




Intel® Server Board S2600GZ/GL

Technical Product Specification



Revision 2.4

October 2014

Revision History

Date	Revision Number	Modifications
January 2012	1.0	First public release.
March 2012	1.1	<ul style="list-style-type: none"> ▪ Deleted section 6.9.2 Fan Profile. ▪ Deleted chapter 11 - Environmental Limits Specification. ▪ Added chapter 12 - BIOS Setup Utility. ▪ Added Figure 17. Intel® Server Board S2600GZ/GL PCI Layout.
August 2013	2.0	<ul style="list-style-type: none"> ▪ Added support for Intel® Xeon® processor E5-2600 v2 product family ▪ Updated memory support tables ▪ Corrected POST code table - E0h – E3h ▪ Updated reference documents table
February 2014	2.1	<ul style="list-style-type: none"> ▪ Corrected Table 39 - pin out for on-board 7-pin SATA connectors. ▪ Updated PCIe Gen 3 support verbiage – section 3.2.5 ▪ Added Phase Shedding support verbiage – section 3.2.4.1.1 ▪ Updated BIOS Setup options to include Phase Shedding option, PCIe Gen3 support option. ▪ Updated BIOS Setup options to include Extended ATR option, PFloor tuning option, Memory Mapped I/O Size option, PCIe AER Support option, Log Correctable Errors option and System Early POST Timeout option.
April 2014	2.2	<ul style="list-style-type: none"> ▪ Corrected Figure 9 and Figure 16 – PCIe Gen1x48GB/s ▪ Updated BIOS Setep options to include Phase Shedding and Memory SPD Override options.
June 2014	2.3	<ul style="list-style-type: none"> ▪ Updated Table 37. System Status LED State Definitions – Remove the “Battery Failure” from Description column. ▪ Updated Table 51. SystemStatus LED State Definitions – Remove the “Battery Failure” from Description column.
October 2014	2.4	<ul style="list-style-type: none"> ▪ Add note for Figure 29.

Disclaimers

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

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

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

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

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

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

Intel and Xeon are trademarks or registered trademarks of Intel Corporation.

*Other brands and names 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. Product Overview	2
2.1 Server Board Component/Feature Identification	4
2.2 Server Board Dimensional Mechanical Drawings	6
3. Product Architecture Overview	11
3.1 Processor Support	12
3.1.1 Processor Socket Assembly	12
3.1.2 Processor Population Rules	13
3.1.3 Processor Initialization Error Summary	14
3.1.4 Processor Thermal Design Power (TDP) Support	16
3.2 Processor Functions Overview	16
3.2.1 Processor Core Features:	17
3.2.2 Supported Technologies:	17
3.2.3 Intel® QuickPath Interconnect	17
3.2.4 Integrated Memory Controller (IMC) and Memory Subsystem	18
3.2.4.1 Supported Memory	19
3.2.4.2 Memory Slot Identification and Population Rules	21
3.2.4.3 Publishing System Memory	25
3.2.4.4 Integrated Memory Controller Operating Modes	25
3.2.4.5 Memory RAS Support	26
3.2.5 Processor Integrated I/O Module (IIO)	29
3.2.5.1 PCIe Interface	30
3.2.5.2 Riser Card Support	32
3.2.5.3 PCIe Add-in card support	34
3.2.5.4 Network Interface	38
3.2.5.5 I/O Module Support	38
3.2.5.6 Intel® Integrated RAID Option	39
3.3 Intel® C602 Chipset Functional Overview	40
3.3.1 Low Pin Count (LPC) Interface	41
3.3.2 Universal Serial Bus (USB) Controller	41
3.3.2.1 eUSB SSD Support	41
3.3.3 Embedded Serial ATA (SATA)/Serial Attached SCSI (SAS)/RAID Support	41
3.3.3.1 Intel® Embedded Server RAID Technology 2 (ESRT2)	43
3.3.3.2 Intel® Rapid Storage Technology (RSTe)	43
3.3.4 Manageability	44
3.4 Integrated Baseboard Management Controller Overview	44
3.4.1 Super I/O Controller	46
3.4.1.1 Keyboard and Mouse Support	46
3.4.1.2 Wake-up Control	46
3.4.2 Graphics Controller and Video Support	46

3.4.3	Baseboard Management Controller	47
3.4.3.1	Remote Keyboard, Video, Mouse, and Storage (KVMS) Support	47
3.4.3.2	Integrated BMC Embedded LAN Channel	48
4.	System Security.....	49
4.1	BIOS Password Protection	49
4.2	Intel® TPM module AXTPME5 Support	50
4.2.1	TPM security BIOS.....	50
4.2.2	Physical Presence.....	50
4.2.3	TPM Security Setup Options	51
4.2.3.1	Security Screen.....	51
4.3	Intel® Trusted Execution Technology.....	53
5.	Technology Support.....	54
5.1	Intel® Virtualization Technology – Intel® VT-x/VT-d/VT-c	54
5.2	Intel® Intelligent Power Node Manager	54
5.2.1	Hardware Requirements	56
6.	Platform Management Functional Overview.....	57
6.1	Baseboard Management Controller (BMC) Firmware Feature Support.....	57
6.1.1	IPMI 2.0 Features.....	57
6.1.2	Non IPMI Features	57
6.2	Advanced Configuration and Power Interface (ACPI)	59
6.3	Power Control Sources	59
6.4	BMC Watchdog	60
6.5	Fault Resilient Booting (FRB)	60
6.6	Sensor Monitoring	61
6.7	Field Replaceable Unit (FRU) Inventory Device	61
6.8	System Event Log (SEL)	61
6.9	System Fan Management	61
6.9.1	Thermal and Acoustic Management.....	62
6.9.2	Thermal Sensor Input to Fan Speed Control	62
6.9.3	Memory Thermal Throttling	63
6.10	Messaging Interfaces	64
6.10.1	User Model.....	64
6.10.2	IPMB Communication Interface.....	65
6.10.3	LAN Interface	65
6.10.3.1	RMCP/ASF Messaging	65
6.10.3.2	BMC LAN Channels	65
6.10.3.3	IPV6 Support.....	67
6.10.3.4	LAN Failover	67
6.10.3.5	BMC IP Address Configuration.....	67
6.10.3.6	DHCP BMC Hostname.....	69
6.10.4	Address Resolution Protocol (ARP).....	70
6.10.5	Internet Control Message Protocol (ICMP)	70
6.10.6	Virtual Local Area Network (VLAN)	70

6.10.7	Secure Shell (SSH)	70
6.10.8	Serial-over-LAN (SOL 2.0)	70
6.10.9	Platform Event Filter (PEF)	71
6.10.10	LAN Alerting	71
6.10.10.1	SNMP Platform Event Traps (PETs)	72
6.10.11	Alert Policy Table	72
6.10.11.1	E-mail Alerting	72
6.10.12	SM-CLP (SM-CLP Lite)	72
6.10.13	Embedded Web Server	72
6.10.14	Virtual Front Panel	74
6.10.15	Embedded Platform Debug	74
6.10.15.1	Output Data Format	75
6.10.15.2	Output Data Availability	75
6.10.15.3	Output Data Categories	75
6.10.16	Data Center Management Interface (DCMI)	76
6.10.17	Lightweight Directory Authentication Protocol (LDAP)	76
7.	Advanced Management Feature Support (RMM4)	77
7.1	Keyboard, Video, Mouse (KVM) Redirection	78
7.1.1	Remote Console	78
7.1.2	Performance	79
7.1.3	Security	79
7.1.4	Availability	79
7.1.5	Usage	79
7.1.6	Force-enter BIOS Setup	79
7.1.7	Availability	80
7.1.8	Network Port Usage	80
8.	On-board Connector/Header Overview	81
8.1	Power Connectors	81
8.1.1	Main Power	81
8.1.2	Riser Card Power Connectors	82
8.1.3	Hot Swap Backplane Power Connector	82
8.1.4	Peripheral Drive Power Connector	83
8.2	Front Panel Headers and Connectors	83
8.2.1	Front Panel Support	83
8.2.1.1	Power/Sleep Button and LED Support	84
8.2.1.2	System ID Button and LED Support	84
8.2.1.3	System Reset Button Support	84
8.2.1.4	NMI Button Support	84
8.2.1.5	NIC Activity LED Support	85
8.2.1.6	Hard Drive Activity LED Support	85
8.2.1.7	System Status LED Support	85
8.2.2	Front Panel USB Connector	87
8.2.3	Front Panel Video Connector	87
8.2.4	Intel® Local Control Panel Connector	87

8.3	On-Board Storage Connectors	87
8.3.1	Single Port SATA Only Connectors	88
8.3.2	Multiport Mini-SAS/SATA Connectors	88
8.3.3	Internal Type-A USB Connector	89
8.3.4	Internal 2mm Low Profile eUSB SSD Connector	89
8.4	Fan Connectors.....	90
8.5	Serial Port Connectors	90
8.6	Other Connectors and Headers.....	91
9.	Reset and Recovery Jumpers.....	92
10.	Light Guided Diagnostics	95
10.1	System ID LED.....	96
10.2	System Status LED	96
10.3	BMC Boot/Reset Status LED Indicators	98
10.4	Post Code Diagnostic LEDs	98
10.5	12 Volt Stand-By Present LED	98
10.6	Fan Fault LEDs	98
10.7	Memory Fault LEDs.....	98
10.8	CPU Fault LEDs.....	99
10.9	System Power Good LED.....	99
10.10	CATERR LED	99
11.	Power Supply Specification Guidelines.....	100
11.1	Power Supply DC Output Connector	100
11.2	Power Supply DC Output Specification	101
11.2.1	Output Power/Currents.....	101
11.2.2	Standby Output	101
11.2.3	Voltage Regulation	101
11.2.4	Dynamic Loading	101
11.2.5	Capacitive Loading.....	102
11.2.6	Grounding	102
11.2.7	Closed loop stability	102
11.2.8	Residual Voltage Immunity in Standby mode	102
11.2.9	Common Mode Noise.....	102
11.2.10	Soft Starting	102
11.2.11	Zero Load Stability Requirements	102
11.2.12	Hot Swap Requirements	102
11.2.13	Forced Load Sharing.....	102
11.2.14	Ripple/Noise.....	103
11.2.15	Timing Requirements	103
12.	BIOS Setup Utility.....	105
12.1	BIOS Setup Operation.....	105
12.1.1	Entering BIOS Setup.....	105
12.1.2	Setup Navigation Keyboard Commands.....	105

12.2	BIOS Setup Utility Screens.....	107
12.2.1	Main Screen (Tab)	107
12.2.2	Advanced Screen (Tab)	110
12.2.2.1	Processor Configuration.....	112
12.2.2.2	Power & Performance	122
12.2.2.3	Memory Configuration.....	123
12.2.2.4	Memory RAS and Performance Configuration.....	128
12.2.2.5	Mass Storage Controller Configuration.....	130
12.2.2.6	PCI Configuration.....	136
12.2.2.7	NIC Configuration.....	138
12.2.2.8	Serial Port Configuration	146
12.2.2.9	USB Configuration.....	148
12.2.2.10	System Acoustic and Performance Configuration.....	151
12.2.3	Security Screen (Tab)	154
12.2.4	Server Management Screen (Tab)	157
12.2.4.1	Console Redirection	164
12.2.4.2	System Information	166
12.2.4.3	BMC LAN Configuration	169
12.2.5	Boot Options Screen (Tab).....	177
12.2.5.1	CDROM Order	183
12.2.5.2	Hard Disk Order	184
12.2.5.3	Floppy Order	185
12.2.5.4	Network Device Order.....	185
12.2.5.5	BEV Device Order.....	186
12.2.5.6	Add EFI Boot Option	187
12.2.5.7	Delete EFI Boot Option	188
12.2.6	Boot Manager Screen (Tab).....	188
12.2.7	Error Manager Screen (Tab)	190
12.2.8	Save & Exit Screen (Tab).....	191
Appendix A: Integration and Usage Tips		195
Appendix B: Integrated BMC Sensor Tables.....		196
Appendix C: Management Engine Generated SEL Event Messages.....		210
Appendix D: POST Code Diagnostic LED Decoder		212
Appendix E: POST Code Errors.....		217
Appendix F: Supported Intel® Server Systems		223

List of Figures

Figure 1. Server Board Component/Features Identification	4
Figure 2. Intel® Light Guided Diagnostic LED Identification	5
Figure 3. Jumper Block Identification	6
Figure 4. Intel® Server Board S2600GZ/S2600GL External I/O Connector Layout	6
Figure 5. Intel® Server Board S2600GZ/GL – Primary Side Keepout Zone.....	7
Figure 6. Intel® Server Board S2600GZ/GL– Hole and Component Positions	8
Figure 7. Intel® Server Board S2600GZ/GL – Secondary Side Keepout Zone.....	9
Figure 8. Intel® Server Board S2600GZ/GL– Primary Side Height Restrictions	10
Figure 9. Intel® Server Boards S2600GZ, S2600GL Functional Block Diagram.....	11
Figure 10. Processor Socket Assembly.....	13
Figure 11. Processor Socket ILM Variations	13
Figure 12. Integrated Memory Controller Functional Block Diagram.....	18
Figure 13. Memory Slots Definition	22
Figure 14. Intel® Server Board S2600GZ Memory Slot Layout	23
Figure 15. Intel® Server Board S2600GL Memory Slot Layout	24
Figure 16. Functional Block Diagram of Processor IIO Sub-system	29
Figure 17. Intel® Server Board S2600GZ/GL PCI Layout	30
Figure 18. PCIe Port Bifurcation Options.....	31
Figure 19. 1U PCIe riser for Intel® Server Board S2600GZ/GL.....	32
Figure 20. 2U three PCIe slots riser for Intel® Server Board S2600GZ/GL	32
Figure 21. 2U two PCIe slots riser for Intel® Server Board S2600GZ/GL.....	33
Figure 22. 2U three PCIx/PCIe slots riser for Intel® Server Board S2600GZ/GL.....	33
Figure 23. Intel® Server Board S2600GZ/GL External RJ45 NIC Port LED Definition.....	38
Figure 24. Server Board Layout - I/O Module Connector.....	39
Figure 25. Server Board Layout – Intel® Integrated RAID Module Option Placement	39
Figure 26. Functional Block Diagram - Chipset Supported Features and Functions	40
Figure 27. Low Profile eUSB SSD Support	41
Figure 28. Intel® RAID C600 Upgrade Key Connector.....	42
Figure 29. Integrated BMC Block Diagram	45
Figure 30. Integrated BMC Functional Block Diagram.....	45
Figure 31. Setup Utility – TPM Configuration Screen	52
Figure 32. Fan Speed Control Process	63
Figure 33. Intel® RMM4 Lite Activation Key Installation	77
Figure 34. Intel® RMM4 Dedicated Management NIC Installation.....	77
Figure 35. Serial-A RJ45 connector pin-out.....	91
Figure 36. Serial A Configuration Jumper Block Location.....	91
Figure 37. Reset and Recovery Jumper Block Location	92
Figure 38. On-Board Diagnostic LED Placement	96
Figure 39. Memory Slot Fault LED Locations	96
Figure 40. Turn On/Off Timing (Power Supply Signals).....	104

Figure 41. Main Screen.....	107
Figure 42. Advanced Screen.....	110
Figure 43. Processor Configuration Screen.....	113
Figure 44. Power & Performance Screen.....	123
Figure 45. Memory Configuration Screen.....	124
Figure 46. Memory RAS and Performance Configuration Screen	129
Figure 47. Mass Storage Controller Configuration Screen	131
Figure 48. PCI Configuration Screen.....	136
Figure 49. NIC Configuration Screen	141
Figure 50. Serial Port Configuration Screen	147
Figure 51. USB Configuration Screen	149
Figure 52. System Acoustic and Performance Configuration.....	151
Figure 53. Security Screen.....	155
Figure 54. Server Management Screen	158
Figure 55. Console Redirection Screen.....	165
Figure 56. System Information Screen	167
Figure 57. BMC LAN Configuration Screen.....	170
Figure 58. Boot Options Screen	179
Figure 59. CDROM Order Screen	184
Figure 60. Hard Disk Order Screen.....	184
Figure 61. Floppy Order Screen.....	185
Figure 62. Network Device Order Screen.....	186
Figure 63. BEV Device Order Screen.....	186
Figure 64. Add EFI Boot Option Screen	187
Figure 65. Delete EFI Boot Option Screen	188
Figure 66. Boot Manager Screen	189
Figure 67. Error Manager Screen.....	190
Figure 68. Save & Exit Screen	191
Figure 69. POST Diagnostic LED Location	212
Figure 70. Intel® Server System R1000GZ/GL	223
Figure 71. Intel® Server System R2000GZ/GL	223

List of Tables

Table 1. Intel® Server Board S2600GZ/S2600GL Feature Set	2
Table 2. Supported Intel® Xeon® processor product family feature comparison table	12
Table 3. Mixed Processor Configurations Error Summary	15
Table 4. UDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family	19
Table 5. UDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family	19
Table 6. RDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family	20
Table 7. RDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family	20
Table 8. LRDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family	21
Table 9. LRDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family	21
Table 10. Intel® Server Board S2600GZ Memory Slot Identification	22
Table 11. Intel® Server Board S2600GL Memory Slot Nomenclature	23
Table 12. Riser Slot #1 – PCIe Port Routing	35
Table 13. Riser Slot #2 – PCIe Port Routing	36
Table 14. Supported Intel® I/O Module Options	39
Table 15. Supported Intel® Integrated RAID Modules	40
Table 16. Intel® RAID C600 Upgrade Key Options	43
Table 17. Video Modes	46
Table 18. BIOS Setup Options for Configuring Video	47
Table 19. TPM Setup Utility – Security Configuration Screen Fields	52
Table 20. Intel® Intelligent Power Node Manager	55
Table 21. ACPI Power States	59
Table 22. Power Control Initiators	59
Table 23. Messaging Interfaces	64
Table 24. Factory Configured PEF Table Entries	71
Table 25. Diagnostic Data	76
Table 26. Additional Diagnostics on Error.	76
Table 27. Intel® Remote Management Module 4 (RMM4) Options	77
Table 28. Enabling Advanced Management Features	77
Table 29. Main Power (Slot 1) Connector Pin-out ("MAIN PWR 1")	81
Table 30. Main Power (Slot 2) Connector Pin-out ("MAIN PWR 2")	82
Table 31. Riser Slot Power Pin-out ("OPT_12V_PWR_1" & "OPT_12V_PWR_2")	82
Table 32. Hot Swap Backplane Power Connector Pin-out ("HSBP PWR")	83
Table 33. Peripheral Drive Power Connector Pin-out ("ODD/SSD_PWR")	83
Table 34. SSI Front Panel Header Pin-out ("Front Panel")	84
Table 35. Power/Sleep LED Functional States	84
Table 36. NMI Signal Generation and Event Logging	85
Table 37. System Status LED State Definitions	85
Table 38. Front Panel USB Connector Pin-out ("FP USB")	87
Table 39. Front Panel Video Connector Pin-out ("FP VIDEO")	87
Table 40. Intel Local Control Panel Connector Pin-out ("LCP")	87

Table 41. Single Port AHCI SATA Controller Connector Pin-out ("SATA 0" & "SATA 1")	88
Table 42. Multiport SAS/SATA Connector Pin-out ("SCU_0 (0-3)")	88
Table 43. Multiport SAS/SATA Connector Pin-out ("SCU_1 (4-7)")	89
Table 44. Internal Type-A USB Connector Pin-out ("USB 2")	89
Table 45. Internal eUSB Connector Pin-out ("eUSB SSD")	89
Table 46. System Fan Connector Pin-out ("SYS_FAN #")	90
Table 47. Serial-B Connector Pin-out	90
Table 48. Serial A Connector Pin-out	91
Table 49. Chassis Intrusion Header Pin-out ("CHAS_INTR")	91
Table 50. Hard Drive Activity Header Pin-out ("HDD_LED")	91
Table 51. SystemStatus LED State Definitions	96
Table 52. BMC Boot/Reset Status LED Indicators	98
Table 53. Power Supply DC Power Output Connector Pinout	100
Table 54. Minimum Load Ratings	101
Table 55. Voltage Regulation Limits	101
Table 56. Transient Load Requirements	101
Table 57. Capacitive Loading Conditions	102
Table 58. Ripples and Noise	103
Table 59. Timing Requirements	103
Table 60. BIOS Setup: Keyboard Command Bar	106
Table 61. BMC Core Sensors	198
Table 62. Server Platform Services Firmware Health Event	210
Table 63. Node Manager Health Event	211
Table 64. POST Progress Code LED Example	212
Table 65. POST Progress Codes	213
Table 66. MRC Progress Codes	215
Table 67. MRC Fatal Error Codes	215
Table 68. POST Error Codes and Messages	217
Table 69. POST Error Beep Codes	221
Table 70. Integrated BMC Beep Codes	222
Table 71. Intel® Server System R1000GZ/GL Product Family Feature Set	224
Table 72. Intel® Server System R2000GZ/GL Product Family Feature Set	226

<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 Boards S2600GZ and S2600GL.

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

1.1 Chapter Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Product Overview
- Chapter 3 – Product Architecture Overview
- 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 – Reset and Recovery Jumpers
- Chapter 10 – Light-Guided Diagnostics
- Chapter 11 – Power Supply Specification Guidelines
- Chapter 12 – BIOS Setup Utility
- Appendix A – Integration and Usage Tips
- Appendix B – Integrated BMC Sensor Tables
- Appendix C – Management Engine Generated SEL Event Messages
- Appendix D – POST Code Diagnostic LED Decoder
- Appendix E – POST Code Errors
- Appendix F – Supported Intel® Server Systems

1.2 Server Board Use Disclaimer

Intel Corporation server boards support add-in peripherals and 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.

2. Product Overview

The Intel® Server Boards S2600GZ and S2600GL are monolithic printed circuit board assemblies with features that are intended for high density 1U and 2U rack mount servers. These server boards are designed to support both the Intel® Xeon® processor E5-2600 product family and Intel® Xeon® processor E5-2600 v2 product family. Previous generation Intel® Xeon® processors are not supported. Many of the features and functions of these two server boards are common. A board will be identified by name when a described feature or function is unique to it.

Table 1. Intel® Server Board S2600GZ/S2600GL Feature Set

Feature	Description
Processor Support	<ul style="list-style-type: none"> Two LGA2011 (Socket R) processor sockets Support for one or two processors: <ul style="list-style-type: none"> Intel® Xeon® processor E5-2600 product family with TDP support up to 135 W Intel® Xeon® processor E5-2600 v2 product family with TDP support up to 130 W
Memory	<ul style="list-style-type: none"> S2600GL - 16 DIMM slots – 2 DIMMs/Channel – 4 memory channels per processor S2600GZ - 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor Unbuffered DDR3 (UDIMM), registered DDR3 (RDIMM), Load Reduced DDR3 (LRDIMM) Memory DDR3 data transfer rates of 800, 1066, 1333, 1600, 1866¹ MT/s DDR3 standard I/O voltage of 1.5V and DDR3 Low Voltage of 1.35V
Chipset	Intel® C602 chipset with support for optional Storage Option Select keys
External (Back Panel) I/O connections	<ul style="list-style-type: none"> DB-15 Video connector RJ-45 Serial Port A connector Four RJ-45 Network Interface Connectors supporting 10/100/1000Mb Three USB 2.0 connectors
Internal I/O connectors/headers	<ul style="list-style-type: none"> One Type-A USB 2.0 connector One 2x5 pin connector providing front panel support for two USB 2.0 ports One 2x15 pin SSI-EEB compliant front panel header One 2x7pin Front Panel Video connector One 1x7pin header for optional Intel® Local Control Panel support One DH-10 Serial Port B connector
I/O Module Options	<p>The following I/O modules utilize a 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"> AXX4P1GBPWLIOM – Quad port 1 GbE based on Intel® Ethernet Controller I350 AXX10GBTWLIOM – Dual RJ-45 port 10GBase-T I/O Module based on Intel® Ethernet Controller x540 AXX10GBNIAIOM – Dual SFP+ port 10GbE module based on Intel® 82599 10 GbE controller AXX1FDRIBIOM – Single Port FDR 56GT/S speed InfiniBand module with QSFP connector AXX2FDRIBIOM – Dual port FDR 56GT/S speed infiniband module with QSFP connector
System Fans	<ul style="list-style-type: none"> Six 10-pin managed system fan headers
Riser Card Support	<p>Two riser card slots.</p> <ul style="list-style-type: none"> Each riser card slot has a total of 24PCIe lanes routed to them Each riser card slot has support for various 1U and 2U riser cards
Video	<ul style="list-style-type: none"> Integrated 2D Video Controller 16 MB DDR3 Memory

¹ Intel® Xeon® processor E5-2600 v2 product family only

Feature	Description
On-board storage controllers and options	<ul style="list-style-type: none"> • One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices • Two 7-pin single port AHCI SATA connectors capable of supporting up to 6 GB/sec • Two SCU 4-port mini-SAS connectors capable of supporting up to 3 GB/sec SAS/SATA <ul style="list-style-type: none"> ○ SCU 0 Port (Enabled standard) ○ SCU 1 Port (Requires Intel RAID C600 Upgrade Key) • Intel® RAID C600 Upgrade Key support providing optional expanded SCU SATA / SAS RAID capabilities • Intel® Integrated RAID module support (Optional)
Security	<ul style="list-style-type: none"> • Intel® TPM module - 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 support (Accessory Option) • Intel® Remote Management Module 4 Lite support (Accessory Option)

2.1 Server Board Component/Feature Identification

The following illustration provides a general overview of the server board, identifying key feature and component locations.

