Intel® Server Board S2600CO Family

Technical Product Specification

Revision History

Date	Revision Number	Modifications
February, 2012	1.0	Initial release.
June, 2012	1.1	Updated Thermal Management and Environmental Limits Specification.
September, 2012	1.2	Updated PCIe slots definition and added NTB support section.

Disclaimers

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

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

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

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

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 © 2012 Intel Corporation

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 S2600CO.

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

1.1   Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Product Overview
- Chapter 3 – Functional Architecture Overview
- Chapter 4 – Technology Support
- Chapter 5 – System Security
- Chapter 6 – Platform Management Functional Overview
- Chapter 7 – Advanced Management Features Support (RMM4)
- Chapter 8 – On-board Connector/Header Overview
- Chapter 9 – Reset and Recovery Jumpers
- Chapter 10 – Light Guided Diagnostics
- Chapter 11 – Environmental Limits Specification
- Chapter 12 – Power Supply Specification Guidelines
- Appendix A: Integration and Usage Tips
- Appendix B: Compatible Intel® Server Chassis
- Appendix C: Integrated BMC Sensor Tables
- Appendix D: Intel® Server Board S2600CO Family Specific Sensors
- Appendix E: Management Engine Generated SEL Event Messages
- Appendix F: POST Code Diagnostic LED Decoder
- Appendix G: POST Code Errors
- Glossary
- Reference Documents

1.2   Server Board Use Disclaimer

Intel Corporation server boards support add-in peripherals and contain a number of high-density VLSI (Very Large Scale Integration) and power delivery components that need adequate airflow to cool. Intel® ensures through its own chassis development and testing that when Intel® server building blocks are used together, the fully integrated system will meet the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel® developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of 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.   Product Overview

The Intel® Server Board S2600CO is a monolithic printed circuit board (PCB) assembly with features designed to support the pedestal server markets. It has two board SKUs, namely S2600CO4 and S2600COE. These server boards are designed to support the Intel® Xeon® Processor E5-2600 product family. Previous generation Intel® Xeon® processors are not supported. Many of the features and functions of these two SKUs are common. A board will be identified by the name when a described feature or a function is unique to it.

Table 1. Intel® Server Board S2600CO Family Feature Set

Feature	Description
Processor Support	▪ Two LGA 2011 (Socket R) processor sockets ▪ **S2600CO4** - Support for one or two Intel® Xeon® Processor(s) E5-2600 product with a Thermal Design Power (TDP) of up to 135W. ▪ **S2600COE -** Support for one or two Intel® Xeon® Processor(s) E5-2600 product family with a Thermal Design Power (TDP) of up to 150W with possible configuration limits.
Memory	▪ 16 DIMM slots – 2 DIMM slots/channel – 4 memory channels per processor ▪ Channels A, B, C, D, E, F, G and H ▪ Support for Registered DDR3 Memory (RDIMM), LV-RDIMM, Unbuffered DDR3 memory ((UDIMM) with ECC and Load Reduced DDR3 memory (LR-DIMM) ▪ Memory DDR3 data transfer rate of 800, 1066, 1333 MT/s, and1600 MT/s ▪ DDR3 standard I/O voltage of 1.5V and DDR3 Low Voltage of 1.35V
Chipset	Intel® C600-A chipset with support for optional Storage Option Select keys
Rear I/O connections	▪ One DB15 Video VGA connector ▪ One DB9 Serial connector ▪ Four USB 2.0 connectors ▪ Four RJ-45 Network Interface Connectors supporting 10/100/1000Mb
Internal I/O connectors/headers	▪ One 2x5 pin connector providing front panel support for two USB ports ▪ One internal type A USB 2.0 connector ▪ One internal low-profile 2mm USB port for USB Solid State Drive ▪ One 2x15 pin SSI-EEB compliant front panel connector ▪ One DH-10 Serial Port B connector ▪ Two internal IEEE 1394b connectors (**Intel® Server Board S2600COE only.)**
Cooling Fan Support	▪ Two 4-pin managed CPU fan headers ▪ Six 6-pin managed system fan headers ▪ One 4-pin managed rear system fan header ▪ One 4-pin I^{2}C header that is intended to be used for third party fan control circuits using a Maxim 72408 controller (**Intel® Server Board S2600CO4 only.**)

Intel order number G42278-004

Feature	Description
Add-in Card Slots	- Support up to six expansion slots - Slot 1: PCIe Gen II x4 electrical with x8 physical connector, routed from Intel® C600 Chipset, support half-length card - Slot 2: PCIe Gen III x16 electrical with x16 physical connector, routed from CPU1, support full length card - Slot 3: PCIe Gen III x16 electrical with x16 physical connector(blue slot), routed from CPU2, support full length, double width card - Slot 4: PCIe Gen III x8 electrical with x8 connector, routed from CPU2, support full length card - Slot 5: PCIe Gen III x16 electrical with x16 connector(blue slot), routed from CPU1, support full length, double width card - Slot 6: PCIe Gen III x16 electrical with x16 connector, routed from CPU2, support half-length card, Intel® designed PCIe riser card
Storage	- One low-profile eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices - Two AHCI SATA connectors capable of supporting up to 6Gb/sec - AHCI SATA 0 supports high profile, vertical SATA DOM with onboard power. - Eight SCU SAS/SATA connectors capable of supporting up to 3Gb/sec - Intel® SAS ROC module support (Optional, PCIe slot form factor) - Intel® RAID C600 Upgrade Key support providing optional expanded SAS/SATA RAID capabilities
RAID Support	- Intel® RSTe SW RAID 0/1/10/5 - LSI* SW RAID 0/1/10
Video Support	Integrated Matrox* G200 2D Video Graphics controller
LAN	Four Gigabit through Intel® I350-AM Quad 10/100/1000 integrated GbE MAC and PHY controller
Security	Trusted Platform Module (Accessory Option)
Server Management	- Integrated Baseboard Management Controller, IPMI 2.0 compliant - Support for Intel® Server Management Software - Intel® Remote Management Module 4 support (Accessory Option) - Intel® Remote Management Module 4 Lite support (Accessory Option)
Form Factor	SSI EEB (12"x13")
Compatible Intel® Server Chassis	Intel® Server Chassis P4000M family

2.1   Server Board Connector and Component Layout

The following illustrations provide a general overview of the server board, identifying key feature and component locations. Each connector and major component is identified by a number or letter, and the description is given below the figure.

AF004732

Callout	Description	Callout	Description
A	Chassis Intrusion	AH	System Fan 1 Connector
B	Slot1, PCI Express* Gen2	AI	LCP
C	RMM4 Lite	AJ	Optional 12V Power Connector
D	Slot 2, PCI Express* Gen3	AK	BIOS Default
E	Slot 3, PCI Express* Gen3 (blue slot), supports	AL	SMB_PMBUS*

Callout	Description	Callout	Description
	double width card		
F	Battery	AM	Main Power
G	Slot 4, PCI Express* Gen3	AN	USB_4 (Internal Type A USB Connector)
H	Slot 5, PCI Express* Gen3 (blue slot), supports double width card	AO	Storage Upgrade key connector
I	Slot 6, PCI Express* Gen3	AP*	FAN BOARD_I^{2}C (Intel® Server Board S2600CO4 only)
J	DIMM E1/E2/F1/F2	AQ	SATA_0 connector
K	Status LED	AR	SATA_1 Connector
L	ID LED	AS	SAS/SATA_7 Connector
M	Diagnostic LED	AT	SAS/SATA_6 Connector
N	NIC 3/4	AU	SAS/SATA_5 Connector
O	USB 0/1/2/3, NIC 1,2	AV	SAS/SATA_4 Connector
P	VGA	AW	SAS/SATA_3 Connector
Q	CPU 2 Power connector	AX	SAS/SATA_2 Connector
R	System Fan 7	AY	SAS/SATA_1 Connector
S	Serial Port A	AZ	SAS/SATA_0 Connector
T	CPU 2 Fan Connector	BA	SAS SGPIO 0
U	DIMM H1/H2/G1/G2	BB	SAS SGPIO 1
V	DIMM A1/A2/B1/B2	BC	IPMB
W	CPU 1 Power connector	BD	HSBP_I^{2}C
X	CPU 1 Fan Connector	BE	USB 5-6 (front panel USB connector)
Y	DIMM C1/C2/D1/D2	BF	ME Force Update
Z	USB port to support SSD	BG	BMC Force Update
AA	HDD LED Header	BH**	IEEE_1394B_0 Connector (**Intel® Server Board S2600COE only**)
AB	TPM Connector	BI	Password Clear
AC	System Fan 6 Connector	BJ	BIOS Recovery
AD	System Fan 5 Connector	BK**	IEEE_1394B_1 Connector (**Intel® Server Board S2600COE only**)
AE	System Fan 4 Connector	BL	Serial B connector
AF	System Fan 3 Connector	BM	RMM4 NIC
AG	System Fan 2 Connector	BN	SSI Front Panel

Figure 1. Major Board Components
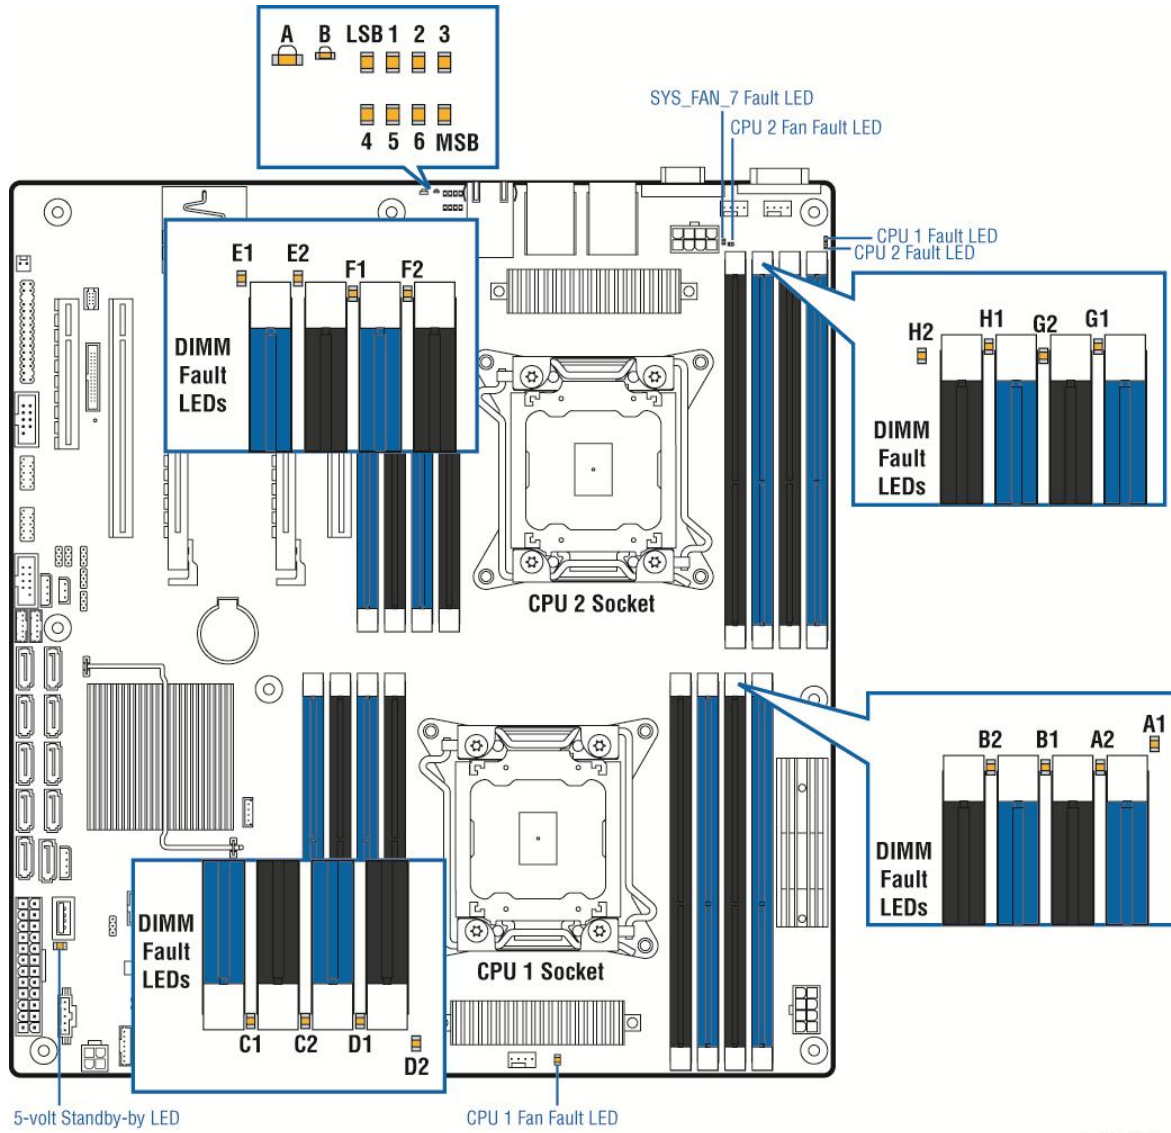
Figure 2. Intel® Light Guided Diagnostic LED Identification

See Chapter 10 for additional details.
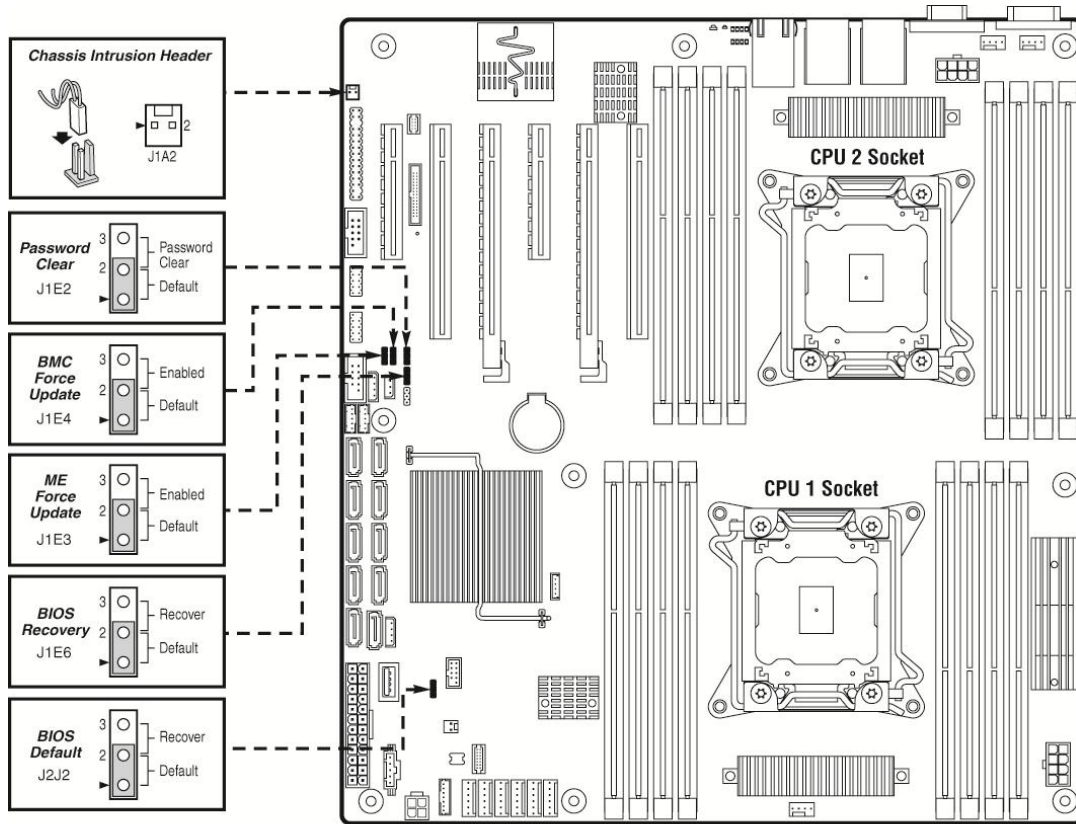
Figure 3. Jumper Block Identification

See Chapter 9 for additional details.

Callout	Description	Callout	Description
A	Serial Port A	E	NIC Port 3 and 4
B	Video	F	Diagnostics LED's
C	NIC Port 1, USB Port 0 (top) and 1 (bottom)	G	ID LED
D	NIC Port 2, USB Port 2 (top) and 3 (bottom)	H	Status LED

Figure 4. Rear I/O Layout

2.2   Server Board Dimensional Mechanical Drawings

Figure 5. Mounting Hole Locations (1 of 2)

Figure 6. Mounting Hole Locations (2 of 2)

Figure 7. Major Connector Pin-1 Locations

Figure 8. Primary Side Keep-out

Figure 9. Secondary Side Keep-out

3.   Functional Architecture Overview

The architecture and design of the Intel® Server Board S2600CO is developed around the integrated features and functions of the Intel® processor E5-2600 product family, the Intel® C600-A chipset, the Intel® Ethernet Controller I350 Quad Port 1GbE chip, and the Server Engines* Pilot-III Server Management Controller.

The following diagram provides an overview of the server board architecture, showing the features and interconnects of each of the major sub-system components.
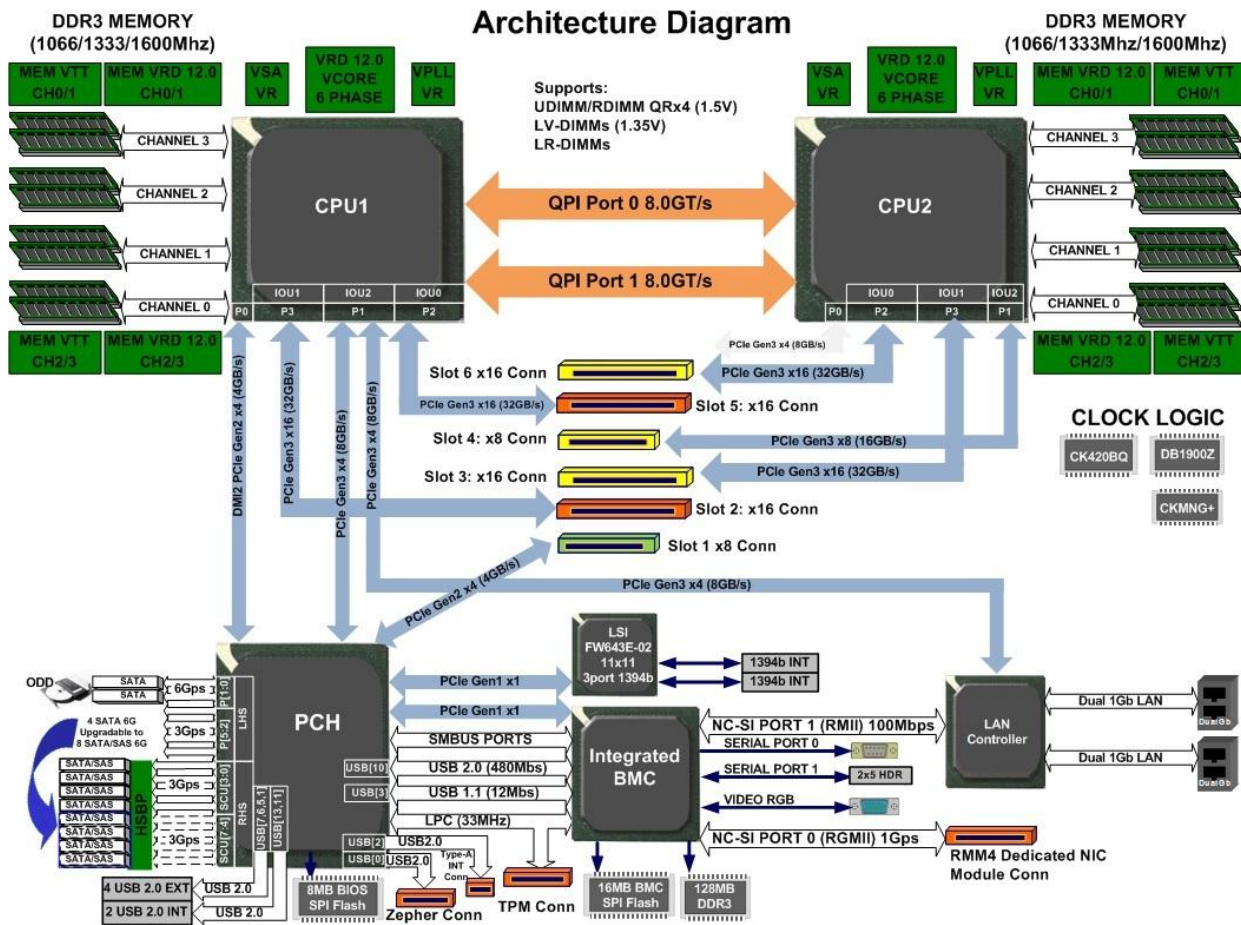
Figure 10. Intel® Server Board S2600CO Functional Block Diagram

3.1   Processor Support

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

- Intel® Xeon® processor E5-2600 product family, with a Thermal Design Power (TDP) of up to 135W.
- Intel® Xeon® processor E5-2600 product family, with a Thermal Design Power (TDP) of up to 150W with Intel® Server Board S2600COE with possible configuration limits.

