<br>▪ Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in a non-redundant mode |
| Amber | Solid on | Critical, non-recoverable – System is halted | Fatal alarm – system has failed or shutdown:<br>▪ CPU CATERR signal asserted<br>▪ MSID mismatch detected (CATERR also asserts for this case).<br>▪ CPU 1 is missing<br>▪ CPU Thermal Trip<br>▪ No power good – power fault<br>▪ DIMM failure when there is only 1 DIMM present and hence no good memory present.<br>▪ Runtime memory uncorrectable error in non-redundant mode.<br>▪ DIMM Thermal Trip or equivalent<br>▪ SSB Thermal Trip or equivalent<br>▪ CPU ERR2 signal asserted<br>▪ BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition)<br>▪ Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition)<br>▪ 240VA fault<br>▪ Fatal Error in processor initialization:<br>▪ Processor family not identical<br>▪ Processor model not identical<br>▪ Processor core/thread counts not identical<br>▪ Processor cache size not identical<br>▪ Unable to synchronize processor frequency<br>▪ Unable to synchronize QPI link frequency |

## 10.5 DIMM Fault LEDs

Each DIMM slot has a DIMM Fault LED near the DIMM slot.

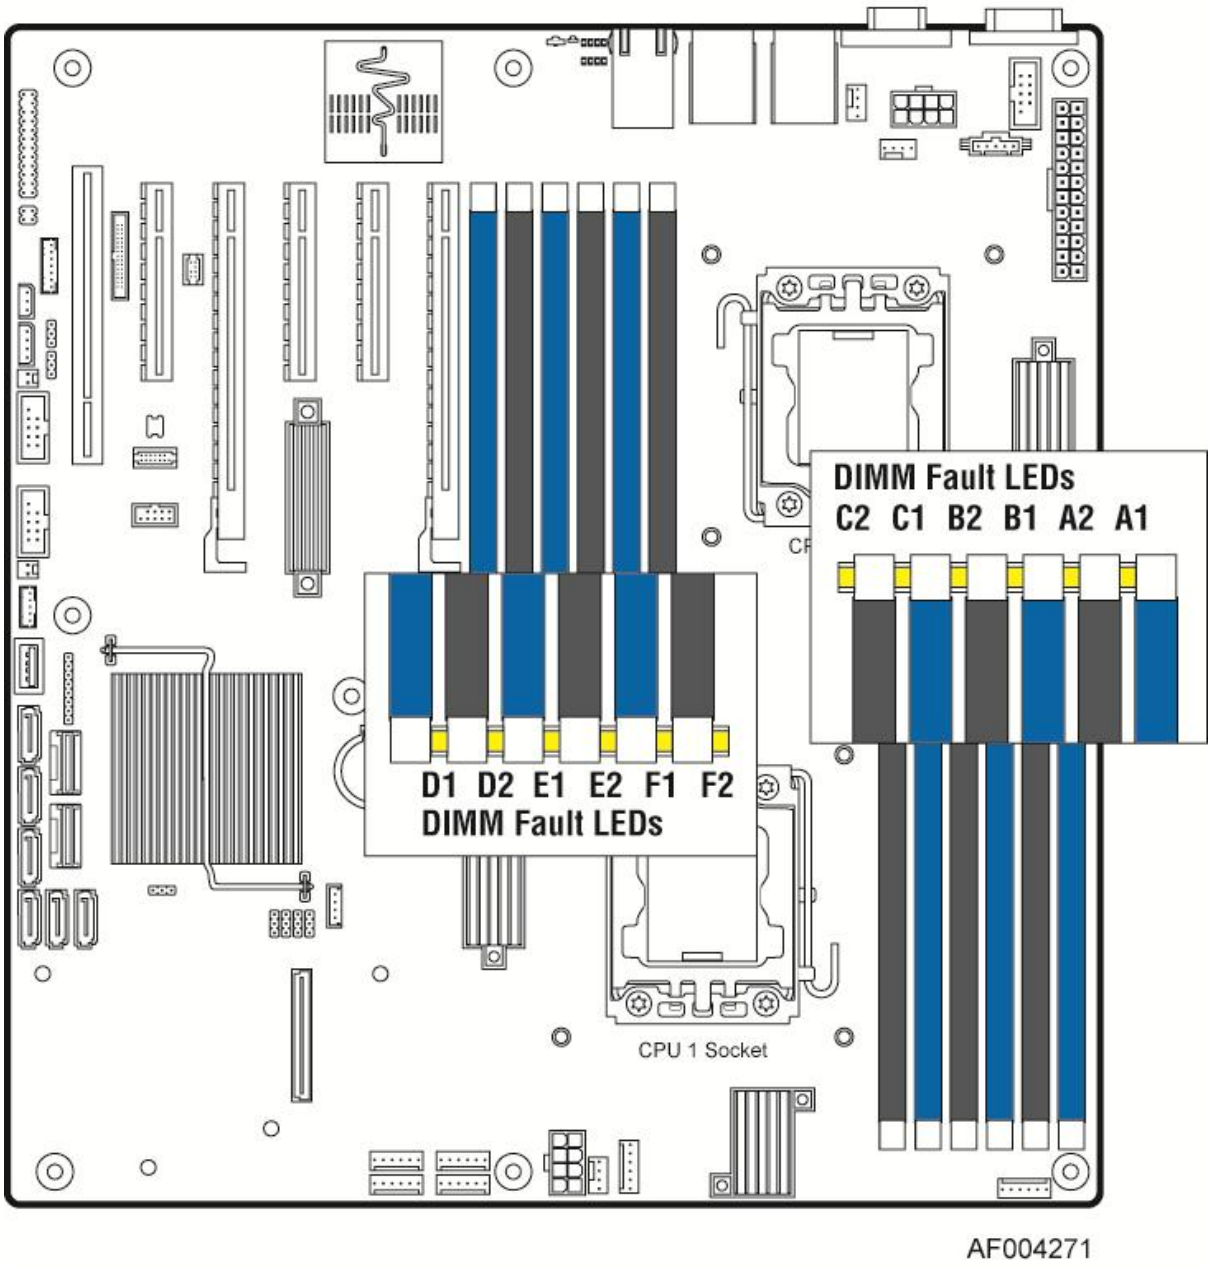


**Figure 29. DIMM Fault LEDs Location**

## 10.6 BMC Boot/Reset Status LED Indicators

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

**Table 47. BMC Boot/Reset Status LED Indicators**

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

## 10.7 Post Code Diagnostic LEDs

Eight amber POST code diagnostic LEDs are located on the back edge of the server boards in the rear I/O area of the server boards by the serial A connector.

During the system boot process, the BIOS executes a number of platform configuration processes, each of which is assigned a specific hex POST code number. As each configuration routine is started, the BIOS displays the given POST code to the POST code diagnostic LEDs on the back edge of the server 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. See Appendix C for a complete description of how these LEDs are read and a list of all supported POST codes.