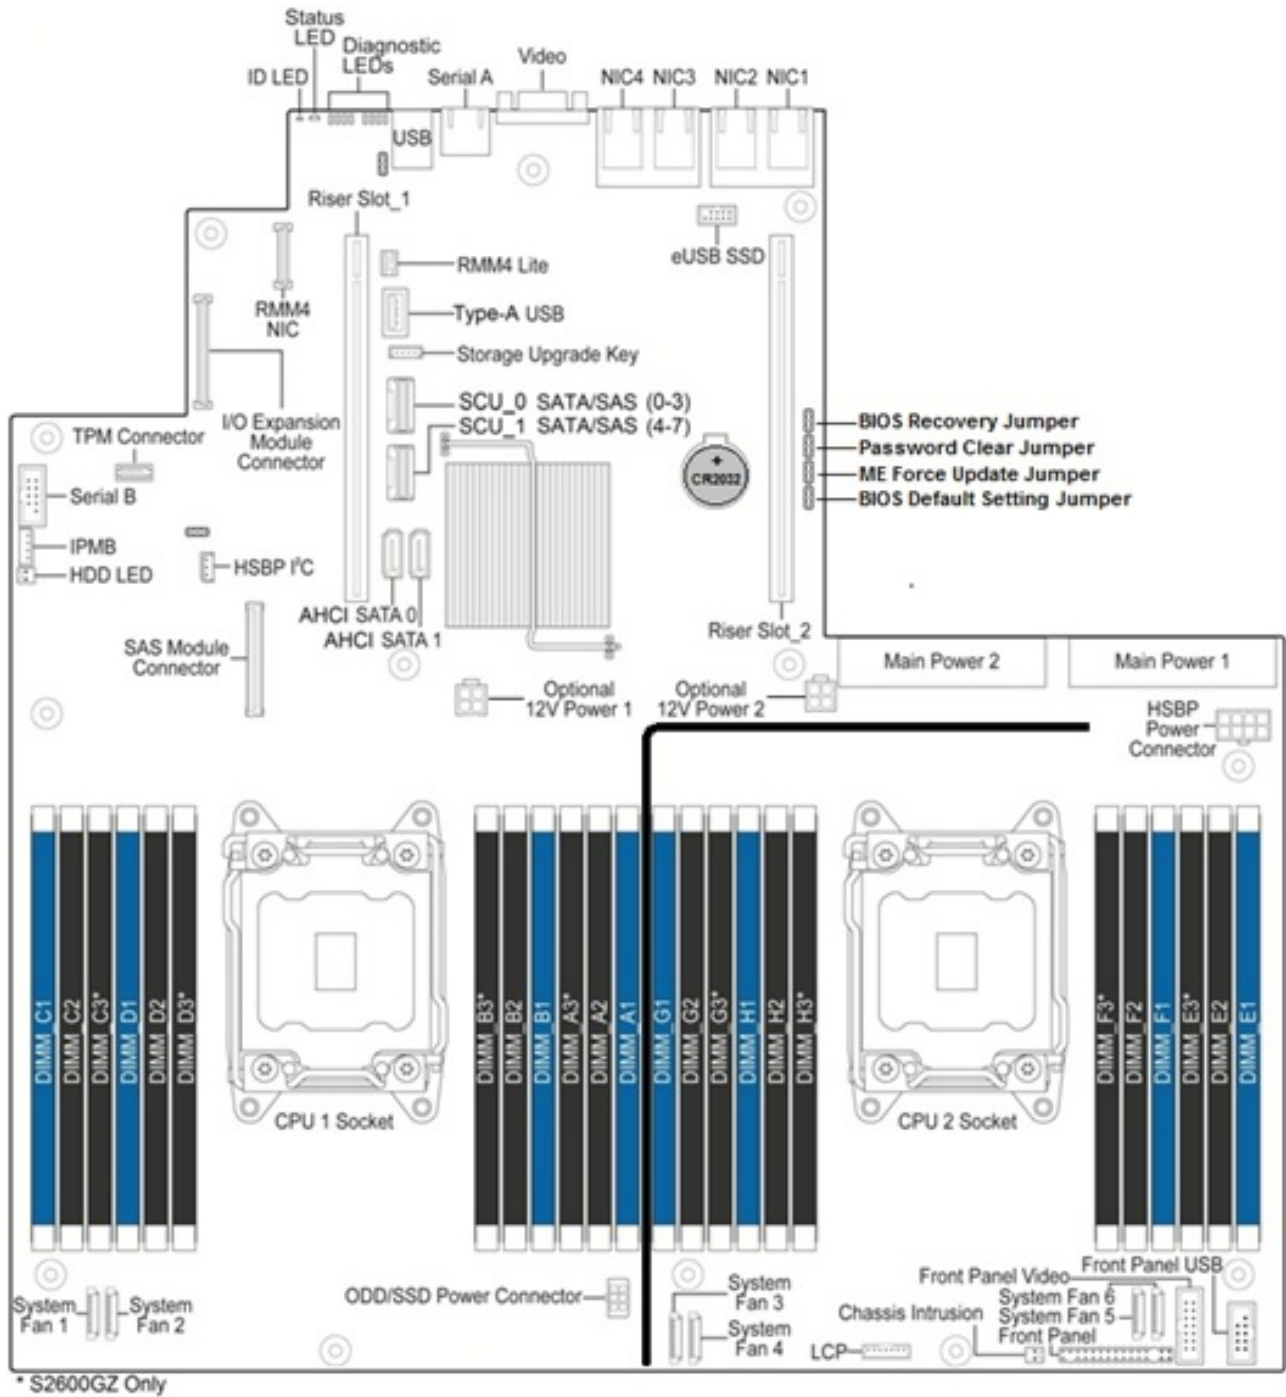



Figure 1. Server Board Component/Features Identification



AF003994

Label	Description	Label	Description
A	System ID	I	System Fan #3 Fan Fault
B	System Status	J	Memory Fault
C	POST Code Diagnostics	K	System Fan #2 Fan Fault
D	12V Stand-by Power Present	L	System Fan #1 Fan Fault
E	CPU-2 Fault	M	CPU-1 Fault
F	System Fan #6 Fan Fault	N	CATERR
G	System Fan #5 Fan Fault	O	System Power Good
H	System Fan #4 Fan Fault		

Figure 2. Intel® Light Guided Diagnostic LED Identification

See Light Guided Diagnostics for additional details.

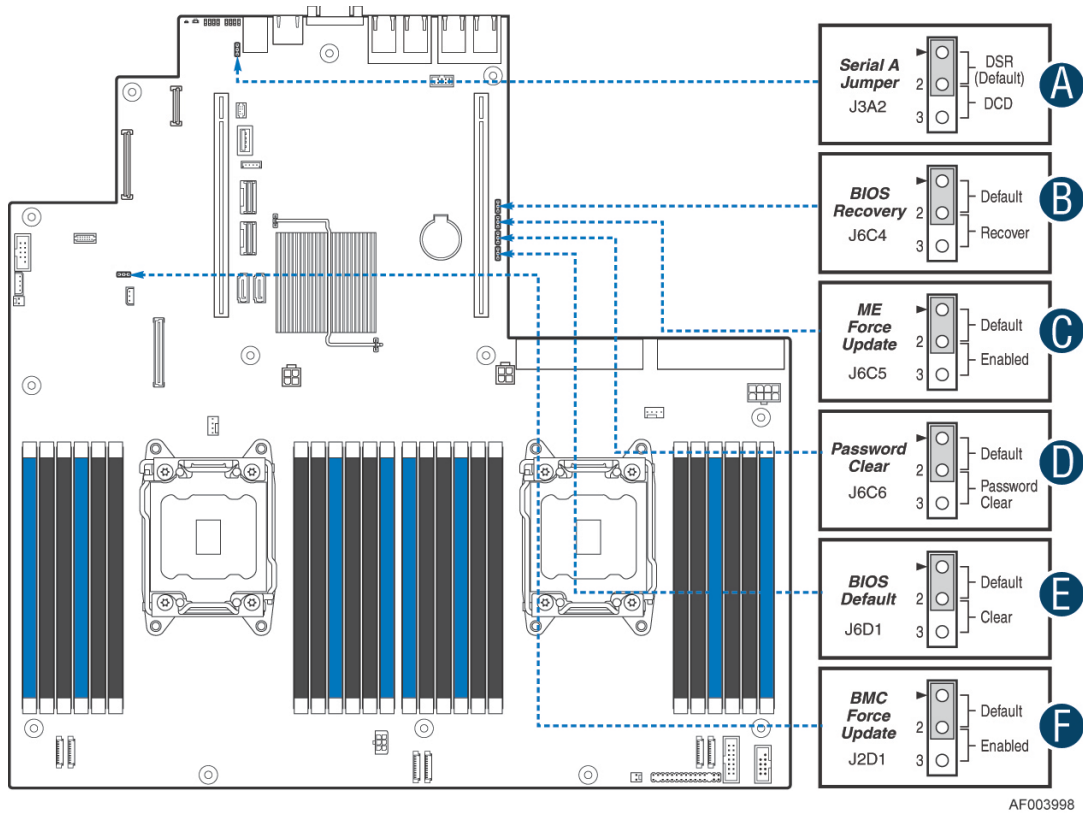


Figure 3. Jumper Block Identification

See Chapter 9 - Reset & Recovery Jumpers for additional details.

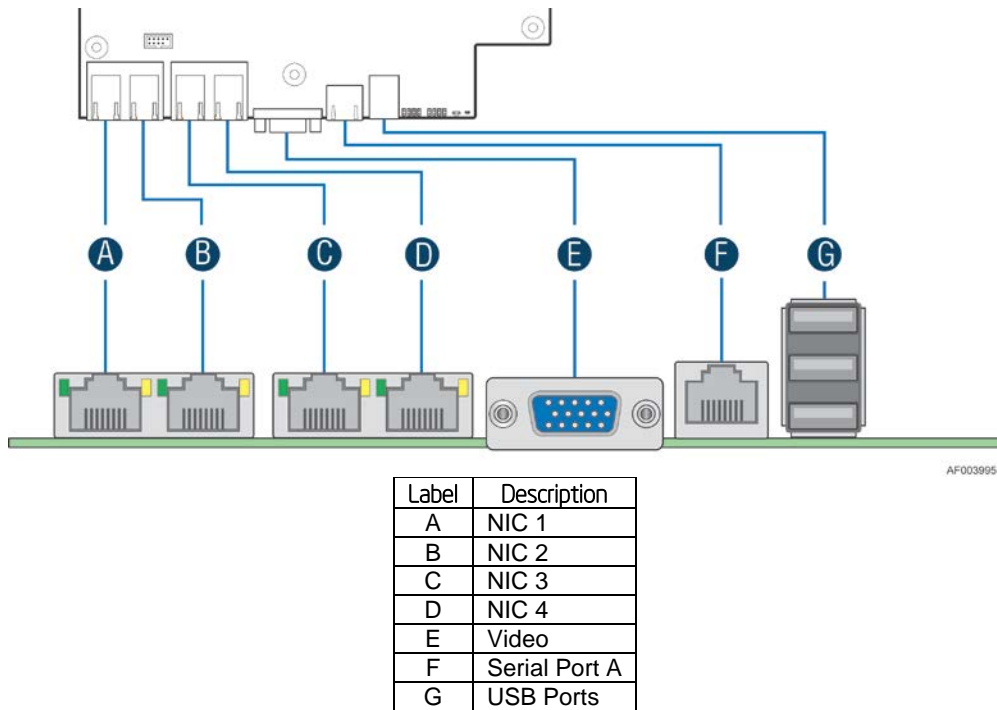


Figure 4. Intel® Server Board S2600GZ/S2600GL External I/O Connector Layout

2.2 Server Board Dimensional Mechanical Drawings




Figure 5. Intel® Server Board S2600GZ/GL – Primary Side Keepout Zone




Figure 6. Intel® Server Board S2600GZ/GL– Hole and Component Positions




Figure 7. Intel® Server Board S2600GZ/GL – Secondary Side Keepout Zone




Figure 8. Intel® Server Board S2600GZ/GL– Primary Side Height Restrictions

3. Product Architecture Overview

The architecture of Intel® Server Boards S2600GZ and S2600GL is developed around the integrated features and functions of the Intel® Xeon® processor E5-2600 product family, the Intel® C602 chipset, the Intel® Ethernet Controller I350 GbE controller chip, and the Emulex* Pilot-III Server Management Controller.

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

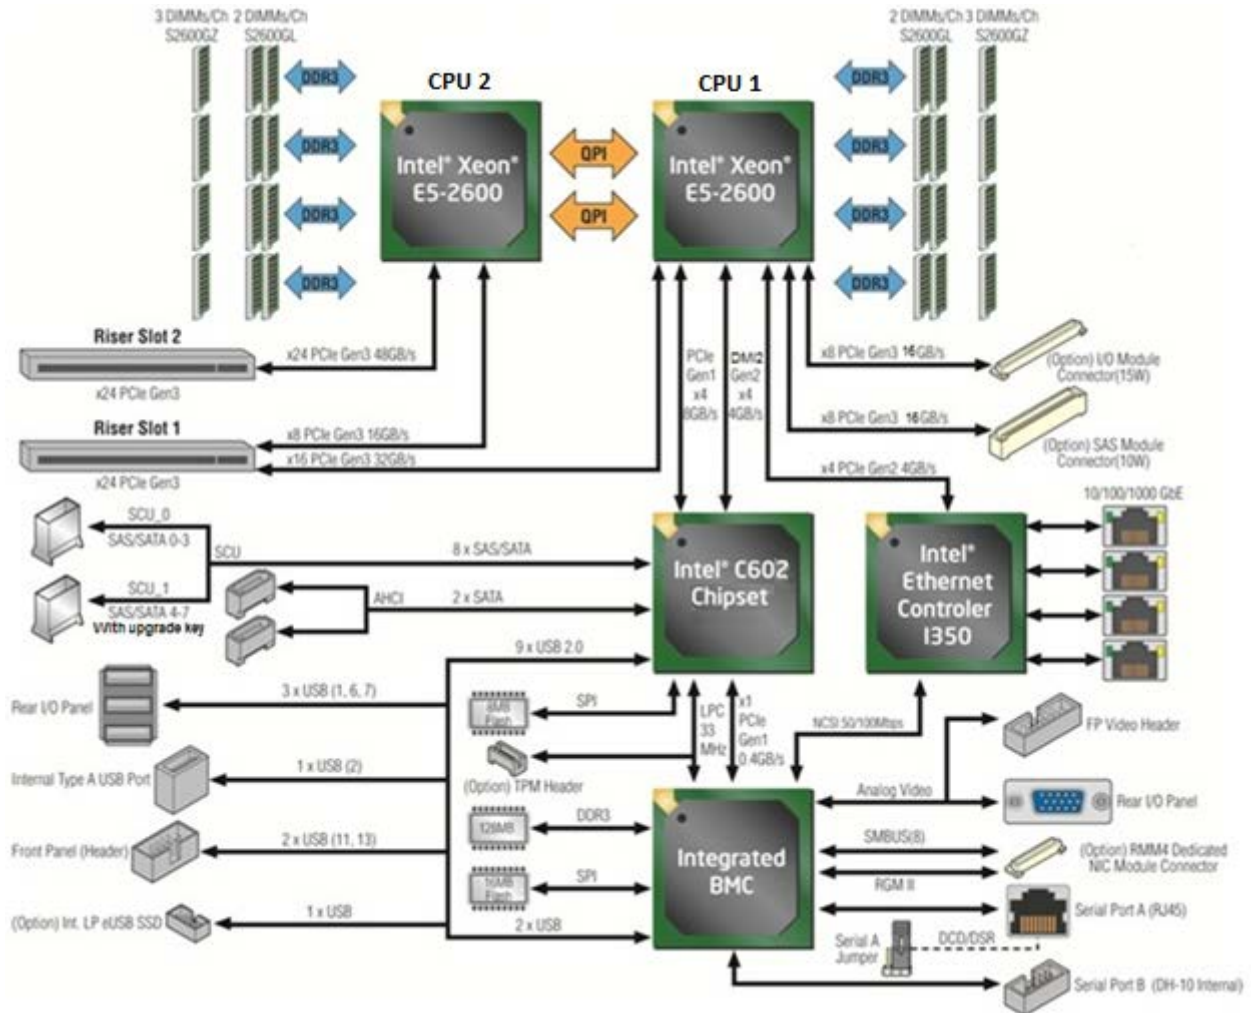


Figure 9. Intel® Server Boards S2600GZ, S2600GL Functional Block Diagram

3.1 Processor Support

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

- Intel® Xeon® processor E5-2600 product family, with a Thermal Design Power (TDP) of up to 135W.
- Intel® Xeon® processor E5-2600 v2 product family, with a Thermal Design Power (TDP) of up to 130W.

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

Table 2. Supported Intel® Xeon® processor product family feature comparison table

Feature	Intel® Xeon® E5-2600 processor product family	Intel® Xeon® E5-2600 processor v2 product family
QPI Speed (GT/s)	8.0, 7.2, and 6.4	
Addressability	46 bits physical, 48 bits virtual	
Cores	Up to 8	Up to 12
Threads Per Socket	Up to 16 threads	Up to 24 threads
Last-Level Cache (LLC)	Up to 20 MB	Up to 30 MB
Intel® Turbo-Boost Technology	YES	
Memory Population	4 Channels of up to 3 RDIMMs, 3 LRDIMMs, or 2 UDIMMs	
Max Memory Speed	Up to 1600	Up to 1866
Memory RAS	ECC, Patrol Scrubbing, Demand Scrubbing, Sparing, Mirroring, Lockstep Mode, x4/x8 SDDC	
PCIe* Lanes / Controllers / Speed (GT/s)	40 / 10 (PCIe* 3.0 at 8 GT/s)	
Idle Power Targets (W)	15 W or higher, 12 W for LV SKUs	10.5 W or higher, 7.5 W for LV SKUs

Visit <http://www.intel.com/support> for a complete list of supported processors.

3.1.1 Processor Socket Assembly

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

The illustration below identifies each sub-assembly component:




Figure 10. Processor Socket Assembly

There are two variations of the ILM: square and narrow.

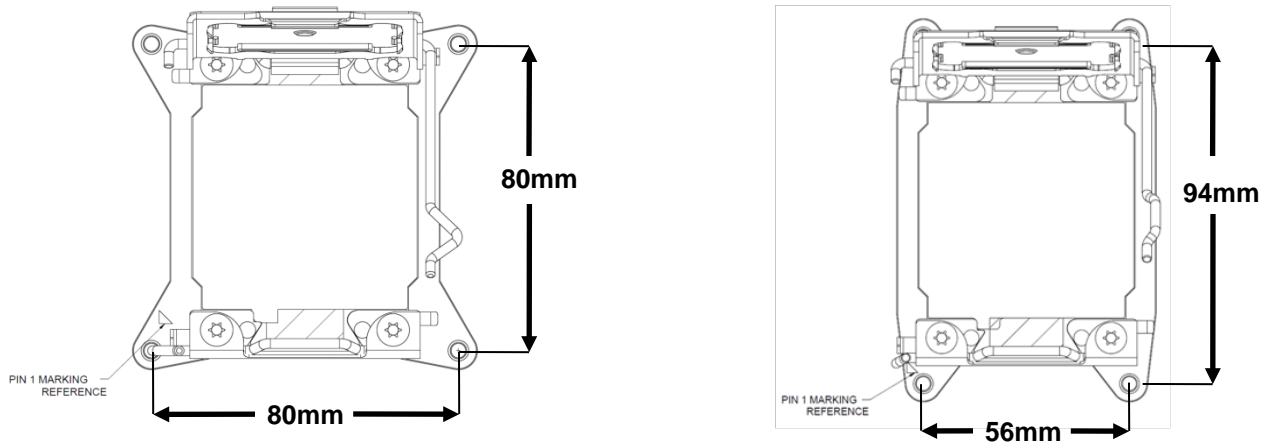


Figure 11. Processor Socket ILM Variations

- The square ILM has an 80x80mm heat sink mounting hole pattern and is used on the Intel® Server Board S2600GL.
- The narrow ILM has a 56x94mm heat sink mounting hole pattern and is used on the Intel® Server Board S2600GZ.

Note: Processor heat sink solutions for the Intel® server boards S2600GL and S2600GZ are NOT the same. Care must be taken when selecting heat sinks for the given server board ensuring the screw layout pattern of the heat sink matches the screw hole pattern of the ILM.

3.1.2 Processor Population Rules

Note: Although the server board does support dual-processor configurations consisting of different processors that meet the defined criteria below, Intel does not perform validation testing of this configuration. For optimal system performance in dual-processor configurations, Intel recommends that identical processors be installed.

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

When two processors are installed, the following population rules apply:

- Both processors must be of the same processor family.
- Both processors must have the same number of cores
- Both processors must have the same cache sizes for all levels of processor cache memory.

Processors with different core frequencies can be mixed in a system, given the prior rules are met. If this condition is detected, all processor core frequencies are set to the lowest common denominator (highest common speed) and an error is reported.

Processors which have different Intel® Quickpath (QPI) Link Frequencies may operate together if they are otherwise compatible and if a common link frequency can be selected. The common link frequency would be the highest link frequency that all installed processors can achieve.

Processor stepping within a common processor family can be mixed as long as it is listed in the processor specification updates published by Intel Corporation.

3.1.3 Processor Initialization Error Summary

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

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

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

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

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

Major: If the “Post Error Pause” setup option is enabled, the system goes directly to the Error Manager to display the error, and logs the POST Error Code to SEL. Operator intervention is required to continue booting the system.

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

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

Table 3. Mixed Processor Configurations Error Summary

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

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

3.1.4 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 E5-2600 product family TDP guidelines up to and including 135W.

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

3.2 Processor Functions Overview

With the release of the Intel® Xeon® processor E5-2600 product family, several key system components, including the CPU, Integrated Memory Controller (IMC), and Integrated IO Module (IIO), have been combined into a single processor package and feature per socket; two Intel® QuickPath Interconnect point-to-point links capable of up to 8.0 GT/s, up to 40 lanes of Gen 3 PCI Express* links capable of 8.0 GT/s, and 4 lanes of DMI2/PCI Express* Gen 2 interface with a peak transfer rate of 5.0 GT/s. The processor supports up to 46 bits of physical address space and 48-bit of virtual address space.

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

3.2.1 Processor Core Features:

- Up to 8 execution cores (Intel® Xeon® processor E5-2600 product family)
- Up to 12 execution cores (Intel® Xeon® processor E5-2600 v2 product family)
- Each core supports two threads (Intel® Hyper-Threading Technology), up to 16 threads per socket
- 46-bit physical addressing and 48-bit virtual addressing
- 1 GB large page support for server applications
- A 32-KB instruction and 32-KB data first-level cache (L1) for each core
- A 256-KB shared instruction/data mid-level (L2) cache for each core
- Up to 20 MB last level cache (LLC): up to 2.5 MB per core instruction/data last level cache (LLC), shared among all cores

3.2.2 Supported Technologies:

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

3.2.3 Intel® QuickPath Interconnect

The Intel® QuickPath Interconnect (QPI) 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 inter-connect 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.