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

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

3.1.1   Processor Socket Assembly

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

The illustration below identifies each sub-assembly component.

Figure 11. Processor Socket Assembly

3.1.2   Processor Population Rules

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

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

3.1.3   Processor Initializion Error Summary

The following table describes mixed processor conditions and recommended actions for all Intel® server boards and Intel® server systems designed around the Intel® Xeon® processor E5-2600 product family and Intel® C600 chipset product family architecture. The errors fall into one of the following categories:

- **Fatal:** If the system can boot, it pauses at a blank screen with the text "**Unrecoverable fatal error found. System will not boot until the error is resolved**" and "**Press <F2> to enter setup**", regardless of whether the "Post Error Pause" setup option is enabled or disabled.

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

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

    For Fatal Errors during processor initialization, the System Status LED will be set to a steady Amber color, indicating an unrecoverable system failure condition

- **Major:** If the "Post Error Pause" setup option is enabled, the system goes directly to the Error Manager screen to display the error and log the error code to SEL. Operator intervention is required to continue booting the system.

    Otherwise, if "POST Error Pause" is disabled, the system continues to boot and no prompt is given for the error, although the error code is logged to the Error Manager and in a SEL message.

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

Table 2. Mixed Processor Configurations Error Summary

Error	Severity	System Action
Processor family not identical	Fatal	The BIOS detects the error condition and responds as follows: - Logs the POST Error Code into the System Event Log (SEL). - Alerts the BMC to set the System Status LED to steady Amber. - Displays "0194: Processor family mismatch detected" message in the Error Manager. - Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor model not identical	Fatal	The BIOS detects the error condition and responds as follows: - Logs the POST Error Code into the System Event Log (SEL). - Alerts the BMC to set the System Status LED to steady Amber. - Displays "0196: Processor model mismatch detected" message in the Error Manager. - Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor cores/threads not identical	Fatal	The BIOS detects the error condition and responds as follows: - Logs the POST Error Code into the SEL. - Alerts the BMC to set the System Status LED to steady Amber. - Displays "0191: Processor core/thread count mismatch detected" message in the Error Manager. Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor cache not identical	Fatal	The BIOS detects the error condition and responds as follows: - Logs the POST Error Code into the SEL. - Alerts the BMC to set the System Status LED to steady Amber. - Displays "0192: Processor cache size mismatch" detected message in the Error Manager. - Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor frequency (speed) not identical	Fatal	The BIOS detects the processor frequency difference, and responds as follows: - Adjusts all processor frequencies to the highest common frequency. - No error is generated – this is not an error condition. - Continues to boot the system successfully. If the frequencies for all processors cannot be adjusted to be the same, then this is an error, and the BIOS responds as follows: - Logs the POST Error Code into the SEL. - Alerts the BMC to set the System Status LED to steady Amber. - Does not disable the processor. - Displays "0197: Processor speeds unable to synchronize" message in the Error Manager. - Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.

Error	Severity	System Action
Processor Intel® QuickPath Interconnect link frequencies not identical	Fatal	The BIOS detects the QPI link frequencies and responds as follows: ▪ Adjusts all QPI interconnect link frequencies to highest common frequency. ▪ No error is generated – this is not an error condition. ▪ Continues to boot the system successfully. If the link frequencies for all QPI links cannot be adjusted to be the same, then this <u>is</u> an error, and the BIOS responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Alerts the BMC to set the System Status LED to steady Amber. ▪ Displays "0195: Processor Intel® QPI link frequencies unable to synchronize" message in the Error Manager. ▪ Does not disable the processor. ▪ Takes Fatal Error action (see above) and will not boot until the fault condition is remedied.
Processor microcode update missing	Minor	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Displays "818x: Processor 0x microcode update not found" message in the Error Manager or on the screen. ▪ The system continues to boot in a degraded state, regardless of the setting of POST Error Pause in the Setup.
Processor microcode update failed	Major	The BIOS detects the error condition and responds as follows: ▪ Logs the POST Error Code into the SEL. ▪ Displays "816x: Processor 0x unable to apply microcode update" message in the Error Manager or on the screen. Takes Major Error action. The system may continue to boot in a degraded state, depending on the setting of POST Error Pause in Setup, or may halt with the POST Error Code in the Error Manager waiting for operator intervention.

3.2   Processor Function Overview

With the release of the Intel® Xeon® processor E5-2600 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; two Intel® QuickPath Interconnect point to point links capable of up to 8.0 GT/s, up to 40 lanes of Gen 3 PCI Express* links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* Gen 2 interface with a peak transfer rate of 5.0 GT/s. The processor supports up to 46 bits of physical address space and 48-bit of virtual address space.

The following sections will provide an overview of the key processor features and functions that help to define the architecture, performance and supported functionality of the server board. For more comprehensive processor specific information, refer to the Intel® Xeon® processor E5-2600 product family documents listed in the Reference Documents list.

Processor Core Features:

- Up to 8 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
- 1 GB large page support for server applications
- A 32-KB instruction and 32-KB data first-level cache (L1) for each core
- A 256-KB shared instruction/data mid-level (L2) cache for each core
- Up to 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 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

3.2.1   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.

3.2.2   Integrated Memory Controller (IMC) and Memory Subsystem
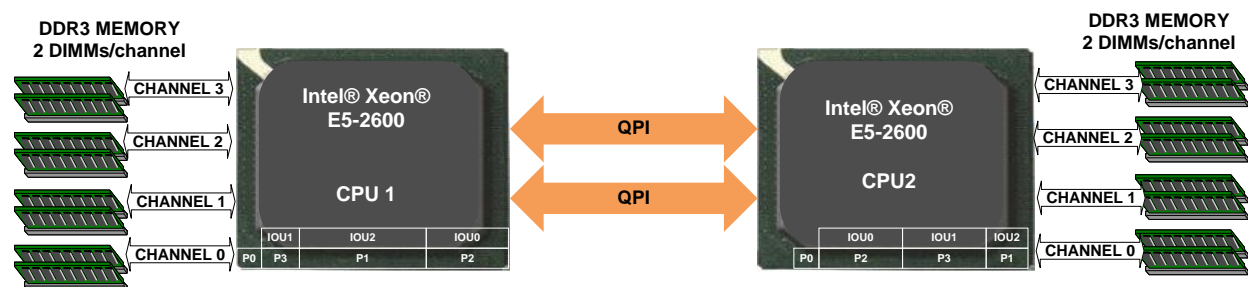
Figure 12. Integrated Memory Controller Functional Block Diagram

Integrated into the processor is a memory controller. Each processor provides four 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
- 1-Gb, 2-Gb, and 4-Gb 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 1 DIMM populated
- Integrated dual SMBus* master controllers
- Command launch modes of 1n/2n
- RAS Support:
    - Rank Level Sparing and Device Tagging

- o Demand and Patrol Scrubbing
- o 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
- o Lockstep mode where channels 0 and 1 and channels 2 and 3 are operated in lockstep mode
- o Data scrambling with address to ease detection of write errors to an incorrect address.
- o Error reporting by the Machine Check Architecture
- o Read Retry during CRC error handling checks by IMC
- o Channel mirroring within a socket
    - • CPU1 Channel Mirror Pairs (A,B) and (C,D)
    - • CPU2 Channel Mirror Pairs (E,F) and (G,H)
- o Error Containment Recovery
- ▪ Improved Thermal Throttling with dynamic Closed Loop Thermal Throttling (CLTT)
- ▪ Memory thermal monitoring support for DIMM temperature

3.2.2.1   Supported Memory

Table 3. UDIMM 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 Slots per Channel			
				1DPC		2DPC	
				1.35V	1.5V	1.35V	1.5V
SRx8 Non-ECC	1GB	2GB	4GB	n/a	1066, 1333	n/a	1066, 1333
DRx8 Non-ECC	2GB	4GB	8GB	n/a	1066, 1333	n/a	1066, 1333
SRx16 Non-ECC	512MB	1GB	2GB	n/a	1066, 1333	n/a	1066, 1333
SRx8 ECC	1GB	2GB	4GB	1066, 1333	1066, 1333	1066	1066, 1333
DRx8 ECC	2GB	4GB	8GB	1066, 1333	1066, 1333	1066	1066, 1333

Table 4. 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]			
				2 Slots per Channel			
				1DPC		2DPC	
				1.35V	1.5V	1.35V	1.5V
SRx8	1GB	2GB	4GB	1066, 1333	1066 1333 1600	1066, 1333	1066, 1333, 1600
DRx8	2GB	4GB	8GB	1066, 1333	1066 1333 1600	1066, 1333	1066, 1333, 1600

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]			
				2 Slots per Channel			
				1DPC		2DPC	
				1.35V	1.5V	1.35V	1.5V
SRx4	2GB	4GB	8GB	1066, 1333	1066, 1333, 1600	1066, 1333	1066, 1333, 1600
DRx4	4GB	8GB	16GB	1066, 1333	1066, 1333, 1600	1066, 1333	1066, 1333, 1600
QRx4	8GB	16GB	32GB	800	1066	800	800
QRx8	4GB	8GB	16GB	800	1066	800	800

Table 5. LRDIMM Support Guidelines

Ranks Per DIMM and Data Width[1]	Memory Capacity Per DIMM[2]		Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)[3,4,5]	
			2 Slots per Channel	
			1DPC and 2DPC	
			1.35V	1.5V
QRx4 (DDP)	16GB	32GB	1066	1066, 1333
QRx8 (P)	8GB	16GB	1066	1066, 1333

3.2.2.2   Memory Slot Identification and Population Rules

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

Each installed processor provides four channels of memory. On the Intel® Server Board S2600CO each memory channel support 2 memory slots, for a total possible 16 DIMMs installed.

- System memory is 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, C and D. The memory channels from processor socket 2 are identified as Channel E, F, G, and H.

- Each memory slot on the server board is identified by channel and slot number within that channel. For example, DIMM_A1 is the first slot on Channel A on processor 1; DIMM_E1 is the first DIMM socket on Channel E on processor 2.

- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.

- A processor may be installed without populating the associated memory slots provided a second processor is installed with associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.

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

- The BLUE memory slots on the server board identify the first memory slot for a given memory channel.

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

On the Intel® Server Board S2600CO a total of 16 DIMM slots is provided (2 CPUs – 4 Channels/CPU, 2 DIMMs /Channel). The nomenclature for DIMM sockets is detailed in the following table:

Table 6. Intel® Server Board S2600CO DIMM Nomenclature

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

Figure 13. Intel® Server Board S2600CO DIMM Slot Layout

The following are generic DIMM population requirements that generally apply to the Intel® Server Board S2600CO.

- All DIMMs must be DDR3 DIMMs
- 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.
- Quad rank RDIMMs are supported but not validated by Intel®.
- A maximum of 8 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.

- Mixing of quad ranks DIMMs (RDIMM Raw Cards F and H) in one channel is not validated.

3.2.2.3 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 display 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

3.2.2.4 Integrated Memory Controller Operating Modes

3.2.2.4.1 Independent Channel Mode

In non-ECC (Error Correction Code) and x4 Single Device Data Correction (SDDC) configuration, each channel runs 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 bursts 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.

3.2.2.4.2   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, and 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).

Channel 0 and Channel 1 can be in lockstep. Channel 2 and Channel 3 can be in lockstep.

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 drawback 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.

3.2.2.4.3   Mirror Mode

Memory mirroring mode is the mechanism by which a component of memory is mirrored. In mirrored mode, when a write is performed to one copy, a write is generated to the target location as well. This guarantees that the target is always updated with the latest data from the main copy. The IMC module supports mirroring across the corresponding mirroring channel within the processor socket but not across sockets. DIMM organization in each slot of one channel must be identical to the DIMM in the corresponding slot of the other channel. This allows a single decode for both channels. When mirroring mode is enabled, memory image in Channel 0 is maintained the same as Channel 1 and Channel 2 is maintained the same as Channel 3.

3.2.2.5   Memory RAS Support

The server board supports the following memory RAS modes:

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

Regardless of RAS mode, the requirements for populating within a channel given in the section 3.2.2.2 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 *Lockstep Channel Mode and Mirroring Mode, processor channels are paired together as a "Domain".*

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

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).

3.2.2.5.1   Singel 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.

3.2.2.5.2   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.

3.2.2.5.2.1   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 copy with the random factor in correctable ECC errors. Rather than reporting every correctable error that occurs, the BIOS have 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.

3.2.2.5.2.2   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 from the 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.

3.2.2.5.3   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.

3.2.2.5.4   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.

3.2.2.5.5   Rank Sparing Mode

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/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 operation. 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.

3.2.2.5.6   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 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 operation. 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.

3.2.3   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 40 lanes of PCI Express*. Following are key attributes of the PCI Express* interface:
    - Gen3 speeds up to 8 GT/s
    - X16 interface bifurcated down to two x8 or four x4 (or combinations)
    - X8 interface bifurcated down to two x4
- **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.
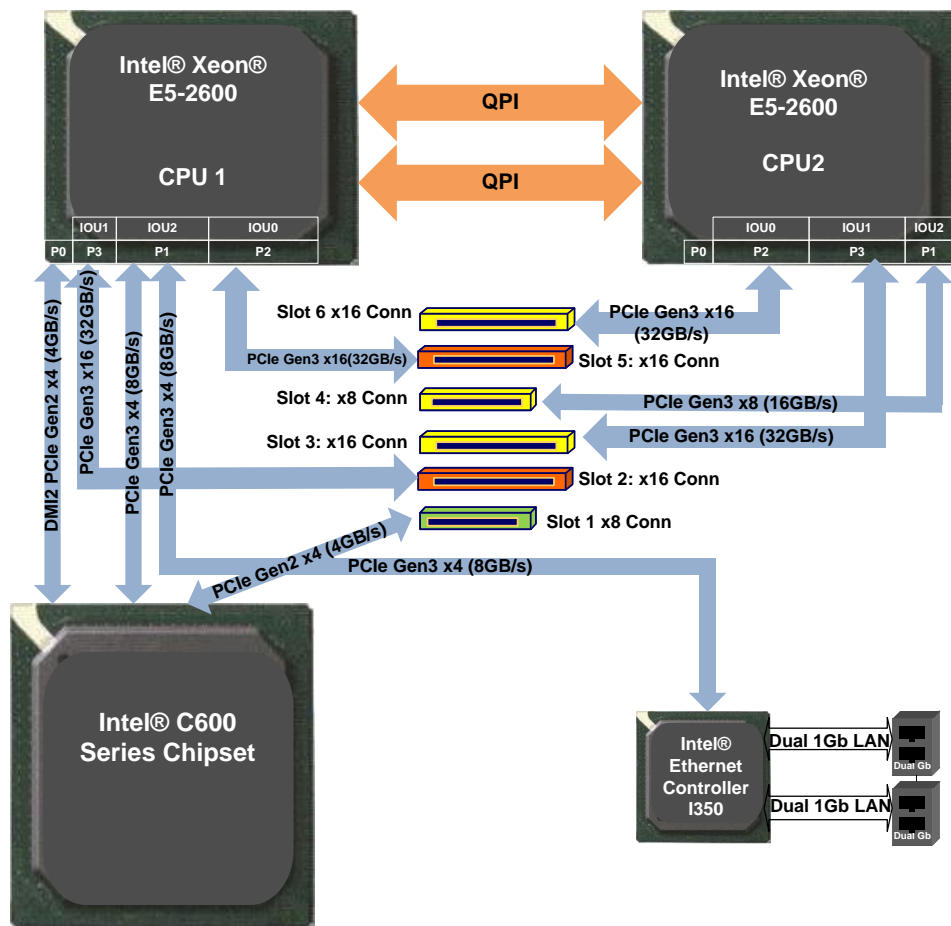
Figure 14. Functional Block Diagram of Processor IIO Sub-system

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

3.2.3.1 PCI Card Support

The server board provides six PCI card slots identified by PCIe Slot 1 to PCIe Slot 6. The PCIe slot 1 is routed from Intel® C600 Series Chipset, The PCIe slot 2 and 5 are routed from CPU 1. The PCIe slot 3, 4, 6 are routed from CPU 2.

- Slot 1: PCIe Gen II x4 electrical with x8 physical connector, routed from Intel® C600 Chipset, support half-length card

- Slot 2: PCIe Gen III x16 electrical with x16 physical connector, routed from CPU1, support full length card

- Slot 3: PCIe Gen III x16 electrical with x16 physical connector, routed from CPU2, support full length, double width card

- Slot 4: PCIe Gen III x8 electrical with x8 connector, routed from CPU2, support full length card

- Slot 5: PCIe Gen III x16 electrical with x16 connector, routed from CPU1, support full length, double width card

- Slot 6: PCIe Gen III x16 electrical with x16 connector, routed from CPU2, support half-length card, Intel® PCIe riser card

Note: PCIe Slot 3, 4, 6 can only be used in dual processor configurations.

3.2.3.2 Intel® Integrated RAID Option

The server board provides support for Intel® Integrated RAID modules with PCIe form factor. These optional modules attach to PCIe slots on the server board and are supported by x8 PCIe Gen3 signals from the IIO module of the CPU processor. Features of this option include:

- SKU options to support full or entry level hardware RAID

- Dual-core 6Gb SAS ROC

- 4 or 8 port and SAS/SATA or SATASKU options to support 512MB or 1GB embedded memory

- Intel® designed flash + optional support for super-cap backup (Maintenance Free Back Up) or improved Lithium Polymer battery

Table 7. Supported Intel® Integrated RAID Modules

External Name	Description	Product Code
Intel® Integrated RAID Module RMS25KB080	8 Port SAS, PCIe Slot RAID Levels 0,1,1E and 10	RMS25KB080
Intel® Integrated RAID Module RMS25KB040	4 Port SAS, PCIe Slot RAID Levels 0,1,1E and10	RMS25KB040
Intel® Integrated RAID Module RMS25PB080	8 Port SAS/SATA, PCIe Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMS25PB080
Intel® Integrated RAID Module RMS25PB040	4 Port SAS/SATA, PCIe Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMS25PB040
Intel® Integrated RAID Module RMT3PB080	8 Port SATA, PCIe Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMT3PB080

For additional product information, please refer the following Intel® documents:

1. *Intel® Integrated RAID Module RMS25PB080, RMS25PB040, RMS25CB080, and RMS25CB040 Hardware/Installation User Guide*
2. *Intel Integrated RAID Module RMS25KB040, RMS25KB080, RMS25JB040, RMS25JB080 –Hardware/Installation User Guide*
3. *Intel Integrated RAID Module RMT3PB080 and RMT3CB080 –Hardware/Installation User Guide*

3.2.3.3 Intel Riser Card Option

The server board provides support for Intel® Riser Card. These optional modules attach to PCIe slot6 on the server board and are supported by x16 PCIe Gen3 signals from the IIO module of the CPU 2 processor.

Table 8. Supported Intel® Riser Card

Product Code	Description
AXX2UPCIEX16	2U PCIe x16 riser, 1 slot x16 electrical/x16 mechanical
AXX2UPCIEX8X4	2U PCIe x16 riser, 3 slots one x8 electrical/x16 mechanical Two x4 electrical/x8 mechanical
AXX1UPCIEX16	1U PCIe x16 riser, 1 slot x16 electrical/x16 mechanical

3.2.3.4 Network Interface

Network connectivity is provided by means of an onboard Intel® Ethernet Controller 1350-AM4 providing up to four 10/100/1000 Mb Ethernet ports. The NIC chip is supported by implementing x4 PCIe Gen3 signals from the IIO module of the CPU 1 processor.

