



Intel® Server Board S1400SP

Technical Product Specification



Intel order number G64248-003

Revision 2.1

May, 2014

Revision History

| Date | Revision Number | Modifications |
|-----------------|-----------------|---|
| September, 2012 | 1.0 | Initial release. |
| December, 2013 | 2.0 | <ul style="list-style-type: none">▪ In Table 1, add a note that SAS Module and eUSB SSD can't be supported at the same time.▪ Added support for Intel® Xeon® processor E5-2400 v2 product family |
| May, 2014 | 2.1 | Added PCI Express* Lane Mapping. |

Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2014 Intel Corporation. All rights reserved.

Table of Contents

| | |
|---|-----------|
| 1. Introduction | 1 |
| 1.1 Chapter Outline..... | 1 |
| 1.2 Server Board Use Disclaimer | 1 |
| 2. Overview | 2 |
| 2.1 Intel® Server Boards S1400SP Feature Set | 2 |
| 2.2 Server Board Layout..... | 4 |
| 2.2.1 Server Board Connector and Component Layout..... | 5 |
| 2.2.2 Server Board Mechanical Drawings | 6 |
| 2.2.3 Server Board Rear I/O Layout | 12 |
| 3. Functional Architecture | 14 |
| 3.1 Processor Support..... | 14 |
| 3.1.1 Processor Socket Assembly | 15 |
| 3.2 Processor Function Overview | 15 |
| 3.2.1 Intel® QuickPath Interconnect | 16 |
| 3.2.2 Integrated Memory Controller (IMC) and Memory Subsystem | 17 |
| 3.2.3 Processor Integrated I/O Module (IIO)..... | 25 |
| 3.3 Intel® C600-A Chipset Functional Overview | 28 |
| 3.3.1 Low Pin Count (LPC) Interface | 29 |
| 3.3.2 Universal Serial Bus (USB) Controller | 29 |
| 3.3.3 On-board Serial Attached SCSI (SAS)/Serial ATA (SATA) Support and Options . | 30 |
| 3.3.4 Manageability..... | 32 |
| 3.4 Integrated Baseboard Management Controller (BMC) Overview | 33 |
| 3.4.1 Super I/O Controller..... | 34 |
| 3.4.2 Graphics Controller and Video Support | 35 |
| 3.4.3 Baseboard Management Controller..... | 36 |
| 4. System Security | 38 |
| 4.1 BIOS Password Protection | 38 |
| 4.2 Trusted Platform Module (TPM) Support..... | 39 |
| 4.2.1 TPM security BIOS | 39 |
| 4.2.2 Physical Presence | 40 |
| 4.2.3 TPM Security Setup Options | 40 |
| 4.3 Intel® Trusted Execution Technology (Intel® TXT)..... | 42 |
| 5. Technology Support | 44 |
| 5.1 Intel® Trusted Execution Technology..... | 44 |
| 5.2 Intel® Virtualization Technology – Intel® VT-x/VT-d/VT-c..... | 44 |
| 5.3 Intel® Intelligent Power Node Manager | 45 |
| 5.3.1 Hardware Requirements | 46 |
| 6. Platform Management Functional Overview..... | 47 |
| 6.1 Baseboard Management Controller (BMC) Firmware Feature Support..... | 47 |

| | | |
|---------|---|----|
| 6.1.1 | IPMI 2.0 Features | 47 |
| 6.1.2 | Non IPMI Features..... | 48 |
| 6.1.3 | New Manageability Features | 49 |
| 6.2 | Basic and Advanced Features..... | 50 |
| 6.3 | Integrated BMC Hardware: Emulex* Pilot III..... | 51 |
| 6.3.1 | Emulex* Pilot III Baseboard Management Controller Functionality | 51 |
| 6.4 | Advanced Configuration and Power Interface (ACPI) | 52 |
| 6.5 | Power Control Sources..... | 52 |
| 6.6 | BMC Watchdog..... | 53 |
| 6.7 | Fault Resilient Booting (FRB) | 53 |
| 6.8 | Sensor Monitoring..... | 54 |
| 6.9 | Field Replaceable Unit (FRU) Inventory Device | 54 |
| 6.10 | System Event Log (SEL) | 55 |
| 6.11 | System Fan Management..... | 55 |
| 6.11.1 | Thermal and Acoustic Management | 55 |
| 6.11.2 | Setting Throttling Mode..... | 56 |
| 6.11.3 | Altitude..... | 56 |
| 6.11.4 | Set Fan Profile | 56 |
| 6.11.5 | Fan PWM Offset | 56 |
| 6.11.6 | Quiet Fan Idle Mode | 56 |
| 6.11.7 | Fan Profiles..... | 57 |
| 6.11.8 | Thermal Sensor Input to Fan Speed Control | 57 |
| 6.11.9 | Memory Thermal Throttling..... | 59 |
| 6.12 | Messaging Interfaces..... | 59 |
| 6.12.1 | User Model | 60 |
| 6.12.2 | IPMB Communication Interface | 60 |
| 6.12.3 | LAN Interface | 60 |
| 6.12.4 | Address Resolution Protocol (ARP)..... | 66 |
| 6.12.5 | Internet Control Message Protocol (ICMP)..... | 66 |
| 6.12.6 | Virtual Local Area Network (VLAN) | 66 |
| 6.12.7 | Secure Shell (SSH)..... | 67 |
| 6.12.8 | Serial-over-LAN (SOL 2.0)..... | 67 |
| 6.12.9 | Platform Event Filter (PEF) | 68 |
| 6.12.10 | LAN Alerting..... | 68 |
| 6.12.11 | Alert Policy Table | 69 |
| 6.12.12 | SM-CLP (SM-CLP Lite) | 69 |
| 6.12.13 | Embedded Web Server | 70 |
| 6.12.14 | Virtual Front Panel | 71 |
| 6.12.15 | Embedded Platform Debug | 72 |
| 6.12.16 | Data Center Management Interface (DCMI)..... | 74 |
| 6.12.17 | Lightweight Directory Authentication Protocol (LDAP) | 74 |

| | |
|---|-----------|
| 7. Advanced Management Feature Support (RMM4)..... | 75 |
| 7.1 Keyboard, Video, Mouse (KVM) Redirection | 76 |
| 7.1.1 Remote Console | 77 |
| 7.1.2 Performance | 77 |
| 7.1.3 Security | 78 |
| 7.1.4 Availability | 78 |
| 7.1.5 Usage | 78 |
| 7.1.6 Force-enter BIOS Setup | 78 |
| 7.2 Media Redirection | 78 |
| 7.2.1 Availability | 79 |
| 7.2.2 Network Port Usage | 79 |
| 8. On-board Connector/Header Overview | 80 |
| 8.1 Board Connector Information | 80 |
| 8.2 Power Connectors | 81 |
| 8.3 Front Panel Headers and Connectors | 82 |
| 8.3.1 Front Panel Support | 82 |
| 8.3.2 Front Panel USB Connector | 84 |
| 8.3.3 Intel® Local Control Panel Connector | 85 |
| 8.4 On-Board Storage Connectors | 85 |
| 8.4.1 SATA Connectors | 85 |
| 8.4.2 Multiport Mini-SAS/SATA Connectors | 85 |
| 8.5 Fan Connectors | 86 |
| 8.6 Serial Port Connector | 87 |
| 8.7 System Management Headers | 88 |
| 8.7.1 Intel® Remote Management Module 4 Connector | 88 |
| 8.7.2 TPM connector | 88 |
| 8.7.3 HSBP Header | 89 |
| 8.7.4 SGPIO Header | 89 |
| 8.8 I/O Connectors | 89 |
| 8.8.1 VGA Connector | 89 |
| 8.8.2 NIC Connectors | 89 |
| 8.8.3 USB Connector | 90 |
| 8.9 Other Connectors and Headers | 91 |
| 9. Jumper Blocks | 92 |
| 9.1 BIOS Recovery Jumper | 93 |
| 9.2 Management Engine (ME) Firmware Force Update Jumper Block | 93 |
| 9.3 Password Clear Jumper Block | 94 |
| 9.4 BIOS Default Jumper Block | 94 |
| 9.5 BMC Force Update Jumper Block | 95 |
| 10. Intel® Light Guided Diagnostics | 96 |
| 10.1 System ID LED | 97 |

| | | |
|------------|---|------------|
| 10.2 | System Status LED..... | 97 |
| 10.3 | BMC Boot/Reset Status LED Indicators | 99 |
| 10.4 | Post Code Diagnostic LEDs | 99 |
| 10.5 | 5 Volt Stand-By Present LED | 100 |
| 10.6 | Fan Fault LEDs..... | 100 |
| 10.7 | Memory Fault LEDs | 100 |
| 11. | Environmental Limits Specification | 101 |
| 11.1 | Processor Thermal Design Power (TDP) Support..... | 101 |
| 11.2 | MTBF | 102 |
| 12. | Server Board Power Distribution | 103 |
| 12.1 | Processor Power Support..... | 103 |
| 12.2 | Power Supply Output Requirements | 104 |
| 12.2.1 | Output Power/Currents | 104 |
| 12.2.2 | Standby Output..... | 104 |
| 12.2.3 | Voltage Regulation | 104 |
| 12.2.4 | Dynamic Loading | 104 |
| 12.2.5 | Capacitive Loading | 105 |
| 12.2.6 | Grounding | 105 |
| 12.2.7 | Closed loop stability | 105 |
| 12.2.8 | Residual Voltage Immunity in Standby mode | 105 |
| 12.2.9 | Common Mode Noise | 105 |
| 12.2.10 | Soft Starting | 106 |
| 12.2.11 | Zero Load Stability Requirements | 106 |
| 12.2.12 | Ripple/Noise | 106 |
| 12.2.13 | Timing Requirements..... | 106 |
| 12.2.14 | Residual Voltage Immunity in Stand-by Mode | 108 |
| | Appendix A: Integration and Usage Tips | 109 |
| | Appendix B: Integrated BMC Sensor Tables | 110 |
| | Appendix C: POST Code Diagnostic LED Decoder | 134 |
| | Appendix D: POST Code Errors | 139 |
| | Appendix E: Supported Intel® Server Chassis..... | 145 |
| | Glossary | 147 |
| | Reference Documents | 150 |

List of Figures

| | |
|---|-----|
| Figure 1. Intel® Server Board S1400SP4 Layout | 4 |
| Figure 2. Intel® Server Board S1400SP Layout | 5 |
| Figure 3. Jumper Block Identification | 5 |
| Figure 4. Intel® Server Board S1400SP – Mounting Hole Locations (1 of 2) | 6 |
| Figure 5. Intel® Server Board S1400SP – Mounting Hole Locations (2 of 2) | 7 |
| Figure 6. Intel® Server Boards S1400SP – Major Connector Pin-1 Locations (1 of 2) | 8 |
| Figure 7. Intel® Server Boards S1400SP – Major Connector Pin-1 Locations (2 of 2) | 9 |
| Figure 8. Intel® Server Boards S1400SP – Primary Side Keepout Zone | 10 |
| Figure 9. Intel® Server Boards S1400SP – Primary Side Card Side Keepout Zone | 11 |
| Figure 10. Intel® Server Boards S1400SP – Primary Side Air Duct Keepout Zone | 11 |
| Figure 11. Intel® Server Boards S1400SP – Second Side Keepout Zone | 12 |
| Figure 12. Intel® Server Boards S1400SP2 Rear I/O Layout..... | 12 |
| Figure 13. Intel® Server Boards S1400SP4 Rear I/O Layout..... | 13 |
| Figure 14. Intel® Server Board S1400SP Functional Block Diagram | 14 |
| Figure 15. Processor Socket Assembly | 15 |
| Figure 16. Intel® Server Board S1400SP DIMM Slot Layout..... | 19 |
| Figure 17. Functional Block Diagram of Processor IIO Sub-system | 26 |
| Figure 18. PCI Express* Lane Mapping..... | 27 |
| Figure 19. Server Board Layout - I/O Module Connector | 27 |
| Figure 20. Functional Block Diagram – Chipset Supported Features and Functions | 29 |
| Figure 21. Intel® RAID C600 Upgrade Key Connector..... | 30 |
| Figure 22. Integrated Baseboard Management Controller (BMC) Overview | 33 |
| Figure 23. Integrated BMC Hardware | 34 |
| Figure 24. Setup Utility – TPM Configuration Screen | 41 |
| Figure 25. Fan Speed Control Process..... | 58 |
| Figure 26. Intel® RMM4 Lite Activation Key Installation | 75 |
| Figure 27. Intel® RMM4 Dedicated Management NIC Installation | 76 |
| Figure 28. Jumper Blocks | 92 |
| Figure 29. On-Board Diagnostic LED Placement | 96 |
| Figure 30. Memory Slot Fault LED Locations | 97 |
| Figure 31. Power Distribution Block Diagram | 103 |
| Figure 32. Output Voltage Timing | 107 |
| Figure 33. Turn On/Off Timing (Power Supply Signals)..... | 108 |
| Figure 34. POST Code Diagnostic LED Layout..... | 134 |

List of Tables

| | |
|--|----|
| Table 1. Intel® Server Board S1400SP Feature Set | 2 |
| Table 2. UDIMM Support Guidelines | 18 |
| Table 3. RDIMM Support Guidelines | 18 |
| Table 4. Intel® Server Board S1400SP DIMM Nomenclature | 19 |
| Table 5. Supported Intel® I/O Module Options | 28 |
| Table 6. External RJ45 NIC Port LED Definition..... | 28 |
| Table 7. Intel® RAID C600 Upgrade Key Options | 30 |
| Table 8. Video Modes | 35 |
| Table 9. Video mode | 36 |
| Table 10. TPM Setup Utility – Security Configuration Screen Fields | 42 |
| Table 11. Intel® Intelligent Power Node Manager | 45 |
| Table 12. Basic and Advanced Features | 50 |
| Table 13. ACPI Power States | 52 |
| Table 14. Power Control Initiators..... | 52 |
| Table 15. Fan Profiles..... | 57 |
| Table 16. Messaging Interfaces..... | 60 |
| Table 17. Factory Configured PEF Table Entries | 68 |
| Table 18. Diagnostic Data..... | 73 |
| Table 19. Additional Diagnostics on Error..... | 74 |
| Table 20. RMM4 Option Kits..... | 75 |
| Table 21. Enabling Advanced Management Features | 76 |
| Table 22. Board Connector Matrix | 80 |
| Table 23. Main Power Connector Pin-out | 81 |
| Table 24. CPU Power Connector Pin-out | 82 |
| Table 25. Power Supply Auxiliary Signal Connector Pin-out | 82 |
| Table 26. Front Panel Header Pin-out | 82 |
| Table 27. Power/Sleep LED Functional States | 83 |
| Table 28. NMI Signal Generation and Event Logging..... | 84 |
| Table 29. Front Panel USB Connector Pin-out | 85 |
| Table 30. Intel® Local Control Panel Connector Pin-out | 85 |
| Table 31. AHCI SATA Controller Connector Pin-out | 85 |
| Table 32. Multiport SAS/SATA Connector Pin-out..... | 86 |
| Table 33. SSI 4-pin Fan Header Pin-out..... | 87 |
| Table 34. 10-pin Fan Header Pin-out..... | 87 |
| Table 35. Serial A Connector Pin-out | 87 |
| Table 36. Intel® RMM4 Connector Pin-out | 88 |
| Table 37. Intel® RMM4 – Lite Connector Pin-out | 88 |
| Table 38. TPM connector Pin-out | 88 |
| Table 39. HSBP_ I ² C Header Pin-out | 89 |

| | |
|---|-----|
| Table 40. SGPIO Header Pin-out | 89 |
| Table 41. VGA Connector Pin-out | 89 |
| Table 42. RJ-45 10/100/1000 NIC Connector Pin-out | 90 |
| Table 43. External USB Connector Pin-out | 90 |
| Table 44. Internal USB Connector Pin-out..... | 90 |
| Table 45. Internal Type-A USB Connector Pin-out | 91 |
| Table 46. Chassis Intrusion Header Pin-out ("CHAS_INTR")..... | 91 |
| Table 47. Hard Drive Activity Header Pin-out ("HDD_LED")..... | 91 |
| Table 48. Server Board Jumpers | 92 |
| Table 49. System Status LED State Definitions..... | 98 |
| Table 50. BMC Boot/Reset Status LED Indicators | 99 |
| Table 51. Server Board Design Specifications | 101 |
| Table 52. MTBF Estimate | 102 |
| Table 53. Intel® Xeon® Processor TDP Guidelines | 104 |
| Table 54. Load Ratings | 104 |
| Table 55. Voltage Regulation Limits | 104 |
| Table 56. Transient Load Requirements..... | 105 |
| Table 57. Capacitive Loading Conditions | 105 |
| Table 58. Ripples and Noise | 106 |
| Table 59. Output Voltage Timing | 106 |
| Table 60. Turn On/Off Timing | 107 |
| Table 61. Integrated BMC Core Sensors | 112 |
| Table 62. POST Progress Code LED Example | 135 |
| Table 63. POST Progress Codes | 135 |
| Table 64. MRC Progress Codes | 137 |
| Table 65. POST Progress LED Codes | 138 |
| Table 66. POST Error Codes and Messages | 140 |
| Table 67. POST Error Beep Codes | 143 |
| Table 68. Integrated BMC Beep Codes | 144 |

<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 S1400SP.

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

1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Overview
- Chapter 3 – Functional Architecture
- Chapter 4 – System Security
- Chapter 5 – Technology Support
- Chapter 6 – Platform Management Functional Overview
- Chapter 7 – Advanced Management Feature Support (RMM4)
- Chapter 8 – On-board Connector/Header Overview
- Chapter 9 – Jumper Blocks
- Chapter 10 – Intel® Light Guided Diagnostics
- Chapter 11 – Environmental Limits Specification
- Chapter 12 – Server Board Power Distribution
- Appendix A – Integration and Usage Tips
- Appendix B – Integrated BMC Sensor Tables
- Appendix C – POST Code Diagnostic LED Decoder
- Appendix D – POST Code Errors
- Appendix E – Supported Intel® Server Chassis
- Glossary
- Reference Documents

1.2 Server Board Use Disclaimer

Intel® Server Boards contain a number of high-density VLSI (Very Large Scale Integration) and power delivery components that require adequate airflow for cooling. Intel® ensures through its own chassis development and testing that when Intel® server building blocks are used together, the fully integrated system meets the intended thermal requirements of these components. It is the responsibility of the system integrator who chooses not to use Intel® developed server building blocks to consult vendor datasheets and operating parameters to determine the amount of airflow required for their specific application and environmental conditions. Intel Corporation cannot be held responsible if components fail or the server board does not operate correctly when used outside any of the published operating or non-operating limits.

2. Overview

The Intel® Server Board S1400SP is monolithic printed circuit boards (PCBs) with features designed to support the 1U rack server markets. This server board is designed to support the Intel® Xeon® processor E5-2400 product family. Previous generation Intel® Xeon® processors are not supported.

2.1 Intel® Server Boards S1400SP Feature Set

Table 1. Intel® Server Board S1400SP Feature Set

| Feature | Description |
|------------------------------|---|
| Processor | Support for one Intel® Xeon® E5-2400 processor or one Intel® Xeon® E5-2400 v2 processor in an FC-LGA 1356 Socket B2 package with Thermal Design Power of up to 95W. |
| Memory | <ul style="list-style-type: none"> ▪ Three memory channels, six memory DIMMs (Two memory DIMMs per channel). ▪ Support for 1066/1333 MT/s Unbuffered (UDIMM) LVDDR3 or DDR3 memory. ▪ Support for 1066/1333/1600 MT/s ECC Registered (RDIMM) DDR3 memory. ▪ Support for 1066/1333 MT/s ECC Registered (RDIMM) LVDDR3 memory. ▪ No support for mixing of RDIMMs and UDIMMs. ▪ No support for LRDIMMs. ▪ No support for Quad Rank DIMMs. |
| Chipset | Intel® C602 (-A) chipset with support for storage option upgrade keys |
| Cooling Fan Support | Support for: <ul style="list-style-type: none"> ▪ One processor fan (4-pin header). ▪ Five front system fans (10-pin headers). ▪ One rear system fan (4-pin header). |
| Riser Card Slots | Slot 6: PCI Express* Gen3 x8 electrical with x16 physical connector, supporting for 1U riser card |
| External I/O control support | <ul style="list-style-type: none"> ▪ One RJ-45 serial port A connector. ▪ One DB-15 video connector. ▪ Two RJ-45 10/100/1000 Mb NIC connectors for S1400SP2. Four RJ-45 10/100/1000 Mb NIC connectors for S1400SP4. ▪ Four USB 2.0 ports in rear side. |
| Internal I/O control support | <ul style="list-style-type: none"> ▪ One 2x7 pin front video header . ▪ One internal 2x5 pin serial port B header¹. ▪ Two 2x5 pin USB headers, each providing front panel support for two USB ports. ▪ One internal Type-A USB 2.0 port. ▪ One 9 pin USB header for eUSB SSD ². ▪ One 1x7 pin header for optional Intel® Local Control Panel support. ▪ One 30 pin front panel connector. ▪ One PCI Express* Gen3 x 4 SAS Module connector ². ▪ One PCI Express* Gen3 x 8 I/O Module connector. |
| Available I/O Module Options | The following I/O modules utilize a single proprietary on-board connector. An installed I/O module can be supported in addition to standard on-board features and any add-in expansion cards. <ul style="list-style-type: none"> ▪ Quad port 1 GbE based on Intel® Ethernet Controller I350 – AXX4P1GBPWLIOM ▪ Dual RJ-45 port 10GBase-T I/O Module based on Intel® Ethernet Controller x540 – AXX10GBTWLIOM |

| Feature | Description |
|--------------------------------------|--|
| | <ul style="list-style-type: none"> ▪ Dual SFP+ port 10GbE module based on Intel® 82599 10 GbE controller – AXX10GBNIAIOM ▪ Single Port FDR 56GT/S speed InfiniBand* module with QSFP connector – AXX1FDRIBIOM ▪ Dual port FDR 56GT/S speed InfiniBand* module with QSFP connector – AXX2FDRIBIOM |
| Hard Drive and Optical Drive Support | <ul style="list-style-type: none"> ▪ Optical devices are supported. ▪ Two AHCI SATA connectors at 6Gbps, ▪ One AHCI 4-port mini-SAS connector, supporting four SATA ports at 3Gbps. ▪ Two SCU 4-port mini-SAS connectors, supporting up to eight SAS/SATA ports total at 3Gbps with optional Intel® C600 RAID Upgrade Keys, |
| RAID Support | <ul style="list-style-type: none"> ▪ Intel® RSTe SW RAID 0/1/10/5 ▪ LSI* SW RAID 0/1/10/5 |
| Video Support | <ul style="list-style-type: none"> ▪ Integrated 2D video controller. ▪ Dual monitor video mode is supported. ▪ 16 MB DDR3 Memory. |
| LAN | <ul style="list-style-type: none"> ▪ Two Gigabit Ethernet Ports through the Intel® Ethernet Controller I350 PHY (for S1400SP2). ▪ Four Gigabit Ethernet Ports through the Intel® Ethernet Controller I350 PHY (for S1400SP4). |
| Security | Intel® TPM module – AXXTPE5 (optional) |
| Server Management | <ul style="list-style-type: none"> ▪ Onboard ServerEngines* LLC Pilot III* Controller. ▪ Support for Intel® Remote Management Module 4 solutions (optional). ▪ Support for Intel® Remote Management Module 4 Lite solutions (optional). ▪ Intel® Light-Guided Diagnostics on field replaceable units. ▪ Support for Intel® System Management Software. ▪ Support for Intel® Intelligent Power Node Manager (Need PMBus*-compliant power supply). |
| Form Factor | SSI ATX 12"x9.6" compliant form factor. |
| Compatible Intel® Server Chassis | Intel® Server Chassis R1000SP Family. |

Note:

1. Internal 2x5 pin serial port B header does not function if power supply cannot support -12V output.
2. SAS Module and eUSB SSD can't be supported at the same time.

2.2 Server Board Layout




Figure 1. Intel® Server Board S1400SP4 Layout

2.2.1 Server Board Connector and Component Layout

The following figures show the layout of the server board.