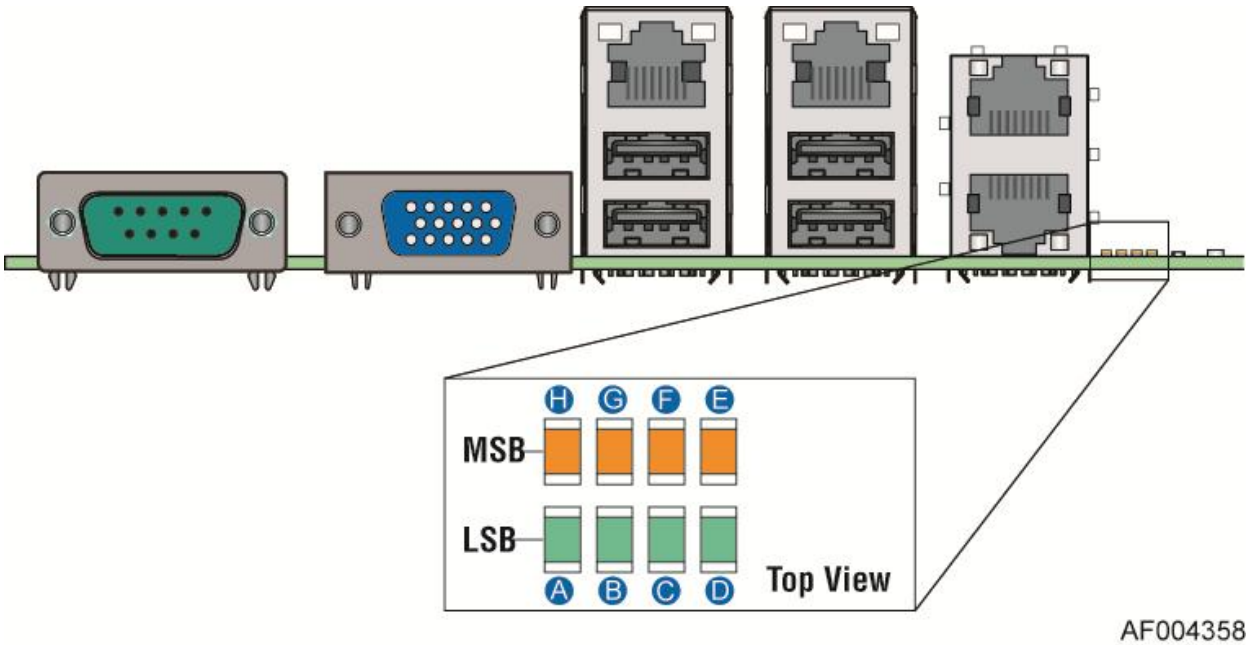


| | |
|---|---|
| A. Diagnostic LED #1 (LSB LED) | E. Diagnostic LED #5 |
| B. Diagnostic LED #2 | F. Diagnostic LED #6 |

| | |
|---|---|
| C. Diagnostic LED #3 | G. Diagnostic LED #7 |
| D. Diagnostic LED #4 | H. Diagnostic LED #8 (MSB LED) |

**Figure 30. POST Code Diagnostic LED Locations**

# 11 Environmental Limits Specification

The following table defines the Intel® Server Board S2400GP operating and non-operating environmental limits. Operation of the Intel® Server Board S2400GP 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 48. Server Board Design Specifications**

| | |
|---|---|
| Operating Temperature | 0° C to 55° C [1] (32° F to 131° F) |
| Non-Operating Temperature | -40° C to 70° C (-40° F to 158° F) |
| DC Voltage | ± 5% of all nominal voltages |
| Shock (Unpackaged) | Trapezoidal, 35g, 170 inches/sec |
| Shock (Packaged)<br>< 20 pounds<br>20 to < 40 pounds<br>40 to < 80 pounds<br>80 to < 100 pounds<br>100 to < 120 pounds<br>120 pounds | <br>36 inches<br>30 inches<br>24 inches<br>18 inches<br>12 inches<br>9 inches |
| Vibration (Unpackaged) | 5 Hz to 500 Hz 3.13 g RMS random |

**Note:**

1. 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.

**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.

## 11.1 Processor Thermal Design Power (TDP) Support

To allow optimal operation and long-term reliability of Intel® processor-based systems, the processor must remain within the defined minimum and maximum case temperature (TCASE) specifications. Thermal solutions not designed to provide sufficient thermal capability may affect the long-term reliability of the processor and system. The server board is designed to support the Intel® Xeon® Processor E5-2400 product family TDP guidelines up to and including 95W.

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

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

## 11.2 MTBF

The following is the calculated Mean Time Between Failures (MTBF) 30 degree 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.

- Calculate standard: Telcordia* issue 2
- Calculate Method: Method I-D
- Temperature = 40 degree C
- Environment = GB, GC – Ground Benign, Controlled
- Model = Serial
- Duty cycle = 100%
- Component Quality: Level II
- Adhere to De-rating data

**Table 49. MTBF Estimate**

| Assembly Name | Temperature (Degree C) | MTBF (hours) |
|---|---|---|
| Intel® Server Board S2400GP | 40 | 226,484 |

## 11.3 Server Board Power Distribution

This section provides power supply design guidelines for a system using the Intel® Server Board S2400GP. The following diagram shows the power distribution implemented on this server board. For power supply data, please refer to the chapter that describes the power system options including 460W, 550W, 750W or 1200W power supply. Please note the intent of 460W/550W/750W/1200W power supply data is to provide customers with a guide to assist in defining and/or selecting a power supply for custom server platform designs that utilize the server boards detailed in this document.