On the Intel® Server Board S2600CO, four external 10/100/1000 Mb RJ45 Ethernet ports are provided. Each Ethernet port drives two 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 9. External RJ45 NIC Port LED Definition

LED	Color	LED State	NIC State
Left	Green	Off	LAN link not established
		On	LAN link is established
		Blinking	LAN activity is occurring
Right	N/A	Off	10 Mb/sec data rate
	Amber	On	100 Mb/sec data rate
	Green	On	1000 Mb /sec data rate

The server board has seven MAC addresses programmed at the factory. 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

3.3   Intel® C600-A Chipset Functional Overview

The following sub-sections will provide an overview of the key features and functions of the Intel® C600-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 Documents list.
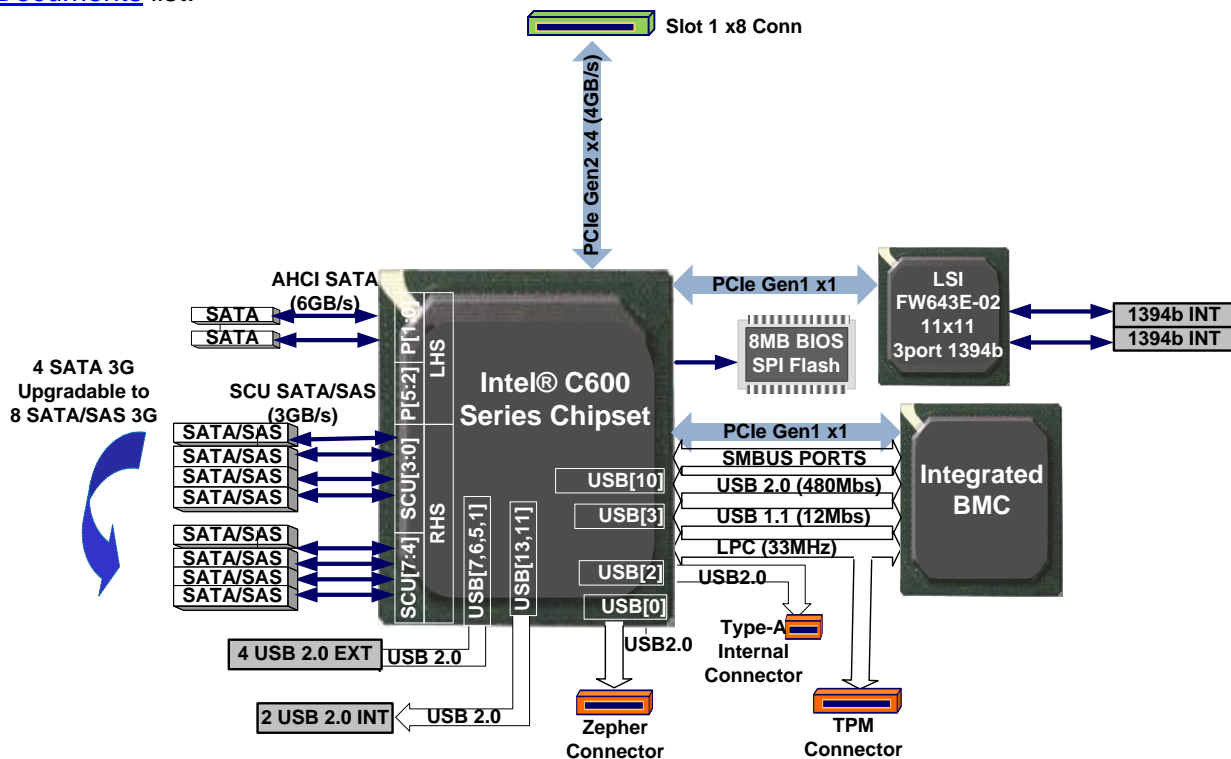
Figure 15. Functional Block Diagram - Chipset Supported Features and Functions

On the Intel® Server Boards S2600CO, the chipset provides support for the following on-board functions:

- Low Pin Count (LPC) interface
- Serial Peripheral Interface (SPI)
- Universal Serial Bus (USB) Controller
- Serial Attached SCSI (SAS)/Serial ATA (SATA) Support
- Manageability Features

3.3.1   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.

The PCI Express* Port 3A of Intel® Xeon® Processor E5-2600 Product Families can be configured to be a transparent bridge or a NTB with x4/x8 link width and Gen1/Gen2/Gen3 link speed. Also this NTB port could be attached to another NTB port or PCI Express* Root Port on another subsystem. NTB supports three 64bit BARs as configuration space or prefetchable memory windows that can access both 32bit and 64bit address space through 64bit BARs.

There are three NTB supported configuration:

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

3.3.2   Low Pin Count (LPC) Interface

The chipset implements an LPC Interface as described in the *LPC 1.1 Specification* and provides support for up to two Master/DMI devices. On the server board, the LPC interface is utilized as an interconnect between the chipset and the Integrated Base Board Management Controller as well as providing support for the optional Trusted Platform Module (TMP).

3.3.3   Universal Serial Bus (USB) Controller

The chipset has 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 server board utilizes ten USB 2.0 ports from the chipset. All ports are high-speed, full- speed, and low-speed capable.
- Four external USB ports are provided in a stacked housing located on the rear I/O section of the server board.
- Two USB ports are routed to an internal 10-pin connector that can be cabled for front panel support.
- One internal Type 'A' USB port.
- One eUSB connector intended for use with an optional eUSB SSD device.
- Two USB ports are routed to the Integrated BMC.

3.3.3.1   eUSB SSD Support

The server board provides support for a low profile eUSB SSD storage device. A 2mm 2x5-pin connector labeled "eUSB SSD" is used to plug this small flash storage device into.

3.3.4   On-board Serial Attached SCSI (SAS)/Serial ATA (SATA)/RAID Support and Options

The Intel® C600-A chipset provides storage support by two integrated controllers: Advanced Host Controller Interface (AHCI) and SCU. By default the server board will support up to 6 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", and four 3Gb/sec SATA ports routed from the SCU to the blue SATA/SAS port connectors labeled "SATA/SAS_0", "SATA/SAS_1", "SATA/SAS_2", "SATA/SAS_3".
AHCI SATA_0 port also supports high profile, vertical SATA DOM header with onboard power.

Note: The connectors labeled "SATA/SAS_4" to "SATA/SAT_7" is NOT functional by default and is only enabled with the addition of an Intel® RAID C600 Upgrade Key option supporting 8 SATA/SAS ports.

The server board is capable of supporting additional chipset embedded SAS, SATA, 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 "STRO_UPG_KEY".

AF004762

Figure 16. Intel® RAID C600 Upgrade Key Connector

The following table identifies available upgrade key options and their supported features.

Table 10. Intel® RAID C600 Upgrade Key Options

Intel® RAID C600 Upgrade Key Options	Key Color	Description
Default – No option key installed	N/A	4 Port SATA with Intel® ESRT RAID 0,1,10 and Intel® RSTe RAID 0,1,5,10
RKSATA4R5	Black	4 Port SATA with Intel® ESRT2 RAID 0,1, **5**, 10 and Intel® RSTe RAID 0,1,5,10
RKSATA8	Blue	8 Port SATA with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,5,10
RKSATA8R5	White	8 Port SATA with Intel® ESRT2 RAID 0,1, **5**, 10 and Intel® RSTe RAID 0,1,5,10

Intel® RAID C600 Upgrade Key Options	Key Color	Description
RKSAS4	Green	4 Port SAS with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,10
RKSAS4R5	Yellow	4 Port SAS with Intel® ESRT2 RAID 0,1, **5**, 10 and Intel® RSTe RAID 0,1,10
RKSAS8	Orange	8 Port SAS with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,10
RKSAS8R5	Purple	8 Port SAS with Intel® ESRT2 RAID 0,1, **5**, 10 and Intel® RSTe RAID 0,1,10

Additional information for the on-board RAID features and functionality can be found in the *Intel®
RAID Software User's Guide*.
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.

3.3.4.1   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 & 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600
      Upgrade Key
    - 4 & 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 = Microsoft Windows 7*, Microsoft Windows 2008*, Microsoft Windows
  2003*, RHEL*, SLES*, other Linux* variants using partial source builds.
- Utilities = Microsoft Windows* GUI and CLI, Linux* GUI and CLI, DOS CLI, and EFI CLI

3.3.4.2   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
    - Four Port SATA RAID 5 available standard (no option key required)
    - Eight 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 = Microsoft Windows 7*, Microsoft Windows 2008*, Microsoft Windows 2003*, RHEL* 6.2 and later, SLES* 11 w/SP2 and later, VMware* 5.x.
- Utilities = Microsoft Windows* GUI and CLI, Linux* CLI, DOS CLI, and EFI CLI
- Uses Matrix Storage Manager for Microsoft Windows*
- MDRAID supported in Linux* (does not require a driver)

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

3.3.5   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 de**tecting 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, and 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.

3.4   Integrated Baseboard Management Controller 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.
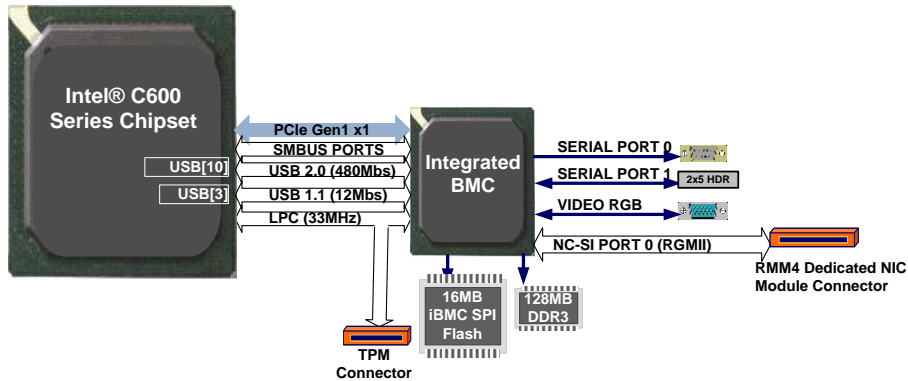
Figure 17. Integrated BMC Functional Block Diagram
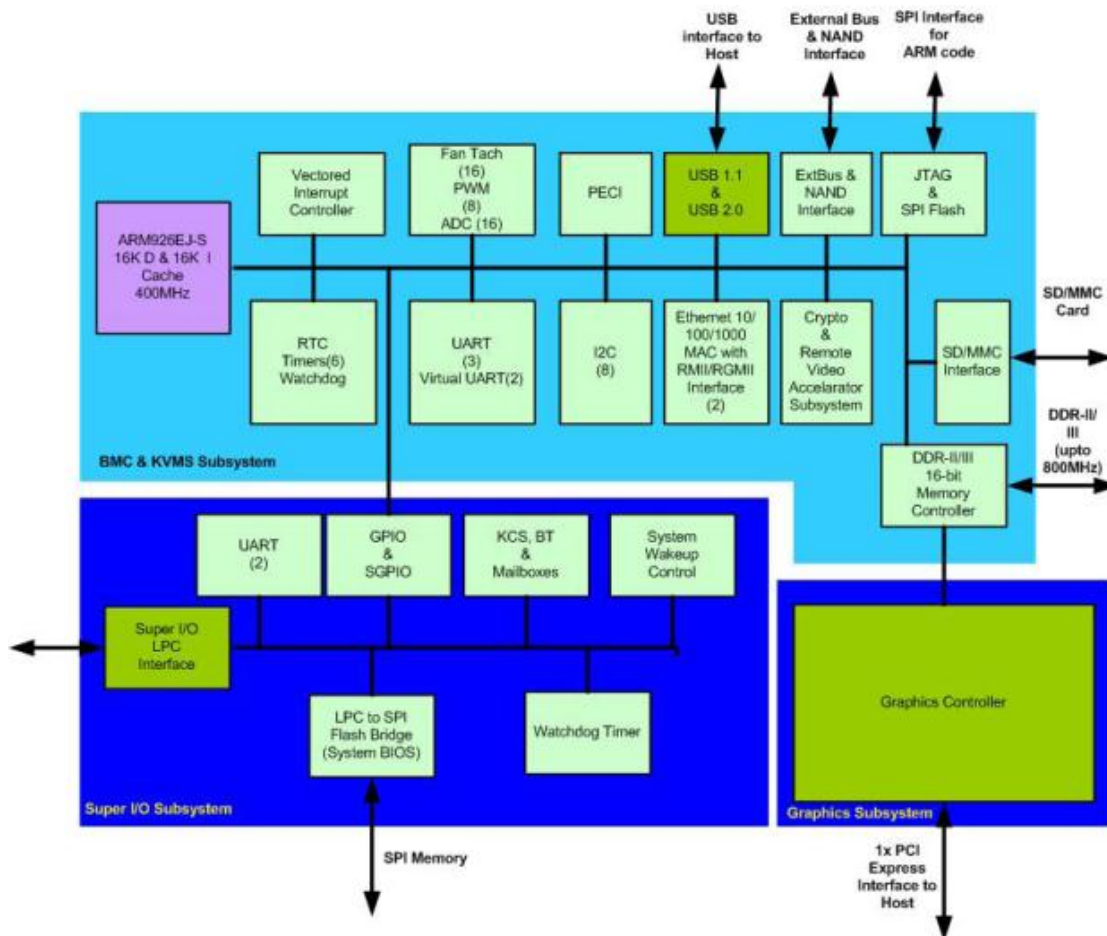
Figure 18. Integrated BMC Hardware

3.4.1 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 GPIOs
- 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

3.4.1.1 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.

3.4.1.2 Wake-up Control

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

3.4.2 Graphics Controller and Video Support

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

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

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

Table 11. Video Modes

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

2D Mode	2D Video Mode Support			
	8 bpp	16 bpp	24 bpp	32 bpp
1600x1200**	Supported	Supported		

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

The server board provides one video interfaces. The video interface is accessed using a standard 15-pin VGA connector found on the back edge of the server 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 12. Video mode

On-board Video	**Enabled**	
	Disabled	
Dual Monitor Video	**Enabled**	Shaded if on-board video is set to "Disabled"
	Disabled	

3.4.3   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 independent 10/100/1000 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^{2}C 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.

3.4.3.1   Remote Keyboard, Video, Mourse, and Storage (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

3.4.3.2   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.

For a functional overview of the baseboard management features, refer to Chapter 6.

4.   Technology Support

4.1   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.

4.2   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*.

4.3   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. Data Center Problems

IT Challenge	Requirement
Over-allocation of power	▪ Ability to monitor actual power consumption. ▪ 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. ▪ 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. ▪ 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.

Following 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.

4.3.1   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.

.

5.   System Security

5.1   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:

<div align="center">

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

</div>

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

5.2   Trusted Platform Module (TPM) Support

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

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

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

5.2.1   TPM security BIOS

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

- Measures and stores the boot process in the TPM microcontroller to allow a TPM enabled operating system to verify system boot integrity.
- Produces EFI and legacy interfaces to a TPM-enabled operating system for using TPM.
- Produces ACPI TPM device and methods to allow a TPM-enabled operating system to send TPM administrative command requests to the BIOS.

- Verifies operator physical presence. Confirms and executes operating system TPM administrative command requests.
- Provides BIOS Setup options to change TPM security states and to clear TPM ownership.

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

5.2.2   Physical Presence

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

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

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

5.2.3   TPM Security Setup Options

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

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.

5.2.3.1   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 S5520URT provides TPM settings through the security screen.

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

| Main | Advanced | **Security** | Server Management | Boot Options | Boot Manager |

Administrator Password Status   <Installed/Not Installed>
User Password Status   <Installed/Not Installed>

Set Administrator Password   [1234aBcD]
Set User Password   [1234aBcD]

Front Panel Lockout   Enabled/**Disabled**

TPM State   <Enabled & Activated/Enabled & Deactivated/Disabled & Activated/Disabled & Deactivated>
TPM Administrative Control   **No Operation/**Turn On/Turn Off/Clear Ownership
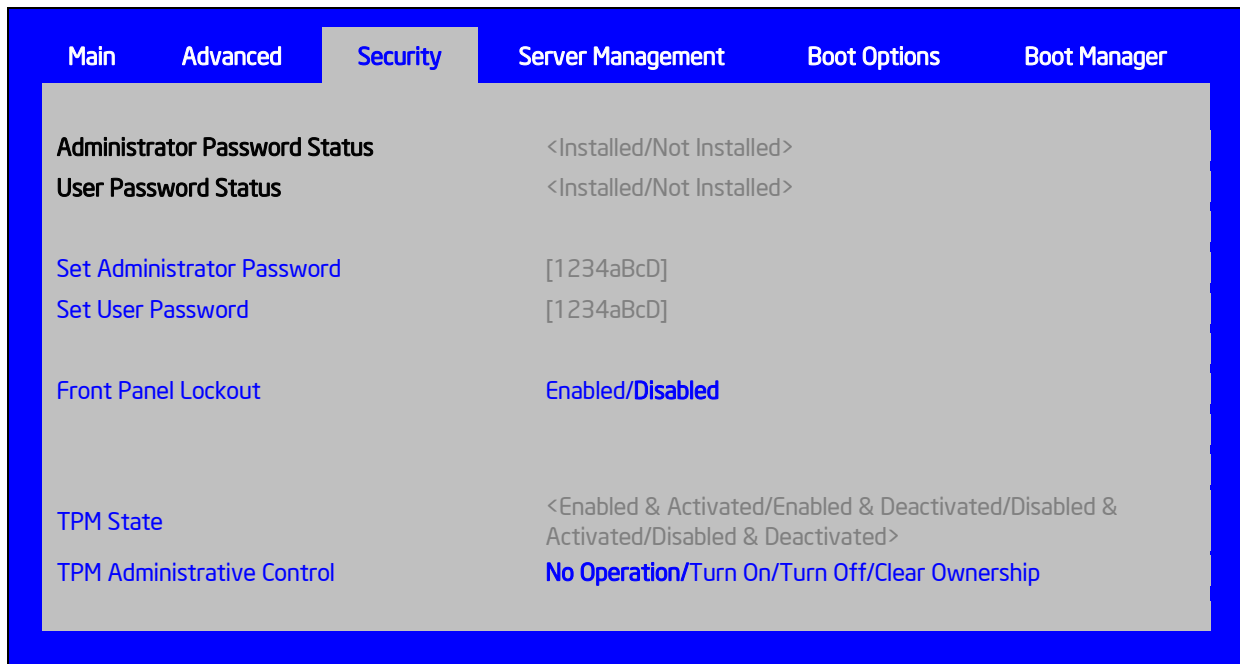
Figure 19. Setup Utility – TPM Configuration Screen

Table 14. Setup Utility – Security Configuration Screen Fields

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

5.3   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 Specifications, Revision 1.2.*

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

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 E5-4600, 2600, 1600 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.
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Chassis intrusion detection (dependent on platform support).
- Basic fan control using Control version two SDRs.
- Fan redundancy monitoring and support.
- Power supply redundancy monitoring and support.
- Hot-swap fan support.
- Acoustic management: Support for multiple fan profiles
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval.
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.
- Power state retention
- Power fault analysis
- Intel® Light-Guided Diagnostics
- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support
- E-mail alerting
- Embedded web server
- Integrated KVM
- Integrated Remote Media Redirection
- Lightweight Directory Access Protocol (LDAP) support
- Intel® Intelligent Power Node Manager support

6.1.3   New Manageability Features

Intel® S2600CO 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 on Intel® Xeon® Processors E5 2600 Families:

- 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:
  - o 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
  - o Human-readable SEL
  - o Additional system configurability
  - o Additional system monitoring capability
  - o Enhanced on-line help
- Enhancements to KVM redirection
  - o 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:
  - o Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC.