Figure 2. Intel® Server Board S1400SP Layout

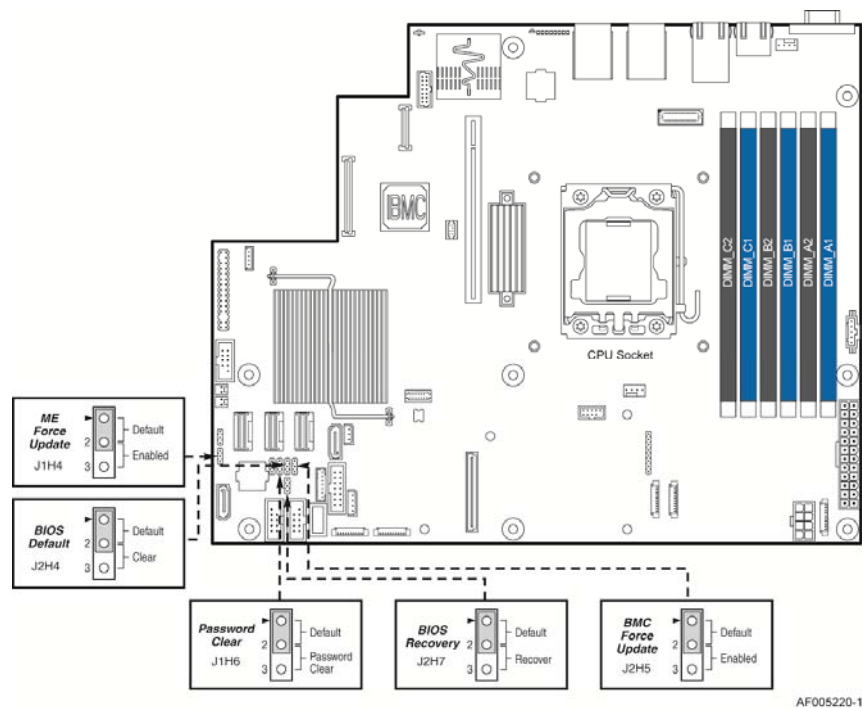


Figure 3. Jumper Block Identification

2.2.2 Server Board Mechanical Drawings




Figure 4. Intel® Server Board S1400SP – Mounting Hole Locations (1 of 2)

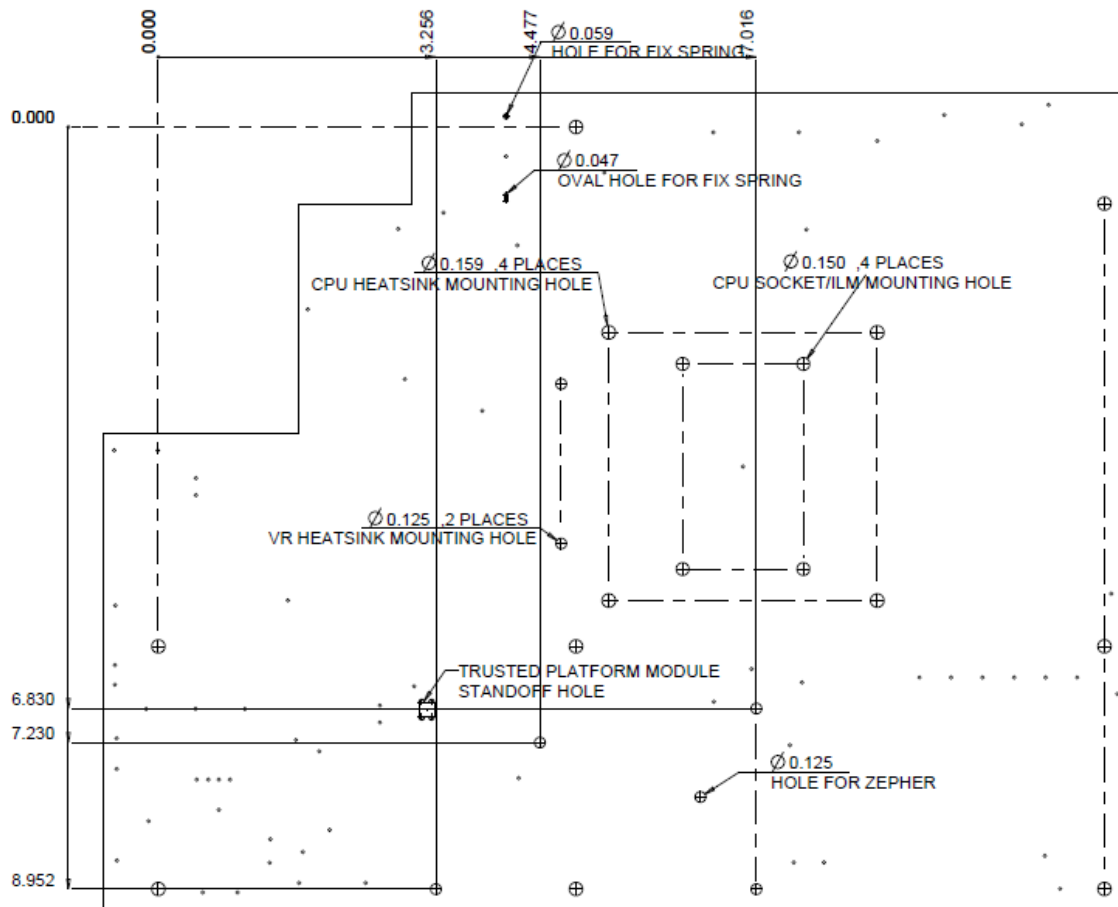


Figure 5. Intel® Server Board S1400SP – Mounting Hole Locations (2 of 2)




Figure 6. Intel® Server Boards S1400SP – Major Connector Pin-1 Locations (1 of 2)




Figure 7. Intel® Server Boards S1400SP – Major Connector Pin-1 Locations (2 of 2)




Figure 8. Intel® Server Boards S1400SP – Primary Side Keepout Zone




Figure 9. Intel® Server Boards S1400SP – Primary Side Card Side Keepout Zone




Figure 10. Intel® Server Boards S1400SP – Primary Side Air Duct Keepout Zone

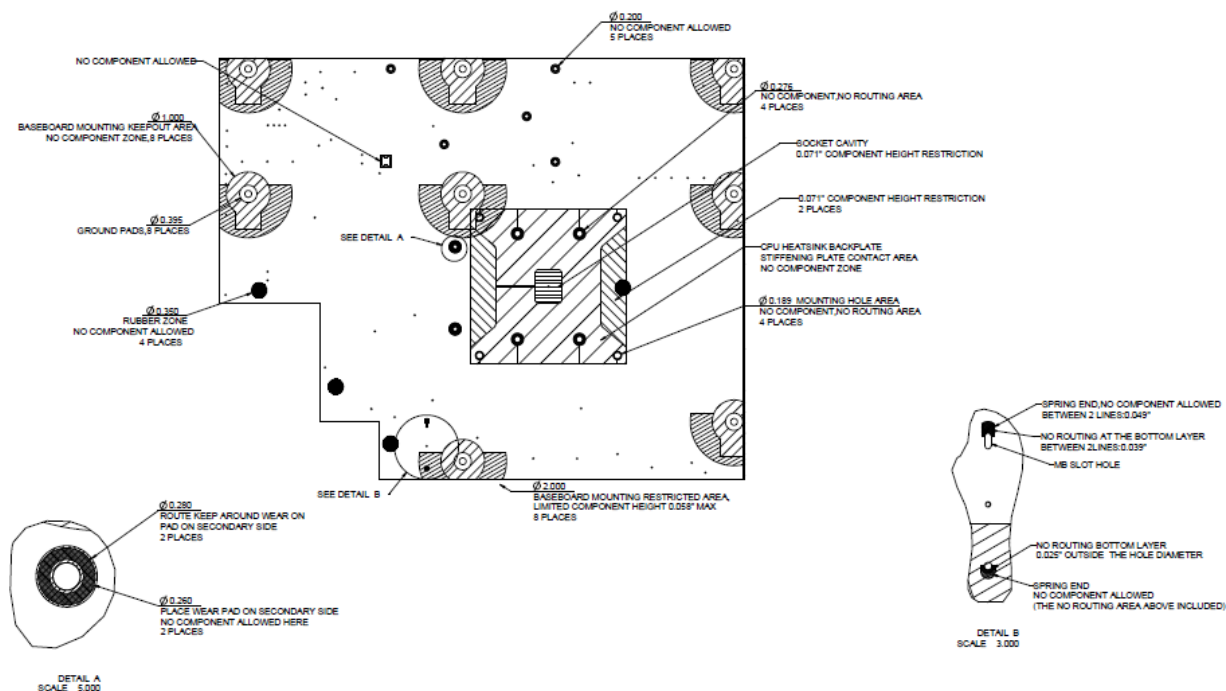


Figure 11. Intel® Server Boards S1400SP – Second Side Keepout Zone

2.2.3 Server Board Rear I/O Layout

The following drawings show the layout of the rear I/O components for the server board.




Figure 12. Intel® Server Boards S1400SP2 Rear I/O Layout

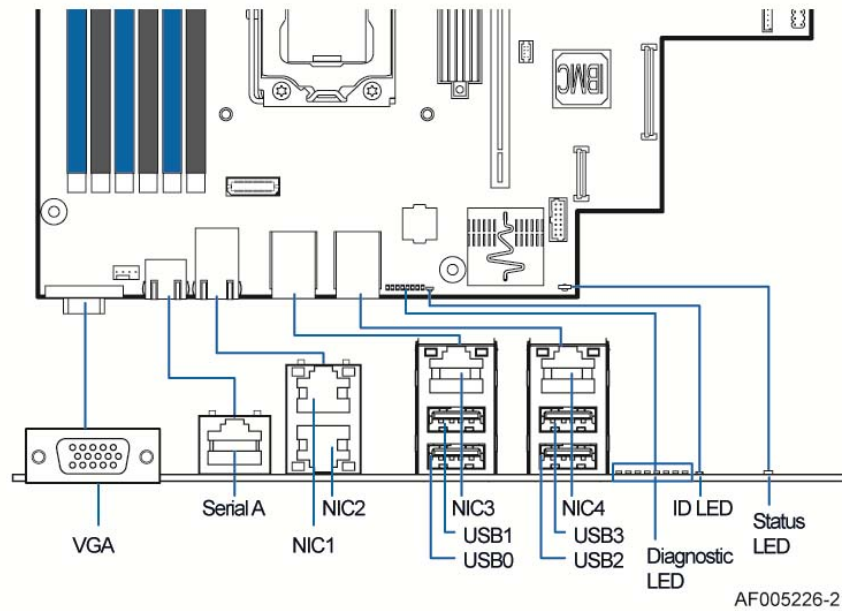


Figure 13. Intel® Server Boards S1400SP4 Rear I/O Layout

3. Functional Architecture

The architecture and design of the Intel® Server Board S1400SP is based on the Intel® C600 chipset family. The chipset is designed for systems based on the Intel® Xeon® processor in an FC-LGA 1356 Socket B2 package with Intel® QuickPath Interconnect (Intel® QPI).

This chapter provides a high-level description of the functionality associated with each chipset component and the architectural blocks that make up the server boards.

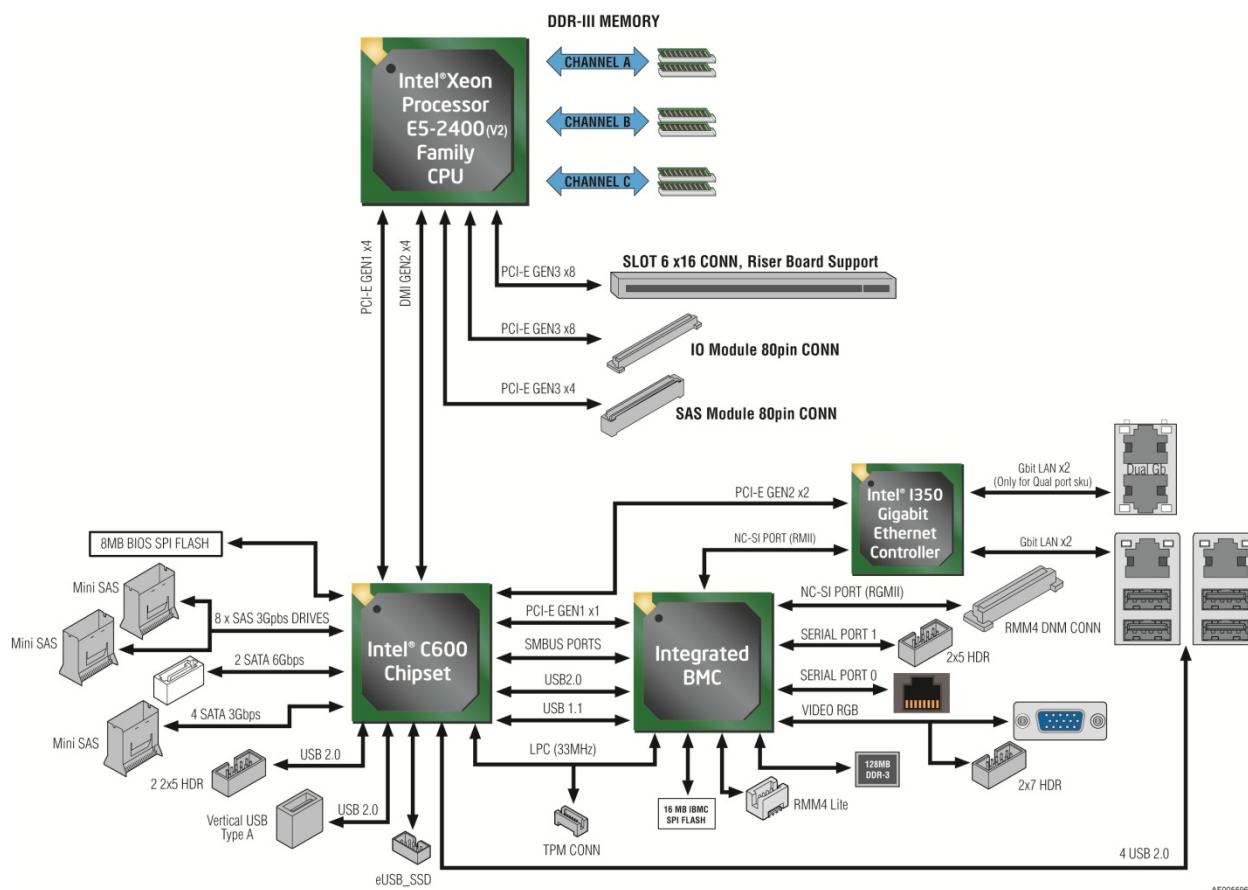


Figure 14. Intel® Server Board S1400SP Functional Block Diagram

3.1 Processor Support

The Intel® Server Board S1400SP includes one Socket-B2 (LGA-1356) processor socket and can support the following processor:

- Intel® Xeon® processor E5-2400 product family, with a Thermal Design Power (TDP) of up to 95W.
- Intel® Xeon® processor E5-2400 v2 product family, with a Thermal Design Power (TDP) of up to 95W.


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

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

3.1.1 Processor Socket Assembly

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

The illustration below identifies each sub-assembly component. (**Note:** The heat sink styles may vary.)



AF004320

Figure 15. Processor Socket Assembly

3.2 Processor Function Overview

With the release of the Intel® Xeon® processor E5-2400 and the Intel® Xeon® processor E5-2400 v2 product families, several key system components, including the CPU, Integrated Memory Controller (IMC), and Integrated IO Module (IIO), have been combined into a single processor package and feature per socket; One Intel® QuickPath Interconnect point-to-point links capable of up to 8.0 GT/s, up to 24 lanes of Gen 3 PCI Express* links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* Gen 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 performance and architecture of the server board.

Processor Feature Details:

- Up to 8 execution cores (Intel® Xeon® processor E5-2400 product family)
- Up to 10 execution cores (Intel® Xeon® processor E5-2400 v2 product family)
- Each core supports two threads (Intel® Hyper-Threading Technology)
- 46-bit physical addressing and 48-bit virtual addressing
- 1GB large page support for server applications
- A 32KB instruction and 32KB data first-level cache (L1) for each core
- A 256KB shared instruction/data mid-level (L2) cache for each core
- Up to 20MB last level cache (LLC): up to 2.5MB 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® Trusted Execution Technology (Intel® TXT)
- Intel® 64 Architecture
- Intel® Streaming SIMD Extensions 4.1 (Intel® SSE4.1)
- Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2)
- Intel® Advanced Vector Extensions (Intel® AVX)
- Intel® Hyper-Threading Technology
- Execute Disable Bit
- Intel® Turbo Boost Technology
- Intel® Intelligent Power Technology
- Enhanced Intel® SpeedStep Technology
- Intel® Data Direct I/O (DDIO) Technology

3.2.1 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

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

- Unbuffered DDR3 and registered DDR3 DIMMs
- 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 and DR – x4 and x8 data widths
- Up to 4 ranks supported per memory channel, 1, 2 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
 - Demand and Patrol Scrubbing
 - DRAM Single Device Data Correction (SDDC) for any single x4 or x8 DRAM device. Independent channel mode supports x4 SDDC. x8 SDDC requires lockstep mode
 - Lockstep mode where channels 0 and 1 and channels 2 and 3 are operated in lockstep mode
 - Data scrambling with address to ease detection of write errors to an incorrect address.
 - Error reporting from Machine Check Architecture
 - Read Retry during CRC error handling checks by iMC
 - Channel mirroring within a socket
 - CPU Channel Mirror Pairs B and C
 - 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 2. UDIMM Support Guidelines

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

Notes:

- Supported DRAM Densities are 1Gb, 2Gb and 4Gb. Only 2Gb and 4Gb are validated by Intel®.
- Command Address Timing is 1N for 1DPC and 2N for 2DPC.

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

Table 3. RDIMM Support Guidelines

| Ranks Per DIMM and Data Width | Memory Capacity Per DIMM1 | | | Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)2,3 | | | | | |
|-------------------------------|---------------------------|-----|------|--|----------------|---------------------|----------------|-----------|----------------|
| | | | | 1 Slot per Channel | | 2 Slots per Channel | | | |
| | | | | 1DPC | | 1DPC | | 2DPC | |
| | | | | 1.35V | 1.5V | 1.35V | 1.5V | 1.35V | 1.5V |
| SRx8 | 1GB | 2GB | 4GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| DRx8 | 2GB | 4GB | 8GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| SRx4 | 2GB | 4GB | 8GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |
| DRx4 | 4GB | 8GB | 16GB | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 | 1066 1333 | 1066 1333 1600 |

- Notes:**
- 1. Supported DRAM Densities are 1Gb, 2Gb and 4Gb. Only 2Gb and 4Gb are validated by Intel®.
 - 2. Command Address Timing is 1N.

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

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.

On the Intel® Server Board S1400SP, a total of 6 DIMM slots is provided. The nomenclature for DIMM sockets is detailed in the following table:

Table 4. Intel® Server Board S1400SP DIMM Nomenclature

| Processor Socket | | | | | |
|------------------|----|------------------|----|------------------|----|
| (0) Channel A | | (1) Channel B | | (2) Channel C | |
| A1 | A2 | B1 | B2 | C1 | C2 |




Figure 16. Intel® Server Board S1400SP DIMM Slot Layout

The following are generic DIMM population requirements that generally apply to the Intel® Server Board S1400SP.

- All DIMMs must be DDR3 DIMMs.
- Registered DIMMs must be ECC only; unbuffered DIMMs can be ECC or non-ECC. However, Intel® only validates and supports ECC memory for its server products.
- Mixing of Registered and Unbuffered DIMMs is not allowed per platform.
- Mixing of 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 DIMMs are NOT supported.
- LR (Load Reduced) DIMMs are NOT supported.
- A maximum of four 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 and dual rank DIMMs are populated for 2DPC, always populate the higher number rank DIMM first (starting from the farthest slot), for example, first dual rank, and then single rank DIMM.

DIMM population rules require that DIMMs within a channel be populated starting with the BLUE DIMM slot or DIMM farthest from the processor in a “fill-farthest” approach. Intel® MRC will check for correct DIMM placement.

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 displays the total system memory on the diagnostic screen at the end of POST. This total is the same as the amount described by the first bullet above.

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

1. ACPI (may vary depending on the number of PCI devices detected in the system)
2. ACPI NVS table
3. Processor microcode
4. Memory Mapped I/O (MMIO)
5. Manageability Engine (ME)
6. BIOS flash

3.2.2.4 RAS Features

The server board supports the following memory RAS modes:

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

Regardless of RAS mode, the requirements for populating within a channel given in the section 3.2.2.2 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 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).

3.2.2.4.1 Independent Channel Mode

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

3.2.2.4.2 Rank Sparing Mode

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

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

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

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

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

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

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

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

3.2.2.4.3 Mirrored Channel Mode

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

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

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

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

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

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

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

3.2.2.4.4 Lockstep Channel Mode

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

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

Performance in lockstep mode cannot be as high as with independent channels. The burst length for DDR3 DIMMs is eight which is shared between two channels that are in lockstep mode. Each channel of the pair provides 32 bytes to produce the 64-byte cache-line. DRAMs on independent channels are configured to deliver a burst length of eight. The maximum read bandwidth for a given Rank is half of peak. There is another 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.

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

3.2.2.4.5 Single Device Data Correction (SDDC)

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

3.2.2.4.6 Error Correction Code (ECC) Memory

ECC uses “extra bits” – 64-bit data in a 72-bit DRAM array – to add an 8-bit calculated “Hamming Code” to each 64 bits of data. This additional encoding enables the memory controller to detect and report single or multiple bit errors when data is read, and to correct single-bit errors.

3.2.2.4.6.1 Correctable Memory ECC Error Handling

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

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

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

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

3.2.2.4.6.2 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 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.

3.2.2.4.7 Demand Scrubbing for ECC Memory

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

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

3.2.2.4.8 Patrol Scrubbing for ECC Memory

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

3.2.3 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 24 lanes of PCI Express*. The key attribute of the PCI Express* interface is Gen3 speeds at 8 GT/s (no 8b/10b encoding)
- **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.
- **Intel® QuickData Technology:** Used for efficient, high bandwidth data movement between two locations in memory or from memory to I/O.

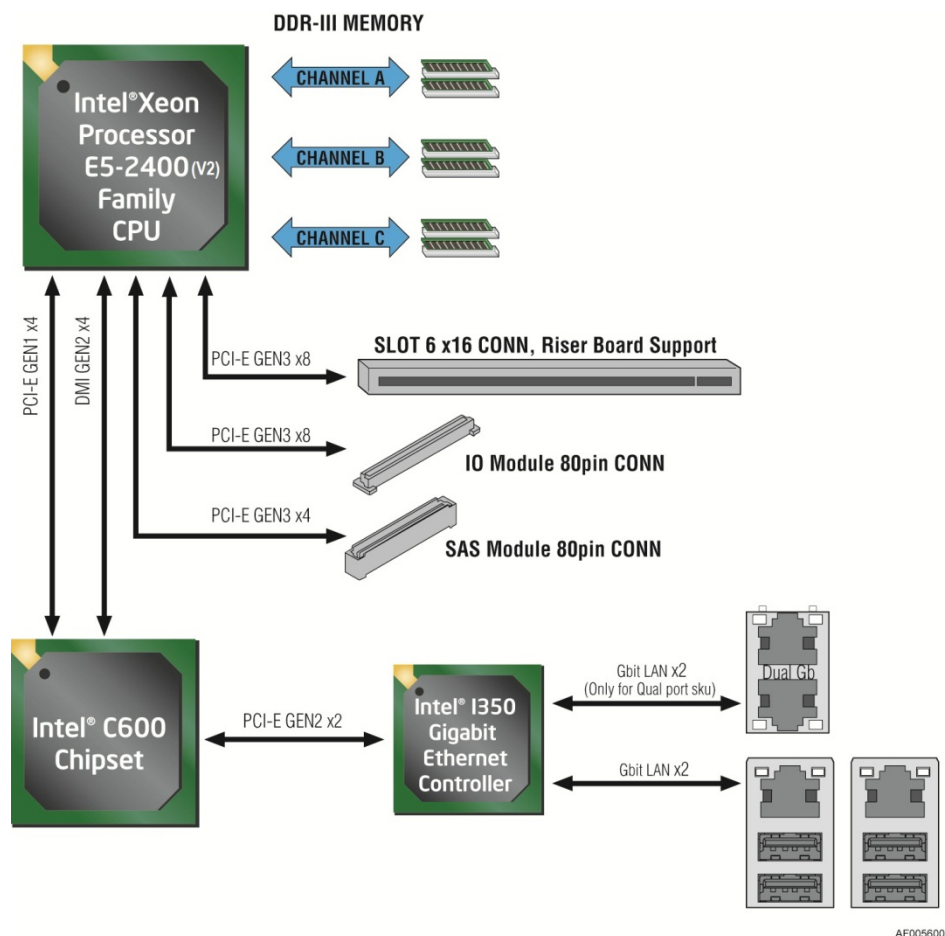


Figure 17. Functional Block Diagram of Processor I/O Sub-system

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

3.2.3.1 PCIe Interface

The integrated I/O module of the processor provides the PCI Express* interface to the server board. The processor supports three PCIe ports, identified as Port 1, Port 2, and Port 3, providing up to 24 PCIe bus lanes that meet the PCIe Gen 1, Gen2, and Gen 3 specifications.