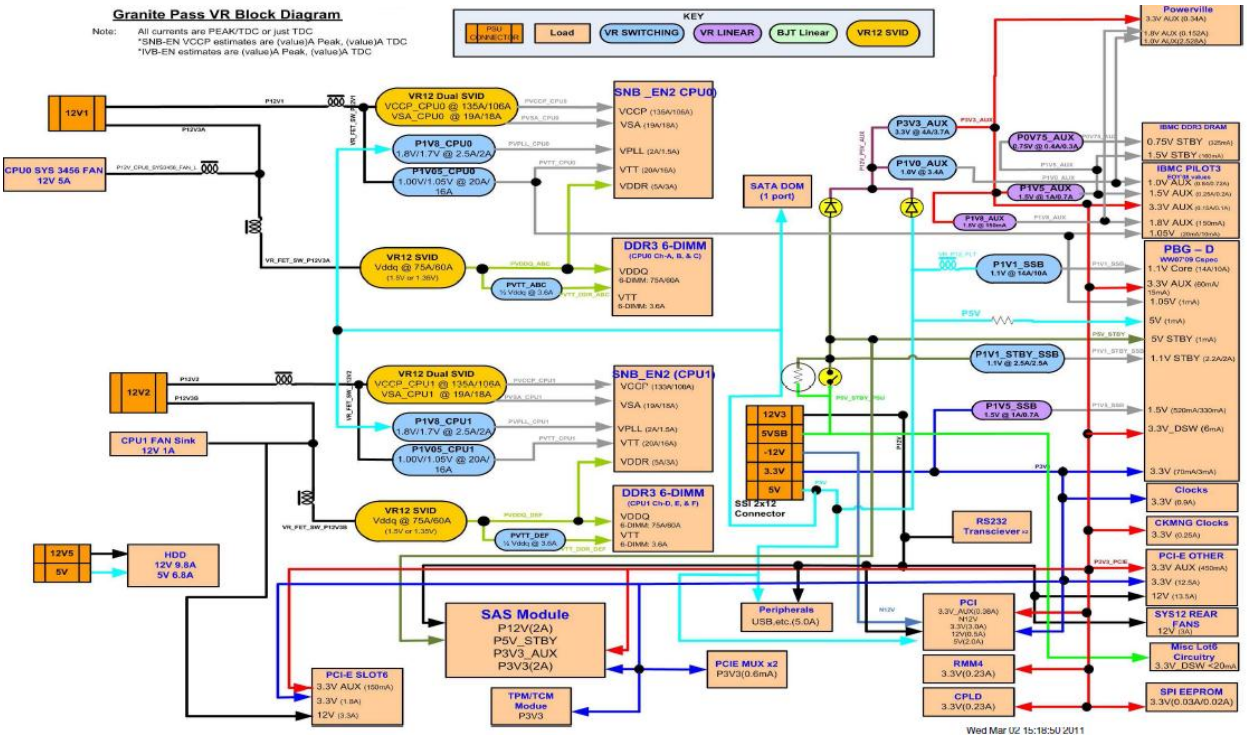


**Figure 31. Power Distribution Block Diagram**

# Appendix A: Integration and Usage Tips

- When adding or removing components or peripherals from the server board, you must remove AC power cord. 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 Intel® Xeon® Processor E5-2400 product family with a Thermal Design Power (TDP) of up to and including 95 Watts. Previous generation Intel® Xeon® processors are not supported.
- You must install processors in order. CPU 1 is located near the back edge of the server board and must be populated to operate the board.
- 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.
- Only Registered DDR3 DIMMs (RDIMMs) and Unbuffered DDR3 DIMMs (UDIMMs) are supported on this server board. Mixing of RDIMMs and UDIMMs is not supported.
- For the best performance, you should balance the number of DDR3 DIMMs installed 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, you should install DIMMs in DIMM sockets A1 and D1. A Six-DIMM configuration (DIMM Sockets A1, B1, C1, D1, E1, and F1) performs better than a Three-DIMM configuration (DIMM Sockets A1, B1, and C1).
- The Intel® RMM4/RMM4 Lite connectors are not compatible with the previous Intel® Remote Management Modules
- Clear CMOS with the AC power cord plugged in. Removing 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 BMC functionality is disabled with the Force BMC Update jumper (J1C2) set to the “enabled” position (pins 2-3). You should never run the server with the Force BMC Update jumper set in this position and should only use the jumper in this position when the standard firmware update process fails. This jumper must remain in the default (disabled) position (pins 1-2) when the server is running normally.
- This server board no longer supports the Rolling BIOS (two BIOS banks). It implements the BIOS Recovery mechanism instead.
- When performing a normal BIOS update procedure, you must set the BIOS Recovery jumper (J3H6) to its default position (pins 1-2).

# Appendix B: BMC Sensor Tables

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

- **Sensor Type**

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

- **Event/Reading Type**

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

- **Event Offset/Triggers**

  Event Thresholds are event-generating thresholds for threshold type sensors.

  - [u,l][nr,c,nc]: upper nonrecoverable, upper critical, upper noncritical, lower nonrecoverable, lower critical, lower noncritical
  - uc, lc: upper critical, lower critical

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

- **Assertion/De-assertion Enables**

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

  - As: Assertions
  - De: De-assertion

- **Readable Value/Offsets**

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

- **Event Data**

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

  - R: Reading value
  - T: Threshold value

- **Rearm Sensors**

  The rearm is a request for the event status of a sensor to be rechecked and updated upon a transition between good and bad states. You can rearm the sensors 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

- **Default Hysteresis**

  The hysteresis setting applies to all thresholds of the sensor. This column provides the count of hysteresis for the sensor, which is one or two (positive or negative hysteresis).

- **Criticality**

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

- **Standby**

  Some sensors operate on standby power. You can access these sensors and/or generate events when the main (system) power is off but AC power is present.

**Table 50. Integrated BMC Core Sensors**

| Sensor Name3 | 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 Stat | 01h | All | Power Unit 09h | Sensor Specific 6Fh | 00 - Power down<br>04 - A/C lost | OK | As and De | – | Trig Offset | A | X |
| – | – | – | – | – | 05 - Soft power control failure<br>06 - Power unit failure | Fatal | – | – | – | – | – |
| Power Redundancy4 | 02h | Chassis-specific | Power Unit 09h | Generic 0Bh | 00 - Fully Redundant | OK | As and De | – | Trig Offset | A | |
| – | – | – | – | – | 01 - Redundancy lost | Degraded | – | – | – | – | – |
| – | – | – | – | – | 02 - Redundancy degraded | Degraded | – | – | – | – | – |
| – | – | – | – | – | 03 - Non-red: suff res from redund | Degraded | – | – | – | – | – |
| – | – | – | – | – | 04 - Non-red: suff from insuff | Degraded | – | – | – | – | – |
| – | – | – | – | – | 05 - Non-red: insufficient | Fatal | – | – | – | – | – |
| – | – | – | – | – | 06 - Redun degrade from fully redun | Degraded | – | – | – | – | – |
| – | – | – | – | – | 07 - Redun degrade from non-redundant | Degraded | – | – | – | – | – |
| IPMI Watchdog | 03h | All | Watchdog 2 23h | Sensor Specific 6Fh | 00 - Timer expired, status only<br>01 - Hard reset<br>02 - Power down<br>03 - Power cycle<br>08 - Timer interrupt | OK | As | – | Trig Offset | A | X |

| Sensor Name3 | 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 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Physical Scrty | 04h | Chassis Intrusion is chassis-specific | Physical Security 05h | Sensor Specific 6Fh | 00 - Chassis intrusion | OK | As and De | – | Trig Offset | A | X |
| | | | | | 04 - LAN least lost | Degraded | | | | | |
| FP Interrupt (NMI) | 05h | All | Critical Interrupt 13h | Sensor Specific 6Fh | 00 - Front panel NMI/diagnostic interrupt | OK | As | – | Trig Offset | A | – |
| SMI Timeout | 06h | All | SMI Timeout F3h | Digital Discrete 03h | 01 – State asserted | Fatal | As and De | – | Trig Offset | A | – |
| System Event Log | 07h | All | Event Logging Disabled 10h | Sensor Specific 6Fh | 02 - Log area reset/cleared | OK | As | – | Trig Offset | A | X |
| BB +1.1V IOH | 10h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +1.1V P1 Vccp | 11h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +1.1V P2 Vccp | 12h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +1.5V P1 DDR3 | 13h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +1.5V P2 DDR3 | 14h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +1.8V AUX | 15h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |

| Sensor Name3 | 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 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| BB +3.3V | 16h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +3.3V STBY | 17h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| BB Vbat | 18h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +5.0V | 19h | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB +5.0V STBY | 1Ah | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| BB +12.0V | 1Bh | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| BB -12.0V | 1Ch | All | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Baseboard Temp | 20h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| Front Panel Temp | 21h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| IOH Temp | 22h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| MEM P1 THRM MRGN | 23h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| MEM P2 THRM MRGN | 24h | Dual processor only | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |

| Sensor Name3 | 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 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Fan Tach Sensors | 30h–39h | Chassis-specific | Fan 04h | Threshold 01h | [l] [c,nc] | nc = Degraded c = Non-fatal2 | As and De | Analog | R, T | M | – |
| Fan Present Sensors | 40h–45h | Chassis-specific | Fan 04h | Generic 08h | 01 - Device inserted | OK | As and De | – | Triggered Offset | Auto | – |
| Fan Redundancy 4 | 46h | 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-red: suff res from redund | Degraded | – | – | – | – | – |
| – | – | – | – | – | 04 - Non-red: suff from insuff | Degraded | – | – | – | – | – |
| – | – | – | – | – | 05 - Non-red: insufficient | Non-fatal | – | – | – | – | – |
| – | – | – | – | – | 06 - Redun degrade from full | Degraded | – | – | – | – | – |
| – | – | – | – | – | 07 - Redun degrade from non-redundant | Degraded | – | – | – | – | – |
| – | – | – | – | – | 07 - Redun degrade from non-redundant | Degraded | – | – | – | – | – |
| PS1 Status 4 | 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 | – | – | – | – | – |

| Sensor Name3 | Sensor # | Platform Applicability | Sensor Type | Event/Reading Type | Event Offset Triggers | Contrib. To System Status | Assert/De-assert | Readable Value/Offsets | Event Data | Rearm | Stand-by |
|---|---|---|---|---|---|---|---|---|---|---|---|
| PS2 Status 4 | 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 | – | – | – | – | – |
| PS1 Power In | 52h | Chassis-specific | Power Supply 08h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| PS2 Power In | 53h | Chassis-specific | Power Supply 08h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| PS1 Current Out | 54h | Chassis-specific | Power Supply 08h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| PS2 Current Out | 55h | Chassis-specific | Power Supply 08h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| PS1 Temperature | 56h | Chassis-specific | Temperature 01h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| PS2 Temperature | 57h | Chassis-specific | Temperature | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| P1 Status | 60h | All | Processor 07h | Sensor Specific 6Fh | 01 - Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| P2 Status | 61h | Dual processor only | Processor 07h | Sensor Specific 6Fh | 01- Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| P1 Therm Margin | 62h | All | Temperature 01h | Threshold 01h | – | – | – | Analog | – | – | – |

| Sensor Name3 | 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 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| P2 Therm Margin | 63h | Dual processor only | Temperature 01h | Threshold 01h | – | – | – | Analog | – | – | – |
| P1 Therm Ctrl % | 64h | All | Temperature 01h | Threshold 01h | [u] [c] | Non-fatal | As and De | Analog | Trig Offset | A | – |
| P2 Therm Ctrl % | 65h | Dual processor only | Temperature 01h | Threshold 01h | [u] [c] | Non-fatal | As and De | Analog | Trig Offset | A | – |
| P1 VRD Temp | 66h | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Fatal | As and De | – | Trig Offset | M | – |
| P2 VRD Temp | 67h | Dual processor only | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Fatal | As and De | – | Trig Offset | M | – |
| CATERR | 68h | All | Processor 07h | Digital Discrete 03h | 01 – State Asserted | Non-fatal | As and De | – | Trig Offset | M | – |
| CPU Missing | 69h | All | Processor 07h | Digital Discrete 03h | 01 – State Asserted | Non-fatal | As and De | – | Trig Offset | M | – |
| IOH Thermal Trip | 6Ah | All | Temperature 01h | Digital Discrete 03h | 01 – State Asserted | Fatal | As and De | – | Trig Offset | M | – |

**Notes:**

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

# Appendix C: POST Code Diagnostic LED Decoder

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

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

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

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

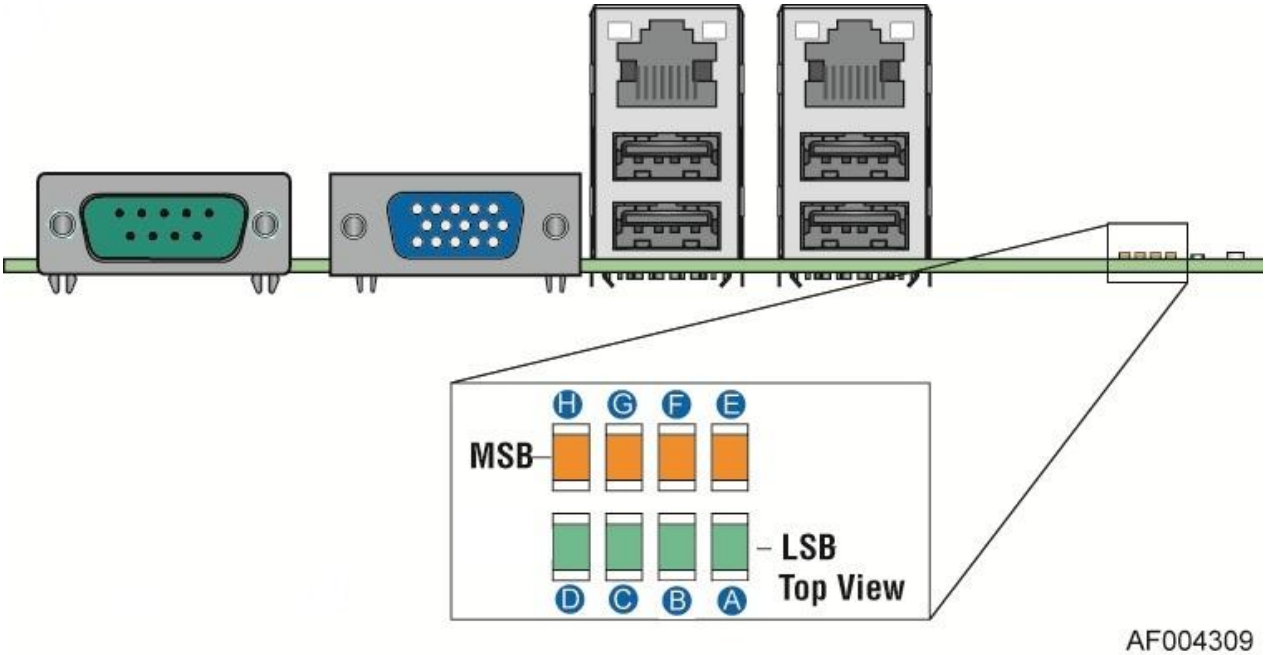


| | |
|---|---|
| A. Diagnostic LED #1(LSB LED) | E. Diagnostic LED #5 |
| B. Diagnostic LED #2 | F. Diagnostic LED #6 |
| C. Diagnostic LED #3 | G. Diagnostic LED #7 |
| D. Diagnostic LED #4 | H. Diagnostic LED #8 (MSB LED) |