> o   BMC System Management Health Monitoring

6.2   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. ▪ The front panel power LED is on (not controlled by the BMC). ▪ The fans spin at the normal speed, as determined by sensor inputs. ▪ Front panel buttons work normally.
S1	Yes	Sleeping. Hardware context is maintained; equates to processor and chipset clocks being stopped. ▪ The front panel power LED blinks at a rate of 1 Hz with a 50% duty cycle (not controlled by the BMC). ▪ The watchdog timer is stopped. ▪ The power, reset, front panel NMI, and ID buttons are unprotected. ▪ Fan speed control is determined by available SDRs. Fans may be set to a fixed state, or basic fan management can be applied. 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	Not supported.
S4	No	Not supported.
S5	Yes	Soft off. ▪ The front panel buttons are not locked. ▪ The fans are stopped. ▪ The power-up process goes through the normal boot process. ▪ The power, reset, front panel NMI, and ID buttons are unlocked.

6.3   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

Source	External Signal Name or Internal Subsystem	Capabilities
CPU Thermal	CPU Thermtrip	Turns power off
WOL (Wake On LAN)	LAN	Turn power on

6.4   BMC Watchdog

The BMC FW is consistently called 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 greater than 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 a 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: There will no SEL event and front panel LED status change for BMC reset due to Linux* "kernel panic".

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 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.

6.5   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.

FSB2 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 it 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 internal.

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.

6.6   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 C – Integrated BMC Sensor Tables for additional sensor information.

6.7   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

6.8   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.

6.9   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 the time the signal is driven high in each pulse.

6.9.1   Thermal and Acoustic Management

The Intel® Server Board S2600CO 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.

6.9.1.1   Set 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].

6.9.1.2   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].

6.9.1.3   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].

6.9.1.4   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].

6.9.1.5   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 *System Technical Product Specification* for the board in Intel® chassis thermal and acoustic management
3. Refer to *Fan Control Whitepaper* for the board in third party chassis fan speed control customization.

6.9.1.6   Thermal Sensor Input for Fan Speed Control

The BMC uses various IPMI sensors as inputs to FSC. Some of the sensors are actual physical sensor and some are "virtual" sensors derived from calculation.

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

- Front Panel Temperature Sensor[1]
- Baseboard Temperature Sensor[2]

- CPU Margin Sensors[3,5,6]
- DIMM Thermal Margin Sensors[3,5]
- Exit Air Temperature Sensor[1, 4, 8]
- PCH Temperature Sensor[4,6]
- On-board Ethernet Controller Temperature Sensors[4, 6]
- Add-In Intel® SAS/IO Module Temperature Sensors[4, 6]
- PSU Thermal Sensor[4, 9]
- CPU VR Temperature Sensors[4, 7]
- DIMM VR Temperature Sensors[4, 7]
- Integrated BMC Temperature Sensor[4, 7]
- Global Aggregate Thermal Margin Sensors[3, 8]

Note:
1. For fan speed control in Intel® chassis
2. For fan speed control in third 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*

A simplified model is shown as below which gives a high level concept of the fan speed control structure.
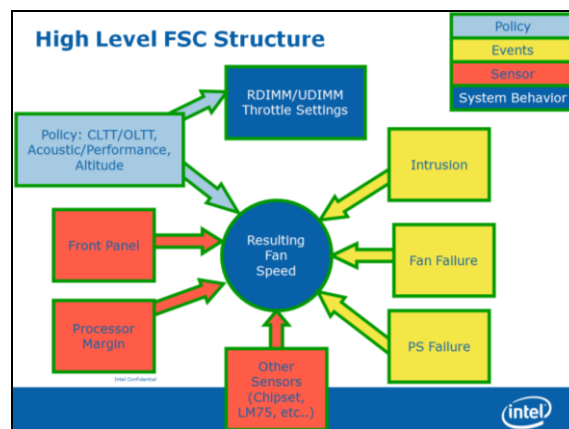
Figure 20. High-level Fan Speed Control Process

6.9.2   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 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.

6.10 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) I^2C interface
- LAN interface using the IPMI-over-LAN protocols

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

Table 17. Standard ID Channel Assignments

Channel ID	Interface	Support Sessions
0	Primary IPMB	No
1	LAN 1	Yes
2	LAN 2	Yes
3	LAN 3[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.

6.10.1   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.

4. 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.

6.10.2   IPMB Communication Interface

The IPMB communication interface used 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.

6.10.3   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.

6.10.3.1 RMCP/Alert Standard Forum (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.*

6.10.3.2 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.

6.10.3.2.1 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 is 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 operations is configured from a BIOS setup option.

6.10.3.2.2 Dedicated Management Channel

An additional LAN channel dedicated to BMC usage and not available to host SW is supported by 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 (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 system 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 purpose. When this is enabled, that port is hidden from the OS.

6.10.3.2.3   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 ports and the optional dedicated RMM4 add-in management NIC
- Two on-board NIC ports and the 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 3 MAC addresses allocated specifically for manageability.

The server board has seven MAC addresses programmed at the factory. 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 printed MAC address on the server board and/or server system is assigned to NIC1 on the server board.
For security reasons, embedded LAN channels have the following default setting:

- 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 LAN2 (Baseboard NIC port) ----- 100Mb (10Mb in DC off state)
- BMC LAN3 (Dedicated NIC) ----- 1000Mb (10Mb in DC off state)

6.10.3.3   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 versus 4 bytes for IPv4.

An IPv6 prefix is 0 to 128 bits whereas IPv4 has a 4 byte subnet mask.

There are two variants of automatic IP Address Source configuration versus 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.

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 a 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.

6.10.3.4   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 by IPMI methods as well as by the BMC's Embedded UI, allowing the user to specify the physical LAN links constitute the redundant network paths or physical LAN links constitute different network paths. BMC will support only 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.

6.10.3.5   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.

6.10.3.5.1   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.

6.10.3.5.2   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.

6.10.3.5.3   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 100ms.

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.

6.10.3.5.4   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."

6.10.3.6   DHCP BMC Hostname

The BMC allows setting a DHCP Hostname using the *Set/Get LAN Configuration Parameter* 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 "Upgrade 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 purpose.

6.10.4   Address Resoluton Protocol (ARP)

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

6.10.5   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.

6.10.6   Virtual Local Area Network (VLAN)

The BMC supports VLAN as defined by *IPMI 2.0 Specification*. 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 IDs 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 IDs are 1 through 4094, VLAN IDs 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 12 (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 bit 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 cod 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 (prior to 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 management application to configure a different IP address if that is not suitable for VLAN.
- If the BMC is configured 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 management application to configure a different IP address if that is not suitable for VLAN.

6.10.7　　　Secure Shell (SSH)

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

6.10.8　　　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 pre-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.

6.10.9　　　Platform Event Filter

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 18. 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

Event Filter Number	Offset Mask	Events
10	Power down, power cycle, and reset	Watchdog timer
11	OEM system boot event	System restart (reboot)
12	Drive Failure, Predicated 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.

6.10.10   LAN Alterting

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.

6.10.10.1   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 Module 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.

6.10.11   Altert 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.

6.10.11.1   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.

6.10.12   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), version 1.0*, 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 from 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 a 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.

6.10.13   Embeded Web Server

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

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

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

A partial list of 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
- 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 Windows 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.
- 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.

6.10.14   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.

- Virtual 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 will not 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 will not reflect the chassis LED software blinking by the software command as there is no mechanism to get the chassis ID Led status.

- Only Infinite chassis ID ON/OFF by the 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 Intel® S1400/S1600/S2400/S2600 Server Platforms. It can be used in future.

6.10.15   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 from 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 timestamped for the last 3 occurrences of the error conditions.
  - PCI error registers
  - MSR registers
  - Integrated Memory Controller (IMC) and Integrated I/O (IIO) modules 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 from the 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 from IPMI interfaces. The platform debug file can be accessed by 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.

6.10.15.1   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.

6.10.15.2   Output Data Availability

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

6.10.15.3   Output Data Categories

The following tables list the data to be provided in the diagnostic output. For items in Table 19, this data is collected on detection of CATERR, ERR2, PERR, SERR, and SMI timeout. The data in Table 20 is accumulated for the three most recent overall errors.

Table 19. 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	
	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

Table 20. Additional Diagnostics on Error.

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

6.10.16   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 spec 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.

6.10.17   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 do on) from 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. Microsoft Windows* and Novell* LDAP are not supported.

7.   Advanced Management Features 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. The Intel® RMM4 is available as two option kits:

Table 21. Intel® RMM4 options kits

Intel® Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE	Intel® Remote Management Module 4 Lite	RMM4 Lite Activation Key	Enable KVM and Media redirection from onboard NIC
AXXRMM4	Intel® Remote Management Module 4	RMM4 Lite Activation Key Dedicated NIC Port Module	Dedicated NIC for management traffic. Higher bandwidth connectivity for KVM and media Redirection.

Table 22. Enabling Advanced Management Features

Manageability Hardware	Benefits
Intel® Integrated BMC	Comprehensive IPMI based base manageability features
Intel® Remote Management Module 4 – Lite Package contains one module – 1- Key for advance Manageability features.	No dedicated NIC for management Enables KVM and media redirection from onboard NIC
Intel® Remote Management Module 4 Package includes 2 modules – 1 - key for advance features 2 - Dedicated NIC for management	Dedicated NIC for management traffic. Higher bandwidth connectivity for KVM and media Redirection.

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

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
- An option to select a mouse configuration based on the OS type.
- Supports user definable keyboard macros.

KVM redirection feature supports the following resolution and refresh rates:

- 640x480 at 60Hz, 72Hz, 85Hz, 100Hz
- 800x600 at 60Hz, 72Hz, 75Hz, 85Hz
- 1024x768 at 60Hz, 72Hz, 75Hz, 85Hz
- 1280X960 at 60Hz
- 1280x1024 at 60Hz
- 1600x1200 at 60Hz
- 1920x1080 at 60Hz
- 1920x1200 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, 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. A 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 shall support redirection for a minimum of two virtual devices concurrently with any combination of devices. As an example, a user could redirect two CD or two USB devices.
- 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 Integrated BMC reset (for example, due to an Integrated BMC reset after Integrated 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

8.   On-board Connector/Header Overview

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

8.1   Power Connectors

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

8.1.1   Main Power

Main server board power is supplied from a 24-pin power connector. The connector is labeled as "MAIN PWR" on the server board. The following tables provide the pin-out of "MAIN PWR" connector.

Table 23. Main Power Connector Pin-out ("MAIN PWR")

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

8.1.2   CPU Power Connectors

On the server board there are two white 8-pin 12V CPU power connectors labeled "CPU_1 PWR" and "CPU_2 PWR". The following table provides the pin-out for both connectors.

Table 24. CPU Power Connector Pin-out ("CPU_1 PWR" and "CPU_2 PWR")

Pin	Signal name	Pin	Signal name
1	GND	5	+12V
2	GND	6	+12V
3	GND	7	+12V
4	GND	8	+12V

8.1.3   PCIe Card Power Connectors

The server board includes one 4-pin power connectors that provide support for add-in cards that require more power than is supported from the PCIe slots direct. The connector is labeled as OPT_12V_PWR.

Table 25. PCIe Card Power Connector Pin-out ("OPT_12V_PWR")

Pin	Signal name	Pin	Signal name
1	GND	3	+12V
2	GND	4	+12V

8.2   Front Panel Headers and Connectors

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

8.2.1   SSI Front Panel Header

Included on the front edge of the server board is a 30-pin SSI compatible front panel header which provides for various front panel features including:

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

On the server board, this header is labeled "SSI FRONT PANEL". The following table provides the pin-out for this header.

Table 26. SSI Front Panel Header Pin-out ("SSI Front Panel")

Pin	Signal Name	Pin	Signal Name
1	SB3.3V	2	SB3.3V
	KEY	4	SB5V
5	Power LED Cathode	6	System ID LED Cathode
7	3.3V	8	System Fault LED Anode
9	HDD Activity LED Cathode	10	System Fault LED Cathode
11	Power Switch	12	NIC#1 Activity LED
13	GND (Power Switch)	14	NIC#1 Link LED
15	Reset Switch	16	I^{2}C SDA
17	GND (Reset/ID/NMI Switch)	18	I^{2}C SCL
19	System ID Switch	20	Chassis Intrusion
21	Pull Down	22	NIC#2 Activity LED
23	NMI to CPU Switch	24	NIC#2 Link LED
	KEY		KEY
27	NIC#3 Activity LED	28	NIC#4 Activity LED

Pin	Signal Name	Pin	Signal Name
29	NIC#3 Link LED	30	NIC#4 Link LED

8.2.1.1 Power/Sleep Button and LED Support

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

Table 27. Power/Sleep LED Functional States

State	Power Mode	LED	Description
Power-off	Non-ACPI	Off	System power is off, and the BIOS has not initialized the chipset.
Power-on	Non-ACPI	On	System power is on
S5	ACPI	Off	Mechanical is off, and the operating system has not saved any context to the hard disk.
S4	ACPI	Off	Mechanical is off. The operating system has saved context to the hard disk.
S3-S1	ACPI	Slow blink[1]	DC power is still on. The operating system has saved context and gone into a level of low-power state.
S0	ACPI	Steady on	System and the operating system are up and running.

8.2.1.2 System ID Button and LED Support

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

8.2.1.3 System Reset Button Support

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

8.2.1.4 NMI Button Support

When the NMI button is pressed, it puts the server in a halt state and causes the BMC to issue a non-maskable interrupt (NMI). This can be useful when performing diagnostics for a given issue where a memory download is necessary to help determine the cause of the problem. Once an NMI has been generated by the BMC, the BMC does not generate another NMI until the system has been reset or powered down.

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

The following table describes behavior regarding NMI signal generation and event logging by the BMC.

Table 28. NMI Signal Generation and Event Logging

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

8.2.1.5   NIC Activity LED Support

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

8.2.1.6   Hard Drive Activity LED Support

The drive activity LED on the front panel indicates drive activity from the on-board hard disk controllers. The server board also provides a header giving access to this LED for add-in controllers.

8.2.1.7   System Status LED Support

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

Table 29. System Status LED State Definitions

Color	State	Criticality	Description
Off	System is not operating	Not ready	1. System is powered off (AC and/or DC). 2. System is in EuP Lot6 Off Mode. 3. System is in S5 Soft-Off State. 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: ▪ Redundancy loss, such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. ▪ Fan warning or failure when the number of fully operational fans is more than minimum number needed to cool the system. ▪ Non-critical threshold crossed – Temperature (including HSBP temp), voltage, input power to power supply, output current for main power rail from power supply and Processor Thermal Control (Therm Ctrl) sensors. ▪ Power supply predictive failure occurred while redundant power

Color	State	Criticality	Description
			supply configuration was present.
			▪ Unable to use all of the installed memory (one or more DIMMs failed/disabled but functional memory remains available)
			▪ 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.
			▪ Uncorrectable memory error has occurred in memory Mirroring Mode, causing Loss of Redundancy.
			▪ Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in fully redundant RAS Mirroring Mode.
			▪ Battery failure.
			▪ 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
			▪ 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.
			▪ BMC Watchdog has reset the BMC.
			▪ Power Unit sensor offset for configuration error is asserted.
			▪ HDD HSC is off-line or degraded.
Amber	~1 Hz blink	Non-critical - System is operating in a degraded state with an impending failure warning, although still functioning	Non-fatal alarm – system is likely to fail: ▪ Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply and PROCHOT (Therm Ctrl) sensors. ▪ VRD Hot asserted. ▪ Minimum number of fans to cool the system not present or failed ▪ Hard drive fault ▪ Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present) ▪ In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window ▪ 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: ▪ CPU CATERR signal asserted ▪ MSID mismatch detected (CATERR also asserts for this case). ▪ CPU 1 is missing ▪ CPU Thermal Trip ▪ No power good – power fault ▪ DIMM failure when there is only 1 DIMM present and hence no good memory present1. ▪ Runtime memory uncorrectable error in non-redundant mode. ▪ DIMM Thermal Trip or equivalent ▪ SSB Thermal Trip or equivalent ▪ CPU ERR2 signal asserted ▪ BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition)

Color	State	Criticality	Description
			▪ Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition)
			▪ 240VA fault
			Fatal Error in processor initialization: ▪ Processor family not identical
			▪ Processor model not identical
			▪ Processor core/thread counts not identical
			▪ Processor cache size not identical
			▪ Unable to synchronize processor frequency
			▪ Unable to synchronize QPI link frequency

8.2.2   Front Panel USB Connector

The server board includes a 10-pin connector, that when cabled, can provide up to two USB ports to a front panel. On the server board the connector is labeled "USB5-6". The following table provides the connector pin-out.

Table 30. Front Panel USB Connector Pin-out (USB5-6)

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

8.2.3   Intel® Local Control Panel Connector

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

Table 31. Intel® Local Control Pane Connector Pin-out (LCP)

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

8.3   On-Board Storage Connectors

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

8.3.1   SATA Only Connectors: 6 Gbps

The server board includes two white single port SATA only connectors capable of transfer rates of up to 6Gb/s. On the server board these connectors are labeled as SATA_0 and SATA_1. The following table provides the pin-out for both connectors.

Table 32. SATA Only Connector Pin-out (SATA_0 and SATA_1)

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

Note: As an option, SATA_0 supports vertical, high profile SATA_DOM.

8.3.2   SATA/SAS Connectors

The server board includes eight SATA/SAS connectors. On the server board, these connectors are labeled as SATA/SAS_ 0 to SATA/SAS_7. By default, only the connector labeled SATA/SAS_0 to SATA/SAS_3 are enabled and has support for up to four SATA ports capable of transfer rates of up to 3Gb/s. The connector labeled SATA/SAS_4 to SATA/SAS_7 is only enabled when an optional Intel® RAID C600 Upgrade Key is installed. See Table 10 for a complete list of supported storage upgrade keys. The following tables provide the pin-out for each connector.

Table 33. SATA/SAS Connector Pin-out (SATA/SAS_0 to SATA/SAS_7)

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

8.3.3   SAS SGPIO Connectors

Table 34. SAS SGPIO Connector Pin-out (SAS_SPGIO_0 and SAS_SPGIO_1)

Pin	Signal Name	Pin#	Signal Name
1	CLOCK	2	LOAD
3	GND	4	DATAOUT
5	DATAIN		

8.3.4   Intel® RAID C600 Upgrade Key Connector

The server board provides one connector to support Intel® RAID C600 Upgrade Key. The Intel® RAID C600 Upgrade Key is a small PCB board that enables different versions of RAID 5 software stack and/or upgrade from SATA to SAS storage functionality. On the server board, the connector is labeled as STRO UPG KEY. The pin configuration of connector is identical and defined in the following table.

Table 35. Intel® RAID C600 Upgrade Key Connector Pin-out (STRO UPG KEY)

Pin	Signal Name
1	GND
2	FM_PBG_DYN_SKU_KEY
3	GND
4	FM_SSB_SAS_SATA_RAID_KEY

8.3.5   HSBP_I²C Header

Table 36. HSBP_I^{2}C Header Pin-out (HSBP_I^{2}C)

Pin	Signal Name
1	SMB_HSBP_3V3STBY_DATA
2	GND
3	SMB_HSBP_3V3STBY_CLK

8.3.6   HDD LED Header

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

Table 37. Hard Drive Acitivity Header Pin-out (HDD_LED)

Pin	Signal Name
1	LED_HDD_ACT_N
2	TP_LED_HDD_ACT

8.3.7   Internal Type-A USB Connector

The server board includes one internal Type-A USB connector labeled USB_4. The following table provides the pin-out for this connector.

Table 38. Internal Type-A USB 2.0 Connector Pin-out (USB_4)

Pin	Signal Description
1	+5V
2	USB_N
3	USB_P
4	GND

8.3.8   Internal 2mm Low Profile eUSB SSD Connector

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

Table 39. Internal eUSB Connector Pin-out (eUSB_SSD)

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

8.4   Management and Security Connectors

8.4.1   RMM4_Lite Connector

A 7-pin Intel® RMM4 Lite Connector is included on the server board to support the optional Intel® Remote Management Module 4. There is no support for third-party management cards on this server board. On the server board this connector is labeled as RMM4_Lite. The following table provides the pin-out for this connector.