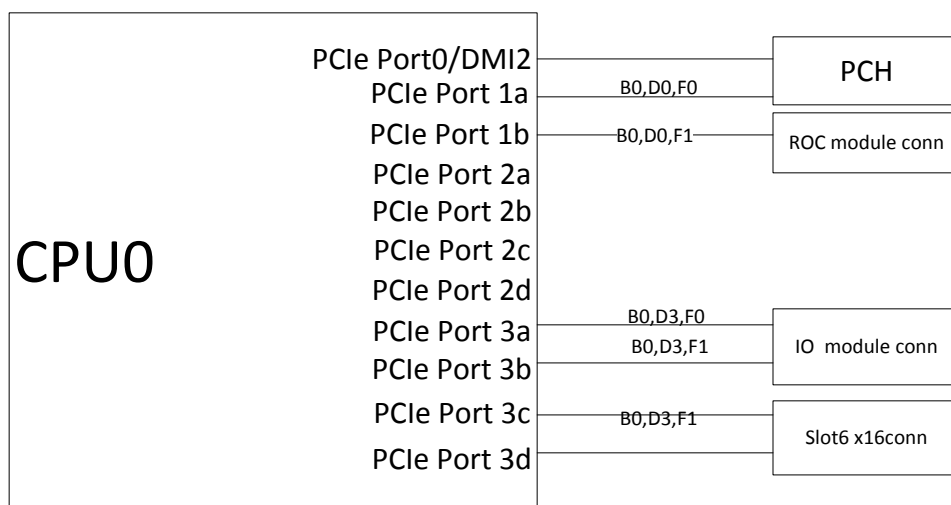


Figure 18. PCI Express* Lane Mapping

3.2.3.2 Riser Card Support

The server board includes one riser card slot labeled **SLOT_6 (PCIe_x16)**. See the *Technical Product Specifications* of the Intel® Server System R1000SP for available riser card options.

3.2.3.3 I/O Module Support

To broaden the standard on-board feature set, the server board provides support for one of several available IO Module options. The I/O module attaches to a high density 80-pin connector on the server board (J2C1) labeled **I/O_MOD** and is supported by x8 PCIe Gen3.




Figure 19. Server Board Layout - I/O Module Connector

Supported I/O modules include:

Table 5. Supported Intel® I/O Module Options

| Description | Intel® Product Code |
|---|---------------------|
| 4-port 1Gb Ethernet Networking IO Module | AXX4P1GBPWL1OM |
| 2-port 10Gb Ethernet Networking IO Module | AXX10GBTWL1OM |
| 2-port 10Gb Ethernet SFP IO Module | AXX10GBNIA1OM |
| FDR InfiniBand* IO Module | AXX1FDRIO1OM |

3.2.3.4 SAS Module


To broaden the standard on-board feature set, the server board provides support for one of several available SAS RAID IO Module options. The SAS module is attached to a mezzanine connector on the server board (J4J1) labeled **SAS_Module** and is supported by x4 PCIe Gen3 signals from the IIO module of the CPU processor.

3.2.3.5 Network Interface

Network connectivity is provided by means of one onboard Intel® Ethernet Controller I350 providing up to four 10/100/1000 Mb Ethernet ports. The NIC chip is supported by implementing x2 PCIe Gen2 signals from the Intel® C600 PCH.

On the Intel® Server Board S1400SP, two for S1400SP2 and four for S1400SP4 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 6. External RJ45 NIC Port LED Definition

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

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.

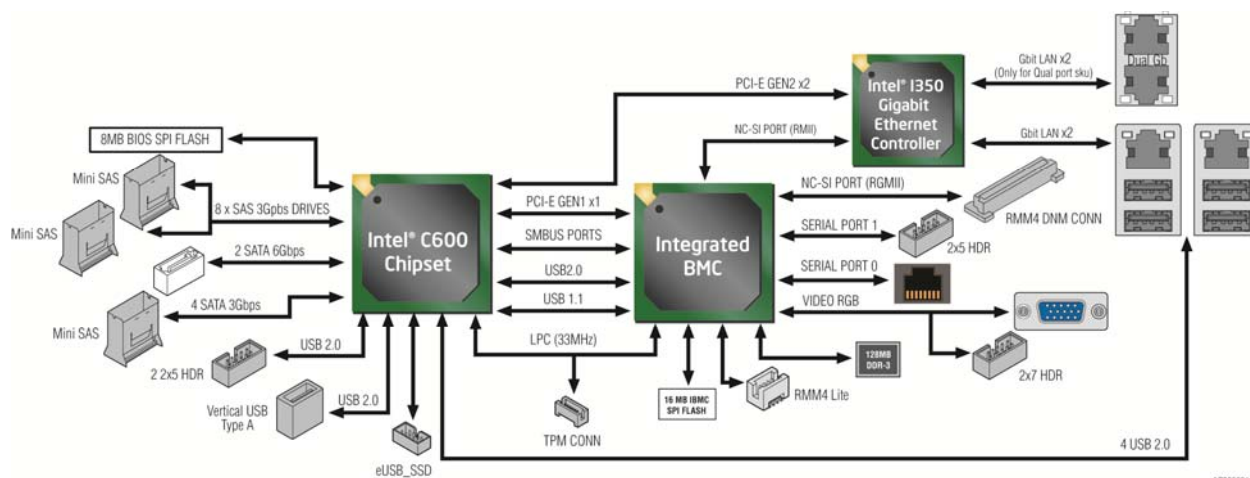


Figure 20. Functional Block Diagram – Chipset Supported Features and Functions

On the Intel® Server Boards S1400SP, the chipset provides support for the following on-board functions:

- PCI Express* root ports
- Low Pin Count (LPC) interface
- Universal Serial Bus (USB) Controller
- Serial Attached SCSI (SAS)/Serial ATA (SATA) Support
- Intel® Rapid Storage Technology
- Manageability Features

3.3.1 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 (Integrated BMC) as well as providing support for the optional Trusted Platform Module (TMP).

3.3.2 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 nine 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.
- Four USB ports are routed to two internal 10-pin connectors that can be cabled for front panel support.
- One internal Type 'A' USB port.


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

The Intel® C600-A chipset provides storage support from two integrated controllers: AHCI and SCU. By default the server board will support two single 6Gb/sec SATA ports routed from the AHCI controller to the two white SATA connectors labeled **SATA_0** and **SATA_1**, four 3Gb/sec SATA ports routed from the AHCI controller to the mini-SAS connector labeled **SATA_2-5**, and four 3Gb/sec SATA ports routed from the SCU to the mini-SAS connector labeled **SCU_0**.

Note: The mini-SAS connector labeled **SCU_1** is NOT functional by default and is only enabled with the addition of an Intel® RAID C600 Upgrade Key option supporting eight SAS/SATA ports.

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

- Intel® Embedded Server RAID Technology 2 (ESRT2) based on LSI* MegaRAID SW RAID technology supporting SATA RAID levels 0,1,10
- Intel® Rapid Storage Technology (RSTe) supporting SATA RAID levels 0,1,5,10



AF005524

Figure 21. Intel® RAID C600 Upgrade Key Connector

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

Table 7. Intel® RAID C600 Upgrade Key Options

| Product Code | Color | On-Server Board SATA/SAS Capable Controller | On-Server Board AHCI Capable SATA Controller |
|--------------|-------|---|--|
|--------------|-------|---|--|

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

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

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.

3.3.3.1 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 and 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key.

- 4 and 8 Port SAS RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key.
- Maximum drive support = Eight (with or without SAS expander option installed).
- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*.
- OS Support = 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.

3.3.3.2 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:
 - 4 Port SATA RAID 5 available standard (no option key required).
 - 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key.
 - No SAS RAID 5 support.
- Maximum drive support = 32 (in arrays with 8 port SAS), 16 (in arrays with 4 port SAS), 128 (JBOD).
- Open Source Compliance = Yes (uses MDRAID).
- OS Support = 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.

3.3.4 Manageability

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

- **TCO Timer.** The chipset's integrated programmable TCO timer is used to detect system locks. The first expiration of the timer generates an SMI# that the system can use to recover from a software lock. The second expiration of the timer causes a system reset to recover from a hardware lock.
- **Processor Present Indicator.** The chipset looks for the processor to fetch the first instruction after reset. If the processor does not fetch the first instruction, the chipset will reboot the system.

- **ECC Error Reporting.** When detecting an ECC error, the host controller has the ability to send one of several messages to the chipset. The host controller can instruct the chipset to generate SMI #, NMI, SERR#, or TCO interrupt.
- **Function Disable.** The chipset provides the ability to disable the following integrated functions: LAN, USB, LPC, SATA, PCI Express* or SMBus*. Once disabled, these functions no longer decode I/O, memory, or PCI configuration space. Also, no interrupts or power management events are generated from the disabled functions.
- **Intruder Detect.** The chipset provides an input signal (INTRUDER#) that can be attached to a switch that is activated by the system case being opened. The chipset can be programmed to generate an SMI# or TCO interrupt due to an active INTRUDER# signal.

3.4 Integrated Baseboard Management Controller (BMC) Overview

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

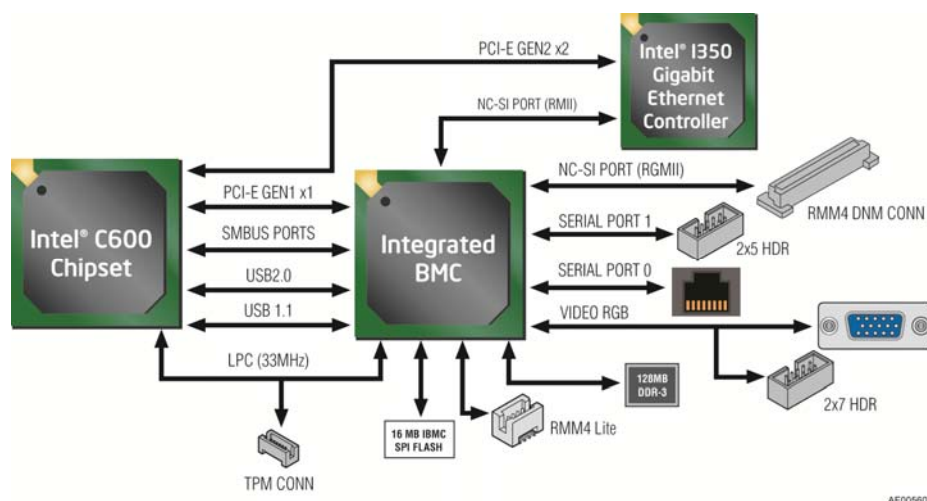


Figure 22. Integrated Baseboard Management Controller (BMC) Overview

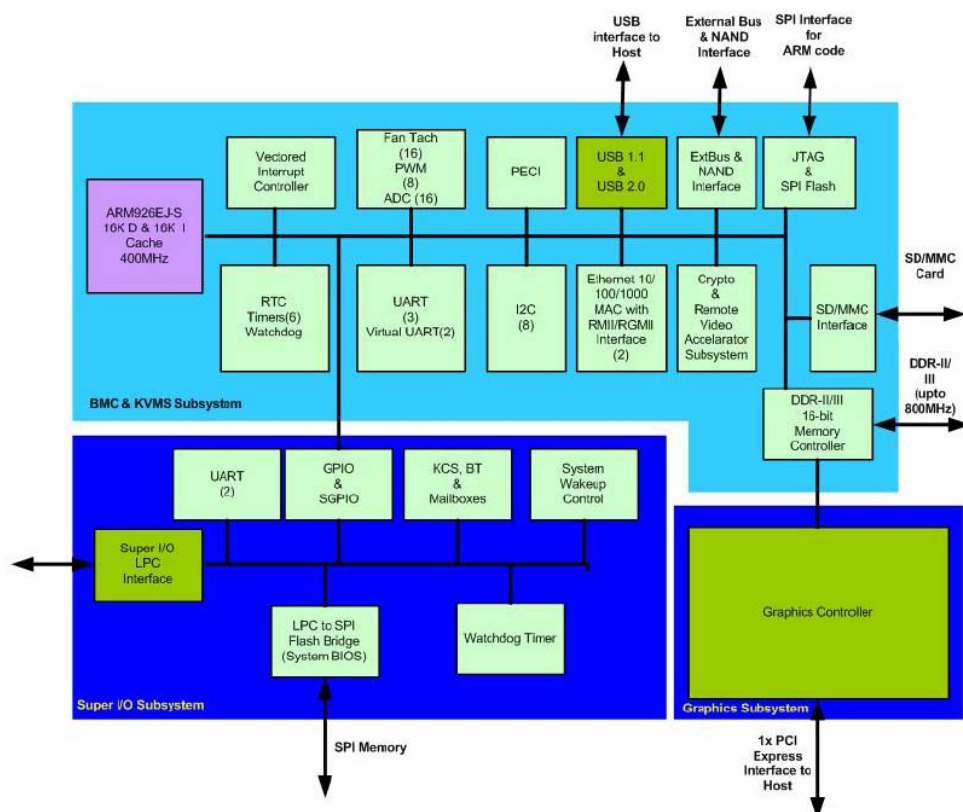


Figure 23. Integrated BMC Hardware

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 mouse.

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 supporting up to 128MB of memory, 16MB allocated to graphic
- Supports display resolutions up to 1600 x 1200 16bpp @ 60Hz
- High speed Integrated 24-bit RAMDAC
- Single lane PCI Express* host interface running at Gen 1 speed

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

Table 8. Video Modes

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

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

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

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

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

Table 9. Video mode

| | | |
|--------------------|---------------------|---|
| On-board Video | Enabled Disabled | - |
| Dual Monitor Video | Enabled Disabled | Shaded if on-board video is set to "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 RMI/IGMII support
- DDR2/3 16-bit interface with up to 800 MHz operation
- 12 10-bit ADCs
- Fourteen fan tachometers
- Eight Pulse Width Modulators (PWM)
- Chassis intrusion logic
- JTAG Master
- Eight I²C interfaces with master-slave and SMBus* timeout support. All interfaces are SMBus* 2.0 compliant.
- Parallel general-purpose I/O Ports (16 direct, 32 shared)
- Serial general-purpose I/O Ports (80 in and 80 out)
- Three UARTs
- Platform Environmental Control Interface (PECI)
- Six general-purpose timers
- Interrupt controller
- Multiple SPI flash interfaces
- NAND/Memory interface
- Sixteen mailbox registers for communication between the BMC and host
- LPC ROM interface
- BMC watchdog timer capability
- SD/MMC card controller with DMA support
- LED support with programmable blink rate controls on GPIOs
- Port 80h snooping capability
- Secondary Service Processor (SSP), which provides the HW capability of offloading time critical processing tasks from the main ARM core.

3.4.3.1 Remote Keyboard, Video, Mouse, 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.

4. System Security

4.1 BIOS Password Protection

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

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

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

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

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

The use of “Strong Passwords” is encouraged, but not required. In order to meet the criteria for a “Strong Password”, the password entered must be at least 8 characters in length, and must include at least one each of alphabetic, numeric, and special characters. If a “weak” password is entered, a popup warning message will be displayed, although the weak password will be accepted.

Once set, a password can be cleared by changing it to a null string. This requires the Administrator password, and must be done through BIOS Setup or other explicit means of changing the passwords. Clearing the Administrator password will also clear the User password.

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

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

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

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

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

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

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

4.2 Trusted Platform Module (TPM) Support

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

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

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

4.2.1 TPM security BIOS

The BIOS TPM support conforms to the *TPM PC Client Specific – Implementation Specification for Conventional BIOS, version 1.2*, and to the *TPM Interface Specification, version 1.2*. The BIOS adheres to the Microsoft Vista* BitLocker requirement. The role of the BIOS for TPM security includes the following:

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

- Produces EFI and legacy interfaces to a TPM-enabled operating system for using TPM.
- Produces ACPI TPM device and methods to allow a TPM-enabled operating system to send TPM administrative command requests to the BIOS.
- Verifies operator physical presence. Confirms and executes operating system TPM administrative command requests.
- Provides BIOS Setup options to change TPM security states and to clear TPM ownership.

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

4.2.2 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).

4.2.3 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.

4.2.3.1 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 S1400SP provides TPM settings through the security screen.

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

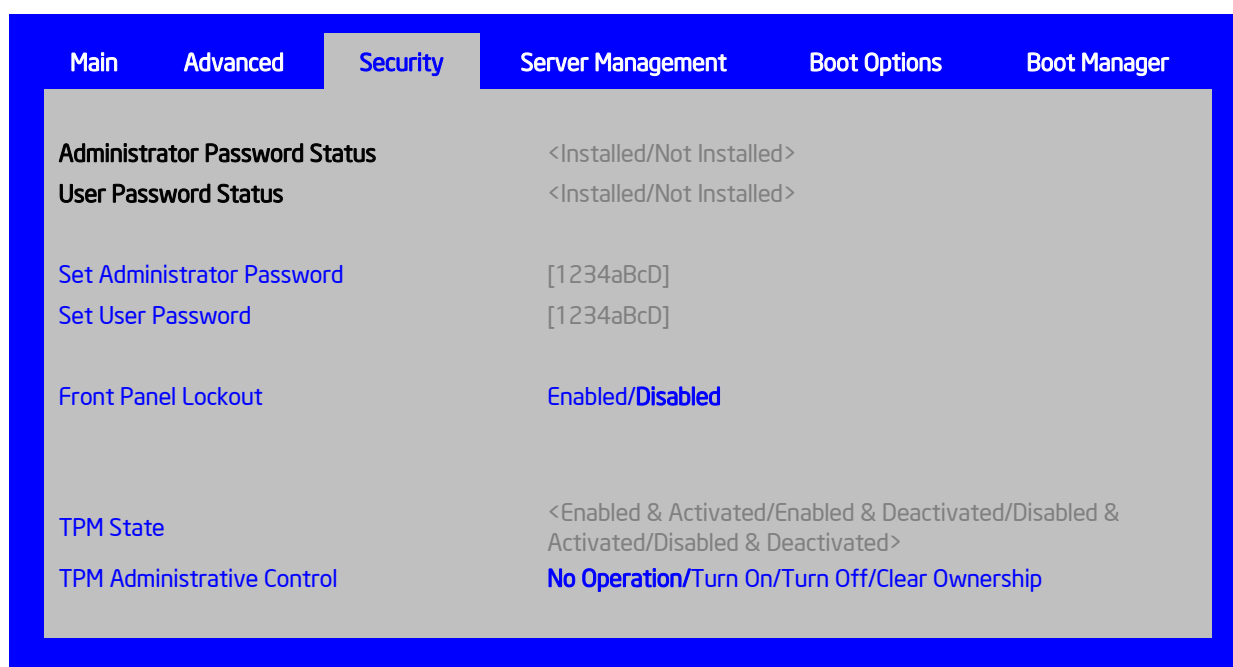


Figure 24. Setup Utility – TPM Configuration Screen

Table 10. TPM Setup Utility – Security Configuration Screen Fields

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

4.3 Intel® Trusted Execution Technology (Intel® TXT)

The Intel® Xeon® Processor E5-4600/2600/2400/1600 and Intel® Xeon® Processor E5 4600/2600/2400/1600 v2 Product Families support Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment. Designed to help protect against software-based attacks, Intel® Trusted Execution Technology integrates new security features and capabilities into the processor, chipset and other platform components. When used in conjunction with Intel® Virtualization Technology, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications.

This hardware-rooted security provides a general-purpose, safer computing environment capable of running a wide variety of operating systems and applications to increase the confidentiality and integrity of sensitive information without compromising the usability of the platform.

Intel® Trusted Execution Technology requires a computer system with Intel® Virtualization Technology enabled (both VT-x and VT-d), an Intel® Trusted Execution Technology-enabled

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

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

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

5. Technology Support

5.1 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5 4600/2600/2400/1600 and Intel® Xeon® Processor E5 4600/2600/2400/1600 v2 Product Families support Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment designed to help protect against software-based attacks. Intel® Trusted Execution Technology integrates new security features and capabilities into the processor, chipset and other platform components. When used in conjunction with Intel® Virtualization Technology and Intel® VT for Directed IO, with an active TPM, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications.

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

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

- **Intel® Virtualization Technology (VT-x)** is processor-related and provides capabilities needed to provide hardware assist to a Virtual Machine Monitor (VMM).
- **Intel® Virtualization Technology for Directed I/O (VT-d)** is primarily concerned with virtualizing I/O efficiently in a VMM environment. This would generally be a chipset I/O feature, but in the Second Generation Intel® Core™ Processor Family there is an Integrated I/O unit embedded in the processor, and the I/O 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 or Intel® Xeon® Processor E5 4600/2600/2400/1600 v2 Product Families and the C600 chipset family. Both support DMA remapping from inbound PCI Express* memory Guest Physical Address (GPA) to Host Physical Address (HPA). PCI devices are directly assigned to a virtual machine leading to a robust and efficient virtualization.

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

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

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

5.3 Intel® Intelligent Power Node Manager

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

Table 11. Intel® Intelligent Power Node Manager

| IT Challenge | Requirement |
|---------------------------------------|--|
| Over-allocation of power | <ul style="list-style-type: none"> 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 | <ul style="list-style-type: none"> Control capability that can maintain a power budget to enable increased rack population. |
| High energy costs | <ul style="list-style-type: none"> Control capability that can maintain a power budget to ensure that a set energy cost can be achieved. |
| Capacity planning | <ul style="list-style-type: none"> 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 | <ul style="list-style-type: none"> 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 family Management Engine (ME) and Intel® Intelligent Power Node Manager (NM) technology. The ME/NM combination is a power and thermal control capability on the platform, which exposes external interfaces that allow IT (through external management software) to query the ME about platform power capability and consumption, thermal characteristics, and specify policy directives (for example, set a platform power budget).

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

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

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

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

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

5.3.1 Hardware Requirements

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

6. Platform Management Functional Overview

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

This chapter provides a high level overview of the platform management features and functionality implemented on the server board. For more in depth and design level Platform Management information, please reference the *BMC Core Firmware External Product Specification (EPS)* and *BIOS Core External Product Specification (EPS)* for Intel® Server products based on the Intel® Xeon® processor E5-2400 and the Intel® Xeon® processor E5-2400 v2 product families.

6.1 Baseboard Management Controller (BMC) Firmware Feature Support

The following sections outline features that the integrated BMC firmware can support. Support and utilization for some features is dependent on the server platform in which the server board is integrated and any additional system level components and options that may be installed.