**Figure 32. POST Code Diagnostic LEDs**

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

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

**Table 51. 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 52. POST Progress Codes**

| Checkpoint | Diagnostic LED Decoder 1 = LED On, 0 = LED Off | | | | | | | | Description |
|---|---|---|---|---|---|---|---|---|---|
| | Upper Nibble | | | | Lower Nibble | | | | |
| | MSB | | | | | | | LSB | |
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED # | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | |
| SEC Phase | | | | | | | | | |
| **01h** | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | First POST code after CPU reset |
| **02h** | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | Microcode load begin |
| **03h** | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 1 | CRAM initialization begin |
| **04h** | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | Pei Cache When Disabled |
| **05h** | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 | SEC Core At Power On Begin. |
| **06h** | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 0 | Early CPU initialization during Sec Phase. |
| **07h** | 0 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | Early SB initialization during Sec Phase. |
| **08h** | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | Early NB initialization during Sec Phase. |
| **09h** | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 1 | End Of Sec Phase. |
| **0Eh** | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 0 | Microcode Not Found. |
| **0Fh** | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 | Microcode Not Loaded. |
| PEI Phase | | | | | | | | | |
| **10h** | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | PEI Core |
| **11h** | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | CPU PEIM |
| **15h** | 0 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | NB PEIM |
| **19h** | 0 | 0 | 0 | 1 | 1 | 0 | 0 | 1 | SB PEIM |
| MRC Process Codes – MRC Progress Code Sequence is executed - See Table 63 | | | | | | | | | |
| PEI Phase continued… | | | | | | | | | |
| **31h** | 0 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | Memory Installed |
| **32h** | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | CPU PEIM (Cpu Init) |
| **33h** | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | CPU PEIM (Cache Init) |
| **34h** | 0 | 0 | 1 | 1 | 0 | 1 | 0 | 0 | CPU PEIM (BSP Select) |
| **35h** | 0 | 0 | 1 | 1 | 0 | 1 | 0 | 1 | CPU PEIM (AP Init) |
| **36h** | 0 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | CPU PEIM (CPU SMM Init) |
| **4Fh** | 0 | 1 | 0 | 0 | 1 | 1 | 1 | 1 | Dxe IPL started |
| DXE Phase | | | | | | | | | |
| **60h** | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 0 | DXE Core started |
| **61h** | 0 | 1 | 1 | 0 | 0 | 0 | 0 | 1 | DXE NVRAM Init |

| Checkpoint | Diagnostic LED Decoder<br>1 = LED On, 0 = LED Off<br>Upper Nibble<br>MSB | | | | Lower Nibble | | | LSB | |
|---|---|---|---|---|---|---|---|---|---|
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED # | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | Description |
| **62h** | 0 | 1 | 1 | 0 | 0 | 0 | 1 | 0 | SB RUN Init |
| **63h** | 0 | 1 | 1 | 0 | 0 | 0 | 1 | 1 | Dxe CPU Init |
| **68h** | 0 | 1 | 1 | 0 | 1 | 0 | 0 | 0 | DXE PCI Host Bridge Init |
| **69h** | 0 | 1 | 1 | 0 | 1 | 0 | 0 | 1 | DXE NB Init |
| **6Ah** | 0 | 1 | 1 | 0 | 1 | 0 | 1 | 0 | DXE NB SMM Init |
| **70h** | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | DXE SB Init |
| **71h** | 0 | 1 | 1 | 1 | 0 | 0 | 0 | 1 | DXE SB SMM Init |
| **72h** | 0 | 1 | 1 | 1 | 0 | 0 | 1 | 0 | DXE SB devices Init |
| **78h** | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | DXE ACPI Init |
| **79h** | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | DXE CSM Init |
| **90h** | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | DXE BDS Started |
| **91h** | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | DXE BDS connect drivers |
| **92h** | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 0 | DXE PCI Bus begin |
| **93h** | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 1 | DXE PCI Bus HPC Init |
| **94h** | 1 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | DXE PCI Bus enumeration |
| **95h** | 1 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | DXE PCI Bus resource requested |
| **96h** | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | DXE PCI Bus assign resource |
| **97h** | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | DXE CON_OUT connect |
| **98h** | 1 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | DXE CON_IN connect |
| **99h** | 1 | 0 | 0 | 1 | 1 | 0 | 0 | 1 | DXE SIO Init |
| **9Ah** | 1 | 0 | 0 | 1 | 1 | 0 | 1 | 0 | DXE USB start |
| **9Bh** | 1 | 0 | 0 | 1 | 1 | 0 | 1 | 1 | DXE USB reset |
| **9Ch** | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 0 | DXE USB detect |
| **9Dh** | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | DXE USB enable |
| **A1h** | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | DXE IDE begin |
| **A2h** | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | DXE IDE reset |
| **A3h** | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 1 | DXE IDE detect |
| **A4h** | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 0 | DXE IDE enable |
| **A5h** | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | DXE SCSI begin |
| **A6h** | 1 | 0 | 1 | 0 | 0 | 1 | 1 | 0 | DXE SCSI reset |
| **A7h** | 1 | 0 | 1 | 0 | 0 | 1 | 1 | 1 | DXE SCSI detect |
| **A8h** | 1 | 0 | 1 | 0 | 1 | 0 | 0 | 0 | DXE SCSI enable |
| **A9h** | 1 | 0 | 1 | 0 | 1 | 0 | 0 | 1 | DXE verifying SETUP password |
| **ABh** | 1 | 0 | 1 | 0 | 1 | 0 | 1 | 1 | DXE SETUP start |
| **ACh** | 1 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | DXE SETUP input wait |
| **ADh** | 1 | 0 | 1 | 0 | 1 | 1 | 0 | 1 | DXE Ready to Boot |
| **AEh** | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | DXE Legacy Boot |
| **AFh** | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | DXE Exit Boot Services |
| **B0h** | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | RT Set Virtual Address Map Begin |
| **B1h** | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | RT Set Virtual Address Map End |
| **B2h** | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | DXE Legacy Option ROM init |
| **B3h** | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | DXE Reset system |
| **B4h** | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 0 | DXE USB Hot plug |
| **B5h** | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 1 | DXE PCI BUS Hot plug |
| **B6h** | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | DXE NVRAM cleanup |
| **B7h** | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 1 | DXE Configuration Reset |
| **00h** | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | INT19 |
| S3 Resume | | | | | | | | | |
| **E0h** | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | S3 Resume PEIM (S3 started) |
| **E1h** | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | S3 Resume PEIM (S3 boot script) |
| **E2h** | 1 | 1 | 0 | 1 | 0 | 0 | 1 | 0 | S3 Resume PEIM (S3 Video Repost) |
| **E3h** | 1 | 1 | 0 | 1 | 0 | 0 | 1 | 1 | S3 Resume PEIM (S3 OS wake) |
| BIOS Recovery | | | | | | | | | |