Table 40. RMM4_Lite Connector Pin-out (RMM4_Lite)

Pin	Signal Name	Pin	Signal Name
1	P3V3_AUX	2	SPI_IBMC_BK_DI
		4	SPI_IBMC_BK_CLK
5	SPI_IBMC_BK_DO	6	GND
7	SPI_IBMC_BK_CS_N	8	GND

8.4.2   RMM4_NIC connector

Table 41. RMM4_NIC Connector Pin-out (RMM4_NIC)

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	GND	24	RXD_3
25	GND	26	TX_CLK
27	GND	28	RX_CLK
29	GND	30	PRESENT#

8.4.3   TPM Connector

Table 42. TPM Connector Pin-out (TPM)

Pin	Signal Name	Pin	Signal Name
		2	LPC_LAD<1>

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

8.4.4   PMBus* Connector

Table 43. PMBus* Connector Pin-out (SMB_PMBUS*)

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

8.4.5   Chassis Intrustion Header

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

Table 44. Chassis Intrusion Header Pin-out (CHAS INTR)

Pin	Signal Description
1	FP_CHASSIS_INTRUSION
2	GND

Table 45. Chassis Instrusion Header State Description

Header state	Description
PINS 1 and 2 CLOSED	IBMC_CHASSIS_N is pulled HIGH. Chassis cover is closed.
PINS 1 and 2 OPEN	IBMC_CHASSIS_N is pulled LOW. Chassis cover is removed.

8.4.6   IPMB Connector

Table 46. IPMB Connector Pin-out (IPMB)

Pin	Signal Description
1	SMB_IPMB_5VSTBY_DATA
2	GND
3	SMB_IPMB_5VSTBY_CLK
4	P5V_STBY

8.5   Fan Connectors

The server board provides support for nine fans. Seven of them are system cooling fans; two of them are CPU fans.

8.5.1   System FAN Connectors

The server board provides support for seven system cooling fans. Each connector is monitored and controlled by on-board platform management. On the server board, each system fan connector is labeled SYS_FAN_#, where # = 1 thru 7. The six system cooling fan connectors near the front edge of the board (SYS_FAN_1 to SYS_FAN_6 are 6-Pin connectors; the one system cooling fan near rear edge of the board is a 4-Pin connectors (SYS_FAN_7). Following table provides the pin-out for all system fan connectors.

Table 47. 6-pin System FAN Connector Pin-out (SYS_FAN_1 to SYS_FAN_6)

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

Table 48. 4-pin System FAN Connector Pin-out (SYS_FAN_7)

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

8.5.2   CPU FAN Connector

The server board also provides support for two CPU cooling fans. Each 4-pin connector is monitored and controlled by platform management. On the server board each CPU fan connector is labeled as CPU_1FAN and CPU_2FAN. The following table provides the pin-out for both fan connectors.

Table 49. CPU Fan Connector Pin-out (CPU_1 FAN and CPU_2 FAN)

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

8.6   Serial Port Connectors

The server board includes two serial port connectors.

8.6.1   Serial Port A connector (DB9)

Serial-A is an external RJ45 type connector labeled as SERIAL_A and has the following pin-out configuration:

Table 50. Serial Port A Connector Pin-out (SERIAL_A)

Pin	Signal Name	Pin	Signal Name
1	SPA_DCD	2	SPA_SIN_N
3	SPA_OUT_N	4	SPA_DTR

Pin	Signal Name	Pin	Signal Name
5	GND	6	SPA_DSR
7	SPA_RTS	8	SPA_CTS
9	SPA_RI	■	

8.6.2   Serial Port B Connector

Serial-B is an internal 10-pin DH-10 connector labeled as SERIAL_B and has the following pin-out:

Table 51. Serial-B Connector Pin-out (SERIAL_B)

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

8.6.3   Video Connector

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

Table 52. Rear VGA Video Connector Pinout (VGA)

Pin	Signal Name
1	RED
2	GREEN
3	BLUE
4	N/C
5	GND
6	GND
7	GND
8	GND
9	P5V (fuse not populated)
10	GND
11	N/C
12	DDC_SDA
13	HSYNC
14	VSYNC
15	DDC_SCL

Figure 21. Other Connectors and Headers

8.7   Other Connectors and Headers
8.7.1   FAN BOARD_I²C Connector

Note: Only Intel® Server Board **S2600CO4** includes this FAN BOARD_I²C connector.

The server board includes a 4-pin I²C header that is intend to be used for third party fan control circuits using a Maxim 72408 controller. On the server board, this header is labeled FAN BOARD_I²C. The header has the following pin-out.

Table 53. HSBP 4-PIN I²C BUS Connector pin out (FAN BOARD_I²C)

Pin	Signal Name
1	SMB_3V3SB_DAT
2	GND
3	SMB_3V3SB_CLK
4	P3V_SB

8.7.2   IEEE 1394b Connector

Note: Only Intel® Server Board **S2600COE** includes these two IEEE 1394b connectors.

The server board includes two red 10-pin IEEE 1394b connectors. On the server board, this header is labeled FRONT_1394B_0 and FRONT_1394B_1. The header has the following pin-out.

Table 54. IEEE 1394b Connector pin out

Pin	Signal Name	Pin	Signal Name
1	GND		
3	+12V (Fused)	4	+12V (Fused)
5	TPB-	6	TPB+
7	TPB_REF	8	TPA_REF
9	TPA_	10	TPA+

9. Reset and Recovery Jumpers

The server board includes several jumper blocks which are used to as part of a process to restore a board function back to a normal functional state. The following diagram and sections identify the location of each jumper block and provides a description of their use.

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

Figure 22. Server Board Jumper Block Locations (J1E2, J1E3, J1E4, J1E6, J2J2)

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

Table 55. Server Board Jumpers (J1E2, J1E3, J1E4, J1E6, J1J2)

Jumper Name	Pins	System Results
J1E6: BIOS Recovery	1-2	Pins 1-2 should be connected for normal system operation. **(Default)**
	2-3	The main system BIOS does not boot with pins 2-3 connected. The system only boots from EFI-bootable recovery media with a recovery BIOS image present.
J2J2: BIOS	1-2	These pins should have a jumper in place for normal system operation. **(Default)**

Jumper Name	Pins	System Results
Default (that is, CMOS Clear)	2-3	If pins 2-3 are connected when AC power unplugged, the CMOS settings clear in 5 seconds. Pins 2-3 should not be connected for normal system operation.
J1E3: ME Force Update	1-2	ME Firmware Force Update Mode – Disabled **(Default)**
	2-3	ME Firmware Force Update Mode – Enabled
J1E4: BMC Force Update	1-2	BMC Firmware Force Update Mode – Disabled **(Default)**
	2-3	BMC Firmware Force Update Mode – Enabled
J1E2: Password Clear	1-2	These pins should have a jumper in place for normal system operation. **(Default)**
	2-3	To clear administrator and user passwords, power on the system with pins 2-3 connected. The administrator and user passwords clear in 5-10 seconds after power on. Pins 2-3 should not be connected for normal system operation.

9.1   BIOS Default (that is, CMOS Clear) and Password Reset Usage Procedure

The BIOS Default (that is, CMOS Clear) and Password Reset recovery features are designed such that the desired operation can be achieved with minimal system downtime. The usage procedure for these two features has changed from previous generation Intel® server boards. The following procedure outlines the new usage model.

9.1.1   Set BIOS to default (that is, Clearing the CMOS)

To clear the CMOS, perform the following steps:

1. Power down the server. Do not unplug the power cord.

2. Open the server chassis. For instructions, see your server chassis documentation.

3. Move jumper from the default operating position (covering pins 1 and 2) to the reset/clear position (covering pins 2 and 3).

4. Wait five seconds.

5. Remove AC power.

6. Move the jumper back to the default position (covering pins 1 and 2).

7. Close the server chassis.

8. Power up the server.

The CMOS is now cleared and can be reset 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 preferred settings.

9.1.2   Clearing the Password

To clear the password, perform the following steps:

1. Power down the server. Do not unplug the power cord.

2. Open the chassis. For instructions, see your server chassis documentation.

3. Move jumper 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 wait 10 seconds or until POST completes.

6. Power down the server.

7. Open the chassis and move the jumper back to the default position (covering pins 1 and 2).

8. Close the server chassis.

9. Power up the server.

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

9.2   Integrated BMC Force Update Procedure

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

1. Power down and remove the AC power cord.

2. Open the server chassis. For instructions, see your server chassis documentation.

3. Move jumper from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).

4. Close the server chassis.

5. Reconnect the AC cord and power up the server.

6. Perform the Integrated BMC firmware update procedure as documented in the *README.TXT* file that is included in the given Integrated BMC firmware update package. After successful completion of the firmware update process, the firmware update utility may generate an error stating that the Integrated BMC is still in update mode.

7. Power down and remove the AC power cord.

8. Open the server chassis.

9. Move jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).

10. Close the server chassis.

11. Reconnect the AC cord and power up the server.

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

9.3   ME Force Update Jumper

When performing the standard ME force update procedure, the update utility places the ME into an update mode, allowing the ME to load safely onto the flash device. In the unlikely event ME

firmware update process fails due to ME not being in the proper update state, the server board provides an Integrated BMC Force Update jumper, which forces the ME into the proper update state. The following procedure should be completed in the event the standard ME firmware update process fails.

1. Power down and remove the AC power cord.

2. Open the server chassis. For instructions, see your server chassis documentation.

3. Move jumper from the default operating position (covering pins 1 and 2) to the enabled position (covering pins 2 and 3).

4. Close the server chassis.

5. Reconnect the AC cord and power up the server.

6. Perform the ME firmware update procedure as documented in the *README.TXT* file that is included in the given ME firmware update package (same package as BIOS).

7. Power down and remove the AC power cord.

8. Open the server chassis.

9. Move jumper from the enabled position (covering pins 2 and 3) to the disabled position (covering pins 1 and 2).

10. Close the server chassis.

11. Reconnect the AC cord and power up the server.

9.4   BIOS Recovery Jumper

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. Light Guided Diagnostics

The server board includes several on-board LED indicators to aid troubleshooting various board level faults. This section provides a description of the location and function of each LED on the server boards.

10.1 5 Volt 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 must be present when the system is powered down. The LED 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 23. 5 Volt Stand-by Status LED Location

10.2 Fan Fault LEDs

Fan fault LEDs are present for the two CPU fans and the one rear system fan. The fan fault LEDs illuminate when the corresponding fan has fault.
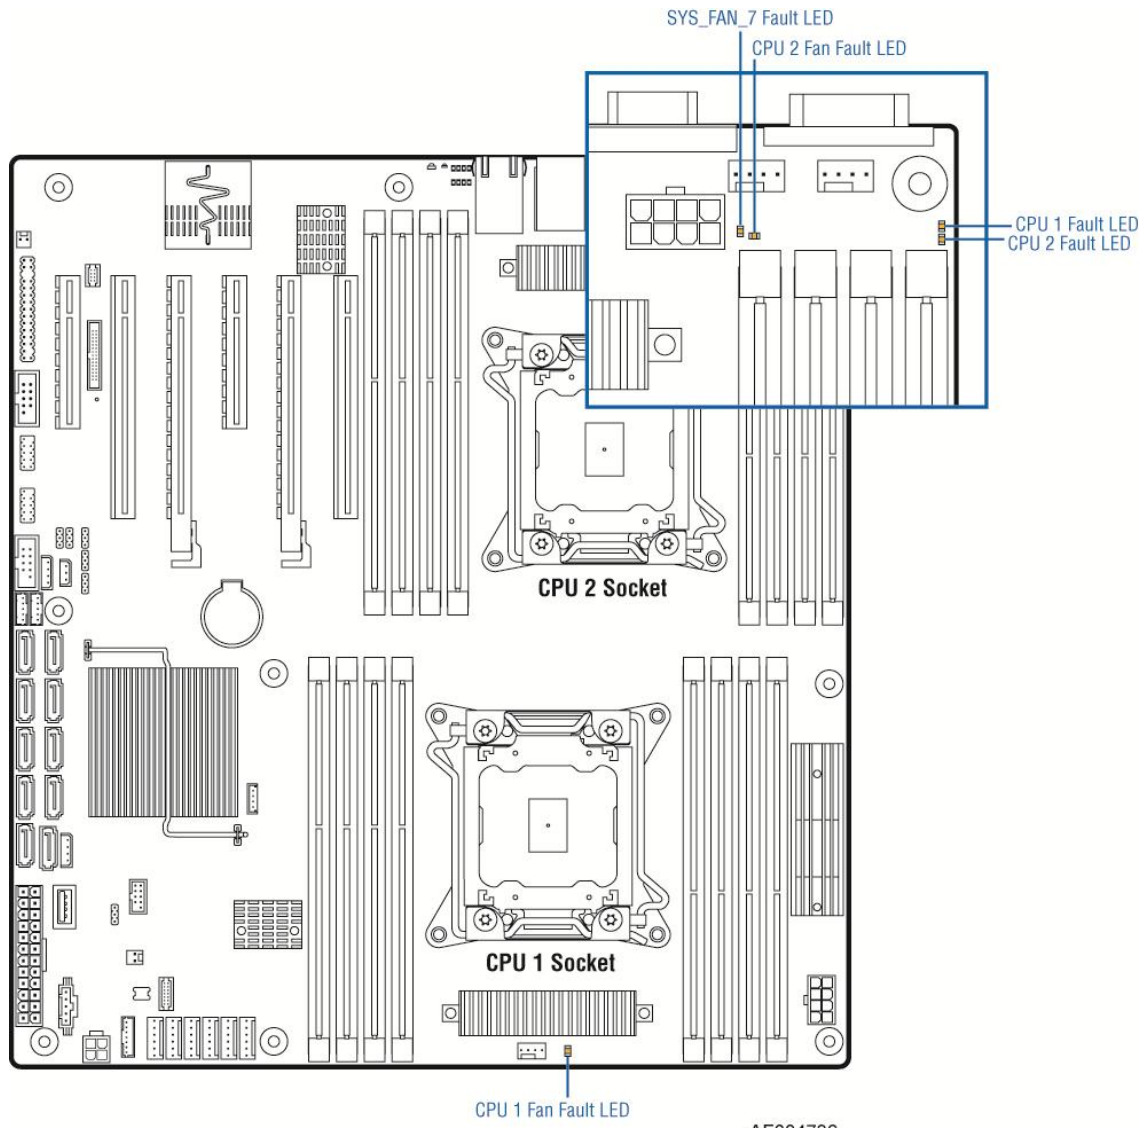
Figure 24. Fan Fault LED's Location

10.3 DIMM Fault LEDs

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

Figure 25. DIMM Fault LED's Location

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

The server boards provide LEDs for system ID, system status and POST code. These LEDs are located in the rear I/O area of the server board as shown in the following figure.
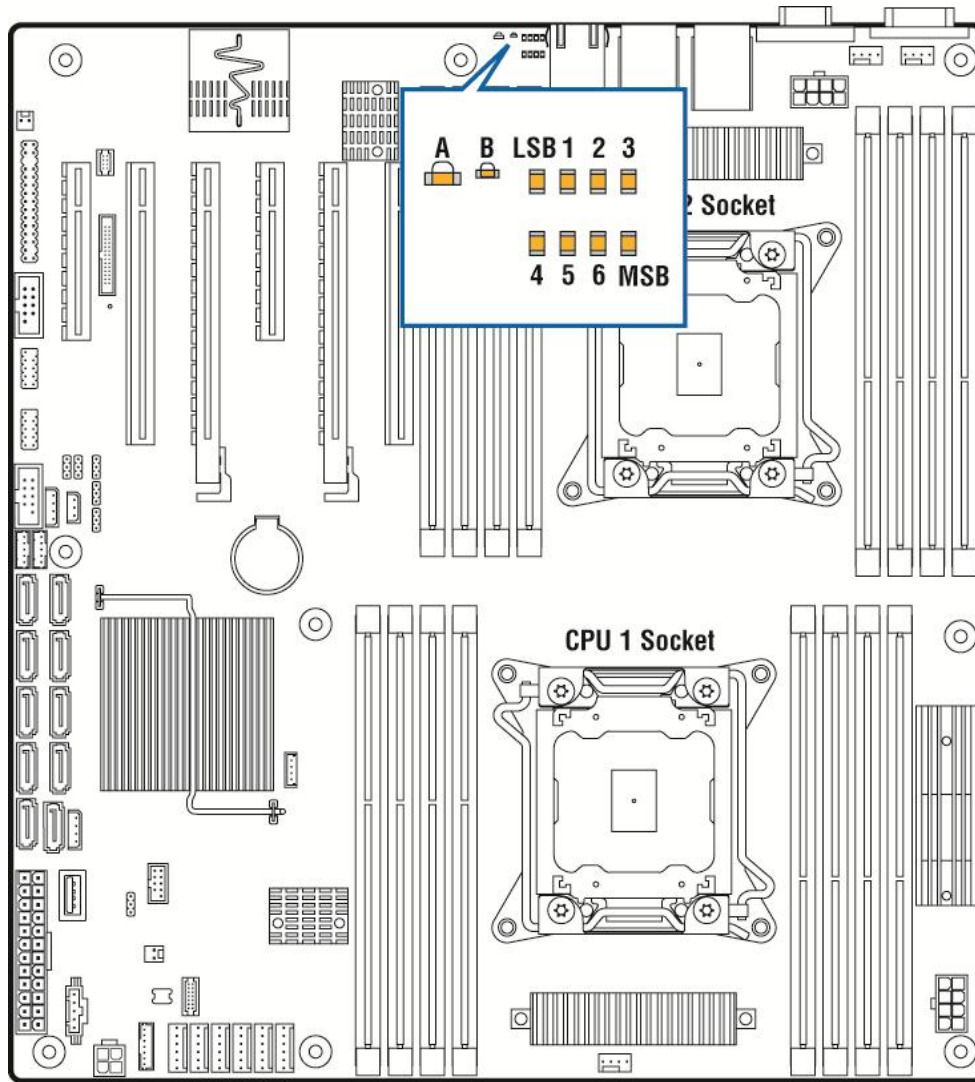
AF004740

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

Figure 26. Location of System Status, System ID and POST Code Diagnostic LEDs

10.4.1   System ID LED

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

▪ By pressing the System ID Button on the system front control panel, the ID LED displays a solid blue color until the button is pressed again.

▪ By issuing the appropriate hex IPMI "Chassis Identify" value, the ID LED will either blink blue for 15 seconds and turn off or will blink indefinitely until the appropriate hex IPMI Chassis Identify value is issued to turn it off.

10.4.2   System Status LED

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

Table 56. System Status LED

Color	State	Criticality	Description
Green	Solid on	System OK	System booted and ready.
Green	Blink	Degraded	System degraded ▪ Non-critical temperature threshold asserted ▪ Non-critical voltage threshold asserted o Non-critical fan threshold asserted o Fan redundancy lost, sufficient system cooling maintained. This does not apply to non-redundant systems. o Power supply predictive failure o Power supply redundancy lost. This does not apply to non-redundant systems. ▪ Correctable errors over a threshold of 10 and migrating to a mirrored DIMM (memory mirroring). This indicates the user no longer has spare DIMMs indicating a redundancy lost condition. The corresponding DIMM LED should light up.
Amber	Blink	Non-critical	Non-fatal alarm – system is likely to fail: ▪ Critical temperature threshold asserted o CATERR asserted o Critical voltage threshold asserted o VRD hot asserted o SMI Timeout asserted
Amber	Solid on	Critical, non-recoverable	Fatal alarm – system has failed or shut down ▪ CPU Missing ▪ Thermal Trip asserted o Non-recoverable temperature threshold asserted o Non-recoverable voltage threshold asserted o Power fault/Power Control Failure o Fan redundancy lost, insufficient system cooling. This does not apply to non-redundant systems. o Power supply redundancy lost insufficient system power. This does not apply to non-redundant systems. **Note:** This state also occurs when AC power is first applied to the system. This indicates the BMC is booting.
Off	N/A	Not ready	▪ AC power off, if no degraded, non-critical, critical, or non-recoverable conditions exist. o System is powered down or S5 states, if no degraded, non-critical, critical, or non-recoverable conditions exist.

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

10.4.3   POST Code Diagnostic LEDs

During the system boot process, the BIOS execute 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 display the given POST code to the POST code diagnostic LEDs on the back edge of the server boards. To assist in troubleshooting a system hang during the POST process, you can use the diagnostic LEDs to identify the last POST process executed.