6.1.1 IPMI 2.0 Features

- Baseboard management controller (BMC)
- IPMI Watchdog timer
- Messaging support, including command bridging and user/session support
- Chassis device functionality, including power/reset control and BIOS boot flags support
- Event receiver device: The BMC receives and processes events from other platform subsystems.
- Field Replaceable Unit (FRU) inventory device functionality: The BMC supports access to system FRU devices using *IPMI FRU* commands.
- System Event Log (SEL) device functionality: The BMC supports and provides access to a SEL.
- Sensor Data Record (SDR) repository device functionality: The BMC supports storage and access of system SDRs.
- Sensor device and sensor scanning/monitoring: The BMC provides IPMI management of sensors. It polls sensors to monitor and report system health.
- IPMI interfaces
 - Host interfaces include system management software (SMS) with receive message queue support, and server management mode (SMM)
 - IPMB interface
 - LAN interface that supports the IPMI-over-LAN protocol (RMCP, RMCP+)
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS.

- BMC self test: The BMC performs initialization and run-time self-tests and makes results available to external entities.

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

6.1.2 Non IPMI Features

The BMC supports the following non-IPMI features:

- In-circuit BMC firmware update
- BMC FW reliability enhancements:
 - Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC.
 - BMC System Management Health Monitoring
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Enable/Disable of System Reset Due CPU Errors
- Chassis intrusion detection
- Fan speed control
- Fan redundancy monitoring and support
- Hot-swap fan support
- Power Supply Fan Sensors
- System Airflow Monitoring
- Exit Air Temperature Monitoring
- Acoustic management: Support for multiple fan profiles
- Ethernet Controller Thermal Monitoring
- Global Aggregate Temperature Margin Sensor
- Platform environment control interface (PECI) thermal management support
- Memory Thermal Management
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Power supply redundancy monitoring and support
- Power unit management: Support for power unit sensor. The BMC handles power-good dropout conditions.
- Intel® Intelligent Power Node Manager support
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.
- Local Control Display Panel support
- Power state retention
- Power fault analysis

- Intel® Light-Guided Diagnostics
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- E-mail alerting
- Embedded web server
 - Support for embedded web server UI in Basic Manageability feature set.
 - Human-readable SEL
 - Additional system configurability
 - Additional system monitoring capability
 - Enhanced on-line help
- Integrated KVM
- Integrated Remote Media Redirection
- Local Directory Access Protocol (LDAP) support
- Sensor and SEL logging additions/enhancements (for example, additional thermal monitoring capability)
- SEL Severity Tracking and the Extended SEL
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
 - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.1 compliance (product-specific).
- Management support for PMBus* rev1.2 compliant power supplies
- Energy Star Server Support
- Smart Ride Through (SmaRT)/Closed Loop System Throttling (CLST)
- Power Supply Cold Redundancy
- Power Supply FW Update
- Power Supply Compatibility Check

6.1.3 New Manageability Features

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

- Sensor and SEL logging additions/enhancements (for example, additional thermal monitoring capability)
- SEL Severity Tracking and the Extended SEL
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
 - Inventory data/system information export (partial SMBIOS table)

- Enhancements to fan speed control.
- DCMI 1.1 compliance (product-specific).
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server
 - Human-readable SEL
 - Additional system configurability
 - Additional system monitoring capability
 - Enhanced on-line help
- Enhancements to KVM redirection
 - Support for higher resolution
- Support for EU Lot6 compliance
- Management support for PMBus* rev1.2 compliant power supplies
- BMC Data Repository (Managed Data Region Feature)
- Local Control Display Panel
- System Airflow Monitoring
- Exit Air Temperature Monitoring
- Ethernet Controller Thermal Monitoring
- Global Aggregate Temperature Margin Sensor
- Memory Thermal Management
- Power Supply Fan Sensors
- Energy Star Server Support
- Smart Ride Through (SmaRT)/ Closed Loop System Throttling (CLST)
- Power Supply Cold Redundancy
- Power Supply FW Update
- Power Supply Compatibility Check
- BMC FW reliability enhancements:
 - Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC.
 - BMC System Management Health Monitoring

6.2 Basic and Advanced Features

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

Table 12. Basic and Advanced Features

| Feature | Basic | Advanced |
|--------------------------------|-------|----------|
| IPMI 2.0 Feature Support | X | X |
| In-circuit BMC Firmware Update | X | X |
| FRB 2 | X | X |
| Chassis Intrusion Detection | X | X |
| Fan Redundancy Monitoring | X | X |
| Hot-Swap Fan Support | X | X |

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

6.3 Integrated BMC Hardware: Emulex* Pilot III

6.3.1 Emulex* Pilot III Baseboard Management Controller Functionality

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

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

- 400MHz 32-bit ARM9 processor with memory management unit (MMU)
- Two independent 10/100/1000 Ethernet Controllers with Reduced Media Independent Interface (RMII)/ Reduced Gigabit Media Independent Interface (RGMI) 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 I²C interfaces with master-slave and SMBus* timeout support. All interfaces are SMBus* 2.0 compliant.
- Parallel general-purpose I/O Ports (16 direct, 32 shared)
- Serial general-purpose I/O Ports (80 in and 80 out)
- Three UARTs
- Platform Environmental Control Interface (PECI)
- Six general-purpose timers
- Interrupt controller
- Multiple Serial Peripheral Interface (SPI) flash interfaces
- NAND/Memory interface
- Sixteen mailbox registers for communication between the BMC and host
- LPC ROM interface
- BMC watchdog timer capability
- SD/MMC card controller with DMA support
- LED support with programmable blink rate controls on GPIOs

- Port 80h snooping capability
- Secondary Service Processor (SSP), which provides the HW capability of offloading time critical processing tasks from the main ARM core.

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

6.4 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states:

Table 13. ACPI Power States

| State | Supported | Description |
|-------|-----------|--|
| S0 | Yes | Working. <ul style="list-style-type: none"> ▪ The front panel power LED is on (not controlled by the BMC). ▪ The fans spin at the normal speed, as determined by sensor inputs. ▪ Front panel buttons work normally. |
| S1 | Yes | Sleeping. Hardware context is maintained; equates to processor and chipset clocks being stopped. <ul style="list-style-type: none"> ▪ 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 | Supported only on Workstation platforms. See appropriate <i>Platform Specific Information</i> for more information. |
| S4 | No | Not supported. |
| S5 | Yes | Soft off. <ul style="list-style-type: none"> ▪ The front panel buttons are not locked. ▪ The fans are stopped. ▪ The power-up process goes through the normal boot process. ▪ The power, reset, front panel NMI, and ID buttons are unlocked. |

6.5 Power Control Sources

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

Table 14. 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 |

| Source | External Signal Name or Internal Subsystem | Capabilities |
|-----------------------|---|--------------------------------------|
| Power state retention | Implemented by means of BMC internal logic | Turns power on when AC power returns |
| Chipset | Sleep S4/S5 signal (same as <i>POWER_ON</i>) | Turns power on or off |
| CPU Thermal | CPU Thermtrip | Turns power off |
| WOL(Wake On LAN) | LAN | Turns power on |

6.6 BMC Watchdog

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

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

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

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

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

6.7 Fault Resilient Booting (FRB)

Fault resilient booting (FRB) is a set of BIOS and BMC algorithms and hardware support that allow a multiprocessor system to boot even if the bootstrap processor (BSP) fails. Only FRB2 is supported using watchdog timer commands.

FRB2 refers to the FRB algorithm that detects system failures during POST. The BIOS uses the BMC watchdog timer to back up its operation during POST. The BIOS configures the watchdog timer to indicate that the BIOS is using the timer for the FRB2 phase of the boot operation. After the BIOS has identified and saved the BSP information, it sets the FRB2 timer use bit and loads the watchdog timer with the new timeout interval.

If the watchdog timer expires while the watchdog use bit is set to FRB2, the BMC (if so configured) logs a watchdog expiration event showing the FRB2 timeout in the event data bytes. The BMC then hard resets the system, assuming the BIOS-selected reset as the watchdog timeout action.

The BIOS is responsible for disabling the FRB2 timeout before initiating the option ROM scan and before displaying a request for a boot password. If the processor fails and causes an FRB2 timeout, the BMC resets the system.

The BIOS gets the watchdog expiration status from the BMC. If the status shows an expired FRB2 timer, the BIOS enters the failure in the system event log (SEL). In the OEM bytes entry in the SEL, the last POST code generated during the previous boot attempt is written. FRB2 failure is not reflected in the processor status sensor value.

The FRB2 failure does not affect the front panel LEDs.

6.8 Sensor Monitoring

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

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

The information gathered from physical sensors is translated into IPMI sensors as part of the “IPMI Sensor Model”. The BMC also reports various system state changes by maintaining virtual sensors that are not specifically tied to physical hardware.

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

6.9 Field Replaceable Unit (FRU) Inventory Device

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

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

6.10 System Event Log (SEL)

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

The BMC allocates 65,502 bytes (approximately 64 KB) of non-volatile storage space to store system events. The SEL timestamps may not be in order. Up to 3,639 SEL records can be stored at a time. Any command that results in an overflow of the SEL beyond the allocated space is rejected with an "Out of Space" IPMI completion code (C4h).

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

6.11 System Fan Management

The BMC controls and monitors the system fans. Each fan is associated with a fan speed sensor that detects fan failure and may also be associated with a fan presence sensor for hot-swap support. For redundant fan configurations, the fan failure and presence status determines the fan redundancy sensor state.

The system fans are divided into fan domains, each of which has a separate fan speed control signal and a separate configurable fan control policy. A fan domain can have a set of temperature and fan sensors associated with it. These are used to determine the current fan domain state.

A fan domain has three states: sleep, nominal, and boost. The sleep and boost states have fixed (but configurable through OEM SDRs) fan speeds associated with them. The nominal state has a variable speed determined by the fan domain policy. An OEM SDR record is used to configure the fan domain policy.

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

6.11.1 Thermal and Acoustic Management

The S1400SP offers multiple thermal and acoustic management features to maintain comprehensive thermal protection as well as intelligent fan speed control. The features can be adjusted in BIOS interface with path **BIOS > Advanced > System Acoustic and Performance Configuration**.

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

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

6.11.2 Setting Throttling Mode

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

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

The default setting is [Auto].

6.11.3 Altitude

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

6.11.4 Set Fan Profile

The [Performance] and [Acoustic] fan profiles in BIOS must be selected in [**BIOS > Advanced > System Acoustic and Performance Configuration > Set Fan Profile**]. The Acoustic mode offers best acoustic experience and appropriate cooling capability covering mainstream and majority of the add-in cards with 100LFM thermal requirements. For any add-in card requiring more than 100LFM, performance mode must be selected to provide sufficient cooling capability.

6.11.5 Fan PWM Offset

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

6.11.6 Quiet Fan Idle Mode

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

Note:

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

6.11.7 Fan Profiles

The server system supports multiple fan control profiles to support acoustic targets and American Society of Heating, Refrigerating, and Air Conditioning Engineers (ASHRAE) compliance. The BIOS Setup utility can be used to choose between meeting the target acoustic level or enhanced system performance. This is accomplished through fan profiles. The BMC supports eight fan profiles, numbered from 0 to 7.

Table 15. Fan Profiles

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

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

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

6.11.8 Thermal Sensor Input to Fan Speed Control

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

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

- Front Panel Temperature Sensor¹
- Baseboard Temperature Sensor²
- CPU Margin Sensors^{3,5,6}
- DIMM Thermal Margin Sensors^{3,5}
- Exit Air Temperature Sensor^{4, 8}
- PCH Temperature Sensor^{4,6}
- On-board Ethernet Controller Temperature Sensors^{4, 6}
- Add-In Intel® SAS/IO Module Temperature Sensors^{4,10}
- 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*
10. On-board of Module sensor

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

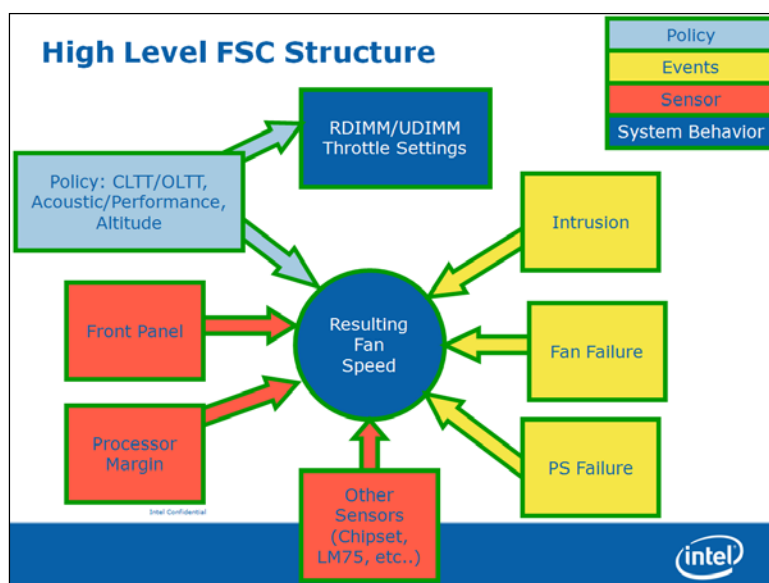


Figure 25. Fan Speed Control Process

6.11.9 Memory Thermal Throttling

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

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

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

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

6.12 Messaging Interfaces

The BMC supports the following communications interfaces:

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

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

Table 16. Messaging Interfaces

| Channel ID | Interface | Supports 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.

6.12.1 User Model

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

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

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

6.12.2 IPMB Communication Interface

The IPMB communication interface uses the 100 KB/s version of an I²C bus as its physical medium. For more information on I²C specifications, see *The I²C Bus and How to Use It*. The IPMB implementation in the BMC is compliant with the *IPMB v1.0, revision 1.0*.

The BMC IPMB slave address is 20h.

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

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

6.12.3 LAN Interface

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

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

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

6.12.3.1 RMCP/ASF Messaging

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

6.12.3.2 BMC LAN Channels

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

6.12.3.2.1 Baseboard NICs

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

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

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

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

6.12.3.2.2 Dedicated Management Channel

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

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

6.12.3.2.3 Concurrent Server Management Use of Multiple Ethernet Controllers

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

- 2 on-board NIC ports
- 1 on-board NIC and the optional dedicated RMM4 add-in management NIC.
- 2 on-board NICs and optional dedicated RMM4 add-in management NIC.

All NIC ports must be on different subnets for the above concurrent usage models. MAC addresses are assigned for management NICs from a pool of up to 3 MAC addresses allocated specifically for manageability.

The Intel® Server Board S1400SP4 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 Intel® Server Board S1400SP2 has five 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)
- BMC LAN channel 1 MAC address = NIC1 MAC address + 2
- BMC LAN channel 2 MAC address = NIC1 MAC address + 3
- BMC LAN channel 3 (RMM) MAC address = NIC1 MAC address + 4

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

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

- IP Address: Static
- All users disabled

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

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

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

Maximum bandwidth supported by BMC LAN channels are as follows:

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

- BMC LAN 3 (Dedicated NIC) ----- 1000Mb

6.12.3.3 IPV6 Support

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

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

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

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

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

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

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

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

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

6.12.3.4 LAN Failover

The BMC FW provides a LAN failover capability such that the failure of the system HW associated with one LAN link will result in traffic being rerouted to an alternate link. This functionality is configurable from IPMI methods as well as from the BMC's Embedded UI, allowing for user to specify the physical LAN links constitute the redundant network paths or physical LAN links constitute different network paths. BMC will support only a "all or nothing" approach – that is, all interfaces bonded together, or none are bonded together.

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

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

6.12.3.5 BMC IP Address Configuration

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

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

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

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

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

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

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

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

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

6.12.3.5.2 Static LAN Configuration Parameters

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

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

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

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

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

6.12.3.5.3 Enabling/Disabling Dynamic Host Configuration (DHCP) Protocol

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

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

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

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

6.12.3.5.4 DHCP-related LAN Configuration Parameters

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

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

To prevent users from disrupting the BMC's LAN configuration, the BMC treats these parameters as read-only while DHCP is enabled for the associated LAN channel. Using the *Set LAN Configuration Parameter* command to attempt to change one of these parameters under

such circumstances has no effect, and the BMC returns error code 0xD5, “Cannot Execute command. Command, or request parameter(s) are not supported in present state.”

6.12.3.6 DHCP BMC Hostname

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

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

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

6.12.4 Address Resolution Protocol (ARP)

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

6.12.5 Internet Control Message Protocol (ICMP)

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

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

6.12.6 Virtual Local Area Network (VLAN)

The BMC supports VLAN as defined by *IPMI 2.0 specifications*. VLAN is supported internally by the BMC, not through switches. VLAN provides a way of grouping a set of systems together so that they form a logical network. This feature can be used to set up a management VLAN where only devices which are members of the VLAN will receive packets related to management and

members of the VLAN will be isolated from any other network traffic. Please note that VLAN does not change the behavior of the host network setting, it only affects the BMC LAN communication.

LAN configuration options are now supported (by means of the *Set LAN Config Parameters* command, parameters 20 and 21) that allow support for 802.1Q VLAN (Layer 2). This allows VLAN headers/packets to be used for IPMI LAN sessions. VLAN ID's are entered and enabled by means of parameter 20 of the *Set LAN Config Parameters IPMI* command. When a VLAN ID is configured and enabled, the BMC only accepts packets with that VLAN tag/ID. Conversely, all BMC generated LAN packets on the channel include the given VLAN tag/ID. Valid VLAN ID's are 1 through 4094, VLAN ID's of 0 and 4095 are reserved, per the *802.1Q VLAN specification*. Only one VLAN can be enabled at any point in time on a LAN channel. If an existing VLAN is enabled, it must first be disabled prior to configuring a new VLAN on the same LAN channel. Parameter 21 (VLAN Priority) of the *Set LAN Config Parameters IPMI* command is now implemented and a range from 0-7 will be allowed for VLAN Priorities. Please note that bits 3 and 4 of Parameter 21 are considered Reserved bits.

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

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

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

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

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

6.12.7 Secure Shell (SSH)

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

6.12.8 Serial-over-LAN (SOL 2.0)

The BMC supports IPMI 2.0 SOL.

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

Three commands are implemented for SOL 2.0 configuration.

- *Get SOL 2.0 Configuration Parameters* and *Set SOL 2.0 Configuration Parameters*: These commands are used to get and set the values of the SOL configuration parameters. The parameters are implemented on a per-channel basis.

- *Activating SOL*: This command is not accepted by the BMC. It is sent by the BMC when SOL is activated to notify a remote client of the switch to SOL.
- Activating a SOL session requires an existing IPMI-over-LAN session. If encryption is used, it should be negotiated when the IPMI-over LAN session is established.

6.12.9 Platform Event Filter (PEF)

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

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

Table 17. Factory Configured PEF Table Entries

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

Additionally, the BMC supports the following PEF actions:

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

The “Diagnostic interrupt” action is not supported.

6.12.10 LAN Alerting

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

The BMC supports a minimum of four LAN alert destinations.

6.12.10.1 SNMP Platform Event Traps (PETs)

This feature enables a target system to send SNMP traps to a designated IP address by means of LAN. These alerts are formatted per the *Intelligent Platform Management Interface Specification Second Generation v2.0*. A Modular Information Block (MIB) file associated with

the traps is provided with the BMC firmware to facilitate interpretation of the traps by external software. The format of the MIB file is covered under RFC 2578.

6.12.11 Alert Policy Table

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

6.12.11.1 E-mail Alerting

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

6.12.12 SM-CLP (SM-CLP Lite)

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

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

The SM-CLP utilized by a remote user by connecting a remote system 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 an Serial Over LAN session.
- Support “help” to provide helpful information
- Get/set the system ID LED.
- Get the system GUID.
- Get/set configuration of user accounts.
- Get/set configuration of LAN parameters.
- Embedded CLP communication should support SSH connection.
- Provide current status of platform sensors including current values. Sensors include voltage, temperature, fans, power supplies, and redundancy (power unit and fan redundancy).

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

6.12.13 Embedded Web Server

BMC Base manageability provides an embedded web server and an OEM-customizable web GUI which exposes the manageability features of the BMC base feature set. It is supported over all on-board NICs that have management connectivity to the BMC as well as an optional RMM4 dedicated add-in management NIC. At least two concurrent web sessions from up to two different users is supported. The embedded web user interface shall support the following client web browsers:

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

The embedded web user interface supports strong security (authentication, encryption, and firewall support) since it enables remote server configuration and control. Embedded web server uses ports #80 and #443. The user interface presented by the embedded web user interface shall authenticate the user before allowing a web session to be initiated. Encryption using 128-bit SSL is supported. User authentication is based on user id and password.

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

Additional features supported by the web GUI includes:

- Presents all the Basic features to the users.
- Power on/off/reset the server and view current power state.
- Displays BIOS, BMC, ME, and SDR version information.
- Display overall system health.
- Configuration of various IPMI over LAN parameters for both IPv4 and IPv6.
- Configuration of alerting (SNMP and SMTP).
- Display system asset information for the product, board, and chassis.
- Display of BMC-owned sensors (name, status, current reading, enabled thresholds), including color-code status of sensors.
- Provides ability to filter sensors based on sensor type (Voltage, Temperature, Fan, and Power supply related).
- Automatic refresh of sensor data with a configurable refresh rate.
- On-line help.
- Display/clear SEL (display is in easily understandable human readable format).
- Supports major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*).
- Automatically logs out after user-configurable inactivity period.

- The GUI session automatically times-out after a user-configurable inactivity period. By default, this inactivity period is 30 minutes.
- Embedded Platform Debug feature - Allow the user to initiate a “diagnostic dump” to a file that can be sent to Intel® for debug purposes.
- Virtual Front Panel. The Virtual Front Panel provides the same functionality as the local front panel. The displayed LEDs match the current state of the local panel LEDs. The displayed buttons (for example, power button) can be used in the same manner as the local buttons.
- Severity level indication of SEL events. The web server UI displays the severity level associated with each event in the SEL. The severity level correlates with the front panel system status LED (“OK”, “Degraded”, “Non-Fatal”, or “Fatal”).
- Display of ME sensor data. Only sensors that have associated SDRs loaded will be displayed.
- Ability to save the SEL to a file.
- Ability to force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display of processor and memory information as is available over IPMI over LAN.
- Ability to get and set Node Manager (NM) power policies.
- Display of power consumed by the server.
- Ability to view and configure VLAN settings.
- Warn user the reconfiguration of IP address will cause disconnect.
- Capability to block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.
- Server Power Control – Ability to force into Setup on a reset.

6.12.14 Virtual Front Panel

- Virtual Front Panel is the module present as “Virtual Front Panel” on the left side in the embedded web server when “remote Control” tab is clicked.
- Main Purpose of the Virtual Front Panel is to provide the front panel functionality virtually.
- 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 from Virtual Front Panel, the reset will be done by a *Chassis control* command.
- For Reset from 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 from software command as there is no mechanism to get the chassis ID Led status.
- Only Infinite chassis ID ON/OFF from software command will reflect in EWS during automatic /manual EWS sync up with BMC.
- Virtual Front Panel help should be available for virtual panel module.
- At present, NMI button in VFP is disabled. It can be used in future.

6.12.15 Embedded Platform Debug

The Embedded Platform Debug feature supports capturing low-level diagnostic data (applicable MSRs, PCI config-space registers, and so on). This feature allows a user to export this data into a file that is retrievable 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
 - MCH registers
- BMC configuration data
 - BMC FW debug log (that is, SysLog) – Captures FW debug messages.
 - *Non-volatile storage of captured data.* Some of the captured data will be stored persistently in the BMC's non-volatile flash memory and preserved across AC power cycles. Due to size limitations of the BMC's flash memory, it is not feasible to store all of the data persistently.
- SMBIOS table data – The entire SMBIOS table is captured from the last boot.
- PCI configuration data for on-board devices and add-in cards – The first 256 bytes of PCI configuration data is captured for each device for each boot.
- System memory map – The system memory map is provided by BIOS on the current boot. This includes the EFI memory map and the Legacy (E820) memory map depending on the current boot.
- Power supplies debug capability:

- Capture of power supply “black box” data and power supply asset information – Power supply vendors are adding the capability to store debug data within the power supply itself. The platform debug feature provides a means to capture this data for each installed power supply. The data can be analyzed by Intel® for failure analysis and possibly provided to the power supply vendor as well. The BMC gets this data from the power supplies from 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 from an external IPMI interface (KCS or LAN).
- POST code sequence for the two most recent boots – This is a best-effort data collection by the BMC as the BMC real-time response cannot guarantee that all POST codes are captured.
- Support for multiple debug files – The platform debug feature provides the ability to save data to 2 separate files that are encrypted with different passwords.
 - File #1 is strictly for viewing by Intel® engineering and may contain BMC log messages (that is, syslog) and other debug data that Intel® FW developers deem useful in addition to the data specified in this document.
 - File #2 can be viewed by Intel® partners who have signed an NDA with Intel® and its contents are restricted to specific data items specified in this with the exception of the BMC syslog messages and power supply “black box” data.