| Checkpoint | Diagnostic LED Decoder 1 = LED On, 0 = LED Off | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | Upper Nibble | | | | Lower Nibble | | | | |
| | MSB | | | | | | | LSB | |
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED # | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | Description |
| **F0h** | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 0 | PEIM which detected forced Recovery condition |
| **F1h** | 1 | 1 | 1 | 1 | 0 | 0 | 0 | 1 | PEIM which detected User Recovery condition |
| **F2h** | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 0 | Recovery PEIM (Recovery started) |
| **F3h** | 1 | 1 | 1 | 1 | 0 | 0 | 1 | 1 | Recovery PEIM (Capsule found) |
| **F4h** | 1 | 1 | 1 | 1 | 0 | 1 | 0 | 0 | Recovery PEIM (Capsule loaded) |
| | | | | | | | | | |

**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 53. MRC Progress Codes**

| Checkpoint | Diagnostic LED Decoder 1 = LED On, 0 = LED Off | | | | | | | | Description |
|---|---|---|---|---|---|---|---|---|---|
| | Upper Nibble | | | | Lower Nibble | | | | |
| | MSB | | | | | | | LSB | |
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | |
| **MRC Progress Codes** | | | | | | | | | |
| B0h | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | Detect DIMM population |
| B1h | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | Set DDR3 frequency |
| B2h | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | Gather remaining SPD data |
| B3h | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | Program registers on the memory controller level |
| B4h | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 0 | Evaluate RAS modes and save rank information |
| B5h | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 1 | Program registers on the channel level |
| B6h | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | Perform the JEDEC defined initialization sequence |
| B7h | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 1 | Train DDR3 ranks |
| B8h | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 0 | Initialize CLTT/OLTT |
| B9h | 1 | 0 | 1 | 1 | 1 | 0 | 0 | 1 | Hardware memory test and init |
| BAh | 1 | 0 | 1 | 1 | 1 | 0 | 1 | 0 | Execute software memory init |
| BBh | 1 | 0 | 1 | 1 | 1 | 0 | 1 | 1 | Program memory map and interleaving |
| BCh | 1 | 0 | 1 | 1 | 1 | 1 | 0 | 0 | Program RAS configuration |
| BFh | 1 | 0 | 1 | 1 | 1 | 1 | 1 | 1 | MRC is done |

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 54. POST Progress LED Codes**

| Checkpoint | Diagnostic LED Decoder | | | | | | | | Description |
|---|---|---|---|---|---|---|---|---|---|
| | 1 = LED On, 0 = LED Off | | | | | | | | |
| | Upper Nibble | | | | Lower Nibble | | | | |
| | MSB | | | | | | | LSB | |
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | |
| MRC Fatal Error Codes | | | | | | | | | |
| E8h | 1 | 1 | 1 | 0 | 1 | 0 | 0 | 0 | No usable memory error<br>01h = No memory was detected through SPD read, or invalid config that causes no operable memory.<br>02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error.<br>3h = No memory installed. All channels are disabled. |
| E9h | 1 | 1 | 1 | 0 | 1 | 0 | 0 | 1 | Memory is locked by Intel® Trusted Execution Technology and is inaccessible |
| EAh | 1 | 1 | 1 | 0 | 1 | 0 | 1 | 0 | DDR3 channel training error<br>01h = Error on read DQ/DQS (Data/Data Strobe) init<br>02h = Error on Receive Enable<br>3h = Error on Write Leveling<br>04h = Error on write DQ/DQS (Data/Data Strobe |
| EBh | 1 | 1 | 1 | 0 | 1 | 0 | 1 | 1 | Memory test failure<br>01h = Software memtest failure.<br>02h = Hardware memtest failed.<br>03h = Hardware Memtest failure in Lockstep Channel mode requiring a channel to be disabled. *This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.* |
| EDh | 1 | 1 | 1 | 0 | 1 | 1 | 0 | 1 | DIMM configuration population error<br>01h = Different DIMM types (UDIMM, RDIMM, LRDIMM) are detected installed in the system.<br>02h = Violation of DIMM population rules.<br>03h = The third DIMM slot cannot be populated when QR DIMMs are installed.<br>04h = UDIMMs are not supported in the third DIMM slot.<br>05h = Unsupported DIMM Voltage. |
| EFh | 1 | 1 | 1 | 0 | 1 | 1 | 1 | 1 | Indicates a CLTT table structure error |

# Appendix D: POST Code Errors

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

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

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

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

  Note that for 0048 “Password check failed”, the system halts, and then after the next reset/reboot will displays the error code on the Error Manager screen.
- **Fatal:** The system halts during post at a blank screen with the text **“Unrecoverable fatal error found. System will not boot until the error is resolved”** and **“Press <F2> to enter setup”** The POST Error Pause option setting in the BIOS setup does not have any effect with this class of error.

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

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

**Table 55. 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® 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 |
| 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 Self Test | Major |
| 83A1 | Management Engine (ME) Failed to respond. | Major |
| 84F2 | Baseboard management controller failed to respond | Major |
| 84F3 | Baseboard management controller in update mode | Major |
| 84F4 | Sensor data record empty | Major |
| 84FF | System event log full | Minor |
| 8500 | Memory component could not be configured in the selected RAS mode | Major |
| 8501 | DIMM Population Error | Major |
| 8520 | DIMM_A1 failed test/initialization | Major |
| 8521 | DIMM_A2 failed test/initialization | Major |
| 8522 | DIMM_A3 failed test/initialization | Major |
| 8523 | DIMM_B1 failed test/initialization | Major |
| 8524 | DIMM_B2 failed test/initialization | Major |
| 8525 | DIMM_B3 failed test/initialization | Major |
| 8526 | DIMM_C1 failed test/initialization | Major |
| 8527 | DIMM_C2 failed test/initialization | Major |
| 8528 | DIMM_C3 failed test/initialization | Major |
| 8529 | DIMM_D1 failed test/initialization | Major |
| 852A | DIMM_D2 failed test/initialization | Major |
| 852B | DIMM_D3 failed test/initialization | Major |
| 852C | DIMM_E1 failed test/initialization | Major |
| 852D | DIMM_E2 failed test/initialization | Major |

| Error Code | Error Message | Response |
|---|---|---|
| 852E | DIMM_E3 failed test/initialization | Major |
| 852F | DIMM_F1 failed test/initialization | Major |
| 8530 | DIMM_F2 failed test/initialization | Major |
| 8531 | DIMM_F3 failed test/initialization | Major |
| 8532 | DIMM_G1 failed test/initialization | Major |
| 8533 | DIMM_G2 failed test/initialization | Major |
| 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 |
| 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 |

| Error Code | Error Message | Response |
|---|---|---|
| 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 |
| 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 |

| Error Code | Error Message | Response |
|---|---|---|
| 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 |
| 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 |
| 8606 | NVRAM variable space was corrupted and has been reinitialized | 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 |
| A6A0 | DXE Boot Service driver: Not enough memory available to shadow a LegacyOption ROM | Minor |