Table 57. POST Code Diagnostic LEDs

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

11. Environmental Limits Specification

The following table defines the Intel® Server Board S2600CO operating and non-operating environmental limits. Operation of the Intel® Server Board S2600CO 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 58. Server Board Design Specifications

Operating Temperature	0ºC to 55ºC [1] (32ºF to 131ºF)
Non-Operating Temperature	-40ºC to 70ºC (-40ºF to 158ºF)
DC Voltage	± 5% of all nominal voltages
Shock (Unpackaged)	Trapezoidal, 35g , 170 inches/sec

Shock (Packaged)
    < 20 pounds
    20 to < 40 pounds
    40 to < 80 pounds
    80 to < 100 pounds
    100 to < 120 pounds
    120 pounds
	
	36 inches
	30 inches
	24 inches
	18 inches
	12 inches
	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 (T_{CASE}) specifications. Thermal solutions not designed to provide sufficient thermal capability may affect the long-term reliability of the processor and system. The Intel® Server board S2600CO4 is designed to support the Intel® Xeon® Processor E5-2600 product family TDP guidelines up to and including 135W. The Intel® Server board S2600COE is designed to support the Intel® Xeon® Processor E5-2600 product family TDP guidelines up to and including 150W with possible configuration limits.

Disclaimer Note: Intel Corporation server boards contain a number of high-density VLSI and power delivery components that need adequate airflow to cool. Intel® ensures through its own chassis development and testing that when Intel® server building blocks are used together, the

fully integrated system will meet the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel® developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible, if components fail or the server board does not operate correctly when used outside any of their published operating or non-operating limits.

11.2 MTBF

The following is the calculated Mean Time Between Failures (MTBF) 30°C (ambient air). These values are derived using a historical failure rate and multiplied by factors for application, electrical and/or thermal stress and for device maturity. You should view MTBF estimates as "reference numbers" only.

- Calculation Model: Telcordia* Issue 1, method I case 3
- Operating Temperature: Server in 30°C ambient air
- Operating Environment: Ground Benign, Controlled
- Duty Cycle: 100%
- Quality Level: II

Table 59. MTBF Estimate

Assembly Name	Failure Rate	MTBF
Mother board	4,447.80	224,830
Integrated Circuits	946.69	1,056,313
Transistor Bipolar	127.02	7,782,934
Transistor_MOSFET	389.76	2,565,714
Diodes	31.49	31,756,974
Diodes_LED	30.00	33,330,661
Resistors	1041.23	960,398
Capacitors	341.86	2,925,173
E-Cap	461.62	2,166,305
Inductors	89.89	11,125,319
Connections	1,176.13	850,245
Misc	80.72	12,388,387

12. Power Supply Specification Guidelines

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

Note: The power supply data provided in this section is for reference purposes only. It reflects Intel®'s own DC power out requirements for a 550W power supply as used in an Intel® designed 4U server platform. The intent of this section is to provide customers with a guide to assist in defining and/or selecting a power supply for custom server platform designs that utilize the server boards detailed in this document.

For additional power supply information which are used in Intel® designed supported chassis please refer to the following Intel® documents: *Intel® Server Chassis P4000M Family Technical Product Specification.*

12.1 Power Supply DC Output Specification

12.1.1   Output Power/Currents

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

Table 60. Over Voltage Protection Limits

Parameter	Min	Max.	Peak	Unit
3.3V	0.5	18.0		A
5V	0.3	15.0		A
12V1	0.7	24.0	28.0	A
12V2	0.7	24.0	28.0	A
12V3	1.5	18.0		
3.3V	0.5	18.0		A
− 12V	0.0	0.5		A
5Vstby	0.0	3.0	3.5	A

Notes:
1.   Max combined power for all output shall not exceed 550W.
2.   Peak combined power for all outputs shall not exceed 630W for 20 seconds.
3.   Max combined power of 12V1, 12V2 and 12V3 shall not exceed 530W.
4.   Max combined power on 3.3V and 5V shall not exceed 120W.

12.1.2   Cross Loading

The power supply maintains voltage regulation limit when operated over the following cross loading conditions.

Table 61. Loading Conditions

	3.3V	5.0V	12V1	12V2	12V3	-12V	5.0Vstby	Total Power	12V Power	3.3V/5V Power
Load1	18	12.1	12	12	11.7	0	0.3	550	428	120
Load2	13.5	15	12	12	11.2	0.5	0.3	549	422	120
Load3	2.5	2	20	20	4.2	0	0.3	550	530	18

	3.3V	5.0V	12V1	12V2	12V3	-12V	5.0Vstby	Total Power	12V Power	3.3V/5V Power
Load4	2.5	2	13.1	13.1	18	0	0.3	550	530	18
Load5	0.5	0.3	15	15	6.5	0.5	3	462	438	3
Load6	16	4	1	1	3.5	0	0.3	140	66	73
Load7	16	13	1	1	9	0.5	3	271	132	118

12.1.3   Standby Output

The 5VSB output is present when an AC input greater than the power supply turn on voltage is applied.

12.1.3.1   Voltage Regulation

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

Table 62. Voltage Regulation Limits

Parameter	Tolerance	Min	Nom	Max	Units
+3.3V	- 3%/+5%	+3.20	+3.30	+3.46	Vrms
+5V	- 4%/+5%	+4.80	+5.00	+5.25	Vrms
+12V1	- 4%/+5%	+11.52	+12.00	+12.60	Vrms
+12V2	- 4%/+5%	+11.52	+12.00	+12.60	Vrms
+12V3	- 4%/+5%	+11.52	+12.00	+12.60	Vrms
- 12V	- 10%/+10%	- 13.20	-12.00	-10.80	Vrms
+5VSB	- 4%/+5%	+4.80	+5.00	+5.25	Vrms

12.1.3.2   Dynamic Loading

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

Table 63. Transient Load Requirements

Output	∆ Step Load Size (See note 2)	Load Slew Rate	Test capacitive Load
+3.3V	6.0A	0.5 A/μsec	970 μF
+5V	4.0A	0.5 A/μsec	400 μF
12V1+12V2+12V3	23.0A	0.5 A/μsec	2200 μF [1,2]
+5VSB	0.5A	0.5 A/μsec	20 μF

Notes:
1.   Step loads on each 12V output may happen simultaneously.
2.   The +12V should be tested with 2200μF evenly split between the four +12V rails
3.   This will be tested over the range of load conditions in section 12.1.2.

12.1.3.3   Capacitive Loading

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

Table 64. Capacitive Loading Conditions

Output	Min	Max	Units
+3.3V	250	5000	μF
+5V	400	5000	μF
+12V	500	8000	μF
-12V	1	350	μF
+5VSB	20	350	μF

12.1.3.4   Grounding

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

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

12.1.3.5   Residual Voltage Immunity in Standby mode

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

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

12.1.3.6   Common Mode Noise

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

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

12.1.3.7   Ripple/Noise

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

Table 65. Ripples and Noise

+3.3V	+5V	+12V 1, 2, 3	-12V	+5VSB
50mVp-p	50mVp-p	120mVp-p	200mVp-p	50mVp-p

The test set-up is as shown below:
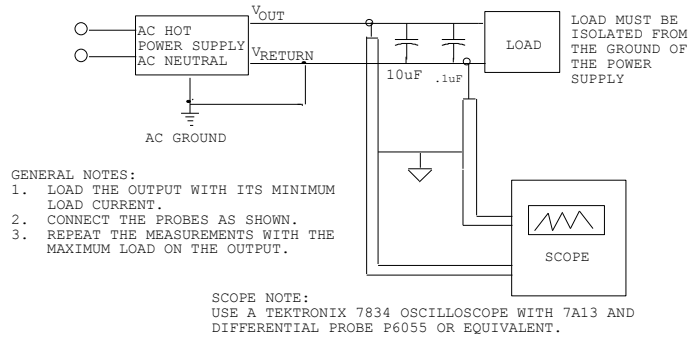
Figure 27. Differential Noise test setup

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

12.1.3.8   Timing Requirements

These are the timing requirements for the power supply operation. The output voltages rise from 10% to within regulation limits (T_{vout_rise}) within 2 to 50ms, except for 5VSB - it is allowed to rise from 1 to 25ms. The +3.3V, +5V and +12V1, +12V2, +12V3 output voltages start to rise approximately at the same time. **All outputs rise monotonically**. Each output voltage reach regulation within 50ms (T_{vout_on}) of each other during turn on the power supply. Each output voltage fall out of regulation within 400ms (T_{vout_off}) of each other during turn off. The table below shows the timing requirements for the power supply being turned on and off from the AC input, with PSON held low and the PSON signal, with the AC input applied. All timing requirements are met for the cross loading condition.

Table 66. Output Voltage Timing

Item	Description	MIN	MAX	UNITS
T_{vout_rise}	Output voltage rise time from each main output.	2	50	ms
	Output rise time for the 5Vstby output.	1	25	ms
T_{vout_on}	All main outputs must be within regulation of each other within this time.		50	ms
T_{vout_off}	All main outputs must leave regulation within this time.		400	ms
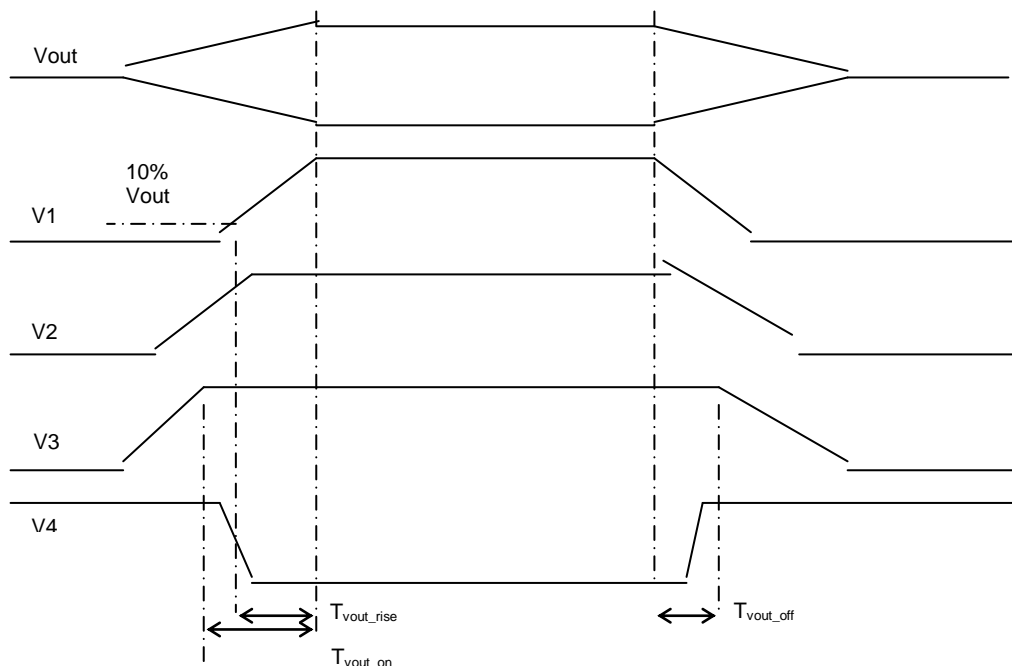
Figure 28. Output Voltage Timing

Table 67. Turn On/Off Timing

Item	Description	MIN	MAX	UNITS
$T_{sb\_on\_delay}$	Delay from AC being applied to 5VSB being within regulation.		1500	ms
$T_{ac\_on\_delay}$	Delay from AC being applied to all output voltages being within regulation.		2500	ms
T_{vout_holdup}	Time all output voltages stay within regulation after loss of AC. Tested at 75% of maximum load.	13		ms
T_{pwok_holdup}	Delay from loss of AC to de-assertion of PWOK. Tested at 75% of maximum load.	12		ms
$T_{pson\_on\_delay}$	Delay from PSON# active to output voltages within regulation limits.	5	400	ms
T_{pson_pwok}	Delay from PSON# deactivate to PWOK being de-asserted.		50	ms
T_{pwok_on}	Delay from output voltages within regulation limits to PWOK asserted at turn on.	100	500	ms
T_{pwok_off}	Delay from PWOK de-asserted to output voltages (3.3V, 5V, 12V, -12V) dropping out of regulation limits.	1		ms
T_{pwok_low}	Duration of PWOK being in the de-asserted state during an off/on cycle using AC or the PSON signal.	100		ms
T_{sb_vout}	Delay from 5VSB being in regulation to O/Ps being in regulation at AC turn on.	10	1000	ms
T_{5VSB_holdup}	Time the 5VSB output voltage stays within regulation after loss of AC.	70		ms

Figure 29. Turn On/Off Timing (Power Supply Signals)

Appendix A: Integration and Usage Tips

- When adding or removing components or peripherals from the server board, AC power must be removed. With AC power plugged into the server board, 5-V standby is still present even though the server board is powered off.

- This server board supports The Intel® Xeon® Processor E5-2600 product family with a Thermal Design Power (TDP) of up to and including 135 W. Previous generations of the Intel® Xeon® processors are not supported. Server systems using this server board may or may not meet the TDP design limits of the server board. Validate the TDP limits of the server system before selecting a processor.

  Processors must be installed in order. CPU 1 must be populated for the server board to operate.

- The server board includes a pre-installed CPU power cable harness. The cable harness must be installed and fully seated in each connector for the server board to operate.

- On the back edge of the server board are eight diagnostic LEDs that display a sequence of amber POST codes during the boot process. If the server board hangs during POST, the LEDs display the last POST event run before the hang.

- This server board only supports registered DDR3 DIMMs (RDIMMs) and unbuffered DDR3 DIMMs (UDIMMs). Mixing of RDIMMs and UDIMMs is not supported.

- For the best performance, the number of DDR3 DIMMs installed should be balanced across both processor sockets and memory channels. For example, a two-DIMM configuration performs better than a one-DIMM configuration. In a two-DIMM configuration, DIMMs should be installed in DIMM sockets A1 and D1. A six-DIMM configuration (DIMM sockets A1, B1, C1, D1, E1, and F1) performs better than a three-DIMM configuration (DIMM sockets A1, B1, and C1).

- The Intel® Remote Management Module 4 (Intel® RMM4) connector is not compatible with any previous versions of the Intel® Remote Management Module (Product Order Code – AXXRMM, AXXRMM2, and AXXRMM3).

- Clear the CMOS with AC power cord plugged. Removing the AC power before performing the CMOS clear operation causes the system to automatically power up and immediately power down after the CMOS clear procedure is followed and AC power is re-applied. If this happens, remove the AC power cord, wait 30 seconds, and then re-connect the AC power cord. Power up the system and proceed to the <F2> BIOS Setup utility to reset the desired settings.

- Normal Integrated BMC functionality is disabled with the BMC Force Update jumper set to the "enabled" position (pins 2-3). The server should never be run with the BMC Force Update jumper set in this position and should only be used when the standard firmware update process fails. This jumper should remain in the default (disabled) position (pins 1-2) when the server is running normally.

- When performing a normal BIOS update procedure, the BIOS recovery jumper must be set to its default position (pins 1-2).

Appendix B: Compatible Intel® Server Chassis

Intel® Server Chassis P4000M Family is 4U pedestal, 25" length server chassis that is designed to support Intel® Server Board S2600CO. It also provides a rackable feature.

A. 550-W Fixed Power supply
B. I/O Ports
C. Alternate RMM4 Knockout
D. PCI Add-in Board Slot Covers
E. AC Input Power Connector
F. Serial Port Knockout
G. A Kensington* Cable Lock Mounting Hole
H. Padlock Loop
I. Alternate RMM4 Knockout
J. Front Control Panel
K. 5.25" Peripheral Bays
L. CPU Zone System Fan (Fixed System Fan 2)
M. Fixed Hard Drive Carrier Tray
N. PCI Zone System Fan (Fixed System Fan 1)
O. PCI Card Retainer

Figure 30. Intel® Server Chassis P4000M with Fixed Power Supply, Fixed Hard Drives and Fixed System Fans

A.   750-W Hot Swap Power Supply (Two)
B.   AC Input Power Connector (Two)
C.   I/O Ports
D.   Alternate RMM4 Knockout
E.   PCI Add-in Board Slot Covers
F.   Serial Port Knockout
G.   A Kensington* Cable Lock Mounting Hole
H.   Padlock Loop
I.   Alternate RMM4 Knockout
J.   Hot-swap System Fan 5
K.   Front Control Panel
L.   Hot-swap System Fan 4
M.   5.25'' Peripheral Bays
N.   Hot-swap System Fan 3
O.   Hot-swap System Fan 2
P.   8x3.5'' Hot-swap HDD Cage
Q.   Hot-swap System Fan 1
R.   PCI Card Retainer

Figure 31. Intel® Server Chassis P4000M with Hot-swap Power Supply, Hot-swap Hard Drives and Hot-swap System Fans

The following figure shows the server chassis/system product code naming conventions:
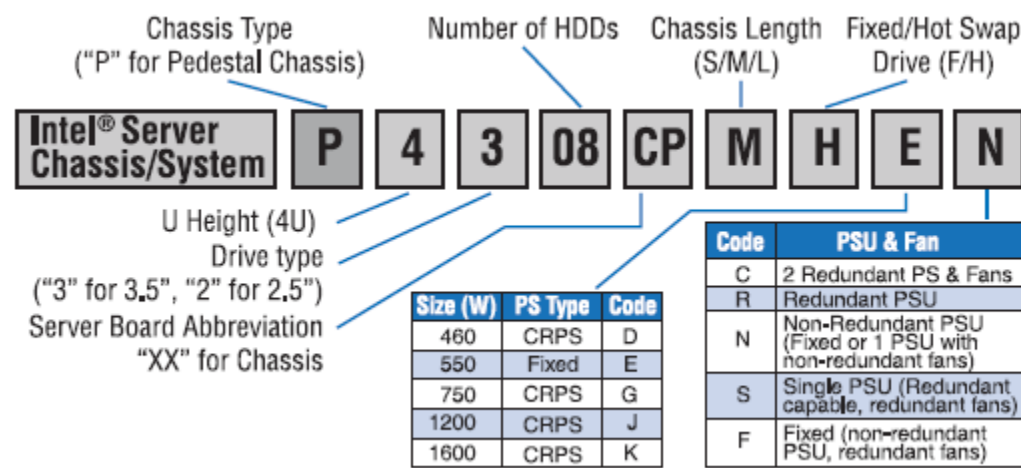
Figure 32. Chassis/System Product Code Naming Conventions

The Intel® Server Chassis P4000M family comes with the following configurations:

1. P4308XXMFEN – one 550-W non-redundant fixed PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" fixed hard drives.

2. P4308XXMHEN – one 550-W non-redundant fixed PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" hot-swap hard drives.

3. P4308XXMFGN – one 750-W hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" fixed hard drives.

4. P4308XXMHGC – two 750-W redundant hot-swap PSU, five redundant hot-swap 80x38mm system fans and up to eight 3.5" hot-swap hard drives.

5. P4308XXMHJC – two 1200-W redundant hot-swap PSU, five redundant hot-swap 80x38mm system fans and up to eight 3.5" hot-swap hard drives.

6. P4208XXMHEN – one 550-W non-redundant fixed PSU, two non-redundant fixed 120x38mm system fans and up to eight 2.5" hot-swap hard drives.

7. P4208XXMHDR – two 460-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 2.5" hot-swap hard drives.

8. P4208XXMHGR – two 750-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 2.5" hot-swap hard drives.

9. P4208XXMHGC – two 750-W redundant hot-swap PSU, five redundant hot-swap 80x38mm system fans and up to eight 2.5" hot-swap hard drives.

10. P4216XXMHJC – two 1200-W redundant hot-swap PSU, five redundant hot-swap 80x38mm system fans and up to sixteen 2.5" hot-swap hard drives.

11. P4216XXMHGC – two 750-W redundant hot-swap PSU, five redundant hot-swap 80x38mm system fans and up to sixteen 2.5" hot-swap hard drives.

12. P4216XXMHGR – two 750-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to sixteen 2.5" hot-swap hard drives.

13. P4216XXMHEN – one 550-W non-redundant fixed PSU, two non-redundant fixed 120x38mm system fans and up to sixteen 2.5" hot-swap hard drives.

14. P4308XXMHGR – two 750-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" hot-swap hard drives.

15. P4308XXMHGN – one 750-W hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" hot-swap hard drives.

16. P4308XXMFGR – two 750-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" fixed hard drives.

17. P4308XXMHDR – two 460-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" hot-swap hard drives.

18. P4308XXMHDN – one 460-W hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" hot-swap hard drives.

19. P4308XXMFDR – two 460-W redundant hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" fixed hard drives.

20. P4308XXMFDN – one 460-W hot-swap PSU, two non-redundant fixed 120x38mm system fans and up to eight 3.5" fixed hard drives.