6.12.15.1 Output Data Format

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

6.12.15.2 Output Data Availability

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

6.12.15.3 Output Data Categories

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

Table 18. 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 |

| Category | Data |
|--------------------|--|
| External BIOS Data | Human-readable sensor listing |
| | BIOS configuration settings |
| | POST codes for the two most recent boots |
| System Data | SMBIOS table for the current boot |
| | 256 bytes of PCI config data for each PCI device |
| | Memory Map (EFI and Legacy) for current boot |

Table 19. Additional Diagnostics on Error

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

6.12.16 Data Center Management Interface (DCMI)

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

6.12.17 Lightweight Directory Authentication Protocol (LDAP)

The Lightweight Directory Access Protocol (LDAP) is an application protocol supported by the BMC for the purpose of authentication and authorization. The BMC user connects with an LDAP server for login authentication. This is only supported for non-IPMI logins including the embedded web UI and SM-CLP. IPMI users/passwords and sessions are not supported over LDAP. LDAP can be configured (IP address of LDAP server, port, and so on) 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 Feature Support (RMM4)

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

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

Table 20. RMM4 Option Kits

| Intel® Product Code | Description | Kit Contents | Benefits |
|---------------------|--|---|---|
| AXXRMM4LITE | Intel® Remote Management Module 4 Lite | RMM4 Lite Activation Key | Enables KVM and media redirection from onboard NIC |
| AXXRMM4R | 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 with 1Gbe NIC |

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




Figure 26. Intel® RMM4 Lite Activation Key Installation




Figure 27. Intel® RMM4 Dedicated Management NIC Installation

Table 21. 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 (1Gbe) for management | Dedicated NIC for management traffic. Higher bandwidth connectivity for KVM and media Redirection with 1Gbe NIC. |

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

7.1 Keyboard, Video, Mouse (KVM) Redirection

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

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

KVM redirection console support the following keyboard layouts: English, Dutch, French, German, Italian, Russian, and Spanish.

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

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

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

KVM redirection feature supports the following resolutions and refresh rates:

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

7.1.1 Remote Console

The Remote Console is the redirected screen, keyboard and mouse of the remote host system. To use the Remote Console window of your managed host system, the browser must include a Java® Runtime Environment plug-in. If the browser has no Java support, such as with a small handheld device, the user can maintain the remote host system using the administration forms displayed by the browser.

The Remote Console window is a Java Applet that establishes TCP connections to the BMC. The protocol that is run over these connections is a unique KVM protocol and not HTTP or HTTPS. This protocol uses ports #7578 for KVM, #5120 for CDROM media redirection, and #5123 for Floppy/USB media redirection. When encryption is enabled, the protocol uses ports #7582 for KVM, #5124 for CDROM media redirection, and #5127 for Floppy/USB media redirection. The local network environment must permit these connections to be made, that is, the firewall and, in case of a private internal network, the NAT (Network Address Translation) settings have to be configured accordingly.

7.1.2 Performance

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

Enabling KVM and/or media encryption will degrade performance. Enabling video compression provides the fastest response while disabling compression provides better video quality.

For the best possible KVM performance, a 2Mb/sec link or higher is recommended. The redirection of KVM over IP is performed in parallel with the local KVM without affecting the local KVM operation.

7.1.3 Security

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

7.1.4 Availability

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

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

7.1.5 Usage

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

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

7.1.6 Force-enter BIOS Setup

KVM redirection can present an option to force-enter BIOS Setup. This enables the system to enter F2 setup while booting which is often missed by the time the remote console redirects the video.

7.2 Media Redirection

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

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

The following capabilities are supported:

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

- The media redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.
- A remote media session is maintained even when the server is powered-off (in standby mode). No restart of the remote media session is required during a server reset or power on/off. An BMC reset (for example, due to an BMC reset after BMC FW update) will require the session to be re-established
- The mounted device is visible to (and useable by) managed system's OS and BIOS in both pre-boot and post-boot states.
- The mounted device shows up in the BIOS boot order and it is possible to change the BIOS boot order to boot from this remote device.
- It is possible to install an operating system on a bare metal server (no OS present) using the remotely mounted device. This may also require the use of KVM-r to configure the OS during install.
- USB storage devices will appear as floppy disks over media redirection. This allows for the installation of device drivers during OS installation.

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

7.2.1 Availability

The default inactivity timeout is 30 minutes and is not user-configurable. Media redirection sessions persist across system reset but not across an AC power loss or BMC reset.

7.2.2 Network Port Usage

The KVM and media redirection features use the following ports:

- 5120 – CD Redirection
- 5123 – FD Redirection
- 5124 – CD Redirection (Secure)
- 5127 – FD Redirection (Secure)
- 7578 – Video Redirection
- 7582 – Video Redirection (Secure)

8. On-board Connector/Header Overview

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

8.1 Board Connector Information

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

Table 22. Board Connector Matrix

| Connector | Quantity | Reference Designators | Connector Type | Pin Count |
|---------------------|----------|--|-------------------------------------|--------------|
| Power supply | 3 | J9H1 J9J2 J9F1 | Main power CPU power PS AUX | 24 8 5 |
| CPU | 1 | U6E1 | CPU socket | 1356 |
| Main memory | 6 | DIMM A1: J9E2 DIMM B1: J8E4 DIMM C1: J8E2 DIMM A2: J9E1 DIMM B2: J8E3 DIMM C2: J8E1 | DIMM sockets | 240 |
| PCI Express* x16 | 1 | J4D1 | Card edge | 164 |
| Intel® RMM4 DNM | 1 | J3B2 | Connector | 40 |
| Intel® RMM4 Lite | 1 | J4D2 | Connector | 7 |
| Storage Upgrade Key | 1 | J1E1 | Header | 4 |
| System fans (front) | 5 | SYS FAN_1: J2J4 SYS FAN_2: J3J1 SYS FAN_3: J7J1 SYS FAN_4: J7J2 SYS FAN_5: J9J1 | Header | 10 |
| System fan (rear) | 1 | J8A2 | Header | 4 |
| CPU fan | 1 | J6G1 | Header | 4 |
| Battery | 1 | BT7G1 | Battery holder | 2 |
| Stacked RJ45/2xUSB | 2 | U7A1, U6A1 | Dual USB/LAN (RJ45) Combo Connector | 22 |
| Stacked RJ45 | 1 | JA7A1 | Dual LAN (RJ45) Connector | 28 |
| Video (rear) | 1 | J9A2 | Connector | 15 |
| Video (front) | 1 | J2J1 | Header | 14 |
| Serial port A | 1 | J8A1 | Connector | 9 |
| Serial port B | 1 | J2J3 | Header | 9 |
| Front panel | 1 | J1E3 | Header | 30 |
| Internal Dual USB | 2 | J1J2, J1F1 | Header | 9 |
| eUSB SSD | 1 | J6G2 | Header | 9 |
| Internal USB | 1 | J2J2 | Type-A USB | 4 |

| Connector | Quantity | Reference Designators | Connector Type | Pin Count |
|-----------------------|----------|--|----------------|-----------|
| HDD LED | 1 | J1G2 | Header | 2 |
| SATA | 2 | SATA_0: J2H3 SATA_1: J1J1 | Connector | 7 |
| Mini-SAS | 2 | SCU_0: J1H1 SCU_1: J1H2 | Connector | 38 |
| Mini-SATA | 1 | J2H1 | Connector | 38 |
| HSBP_I2C | 1 | J2H2 | Header | 3 |
| LCP | 1 | J2H8 | Header | 7 |
| IPMB | 1 | J2H6 | Header | 4 |
| Configuration jumpers | 5 | J2H5: BMC Force Update J2H7: BIOS Recovery J2H4: BIOS Default J1H4: ME Force Update J1H6: Password Clear | Jumper | 3 |
| TPM | 1 | J3G1 | Connector | 14 |
| SAS Module | 1 | J4J1 | Connector | 80 |
| IO Module | 1 | J2C1 | Connector | 80 |
| Chassis Intrusion | 1 | J1G1 | Header | 2 |

8.2 Power Connectors

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

Two additional power-related connectors also exist:

- One SSI-compliant 2x4 pin power connectors to provide 12-V power to the CPU voltage regulators and memory.
- One SSI-compliant 1x5 pin connector to provide I²C monitoring of the power supply.

The following tables define these connector pin-outs:

Table 23. Main Power Connector Pin-out

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

Table 24. CPU Power Connector Pin-out

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

Table 25. Power Supply Auxiliary Signal Connector Pin-out

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

8.3 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.

8.3.1 Front Panel Support

The server board provides a 30-pin front panel connector (J1C3) for use with Intel® and third-party chassis. The 30-pin connector consists of a 24-pin SSI compatible front panel connector and a 4-pin connector to support optional NIC 3/4 LEDs. The 24-pin SSI front panel connector provides 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. Front Panel Header Pin-out

| Signal Description | Pin# | Pin# | Signal Description |
|--------------------|------|------|--------------------|
| P3V3_AUX | 1 | 2 | P3V3_AUX |
| KEY | | 4 | P5V_STBY |
| FP_PWR_LED_BUF_R_N | 5 | 6 | FP_ID_LED_BUF_R_N |

| Signal Description | Pin# | Pin# | Signal Description |
|--------------------------|------|------|----------------------------|
| P3V3 | 7 | 8 | FP_LED_STATUS_GREEN_R_N |
| LED_HDD_ACTIVITY_R_N | 9 | 10 | FP_LED_STATUS_AMBER_R_N |
| FP_PWR_BTN_N | 11 | 12 | LED_NIC_LINK0_ACT_FP_N |
| GROUND | 13 | 14 | LED_NIC_LINK0_LNKUP_FP_N |
| FP_RST_BTN_R_N | 15 | 16 | SMB_SENSOR_3V3STBY_DATA_R0 |
| GROUND | 17 | 18 | SMB_SENSOR_3V3STBY_CLK |
| FP_ID_BTN_R_N | 19 | 20 | FP_CHASSIS_INTRUSION |
| PU_FM_SIO_TEMP_SENSOR | 21 | 22 | LED_NIC_LINK1_ACT_FP_N |
| FP_NMI_BTN_R_N | 23 | 24 | LED_NIC_LINK1_LNKUP_FP_N |
| KEY | | | KEY |
| LED_NIC_LINK2_ACT_FP_N | 27 | 28 | LED_NIC_LINK3_ACT_FP_N |
| LED_NIC_LINK2_LNKUP_FP_N | 29 | 30 | LED_NIC_LINK3_LNKUP_FP_N |

8.3.1.1 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. |

8.3.1.2 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.

8.3.1.3 System Reset Button Support

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

8.3.1.4 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

| Casual 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 |

8.3.1.5 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.

8.3.1.6 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.

8.3.1.7 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.

8.3.2 Front Panel USB Connector

The server board includes two 10-pin connector, that when cabled, can provide up to two USB ports to a front panel respectively. On the server board the connectors are labeled “USB_4-5” and “USB_6-7”. The following table provides the connector pin-out.

Table 29. Front Panel USB Connector Pin-out

| Signal Description | Pin# | Pin# | Signal Description |
|--------------------|------|------|--------------------|
| P5V_USB_FP | 1 | 2 | P5V_USB_FP |
| USB2_P11_F_DN | 3 | 4 | USB2_P13_F_DN |
| USB2_P11_F_DP | 5 | 6 | USB2_P13_F_DP |
| GROUND | 7 | 8 | GROUND |
| | | 10 | TP_USB2_FP_10 |

8.3.3 Intel® Local Control Panel Connector

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

Table 30. Intel® Local Control Panel Connector Pin-out

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

8.4 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.

8.4.1 SATA Connectors

The server board includes two white ports SATA connectors capable of transfer rates of up to 6Gb/s. On the server board this connector is labeled as “SATA_0” and “SATA_1”. The following table provides the pin-out for these connectors:

Table 31. AHCI SATA Controller Connector Pin-out

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

8.4.2 Multiport Mini-SAS/SATA Connectors

The server board includes three 40-pin high density multiport mini-SAS/SATA connectors labeled as “SCU_0”, “SCU_1” and “SATA_2-5”.

“SCU_0” and “SCU_1” connectors support up to four SATA or SAS ports each. By default, only the connector labeled “SCU_0” is enabled and has support for up to four SATA ports capable of

transfer rates of up to 3Gb/s. The connector labeled “SCU_1” is only enabled when an optional 8-port SAS or SATA Intel® RAID C600 Upgrade Key is installed. See Table 7 for a complete list of supported storage upgrade keys.

“SATA_2-5” connector supports up to four SATA ports capable of transfer rates of up to 3Gb/s. The following tables provide the pin-out for each connector.

Table 32. Multiport SAS/SATA Connector Pin-out

| Signal Description | Pin# | Pin# | Signal Description |
|------------------------|------|------|-------------------------|
| GROUND | A1 | B1 | GROUND |
| SAS0_RX_C_DP | A2 | B2 | SAS0_TX_C_DP |
| SAS0_RX_C_DN | A3 | B3 | SAS0_TX_C_DN |
| GROUND | A4 | B4 | GROUND |
| SAS1_RX_C_DP | A5 | B5 | SAS1_TX_C_DP |
| SAS1_RX_C_DN | A6 | B6 | SAS1_TX_C_DN |
| GROUND | A7 | B7 | GROUND |
| TP_SAS1_BACKPLANE_TYPE | A8 | B8 | SGPIO_SAS1_CLOCK |
| GROUND | A9 | B9 | SGPIO_SAS1_LOAD |
| SGPIO_SAS1_DATAOUT | A10 | B10 | GROUND |
| SGPIO_SAS1_DATAIN | A11 | B11 | PD_SAS1_CONTROLLER_TYPE |
| GROUND | A12 | B12 | GROUND |
| SAS2_RX_C_DP | A13 | B13 | SAS2_TX_C_DP |
| SAS2_RX_C_DN | A14 | B14 | SAS2_TX_C_DN |
| GROUND | A15 | B15 | GROUND |
| SAS3_RX_C_DP | A16 | B16 | SAS3_TX_C_DP |
| SAS3_RX_C_DN | A17 | B17 | SAS3_TX_C_DN |
| GROUND | A18 | B18 | GROUND |
| GROUND | MTH1 | MTH5 | GROUND |
| GROUND | MTH2 | MTH6 | GROUND |
| GROUND | MTH3 | MTH7 | GROUND |
| GROUND | MTH4 | MTH8 | GROUND |

8.5 Fan Connectors

The server board provides two SSI-compliant 4-pin and five 10-pin fan headers to use as CPU and I/O cooling fans. Each connector is monitored and controlled by on-board platform management.

The following table provides the pin-out for all fan connectors.

- One 4-pin fan header is designated as processor cooling fans:
 - CPU FAN
- One 4-pin fan header is designated as rear system fan:
 - SYS FAN_6
- Five 10-pin fan headers are designated as redundant system fans:
 - SYS FAN_1
 - SYS FAN_2
 - SYS FAN_3
 - SYS FAN_4
 - SYS FAN_5

Table 33. SSI 4-pin Fan Header Pin-out

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

Table 34. 10-pin Fan Header Pin-out

| SYS_FAN 1 | | SYS_FAN 2 | | SYS_FAN 3 | |
|-----------------------|------|-----------------------|------|-----------------------|------|
| Signal Description | Pin# | Signal Description | Pin# | Signal Description | Pin# |
| FAN_TACH1_IN | 1 | FAN_TACH3_IN | 1 | FAN_TACH5_IN | 1 |
| FAN_PWM_OUT_SYS1 | 2 | FAN_PWM_OUT_SYS2 | 2 | FAN_PWM_OUT_SYS3 | 2 |
| P12V_SYS_FAN12 | 3 | P12V_SYS_FAN12 | 3 | P12V_SYS_FAN34 | 3 |
| P12V_SYS_FAN12 | 4 | P12V_SYS_FAN12 | 4 | P12V_SYS_FAN34 | 4 |
| FAN_TACH0_IN | 5 | FAN_TACH2_IN | 5 | FAN_TACH4_IN | 5 |
| GROUND | 6 | GROUND | 6 | GROUND | 6 |
| GROUND | 7 | GROUND | 7 | GROUND | 7 |
| FAN_SYS1_PRSENT_BAT_N | 8 | FAN_SYS2_PRSENT_BAT_N | 8 | FAN_SYS3_PRSENT_BAT_N | 8 |
| LED_FAN_FAULT0_R | 9 | LED_FAN_FAULT1_R | 9 | LED_FAN_FAULT2_R | 9 |
| LED_FAN0 | 10 | LED_FAN1 | 10 | LED_FAN2 | 10 |
| SYS_FAN 4 | | SYS_FAN 5 | | | |
| Signal Description | Pin# | Signal Description | | | |
| FAN_TACH7_IN | 1 | FAN_TACH9_IN | 1 | | |
| FAN_PWM_OUT_SYS4 | 2 | FAN_PWM_OUT_SYS5 | 2 | | |
| P12V_SYS_FAN34 | 3 | P12V_SYS_FAN5 | 3 | | |
| P12V_SYS_FAN34 | 4 | P12V_SYS_FAN5 | 4 | | |
| FAN_TACH6_IN | 5 | FAN_TACH8_IN | 5 | | |
| GROUND | 6 | GROUND | 6 | | |
| GROUND | 7 | GROUND | 7 | | |
| FAN_SYS4_PRSENT_BAT_N | 8 | FAN_SYS5_PRSENT_BAT_N | 8 | | |
| LED_FAN_FAULT3_R | 9 | LED_FAN_FAULT4_R | 9 | | |
| LED_FAN3 | 10 | LED_FAN4 | 10 | | |

8.6 Serial Port Connector

The server board includes two serial port connectors. Serial-A is an external RJ45 type connector and has the following pin-out configuration.

**Table 35. Serial A Connector Pin-out**

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

8.7 System Management Headers

8.7.1 Intel® Remote Management Module 4 Connector

A 40-pin Intel® RMM4 connector and a 7-pin Intel® RMM4 Lite connector are included on the server board to support the optional Intel® Remote Management Module 4 or Intel® Remote Management Module 4 Lite. This server board does not support third-party management cards.

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

Table 36. Intel® RMM4 Connector Pin-out

| Pin | Signal Name | Pin | Signal Name |
|-----|-------------|-----|-------------|
| 1 | 3V3_AUX | 2 | MDIO |
| 3 | 3V3_AUX | 4 | MDC |
| 5 | GND | 6 | TXD_0 |
| 7 | GND | 8 | TXD_1 |
| 9 | GND | 10 | TXD_2 |
| 11 | GND | 12 | TXD_3 |
| 13 | GND | 14 | TX_CTL |
| 15 | GND | 16 | RX_CTL |
| 17 | GND | 18 | RXD_0 |
| 19 | GND | 20 | RXD_1 |
| 21 | GND | 22 | RXD_2 |
| 23 | GND | 24 | RXD_3 |
| 25 | GND | 26 | TX_CLK |
| 27 | GND | 28 | RX_CLK |
| 29 | GND | 30 | PRESENT# |
| 31 | Reserved | 32 | Reserved |
| 33 | Reserved | 34 | Reserved |
| 35 | Reserved | 36 | Reserved |
| 37 | Reserved | 38 | Reserved |
| 39 | Reserved | 40 | Reserved |

Table 37. Intel® RMM4 – Lite Connector Pin-out

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

8.7.2 TPM connector

Table 38. TPM connector Pin-out

| Pin | Signal Name | Pin | Signal Name |
|-----|-------------|-----|-------------|
| 1 | No pin | 2 | LPC_LAD<1> |
| 3 | LPC_LAD<0> | 4 | GND |
| 5 | IRQ_SERIAL | 6 | LPC_FRAME_N |

| Pin | Signal Name | Pin | Signal Name |
|-----|----------------|-----|------------------|
| 7 | P3V3 | 8 | GND |
| 9 | RST_IBMC_NIC_N | 10 | CLK_33M_TPM_CONN |
| 11 | LPC_LAD<3> | 12 | GND |
| 13 | GND | 14 | LPC_LAD<2> |

8.7.3 HSBP Header

Table 39. HSBP_ I²C Header Pin-out

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

8.7.4 SGPIO Header

Table 40. SGPIO Header Pin-out

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

8.8 I/O Connectors

8.8.1 VGA Connector

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

Table 41. VGA Connector Pin-out

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

8.8.2 NIC Connectors

The pin-out for NIC RJ-45 connectors is identical and defined in the following table.

Table 42. RJ-45 10/100/1000 NIC Connector Pin-out

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

8.8.3 USB Connector

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

Table 43. External USB Connector Pin-out

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

Two 2x5 connectors on the server board provide support for four additional USB ports.

Table 44. Internal USB Connector Pin-out

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

The server board includes one internal Type-A USB connector labeled “USB_8” to support the installation of a USB device inside the server chassis. The following table provides the pin-out for this connector.

Table 45. Internal Type-A USB Connector Pin-out

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

8.9 Other Connectors and Headers

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” and is located on the front edge of the server board. The header has the following pin-out.

Table 46. Chassis Intrusion Header Pin-out ("CHAS_INTR")

| Signal Description | Pin# |
|----------------------|------|
| FP_CHASSIS_INTRUSION | 1 |
| GROUND | 2 |

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

Table 47. Hard Drive Activity Header Pin-out ("HDD_LED")

| Signal Description | Pin# |
|--------------------|------|
| LED_HDD_ACT_N | 1 |
| TP_LED_HDD_ACT | 2 |

9. Jumper Blocks

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

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

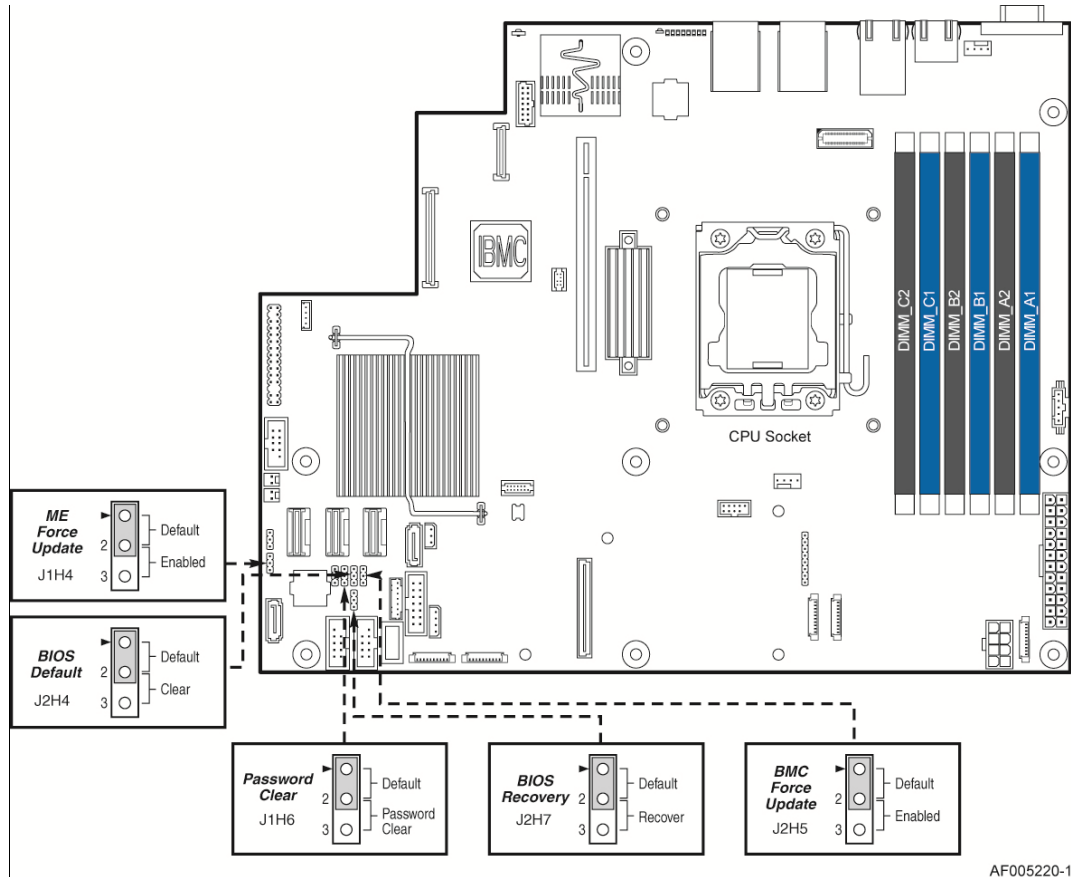


Figure 28. Jumper Blocks

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 website.