3.2.4 Integrated Memory Controller (IMC) and Memory Subsystem

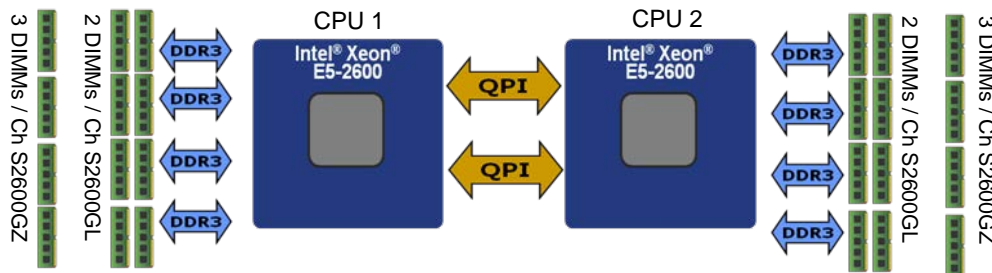


Figure 12. Integrated Memory Controller Functional Block Diagram

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

- Unbuffered DDR3 and registered DDR3 DIMMs
- LR DIMM (Load Reduced DIMM) for buffered memory solutions demanding higher capacity memory subsystems
- Independent channel mode or lockstep mode
- Data burst length of eight cycles for all memory organization modes
- Memory DDR3 data transfer rates of 800, 1066, 1333, 1600, 1866² MT/s,
- 64-bit wide channels plus 8-bits of ECC support for each channel
- DDR3 standard I/O Voltage of 1.5 V and DDR3 Low Voltage of 1.35 V
- 1-Gb, 2-Gb, and 4-Gb DDR3 DRAM technologies supported for these devices:
 - UDIMM DDR3 – SR x8 and x16 data widths, DR – x8 data width
 - RDIMM DDR3 – SR, DR, and QR – x4 and x8 data widths
 - LRDIMM DDR3 – QR – x4 and x8 data widths with direct map or with rank multiplication
- Up to 8 ranks supported per memory channel, 1, 2 or 4 ranks per DIMM
- Open with adaptive idle page close timer or closed page policy
- Per channel memory test and initialization engine can initialize DRAM to all logical zeros with valid ECC (with or without data scrambler) or a predefined test pattern
- Isochronous access support for Quality of Service (QoS)
- Minimum memory configuration: independent channel support with 1 DIMM populated
- Integrated dual SMBus* master controllers
- Command launch modes of 1n/2n
- RAS Support:
 - Rank Level Sparring 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 & 1 and channels 2 & 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
 - CPU1 Channel Mirror Pairs (A,B) and (C,D)
 - CPU2 Channel Mirror Pairs (E,F) and (G,H)
 - Error Containment Recovery

² Intel® Xeon® processor E5-2600 v2 product family only

- Improved Thermal Throttling with dynamic Closed Loop Thermal Throttling (CLTT)
- Memory thermal monitoring support for DIMM temperature

3.2.4.1 Supported Memory

Table 4. UDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ¹			Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)										
				2 Slots per Channel Intel® Server Board S2600GL				3 Slots per Channel Intel® Sever Board S2600GZ						
				1DPC		2DPC		1DPC		2DPC		3DPC		
				1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V			
SRx8 ECC	1GB	2GB	4GB	1066, 1333	1066, 1333	1066	1066, 1333	1066	1066, 1333	1066	1066, 1333	1066	1066, 1333	n/a
DRx8 ECC	2GB	4GB	8GB	1066, 1333	1066, 1333	1066	1066, 1333	1066	1066, 1333	1066	1066, 1333	1066	1066, 1333	n/a

Table 5. UDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ¹			Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)									
				2 Slots per Channel Intel® Server Board S2600GL				3 Slots per Channel Intel® Sever Board S2600GZ					
				1DPC		2DPC		1DPC		2DPC		3DPC	
				1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V		
SRx8 ECC	1GB	2GB	4GB	1066, 1333	1066, 1333, 1600, 1866	1066, 1333	1066, 1333, 1600	1066, 1333	1066, 1333, 1600, 1866	1066, 1333	1066, 1333, 1600	1066, 1333, 1600	Not supported
DRx8 ECC	2GB	4GB	8GB	1066, 1333	1066, 1333, 1600, 1866	1066, 1333	1066, 1333, 1600	1066, 1333	1066, 1333, 1600, 1866	1066, 1333	1066, 1333, 1600	1066, 1333, 1600	Not supported

Table 6. RDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ¹			Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)									
				2 Slots per Channel Intel® Server Board S2600GL					3 Slots per Channel Intel® Server Board S2600GZ				
				1DPC		2DPC		1DPC		2DPC		3DPC	
				1.35V	1.5V	1.35V	1.5V	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	1066 1333 1600	1066 1333 1600	n/a	800 1066
DRx8	2GB	4GB	8GB	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600	1066 1333 1600	1066 1333 1600	n/a	800 1066
SRx4	2GB	4GB	8GB	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600	1066, 1333	1066 1333 1600	1066 1333 1600	1066 1333 1600	n/a	800 1066
DRx4	4GB	8GB	16GB	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600	1066 1333 1600	1066 1333 1600	n/a	800 1066
QRx4	8GB	16GB	32GB	800	1066	800	800	800	1066	800	800	n/a	n/a
QRx8	4GB	8GB	16GB	800	1066	800	800	800	1066	800	800	n/a	n/a

Table 7. RDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ¹			Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)									
				2 Slots per Channel Intel® Server Board S2600GL					3 Slots per Channel Intel® Server Board S2600GZ				
				1DPC		2DPC		1DPC		2DPC		3DPC	
				1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V
SRx8	1GB	2GB	4GB	1066 1333	1066 1333 1600 1866	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	800	800 1066
DRx8	2GB	4GB	8GB	1066 1333	1066 1333 1600 1866	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	800	800 1066
SRx4	2GB	4GB	8GB	1066 1333	1066 1333 1600 1866	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	800	800 1066
DRx4	4GB	8GB	16GB	1066 1333	1066 1333 1600 1866	1066 1333	1066 1333 1600	1066 1333	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	800	800 1066
QRx8	4GB	8GB	16GB	800	800 1066	800	800	800	800 1066	800	800	n/a	n/a
QRx4	8GB	16GB	32GB	800	800 1066	800	800	800	800 1066	800	800	n/a	n/a

Table 8. LRDIMM Support Guidelines – Intel® Xeon® processor E5-2600 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ²		Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)							
			2 Slots per Channel Intel® Server Board S2600GL			3 Slots per Channel Intel® Server Board S2600GZ				
			1DPC and 2DPC		1DPC and 2DPC		3DPC			
			1.35V	1.5V	1.35V	1.5V	1.35V	1.5V		
QRx4 (DDP)	16GB	32GB	1066	1066, 1333	1066	1066, 1333	1066	1066	1066	1066
QRx8 (P)	8GB	16GB	1066	1066, 1333	1066	1066, 1333	1066	1066	1066	1066

Table 9. LRDIMM Support Guidelines – Intel® Xeon® processor E5-2600 v2 product family

Ranks Per DIMM & Data Width	Memory Capacity Per DIMM ²		Speed (MT/s) and Voltage Validated by Slot per Channel (SPC) and DIMM Per Channel (DPC)											
			2 Slots per Channel Intel® Server Board S2600GL					3 Slots per Channel Intel® Server Board S2600GZ						
			1DPC		2DPC			1DPC		2DPC			3DPC	
			1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V	1.35V	1.5V
QRx4 (DDP)	16GB	32GB	1066 1333 1600	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	1066 1333 1600	1066 1333 1600 1866	1066 1333 1600	1066 1333 1600	1066	1066		
8Rx4 (QDP)	32GB	64GB	1066	1066	1066	1066	1066	1066	1066	1066	1066	1066	1066	

3.2.4.1.1 Phase Shedding Configuration Option in BIOS Setup

Intel has added a user configurable option called “Phase Shedding” to BIOS Setup. By default, this option is configured to “Disabled”, providing maximum VR power which mitigates possible instability issues when using LRDIMMs greater than 16 GB or 3 DPC memory configurations. This option should only be “Enabled” for optimized power in 1DPC and 2DPC configurations and when using LRDIMM sizes 16GB and smaller.

Note: The Phase Shedding BIOS setup option can be found in BIOS revision 02.02.0002 and later.

3.2.4.2 Memory Slot Identification and Population Rules

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

- Each installed processor provides four channels of memory. On the Intel® Server Board S2600GZ each memory channel supports three memory slots, for a total possible 24 DIMMs installed. On the Intel® Server Board S2600GL each memory channel supports 2 memory slots, for a total possible 16 DIMMs installed.
- System memory is organized into physical slots on DDR3 memory channels that belong to processor sockets.
- The memory channels from processor socket 1 are identified as Channel A, B, C and D. The memory channels from processor socket 2 are identified as Channel E, F, G, and H.

- Each memory slot on the server board is identified by channel and slot number within that channel. For example, DIMM_A1 is the first slot on Channel A on processor 1; DIMM_E1 is the first DIMM socket on Channel E on processor 2.
- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- A processor may be installed without populating the associated memory slots provided a second processor is installed with associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS, Error Management,) in the BIOS setup are applied commonly across processor sockets.
- The BLUE memory slots on the server board identify the first memory slot for a given memory channel.

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

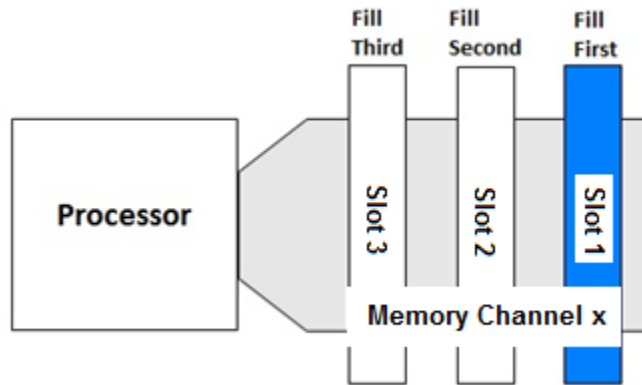


Figure 13. Memory Slots Definition

On the Intel® Server Board S2600GZ, a total of 24 DIMM slots is provided (2 CPUs – 4 Channels/CPU, 3 DIMMs/Channel). The nomenclature for memory slots is detailed in the following table:

Table 10. Intel® Server Board S2600GZ Memory Slot Identification

Processor Socket 1												Processor Socket 2											
(0) Channel A			(1) Channel B			(2) Channel C			(3) Channel D			(0) Channel E			(1) Channel F			(2) Channel G			(3) Channel H		
A1	A2	A3	B1	B2	B3	C1	C2	C3	D1	D2	D3	E1	E2	E3	F1	F2	F3	G1	G2	G3	H1	H2	H3




Figure 14. Intel® Server Board S2600GZ Memory Slot Layout

On the Intel® Server Board S2600GL a total of 16 DIMM slots is provided (2 CPUs – 4 Channels/CPU, 2 DIMMs /Channel). The nomenclature for memory slots is detailed in the following table:

Table 11. Intel® Server Board S2600GL Memory Slot Nomenclature

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




Figure 15. Intel® Server Board S2600GL Memory Slot Layout

The following are generic DIMM population requirements that generally apply to both the Intel® Server Board S2600GZ and Intel® Server Board S2600GL.

- All DIMMs must be DDR3 DIMMs
- Intel only validates and supports ECC memory for its server products.
- Mixing of Registered and Unbuffered DIMMs is not allowed per platform.
- Mixing of LRDIMM with any other DIMM type is not allowed per platform.
- Mixing of DDR3 voltages is not validated within a socket or across sockets by Intel. If 1.35V (DDR3L) and 1.50V (DDR3) DIMMs are mixed, the DIMMs will run at 1.50V.
- Mixing of DDR3 operating frequencies is not validated within a socket or across sockets by Intel. If DIMMs with different frequencies are mixed, all DIMMs will run at the common lowest frequency.
- Quad rank RDIMMs are supported but not validated by Intel.
- A maximum of 8 logical ranks (ranks seen by the host) per channel is allowed.
- DIMMs with different timing parameters can be installed on different slots within the same channel, but only timings that support the slowest DIMM will be applied to all. As a consequence, faster DIMMs will be operated at timings supported by the slowest DIMM populated.
- When one DIMM is used, it must be populated in the BLUE DIMM slot (farthest away from the CPU) of a given channel.
- When single, dual and quad rank DIMMs are populated for 2DPC or 3DPC, always populate the higher number rank DIMM first (starting from the farthest slot), for example, first quad rank, then dual rank, and last single rank DIMM.
- Mixing of quad ranks DIMMs (RDIMM Raw Cards F and H) in one channel and three DIMMs in other channel (3DPC) on the same CPU socket is not validated.

3.2.4.3 Publishing System Memory

- The BIOS displays the “Total Memory” of the system during POST if Quiet Boot is disabled in BIOS setup. This is the total size of memory discovered by the BIOS during POST, and is the sum of the individual sizes of installed 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 Quiet Boot is disabled, the BIOS displays the total system memory on the diagnostic screen at the end of POST. This total is the same as the amount described by the first bullet above.
-

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

These BIOS components include, but are not limited to:

- ACPI (may vary depending on the number of PCI devices detected in the system)
- ACPI NVS table
- Processor microcode
- Memory Mapped I/O (MMIO)
- Manageability Engine (ME)
- BIOS flash

3.2.4.4 Integrated Memory Controller Operating Modes

3.2.4.4.1 *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.

3.2.4.4.2 *Lockstep Channel Mode*

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

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

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

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

3.2.4.4.3 *Mirror Mode*

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

3.2.4.5 **Memory RAS Support**

The server board supports the following memory RAS modes:

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

Regardless of RAS mode, the requirements for populating within a channel given in the section 3.2.2.2 must be met at all times. Note that support of RAS modes that require matching DIMM population between channels (Mirrored and Lockstep) require that ECC DIMMs be populated. Independent Channel Mode is the only mode that supports non-ECC DIMMs in addition to ECC DIMMs.

For Lockstep Channel Mode and Mirroring Mode, processor channels are paired together as a “Domain”.

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

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

3.2.4.5.1 *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.

3.2.4.5.2 *Error Correction Code (ECC) Memory*

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

3.2.4.5.2.1 *Correctable Memory ECC Error Handling*

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

The system BIOS has logic to 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.

3.2.4.5.2.2 *Uncorrectable Memory ECC Error Handling*

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

However, before returning control to the OS drivers from 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 a solid Amber.

3.2.4.5.3 *Demand Scrubbing for ECC Memory*

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

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

3.2.4.5.4 *Patrol Scrubbing for ECC Memory*

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

3.2.4.5.5 *Rank Sparing Mode*

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

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

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

When Sparing Mode is operational, for each channel, the largest size memory rank is reserved as a “spare” and is not used during normal 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.

3.2.4.5.6 Mirrored Channel Mode

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

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

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

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

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

What Mirroring primarily protects against is the possibility of an Uncorrectable ECC Error occurring with critical data “in process”. Without Mirroring, the system would be expected to “Blue Screen” and halt, possibly with serious impact to 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.

3.2.5 Processor Integrated I/O Module (IIO)

The processor 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 supporting up to 40 PCI Express bus lanes via three PCIe ports.
- **DMI2 Interface to the PCH:** The platform requires an interface to the legacy Southbridge (PCH) which provides basic, legacy functions required for the server platform and operating systems. Since only one PCH is required and allowed for the system, any sockets which do not connect to PCH would use this port as a standard x4 PCI Express 2.0 interface.
- **Integrated IOAPIC:** Provides support for PCI Express devices implementing legacy interrupt messages without interrupt sharing
- **Non Transparent Bridge:** PCI Express non-transparent bridge (NTB) acts as a gateway that enables high performance, low overhead communication between two intelligent subsystems; the local and the remote subsystems. The NTB allows a local processor to independently configure and control the local subsystem, provides isolation of the local host memory domain from the remote host memory domain while enabling status and data exchange between the two domains.
- **Intel® QuickData Technology:** Used for efficient, high bandwidth data movement between two locations in memory or from memory to I/O.

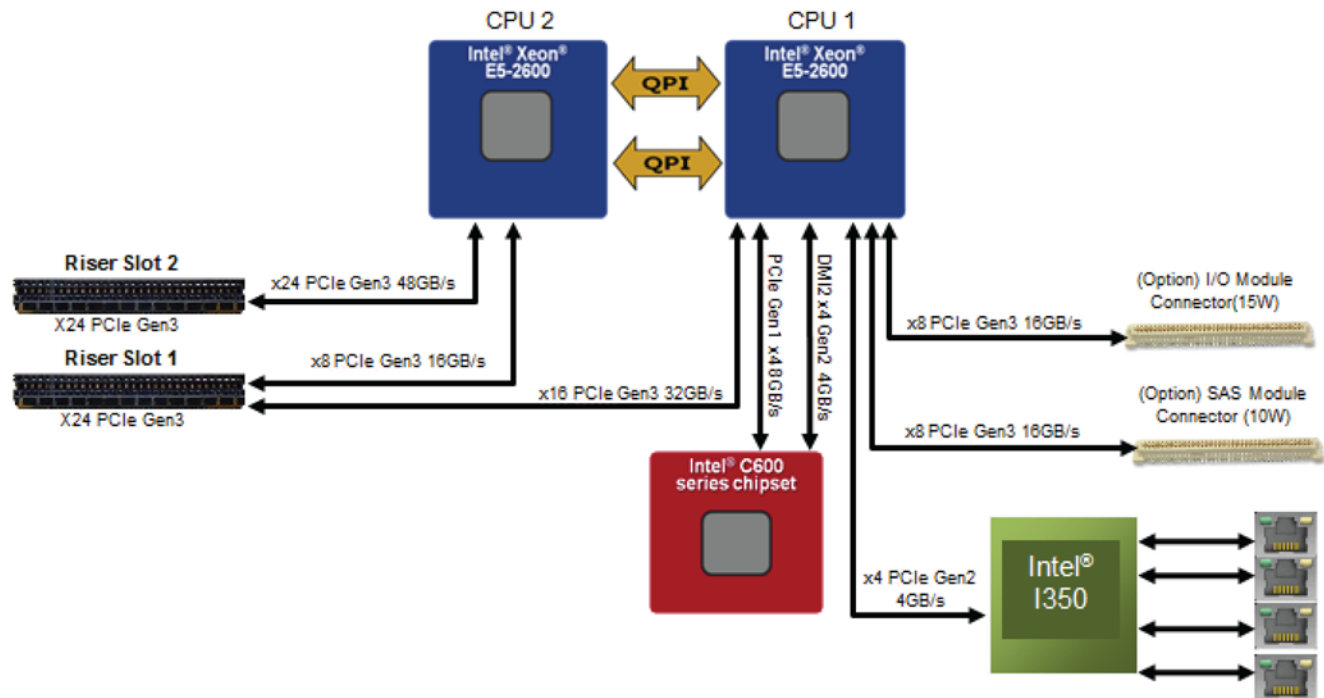


Figure 16. Functional Block Diagram of Processor IIO Sub-system

The following sub-sections describe the server board I/O features supported by the PCIe interface of the processor IIO module. Features and functions of the Intel® C600 Series chipset will be described in its own dedicated section.

3.2.5.1 PCIe Interface

The integrated I/O module of each installed processor provides the PCI Express interface to the server board. Each installed processor supports three PCIe ports, identified as Port 1 (x8), Port 2 (x16), and Port 3 (x16), providing up to 40 PCIe bus lanes that meet the PCIe Gen 1, Gen2, and Gen 3 specifications.

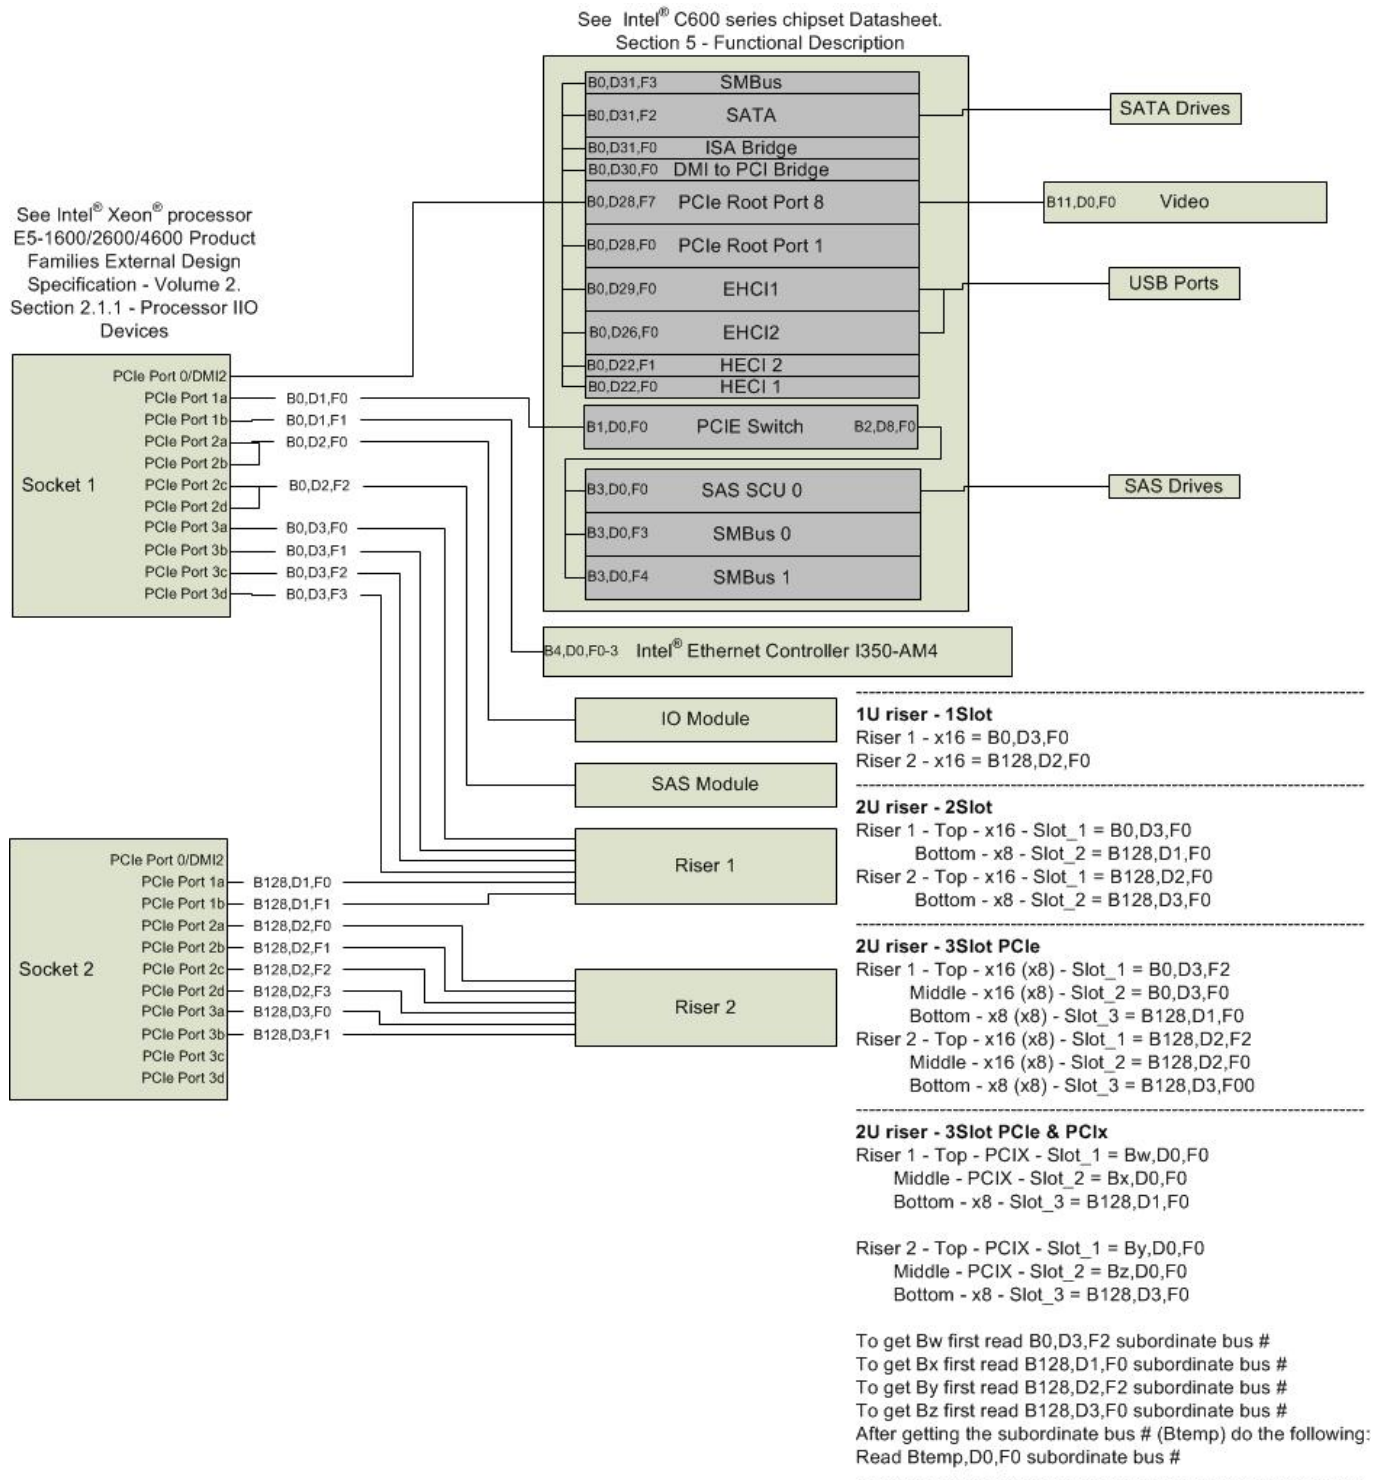


Figure 17. Intel® Server Board S2600GZ/GL PCI Layout

Each PCIe port from the given processor I/O module can be bifurcated down to smaller segments as needed. The number of PCIe port segments required is dependent on the number and type of PCIe devices that need to be supported for the given system configuration. The root port is identified by the port number followed by the letter “A”. As each port is bifurcated, the smaller segments are identified by the port number followed by sequential letters starting with “A”. The following diagrams illustrate possible PCIe port bifurcation options and port identification.