21. P4304XXMHEN – one 550-W non-redundant fixed PSU, two non-redundant fixed 120x38mm system fans and up to four 3.5" hot-swap hard drives.

The following table summarizes the features for all chassis combinations:

Table 68. Intel® Server Chassis P4000M family Features

Configuration	P4308XXMFEN	P4308XXMHEN	P4308XXMFGN	P4308XXMHGC	P4308XXMHJC
Intel® Server Board Support	▪ Intel® Server Board S2600CP ▪ Intel® Server Board S2600CO ▪ Intel® Server Board S2400SC ▪ Intel® Server Board S2400GP				
Power	550W non-redundant fixed power supply	One 750W hot-swap power supply with high current PDB	Two 750W redundant hot-swap power supply with high current PDB	Two 1200W redundant hot-swap power supply with high current PDB	
System Cooling	Two 120x38mm non-redundant fans			Five 80x38mm redundant hot swap fans	
Peripherals Bays	Three (3) half height 5-1/4" bays for optical devices.				
Drive Bays	Includes one fixed drive bay. Supports up to eight 3.5" fixed hard drives	Includes one 8x3.5" hot-swap hard drive cage. Supports up to eight 3.5" hot-swap hard drives.	Includes one fixed drive bay. Supports up to eight 3.5" fixed hard	Includes one 8x3.5" hot-swap hard drive cage. Supports up to eight 3.5" hot-swap hard drives.	
Expansion Slots	Support up to six (6) full height, full length PCI form factor cards mechanically.				
Front Panel	Power Button with LED, Reset Button, NMI Button, ID Button with LED, Four NIC LEDs, Hard drive				

Configuration	P4308XXMFEN	P4308XXMHEN	P4308XXMFGN	P4308XXMHGC	P4308XXMHJC
	activity LED, System status LED, two USB ports, Optional front serial port/VGA port				
Appearance	Color: Cosmetic black (GE 701 or equivalent), service Intel® blue, hot swap Intel® green. Support for Intel® standard front panel or LCD				
Dimensions Pedestal	17.24 in (438 mm) x 6.81 in (173mm) x 25 in (612 mm) (Height X Width X Depth)				
Optional Accessory Kits	Zephyr flash storage, RMM4-lite modules, TPM module, dedicated NIC module, Expander Card module,				
Configuration	P4208XXMHEN	P4208XXMHDR	P4208XXMHGR	P4208XXMHGC	P4216XXMHJC
Intel® Server Board Support	▪ Intel® Server Board S2600CP ▪ Intel® Server Board S2600CO ▪ Intel® Server Board S2400SC ▪ Intel® Server Board S2400GP				
Power	550W non-redundant fixed power supply	Two 460W redundant hot-swap power supply with low current PDB	Two 750W redundant hot-swap power supply with high current PDB		Two 1200W redundant hot-swap power supply with high current PDB
System Cooling	Two 120x38mm non-redundant fans			Five 80x38mm redundant hot swap fans	
Peripherals Bays	Three (3) half height 5-1/4" bays for optical devices.				
Drive Bays	Include one 8x2.5" hot-swap hard drive cage. Supports up to eight 2.5" hot-swap hard drives.				Include two 8x2.5" hot-swap hard drive cage. Supports up to sixteen 2.5" hot-swap hard drives.
Expansion Slots	Support up to six (6) full height, full length PCI form factor cards mechanically				
Front Panel	Power Button with LED, Reset Button, NMI Button, ID Button with LED, Four NIC LEDs, Hard drive activity LED, System status LED, two USB ports, Optional front serial port/VGA port				
Appearance	Color: Cosmetic black (GE 701 or equivalent), service Intel® blue, hot swap Intel® green. Support for Intel® standard front panel or LCD				
Dimensions Pedestal	17.24 in (438 mm) x 6.81 in (173mm) x 25 in (612 mm) (Height X Width X Depth)				
Optional Accessory Kits	Zephyr flash storage, RMM4-lite modules, TPM module, dedicated NIC module, Expander Card module				
Configuration	P4216XXMHGC	P4216XXMHGR	P4308XXMHGR	P4308XXMFGR	P4308XXMHGN
Intel® Server Board Support	▪ Intel® Server Board S2600CP ▪ Intel® Server Board S2600CO ▪ Intel® Server Board S2400SC ▪ Intel® Server Board S2400GP				
Power	Two 750W redundant hot-swap power supply with high current PDB				One 750W hot-swap power supply with high current PDB
System Cooling	Five 80x38mm redundant hot swap fans	Two 120x38mm non-redundant fans			
Peripherals Bays	Three (3) half height 5-1/4" bays for optical devices.				
Drive Bays	Include two 8x2.5" hot-swap hard drive cage. Supports up to	Include one 8x3.5" hot-	Includes one fixed drive		Include one 8x3.5" hot-swap hard drive cage. Supports up to

Configuration	P4308XXMFEN	P4308XXMHEN	P4308XXMFGN	P4308XXMHGC	P4308XXMHJC
	sixteen 2.5" hot-swap hard drives.		swap hard drive cage. Supports up to eight 3.5" hot-swap hard drives.	bay. Supports up to eight 3.5" fixed hard	eight 3.5" hot-swap hard drives.
Expansion Slots	Support up to six (6) full height, full length PCI form factor cards mechanically.				
Front Panel	Power Button with LED, Reset Button, NMI Button, ID Button with LED, Four NIC LEDs, Hard drive activity LED, System status LED, two USB ports, Optional front serial port/VGA port				
Appearance	Color: Cosmetic black (GE 701 or equivalent), service Intel® blue, hot swap Intel® green. Support for Intel® standard front panel or LCD				
Dimensions Pedestal	17.24 in (438 mm) x 6.81 in (173mm) x 25 in (612 mm) (Height X Width X Depth)				
Optional Accessory Kits	Zephyr flash storage, RMM4-lite modules, TPM module, dedicated NIC module, Expander Card module,				

Configuration	P4216XXMHEN	P4308XXMHDR	P4308XXMHDN	P4308XXMFDR	P4308XXMFDN	P4304XXMHEN
Intel® Server Board Support	- Intel® Server Board S2600CP - Intel® Server Board S2600CO - Intel® Server Board S2400SC - Intel® Server Board S2400GP					
Power	550W non-redundant fixed power supply	Two 460W redundant hot-swap power supply with low current PDB	One 460W hot-swap power supply with low current PDB	Two 460W redundant hot-swap power supply with low current PDB	One 460W hot-swap power supply with low current PDB	550W non-redundant fixed power supply
System Cooling	Two 120x38mm non-redundant fans					
Peripherals Bays	Three (3) half height 5-1/4" bays for optical devices.					
Drive Bays	Include two 8x2.5" hot-swap hard drive cage. Supports up to sixteen 2.5" hot-swap hard drives.	Include one 8x3.5" hot-swap hard drive cage. Supports up to eight 3.5" hot-swap hard drives.		Includes one fixed drive bay. Supports up to eight 3.5" fixed hard		Include one 4x3.5" hot-swap hard drive cage. Supports up to four 3.5" hot-swap hard drives.
Expansion Slots	Support up to six (6) full height, full length PCI form factor cards mechanically.					
Front Panel	Power Button with LED, Reset Button, NMI Button, ID Button with LED, Four NIC LEDs, Hard drive activity LED, System status LED, two USB ports, Optional front serial port/VGA port					
Appearance	Color: Cosmetic black (GE 701 or equivalent), service Intel® blue, hot swap Intel® green. Support for Intel® standard front panel or LCD					
Dimensions Pedestal	17.24 in (438 mm) x 6.81 in (173mm) x 25 in (612 mm) (Height X Width X Depth)					
Optional Accessory Kits	Zephyr flash storage, RMM4-lite modules, TPM module, dedicated NIC module, Expander Card module,					

For additional support on Intel® Server Chassis, please refer *Intel® Server Chassis P4000M Family Technical Product Specification.*

Appendix C: Integrated BMC Sensor Tables

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

- **Sensor Type**

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

- **Event/Reading Type**

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

- **Event Offset/Triggers**

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

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

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

- **Assertion/De-assertion Enables**

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

  - As: Assertions
  - De: De-assertion

- **Readable Value/Offsets**

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

- **Event Data**

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

  - R: Reading value
  - T: Threshold value

▪ **Rearm Sensors**

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

- o A: Auto-rearm
- o M: Manual rearm

▪ **Default Hysteresis**

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

▪ **Criticality**

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

▪ **Standby**

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

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

Table 69. BMC Core Sensors

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers		Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Power Unit Status (Pwr Unit Status)	01h	All	Power Unit 09h	Sensor Specific 6Fh	00 - Power down		OK	As and De	–	Trig Offset	A	X
					02 - 240 VA power down		Fatal					
					04 - A/C lost		OK					
					05 - Soft power control failure		Fatal					
					06 - Power unit failure							
Power Unit Redundancy[Note1] (Pwr Unit Redund)	02h	Chassis-specific	Power Unit 09h	Generic 0Bh	00 - Fully Redundant		OK	As and De	–	Trig Offset	A	X
					01 - Redundancy lost		Degraded					
					02 - Redundancy degraded		Degraded					
					03 - Non-redundant: sufficient resources. Transition from full redundant state.		Degraded					
					04 – Non-redundant: sufficient resources. Transition from insufficient state.		Degraded					
					05 - Non-redundant: insufficient resources		Fatal					
					06 – Redundant: degraded from fully redundant state.		Degraded					
					07 – Redundant: Transition from non-redundant state.		Degraded					

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
IPMI Watchdog (IPMI Watchdog)	03h	All	Watchdog 2 23h	Sensor Specific 6Fh	00 - Timer expired, status only 01 - Hard reset 02 - Power down 03 - Power cycle 08 - Timer interrupt	OK	As	–	Trig Offset	A	X
Physical Security (Physical Scrty)	04h	Chassis Intrusion is chassis-specific	Physical Security 05h	Sensor Specific 6Fh	00 - Chassis intrusion 04 - LAN leash lost	OK	As and De	–	Trig Offset	A	X
FP Interrupt (FP NMI Diag Int)	05h	Chassis - specific	Critical Interrupt 13h	Sensor Specific 6Fh	00 - Front panel NMI/diagnostic interrupt	OK	As	–	Trig Offset	A	–
SMI Timeout (SMI Timeout)	06h	All	SMI Timeout F3h	Digital Discrete 03h	01 – State asserted	Fatal	As and De	–	Trig Offset	A	–
System Event Log (System Event Log)	07h	All	Event Logging Disabled 10h	Sensor Specific 6Fh	02 - Log area reset/cleared	OK	As	–	Trig Offset	A	X
System Event (System Event)	08h	All	System Event 12h	Sensor Specific 6Fh	02 - Undetermined system H/W failure 04 – PEF action	Fatal OK	As and De As	-	Trig Offset	A	X
Button Sensor (Button)	09h	All	Button/Switch 14h	Sensor Specific 6Fh	00 – Power Button 02 – Reset Button	OK	AS	_	Trig Offset	A	X
BMC Watchdog	0Ah	All	Mgmt System Health 28h	Digital Discrete 03h	01 – State Asserted	Degraded	As	–	Trig Offset	A	-
Voltage Regulator Watchdog (VR Watchdog)	0Bh	All	Voltage 02h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Fan Redundancy[Note1]	0Ch	Chassis-	Fan	Generic	00 - Fully redundant	OK	As	–	Trig Offset	A	–

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

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Baseboard Temperature 1 (Platform Specific)	20h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Front Panel Temperature (Front Panel Temp)	21h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SSB Temperature (SSB Temp)	22h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 2 (Platform Specific)	23h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 3 (Platform Specific)	24h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard Temperature 4 (Platform Specific)	25h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Module Temperature (I/O Mod Temp)	26h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 1 Temperature (PCI Riser 1 Temp)	27h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Riser Temperature (IO Riser Temp)	28h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 1 Temperature (HSBP 1 Temp)	29h	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 2 Temperature (HSBP 2 Temp)	2Ah	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers		Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Hot-swap Backplane 3 Temperature (HSBP 3 Temp)	2Bh	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 2 Temperature (PCI Riser 2 Temp)	2Ch	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SAS Module Temperature (SAS Mod Temp)	2Dh	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Exit Air Temperature (Exit Air Temp)	2Eh	Chassis and Platform Specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Network Interface Controller Temperature (LAN NIC Temp)	2Fh	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Fan Tachometer Sensors (Chassis specific sensor names)	30h–3Fh	Chassis and Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]		nc = Degraded c = Non-fatal[Note2]	As and De	Analog	R, T	M	-
Fan Present Sensors (Fan x Present)	40h–4Fh	Chassis and Platform Specific	Fan 04h	Generic 08h	01 - Device inserted		OK	As and De	-	Triggered Offset	Auto	-
Power Supply 1 Status (PS1 Status)	50h	Chassis-specific	Power Supply 08h	Sensor Specific 6Fh	00 - Presence	OK		As and De	–	Trig Offset	A	X
					01 - Failure	Degraded						
					02 – Predictive Failure	Degraded						
					03 - A/C lost	Degraded						
					06 – Configuration error	OK						
Power Supply 2 Status (PS2 Status)	51h	Chassis-specific	Power Supply 08h	Sensor Specific 6Fh	00 - Presence	OK		As and De	–	Trig Offset	A	X
					01 - Failure	Degraded						
					02 – Predictive Failure	Degraded						
					03 - A/C lost	Degraded						

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers		Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
					06 – Configuration error		OK					
Power Supply 1 AC Power Input (PS1 Power In)	54h	Chassis-specific	Other Units 0Bh	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 AC Power Input (PS2 Power In)	55h	Chassis-specific	Other Units 0Bh	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 1 +12V % of Maximum Current Output (PS1 Curr Out %)	58h	Chassis-specific	Current 03h	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 +12V % of Maximum Current Output (PS2 Curr Out %)	59h	Chassis-specific	Current 03h	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 1 Temperature (PS1 Temperature)	5Ch	Chassis-specific	Temperature 01h	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 Temperature (PS2 Temperature)	5Dh	Chassis-specific	Temperature	Threshold 01h	[u] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hard Disk Drive 16 - 24 Status (HDD 16 - 24 Status)	60h – 68h	Chassis-specific	Drive Slot 0Dh	Sensor Specific 6Fh	00 - Drive Presence	OK		As and De	–	Trig Offset	A	X
					01- Drive Fault	Degraded						
					07 - Rebuild/Remap in progress	Degraded						
Processor 1 Status (P1 Status)	70h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal		As and De	–	Trig Offset	M	X
					07 - Presence	OK						
Processor 2 Status (P2 Status)	71h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal		As and De	–	Trig Offset	M	X
					07 - Presence	OK						

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De- assert	Readable Value/ Offsets	Event Data	Rearm	Stand- by
Processor 1 Thermal Margin (*P1 Therm Margin*)	74h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 2 Thermal Margin (*P2 Therm Margin*)	75h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Processor 1 Thermal Control % (*P1 Therm Ctrl %*)	78h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor 2 Thermal Control % (*P2 Therm Ctrl %*)	79h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	–
Processor 1 ERR2 Timeout (*P1 ERR2*)	7Ch	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	A	–
Processor 2 ERR2 Timeout (*P2 ERR2*)	7Dh	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	A	–
Catastrophic Error (*CATERR*)	80h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor0 MSID Mismatch (*P0 MSID Mismatch*)	81h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor Population Fault (*CPU Missing*)	82h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	–
Processor1 MSID Mismatch (*P1 MSID Mismatch*)	87h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor 1 VRD Temperature (*P1 VRD Hot*)	90h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	M	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Processor 2 VRD Temperature (P2 VRD Hot)	91h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	M	–
Processor 1 Memory VRD Hot 0-1 (P1 Mem01 VRD Hot)	94h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Processor 1 Memory VRD Hot 2-3 (P1 Mem23 VRD Hot)	95h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Processor 2 Memory VRD Hot 0-1 (P2 Mem01 VRD Hot)	96h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Processor 2 Memory VRD Hot 2-3 (P2 Mem23 VRD Hot)	97h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Power Supply 1 Fan Tachometer 1 (PS1 Fan Tach 1)	A0h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 1 Fan Tachometer 2 (PS1 Fan Tach 2)	A1h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 1 (PS2 Fan Tach 1)	A4h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 2 (PS2 Fan Tach 2)	A5h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Processor 1 DIMM Aggregate Thermal Margin (P1 DIMM Thrm Mrgn)	B0h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Processor 2 DIMM Aggregate Thermal Margin (P2 DIMM Thrm Mrgn)	B1h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Processor 1 DIMM Thermal Trip (P1 Mem Thrm Trip)	C0h	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Processor 2 DIMM Thermal Trip (P2 Mem Thrm Trip)	C1h	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Global Aggregate Temperature Margin 1 (Agg Therm Mrgn 1)	C8h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 2 (Agg Therm Mrgn 2)	C9h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 3 (Agg Therm Mrgn 3)	CAh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 4 (Agg Therm Mrgn 4)	CBh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 5 (Agg Therm Mrgn 5)	CCh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 6 (Agg Therm Mrgn 6)	CDh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 7 (Agg Therm Mrgn 7)	CEh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Global Aggregate Temperature Margin 8 (Agg Therm Mrgn 8)	CFh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	–
Baseboard +12V (BB +12.0V)	D0h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +5V (BB +5.0V)	D1h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Baseboard +3.3V (BB +3.3V)	D2h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +5V Stand-by (BB +5.0V STBY)	D3h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard +3.3V Auxiliary (BB +3.3V AUX)	D4h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Baseboard +1.05V Processor 1 Vccp (BB +1.05Vccp P1)	D6h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.05V Processor 1 Vccp (BB +1.05Vccp P2)	D7h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.5V P1 Memory AB VDDQ (BB +1.5 P1MEM AB)	D8h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.5V P1 Memory CD VDDQ (BB +1.5 P1MEM CD)	D9h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.5V P2 Memory AB VDDQ (BB +1.5 P2MEM AB)	DAh	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.5V P2 Memory CD VDDQ (BB +1.5 P2MEM CD)	DBh	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.8V Aux (BB +1.8V AUX)	DCh	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.1V Stand-by (BB +1.1V STBY)	DDh	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/ Reading Type	Event Offset Triggers		Contrib. To System Status	Assert/ De-assert	Readable Value/ Offsets	Event Data	Rearm	Stand-by
Baseboard CMOS Battery (BB +3.3V Vbat)	DEh	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.35V P1 Low Voltage Memory AB VDDQ (BB +1.35 P1LV AB)	E4h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.35V P1 Low Voltage Memory CD VDDQ (BB +1.35 P1LV CD)	E5h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.35V P2 Low Voltage Memory AB VDDQ (BB +1.35 P2LV AB)	E6h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +1.35V P2 Low Voltage Memory CD VDDQ (BB +1.35 P2LV CD)	E7h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +3.3V Riser 1 Power Good (BB +3.3 RSR1 PGD)	EAh	Platform Specific	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Baseboard +3.3V Riser 2 Power Good (BB +3.3 RSR2 PGD)	EBh	Platform Specific	Voltage 02h	Threshold 01h	[u,l] [c,nc]		nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	–
Hard Disk Drive 1 -15 Status (HDD 1 - 15 Status)	F0h - FEh	Chassis-specific	Drive Slot 0Dh	Sensor Specific 6Fh	00 - Drive Presence	OK		As and De	–	Trig Offset	A	X
					01- Drive Fault	Degraded						
					07 - Rebuild/Remap in progress	Degraded						

Notes:

1. Redundancy sensors will be only present on systems with appropriate hardware to support redundancy (for instance, fan or power supply).

2. This is only applicable when the system does not support redundant fans. When fan redundancy is supported, then the contribution to system state is driven by the fan redundancy sensor.

Appendix D: Intel® Server Board S2600CO Family Specific Sensors

Product ID

Bytes 11:12 (product ID) of Get Device ID command response: 5Dh 00h

ACPI S3 Sleep State Support

Not supported.