Table 48. Server Board Jumpers

| Jumper Name | Pins | System Results |
|------------------------|------|---|
| J2H5: BMC Force Update | 1-2 | BMC Firmware Force Update Mode – Disabled (Default) |
| | 2-3 | BMC Firmware Force Update Mode – Enabled |
| J2H7: BIOS Recovery | 1-2 | Pins 1-2 should be jumpered for normal system operation. (Default) |
| | 2-3 | The main system BIOS does not boot with pins 2-3 jumpered. The system only boots from EFI-bootable recovery media with a recovery BIOS image present. |
| J2H4: BIOS | 1-2 | These pins should have a jumper in place for normal system operation. (Default) |

| Jumper Name | Pins | System Results |
|-----------------------|------|--|
| Default | 2-3 | If pins 2-3 are jumpered with AC power plugged in, the CMOS settings clear in 5 seconds. Pins 2-3 should not be jumpered for normal system operation. |
| J1H4: ME Force Update | 1-2 | ME Firmware Force Update Mode – Disabled (Default) |
| | 2-3 | ME Firmware Force Update Mode – Enabled |
| J1H6: Password Clear | 1-2 | These pins should have a jumper in place for normal system operation. |
| | 2-3 | To clear administrator and user passwords, power on the system with pins 2-3 connected. The administrator and user passwords clear in 5-10 seconds after power on. Pins 2-3 should not be connected for normal system operation. |

9.1 BIOS Recovery Jumper

When the BIOS Recovery jumper block is moved from its default pin position, the system will boot into a BIOS Recovery Mode. It is used when the system BIOS has become corrupted and is non-functional, requiring a new BIOS image to be loaded on to the server board.

Note: The BIOS Recovery jumper is ONLY used to re-install a BIOS image in the event the BIOS has become corrupted. This jumper is NOT used when the BIOS is operating normally and you need to update the BIOS from one version to another.

The following steps demonstrate the BIOS recovery process:

1. After downloading the latest System Update Package (SUP) from the Intel® website, copy the following files to the root directory of a USB media device:
 - IPMI.EFI
 - IFlash32.EFI
 - RML.ROM
 - #####REC.CAP (where ##### = BIOS revision number)
 - STARTUP.NSH
2. Power OFF the system.
3. Locate the BIOS Recovery Jumper on the server board and move the jumper block from pins 1-2 (default) to pins 2-3 (recovery setting).
4. Insert the recovery media into a USB port.
5. Power ON the system.
6. The system will automatically boot into the embedded EFI Shell.
7. The STARTUP.NSH file automatically executes and initiates the flash update. When complete, the IFlash utility will display a message.
8. Power OFF the system and return the BIOS Recovery jumper to its default position.
9. Power ON the system.
10. Do ***NOT*** interrupt the BIOS POST during the first boot.
11. Configure desired BIOS settings.

9.2 Management Engine (ME) Firmware Force Update Jumper Block

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

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

1. Turn off the system and remove power cords.
2. Move the ME FRC UPD Jumper from the default (pins 1 and 2) operating position to the Force Update position (pins 2 and 3).
3. Re-attach system power cords.
4. Power on the system.

Note: System Fans will boost and the BIOS Error Manager should report an 83A0 error code (ME in recovery mode).

5. Boot to the EFI shell and update the ME firmware using the "MEComplete####.cap" file (where #### = ME revision number) using the following command: `iflash32 /u /ni MEComplete####.cap`.
6. When update has successfully completed, power off system.
7. Remove AC power cords.
8. Move ME FRC UPD jumper back to the default position.

Note: If the ME FRC UPD jumper is moved with AC power applied, the ME will not operate properly. The system will need have the AC power cords removed, wait for at least 10 seconds and then reinstalled to ensure proper operation.

9. Install AC power cords.
10. Power on system.

9.3 Password Clear Jumper Block

This jumper causes both the User password and the Administrator password to be cleared if they were set. The operator should be aware that this creates a security gap until passwords have been installed again through the BIOS Setup utility. This is the only method by which the Administrator and User passwords can be cleared unconditionally. Other than this jumper, passwords can only be set or cleared by changing them explicitly in BIOS Setup or by similar means. No method of resetting BIOS configuration settings to default values will affect either the Administrator or User passwords.

1. Power down the server and unplug the power cords.
2. Move jumper from the default (pins 1 and 2) operating position to the password clear position (pins 2 and 3).
3. Close the server chassis and reattach the power cords.
4. Power up the server and wait until POST completes.

Note: BIOS Error Manager should report a 5224 and 5221 error codes (Password clear jumper is set and Passwords cleared by jumper).

5. Power down the server and unplug the power cords.
6. Move the jumper back to the default position (covering pins 1 and 2).
7. Close the server chassis and reattach the power cords.
8. Power up the server.

9.4 BIOS Default Jumper Block

This jumper resets BIOS Setup options to their default factory settings.

1. Power down the server and unplug the power cords.
2. Move BIOS DFLT jumper from the default (pins 1 and 2) position to the Set BIOS Defaults position (pins 2 and 3).
3. Wait 5 seconds then move the jumper back to the default position of pins 1 and 2.
4. Install Power Cords.
5. Power on system.

Note: BIOS Error Manager should report a 5220 error code (BIOS Settings reset to default settings)

9.5 BMC Force Update Jumper Block

The BMC Force Update jumper is used to put the BMC in Boot Recovery mode for a low-level update.

It is used when the BMC has become corrupted and is non-functional, requiring a new BMC image to be loaded on to the server board.

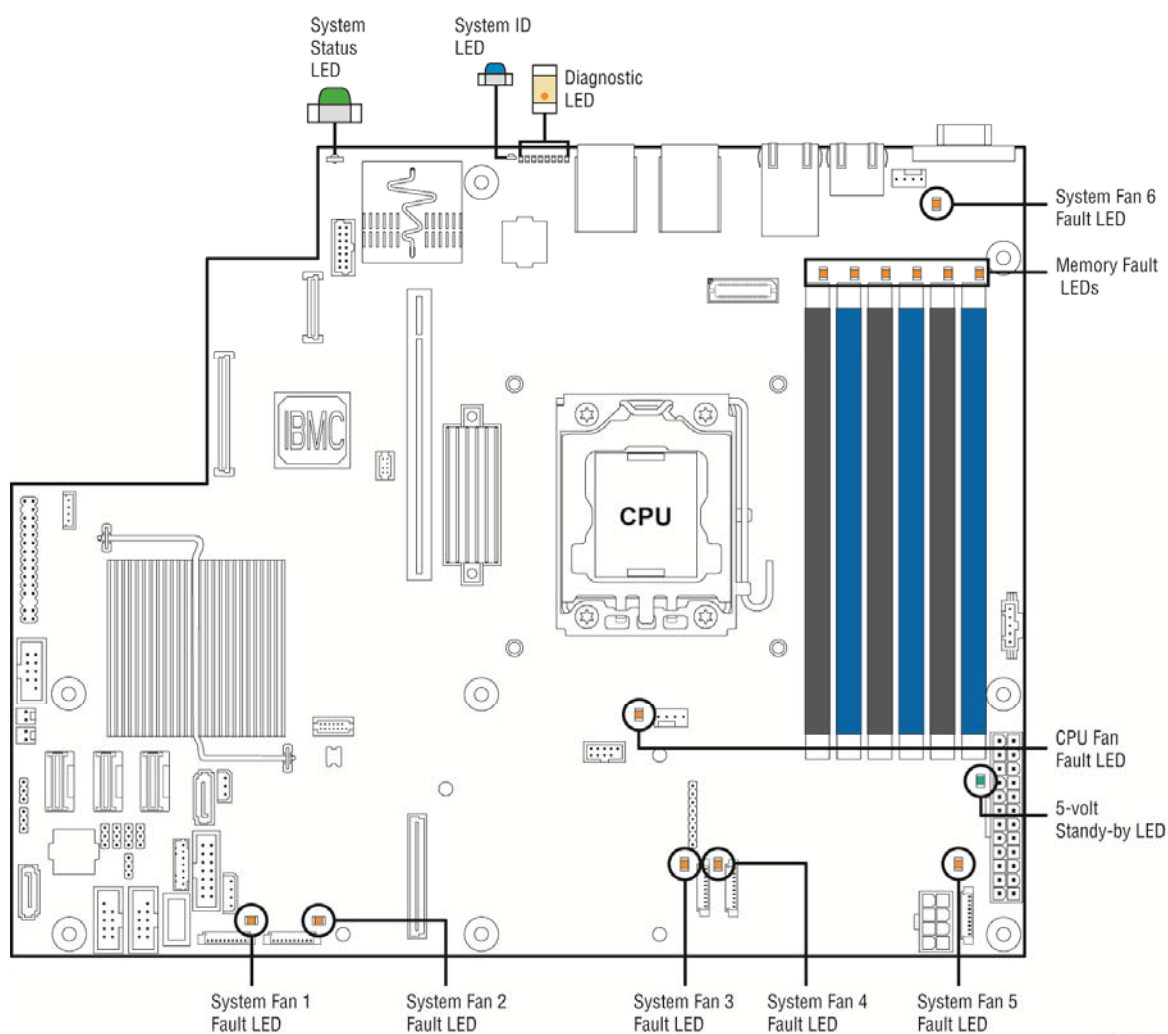
1. Turn off the system and remove power cords.
2. Move the *BMC FRC UPDT* Jumper from the default (pins 1 and 2) operating position to the Force Update position (pins 2 and 3).
3. Re-attach system power cords.
4. Power on the system.

Note: System Fans will boost and the BIOS Error Manager should report an 84F3 error code (Baseboard Management Controller in update mode).

5. Boot to the EFI shell and update the BMC firmware using BMC####.NSH (where #### is the version number of the BMC).
6. When update has successfully completed, power off system.
7. Remove AC power cords.
8. Move BMC FRC UPDT jumper back to the default position
9. Install AC power cords.
10. Power on system.
11. Boot to the EFI shell and update the FRU and SDR data using FRUSDR####.nsh (where #### is the version number of the FRUSDR package).
12. Reboot the system.
13. Configure desired BMC configuration settings.

10. Intel® Light Guided Diagnostics

The server board includes several on-board LED indicators to aid troubleshooting various board level faults. The following diagram shows the location for each.



AF005656

Figure 29. On-Board Diagnostic LED Placement




Figure 30. Memory Slot Fault LED Locations

10.1 System ID LED

The server board includes a blue system ID LED which is used to visually identify a specific server installed among many other similar servers. There are two options available for illuminating the System ID LED:

1. The front panel ID LED Button is pushed, which causes the LED to illuminate to a solid on state until the button is pushed again.
2. An *IPMI Chassis Identify* command is remotely entered, which causes the LED to blink.

The System ID LED on the server board is tied directly to the System ID LED on system front panel if present.

10.2 System Status LED

The server board includes a bi-color System Status LED. The System Status LED on the server board is tied directly to the System Status LED on the front panel (if present). This LED indicates the current health of the server. Possible LED states include solid green, blinking green, blinking amber, and solid amber.

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

When AC power is first applied to the system, the status LED turns solid amber and then immediately changes to blinking green to indicate that the BMC is booting. If the BMC boot process completes with no errors, the status LED will change to solid green.

Table 49. System Status LED State Definitions

| Color | State | Criticality | Description |
|-------|-------------|--------------|---|
| Green | Solid on | Ok | Indicates that the System Status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running. |
| Green | ~1 Hz blink | Degraded | <p>System degraded:</p> <ol style="list-style-type: none"> 1. Redundancy loss such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. 2. Fan warning or failure when the number of fully operational fans is more than minimum number needed to cool the system. 3. Non-critical threshold crossed – Temperature (including HSBP temp), voltage, input power to power supply, output current for main power rail from power supply and Processor Thermal Control (Therm Ctrl) sensors. 4. Power supply predictive failure occurred while redundant power supply configuration was present. 5. Unable to use all of the installed memory (more than 1 DIMM installed) 6. Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the user no longer has spared DIMMs indicating a redundancy lost condition. Corresponding DIMM LED lit. 7. In mirrored configuration, when memory mirroring takes place and system loses memory redundancy. 8. Battery failure. 9. BMC executing in uBoot. (Indicated by Chassis ID blinking at Blinking at 3Hz). System in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux*. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux* image into flash 10. BMC booting Linux*. (Indicated by Chassis ID solid ON). System in degraded state (no manageability). Control has been passed from BMC uBoot to BMC Linux* itself. It will be in this state for ~10~20 seconds. 11. BMC Watchdog has reset the BMC. 12. Power Unit sensor offset for configuration error is asserted. 13. HDD HSC is off-line or degraded. |
| Amber | ~1 Hz blink | Non-critical | <p>Non-fatal alarm – system is likely to fail:</p> <ol style="list-style-type: none"> 1. Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply and PROCHOT (Therm Ctrl) sensors. 2. VRD Hot asserted. 3. Minimum number of fans to cool the system not present or failed 4. Hard drive fault 5. Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present) 6. In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window1 |

| | | | |
|-------|----------|---------------------------|--|
| Amber | Solid on | Critical, non-recoverable | Fatal alarm – system has failed or shutdown: <ol style="list-style-type: none"> 1. CPU CATERR signal asserted 2. MSID mismatch detected (CATERR also asserts for this case). 3. CPU 1 is missing 4. CPU ThermalTrip 5. No power good – power fault 6. DIMM failure when there is only 1 DIMM present and hence no good memory present¹. 7. Runtime memory uncorrectable error in non-redundant mode¹. 8. DIMM Thermal Trip or equivalent 9. SSB Thermal Trip or equivalent 10. CPU ERR2 signal asserted 11. BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition) 12. Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition) 13. 240VA fault |
| Off | N/A | Not ready | AC power off |

10.3 BMC Boot/Reset Status LED Indicators

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

Table 50. BMC Boot/Reset Status LED Indicators

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

10.4 Post Code Diagnostic LEDs

A bank of eight POST code diagnostic LEDs are located on the back edge of the server next to the stacked USB connectors. During the system boot process, the BIOS executes a number of platform configuration processes, each of which is assigned a specific hex POST code number. As each configuration routine is started, the BIOS displays the given POST code to the POST code diagnostic LEDs. The purpose of these LEDs is to assist in troubleshooting a system hang

condition during the POST process. The diagnostic LEDs can be used to identify the last POST process to be executed. See Appendix D for a complete description of how these LEDs are read, and for a list of all supported POST codes.

10.5 5 Volt Stand-By Present LED

This LED is illuminated when a power cord (AC or DC) is connected to the server and the power supply is supplying 5 Volt Stand-by power to the server board. This LED is intended as a service caution indicator to anyone accessing the inside of the server system.

10.6 Fan Fault LEDs

The server board includes a Fan Fault LED next to each of the six system fans and both CPU fans. The LED has two states: On and Off. The BMC lights a fan fault LED if the associated fan-tach sensor has a lower critical threshold event status asserted. Fan-tach sensors are manual re-arm sensors. Once the lower critical threshold is crossed, the LED remains lit until the sensor is rearmed. These sensors are rearmed at system DC power-on and system reset.

10.7 Memory Fault LEDs

The server board includes a Memory Fault LED for each DIMM slot. When the BIOS detects a memory fault condition, it sends an *IPMI OEM* command (*Set Fault Indication*) to the BMC to instruct the BMC to turn on the associated Memory Slot Fault LED. These LEDs are only active when the system is in the 'on' state. The BMC will not activate or change the state of the LEDs unless instructed by the BIOS.

11. Environmental Limits Specification

The following table defines the Intel® Server Board S1400SP operating and non-operating environmental limits. Operation of the Intel® Server Board S1400SP 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 51. Server Board Design Specifications

| | |
|---------------------------|---|
| Operating Temperature | 0°C to 55°C ¹ (32°F to 131°F) |
| Non-Operating Temperature | -40°C to 70°C (-40°F to 158°F) |
| DC Voltage | ± 5% of all nominal voltages |
| Shock (Unpackaged) | Trapezoidal, 25g , 170 inches/sec |
| Shock (Packaged) | |
| < 20 pounds | 36 inches |
| 20 to < 40 pounds | 30 inches |
| 40 to < 80 pounds | 24 inches |
| 80 to < 100 pounds | 18 inches |
| 100 to < 120 pounds | 12 inches |
| 120 pounds | 9 inches |
| Vibration (Unpackaged) | 5 Hz to 500 Hz 3.13 g RMS random |

Note:

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

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 server board is designed to support the Intel® Xeon® Processor E5-2400 product family TDP guidelines up to and including 95W.

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

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

11.2 MTBF

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

- Telcordia* SR_332 Issue II: Reliability Prediction Procedure
- Method 1: Parts Count Prediction
- Case III: Generic Value + Quality + Stress + Temperature
- Confidence Level: 90%
- Quality Level: II
- Temperature: Customer Specified (default 40°C)
- Duty Cycle: Continuous, 100%
- Operating Environment: Ground Benign, Fixed, Controlled

Table 52. MTBF Estimate

| Assembly Name | Temperature (Degree C) | MTBF (hours) |
|-----------------------------|------------------------|--------------|
| Intel® Server Board S1400SP | 40 | 176414 |

12. Server Board Power Distribution

This section provides power supply design guidelines for a system using the Intel® Server Board S1400SP. The following diagram shows the power distribution implemented on this server board.

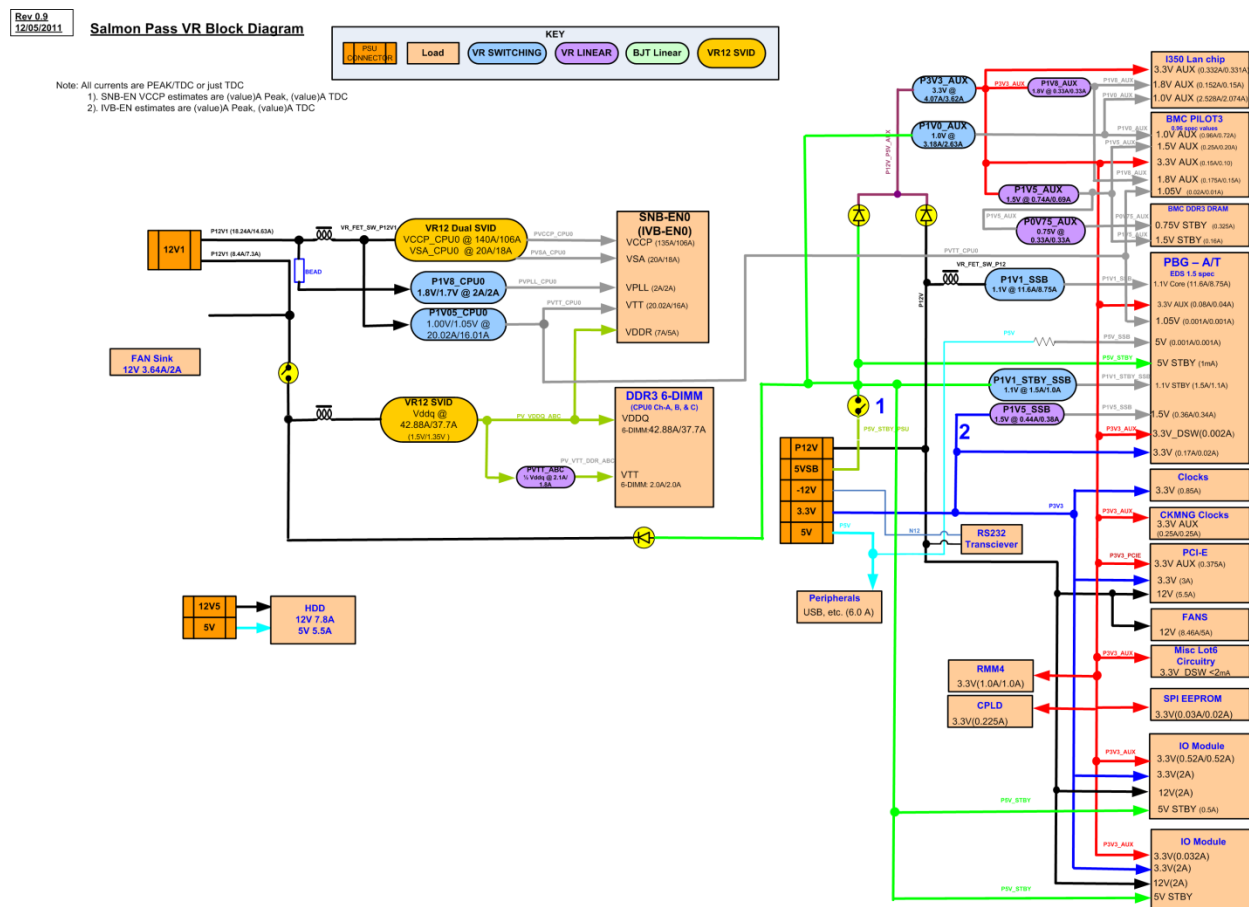


Figure 31. Power Distribution Block Diagram

12.1 Processor Power Support

The server boards support the Thermal Design Power (TDP) guideline for Intel® Xeon® processors. The Flexible Motherboard Guidelines (FMB) were also followed to determine the suggested thermal and current design values for anticipating future processor needs. The following table provides maximum values for I_{CC}, TDP power and T_{CASE} for the compatible Intel® Xeon® Processor E5-2400 and Intel® Xeon® processor E5-2400 v2 product families.

Table 53. Intel® Xeon® Processor TDP Guidelines

| TDP Power | Max Tcase | Icc Max |
|----------------|-----------|---------|
| 95 W | 78°C | 130 A |
| 80W | 75°C | 85A |
| 80W (1 socket) | 71°C | 80A |
| 70W | 70°C | 110A |
| 60W | 67°C | 90A |
| 50W | 65°C | 65A |

12.2 Power Supply Output Requirements

This section is for reference purposes only. The intent is to provide guidance to system designers to determine a power supply to use with these server boards. This section specifies the power supply requirements Intel® used to develop a mini-ERP 450W power supply for its 1U server system.

12.2.1 Output Power/Currents

The following table defines the output current ratings. Each output has a maximum and minimum current rating shown below. The power supply shall meet both static and dynamic voltage regulation requirements for the minimum dynamic load conditions. The power supply must meet only the static load voltage regulation requirements for the minimum static load conditions.

Table 54. Load Ratings

| | +12V | +5VSB |
|-------------|--------|-------|
| MAX | 36.26A | 3A |
| MIN DYNAMIC | 2A | 0.5A |
| MIN STATIC | 0.5A | 0.1A |

Note: *The combined output power of all outputs shall not exceed 405W @100~127VAC.

12.2.2 Standby Output

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

12.2.3 Voltage Regulation

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

Table 55. Voltage Regulation Limits

| PARAMETER | TOLERANCE | MIN | NOM | MAX | UNITS |
|-----------|-----------|--------|--------|--------|------------------|
| +12V | - 4%/+5% | +11.52 | +12.00 | +12.60 | V _{rms} |
| +5VSB | - 4%/+5% | +4.80 | +5.00 | +5.25 | V _{rms} |

12.2.4 Dynamic Loading

The output voltages shall remain within limits specified for the step loading and capacitive loading specified in the table below. The load transient repetition rate shall be tested between

50Hz and 5kHz at duty cycles ranging from 10%-90%. The load transient repetition rate is only a test specification. The Δ step load may occur anywhere within the MIN load to the MAX load conditions.