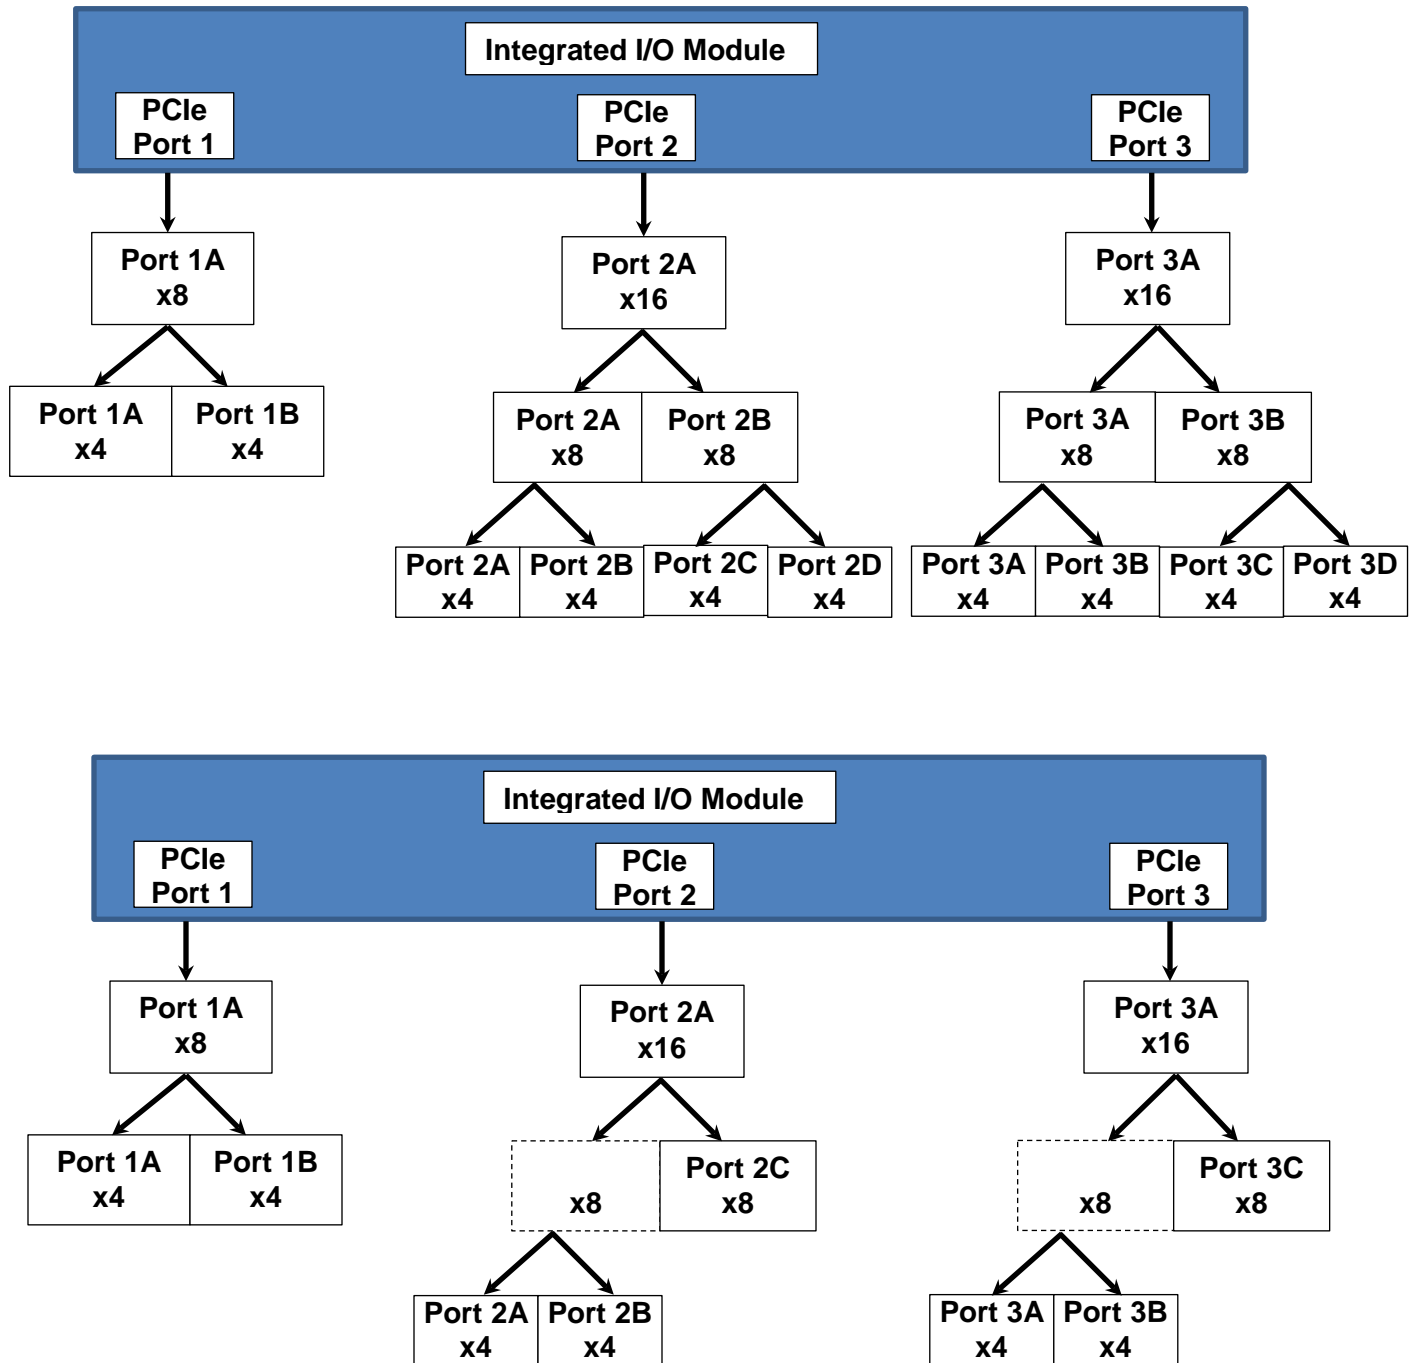


Figure 18. PCIe Port Bifurcation Options

3.2.5.2 Riser Card Support

The server board provides two riser card slots identified as Riser Slot 1 and Riser Slot 2.

Each riser slot has support for a total of x24 PCIe bus lanes. For Riser Slot 1, x16 PCIe lanes are routed from CPU 1 + x8 PCIe lanes are routed from CPU 2. For Riser Slot 2, all x24 PCIe lanes are routed from CPU 2.

Note: Riser Slot 2 can only be used in dual processor configurations. In addition, some 2U riser cards installed in Riser Slot 1 may also require two processors to be installed in order to support all available add-in card slots.

Available riser cards are common between both riser slots. Supported 1U and 2U riser cards include:

- 1U Riser Card – 1 PCIe add-in card slot – PCIe x16, x16 mechanical





Figure 19. 1U PCIe riser for Intel® Server Board S2600GZ/GL


- 2U Riser Card – 3 PCIe add-in card slots



Slot #	Description
Slot-1 (Top)	PCIe x8, x16 mechanical
Slot-2 (Middle)	PCIe x8, x16 mechanical
Slot-3 (Bottom)	PCIe x8, x8 mechanical

Figure 20. 2U three PCIe slots riser for Intel® Server Board S2600GZ/GL


- 2U Riser Card – 2 PCIe add-in card slots



Slot #	Description
Slot-1 (Top)	PCIe x16, x16 mechanical
Slot-2 (Bottom)	PCIe x8, x8 mechanical

Figure 21. 2U two PCIe slots riser for Intel® Server Board S2600GZ/GL

- 2U Riser Card – 2 PCIx + 1 PCIe add-in card slots –



Slot #	Description
Slot-1 (Top)	PCIx 133MHz
Slot-2 (Middle)	PCIx 133 MHz
Slot-3 (Bottom)	PCIe x8, x8 mechanical

Figure 22. 2U three PCIx/PCIe slots riser for Intel® Server Board S2600GZ/GL

3.2.5.3 PCIe Add-in card support

The PCIe sub-system of the server board has support for PCIe add-in cards that follow the PCIe Gen1, Gen2 and Gen3 specifications. However, the performance of some PCIe Gen3 cards may be forced to operate at Gen2 speeds due to electrical signaling characteristic limitations that exist between the server board and some PCIe Gen 3 add-in cards.

Intel has implemented the following PCIe Gen 3 support model for this generation of its server boards and server systems.

3.2.5.3.1 PCIe Gen3 support - Systems configured with an Intel® Xeon® processor E5-2600 product family

For a server system configured with one or more Intel® Xeon® processor E5-2600 product family, the system BIOS will use an embedded PCIe Gen 3 compatibility list which identifies all PCIe Gen 3 add-in cards tested by Intel to operate reliably at Gen 3 speeds on the given server system. During POST, the system BIOS will compare installed PCIe Gen 3 add-in cards with those included in the embedded compatibility list. If BIOS matches an installed card to one listed on the compatibility list, the BIOS will configure the device to operate at PCIe Gen 3 speeds. If the BIOS cannot match an installed PCIe add-in card with any device included in the list, the BIOS will force the device to operate at PCIe Gen2 speeds.

Note: The latest available BIOS should be installed on the system to ensure the most up to date embedded PCIe Gen 3 compatibility list is being used.

Visit the following Intel web site for a list of Intel tested PCIe Gen 3 compatible cards included in the BIOS embedded compatibility list – <http://intel.com/support/motherboards/server/sb/CS-034157.htm>

3.2.5.3.2 PCIe Gen3 support - Systems configured with an Intel® Xeon® processor E5-2600 V2 product family

For a server system configured with one or more Intel® Xeon® processor E5-2600 V2 product family, the system BIOS will configure all installed PCIe Gen 3 compatible add-in cards to operate at PCIe Gen 3 speeds by default. For a list of Intel tested PCIe Gen 3 add-in cards, review the Tested Hardware and OS list (THOL) using Intel's Server Configurator Tool at the following web site:

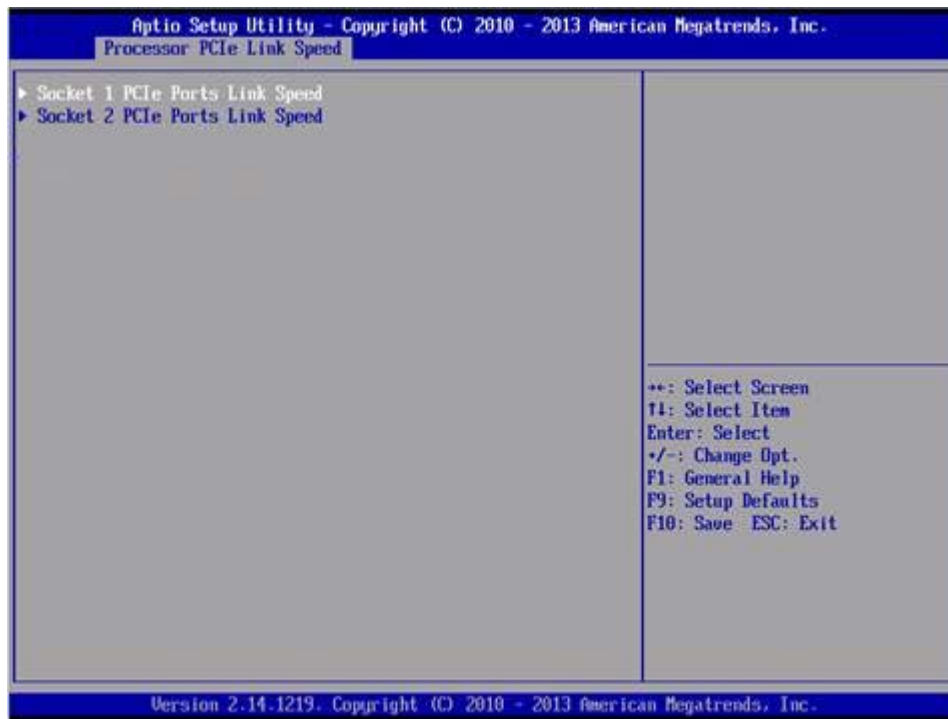
<https://serverconfigurator.intel.com>

The following tables identify the PCIe port routing for each add-in card slot on all supported riser cards as installed in either Riser Slot # 1 or Riser Slot #2. Note the specific processor, PCIe port ID, and number of PCIe bus lanes supporting each add-in card slot.

Depending on the riser card installed, specific PCIe ports routed from the processor IIO module provide the PCIe interface to each riser card slot.

For a system configuration that includes a PCIe add-in card that is not included on Intel's THOL and/or may be operating erratically, BIOS setup includes an option to force specified PCIe ports to operate at a different PCIe level than the card is optimally designed to support. To configure this option, access the <F2> BIOS setup utility.

From the BIOS Setup “Main Menu”, navigate to the following menus: “Advanced Menu”, “PCI Configuration”, “Processor PCIe Link Speed”



The “Processor PCIe Link Speed” menu will display selectable options for each installed processor, identified as “Socket #”, where # identifies the CPU number.

Using information provided in *the following tables*, select the processor associated with the PCIe port to be configured.

Table 12. Riser Slot #1 – PCIe Port Routing

Riser Slot #1 – Riser Card Options			
2U - PCIx Riser Card	2U – 3 PCIe Slot Riser Card	2U – 2 PCIe Slot Riser Card	1U – 1 PCIe Slot Riser Card
Top PCIx Slot (CPU #1 – Port 3C)	Top PCIe Slot (CPU #1 – Port 3C) (x8 elec, x16 mech)	Top PCIe Slot (CPU #1 – Port 3A) (x16 elec, x16mech)	PCIe Slot (CPU #1 – Port 3A) (x16 elec, x16mech)
Middle PCIx Slot (CPU #2 – Port 1B)	Middle PCIe Slot (CPU #1 – Port 3A) (x8 elec, x16 mech)		
Bottom PCIe Slot (CPU # 1 – Port 3A) (x8 elec, x8 mech)	Bottom PCIe Slot (CPU #2 – Port 1A) (x8 elec, x8 mech)	Bottom PCIe Slot (CPU #2 – Port 1A) (x8 elec, x8 mech)	

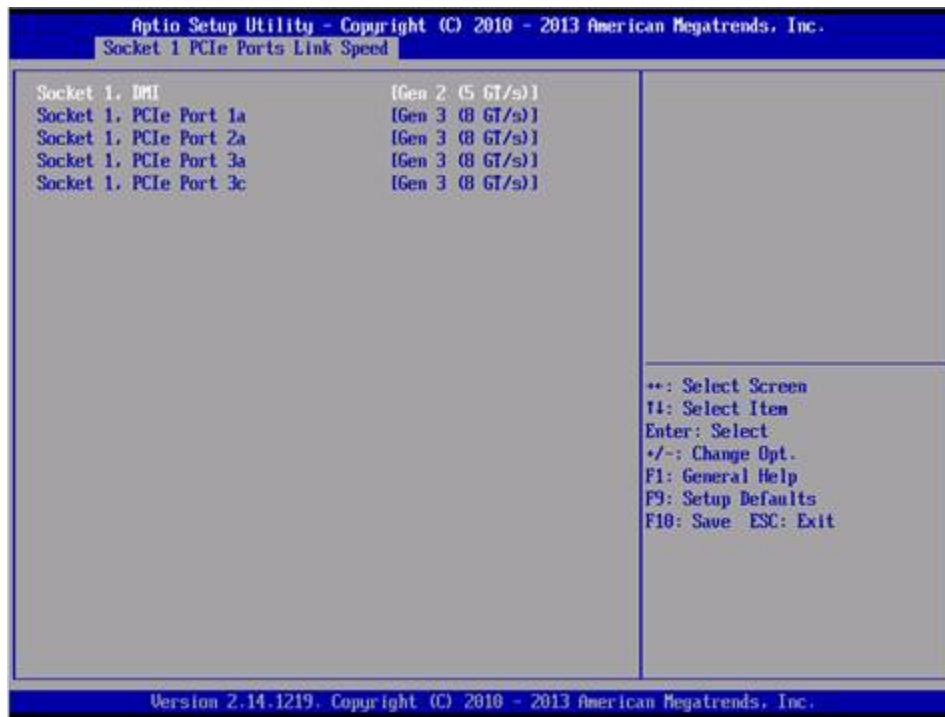
Table 13. Riser Slot #2 – PCIe Port Routing

Riser Slot #2 – Riser Card Options			
2U - PCIx Riser Card	2U – 3 PCIe Slot Riser Card	2U – 2 PCIe Slot Riser Card	1U – 1 PCIe Slot Riser Card
Top PCIx Slot (CPU #2 – Port 2C)	Top PCIe Slot (CPU #2 – Port 2C) (x8 elec, x16 mech)	Top PCIe Slot (CPU #2 – Port 2A) (x16 elec, x16 mech)	PCIe Slot (CPU #2 – Port 2A) (x16 elec, x16 mech)
Middle PCIx Slot (CPU #2 –Port 3A)	Middle PCIe Slot (CPU #2 – Port 2A) (x8 elec, x16 mech)		
Bottom PCIe Slot (CPU #2 – Port 2A) (x8 elec, x8 mech)	Bottom PCIe Slot (CPU #2 – Port 3A) (x8 elec, x8 mech)	Bottom PCIe Slot (CPU #2 – Port 3A) (x8 elec, x8 mech)	

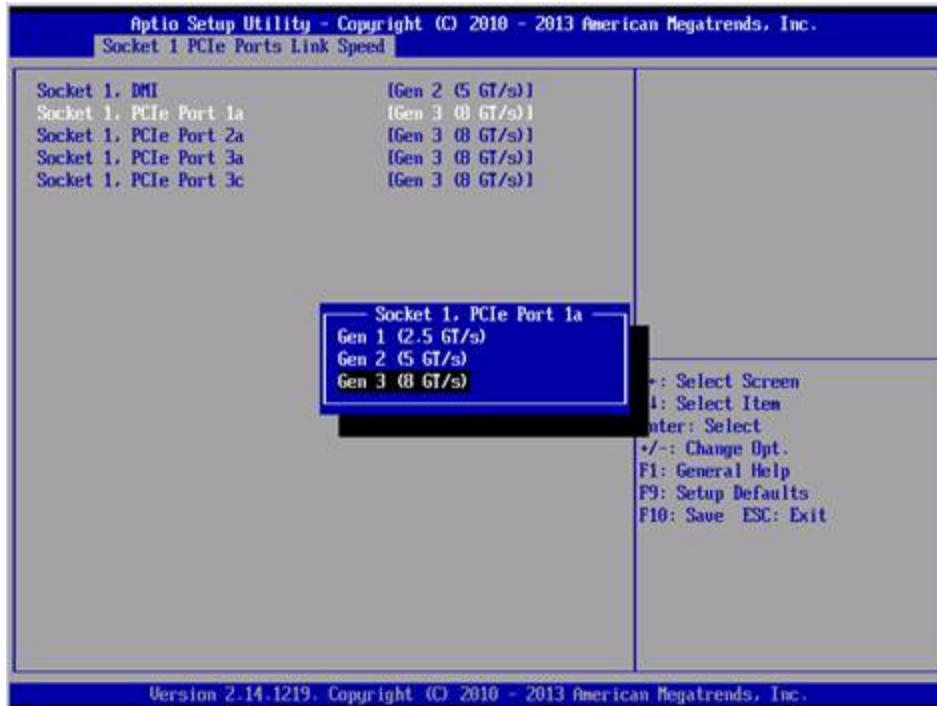
The “Socket # PCIe Ports Link Speed” window displays selectable options for each configurable PCIe port associated with the current system configuration.

Note: The illustrations below are for reference purposes only. Actual PCIe port data displayed in the “Socket # PCIe Ports Link Speed” window may be different than what is shown here.

Using the arrow keys, move the cursor down to the PCIe port to be changed.



Once a port is selected, a port configuration window appears and provides options to configure the specified PCIe port to operate at a specified PCIe Gen level.




Select the desired PCIe Gen level. After making all desired changes in BIOS setup, be sure to save the changes and reboot the system.

3.2.5.4 Network Interface

On-board network connectivity is provided by means of an Intel® Ethernet Controller I350-AM4 providing up to four 10/100/1000 Mb Ethernet ports. The NIC chip is supported by x4 PCIe Gen2 signals from the IIO module of the CPU 1 processor.

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:



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

Figure 23. Intel® Server Board S2600GZ/GL External RJ45 NIC Port LED Definition

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

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

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

3.2.5.5 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 (J2B1) labeled "IO_Module" and is supported by x8 PCIe Gen3 signals from the IIO module of the CPU 1 processor.




Figure 24. Server Board Layout - I/O Module Connector

Supported I/O modules include:

Table 14. Supported Intel® I/O Module Options

Quad Port Intel® I350 GbE I/O Module	AXX4P1GBPWL1OM
Dual Port Intel® X540 10GbE I/O Module	AXX10GBTWL1OM
Dual Port Intel® 82599 10GbE I/O Module	AXX10GBNIA1OM
Single Port FDR InfiniBand* ConnectX*-3 I/O Module	AXX1FDRIB1OM
Dual Ports FDR InfiniBand* ConnectX*-3 I/O Module	AXX2FDRIB1OM

3.2.5.6 Intel® Integrated RAID Option

The server board provides support for Intel® Integrated RAID modules. These optional modules attach to a high density 80-pin connector labeled “SAS Module” on the server board and are supported by x8 PCIe Gen3 signals from the IIO module of the CPU 1 processor.




Figure 25. Server Board Layout – Intel® Integrated RAID Module Option Placement

Features of this option include:

- SKU options to support full or entry level hardware RAID
- Dual-core 6Gb SAS ROC/IOC (LSI* 2208)
- 4 or 8 port and SAS/SATA or SATA –only ROC options
- SKU options to support 512MB or 1GB embedded memory

- Intel designed flash + optional support for Intel® RAID Maintenance Free Backup Units (AXXRMFBU2) or improved Lithium Polymer battery

Table 15. Supported Intel® Integrated RAID Modules

External Name	Description	Product Code
Intel® Integrated RAID Module RMS25CB080	8 Port SAS-2.1, Full HW RAID, 1GB, IOM Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMS25CB080
Intel® Integrated RAID Module RMS25CB040	4 Port SAS-2.1, Full HW RAID, 1GB, IOM Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMS25CB040
Intel® Integrated RAID Module RMT3CB080	8 Port SATA-3, Full HW RAID, 512MB, IOM Slot RAID Levels 0,1,10, 5, 50, 6, 60	RMT3CB080
Intel® Integrated RAID Module RMS25KB080	8 Port SAS-2.1, Entry-level HW RAID, IOM Slot RAID Levels 0,1,1E	RMS25KB080
Intel® Integrated RAID Module RMS25KB040	4 Port SAS-2.1, Entry-level HW RAID, IOM Slot RAID Levels 0,1,1E	RMS25KB040

For additional product information, please reference the following Intel document:
Intel Integrated RAID Module RMS25PB080, RMS25PB040, RMS25CB080, and RMS25C040 Hardware Users Guide.

3.3 Intel® C602 Chipset Functional Overview

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

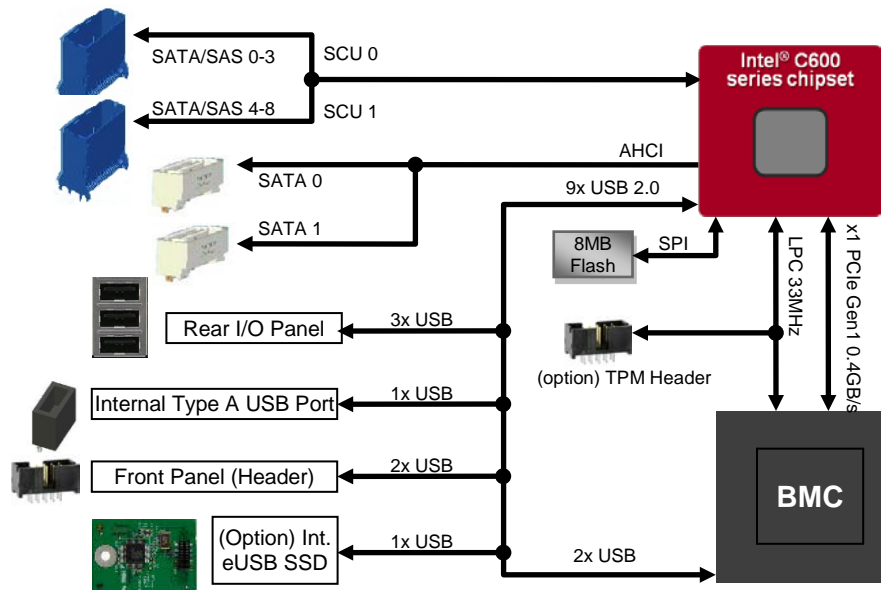


Figure 26. Functional Block Diagram - Chipset Supported Features and Functions

On the Intel® Server Boards S2600GZ and S2600GL, the chipset provides support for the following on-board functions:

- Low Pin Count (LPC) interface
- Universal Serial Bus (USB) Controller
- Serial Attached SCSI (SAS)/Serial ATA (SATA) Support
- 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 as well as providing support for the optional Trusted Platform Module (TPM).

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.

- Three external USB ports are provided in a stacked housing located on the rear I/O section of the server board.
- Two USB ports are routed to an internal 10-pin connector (Labeled “FP_USB” on the server board) that can be cabled for front panel support.
- One internal Type ‘A’ USB port.
- One eUSB connector intended for use with an optional eUSB SSD device.
- Two USB ports are routed to the integrated BMC.