## POST Error Beep Codes

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

**Table 56. POST Error Beep Codes**

| Beeps | Error Message | POST Progress Code | Description |
|---|---|---|---|
| 1 | USB device action | NA | Short beep sounded whenever a USB device is discovered in POST, or inserted or removed during runtime |
| 1 long | Intel® TXT security violation | 0xAE, 0xAF | System halted because Intel® Trusted Execution Technology detected a potential violation of system security. |
| 3 | Memory error | See Tables 28 and 29 | System halted because a fatal error related to the memory was detected. |
| 2 | BIOS Recovery started | NA | Recovery boot has been initiated |

| Beeps | Error Message | POST Progress Code | Description |
|---|---|---|---|
| 4 | BIOS Recovery failure | NA | BIOS recovery has failed. This typically happens so quickly after recovery us initiated that it sounds like a 2-4 beep code. |

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

**Table 57. Integrated BMC Beep Codes**

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

# Appendix E: Supported Intel® Server Chassis

The Intel® Server Board S2400GP requires a passive processor heat sink solution when integrated in the Intel® pedestal server chassis Intel® Server Chassis P4000 series. The Intel® Server Board S2400GP supports up to 95W TDP Intel® Xeon® Processor.

**Table 58. Intel® Server Chassis P4000M family**

| Intel® Server Chassis SKU | System Fans | Storage Drives | Power Supply(s) |
|---|---|---|---|
| P4308XXMFEN | Two Fixed Fans | Eight 3.5” Fixed Drive Trays | 550W Fixed PSU |
| P4308XXMHEN | Two Fixed Fans | Eight 3.5” Hotswap Drive Bay | 550W Fixed PSU |
| P4308XXMFGN | Two Fixed Fans | Eight 3.5” Fixed Drive Trays | One 750W CRPS |
| P4308XXMHGC | Five Redundant Fans | Eight 3.5” Hotswap Drive Bay | One 750W CRPS |
| P4308XXMHJC | Five Redundant Fans | Eight 3.5” Hotswap Drive Bay | Two 1200W CRPS |
| P4208XXMHEN | Two Fixed Fans | Eight 2.5” Hotswap Drive Bay | 550W Fixed PSU |
| P4208XXMHDR | Two Fixed Fans | Eight 2.5” Hotswap Drive Bay | Two 460W CRPS |
| P4208XXMHGR | Two Fixed Fans | Eight 2.5” Hotswap Drive Bay | Two 750W CRPS |
| P4208XXMHGC | Five Redundant Fans | Eight 2.5” Hotswap Drive Bay | Two 750W CRPS |
| P4216XXMHJC | Five Redundant Fans | Sixteen 2.5” Hotswap Drive Bay | Two 1200W CRPS |

You must install the active processor heat sink with the airflow direction as shown in the following figure:

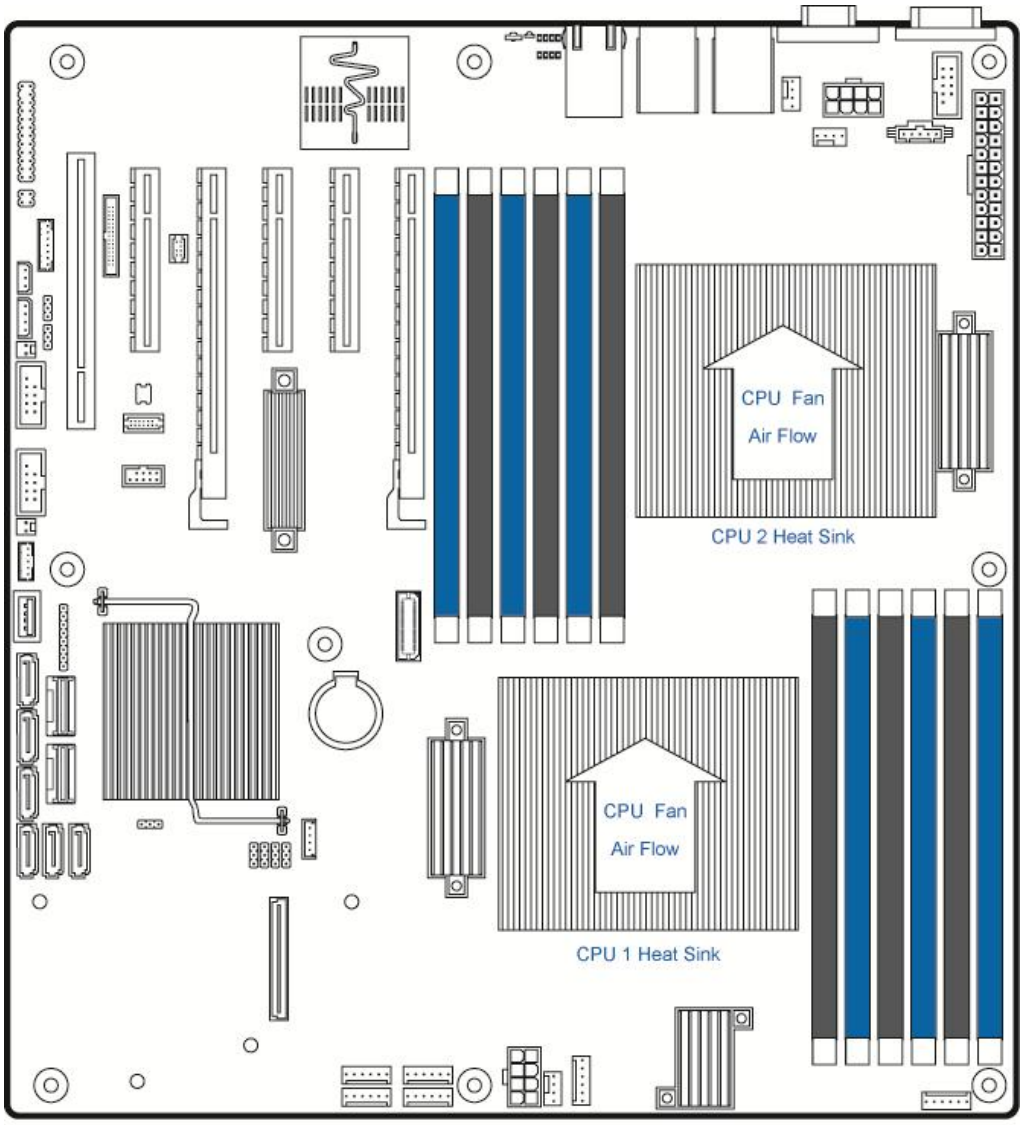


**Figure 33. Processor Heatsink Installation**

# 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 |
| AGP | Accelerated Graphics Port |
| AP | Application Processor |
| APIC | Advanced Programmable Interrupt Control |
| ARP | Address Resolution Protocol |
| ASIC | Application Specific Integrated Circuit |
| ASMI | Advanced Server Management Interface |
| BIOS | Basic Input/Output System |
| BIST | Built-In Self Test |
| BMC | Baseboard Management Controller |
| Bpp | Bits per pixel |
| Bridge | Circuitry connecting one computer bus to another, allowing an agent on one to access the other |
| BSP | Bootstrap Processor |
| Byte | 8-bit quantity. |
| CBC | Chassis Bridge Controller (A microcontroller connected to one or more other CBCs, together they bridge the IPMB buses of multiple chassis.) |
| CEK | Common Enabling Kit |
| CHAP | Challenge Handshake Authentication Protocol |
| CMOS | In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory, which normally resides on the server board. |
| DHCP | Dynamic Host Configuration Protocol |
| DIMM | Dual In-Line Memory Module |
| DPC | Direct Platform Control |
| DMA | Direct Memory Access |
| EEPROM | Electrically Erasable Programmable Read-Only Memory |
| EHCI | Enhanced Host Controller Interface |
| EMP | Emergency Management Port |
| EPS | External Product Specification |
| EVRD | Enterprise Voltage Regulator-Down |
| ESB2 | Enterprise South Bridge 2 |
| FBD | Fully Buffered DIMM |
| FMB | Flexible Mother Board |
| FRB | Fault Resilient Booting |
| FRU | Field Replaceable Unit |
| FSB | Front Side Bus |
| GB | 1024 MB |

| Term | Definition |
|---|---|
| GPIO | General Purpose I/O |
| GTL | Gunning Transceiver Logic |
| GPA | Guest Physical Address |
| HSC | Hot-Swap Controller |
| HPA | Host Physical Address |
| Hz | Hertz (1 cycle/second) |
| I2C | Inter-Integrated Circuit Bus |
| IA | Intel® Architecture |
| IBF | Input Buffer |
| ICH | I/O Controller Hub |
| IC MB | Intelligent Chassis Management Bus |
| IERR | Internal Error |
| IFB | I/O and Firmware Bridge |
| ILM | Independent Loading Mechanism |
| IMC | Integrated Memory Controller |
| INTR | Interrupt |
| IOAT | I/O Acceleration Technology |
| IP | Internet Protocol |
| IPMB | Intelligent Platform Management Bus |
| IPMI | Intelligent Platform Management Interface |
| IR | Infrared |
| ITP | In-Target Probe |
| JTAG | Joint Test Action Group |
| KB | 1024 bytes |
| KCS | Keyboard Controller Style |
| KVM | Keyboard, Video, and Mouse (Also referred to as Keyboard, Video or Video Display Unit, and Mouse) |
| LAN | Local Area Network |
| LCD | Liquid Crystal Display |
| LDAP | Local Directory Authentication Protocol |
| LED | Light Emitting Diode |
| LPC | Low Pin Count |
| LSB | Least Significant Bit |
| LUN | Logical Unit Number |
| MAC | Media Access Control |
| MB | 1024 KB |
| MCH | Memory Controller Hub |
| MD2 | Message Digest 2 – Hashing Algorithm |
| MD5 | Message Digest 5 – Hashing Algorithm – Higher Security |
| ME | Management Engine |
| ms | Milliseconds |
| MSB | Most Significant Bit |
| MTTR | Memory Type Range Register |
| Mux | Multiplexer |

| Term | Definition |
|---|---|
| NIC | Network Interface Controller |
| Nm | Nanometer |
| NMI | Non-maskable Interrupt |
| NUMA | Non-Uniform Memory Architecture |
| NVSRAM | Non-volatile Static Random Access Memory |
| OBF | Output Buffer |
| OEM | Original Equipment Manufacturer |
| Ohm | Unit of electrical resistance |
| PAE | Physical Address Extension |
| PECI | Platform Environment Control Interface |
| PEF | Platform Event Filtering |
| PEP | Platform Event Paging |
| PIA | Platform Information Area (This feature configures the firmware for the platform hardware) |
| PLD | Programmable Logic Device |
| PMI | Platform Management Interrupt |
| POST | Power-On Self Test |
| PSMI | Power Supply Management Interface |
| PWM | Pulse-Width Modulation |
| QPI | QuickPath Interconnect |
| RAM | Random Access Memory |
| RAS | Reliability, Availability, and Serviceability |
| RASUM | Reliability, Availability, Serviceability, Usability, and Manageability |
| RDIMM | Registered Dual In-Line Memory Module |
| RISC | Reduced Instruction Set Computing |
| RMII | Reduced Media Independent Interface |
| ROM | Read Only Memory |
| RTC | Real-Time Clock (Component of ICH peripheral chip on the server board) |
| SAS | Serial Attached SCSI |
| SDR | Sensor Data Record |
| SECC | Single Edge Connector Cartridge |
| SEEPROM | Serial Electrically Erasable Programmable Read-Only Memory |
| SEL | System Event Log |
| SES | SCSI Enclosure Services |
| SGPIO | Serial General Purpose Input/Output |
| SIO | Server Input/Output |
| SMBUS | System Management BUS |
| SMI | Server Management Interrupt (SMI is the highest priority nonmaskable interrupt) |
| SMM | Server Management Mode |
| SMS | Server Management Software |
| SNMP | Simple Network Management Protocol |
| SPD | Serial Presence Detect |
| SPS | Server Platforms Services (as in Intel® Server Platform Services) |
| TBD | To Be Determined |

| Term | Definition |
|---|---|
| TDP | Thermal Design Power |
| TIM | Thermal Interface Material |
| UART | Universal Asynchronous Receiver/Transmitter |
| UDIMM | Unbuffered Dual In-Line Memory Module |
| UDP | User Datagram Protocol |
| UHCI | Universal Host Controller Interface |
| URS | Unified Retention System |
| UTC | Universal time coordinate |
| VID | Voltage Identification |
| VLSI | Very-large-scale integration |
| VRD | Voltage Regulator Down |
| VT | Virtualization Technology |
| Word | 16-bit quantity |
| ZIF | Zero Insertion Force |

# Reference Documents

See the following documents for additional information:

- *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 20*03, © 2000-2003, Di*stributed Management Task Force, Inc.,* http://www.dmtf.org.
- *BIOS for EPSD Platforms Based on Intel® Xeon Processor E5-4600/2600/2400/1600 Product Families External Product Specification.*
- *EPSD Platforms Based on Intel Xeon® Processor E5 4600/2600/2400/1600 Product Families BMC Core Firmware External Product Specification.*

Preliminary