Processor Support for Intel® Server Board S2600CO

- Intel® Xeon® processor E5-2600 product family up to 150 W
- Intel® Xeon® processor E5-2600 v2 product family up to 150 W

Supported Chassis

- Intel® Server Chassis P4208 (UPM with fixed fan 4U/general SKU)
- Intel® Server Chassis P4216 (UPM with fixed fan 4U/general SKU)
- Intel® Server Chassis P4304 (UPM with fixed fan 4U/general SKU)
- Intel® Server Chassis P4308 (UPM with fixed fan 4U/general SKU)
- Intel® Server Chassis P4208 (UPM with fixed fan 4U/150W SKU)
- Intel® Server Chassis P4216 (UPM with fixed fan 4U/150W SKU)
- Intel® Server Chassis P4308 (UPM with fixed fan 4U/150W SKU)
- Intel® Server Chassis P4208 (UPM with redundant fan 4U)
- Intel® Server Chassis P4216 (UPM with redundant fan 4U)
- Intel® Server Chassis P4308 (UPM with redundant fan 4U)

Table 70. Chassis-specific Sensors

Intel® Server Chassis	Fan Tachometer sensors	Fan Presence sensors	Physical security (Chassis intrusion) Sensor	FP interrupt (FP NMI Diag Int)
P4208/P4216/P4304/P4308 (UPM with fixed fan 4U/general SKU)	System Fan 1 (30h)	NA	Physical Scrty (04h)	FP NMI Diag Int (05h)
	System Fan 2 (31h)	NA	NA	NA
P4208/P4216/P4308 (UPM with fixed fan 4U/150W SKU)	System Fan 1 (30h)	NA	Physical Scrty (04h)	FP NMI Diag Int (05h)
	System Fan 2 (31h)	NA	NA	NA
	Processor 1 Fan (36h)	NA	NA	NA
	Processor 2 Fan (37h)detected	NA	NA	NA
P4208/P4216/P4308 (UPM with redundant fan 4U)	System Fan 1 (30h)	Fan 1 Present (40h)	Physical Scrty (04h)	FP NMI Diag Int (05h)
	System Fan 2 (31h)	Fan 2 Present (41h)	NA	NA
	System Fan 3 (32h)	Fan 3 Present (42h)	NA	NA
	System Fan 4 (33h)	Fan 4 Present (43h)	NA	NA
	System Fan 5 (34h)	Fan 5 Present (44h)	NA	NA

Hot-plug fan support

Supported on:

- Intel® Server Chassis P4208 (UPM with redundant fan 4U)
- Intel® Server Chassis P4216 (UPM with redundant fan 4U)
- Intel® Server Chassis P4308 (UPM with redundant fan 4U)

Fan redundancy support

Supported on:

- Intel® Server Chassis P4208 (UPM with redundant fan 4U)
- Intel® Server Chassis P4216 (UPM with redundant fan 4U)
- Intel® Server Chassis P4308 (UPM with redundant fan 4U)

Table 71. Fan Domain Definition

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
P4208/P4216/P4304/P4308 (UPM with fixed fan 4U/general SKU)	0	▪ SSB Temp (22h) ▪ BB BMC Temp (23h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ LAN NIC Temp (2Fh) ▪ DIMM Thrm Mrgn 2 (B1h) ▪ DIMM Thrm Mrgn 3 (B2h)	System Fan 1 (30h)
	1	▪ BB P2 VR Temp (24h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ P1 Therm Margin (74h) ▪ P2 Therm Margin (75h) ▪ DIMM Thrm Mrgn 1 (B0h) ▪ DIMM Thrm Mrgn 2 (B1h) ▪ DIMM Thrm Mrgn 3 (B2h) ▪ DIMM Thrm Mrgn 4 (B3h)	System Fan 2 (31h)
P4208/P4216/P4308 (UPM with fixed fan 4U/150W SKU)	0	▪ SSB Temp (22h) ▪ BB BMC Temp (23h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ LAN NIC Temp (2Fh) ▪ DIMM Thrm Mrgn 2 (B1h) ▪ DIMM Thrm Mrgn 3 (B2h)	System Fan 1 (30h)
	1	▪ BB P2 VR Temp (24h) ▪ BB MEM VR Temp (25h)	System Fan 2 (31h)

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
		▪ HSBP 1 Temp (29h)	
		▪ HSBP 2 Temp (2Ah)	
		▪ Exit Air Temp (2Eh)	
		▪ P1 Therm Margin (74h)	
		▪ P2 Therm Margin (75h)	
		▪ DIMM Thrm Mrgn 1 (B0h)	
		▪ DIMM Thrm Mrgn 2 (B1h)	
		▪ DIMM Thrm Mrgn 3(B2h)	
		▪ DIMM Thrm Mrgn 4 (B3h)	
	6	▪ P1 MEM01 VRD Hot (94h)	Processor 1 Fan (36h)
	7	▪ P1 MEM23 VRD Hot (95h) ▪ BB P2 VR Temp (24h)	Processor 2 Fan (37h)
P4208/ P4308/ P4216 (UPM with redundant fan 4U)	0	▪ SSB Temp (22h) ▪ BB BMC Temp (23h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ LAN NIC Temp (2Fh)	System Fan 1 (30h)
	1	▪ SSB Temp (22h) ▪ BB BMC Temp (23h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ LAN NIC Temp (2Fh) ▪ P1 Therm Margin (74h) ▪ P2 Therm Margin (75h) ▪ DIMM Thrm Mrgn 2 (B1h) ▪ DIMM Thrm Mrgn 3 (B2h)	System Fan 2 (31h)
	2	▪ BB P2 VR Temp (24h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ P1 Therm Margin (74h) ▪ P2 Therm Margin (75h) ▪ DIMM Thrm Mrgn 1 (B0h) ▪ DIMM Thrm Mrgn 2 (B1h) ▪ DIMM Thrm Mrgn 3 (B2h) ▪ DIMM Thrm Mrgn 4 (B3h)	System Fan 3 (32h)
	3	▪ BB P2 VR Temp (24h) ▪ BB MEM VR Temp (25h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh)	System Fan 4 (33h)

Chassis	Fan Domain	Major Components Cooled (Temperature sensor number)	Fans (Sensor number)
		▪ P1 Therm Margin (74h) ▪ P2 Therm Margin (75h) ▪ DIMM Thrm Mrgn 1 (B0h) ▪ DIMM Thrm Mrgn 2 B1h) ▪ DIMM Thrm Mrgn 3 (B2h) ▪ DIMM Thrm Mrgn 4 (B3h)	
	4	▪ BB P2 VR Temp (24h) ▪ HSBP 1 Temp (29h) ▪ HSBP 2 Temp (2Ah) ▪ Exit Air Temp (2Eh) ▪ PS1 Temperature (5Ch) ▪ PS2 Temperature (5Dh) ▪ P2 Therm Margin (75h) ▪ DIMM Thrm Mrgn 1 (B0h) ▪ DIMM Thrm Mrgn 4 (B3h)	System Fan 5 (34h)

HSC Availability

- Intel® Server Chassis P4208(UPM)
  - 8-bay 2.5'' HDD – FXX8X25HSBP
- Intel® Server Chassis P4216 (UPM)
  - 8-bay 2.5'' HDD – FXX8X25HSBP
- Intel® Server Chassis P4304(UPM)
  - 4-bay 3.5'' HDD – FUP4X35HSBP
- Intel® Server Chassis P4308 (UPM)
  - 8-bay 3.5'' HDD – FUP8X35HSBP

Power unit support

- Intel® Server Chassis P4208/P4216/P4304/P4308 (Intel® Server Chassis P4000M Family)

Table 72. Power Supply Support

PS Module Number	PMBus*	Product Name (in product area of the FRU)	PSU Redundant	Cold Redundant	Fans in each PS
460W HS Power Supply	Supported	DPS-460KB A	Not supported (One PSU configuration)	Not supported	1 PS fan
460W HS Power Supply	Supported	DPS-460KB A	Supported	Supported	1 PS fan
550W Fixed Power Supply	Not supported	No FRU	Not supported	Not supported	NA
750W HS Power Supply	Supported	DPS-750XB A	Not supported (One PSU configuration)	Not supported	1 PS fan
750W HS Power Supply	Supported	DPS-750XB A	Supported	Supported	1 PS fan

PS Module Number	PMBus*	Product Name (in product area of the FRU)	PSU Redundant	Cold Redundant	Fans in each PS
1200W HS Power Supply	Supported	DPS-1200TB A	Supported	Supported	2 PS fan

Redundant Fans only for Intel® Server Chassis

- Intel® Server Chassis P4208 (UPM with redundant fan 4U)
- Intel® Server Chassis P4216 (UPM with redundant fan 4U)
- Intel® Server Chassis P4308 (UPM with redundant fan 4U)

Fan Fault LED support

Fan fault LEDs are available on the hot-swap redundant fans available on the on below chassis:

- Intel® Server Chassis P4208 (UPM with redundant fan 4U)
- Intel® Server Chassis P4216 (UPM with redundant fan 4U)
- Intel® Server Chassis P4308 (UPM with redundant fan 4U)

Memory Throttling support

- Baseboard supports this feature.

Appendix E: Management Engine Generated SEL Event Messages

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

Table 73. Server Platform Services Firmware Health Event

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

Table 74. Node Manager Health Event

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

Appendix F: POST Code Diagnostic LED Decoder

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

During the system boot process, Memory Reference Code (MRC) and System BIOS execute a number of memory initialization and platform configuration processes, each of which is assigned a specific hex POST code number. As each 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, and #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, then the corresponding LED is lit. If the bit is clear, then the corresponding LED is off.

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

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

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

Table 75. 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 76. Diagnostic LED POST Code Decoder

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	
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 77									
PEI Phase continued…									
31h	0	0	1	1	0	0	0	1	Memory Installed
32h	0	0	1	1	0	0	1	0	CPU PEIM (Cpu Init)
33h	0	0	1	1	0	0	1	1	CPU PEIM (Cache Init)
34h	0	0	1	1	0	1	0	0	CPU PEIM (BSP Select)
35h	0	0	1	1	0	1	0	1	CPU PEIM (AP Init)
36h	0	0	1	1	0	1	1	0	CPU PEIM (CPU SMM Init)
4Fh	0	1	0	0	1	1	1	1	Dxe IPL started
DXE Phase									
60h	0	1	1	0	0	0	0	0	DXE Core started
61h	0	1	1	0	0	0	0	1	DXE NVRAM Init
62h	0	1	1	0	0	0	1	0	SB RUN Init
63h	0	1	1	0	0	0	1	1	Dxe CPU Init
68h	0	1	1	0	1	0	0	0	DXE PCI Host Bridge Init
69h	0	1	1	0	1	0	0	1	DXE NB Init
6Ah	0	1	1	0	1	0	1	0	DXE NB SMM Init
70h	0	1	1	1	0	0	0	0	DXE SB Init
71h	0	1	1	1	0	0	0	1	DXE SB SMM Init
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

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	
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									
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 77. MRC Progress Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off								Description
	Upper Nibble				Lower Nibble				
	MSB 8h	4h	2h	1h	8h	4h	2h	LSB 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 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 do NOT get logged as SEL events. The following table lists all MRC fatal errors that are displayed to the Diagnostic LEDs.

Table 78. MRC Fatal Error Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off								Description (with MRC Internal Minor Code)
	Upper Nibble				Lower Nibble				
	MSB 8h	4h	2h	1h	8h	4h	2h	LSB 1h	
LED	#7	#6	#5	#4	#3	#2	#1	#0	
MRC Fatal Error Codes									
E8h	1	1	1	0	1	0	0	0	No usable memory error 01h = No memory was detected from the SPD read, or invalid config that causes no operable memory. 02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error.

Checkpoint	Diagnostic LED Decoder								Description (with MRC Internal Minor Code)
	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	
									03h = No memory installed. All channels are disabled.
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel® Trusted Execution Technology and is inaccessible
EAh	1	1	1	0	1	0	1	0	DDR3 channel training error 01h = Error on read DQ/DQS (Data/Data Strobe) init 02h = Error on Receive Enable 03h = Error on Write Leveling 04h = Error on write DQ/DQS (Data/Data Strobe
EBh	1	1	1	0	1	0	1	1	Memory test failure 01h = Software memtest failure. 02h = Hardware memtest failed. 03h = Hardware memtest failure in Lockstep Channel mode requiring a channel to be disabled. *This is a fatal error which requires a reset and calling MRC with a different RAS mode to retry.*
EDh	1	1	1	0	1	1	0	1	DIMM configuration population error 01h = Different DIMM types (UDIMM, RDIMM, LRDIMM) are detected installed in the system. 02h = Violation of DIMM population rules. 03h = The third DIMM slot cannot be populated when QR DIMMs are installed. 04h = UDIMMs are not supported in the third DIMM slot. 05h = Unsupported DIMM Voltage.
EFh	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error

Appendix G: POST Code Errors

Most error conditions encountered during POST are reported using **POST Error Codes**. These codes represent specific failures, warnings, or informational messages. 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 condition 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 79. 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

Error Code	Error Message	Response
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 Selftest	Major
83A1	Management Engine (ME) Failed to respond.	Major
84F2	Baseboard management controller failed to respond	Major
84F3	Baseboard management controller in update mode	Major
84F4	Sensor data record empty	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8501	DIMM Population Error	Major
8520	DIMM_A1 failed test/initialization	Major
8521	DIMM_A2 failed test/initialization	Major
8522	DIMM_A3 failed test/initialization	Major
8523	DIMM_B1 failed test/initialization	Major
8524	DIMM_B2 failed test/initialization	Major
8525	DIMM_B3 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8527	DIMM_C2 failed test/initialization	Major
8528	DIMM_C3 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major

Error Code	Error Message	Response
852A	DIMM_D2 failed test/initialization	Major
852B	DIMM_D3 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852D	DIMM_E2 failed test/initialization	Major
852E	DIMM_E3 failed test/initialization	Major
852F	DIMM_F1 failed test/initialization	Major
8530	DIMM_F2 failed test/initialization	Major
8531	DIMM_F3 failed test/initialization	Major
8532	DIMM_G1 failed test/initialization	Major
8533	DIMM_G2 failed test/initialization	Major
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

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

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

Error Code	Error Message	Response
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 Legacy Option ROM	Minor

POST Error Beep Codes

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

Table 80. 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	Multiple	System halted because a fatal error related to the memory was detected.
The following Beep Codes are sounded during BIOS Recovery.			
2	BIOS Recovery started	NA	Recovery boot has been initiated
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 81. Integrated BMC Beep Codes

Code	Reason for Beep	Associated Sensors
1-5-2-1	CPU socket population error	CPU1 socket is empty, or sockets are populated incorrectly CPU1 must be populated before CPU2.
1-5-2-4	MSID Mismatch	MSID mismatch occurs if a processor is installed into a system board that has incompatible power capabilities.
1-5-4-2	Power fault	DC power unexpectedly lost (power good dropout) – Power unit sensors report power unit failure offset
1-5-4-4	Power control fault	Power good assertion timeout – Power unit sensors report soft power control failure offset
1-5-1-2	VR Watchdog Timer	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.

Glossary

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

Term	Definition
ACPI	Advanced Configuration and Power Interface
AP	Application Processor
APIC	Advanced Programmable Interrupt Control
ASIC	Application Specific Integrated Circuit
BIOS	Basic Input/Output System
BIST	Built-In Self Test
BMC	Baseboard Management Controller
Bridge	Circuitry connecting one computer bus to another, allowing an agent on one to access the other
BSP	Bootstrap Processor
byte	8-bit quantity.
CBC	Chassis Bridge Controller (A microcontroller connected to one or more other CBCs, together they bridge the IPMB buses of multiple chassis.
CEK	Common Enabling Kit
CHAP	Challenge Handshake Authentication Protocol
CMOS	In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory, which normally resides on the server board.
DPC	Direct Platform Control
EEPROM	Electrically Erasable Programmable Read-Only Memory
EHCI	Enhanced Host Controller Interface
EMP	Emergency Management Port
EPS	External Product Specification
FMB	Flexible Mother Board
FMC	Flex Management Connector
FMM	Flex Management Module
FRB	Fault Resilient Booting
FRU	Field Replaceable Unit
FSB	Front Side Bus
GB	1024MB
GPIO	General Purpose I/O
GTL	Gunning Transceiver Logic
HSC	Hot-Swap Controller
Hz	Hertz (1 cycle/second)
I^{2}C	Inter-Integrated Circuit Bus
IA	Intel® Architecture
IBF	Input Buffer
ICH	I/O Controller Hub
ICMB	Intelligent Chassis Management Bus
IERR	Internal Error
IFB	I/O and Firmware Bridge

Term	Definition
INTR	Interrupt
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IR	Infrared
ITP	In-Target Probe
KB	1024 bytes
KCS	Keyboard Controller Style
LAN	Local Area Network
LCD	Liquid Crystal Display
LED	Light Emitting Diode
LPC	Low Pin Count
LUN	Logical Unit Number
MAC	Media Access Control
MB	1024KB
mBMC	National Semiconductor* PC87431x mini BMC
MCH	Memory Controller Hub
MD2	Message Digest 2 – Hashing Algorithm
MD5	Message Digest 5 – Hashing Algorithm – Higher Security
ms	milliseconds
MTTR	Memory Tpe Range Register
Mux	Multiplexor
NIC	Network Interface Controller
NMI	Nonmaskable Interrupt
OBF	Output Buffer
OEM	Original Equipment Manufacturer
Ohm	Unit of electrical resistance
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
RAM	Random Access Memory
RASUM	Reliability, Availability, Serviceability, Usability, and Manageability
RISC	Reduced Instruction Set Computing
ROM	Read Only Memory
RTC	Real-Time Clock (Component of ICH peripheral chip on the server board)
SDR	Sensor Data Record
SECC	Single Edge Connector Cartridge
SEEPROM	Serial Electrically Erasable Programmable Read-Only Memory
SEL	System Event Log
SIO	Server Input/Output

Term	Definition
SMI	Server Management Interrupt (SMI is the highest priority nonmaskable interrupt)
SMM	Server Management Mode
SMS	Server Management Software
SNMP	Simple Network Management Protocol
TBD	To Be Determined
TIM	Thermal Interface Material
UART	Universal Asynchronous Receiver/Transmitter
UDP	User Datagram Protocol
UHCI	Universal Host Controller Interface
UTC	Universal time coordinare
VID	Voltage Identification
VRD	Voltage Regulator Down
Word	16-bit quantity
ZIF	Zero Insertion Force

Reference Documents

1. *Advanced Configuration and Power Interface Specification*, Revision 3.0, http://www.acpi.info/

2. *Intelligent Platform Management Bus Communications Protocol Specification, Version 1.0*. 1998. Intel Corporation, Hewlett-Packard* Company, NEC* Corporation, Dell* Computer Corporation

3. *Intelligent Platform Management Interface Specification, Version 2.0*. 2004. Intel Corporation, Hewlett-Packard* Company, NEC* Corporation, Dell* Computer Corporation

4. *Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification, Version 1.1*, 02/01/02, Intel Corporation

5. *Intel® Remote Management Module User's Guide*, Intel Corporation

6. *Alert Standard Format (ASF) Specification, Version 2.0, 23 April 20*03, © 2000-2003, Distributed Management Task Force, Inc., http://www.dmtf.org

7. *BIOS for EPSD Platforms Based on Intel® Xeon Processor E5-4600/2600/2400/1600 Product Families External Product Specification*

8. *EPSD Platforms Based On Intel Xeon® Processor E5 4600/2600/2400/1600 Product Families BMC Core Firmware External Product Specification*

9. *Intel® Integrated RAID Module RMS25PB080, RMS25PB040, RMS25CB080, and RMS25CB040 Hardware/Installation User Guide*

10. *Intel® Integrated RAID Module RMS25KB040, RMS25KB080, RMS25JB040, RMS25JB080 – Hardware/Installation User Guide*

11. *Intel Integrated RAID Module RMT3PB080 and RMT3CB080 – Hardware/Installation User Guide*

12. *Intel® Server Chassis P4000M Family Technical Product Specification*

13. *Intel® Server Chassis P4000M Family Service Guide*

14. *Intel® Remote Management Module 4 Technical Product Specification*

15. *Intel® Remote Management Module 4 and Integrated BMC Web Console User's Guide*