3.3.2.1 eUSB SSD Support

The server board provides support for a low profile eUSB SSD storage device. A 2mm 2x5-pin connector labeled “eUSB SSD” near the rear I/O section of the server board is used to plug this small flash storage device into.




Figure 27. Low Profile eUSB SSD Support

eUSB SSD features include:

- 2 wire small form factor Universal Serial Bus 2.0 (Hi-Speed USB) interface to host
- Read Speed up to 35 MB/s and write Speed up to 24 MB/s.
- Capacity range from 256 MB to 32 GB.
- Support USB Mass Storage Class requirements for Boot capability.

3.3.3 Embedded Serial ATA (SATA)/Serial Attached SCSI (SAS)/RAID Support

The Intel® C602 chipset provides storage support from two integrated controllers: AHCI and SCU. By default the server board will support up to 6 SATA ports: Two single 6Gb/sec SATA ports routed from the AHCI controller to the two white SATA connectors labeled “SATA-0” and “SATA-1”, and four 3Gb/sec SATA ports routed from the SCU to the mini-SAS connector labeled “SCU_0 (0-3)”.

Note: The mini-SAS connector labeled “SCU_1 (4-7)” is NOT functional by default and is only enabled with the addition of an Intel® RAID C600 Upgrade Key option supporting 8 SAS/SATA ports.

The server board is capable of supporting additional chipset embedded SAS, SATA, and RAID options 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”.




Figure 28. Intel® RAID C600 Upgrade Key Connector

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

Table 16. Intel® RAID C600 Upgrade Key Options

Intel® RAID C600 Upgrade Key Options (Intel Product Codes)	Key Color	Description
Default – No option key installed	N/A	4 Port SATA with Intel® ESRT RAID 0,1,10 and Intel® RSTe RAID 0,1,5,10
RKSATA4R5	Black	4 Port SATA with Intel® ESRT2 RAID 0,1, 5, 10 and Intel® RSTe RAID 0,1,5,10
RKSATA8	Blue	8 Port SATA with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,5,10
RKSATA8R5	White	8 Port SATA with Intel® ESRT2 RAID 0,1, 5, 10 and Intel® RSTe RAID 0,1,5,10
RKSAS4	Green	4 Port SAS with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,10
RKSAS4R5	Yellow	4 Port SAS with Intel® ESRT2 RAID 0,1, 5, 10 and Intel® RSTe RAID 0,1,10
RKSAS8	Orange	8 Port SAS with Intel® ESRT2 RAID 0,1, 10 and Intel® RSTe RAID 0,1,10
RKSAS8R5	Purple	8 Port SAS with Intel® ESRT2 RAID 0,1, 5, 10 and Intel® RSTe RAID 0,1,10

Additional information for the on-board RAID features and functionality can be found in the *Intel® RAID Software Users Guide* (Intel Document Number D29305-015).

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 & 8 Port SATA RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key
 - 4 & 8 Port SAS RAID 5 support provided with appropriate Intel® RAID C600 Upgrade Key
- Maximum drive support = 8
- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*.
- OS Support = Windows 7*, Windows 2008*, Windows 2003*, RHEL*, SLES, other Linux* variants using partial source builds.
- Utilities = Windows* GUI and CLI, Linux GUI and CLI, DOS CLI, and EFI CLI

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 = Windows 7*, Windows 2008*, Windows 2003*, RHEL* 6.2 ¹ and later, SLES* 11 w/SP2¹ and later, VMWare* 5.x.
- Utilities = Windows* GUI and CLI, Linux CLI, DOS CLI, and EFI CLI
- Uses Matrix Storage Manager for Windows*
- MDRAID supported in Linux (Does not require a driver)

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

1) See latest product errata list for support status.

Product Errata are documented in the *Intel® Server Board S2600GZ/GL, Intel® Server System R1000GZ/GL, Intel® Server System R2000GZ/GL Monthly Specification Update* which can be downloaded from www.intel.com.

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 either 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.

3.4 Integrated Baseboard Management Controller Overview

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

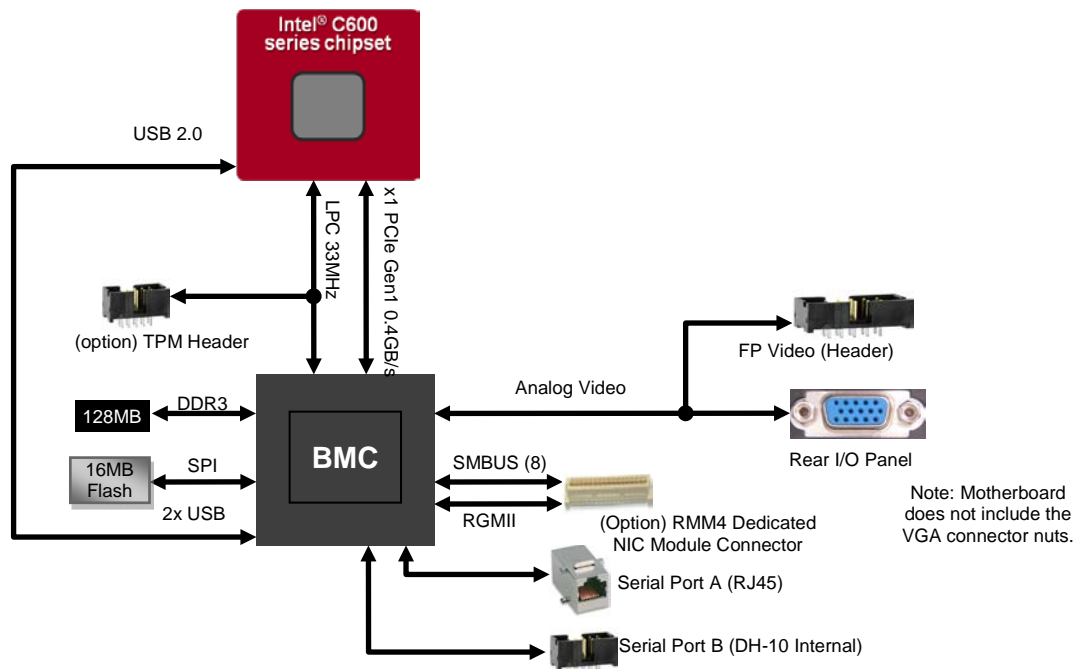


Figure 29. Integrated BMC Block Diagram

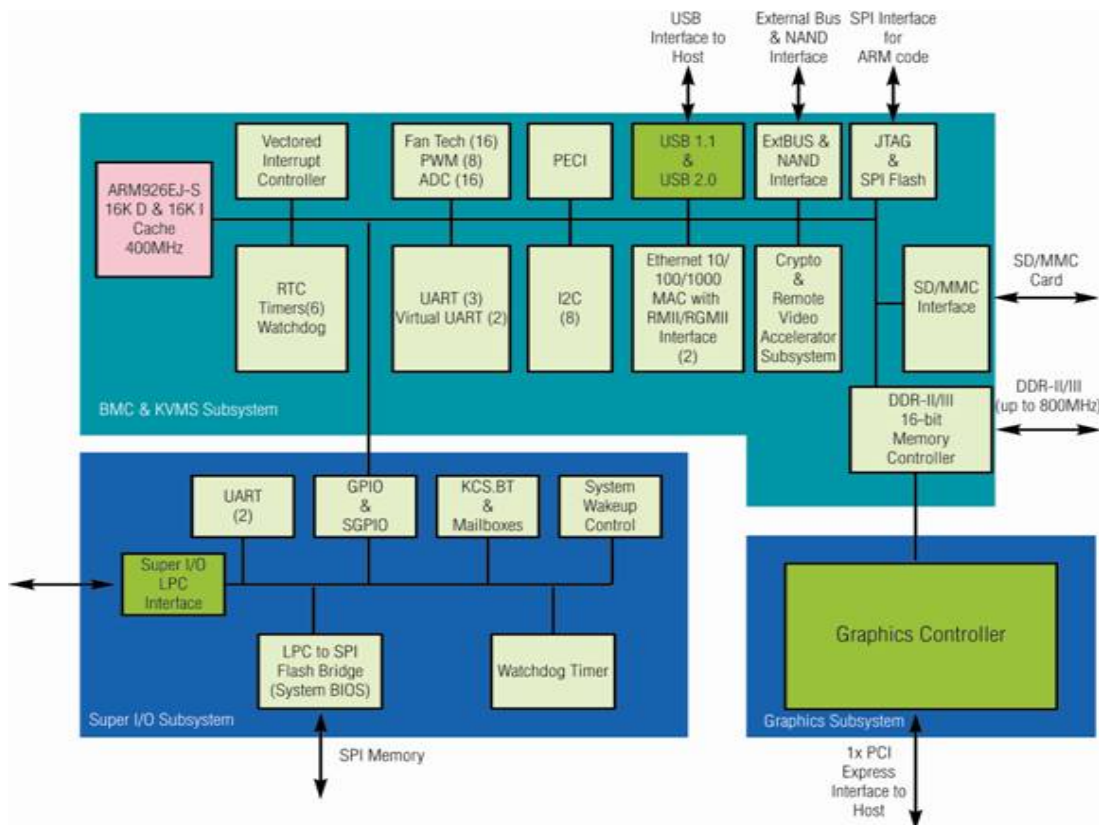


Figure 30. Integrated BMC Functional Block Diagram

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 GPIO's
- Serial GPIO support for 80 general purpose inputs and 80 general purpose outputs available for host processor
- Programmable Wake-up Event Support
- Plug and Play Register Set
- Power Supply Control
- Host SPI bridge for system BIOS support

3.4.1.1 Keyboard and Mouse Support

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

3.4.1.2 Wake-up Control

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

3.4.2 Graphics Controller and Video Support

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

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

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

Table 17. 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 18. BIOS Setup Options for Configuring Video

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

3.4.3.1 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.

For a functional overview of the baseboard management features, refer to Platform Management Functional Overview.

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, anyone with access to the system 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 Intel® TPM module AXTPME5 Support

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

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

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

4.2.1 TPM security BIOS

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

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

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

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 S2600GZ/GL provides TPM settings through the security screen.

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

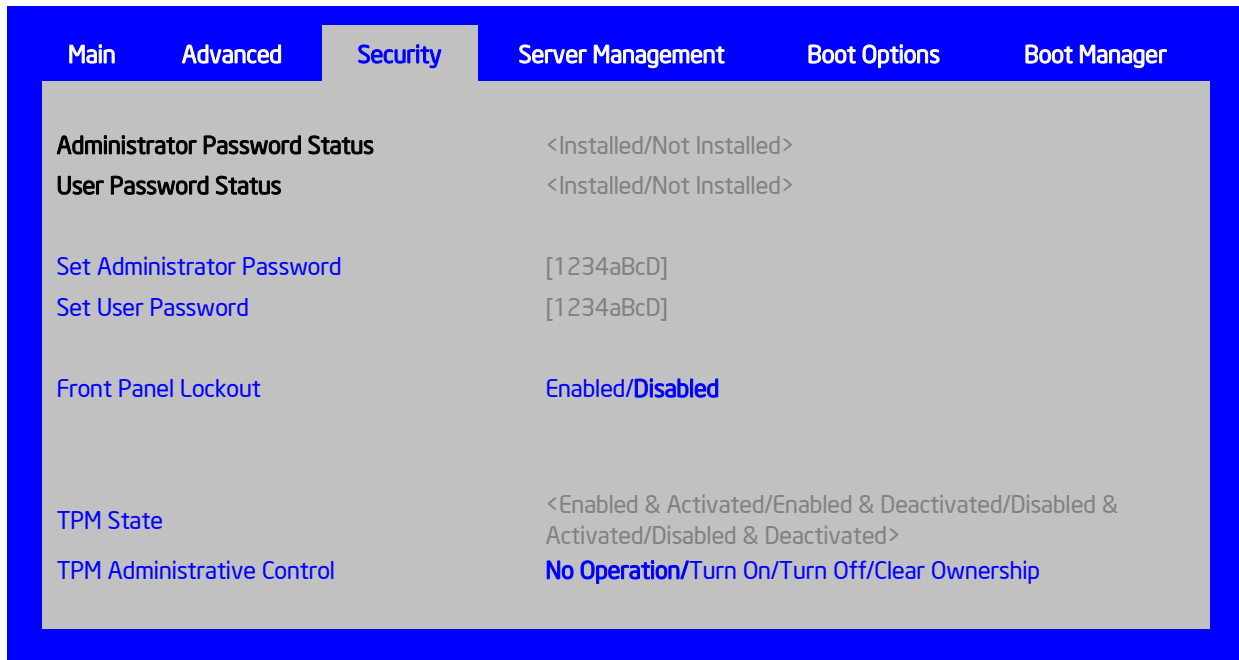


Figure 31. Setup Utility – TPM Configuration Screen

Table 19. 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		<p>Information only. Shows the current TPM device state.</p> <p>A disabled TPM device will not execute commands that use TPM functions and TPM security operations will not be available.</p> <p>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.</p> <p>An enabled and activated TPM executes all commands that use TPM functions and TPM security operations will be available.</p>
TPM Administrative Control**	No Operation Turn On Turn Off Clear Ownership	<p>[No Operation] - No changes to current state.</p> <p>[Turn On] - Enables and activates TPM.</p> <p>[Turn Off] - Disables and deactivates TPM.</p> <p>[Clear Ownership] - Removes the TPM ownership authentication and returns the TPM to a factory default state.</p> <p>Note: The BIOS setting returns to [No Operation] on every boot cycle by default.</p>	

4.3 Intel® Trusted Execution Technology

The Intel® Xeon® Processor E5-4600/2600/2400/1600 Product Families support Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment. Designed to help protect against software-based attacks, Intel® Trusted Execution Technology integrates new security features and capabilities into the processor, chipset and other platform components. When used in conjunction with Intel® Virtualization Technology, Intel® Trusted Execution Technology provides hardware-rooted trust for your virtual applications. This hardware-rooted security provides a general-purpose, safer computing environment capable of running a wide variety of operating systems and applications to increase the confidentiality and integrity of sensitive information without compromising the usability of the platform.

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

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

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

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

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

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

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

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

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

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

5.2 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 20. 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	Control capability that can maintain a power budget to enable increased rack population.
High energy costs	Control capability that can maintain a power budget to ensure that a set energy cost can be achieved
Capacity planning	<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 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.

5.2.1 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 refer the *BMC Core Firmware External Product Specification (EPS)* and *BIOS Core External Product Specification (EPS)* for Intel® Server products based on the Intel® Xeon® processor E5-4600,2600,1600 product families.

6.1 Baseboard Management Controller (BMC) Firmware Feature Support

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

6.1.1 IPMI 2.0 Features

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

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

6.1.2 Non IPMI Features

The BMC supports the following non-IPMI features.

- In-circuit BMC firmware update
- 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

6.2 Advanced Configuration and Power Interface (ACPI)

The server board has support for the following ACPI states:

Table 21. 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 Platform Specific Information 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.

6.3 Power Control Sources

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

Table 22. Power Control Initiators

Source	External Signal Name or Internal Subsystem	Capabilities
Power button	Front panel power button	Turns power on or off
BMC watchdog timer	Internal BMC timer	Turns power off, or power cycle

Command	Routed through command processor	Turns power on or off, or power cycle
Power state retention	Implemented by means of BMC internal logic	Turns power on when AC power returns
Chipset	Sleep S4/S5 signal (same as <i>POWER_ON</i>)	Turns power on or off
CPU Thermal	CPU Thermtrip	Turns power off
WOL(Wake On LAN)	LAN	Turns power on

6.4 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 safeguard 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 a BMC HW reset or FW stack reset. This state (which follows the state of the associated sensor) will be cleared upon system reset or (AC or DC) power cycle.

Note: 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 parameterPotentially incorrect ACPI Power State reported by the BMC.
2. Reversion of temporary test modes for the BMC back to normal operational modes.
3. FP status LED and DIMM fault LEDs may not reflect BIOS detected errors.

6.5 Fault Resilient Booting (FRB)

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

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.

6.6 Sensor Monitoring

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

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

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

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

6.7 Field Replaceable Unit (FRU) Inventory Device

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

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

6.8 System Event Log (SEL)

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

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

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

6.9 System Fan Management

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

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

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

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

6.9.1 Thermal and Acoustic Management

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

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

In order to maintain comprehensive thermal protection, deliver the best system acoustics, and fan power efficiency, an intelligent Fan Speed Control (FSC) and thermal management technology (mechanism) is used. Options in <F2> BIOS Setup (**BIOS > Advanced > System Acoustic and Performance Configuration**) allow for parameter adjustments based on the actual system configuration and usage. Refer to System Acoustic and Performance Configuration for a description of each setting.

- Set Throttling Mode
- Altitude
- Set Fan Profile
- Fan PWM Offset
- Quiet Fan Idle Mode

Note: The above features may or may not be in effective depends on the actual thermal characters of a specific system. Refer to Intel® Server System R1000GZ/GL product family Technical Product Specification and Intel® Server System R2000GZ/GL product family Technical Product Specification for system thermal and acoustic management.

6.9.2 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^{1, 4, 8}
- PCH Temperature Sensor^{4,6}
- On-board Ethernet Controller Temperature Sensors^{4, 6}
- Add-In Intel SAS/IO Module Temperature Sensors^{4, 6}

- PSU Thermal Sensor^{4, 9}
- CPU VR Temperature Sensors^{4, 7}
- DIMM VR Temperature Sensors^{4, 7}
- Integrated BMC Temperature Sensor^{4, 7}
- Global Aggregate Thermal Margin Sensors⁸

Notes:

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

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

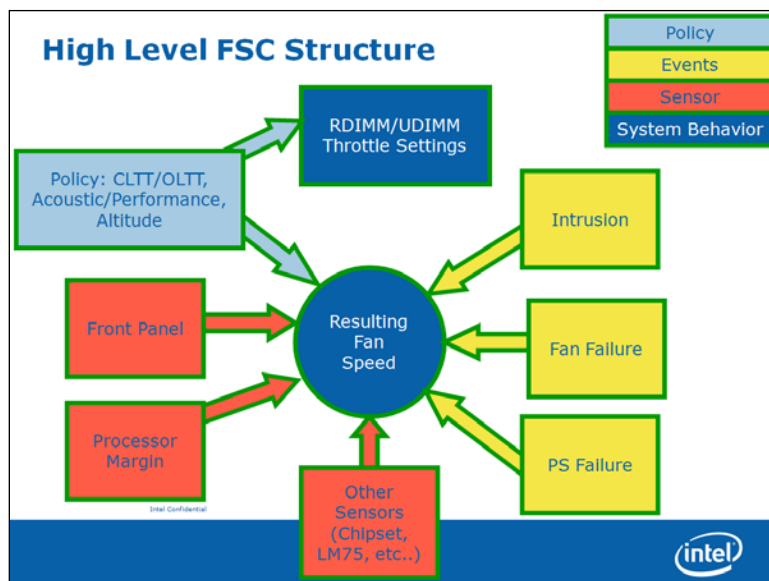


Figure 32. Fan Speed Control Process

6.9.3 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 Closed Loop Thermal Throttling (Dynamic-CLTT):** CLTT control registers are configured by BIOS MRC during POST. The memory throttling is run as a closed-loop system with the DIMM temperature sensors as the control input. Adjustments are made to the throttling during runtime based on changes in system cooling (fan speed).

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

6.10 Messaging Interfaces

The BMC supports the following communications interfaces:

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

Every messaging interface is assigned an IPMI channel ID by IPMI 2.0. The following tables shows the standard channel assignments.

Table 23. Mesaaging Interfaces

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

Notes:

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

6.10.1 User Model

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

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

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

6.10.2 IPMB Communication Interface

The IPMB communication interface 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.

6.10.3 LAN Interface

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

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

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

6.10.3.1 RMCP/ASF Messaging

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

6.10.3.2 BMC LAN Channels

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

6.10.3.2.1 Baseboard NICs

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

6.10.3.2.2 *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.

6.10.3.2.3 *Concurrent Server Management Use of Multiple Ethernet Controllers*

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

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

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

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

For security reasons, embedded LAN channels have the following default 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

6.10.3.3 IPV6 Support

In addition to IPv4, the server board has support for IPv6 for manageability channels. Configuration of IPv6 is provided by extensions to the IPMI Set & 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.

6.10.3.4 LAN Failover

The BMC FW provides a LAN failover capability such that the failure of the system HW associated with one LAN link will result in traffic being rerouted to an alternate link. This functionality is configurable 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 (the next active LAN link will be chosen randomly from the pool of backup LAN links with link status as "UP"). Traffic immediately resumes on the new active connection.

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

6.10.3.5 BMC IP Address Configuration

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

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

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

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

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

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

The BMC takes no special action when the following IP address source is specified as the IP address source for any management NIC:

- 1h – Static address (manually configured)

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

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

6.10.3.5.2 Static LAN Configuration Parameters

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

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

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

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

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

6.10.3.5.3 Enabling/Disabling Dynamic Host Configuration (DHCP) Protocol

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

6.10.3.5.4 DHCP-related LAN Configuration Parameters

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

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

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

6.10.3.6 DHCP BMC Hostname

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

6.10.4 Address Resolution Protocol (ARP)

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

6.10.5 Internet Control Message Protocol (ICMP)

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

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

6.10.6 Virtual Local Area Network (VLAN)

The BMC supports VLAN as defined by IPMI 2.0 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.

6.10.7 Secure Shell (SSH)

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

6.10.8 Serial-over-LAN (SOL 2.0)

The BMC supports IPMI 2.0 SOL.

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

Three commands are implemented for SOL 2.0 configuration.

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

6.10.9 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 24. 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.

6.10.10 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.

6.10.10.1 SNMP Platform Event Traps (PETs)

This feature enables a target system to send SNMP traps to a designated IP address by means of LAN. These alerts are formatted per the *Intelligent Platform Management Interface Specification Second Generation v2.0*. A 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.

6.10.11 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.

6.10.11.1 E-mail Alerting

The Embedded Email Alerting feature allows the user to receive e-mails alerts indicating issues with the server. This allows e-mail alerting in an OS-absent (for example, Pre-OS and OS-Hung) situation. This feature provides support for sending e-mail by means of SMTP, 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.

6.10.12 SM-CLP (SM-CLP Lite)

SMASH refers to Systems Management Architecture for Server Hardware. SMASH is defined by a suite of specifications, managed by the DMTF, that standardize the manageability interfaces for server hardware. CLP refers to Command Line Protocol. SM-CLP is defined by the *Server Management Command Line Protocol Specification (SM-CLP) 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.

6.10.13 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 & Power supply related)
- Automatic refresh of sensor data with a configurable refresh rate.
- On-line help.
- Display/clear SEL (display is in easily understandable human readable format).
- Supports major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*).
- The GUI session automatically times-out after a user-configurable inactivity period. By default, this inactivity period is 30 minutes.
- Embedded Platform Debug feature - Allow the user to initiate a "diagnostic dump" to a file that can be sent to Intel for debug purposes.
- Virtual Front Panel. The Virtual Front Panel provides the same functionality as the local front panel. The displayed LEDs match the current state of the local panel LEDs. The displayed buttons (for example, power button) can be used in the same manner as the local buttons.
- Display of ME sensor data. Only sensors that have associated SDRs loaded will be displayed.
- Ability to save the SEL to a file.
- Ability to force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display of processor and memory information as is available over IPMI over LAN.
- Ability to get and set Node Manager (NM) power policies.
- Display of power consumed by the server.
- Ability to view and configure VLAN settings.
- Warn user the reconfiguration of IP address will cause disconnect.

- Capability to block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.
- Server Power Control - Ability to force into Setup on a reset.

6.10.14 Virtual Front Panel

- Virtual Front Panel is the module present as “Virtual Front Panel” on the left side in the embedded web server when “remote Control” tab is clicked.
- Main Purpose of the Virtual Front Panel is to provide the front panel functionality virtually.
- Virtual Front Panel (VFP) will mimic the status LED and Power LED status and Chassis ID alone. It is automatically in sync with BMC every 40 seconds.
- For any abnormal status LED state, Virtual Front Panel will get the reason behind the abnormal or status LED changes and displayed in VFP side.
- As Virtual Front Panel uses the chassis control command for power actions. It won’t log the Front button press event since Logging the front panel press event for Virtual Front Panel press will mislead the administrator.
- For Reset 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 won’t 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 available for virtual panel module.
- At present, NMI button in VFP is disabled in Intel® S1400/S1600/S2400/S2600 Server Platforms. It can be used in future.

6.10.15 Embedded Platform Debug

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

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

- Platform sensor readings – This includes all “readable” sensors that can be accessed by the BMC FW and have associated SDRs populated in the SDR repository. This does not include any “event-only” sensors. (All BIOS sensors and some BMC and ME sensors are “event-only”; meaning that they are not readable using an IPMI Get Sensor Reading command but rather are used just for event logging purposes).
- SEL – The current SEL contents are saved in both hexadecimal and text format.