Table 56. Transient Load Requirements

| Output | Δ Step Load Size (See note 2) | Load Slew Rate | Test capacitive Load |
|--------|---|------------------|-----------------------------|
| 12V | 18A | 0.5 A/ μ sec | 2200 μ F ^{1,2} |
| +5VSB | 0.5A | 0.5 A/ μ sec | 20 μ F |

12.2.5 Capacitive Loading

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

Table 57. Capacitive Loading Conditions

| Output | MIN | MAX | Units |
|--------|-----|------|---------|
| +12V | 500 | 8000 | μ F |
| +5VSB | 20 | 350 | μ F |

12.2.6 Grounding

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

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

12.2.7 Closed loop stability

The power supply shall be unconditionally stable under all line/load/transient load conditions including capacitive load ranges. A minimum of: 45 degrees phase margin and -10dB-gain margin is required. The power supply manufacturer shall provide proof of the unit's closed-loop stability with local sensing through the submission of Bode plots. Closed-loop stability must be ensured at the maximum and minimum loads as applicable.

12.2.8 Residual Voltage Immunity in Standby mode

The power supply should be immune to any residual voltage placed on its outputs (Typically a leakage voltage through the system from standby output) up to 500mV. There shall be no additional heat generated, nor stressing of any internal components with this voltage applied to any individual or all outputs simultaneously. It also should not trip the protection circuits during turn on.

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

12.2.9 Common Mode Noise

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

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

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

12.2.10 Soft Starting

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

12.2.11 Zero Load Stability Requirements

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

12.2.12 Ripple/Noise

The maximum allowed ripple/noise output of the power supply is defined 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 58. Ripples and Noise

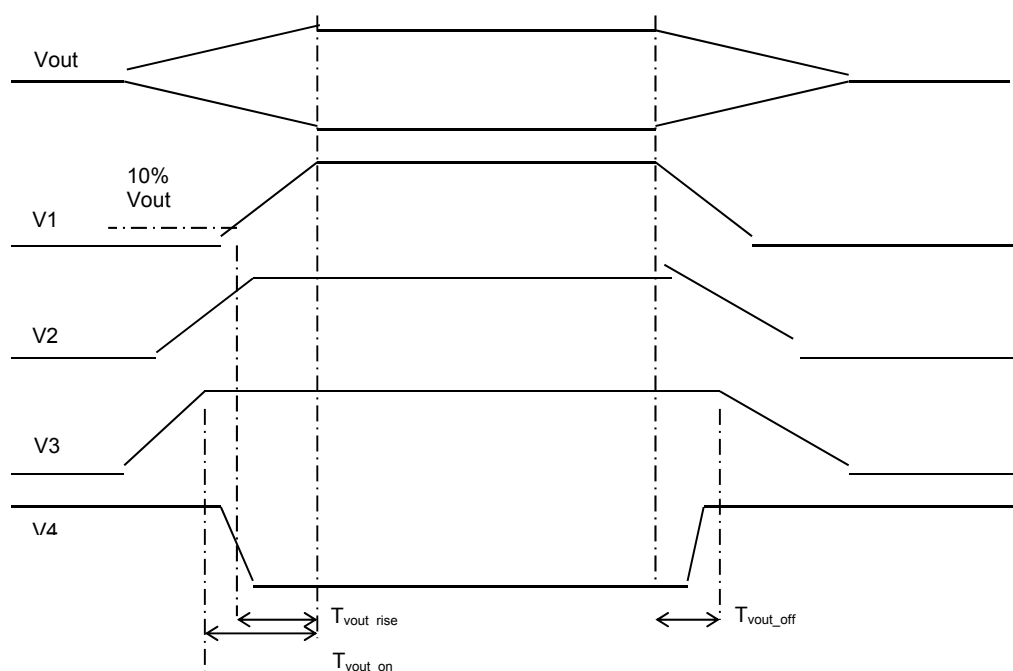
| | |
|----------|---------|
| +12V | +5VSB |
| 120mVp-p | 50mVp-p |

12.2.13 Timing Requirements

These are the timing requirements for the power supply operation. The output voltages must rise from 10% to within regulation limits ($T_{\text{vout_rise}}$) within 2 to 50ms, except for 5VSB - it is allowed to rise from 1 to 25ms. The +12V, output voltage should start to rise approximately at the same time. All outputs must rise monotonically. Each output voltage shall reach regulation within 50ms ($T_{\text{vout_on}}$) of each other during turn on of the power supply. Each output voltage shall fall out of regulation within 400ms ($T_{\text{vout_off}}$) of each other during turn off. 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 must be met for the cross loading condition.

Table 59. Output Voltage Timing

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

**Figure 32. Output Voltage Timing****Table 60. 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 |

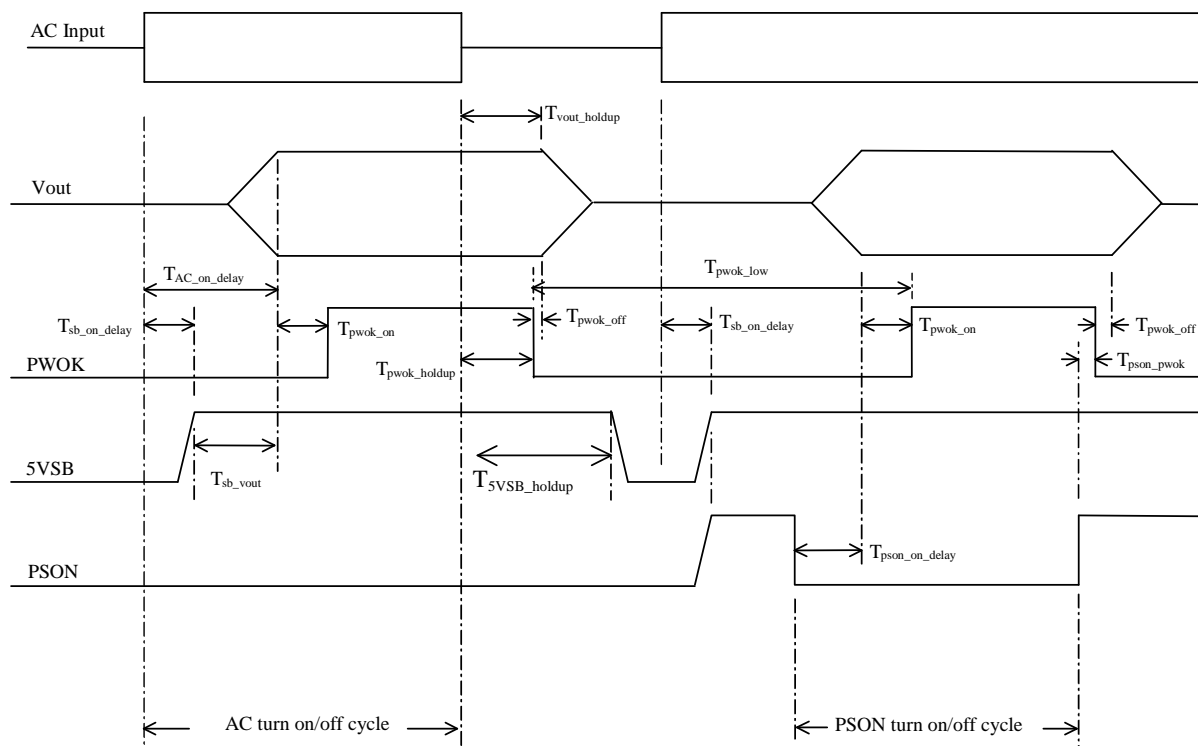


Figure 33. Turn On/Off Timing (Power Supply Signals)

12.2.14 Residual Voltage Immunity in Stand-by 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.

Appendix A: Integration and Usage Tips

- When adding or removing components or peripherals from the server board, you must remove AC power cord. With AC power plugged into the server board, 5-V standby is still present even though the server board is powered off.
- This server board supports Intel® Xeon® Processor E5-2400 product family or Intel® Xeon® Processor E5-2400 v2 product family with a Thermal Design Power (TDP) of up to and including 95 Watts. Previous generation Intel® Xeon® processors are not supported.
- The onboard SATA connectors 0 and 1 are designed to support the ODD or SSD, do NOT connect SATA connector 0 and 1 to the backplane. The onboard SATA connectors 2-5 are designed to support the backplane.
- On the back edge of the server board are EIGHT diagnostic LEDs that display a sequence of amber POST codes during the boot process. If the server board hangs during POST, the LEDs display the last POST event run before the hang.
- Only Registered DDR3 DIMMs (RDIMMs) and Unbuffered DDR3 DIMMs (UDIMMs) are supported on this server board. Mixing of RDIMMs and UDIMMs is not supported.
- The Intel® RMM4/RMM4 Lite connectors are not compatible with the previous Intel® Remote Management Modules
- Clear CMOS with the AC power cord plugged in. Removing AC power before performing the CMOS Clear operation causes the system to automatically power up and immediately power down after the CMOS Clear procedure is followed and AC power is re-applied. If this happens, remove the AC power cord, wait 30 seconds, and then re-connect the AC power cord. Power up the system and proceed to the <F2> BIOS Setup Utility to reset the desired settings.
- Normal BMC functionality is disabled with the Force BMC Update jumper set to the “enabled” position (pins 2-3). You should never run the server with the Force BMC Update jumper set in this position and should only use the jumper in this position when the standard firmware update process fails. This jumper must remain in the default (disabled) position (pins 1-2) when the server is running normally.
- This server board no longer supports the Rolling BIOS (two BIOS banks). It implements the BIOS Recovery mechanism instead.
- When performing a normal BIOS update procedure, you must set the BIOS Recovery jumper to its default position (pins 1-2).

Appendix B: Integrated BMC Sensor Tables

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

- **Sensor Type**

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

- **Event/Reading Type**

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

- **Event Offset/Triggers**

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

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

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

- **Assertion/De-assertion Enables**

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

- As: Assertions
- De: De-assertion

- **Readable Value/Offsets**

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

- **Event Data**

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

- R: Reading value
- T: Threshold value

- **Rearm Sensors**

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

- A: Auto-rearm
- M: Manual rearm

- **Default Hysteresis**

The hysteresis setting applies to all thresholds of the sensor. This column provides the count of hysteresis for the sensor, which is 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. You can access these sensors and/or generate events when the main (system) power is off but AC power is present.

Table 61. Integrated BMC Core Sensors

| Full Sensor Name (Sensor name in SDR) | Sensor # | Platform Applicability | Sensor Type | Event/Reading Type | Event Offset Triggers | Contrib. To System Status | Assert/De-assert | Readable Value/ Offsets | Event Data | Rearm | Standby |
|---|----------|------------------------|----------------|---------------------|---|---------------------------|------------------|-------------------------|-------------|-------|---------|
| 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 Redundancy1 (Pwr Unit Redund) | 02h | Chassis-specific | Power Unit 09h | Generic 0Bh | 00 - Fully Redundant | OK | As and De | – | Trig Offset | M | X |
| | | | | | 01 - Redundancy lost | Degraded | | | | | |
| | | | | | 02 - Redundancy degraded | Degraded | | | | | |
| | | | | | 03 - Non-redundant: sufficient resources. Transition from full redundant state. | Degraded | | | | | |
| | | | | | 04 – Non-redundant: sufficient resources. Transition from insufficient state. | Degraded | | | | | |
| | | | | | 05 - Non-redundant: insufficient resources | Fatal | | | | | |

| 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 | Standby |
|--|----------|---------------------------------------|-------------------------------|-------------------------|---|---------------------------|------------------|------------------------|-------------|-------|---------|
| | | | | | 06 – Redundant: degraded from fully redundant state. | Degraded | | | | | |
| | | | | | 07 – Redundant: Transition from non-redundant state. | Degraded | | | | | |
| IPMI Watchdog (IPMI Watchdog) | 03h | All | Watchdog 2 23h | Sensor Specific 6Fh | 00 - Timer expired, status only | OK | As | – | Trig Offset | A | X |
| | | | | | 01 - Hard reset | | | | | | |
| | | | | | 02 - Power down | | | | | | |
| | | | | | 03 - Power cycle | | | | | | |
| | | | | | 08 - Timer interrupt | | | | | | |
| Physical Security (Physical Scrtcy) | 04h | Chassis Intrusion is chassis-specific | Physical Security 05h | Sensor Specific 6Fh | 00 - Chassis intrusion | Degraded | As and De | – | Trig Offset | A | X |
| | | | | | 04 - LAN leash lost | OK | | | | | |
| 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 |

| 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 | Standby |
|--|----------|------------------------|------------------------|----------------------|---|---------------------------|------------------|-------------------------|-------------|-------|---------|
| 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 Redundancy1 (Fan Redundancy) | 0Ch | Chassis-specific | Fan 04h | Generic 0Bh | 00 - Fully redundant | OK | As and De | – | Trig Offset | A | – |
| | | | | | 01 - Redundancy lost | Degraded | | | | | |
| | | | | | 02 - Redundancy degraded | Degraded | | | | | |
| | | | | | 03 - Non-redundant: Sufficient resources. Transition from redundant | Degraded | | | | | |
| | | | | | 04 - Non-redundant: Sufficient resources. Transition from insufficient. | Degraded | | | | | |
| | | | | | 05 - Non-redundant: insufficient resources. | Non-Fatal | | | | | |

| 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 | Standby |
|---|----------|------------------------|-----------------------|-------------------------|--|---------------------------|------------------|------------------------|-------------|-------|---------|
| | | | | | 06 – Non-Redundant: degraded from fully redundant. | Degraded | | | | | |
| | | | | | 07 - Redundant degraded from non-redundant | Degraded | | | | | |
| SSB Thermal Trip (SSB Therm Trip) | 0Dh | All | Temperature 01h | Digital Discrete 03h | 01 – State Asserted | Fatal | As and De | – | Trig Offset | M | X |
| IO Module Presence (IO Mod Presence) | 0Eh | Platform-specific | Module/Board 15h | Digital Discrete 08h | 01 – Inserted/Present | OK | As and De | – | Trig Offset | M | - |
| SAS Module Presence (SAS Mod Presence) | 0Fh | Platform-specific | Module/Board 15h | Digital Discrete 08h | 01 – Inserted/Present | OK | As and De | – | Trig Offset | M | X |
| BMC Firmware Health (BMC FW Health) | 10h | All | Mgmt Health 28h | Sensor Specific 6Fh | 04 – Sensor Failure | Degraded | As | - | Trig Offset | A | X |
| System Airflow (System Airflow) | 11h | All | Other Units 0Bh | Threshold 01h | – | – | – | Analog | – | – | – |
| FW Update Status | 12h | All | Version Change 2Bh | OEM defined x70h | 00h→Update started 01h→Update completed successfully. 02h→Update failure | OK | As | – | Trig Offset | A | – |

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

| 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 | Standby |
|--|----------|------------------------|-----------------|--------------------|-----------------------|--------------------------------|------------------|------------------------|------------|-------|---------|
| Baseboard +1.05V Processor4 Vccp (BB +1.05Vccp P4) | 1Ah | Platform-specific | Voltage 02h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Baseboard Temperature 1 (Platform Specific) | 20h | Platform-specific | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |
| Front Panel Temperature (Front Panel Temp) | 21h | 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 |

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

| 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 | Standby |
|--|----------|-------------------------------|------------------|---------------------|--------------------------|---------------------------------|------------------|------------------------|------------------|-------|---------|
| 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-fatal2 | As and De | Analog | R, T | M | - |
| Fan Present Sensors (Fan x Present) | 40h–4Fh | Chassis and Platform Specific | Fan 04h | Generic 08h | 01 - Device inserted | OK | As and De | - | Triggered Offset | Auto | - |
| Power Supply 1 Status (PS1 Status) | 50h | Chassis-specific | Power Supply 08h | Sensor Specific 6Fh | 00 - Presence | OK | As and De | – | Trig Offset | A | X |
| | | | | | 01 - Failure | Degraded | | | | | |
| | | | | | 02 – Predictive Failure | Degraded | | | | | |
| | | | | | 03 - A/C lost | Degraded | | | | | |
| | | | | | 06 – Configuration error | OK | | | | | |
| Power Supply 2 Status (PS2 Status) | 51h | Chassis-specific | Power Supply 08h | Sensor Specific 6Fh | 00 - Presence | OK | As and De | – | Trig Offset | A | X |
| | | | | | 01 - Failure | Degraded | | | | | |
| | | | | | 02 – Predictive Failure | Degraded | | | | | |
| | | | | | 03 - A/C lost | Degraded | | | | | |
| | | | | | 06 – Configuration error | OK | | | | | |
| Power Supply 1 AC Power Input (PS1 Power In) | 54h | Chassis-specific | Other Units 0Bh | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | X |

| 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 | Standby |
|--|-----------|------------------------|-----------------|---------------------|--------------------------------|--------------------------------|------------------|------------------------|-------------|-------|---------|
| 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 | | | | | |

| 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 | Standby |
|---|-----------------|------------------------|---------------------|---------------------|----------------------------|---------------------------|------------------|------------------------|-------------|-------|---------|
| | 69h - 6Bh | Chassis-specific | Microcontroller 16h | Discrete 0Ah | 04- transition to Off Line | Degraded | | – | Trig Offset | | X |
| Processor 1 Status (P1 Status) | 70h | All | Processor 07h | Sensor Specific 6Fh | 01 - Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| Processor 2 Status (P2 Status) | 71h | All | Processor 07h | Sensor Specific 6Fh | 01 - Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| Processor 3 Status (P3 Status) | 72h | Platform-specific | Processor 07h | Sensor Specific 6Fh | 01 - Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| Processor 4 Status (P4 Status) | 73h | Platform-specific | Processor 07h | Sensor Specific 6Fh | 01 - Thermal trip | Fatal | As and De | – | Trig Offset | M | X |
| | | | | | 07 - Presence | OK | | | | | |
| Processor 1 Thermal Margin (P1 Therm Margin) | 74h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 2 Thermal Margin (P2 Therm Margin) | 75h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 3 Thermal Margin (P3 Therm Margin) | 76h | Platform-specific | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 4 Thermal Margin (P4 Therm Margin) | 77h | Platform-specific | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |

| 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 | Standby |
|--|----------|------------------------|-----------------|----------------------|-----------------------|--------------------------------|------------------|------------------------|-------------|-------|---------|
| Processor 1 Thermal Control % (P1 Therm Ctrl %) | 78h | All | Temperature 01h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | Trig Offset | A | – |
| Processor 2 Thermal Control % (P2 Therm Ctrl %) | 79h | All | Temperature 01h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | Trig Offset | A | – |
| Processor 3 Thermal Control % (P3 Therm Ctrl %) | 7Ah | Platform-specific | Temperature 01h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | Trig Offset | A | – |
| Processor 4 Thermal Control % (P4 Therm Ctrl %) | 7Bh | Platform-specific | Temperature 01h | Threshold 01h | [u] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | Trig Offset | A | – |
| Processor 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 | – |
| Processor 3 ERR2 Timeout (P3 ERR2) | 7Eh | Platform-specific | Processor 07h | Digital Discrete 03h | 01 – State Asserted | fatal | As and De | – | Trig Offset | A | – |
| Processor 4 ERR2 Timeout (P4 ERR2) | 7Fh | Platform-specific | Processor 07h | Digital Discrete 03h | 01 – State Asserted | fatal | As and De | – | Trig Offset | A | – |

| Full Sensor Name (Sensor name in SDR) | Sensor # | Platform Applicability | Sensor Type | Event/Reading Type | Event Offset Triggers | Contrib. To System Status | Assert/De-assert | Readable Value/Offsets | Event Data | Rearm | Standby |
|---|----------|------------------------|-----------------|----------------------|-----------------------|---------------------------|------------------|------------------------|-------------|-------|---------|
| Catastrophic Error (CATERR) | 80h | All | Processor 07h | Digital Discrete 03h | 01 – State Asserted | fatal | As and De | – | Trig Offset | M | – |
| Processor1 MSID Mismatch (P1 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 | – |
| Processor 1 DTS Thermal Margin (P1 DTS Therm Mgn) | 83h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 2 DTS Thermal Margin (P2 DTS Therm Mgn) | 84h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 3 DTS Thermal Margin (P3 DTS Therm Mgn) | 85h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor 4 DTS Thermal Margin (P4 DTS Therm Mgn) | 86h | All | Temperature 01h | Threshold 01h | - | - | - | Analog | R, T | A | – |
| Processor2 MSID Mismatch (P2 MSID Mismatch) | 87h | All | Processor 07h | Digital Discrete 03h | 01 – State Asserted | fatal | As and De | – | Trig Offset | M | – |

| Full Sensor Name (Sensor name in SDR) | Sensor # | Platform Applicability | Sensor Type | Event/Reading Type | Event Offset Triggers | Contrib. To System Status | Assert/De-assert | Readable Value/Offsets | Event Data | Rearm | Standby |
|--|----------|------------------------|-----------------|----------------------|-----------------------|---------------------------|------------------|------------------------|-------------|-------|---------|
| 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 | – |
| 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 3 VRD Temperature (P3 VRD Hot) | 92h | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Fatal | As and De | – | Trig Offset | M | – |
| Processor 4 VRD Temperature (P4 VRD Hot) | 93h | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | 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 | – |

| 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 | Standby |
|--|----------|------------------------|-----------------|----------------------------|-----------------------|---------------------------|------------------|------------------------|-------------|-------|---------|
| 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 | – |
| Processor 3 Memory VRD Hot 0-1 (P3 Mem01 VRD Hot) | 98h | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Non-fatal | As and De | – | Trig Offset | A | – |
| Processor 3 Memory VRD Hot 2-3 (P4 Mem23 VRD Hot) | 99h | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Non-fatal | As and De | – | Trig Offset | A | – |
| Processor 4 Memory VRD Hot 0-1 (P4 Mem01 VRD Hot) | 9Ah | All | Temperature 01h | Digital Discrete 05h | 01 - Limit exceeded | Non-fatal | As and De | – | Trig Offset | A | – |
| Processor 4 Memory VRD Hot 2-3 (P4 Mem23 VRD Hot) | 9Bh | 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 | - |

| 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 | Standby |
|---|----------|------------------------|-----------------|----------------------------|-----------------------|--------------------------------|------------------|------------------------|-------------|-------|---------|
| 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 1 (P1 DIMM Thrm Mrgn1) | B0h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 1 DIMM Aggregate Thermal Margin 2 (P1 DIMM Thrm Mrgn2) | B1h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 2 DIMM Aggregate Thermal Margin 1 (P2 DIMM Thrm Mrgn1) | B2h | All | Temperature 01h | Threshold 01h | [u,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 | Standby |
|---|----------|-------------------------------|-----------------|--------------------|-----------------------|---------------------------------|------------------|------------------------|------------|-------|---------|
| Processor 2 DIMM Aggregate Thermal Margin 2 (P2 DIMM Thrm Mrgn2) | B3h | All | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 3 DIMM Aggregate Thermal Margin 1 (P3 DIMM Thrm Mrgn1) | B4h | Platform Specific | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 3 DIMM Aggregate Thermal Margin 2 (P3 DIMM Thrm Mrgn2) | B5h | Platform Specific | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 4 DIMM Aggregate Thermal Margin 1 (P4 DIMM Thrm Mrgn1) | B6h | Platform Specific | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Processor 4 DIMM Aggregate Thermal Margin 2 (P4 DIMM Thrm Mrgn2) | B7h | Platform Specific | Temperature 01h | Threshold 01h | [u,l] [c,nc] | nc = Degraded c = Non-fatal | As and De | Analog | R, T | A | – |
| Fan Tachometer Sensors (Chassis specific sensor names) | BAh–BFh | Chassis and Platform Specific | Fan 04h | Threshold 01h | [l] [c,nc] | nc = Degraded c = Non-fatal2 | As and De | Analog | R, T | M | - |

| Full Sensor Name (Sensor name in SDR) | Sensor # | Platform Applicability | Sensor Type | Event/Reading Type | Event Offset Triggers | Contrib. To System Status | Assert/De-assert | Readable Value/Offsets | Event Data | Rearm | Standby |
|--|----------|------------------------|-----------------|----------------------|-------------------------------|---------------------------|------------------|------------------------|-------------|-------|---------|
| Processor 1 DIMM Thermal Trip (P1 Mem Thrm Trip) | C0h | All | Memory 0Ch | Digital Discrete 03h | 0A- Critical overtemperature | Fatal | As and De | – | Trig Offset | M | - |
| Processor 2 DIMM Thermal Trip (P2 Mem Thrm Trip) | C1h | All | Memory 0Ch | Digital Discrete 03h | 0A- Critical over temperature | Fatal | As and De | – | Trig Offset | M | - |
| Processor 3 DIMM Thermal Trip (P3 Mem Thrm Trip) | C2h | All | Memory 0Ch | Digital Discrete 03h | 0A- Critical overtemperature | Fatal | As and De | – | Trig Offset | M | X |
| Processor 4 DIMM Thermal Trip (P4 Mem Thrm Trip) | C3h | All | Memory 0Ch | Digital Discrete 03h | 0A- Critical overtemperature | 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 | – |

| 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 | Standby |
|---|----------|------------------------|-----------------|--------------------|-----------------------|---------------------------|------------------|------------------------|------------|-------|---------|
| 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 | – |

| 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 | Standby |
|---|----------|------------------------|-----------------|--------------------|-----------------------|--------------------------------|------------------|------------------------|------------|-------|---------|
| 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 | – |
| 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 Processor1 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 | – |

| 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 | Standby |
|--|----------|------------------------|-------------|--------------------|-----------------------|--------------------------------|------------------|------------------------|------------|-------|---------|
| Baseboard +1.05V Processor2 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 | – |

| 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 | Standby |
|---|----------|------------------------|-------------|--------------------|-----------------------|--------------------------------|------------------|------------------------|------------|-------|---------|
| 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 | – |
| 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 | – |

| 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 | Standby |
|---|-----------|------------------------|----------------|---------------------|--------------------------------|--------------------------------|------------------|------------------------|-------------|-------|---------|
| 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 C: POST Code Diagnostic LED Decoder