- CPU/memory register data – useful for diagnosing the cause of the following system errors: CATERR, ERR[2], SMI timeout, PERR, and SERR. The debug data is saved and timestamped for the last 3 occurrences of the error conditions.
 - PCI error registers
 - MSR registers
 - 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.

6.10.15.1 Output Data Format

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

6.10.15.2 Output Data Availability

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

6.10.15.3 Output Data Categories

The following tables list the data to be provided in the diagnostic output. For items in **Table 25**, this data is collected on detection of CATERR, ERR2, PERR, SERR, and SMI timeout. The data in Table 26 is accumulated for the three most recent overall errors.

Table 25. Diagnostic Data

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

Table 26. 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

6.10.16 Data Center Management Interface (DCMI)

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

6.10.17 Lightweight Directory Authentication Protocol (LDAP)

The Lightweight Directory Access Protocol (LDAP) is an application protocol supported by the BMC for the purpose of authentication and authorization. The BMC user connects with an LDAP server for login authentication. This is only supported for non-IPMI logins including the embedded web UI and SM-CLP. IPMI users/passwords and sessions are not supported over LDAP. LDAP can be configured (IP address of LDAP server, port, and 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. 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 27. Intel® Remote Management Module 4 (RMM4) Options

Intel Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE	Intel® Remote Management Module 4 Lite	RMM4 Lite Activation Key	Enables KVM & 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 & media Redirection with 1Gbe NIC.

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




Figure 33. Intel® RMM4 Lite Activation Key Installation




Figure 34. Intel® RMM4 Dedicated Management NIC Installation

Table 28. 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 & media redirection from onboard NIC
Intel® Remote Management Module 4 Package includes 2 modules – 1 - key for advance features 2 - Dedicated NIC (1 Gbe) for management	Dedicated NIC for management traffic. Higher bandwidth connectivity for KVM & 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 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.

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.

7.1.7 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.

7.1.8 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

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

8.1 Power Connectors

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

8.1.1 Main Power

Main server board power is supplied from two slot connectors, which allow for one or two (redundant) power supplies to dock directly to the server board. Each connector is labeled as “MAIN PWR 1” or “MAIN PWR 2” on the server board. The server board provides no option to support power supplies with cable harnesses. In a redundant power supply configuration, a failed power supply module is hot-swappable. The following tables provide the pin-out for both “MAIN PWR 1” and “MAIN PWR 2” connectors.

Table 29. Main Power (Slot 1) Connector Pin-out (“MAIN PWR 1”)

Signal Description	Pin #	Pin#	Signal Description
GROUND	B1	A1	GROUND
GROUND	B2	A2	GROUND
GROUND	B3	A3	GROUND
GROUND	B4	A4	GROUND
GROUND	B5	A5	GROUND
GROUND	B6	A6	GROUND
GROUND	B7	A7	GROUND
GROUND	B8	A8	GROUND
GROUND	B9	A9	GROUND
P12V	B10	A10	P12V
P12V	B11	A11	P12V
P12V	B12	A12	P12V
P12V	B13	A13	P12V
P12V	B14	A14	P12V
P12V	B15	A15	P12V
P12V	B16	A16	P12V
P12V	B17	A17	P12V
P12V	B18	A18	P12V
P3V3_AUX: PD_PS1_FRU_A0	B19	A19	SMB_PMBUS_DATA_R
P3V3_AUX: PD_PS1_FRU_A1	B20	A20	SMB_PMBUS_CLK_R
P12V_STBY	B21	A21	FM_PS_EN_PSU_N
FM_PS_CR1	B22	A22	IRQ_SML1_PMBUS_ALERTR2_N
P12V_SHARE	B23	A23	ISENSE_P12V_SENSE_RTN
TP_1_B24	B24	A24	ISENSE_P12V_SENSE
FM_PS_COMPATIBILITY_BUS	B25	A25	PWRGD_PS_PWROK

Table 30. Main Power (Slot 2) Connector Pin-out ("MAIN PWR 2")

Signal Description	Pin #	Pin#	Signal Description
GROUND	B1	A1	GROUND
GROUND	B2	A2	GROUND
GROUND	B3	A3	GROUND
GROUND	B4	A4	GROUND
GROUND	B5	A5	GROUND
GROUND	B6	A6	GROUND
GROUND	B7	A7	GROUND
GROUND	B8	A8	GROUND
GROUND	B9	A9	GROUND
P12V	B10	A10	P12V
P12V	B11	A11	P12V
P12V	B12	A12	P12V
P12V	B13	A13	P12V
P12V	B14	A14	P12V
P12V	B15	A15	P12V
P12V	B16	A16	P12V
P12V	B17	A17	P12V
P12V	B18	A18	P12V
P3V3_AUX: PU_PS2FRU_A0	B19	A19	SMB_PMBUS_DATA_R
P3V3_AUX: PD_PS2_FRU_A1	B20	A20	SMB_PMBUS_CLK_R
P12V_STBY	B21	A21	FM_PS_EN_PSU_N
FM_PS_CR1	B22	A22	IRQ_SML1_PMBUS_ALERTR3_N
P12V_SHARE	B23	A23	ISENSE_P12V_SENSE_RTN
TP_2_B24	B24	A24	ISENSE_P12V_SENSE
FM_PS_COMPATIBILITY_BUS	B25	A25	PWRGD_PS_PWROK

8.1.2 Riser Card Power Connectors

The server board includes two white 2x2-pin power connectors that provide supplemental power to high power PCIe x16 add-in cards (Video or GPGPU) that have power requirements that exceed the 75W maximum power supplied by the PCIe x16 riser slot. A cable from this connector may be routed to a power connector on the given add-in card. Maximum power draw for each connector is 225W, but is also limited by available power provided by the power supply and the total power draw of the rest of the system. A power budget for the complete system should be performed to determine how much supplemental power is available to support any high power add-in cards.

Table 31. Riser Slot Power Pin-out ("OPT_12V_PWR_1" & " OPT_12V_PWR_2")

Signal Description	Pin#	Pin#	Signal Description
P12V	3	1	GROUND
P12V	4	2	GROUND

8.1.3 Hot Swap Backplane Power Connector

The server board includes one 8-pin power connector that can be cabled to provide power for hot swap backplanes. On the server board, this connector is labeled as "HSBP PWR". The following table provides the pin-out for this connector.

Table 32. Hot Swap Backplane Power Connector Pin-out ("HSBP_PWR")

Signal Description	Pin#	Pin#	Signal Description
P12V_240VA1	5	1	GROUND
P12V_240VA1	6	2	GROUND
P12V_240VA2	7	3	GROUND
P12V_240VA2	8	4	GROUND

8.1.4 Peripheral Drive Power Connector

The server board includes one 6-pin power connector intended to provide power for peripheral devices such as Optical Disk Drives (ODD) and/or Solid State Devices (SSD). On the server board this connector is labeled as "ODD/SSD_PWR". The following table provides the pin-out for this connector.

Table 33. Peripheral Drive Power Connector Pin-out ("ODD/SSD_PWR")

Signal Description	Pin#	Pin#	Signal Description
P12V	4	1	P5V
P3V3	5	2	P5V
GROUND	6	3	GROUND

8.2 Front Panel Headers and Connectors

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

8.2.1 Front Panel Support

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

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

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

Table 34. SSI Front Panel Header Pin-out ("Front Panel")

Signal Description	Pin#	Pin#	Signal Description
P3V3_AUX	1	2	P3V3_AUX
KEY		4	P5V_STBY
FP_PWR_LED_BUF_R_N	5	6	FP_ID_LED_BUF_R_N
P3V3	7	8	FP_LED_STATUS_GREEN_R_N
LED_HDD_ACTIVITY_R_N	9	10	FP_LED_STATUS_AMBER_R_N
FP_PWR_BTN_N	11	12	LED_NIC_LINK0_ACT_FP_N
GROUND	13	14	LED_NIC_LINK0_LNKUP_FP_N
FP_RST_BTN_R_N	15	16	SMB_SENSOR_3V3STBY_DATA_R0
GROUND	17	18	SMB_SENSOR_3V3STBY_CLK
FP_ID_BTN_R_N	19	20	FP_CHASSIS_INTRUSION
PU_FM_SIO_TEMP_SENSOR	21	22	LED_NIC_LINK1_ACT_FP_N
FP_NMI_BTN_R_N	23	24	LED_NIC_LINK1_LNKUP_FP_N
KEY			KEY
LED_NIC_LINK2_ACT_FP_N	27	28	LED_NIC_LINK3_ACT_FP_N
LED_NIC_LINK2_LNKUP_FP_N	29	30	LED_NIC_LINK3_LNKUP_FP_N

8.2.1.1 Power/Sleep Button and LED Support

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

Table 35. Power/Sleep LED Functional States

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

8.2.1.2 System ID Button and LED Support

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

8.2.1.3 System Reset Button Support

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

8.2.1.4 NMI Button Support

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

The following actions cause the BMC to generate an NMI pulse:

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

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

Table 36. NMI Signal Generation and Event Logging

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

8.2.1.5 NIC Activity LED Support

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

8.2.1.6 Hard Drive Activity LED Support

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

8.2.1.7 System Status LED Support

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

Table 37. System Status LED State Definitions

Color	State	Criticality	Description
Off	System is not operating	Not ready	<ol style="list-style-type: none"> 1. System is powered off (AC and/or DC). 2. System is in EuP Lot6 Off Mode. 3. System is in S5 Soft-Off State. 4. System is in S4 Hibernate Sleep State.
Green	Solid on	Ok	Indicates that the System is running (in S0 State) and its status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running.
Green	~1 Hz blink	Degraded - system is operating in a degraded state although still functional, or system is operating in a redundant state but with an impending failure	<p>System degraded:</p> <p>Redundancy loss, such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities.</p> <p>Fan warning or failure when the number of fully operational fans is more than minimum number needed to cool the system.</p> <p>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.</p> <p>Power supply predictive failure occurred while redundant power supply configuration was present.</p>

Color	State	Criticality	Description
		warning	<p>Unable to use all of the installed memory (one or more DIMMs failed/disabled but functional memory remains available)</p> <p>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.</p> <p>Uncorrectable memory error has occurred in memory Mirroring Mode, causing Loss of Redundancy.</p> <p>Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in fully redundant RAS Mirroring Mode. BMC executing in uBoot. (Indicated by Chassis ID blinking at Blinking at 3Hz). System in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux image into flash BMC booting Linux. (Indicated by Chassis ID solid ON). System in degraded state (no manageability). Control has been passed from BMC uBoot to BMC Linux itself. It will be in this state for ~10--20 seconds. BMC Watchdog has reset the BMC.</p> <p>Power Unit sensor offset for configuration error is asserted.</p> <p>HDD HSC is off-line or degraded.</p>
Amber	~1 Hz blink	Non-critical - System is operating in a degraded state with an impending failure warning, although still functioning	<p>Non-fatal alarm – system is likely to fail:</p> <p>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.</p> <p>VRD Hot asserted.</p> <p>Minimum number of fans to cool the system not present or failed</p> <p>Hard drive fault</p> <p>Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present)</p> <p>In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window</p> <p>Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in a non-redundant mode</p>
Amber	Solid on	Critical, non-recoverable – System is halted	<p>Fatal alarm – system has failed or shutdown:</p> <p>CPU CATERR signal asserted</p> <p>MSID mismatch detected (CATERR also asserts for this case).</p> <p>CPU 1 is missing</p> <p>CPU Thermal Trip</p> <p>No power good – power fault</p> <p>DIMM failure when there is only 1 DIMM present and hence no good memory present¹.</p> <p>Runtime memory uncorrectable error in non-redundant mode.</p> <p>DIMM Thermal Trip or equivalent</p> <p>SSB Thermal Trip or equivalent</p> <p>CPU ERR2 signal asserted</p> <p>BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition)</p> <p>Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition)</p> <p>240VA fault</p> <p>Fatal Error in processor initialization:</p> <p>Processor family not identical</p> <p>Processor model not identical</p> <p>Processor core/thread counts not identical</p> <p>Processor cache size not identical</p> <p>Unable to synchronize processor frequency</p> <p>Unable to synchronize QPI link frequency</p>

8.2.2 Front Panel USB Connector

The server board includes a 10-pin connector, that when cabled, can provide up to two USB ports to a front panel. On the server board the connector is labeled “FP USB” and is located on the front edge of the board. The following table provides the connector pin-out.

Table 38. Front Panel USB Connector Pin-out ("FP USB")

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

8.2.3 Front Panel Video Connector

The server board includes a 14-pin header, that when cabled, can provide an alternate video connector to the front panel. On the server board this connector is labeled “FP VIDEO” and is located on the front edge of the board. When a monitor is attached to the front panel video connector, the external video connector located on the back edge of the board is disabled. The following table provides the pin-out for this connector.

Table 39. Front Panel Video Connector Pin-out ("FP VIDEO")

Signal Description	Pin#	Pin#	Signal Description
V_IO_FRONT_R_CONN	1	2	GROUND
V_IO_FRONT_G_CONN	3	4	GROUND
V_IO_FRONT_B_CONN	5	6	GROUND
V_BMC_GFX_FRONT_VSYN	7	8	GROUND
V_BMC_GFX_FRONT_HSYN	9		KEY
V_BMC_FRONT_DDC_SDA_CONN	11	12	V_FRONT_PRES_N
V_BMC_FRONT_DDC_SCL_CONN	13	14	P5V_VID_CONN_FNT

8.2.4 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 40. Intel Local Control Panel Connector Pin-out ("LCP")

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

8.3 On-Board Storage Connectors

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

8.3.1 Single Port SATA Only Connectors

The server board includes two white single port SATA only connectors capable of transfer rates of up to 6Gb/s. On the server board these connectors are labeled as “SATA 0” and “SATA 1”. The following table provides the pin-out for both connectors.

Table 41. Single Port AHCI SATA Controller Connector Pin-out ("SATA 0" & "SATA 1")

Signal Description	Pin#
GROUND	1
SATA_TXP	2
SATA_TXN	3
GROUND	4
SATA_RXN	5
SATA_RXP	6
P5V**	7

**SATA-DOM power support

8.3.2 Multiport Mini-SAS/SATA Connectors

The server board includes two 40-pin high density multiport mini-SAS/SATA connectors. On the server board, these connectors are labeled as “SCU_0 (0-3)” supporting the chipset embedded SCU 0 controller, and “SCU_1 (4-7)”, supporting the embedded SCU 1 controller. Both connectors can support up to four SATA or SAS ports each. By default, only the connector labeled “SCU_0 (0-3)” is enabled and has support for up to four SATA ports capable of transfer rates of up to 6Gb/s. The connector labeled “SCU_1 (4-7)” is only enabled when an optional 8-port SAS or SATA Intel® RAID C600 Upgrade Key is installed. See Table 16. Intel® RAID C600 Upgrade Key Options for a complete list of supported storage upgrade keys. The following tables provide the pin-out for each connector.

Table 42. Multiport SAS/SATA Connector Pin-out ("SCU_0 (0-3)")

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

Table 43. Multiport SAS/SATA Connector Pin-out ("SCU_1 (4-7)")

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

8.3.3 Internal Type-A USB Connector

The server board includes one internal Type-A USB connector labeled "USB 2" and is located near the back edge of the board next to the Riser 1 slot. The following table provides the pin-out for this connector.

Table 44. Internal Type-A USB Connector Pin-out ("USB 2")

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

8.3.4 Internal 2mm Low Profile eUSB SSD Connector

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

Table 45. Internal eUSB Connector Pin-out ("eUSB SSD")

Signal Description	Pin#	Pin#	Signal Description
P5V	1	2	NOT USED
USB2_P0_DN	3	4	NOT USED
USB2_P0_DP	5	6	NOT USED
GROUND	7	8	NOT USED
NOT USED	9	10	LED_HDD_ACT_N

8.4 Fan Connectors

The server board provides support for six system fans. Each 10-pin connector is monitored and controlled by on-board platform management. On the server board, each system fan connector is labeled “SYS_FAN #”, where # = 1 thru 6. The following table provides the pin-out for all six system fan connectors.

Table 46. System Fan Connector Pin-out ("SYS_FAN #")

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_BMC_PWM0_R_BUF	2	FAN_BMC_PWM1_R_BUF	2	FAN_BMC_PWM2_R_BUF	2
P12V_FAN	3	P12V_FAN	3	P12V_FAN	3
P12V_FAN	4	P12V_FAN	4	P12V_FAN	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_SYS0_PRSENT_N	8	FAN_SYS1_PRSENT_N	8	FAN_SYS2_PRSENT_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		SYS_FAN 6	
Signal Description	Pin#	Signal Description	Pin#	Signal Description	Pin#
FAN_TACH7_IN	1	FAN_TACH9_IN	1	FAN_TACH11_IN	1
FAN_BMC_PWM3_R_BUF	2	FAN_BMC_PWM4_R_BUF	2	FAN_BMC_PWM5_R_BUF	2
P12V_FAN	3	P12V_FAN	3	P12V_FAN	3
P12V_FAN	4	P12V_FAN	4	P12V_FAN	4
FAN_TACH6_IN	5	FAN_TACH8_IN	5	FAN_TACH10_IN	5
GROUND	6	GROUND	6	GROUND	6
GROUND	7	GROUND	7	GROUND	7
FAN_SYS3_PRSENT_N	8	FAN_SYS4_PRSENT_N	8	FAN_SYS5_PRSENT_N	8
LED_FAN_FAULT3_R	9	LED_FAN_FAULT4_R	9	LED_FAN_FAULT5_R	9
LED_FAN3	10	LED_FAN4	10	LED_FAN5	10

8.5 Serial Port Connectors

The server board includes two serial port connectors.

Serial-B is an internal 10-pin DH-10 connector labeled “Serial_B” and has the following pin-out.

Table 47. Serial-B Connector Pin-out

Signal Description	Pin#	Pin#	Signal Description
DCD	1	2	DSR
SIN	3	4	RTS
SOUT	5	6	CTS
DTR	7	8	RI
GROUND	9		KEY

Serial-A is an external RJ45 type connector and has the following pin-out configuration.




Figure 35. Serial-A RJ45 connector pin-out**Table 48. Serial A Connector Pin-out**

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

** Pin 7 of the RJ45 Serial A connector is configurable to support either a DSR (Default) signal or a DCD signal by switching jumper locations on the 3-pin jumper block labeled “J3A2” on the server board which is located next to the stacked external USB connectors near the back edge of the board.

Serial-A configuration jumper block setting:

Signal	Pins
DSR (Default)	1-2
DCD	2-3

**Figure 36. Serial A Configuration Jumper Block Location**

8.6 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 49. 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 50. Hard Drive Activity Header Pin-out ("HDD_LED")

Signal Description	Pin#
LED_HDD_ACT_N	1
TP_LED_HDD_ACT	2

9. Reset and Recovery Jumpers

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

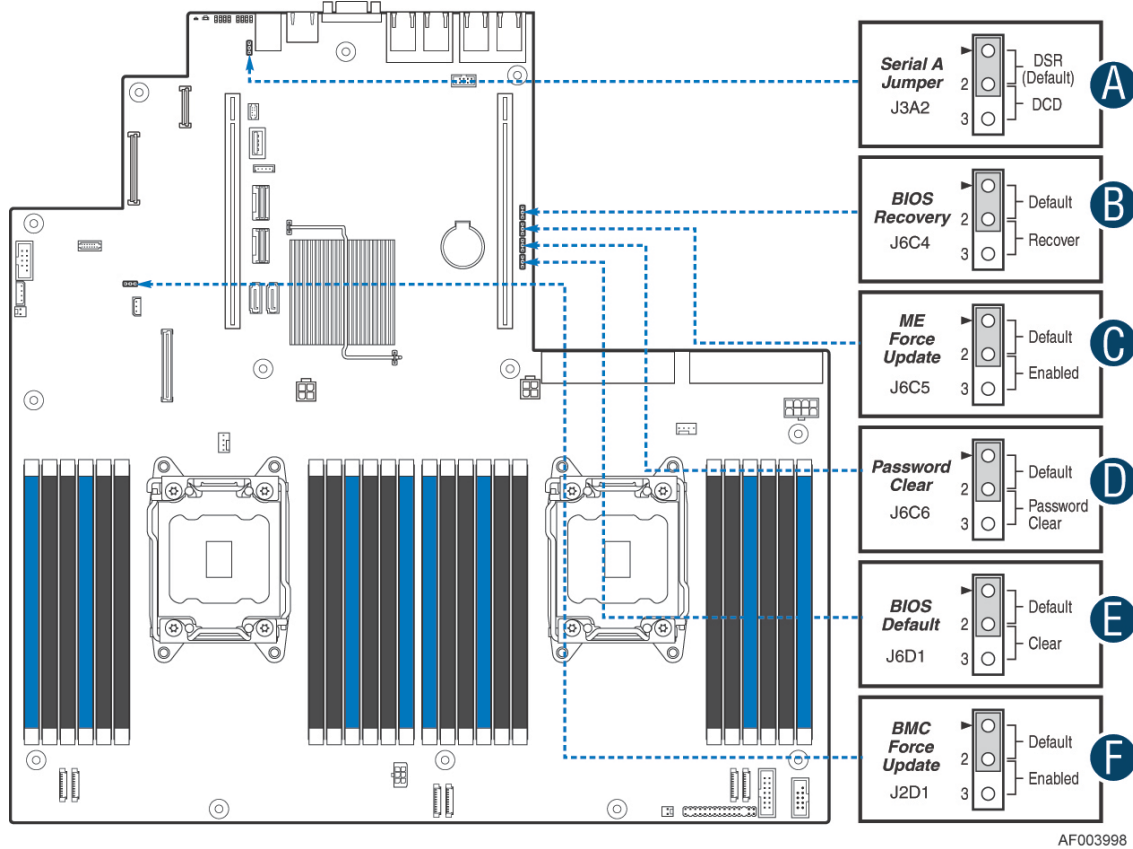


Figure 37. Reset and Recovery Jumper Block Location

Note:

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

A – Serial Port ‘A’ Configuration Jumper – See section 8.5.

B – 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 Web site, 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

Note: It may be necessary to edit the STARTUP.NSH file to ensure the ####REC.CAP file is called in the shell script.

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.

C – 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 web site.

1. Turn off the system and remove power cords.
2. Remove Riser Card Assembly #2.
3. Move the *ME FRC UPD* Jumper (J6C5) from the default (pins 1 and 2) operating position to the Force Update position (pins 2 and 3).
4. Re-attach system power cords.
5. Power on the system.

Note: System Fans will boost and the BIOS Error Manager should report an 83A0 error code (ME in recovery mode).
6. Boot to the EFI shell and update the ME firmware using the "MEComplete.cap" file using the following command: `iflash32 /u /ni Mecomplete.cap`
7. When update has successfully completed, power off system
8. Remove AC power cords
9. Move ME FRC UPD jumper back to the default position

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

10. Install PCI Riser
11. Install AC power cords
12. Power on system

D – 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. Open the chassis and remove the Riser #2 assembly
3. Move jumper (J6C6) from the default (pins 1 and 2) operating position to the password clear position (pins 2 and 3).
4. Close the server chassis and reattach the power cords
5. Power up the server and wait until POST completes.
6. Power down the server and unplug the power cords.
7. Open the chassis, remove the Riser #2 assembly, and move the jumper back to the default position (covering pins 1 and 2).
8. Reinstall the Riser #2 assembly
9. Close the server chassis and reattach the power cords.
10. Power up the server

E – 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. Open the chassis and remove the Riser #2 assembly
3. Move BIOS DFLT jumper (J6D1) from the default (pins 1 and 2) position to the Set BIOS Defaults position (pins 2 and 3)
4. Wait 5 seconds then move the jumper back to the default position of pins 1 and 2
5. Install riser card assembly
6. Install Power Cords
7. Power on system


F – 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 causes the BMC to abort its normal boot process and stay in the boot loader without executing any Linux code.

The jumper is normally in the de-asserted position. The system must be completely powered off (A/C power removed) before the jumper is moved. After power is re-applied and the firmware update is complete, the system must be powered off again and the jumper must be returned to the de-asserted position before normal operation can begin.


10. 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 LED.



AF003994

Label	Description	Label	Description
A	System ID	I	System Fan #3 Fan Fault
B	System Status	J	Memory Fault
C	POST Code Diagnostics	K	System Fan #2 Fan Fault
D	12V Stand-by Power Present	L	System Fan #1 Fan Fault
E	CPU-2 Fault	M	CPU-1 Fault
F	System Fan #6 Fan Fault	N	CATERR
G	System Fan #5 Fan Fault	O	System Power Good
H	System Fan #4 Fan Fault		

Figure 38. On-Board Diagnostic LED Placement**Figure 39. 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 51. SystemStatus LED State Definitions

Color	State	Criticality	Description
Off	System is	Not ready	1. System is powered off (AC and/or DC).

Color	State	Criticality	Description
	not operating		<ol style="list-style-type: none"> 2. System is in EuP Lot6 Off Mode. 3. System is in S5 Soft-Off State. 4. System is in S4 Hibernate Sleep State.
Green	Solid on	Ok	Indicates that the System is running (in S0 State) and its status is 'Healthy'. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running.
Green	~1 Hz blink	Degraded - system is operating in a degraded state although still functional, or system is operating in a redundant state but with an impending failure warning	<p>System degraded:</p> <p>Redundancy loss, such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities.</p> <p>Fan warning or failure when the number of fully operational fans is more than minimum number needed to cool the system.</p> <p>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.</p> <p>Power supply predictive failure occurred while redundant power supply configuration was present.</p> <p>Unable to use all of the installed memory (one or more DIMMs failed/disabled but functional memory remains available)</p> <p>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.</p> <p>Uncorrectable memory error has occurred in memory Mirroring Mode, causing Loss of Redundancy.</p> <p>Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in fully redundant RAS Mirroring Mode.</p> <p>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</p> <p>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.</p> <p>BMC Watchdog has reset the BMC.</p> <p>Power Unit sensor offset for configuration error is asserted.</p> <p>HDD HSC is off-line or degraded.</p>
Amber	~1 Hz blink	Non-critical - System is operating in a degraded state with an impending failure warning, although still functioning	<p>Non-fatal alarm – system is likely to fail:</p> <p>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.</p> <p>VRD Hot asserted.</p> <p>Minimum number of fans to cool the system not present or failed</p> <p>Hard drive fault</p> <p>Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present)</p> <p>In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window</p> <p>Correctable memory error threshold has been reached for a failing DDR3 DIMM when the system is operating in a non-redundant mode</p>
Amber	Solid on	Critical, non-recoverable – System is halted	<p>Fatal alarm – system has failed or shutdown:</p> <p>CPU CATERR signal asserted</p> <p>MSID mismatch detected (CATERR also asserts for this case).</p> <p>CPU 1 is missing</p> <p>CPU Thermal Trip</p> <p>No power good – power fault</p> <p>DIMM failure when there is only 1 DIMM present and hence no good memory present.</p> <p>Runtime memory uncorrectable error in non-redundant mode.</p> <p>DIMM Thermal Trip or equivalent</p> <p>SSB Thermal Trip or equivalent</p> <p>CPU ERR2 signal asserted</p> <p>BMC\Video memory test failed. (Chassis ID shows blue/solid-on for this condition)</p> <p>Both uBoot BMC FW images are bad. (Chassis ID shows blue/solid-on for this condition)</p> <p>240VA fault</p> <p>Fatal Error in processor initialization:</p> <p>Processor family not identical</p>

Color	State	Criticality	Description
			Processor model not identical Processor core/thread counts not identical Processor cache size not identical Unable to synchronize processor frequency Unable to synchronize QPI link frequency

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 52. BMC Boot/Reset Status LED Indicators

BMC Boot/Reset State	Chassis ID LED	Status LED	Comment
BMC/Video memory test failed	Solid Blue	Solid Amber	Nonrecoverable condition. Contact your Intel® representative for information on replacing this motherboard.
Both Universal Bootloader (u-Boot) images bad	Solid Blue	Solid Amber	Nonrecoverable 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 12 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 12 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.

10.8 CPU Fault LEDs

The server board includes a CPU fault LED for each CPU socket. The CPU Fault LED is lit if there is an MSID mismatch error is detected (that is, CPU power rating is incompatible with the board).

10.9 System Power Good LED

The Power Good LED is illuminated when system power is turned on and DC power to the server board is within acceptable limits.

10.10 CATERR LED

For processors that support the Intel® Quick Path Interconnect interface, the CATERR# signal is used to indicate that a catastrophic error condition has been experienced by the processor. Should such an error occur, the CATERR LED on the server board will illuminate.

11. Power Supply Specification Guidelines

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

Note: The power supply data provided in this section is for reference purposes only. It reflects Intel's own DC power out requirements for a 750W power supply as used in an Intel designed 2U server platform. The intent of this section is to provide customers with a guide to assist in defining and/or selecting a power supply for custom server platform designs that utilize the server boards detailed in this document.

11.1 Power Supply DC Output Connector

The server board includes two main power slot connectors allowing for power supplies to attach directly to the server board. Power supplies must utilize a card edge output connection for power and signal that is compatible with a 2x25 Power Card Edge connector (equivalent to 2x25 pin configuration of the FCI power card connector 10035388-102LF).

Table 53. Power Supply DC Power Output Connector Pinout

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

11.2 Power Supply DC Output Specification

11.2.1 Output Power/Currents

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

Table 54. Minimum Load Ratings

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

Notes:

- 12Vstby must provide 4.0A with two power supplies in parallel. The Fan may start to work when stby current >1.5A
- Peak combined power for all outputs shall not exceed 850W.
- Length of time peak power can be supported is based on thermal sensor and assertion of the SMBAlert# signal. Minimum peak power duration shall be 20 seconds without asserting the SMBAlert# signal at maximum operating temperature.

11.2.2 Standby Output

The 12VSB output shall be present when an AC input greater than the power supply turn on voltage is applied. There should be load sharing in the standby rail. Two PSU modules should be able to support 4A standby current.

11.2.3 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	- 5%/+5%	+11.40	+12.00	+12.60	V _{rms}
+12V stby	- 5%/+5%	+11.40	+12.00	+12.60	V _{rms}

11.2.4 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	Load Slew Rate	Test capacitive Load
+12VSB	1.0A	0.25 A/ μ sec	20 μ F
+12V	60% of max load	0.25 A/ μ sec	2000 μ F

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

11.2.5 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
+12VSB	20	3100	μF
+12V	500	25000	μF

11.2.6 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 mΩ. This path may be used to carry DC current.

11.2.7 Closed loop stability

The power supply shall be unconditionally stable under all line/load/transient load conditions including specified capacitive load ranges. A minimum of: **45 degrees phase margin** and **-10dB-gain margin** is required. Closed-loop stability must be ensured at the maximum and minimum loads as applicable.

11.2.8 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.

11.2.9 Common Mode Noise

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

11.2.10 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.

11.2.11 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.

11.2.12 Hot Swap Requirements

Hot swapping a power supply is the process of inserting and extracting a power supply from an operating power system. During this process the output voltages shall remain within the limits with the capacitive load specified. The hot swap test must be conducted when the system is operating under static, dynamic, and zero loading conditions.

11.2.13 Forced Load Sharing

The +12V output will have active load sharing. The output will share within 10% at full load. The failure of a power supply should not affect the load sharing or output voltages of the other supplies still operating. The supplies must be able to load share in parallel and operate in a hot-swap/redundant **1+1** configurations. The 12VSBoutput is not required to actively share current between power supplies (passive sharing). The

12VSB output of the power supplies are connected together in the system so that a failure or hot swap of a redundant power supply does not cause these outputs to go out of regulation in the system.

11.2.14 Ripple/Noise

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

Table 58. Ripples and Noise

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

11.2.15 Timing Requirements

These are the timing requirements for the power supply operation. The output voltages must rise from 10% to within regulation limits (T_{vout_rise}) within 5 to 70ms. For 12VSB, it is allowed to rise from 1.0 to 25ms. **All outputs must rise monotonically.** The following table 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.

Table 59. Timing Requirements

Item	Description	MIN	MAX	UNITS
T_{vout_rise}	Output voltage rise time	5.0 *	70 *	ms
Tsb_on_delay	Delay from AC being applied to 12VSB being within regulation.		1500	ms
T ac_on_delay	Delay from AC being applied to all output voltages being within regulation.		3000	ms
Tvout_holdup	Time 12V output voltage stay within regulation after loss of AC.	13		ms
Tpwok_holdup	Delay from loss of AC to de-assertion of PWOK	12		ms
Tpson_on_delay	Delay from PSON# active to output voltages within regulation limits.	5	400	ms
Tpson_pwok	Delay from PSON# deactivate to PWOK being de-asserted.		5	ms
Tpwok_on	Delay from output voltages within regulation limits to PWOK asserted at turn on.	100	500	ms
Tpwok_off	Delay from PWOK de-asserted to output voltages dropping out of regulation limits.	1		ms
Tpwok_low	Duration of PWOK being in the de-asserted state during an off/on cycle using AC or the PSON signal.	100		ms
Tsb_vout	Delay from 12VSB being in regulation to O/Ps being in regulation at AC turn on.	50	1000	ms
T12VSB_holdup	Time the 12VSB output voltage stays within regulation after loss of AC.	70		ms

* The 12VSB output voltage rise time shall be from 1.0ms to 25ms

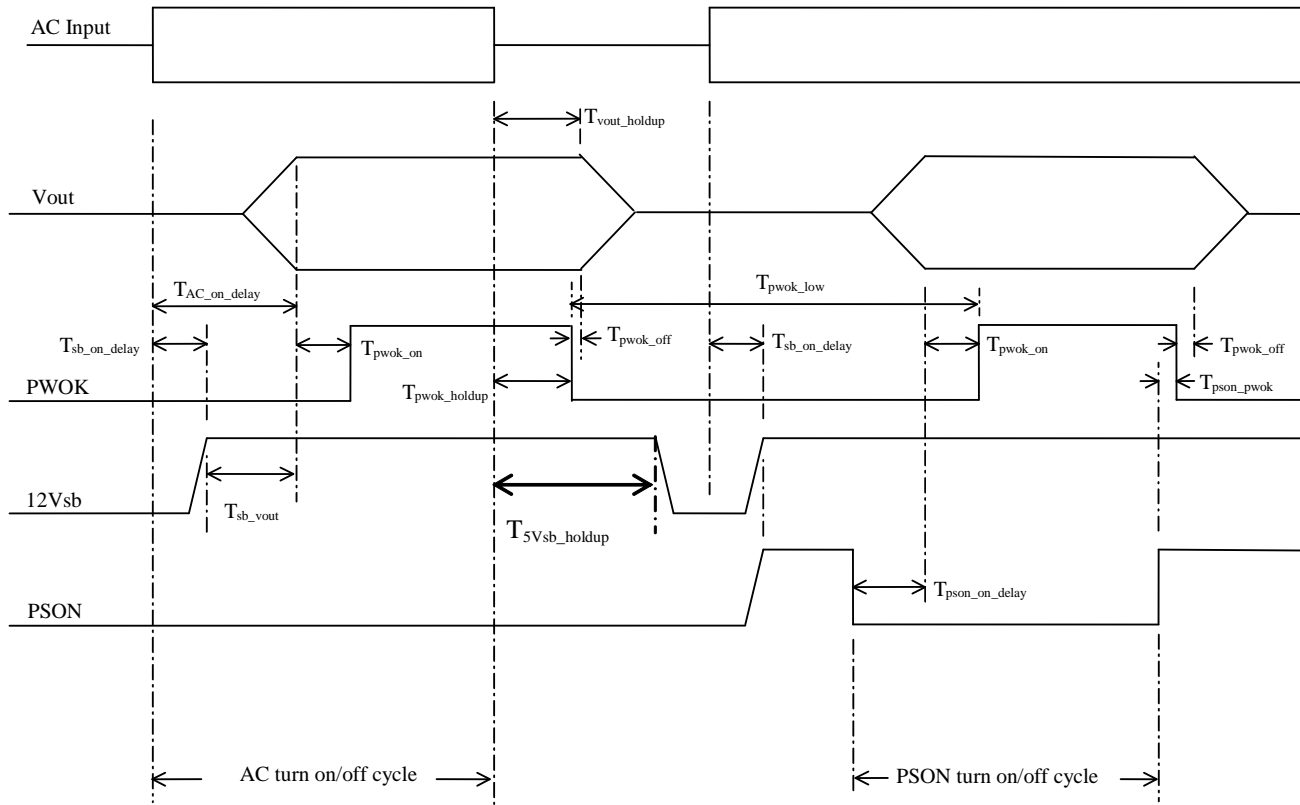


Figure 40. Turn On/Off Timing (Power Supply Signals)

12. BIOS Setup Utility

The BIOS Setup utility is a text-based utility that allows the user to configure the system and view current settings and environment information for the platform devices. The Setup utility controls the platform's built-in devices, the boot manager, and error manager.

The BIOS Setup interface consists of a number of pages or screens. Each page contains information or links to other pages. The advanced tab in Setup displays a list of general categories as links. These links lead to pages containing a specific category's configuration.

The following sections describe the look and behavior for the platform setup utility.

12.1 BIOS Setup Operation

The BIOS Setup Utility has the following features:

- Localization – The Intel® Server Board BIOS is only available in English. However, BIOS Setup uses the Unicode standard and is capable of displaying data and input in Setup fields in all languages currently included in the Unicode standard.
- Console Redirection – BIOS Setup is functional from Console Redirection over various terminal emulation standards. When Console Redirection is enabled, the POST display out is in purely Text Mode due to Redirection data transfer in a serial port data terminal emulation mode. This may limit some functionality for compatibility, for example, usage of colors or some keys or key sequences or support of pointing devices.
- Setup screens are designed to be displayable in an 80-character x 24-line format in order to work with Console Redirection, although that screen layout should display correctly on any format with longer lines or more lines on the screen.
- Password protection – BIOS Setup may be protected from unauthorized changes by setting an Administrative Password in the Security screen. When an Administrative Password has been set, all selection and data entry fields in Setup (except System Time and Date) are grayed out and cannot be changed unless the Administrative Password has been entered.
 - **Note:** If an Administrative Password has not been set, anyone who boots the system to Setup has access to all selection and data entry fields in Setup and can change any of them.

12.1.1 Entering BIOS Setup

To enter the BIOS Setup using a keyboard (or emulated keyboard), press the <F2> function key during boot time when the OEM or Intel Logo Screen or the POST Diagnostic Screen is displayed.

The following message is displayed on the Diagnostic Screen or under the Quiet Boot Logo Screen:

Press <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot

Note: With a USB keyboard, it is important to wait until the BIOS “discovers” the keyboard and beeps – until the USB Controller has been initialized and the USB keyboard activated.

When the Setup Utility is entered, the Main screen is displayed initially. However, serious errors cause the system to display the Error Manager screen instead of the Main screen.

It is also possible to cause a boot directly to Setup using an IPMI 2.0 command “Get/Set System Boot Options”. For details on that capability, see the explanation in the IPMI description.

12.1.2 Setup Navigation Keyboard Commands

The bottom right portion of the Setup screen provides a list of commands that are used to navigate through the Setup utility. These commands are displayed at all times.

Each Setup menu page contains a number of features. Each feature is associated with a value field, except those used for informative purposes. Each value field contains configurable parameters. Depending on the security option chosen and in effect by the password, a menu feature's value may or may not be changed. If a value cannot be changed, its field is made inaccessible and appears grayed out.

Table 60. BIOS Setup: Keyboard Command Bar

Key	Option	Description
<Enter>	Execute Command	The <Enter> key is used to activate submenus when the selected feature is a submenu, or to display a pick list if a selected option has a value field, or to select a subfield for multi-valued features like time and date. If a pick list is displayed, the <Enter> key selects the currently highlighted item, undoes the pick list, and returns the focus to the parent menu.
<Esc>	Exit	The <Esc> key provides a mechanism for backing out of any field. When the <Esc> key is pressed while editing any field or selecting features of a menu, the parent menu is re-entered. When the <Esc> key is pressed in any submenu, the parent menu is re-entered. When the <Esc> key is pressed in any major menu, the exit confirmation window is displayed and the user is asked whether changes can be discarded. If "No" is selected and the <Enter> key is pressed, or if the <Esc> key is pressed, the user is returned to where they were before <Esc> was pressed, without affecting any existing settings. If "Yes" is selected and the <Enter> key is pressed, the setup is exited and the BIOS returns to the main System Options Menu screen.
↑	Select Item	The up arrow is used to select the previous value in a pick list, or the previous option in a menu item's option list. The selected item must then be activated by pressing the <Enter> key.
↓	Select Item	The down arrow is used to select the next value in a menu item's option list, or a value field's pick list. The selected item must then be activated by pressing the <Enter> key.
← →	Select Menu	The left and right arrow keys are used to move between the major menu pages. The keys have no effect if a sub-menu or pick list is displayed.
<Tab>	Select Field	The <Tab> key is used to move between fields. For example, <Tab> can be used to move from hours to minutes in the time item in the main menu.
-	Change Value	The minus key on the keypad is used to change the value of the current item to the previous value. This key scrolls through the values in the associated pick list without displaying the full list.
+	Change Value	The plus key on the keypad is used to change the value of the current menu item to the next value. This key scrolls through the values in the associated pick list without displaying the full list. On 106-key Japanese keyboards, the plus key has a different scan code than the plus key on the other keyboards, but will have the same effect.
<F9>	Setup Defaults	Pressing the <F9> key causes the following to display: <div style="border: 1px solid black; padding: 5px; text-align: center;"> Load Optimized Defaults? Yes No </div> If "Yes" is highlighted and <Enter> is pressed, all Setup fields are set to their default values. If "No" is highlighted and <Enter> is pressed, or if the <Esc> key is pressed, the user is returned to where they were before <F9> was pressed without affecting any existing field values.
<F10>	Save and Exit	Pressing the <F10> key causes the following message to display: <div style="border: 1px solid black; padding: 5px; text-align: center;"> Save configuration and reset? Yes No </div> If "Yes" is highlighted and <Enter> is pressed, all changes are saved and the Setup is exited. If "No" is highlighted and <Enter> is pressed, or the <Esc> key is pressed, the user is returned to where they were before <F10> was pressed without affecting any existing values.

12.2 BIOS Setup Utility Screens

The following sections describe the screens available in the BIOS Setup utility for the configuration of the server platform.

For each of these screens, there is an image of the screen with a list of Field Descriptions which describe the contents of each item on the screen. Each item on the screen is hyperlinked to the relevant Field Description. Each Field Description is hyperlinked back to the screen image.

There are a number of screens in the entire Setup collection. They are organized into major categories. Each category has a hierarchy beginning with a top-level screen from which lower-level screens may be selected. Each top-level screen appears as a tab, arranged across the top of the Setup screen image of all top-level screens.

There are more categories than will fit across the top of the screen, so at any given time there will be some categories which will not appear until the user has scrolled across the tabs which are present.

The categories and the screens included in each category are listed below, with links to each of the screens named.

12.2.1 Main Screen (Tab)

The Main Screen is the first screen that appears when the BIOS Setup configuration utility is entered, unless an error has occurred. If an error has occurred, the Error Manager Screen appears instead.

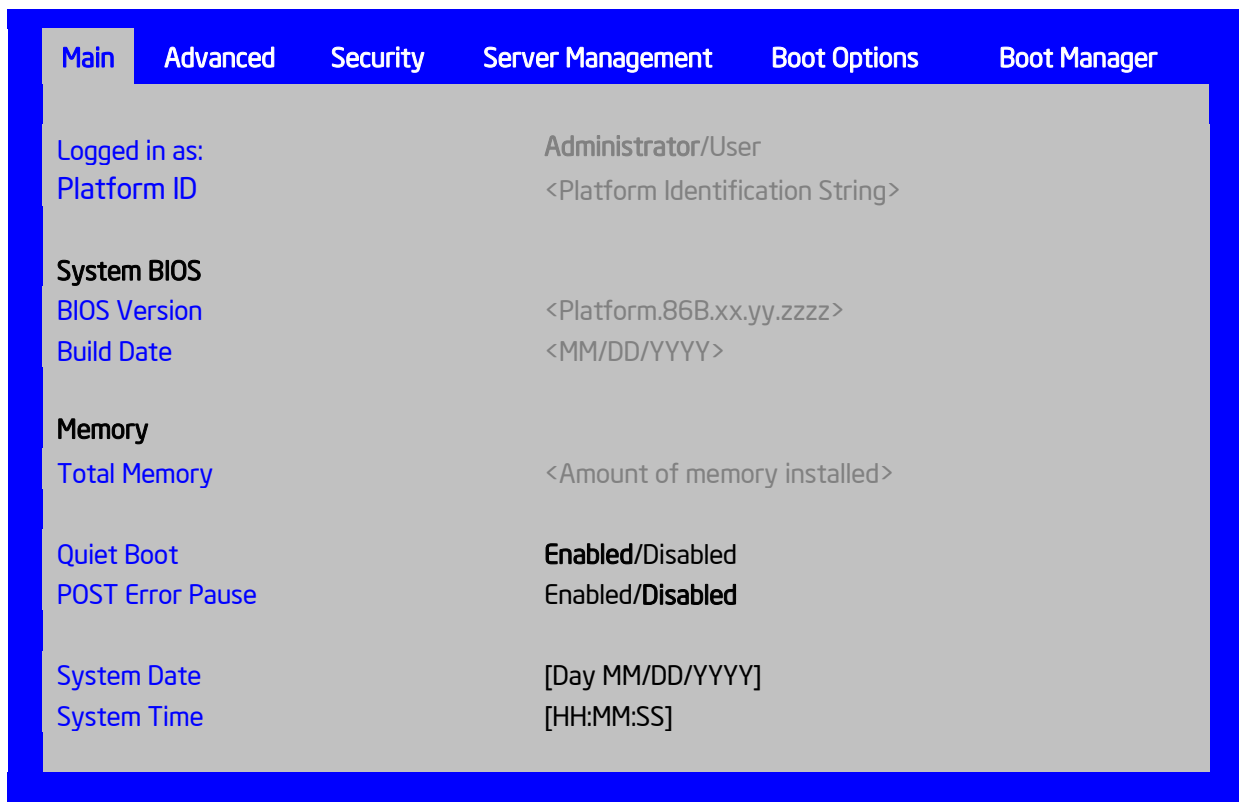


Figure 41. Main Screen

Screen Field Descriptions:

Logged in as:

Option Values: <Administrator/User>

Help Text: <None>

Comments: Information only. Displays password level that setup is running in: Administrator or User. With no passwords set, Administrator is the default mode.

[Back to \[Main Screen\]](#)

Platform ID

Option Values: < Platform ID>

Help Text: <None>

Comments: Information only. Displays the Platform ID (Board ID) for the board on which the BIOS is executing POST.

The Platform ID is limited to 8 characters, because it is also used in the ACPI Tables which have that limitation. In some cases, this means that the Platform ID is abbreviated from the marketing designation (for example, MFS2600KI is abbreviated to S2600KI).

[Back to \[Main Screen\]](#)

BIOS Version

Option Values: <Current BIOS version ID>

Help Text: <None>

Comments: Information only. The BIOS version displayed uniquely identifies the BIOS that is currently installed and operational on the board. The version information displayed is taken from the BIOS ID String, with the timestamp segment dropped off. The segments displayed are:

Platform:	Identifies that this is the correct platform BIOS
86B:	Identifies this BIOS as being an EPSD Server BIOS
xx:	Major Revision level of the BIOS
yy:	Release Revision level for this BIOS
zzzz:	Release Number for this BIOS

[Back to \[Main Screen\]](#)

Build Date

Option Values: <Date and time when the currently installed BIOS was created (built)>

Help Text: <None>

Comments: Information only. The time and date displayed are taken from the timestamp segment of the BIOS ID String.

[Back to \[Main Screen\]](#)

Total Memory

Option Values: <Amount of memory installed in the system>

Help Text: <None>

Comments: Information only. Displays the total physical memory installed in the system, in MB or GB. The term physical memory indicates the total memory discovered in the form of installed DDR3 DIMMs.

[Back to \[Main Screen\]](#)

Quiet Boot

Option Values: **Enabled**
 Disabled

Help Text:

[Enabled] – Display the logo screen during POST.

[Disabled] – Display the diagnostic screen during POST.

Comments: This field controls whether the full diagnostic information is displayed on the screen during POST. When Console Redirection is enabled, the Quiet Boot setting is disregarded and the text mode Diagnostic Screen is displayed unconditionally.

[Back to \[Main Screen\]](#)

POST Error Pause

Option Values: Enabled
 Disabled

Help Text:

[Enabled] – Go to the Error Manager for critical POST errors.

[Disabled] – Attempt to boot and do not go to the Error Manager for critical POST errors.

Comments: If enabled, the POST Error Pause option takes the system to the error manager to review the errors when major errors occur. Minor and fatal error displays are not affected by this setting.

[Back to \[Main Screen\]](#)

System Date

Option Values: <System Date initially displays the current system calendar date,
including the day of the week>

Help Text:

System Date has configurable fields for the current Month, Day, and Year.

The year must be between 2005 and 2099.

Use [Enter] or [Tab] key to select the next field.

Use [+] or [-] key to modify the selected field.

Comments: This field will initially display the current system day of week and date. It may be edited to change the system date. When the System Date is reset by the “BIOS Defaults” jumper, BIOS Recovery Flash Update, or other method, the date will be the earliest date in the allowed range – **Saturday 01/01/2005.**

[Back to \[Main Screen\]](#)

System Time

Option Values: <System Time initially displays the current system time of day, in 24-hour
format>

Help Text:

System Time has configurable fields for Hours, Minutes, and Seconds.

Hours are in 24-hour format.

Use the [Enter] or [Tab] key to select the next field.

Use the [+] or [-] key to modify the selected field.

Comments: This field will initially display the current system time (24 hour time). It may be edited to change the system time. When the System Time is reset by the “BIOS Defaults” jumper, BIOS Recovery Flash Update, or other method, the time will be the earliest time of day in the allowed range – **00:00:00** (although the time will updated beginning from when it is reset early in POST).

[Back to \[Main Screen\]](#)

12.2.2 Advanced Screen (Tab)

The Advanced screen provides an access point to configure several groups of options. On this screen, the user can select the option group to be configured. Configuration actions are performed on the selected screen, and not directly on the Advanced screen.