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

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

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

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



AF005227

Figure 34. POST Code Diagnostic LED Layout

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

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

Table 62. 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 63. POST Progress Codes

| Checkpoint | Diagnostic LED Decoder | | | | | | | | Description |
|---|-------------------------|----|----|----|--------------|----|----|-----|--|
| | 1 = LED On, 0 = LED Off | | | | | | | | |
| | Upper Nibble | | | | Lower Nibble | | | | |
| | MSB | | | | | | | LSB | |
| | 8h | 4h | 2h | 1h | 8h | 4h | 2h | 1h | |
| LED # | #7 | #6 | #5 | #4 | #3 | #2 | #1 | #0 | |
| 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 64 | | | | | | | | | |
| 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 |

| 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 | |
| 78h | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 0 | DXE ACPI Init |
| 79h | 0 | 1 | 1 | 1 | 1 | 0 | 0 | 1 | DXE CSM Init |
| 90h | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | DXE BDS Started |
| 91h | 1 | 0 | 0 | 1 | 0 | 0 | 0 | 1 | DXE BDS connect drivers |
| 92h | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 0 | DXE PCI Bus begin |
| 93h | 1 | 0 | 0 | 1 | 0 | 0 | 1 | 1 | DXE PCI Bus HPC Init |
| 94h | 1 | 0 | 0 | 1 | 0 | 1 | 0 | 0 | DXE PCI Bus enumeration |
| 95h | 1 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | DXE PCI Bus resource requested |
| 96h | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 0 | DXE PCI Bus assign resource |
| 97h | 1 | 0 | 0 | 1 | 0 | 1 | 1 | 1 | DXE CON_OUT connect |
| 98h | 1 | 0 | 0 | 1 | 1 | 0 | 0 | 0 | DXE CON_IN connect |
| 99h | 1 | 0 | 0 | 1 | 1 | 0 | 0 | 1 | DXE SIO Init |
| 9Ah | 1 | 0 | 0 | 1 | 1 | 0 | 1 | 0 | DXE USB start |
| 9Bh | 1 | 0 | 0 | 1 | 1 | 0 | 1 | 1 | DXE USB reset |
| 9Ch | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 0 | DXE USB detect |
| 9Dh | 1 | 0 | 0 | 1 | 1 | 1 | 0 | 1 | DXE USB enable |
| A1h | 1 | 0 | 1 | 0 | 0 | 0 | 0 | 1 | DXE IDE begin |
| A2h | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 0 | DXE IDE reset |
| A3h | 1 | 0 | 1 | 0 | 0 | 0 | 1 | 1 | DXE IDE detect |
| A4h | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 0 | DXE IDE enable |
| A5h | 1 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | DXE SCSI begin |
| A6h | 1 | 0 | 1 | 0 | 0 | 1 | 1 | 0 | DXE SCSI reset |
| A7h | 1 | 0 | 1 | 0 | 0 | 1 | 1 | 1 | DXE SCSI detect |
| A8h | 1 | 0 | 1 | 0 | 1 | 0 | 0 | 0 | DXE SCSI enable |
| A9h | 1 | 0 | 1 | 0 | 1 | 0 | 0 | 1 | DXE verifying SETUP password |
| ABh | 1 | 0 | 1 | 0 | 1 | 0 | 1 | 1 | DXE SETUP start |
| ACH | 1 | 0 | 1 | 0 | 1 | 1 | 0 | 0 | DXE SETUP input wait |
| ADh | 1 | 0 | 1 | 0 | 1 | 1 | 0 | 1 | DXE Ready to Boot |
| AEh | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 0 | DXE Legacy Boot |
| AFh | 1 | 0 | 1 | 0 | 1 | 1 | 1 | 1 | DXE Exit Boot Services |
| B0h | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 0 | RT Set Virtual Address Map Begin |
| B1h | 1 | 0 | 1 | 1 | 0 | 0 | 0 | 1 | RT Set Virtual Address Map End |
| B2h | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 0 | DXE Legacy Option ROM init |
| B3h | 1 | 0 | 1 | 1 | 0 | 0 | 1 | 1 | DXE Reset system |
| B4h | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 0 | DXE USB Hot plug |
| B5h | 1 | 0 | 1 | 1 | 0 | 1 | 0 | 1 | DXE PCI BUS Hot plug |
| B6h | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 0 | DXE NVRAM cleanup |
| B7h | 1 | 0 | 1 | 1 | 0 | 1 | 1 | 1 | DXE Configuration Reset |
| 00h | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | INT19 |
| S3 Resume | | | | | | | | | |
| E0h | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 0 | S3 Resume PEIM (S3 started) |
| E1h | 1 | 1 | 0 | 1 | 0 | 0 | 0 | 1 | S3 Resume PEIM (S3 boot script) |
| E2h | 1 | 1 | 0 | 1 | 0 | 0 | 1 | 0 | S3 Resume PEIM (S3 Video Repost) |
| E3h | 1 | 1 | 0 | 1 | 0 | 0 | 1 | 1 | S3 Resume PEIM (S3 OS wake) |
| BIOS Recovery | | | | | | | | | |
| 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 displayed to the Diagnostic LEDs that show the execution point in the MRC operational path at each step.

Table 64. MRC Progress Codes

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

Memory Initialization at the beginning of POST includes multiple functions, including: discovery, channel training, validation that the DIMM population is acceptable and functional, initialization of the IMC and other hardware settings, and initialization of applicable RAS configurations.

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

Table 65. POST Progress LED Codes

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

Appendix D: POST Code Errors

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

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

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

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

Note that for 0048 “Password check failed”, the system halts, and then after the next reset/reboot will display the error code on the Error Manager screen.

- **Fatal:** The system halts during post at a blank screen with the text “**Unrecoverable fatal error found. System will not boot until the error is resolved**” and “**Press <F2> to enter setup**” The POST Error Pause option setting in the BIOS setup does not have any effect with this class of error.

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

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

Table 66. POST Error Codes and Messages

| Error Code | Error Message | Response |
|------------|---|----------|
| 0012 | System RTC date/time not set | Major |
| 0048 | Password check failed | Major |
| 0140 | PCI component encountered a PERR error | Major |
| 0141 | PCI resource conflict | Major |
| 0146 | PCI out of resources error | Major |
| 0191 | Processor core/thread count mismatch detected | Fatal |
| 0192 | Processor cache size mismatch detected | Fatal |
| 0194 | Processor family mismatch detected | Fatal |
| 0195 | Processor Intel(R) QPI link frequencies unable to synchronize | Fatal |
| 0196 | Processor model mismatch detected | Fatal |
| 0197 | Processor frequencies unable to synchronize | Fatal |
| 5220 | BIOS Settings reset to default settings | Major |
| 5221 | Passwords cleared by jumper | Major |
| 5224 | Password clear jumper is Set | Major |
| 8130 | Processor 01 disabled | Major |
| 8131 | Processor 02 disabled | Major |
| 8132 | Processor 03 disabled | Major |
| 8133 | Processor 04 disabled | Major |
| 8160 | Processor 01 unable to apply microcode update | Major |
| 8161 | Processor 02 unable to apply microcode update | Major |
| 8162 | Processor 03 unable to apply microcode update | Major |
| 8163 | Processor 04 unable to apply microcode update | Major |
| 8170 | Processor 01 failed Self Test (BIST) | Major |
| 8171 | Processor 02 failed Self Test (BIST) | Major |
| 8172 | Processor 03 failed Self Test (BIST) | Major |
| 8173 | Processor 04 failed Self Test (BIST) | Major |
| 8180 | Processor 01 microcode update not found | Minor |
| 8181 | Processor 02 microcode update not found | Minor |
| 8182 | Processor 03 microcode update not found | Minor |
| 8183 | Processor 04 microcode update not found | Minor |
| 8190 | Watchdog timer failed on last boot | Major |
| 8198 | OS boot watchdog timer failure | Major |
| 8300 | Baseboard management controller failed self-test | Major |
| 8305 | Hot Swap Controller failure | Major |
| 83A0 | Management Engine (ME) failed Self Test | Major |
| 83A1 | Management Engine (ME) Failed to respond. | Major |
| 84F2 | Baseboard management controller failed to respond | Major |
| 84F3 | Baseboard management controller in update mode | Major |
| 84F4 | Sensor data record empty | Major |
| 84FF | System event log full | Minor |
| 8500 | Memory component could not be configured in the selected RAS mode | Major |
| 8501 | DIMM Population Error | Major |
| 8520 | DIMM_A1 failed test/initialization | Major |
| 8521 | DIMM_A2 failed test/initialization | Major |
| 8522 | DIMM_A3 failed test/initialization | Major |
| 8523 | DIMM_B1 failed test/initialization | Major |
| 8524 | DIMM_B2 failed test/initialization | Major |
| 8525 | DIMM_B3 failed test/initialization | Major |
| 8526 | DIMM_C1 failed test/initialization | Major |
| 8527 | DIMM_C2 failed test/initialization | Major |
| 8528 | DIMM_C3 failed test/initialization | Major |
| 8529 | DIMM_D1 failed test/initialization | Major |
| 852A | DIMM_D2 failed test/initialization | Major |
| 852B | DIMM_D3 failed test/initialization | Major |
| 852C | DIMM_E1 failed test/initialization | Major |

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

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

| Error Code | Error Message | Response |
|------------|--|----------|
| 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 |
| 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 67. POST Error Beep Codes

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

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

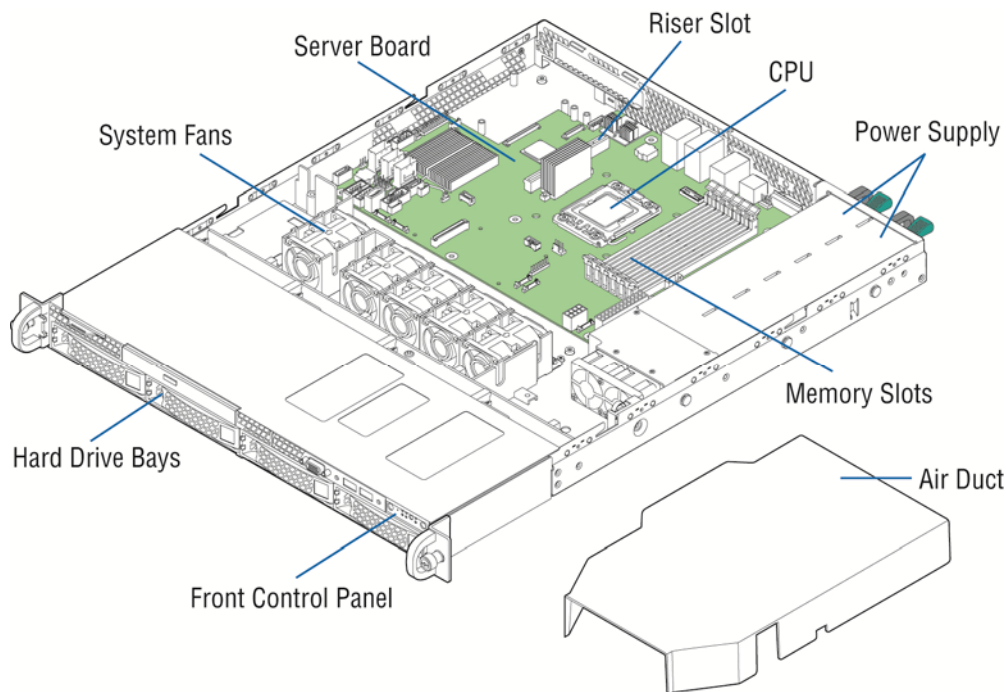
The Integrated BMC may generate beep codes upon detection of failure conditions. Beep codes are sounded each time the problem is discovered, such as on each power-up attempt, but are not sounded continuously. Codes that are common across all Intel® server boards and systems that use same generation chipset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

Table 68. Integrated BMC Beep Codes

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

Appendix E: Supported Intel® Server Chassis

The Intel® Server System R1000SP product family is comprised of several available 1U rack mount server systems that are all integrated with an Intel® Server Board S1400SP.



AF005707

| Server System | Integrated Server Board |
|---|-----------------------------|
| Intel® Server System R1000SP product family | Intel® Server Board S1400SP |

| Feature | Description |
|--------------------------|---|
| Processor Support | <ul style="list-style-type: none"> Support for one Intel® Xeon® E5-2400 processor or one Intel® Xeon® E5-2400 v2 processor in an FC-LGA 1356 Socket B2 package with Thermal Design Power of up to 95W. |
| Memory | <ul style="list-style-type: none"> Three memory channels, six memory DIMMs (Two memory DIMMs per channel). Support for 1066/1333 MT/s Unbuffered (UDIMM) LVDDR3 or DDR3 memory. Support for 1066/1333/1600 MT/s ECC Registered (RDIMM) DDR3 memory. Support for 1066/1333 MT/s ECC Registered (RDIMM) LVDDR3 memory. No support for mixing of RDIMMs and UDIMMs. No support for LRDIMMs. No support for Quad Rank DIMMs. |
| Chipset | Intel® C602 (-A) chipset with support for storage option upgrade keys |
| External I/O connections | <ul style="list-style-type: none"> Video (back and front video connectors) RJ-45 Serial-A Port Two RJ-45 Network Interface Connectors supporting 10/100/1000Mb for system with S1400SP2, four RJ-45 Network Interface Connectors supporting 10/100/1000Mb for system with S1400SP4 USB 2.0 connectors - 4 on back panel + 2 on front panel |

| Feature | Description |
|--|--|
| Internal I/O connectors/headers | <ul style="list-style-type: none"> One Type-A USB 2.0 connector One internal 2x5 pin serial port B header |
| I/O Module Accessory Options | <p>The following I/O modules utilize a single proprietary on-board connector. An installed I/O module can be supported in addition to standard on-board features and any add-in expansion cards.</p> <ul style="list-style-type: none"> Quad port 1 GbE based on Intel® Ethernet Controller I350 – AXX4P1GBPWL1OM Dual RJ-45 port 10GBase-T I/O Module based on Intel® Ethernet Controller x540 – AXX10GBTWL1OM Dual SFP+ port 10GbE module based on Intel® 82599 10 GbE controller – AXX10GBN1A1OM Single Port FDR 56GT/S speed InfiniBand* module with QSFP connector – AXX1FDRIB1OM Dual port FDR 56GT/S speed InfiniBand* module with QSFP connector – AXX2FDRIB1OM |
| System Fan Options | <p>Fixed system fan option:</p> <ul style="list-style-type: none"> Four 40mmx28mm single-rotor fans <p>Redundant system fan option:</p> <ul style="list-style-type: none"> Five 40mmx56mm dual-rotor fans |
| Riser Cards Options | <p>Support for one 1U PCI Express* riser card:</p> <ul style="list-style-type: none"> Single add-in card slot – PCI Express* Gen3 x8 electrical with x16 physical |
| Video | <ul style="list-style-type: none"> Integrated 2D Video Controller 16 MB DDR3 Memory |
| On-board storage controllers and options | <ul style="list-style-type: none"> Two AHCI SATA connectors at 6Gbps. One AHCI 4-port mini-SAS connector, supporting four SATA ports at 3Gbps. Two SCU 4-port mini-SAS connectors, supporting up to eight SAS/SATA ports total at 3Gbps with optional Intel® C600 RAID Upgrade Keys. One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices |
| Security | Intel® Trusted Platform Module (TPM) - AXXTPME5 (Accessory Option) |
| Server Management | <ul style="list-style-type: none"> Integrated Baseboard Management Controller, IPMI 2.0 compliant Support for Intel® Server Management Software Intel® Remote Management Module 4 Lite – Accessory option Intel® Remote Management Module 4 Management NIC – Accessory option |
| Power Supply Options | <p>Fixed power supply option:</p> <ul style="list-style-type: none"> AC 350W Silver <p>Redundant power supply option:</p> <ul style="list-style-type: none"> AC 450W Gold |
| Storage Bay Options | <ul style="list-style-type: none"> 4x 3.5" SATA fixed hard drive bays + Optical Drive support 4x 3.5" SATA/SAS hot swap hard drive bays + Optical Drive support 8x 2.5" SATA/SAS hot swap hard drive bays |
| Supported Rack Mount Kit Accessory Options | <ul style="list-style-type: none"> Value rack mount rail kit – Intel® Product Code – AXXVRAIL Tool-less rack mount rail kit – Intel® Product Code – AXXPRAIL Cable Management Arm – Intel® Product Code – AXX1U2UCMA (*supported with AXXPRAIL only) 2-post fixed mount bracket kit – Intel® Product Code – AXX2POSTBRCKT |

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 |
| AES | Advanced Encryption Standard |
| AMB | Advanced Memory Buffer (there is an AMB on each FBDIMM) |
| APIC | Advanced Programmable Interrupt Controller |
| ARP | Address Resolution Protocol |
| ASF | Alert Standards Forum |
| ASIC | Application specific integrated circuit |
| 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 |
| CBC | Chassis bridge controller. A microcontroller connected to one or more other CBCs. Together they bridge the IPMB buses of multiple chassis. |
| CLI | Command-line interface |
| CLTT | Closed-loop thermal throttling (memory throttling mode) |
| CMOS | In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory on the server board. |
| CSR | Control and status register |
| D-cache | Data cache. Processor-local cache dedicated for memory locations explicitly loaded and stored by running code. |
| DHCP | Dynamic Host Configuration Protocol |
| DIB | Device Information Block |
| DPC | Direct Platform Control |
| EEPROM | Electrically erasable programmable read-only memory |
| EMP | Emergency management port |
| EPS | External Product Specification |
| FML | Fast management link |
| FNI | Fast management link network interface |
| FRB | Fault resilient booting |
| FRU | Field replaceable unit |
| FSB | Front side bus |
| FTM | Firmware transfer mode |
| GPIO | General-purpose input/output |
| HSBP | Hot-swap backplane |
| HSC | Hot-swap controller |
| I-cache | Instruction cache. Processor-local cache dedicated for memory locations retrieved through instruction fetch operations. |

| Term | Definition |
|------------------|---|
| I ² C | Inter-integrated circuit bus |
| IA | Intel® architecture |
| IBF | Input buffer |
| ICH | I/O controller hub |
| IERR | Internal error |
| INIT | Initialization signal |
| IPMB | Intelligent Platform Management Bus |
| IPMI | Intelligent Platform Management Interface |
| ITP | In-target probe |
| KCS | Keyboard controller style |
| KT | Keyboard text |
| KVM | Keyboard, video, mouse |
| LAN | Local area network |
| LCD | Liquid crystal display |
| LPC | Low pin count |
| LUN | Logical unit number |
| MAC | Media Access Control |
| MD5 | Message Digest 5. A hashing algorithm that provides higher security than MD2. |
| MIB | Modular information block. A descriptive text translation of a PET event, contained in a MIB file for use by an SNMP agent when decoding SEL entries. |
| ms | Millisecond |
| MUX | Multiplexer |
| NIC | Network interface card |
| NMI | Non-maskable interrupt |
| OBF | Output buffer |
| OEM | Original equipment manufacturer |
| OLTT | Open-loop thermal throttling (memory throttling mode) |
| PCI | Peripheral Component Interconnect |
| PECI | Platform Environmental Control Interface |
| PEF | Platform event filtering |
| PET | Platform event trap |
| PIA | Platform information area |
| PLD | Programmable logic device |
| POST | Power-on self-test |
| PROM | Programmable read-only memory |
| PSMI | Power Supply Management Interface |
| PWM | Pulse Width Modulation. The mechanism used to control the speed of system fans. |
| RAM | Random Access Memory |
| RAS | Reliability, availability, and serviceability |
| RC4 | Rivest Cipher 4*. A stream cipher designed by Rivest* for RSA* data security, now RSA* security. It is a variable key-size stream cipher with byte-oriented operations. The algorithm is based on a random permutation. |
| RMCP+ | Remote Management Control Protocol |

| Term | Definition |
|--------|---|
| ROM | Read-only memory |
| RTC | Real-time clock |
| SCI | System Control Interrupt. A system interrupt used by hardware to notify the operating system of ACPI events. |
| SDR | Sensor data record |
| SDRAM | Synchronous dynamic random access memory |
| SEL | System event log |
| SHA1 | Secure Hash Algorithm 1 |
| SIO | Server Input/Output |
| SMBus* | A two-wire interface based on the I ² C protocol. The SMBus* is a low-speed bus that provides positive addressing for devices and bus arbitration. |
| SMI | Server management interrupt. SMI is the highest priority non-maskable interrupt. |
| SMM | Server management mode |
| SMS | Server management software |
| SNMP | Simple Network Management Protocol |
| SOL | Serial-over-LAN |
| SPT | Straight pass-through |
| SRAM | Static random access memory |
| UART | Universal asynchronous receiver and transmitter |
| UDP | User Datagram Protocol |
| UHCI | Universal Host Controller Interface |
| VLAN | Virtual local area network |

Reference Documents

See the following documents for additional information:

- *Advanced Configuration and Power Interface Specification, Revision 3.0*, <http://www.acpi.info/>.
- *Intelligent Platform Management Bus Communications Protocol Specification, Version 1.0*. 1998. Intel Corporation, Hewlett-Packard* Company, NEC* Corporation, Dell* Computer Corporation.
- *Intelligent Platform Management Interface Specification, Version 2.0*. 2004. Intel Corporation, Hewlett-Packard* Company, NEC* Corporation, Dell* Computer Corporation.
- *Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification, Version 1.1*, 02/01/02, Intel Corporation.
- *Intel® Remote Management Module User's Guide*, Intel Corporation.
- *Alert Standard Format (ASF) Specification, Version 2.0*, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., <http://www.dmtf.org>.
- *BIOS for PCSD Platforms Based on Intel® Xeon Processor E5-4600/2600/2400/1600 Product Families External Product Specification*
- *PCSD Platforms Based On Intel Xeon® Processor E5 4600/2600/2400/1600 Product Families BMC Core Firmware External Product Specification*

Mouser Electronics

Authorized Distributor

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

[Intel:](#)

[BBS1400SP2](#) [BBS1400SP4](#) [DBS1400SP2](#) [DBS1400SP4](#)