This screen is the same for all board series, selecting between the same groups of options, although the options for different boards are not necessarily identical.

To access this screen from the **Main** screen or other top-level “Tab” screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Advanced** screen is selected.

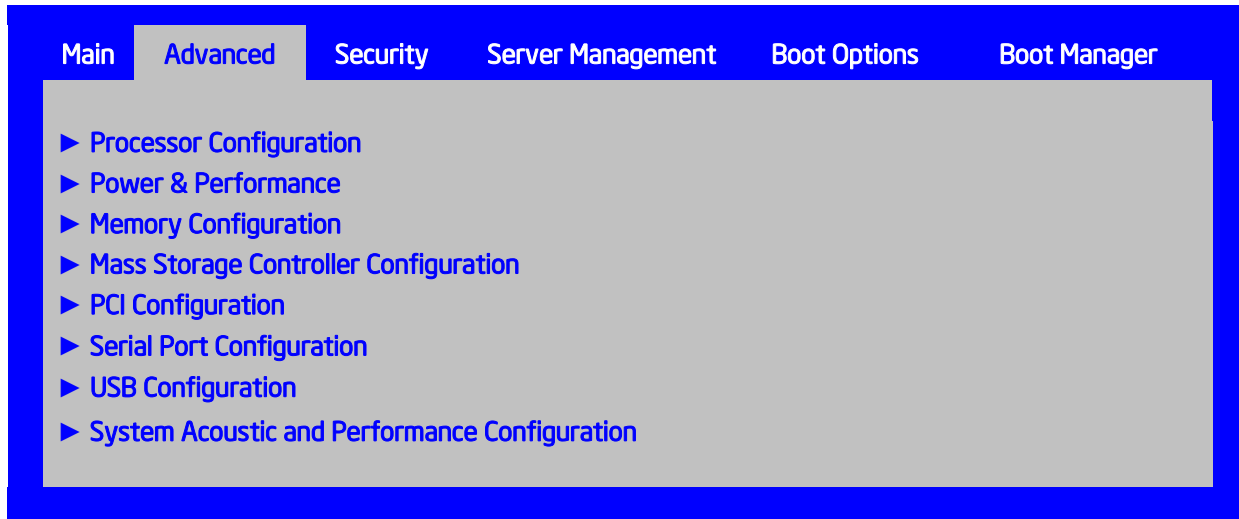


Figure 42. Advanced Screen

Screen Field Descriptions:

Processor Configuration

Option Values: <None>

Help Text: *View/Configure processor information and settings.*

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Processor Configuration** group of configuration settings.

[Back to \[Advanced Screen\]](#)

1. Power & Performance

Option Values: <None>

Help Text: *View/Configure processor information and settings.*

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Power & Performance** group of configuration settings.

[Back to \[Advanced Screen\]](#)

Memory Configuration

Option Values: <None>

Help Text:

View/Configure memory information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Memory Configuration** group of configuration settings.

Back to [Advanced Screen]

Mass Storage Controller Configuration

Option Values: <None>

Help Text:

View/Configure mass storage controller information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Mass Storage Controller Configuration** group of configuration settings.

Back to [Advanced Screen]

PCI Configuration

Option Values: <None>

Help Text:

View/Configure PCI information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **PCI Configuration** group of configuration settings.

Back to [Advanced Screen]

Serial Port Configuration

Option Values: <None>

Help Text:

View/Configure serial port information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Serial Port Configuration** group of configuration settings.

Back to [Advanced Screen]

USB Configuration

Option Values: <None>

Help Text:

View/Configure USB information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **USB Configuration** group of configuration settings.

Back to [Advanced Screen]

System Acoustic and Performance Configuration

Option Values: <None>

Help Text:

View/Configure system acoustic and performance information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **System Acoustic and Performance Configuration** group of configuration settings.

Back to [Advanced Screen]

12.2.2.1 Processor Configuration

The Processor Configuration screen displays the processor identification and microcode level, core frequency, cache sizes, Intel® QuickPath Interconnect (QPI) information for all processors currently installed. It also allows the user to enable or disable a number of processor options.

To access this screen from the **Main** screen, select **Advanced > Processor Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

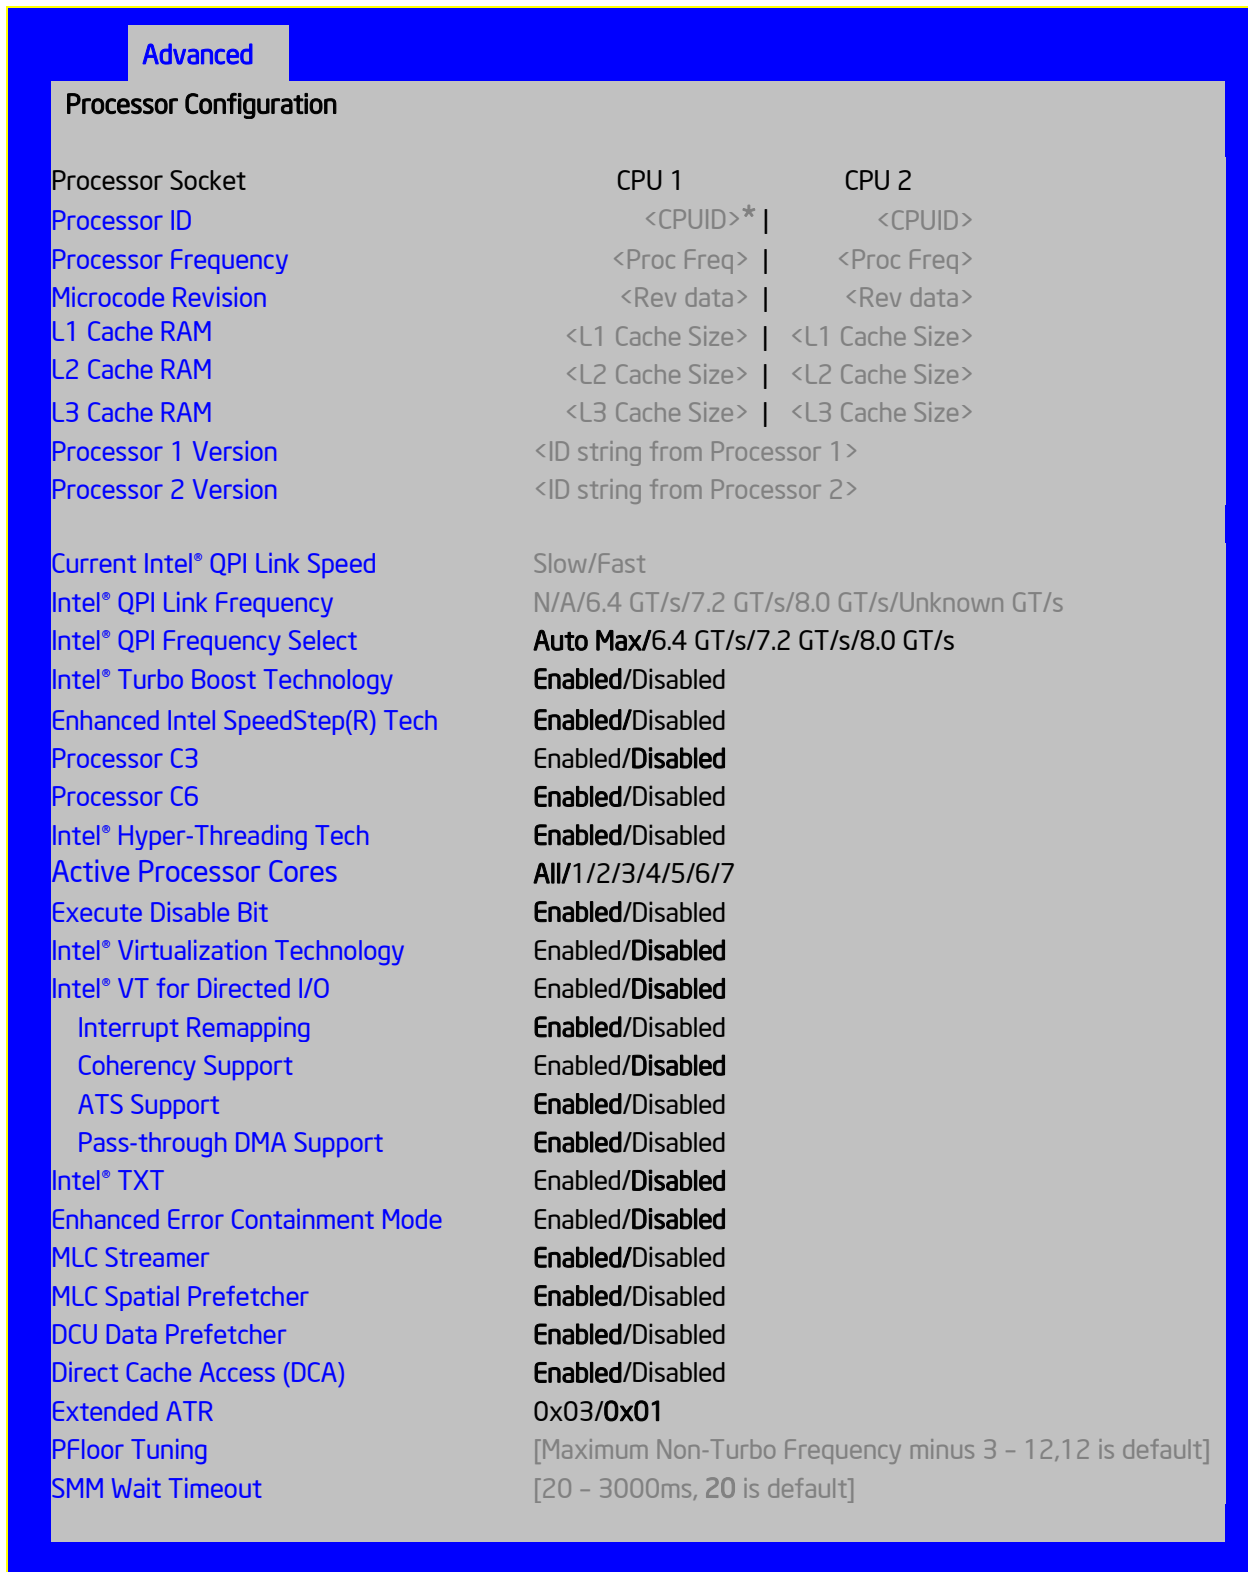


Figure 43. Processor Configuration Screen

Screen Field Descriptions:

Processor ID

Option Values: <CPUID>

Help Text: <None>

Comments: Information only. Displays the Processor Signature value (from the CPUID instruction) identifying the type of processor and the stepping.

For multi-socket boards, the processor selected as the Bootstrap Processor (BSP) has an asterisk (“*”) displayed beside the Processor ID. “N/A” will be displayed for a processor if not installed.

S1400 or S1600 series boards have a single Processor ID display

S2400 or S2600 series boards have 2 Processor ID displays, regardless of whether the second CPU socket has a processor installed. If the socket does not have a processor installed, “N/A” will be displayed for the processor data.

S4600 series boards have 4 Processor ID displays, regardless of whether the 2nd through 4th CPU sockets have a processor installed. For empty CPU sockets, “N/A” will be displayed for processor data.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Processor Frequency

Option Values: <Current Processor Operating Frequency>

Help Text: <None>

Comments: Information only. Displays current operating frequency of the processor.

Single socket boards have a single processor display, 2 socket or 4 socket. boards have a display column for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Microcode Revision

Option Values: <Microcode Revision Number>

Help Text: <None>

Comments: Information only. Displays Revision Level of the currently loaded processor microcode.

Single socket boards have a single processor display, 2 socket or 4 socket. boards have a display column for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

L1 Cache RAM

Option Values: <L1 cache size>

Help Text: <None>

Comments: Information only. Displays size in KB of the processor L1 Cache. Since L1 cache is not shared between cores, this is shown as the amount of L1 cache per core. There are two types of L1 cache, so this amount is the total of L1 Instruction Cache plus L1 Data Cache for each core.

Single socket boards have a single processor display, 2 socket or 4 socket. boards have a display column for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

L2 Cache RAM

Option Values: <L2 cache size>

Help Text: <None>

Comments: Information only. Displays size in KB of the processor L2 Cache. Since L2 cache is not shared between cores, this is shown as the amount of L2 cache per core.

Single socket boards have a single processor display, 2 socket or 4 socket. boards have a display column for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

L3 Cache RAM

Option Values: <L3 cache size>

Help Text: <None>

Comments: Information only. Displays size in MB of the processor L3 Cache. Since L3 cache is shared between all cores in a processor package, this is shown as the total amount of L3 cache per processor package.

Single socket boards have a single processor display, 2 socket or 4 socket boards have a display column for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Processor Version

See following...

Processor 1 Version

See following...

Processor 2 Version

Option Values: <ID string from processor>

Help Text: <None>

Comments: Information only. Displays Brand ID string read from processor with CPUID instruction.

Single socket boards have a single processor display, 2 socket or 4 socket. boards have a display line for each socket, showing “N/A” for empty sockets where processors are not installed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Current Intel® QPI Link Speed

Option Values: *Slow*
Fast

Help Text: <None>

Comments: Information only. Displays current Link Speed setting for the QPI Links. Appears only on multi-socket boards.

QPI Link Speed should display as “Slow” only when running at the “Boot Speed” of 50 MT/s, or when a multi-socket board has only one processor installed, so QPI is not functional. It should always be “Fast” when the QPI Link Frequency is in the normal functional range of 6.4 GT/s or above.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® QPI Link Frequency

Option Values: N/A
 6.4 GT/s
 7.2 GT/s
 8.0 GT/s
 Unknown GT/s

Help Text: <None>

Comments: *Information only.* Displays current frequency at which the QPI Links are operating.
Appears only on multi-socket boards.

When a multi-socket board has only one processor installed, QPI Link Frequency will be shown as "N/A".

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® QPI Frequency Select

Option Values: **Auto Max**
 6.4 GT/s
 7.2 GT/s
 8.0 GT/s

Help Text:

Allows for selecting the Intel® QuickPath Interconnect Frequency. Recommended to leave in [Auto Max] so that BIOS can select the highest common Intel® QuickPath Interconnect frequency.

Comments: Lowering the QPI frequency may improve performance per watt for some processing loads and on certain benchmarks. [Auto Max] will give the maximum QPI performance available. Appears only on multi-socket boards.

When a multi-socket board has only one processor installed, this will be grayed out, with the previous value remaining displayed.

Changes in QPI Link Frequency will not take effect until the system reboots, so this will not immediately change the QPI Link Frequency display. Changing QPI Link Frequency does not affect the QPI Link Speed.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® Turbo Boost Technology

Option Values: **Enabled**
 Disabled

Help Text:

Intel® Turbo Boost Technology allows the processor to automatically increase its frequency if it is running below power, temperature, and current specifications.

Comments: This option is only visible if all processors installed in the system support Intel® Turbo Boost Technology. In order for this option to be available, Enhanced Intel® SpeedStep® Technology must be **Enabled**.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Enhanced Intel SpeedStep(R) Tech

Option Values: **Enabled**
 Disabled

Help Text:

Enhanced Intel SpeedStep (R) Technology allows the system to dynamically adjust processor voltage and core frequency, which can result in decreased average power consumption and decreased average heat production.

Contact your OS vendor regarding OS support of this feature.

Comments: When Disabled, the processor setting reverts to running at Max TDP Core Frequency (rated frequency).

This option is only visible if all processors installed in the system support Enhanced Intel® SpeedStep® Technology. In order for the Intel® Turbo Boost option to be available, Enhanced Intel® SpeedStep® Technology must be **Enabled**.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Processor C3

Option Values: Enabled
Disabled

Help Text:
Enable/Disable Processor C3 (ACPI C2/C3) report to OS

Comments: This is normally **Disabled**, but can be **Enabled** for improved performance on certain benchmarks and in certain situations.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Processor C6

Option Values: **Enabled**
Disabled

Help Text:
Enable/Disable Processor C6 (ACPI C3) report to OS

Comments: This is normally **Enabled** but can be **Disabled** for improved performance on certain benchmarks and in certain situations.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® Hyper-Threading Tech

Option Values: **Enabled**
Disabled

Help Text:
Intel® Hyper-Threading Technology allows multithreaded software applications to execute threads in parallel within each processor.

Contact your OS vendor regarding OS support of this feature.

Comments: This option is only visible if all processors installed in the system support Intel® Hyper-Threading Technology.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

1. Active Processor Cores

Option Values: All

1
2
3
4
5
6
7

Help Text:

Number of cores to enable in each processor package.

Comments: The numbers of cores that appear as selections depends on the number of cores available in the processors installed. Boards may have as many as 8 cores in each of 1, 2, or 4 processors. The same number of cores must be active in each processor package.

This Setup screen should begin with the number of currently-active cores as the number displayed. See note below – this may be different from the number previously set by the user.

Note: The ME can control the number of active cores independently of the BIOS Setup setting. If the ME disables or enables processor cores, that will override the BIOS setting, and the number selected by BIOS will be disregarded.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Execute Disable Bit

Option Values: Enabled

Disabled

Help Text:

Execute Disable Bit can help prevent certain classes of malicious buffer overflow attacks.

Contact your OS vendor regarding OS support of this feature.

Comments: This option is only visible if all processors installed in the system support the Execute Disable Bit. The OS and applications installed must support this feature in order for it to be enabled.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® Virtualization Technology

Option Values: Enabled

Disabled

Help Text:

Intel® Virtualization Technology allows a platform to run multiple operating systems and applications in independent partitions.

Note: A change to this option requires the system to be powered off and then back on before the setting takes effect.

Comments: This option is only visible if all processors installed in the system support Intel® VT. The software configuration installed on the system must support this feature in order for it to be enabled.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® VT for Directed I/O

Option Values: Enabled
Disabled

Help Text:
*Enable/Disable Intel® Virtualization Technology for Directed I/O (Intel® VT-d).
Report the I/O device assignment to VMM through DMAR ACPI Tables.*

Comments: This option is only visible if all processors installed in the system support Intel® VT-d. The software configuration installed on the system must support this feature in order for it to be enabled.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Interrupt Remapping

Option Values: **Enabled**
Disabled

Help Text:
Enable/Disable Intel® VT-d Interrupt Remapping support. For some processors, this option may be "always enabled".

Comments: This option only appears when Intel® Virtualization Technology for Directed I/O is **Enabled**. For some processors this will be enabled unconditionally whenever Intel® VT-d is enabled. In that case, this option will be shown as "Enabled", and grayed out and not changeable.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Coherency Support

Option Values: Enabled
Disabled

Help Text:
Enable/Disable Intel® VT-d Coherency support.

Comments: This option only appears when Intel® Virtualization Technology for Directed I/O is **Enabled**.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

ATS Support

Option Values: **Enabled**
Disabled

Help Text:
Enable/Disable Intel® VT-d Address Translation Services (ATS) support.

Comments: This option only appears when Intel® Virtualization Technology for Directed I/O is **Enabled**.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Pass-through DMA Support

Option Values: **Enabled**
Disabled

Help Text:
Enable/Disable Intel® VT-d Pass-through DMA support. For some processors, this option may be "always enabled".

Comments: This option only appears when Intel® Virtualization Technology for Directed I/O is **Enabled**. For some processors this will be enabled unconditionally whenever Intel® VT-d is enabled. In that case, this option will be shown as "Enabled", and grayed out and not changeable.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Intel® TXT

Option Values: Enabled
Disabled

Help Text:
Enable/Disable Intel® Trusted Execution Technology. Takes effect after reboot.

Comments: Intel® TXT only appears when both Intel® Virtualization Technology and Intel® VVT for Directed IO are enabled.

This option appears only on models equipped with a TPM. The TPM must be active in order to support Intel® TXT.

Note: Changing the setting for Intel® TXT will require the system to perform a Hard Reset in order for the new setting to become effective.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

2. Enhanced Error Containment Mode

Option Values: Enabled
Disabled

Help Text:
Enable Enhanced Error Containment Mode (Data Poisoning) - Erroneous data coming from memory will be poisoned. If disabled (default), will be in Legacy Mode - No data poisoning support available.

Comments: Enhanced Error Containment (Data Poisoning) is not supported by all models of processors, and this option will not appear unless all installed processors support Enhanced Error Containment. This option globally enables or disables both Core and Uncore Data Poisoning, for processors which support them.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

MLC Streamer

Option Values: **Enabled**
Disabled

Help Text:
MLC Streamer is a speculative prefetch unit within the processor(s).

Note: Modifying this setting may affect performance.

Comments: MLC Streamer is normally **Enabled**, for best efficiency in L2 Cache and Memory Channel use, but disabling it may improve performance for some processing loads and on certain benchmarks.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

MLC Spatial Prefetcher

Option Values: **Enabled**
 Disabled

Help Text:

[Enabled] – Fetches adjacent cache line (128 bytes) when required data is not currently in cache.

[Disabled] – Only fetches cache line with data required by the processor (64 bytes).

Comments: MLC Spatial Prefetcher is normally **Enabled**, for best efficiency in L2 Cache and Memory Channel use, but disabling it may improve performance for some processing loads and on certain benchmarks.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

DCU Data Prefetcher

Option Values: **Enabled**
 Disabled

Help Text:

The next cache line will be prefetched into L1 data cache from L2 or system memory during unused cycles if it sees that the processor core has accessed several bytes sequentially in a cache line as data.

[Disabled] – Only fetches cache line with data required by the processor (64 bytes).

Comments: DCU Data Prefetcher is normally **Enabled**, for best efficiency in L1 Data Cache and Memory Channel use, but disabling it may improve performance for some processing loads and on certain benchmarks.

[Back to Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

DCU Instruction Prefetcher

Option Values: **Enabled**
 Disabled

Help Text:

The next cache line will be prefetched into L1 instruction cache from L2 or system memory during unused cycles if it sees that the processor core has accessed several bytes sequentially in a cache line as data.

Comments: DCU Data Prefetcher is normally **Enabled**, for best efficiency in L1 I Cache and Memory Channel use, but disabling it may improve performance for some processing loads and on certain benchmarks.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Direct Cache Access (DCA)

Option Values: **Enabled**
 Disabled

Help Text:

Allows processors to increase the I/O performance by placing data from I/O devices directly into the processor cache.

Comments: System performance is usually best with Direct Cache Access Enabled. In certain unusual cases, disabling this may give improved results.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

Extended ATR

Option Values: **0x01**
 0x03

Help Text:
 Extended *Timeout* value for PCIe tuning.

Comments: Adjust Aging Timer Rollover (ATR) value for PCIe performance improvement.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

PFloor Tuning

Option Values: [Entry Field maximum Non-Turbo Frequency -3 of installed Processors –
 12, **12** is default]

Help Text:
 Adjustment of CPU idle frequency for PCIe Bus tuning.

Comments: Changing processor power state configuration may improve I/O performance. If processor low power states (Processor C-state) appear to be involved in the I/O performance issue experimentally, the next step is to see if establishing a minimum operating frequency resolves the problem by setting the value of the BIOS Setup option "PFloor Tuning". This will set the minimum operating frequency to be the same as the maximum non-Turbo frequency. The BIOS provides the Setup option "PFloor Tuning" with allowable range between the maximum Non-Turbo operating frequency minus 3 and 12, and the value is the number of 100 MHz step.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

3. SMM Wait Timeout

Option Values: [Entry Field 20 – 3000ms, **20** is default]

Help Text:
 Millisecond timeout waiting for BSP and APs to enter SMM. Range is 20ms to 3000ms.

Comments: Amount of time to allow for the SMI Handler to respond to an SMI. If exceeded, BMC generates an SMI Timeout and resets the system.

Note: this field is temporary, and will be removed when no longer required.

[Back to \[Processor Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12.2.2.2 Power & Performance

The Power & Performance screen allows the user to specify a profile which is optimized in the direction of either reduced power consumption or increased performance.

To access this screen from the **Main** screen, select **Advanced > Power and Performance**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

There are four possible profiles from which to choose. When a Power and Performance Profile is chosen, that in turn will cause the system to implement a defined list of Setup option settings and internal (non-visible) settings.

There is an explanation displayed on the screen, because of the fact that other settings may be adjusted without specifically notifying the user.

Advanced

Power & Performance

CPU Power and Performance Policy

Performance/**Balanced Performance**/
Balanced Power/Power

[Performance] Optimization is strongly toward performance, even at the expense of energy efficiency.

[Balanced Performance] Weights optimization toward performance, while conserving energy.

[Balanced Power] Weights optimization toward energy conservation, with good performance.

[Power] Optimization is strongly toward energy efficiency, even at the expense of performance.

Figure 44. Power & Performance Screen

Screen Field Descriptions:

1. CPU Power and Performance Policy

Option Values: *Performance*
 Balanced Performance
 Balanced Power
 Power

Help Text:

Allows the user to set an overall power and performance policy for the system, and when changed will modify a selected list of options to achieve the policy. These options are still changeable outside of the policy, but do reflect the changes that the policy makes when a new policy is selected.

Comments: Choosing one of these four Power and Performance Profiles implements a number of changes in BIOS settings, both visible settings in the Setup screens and non-visible internal settings.

Back to [Power & Performance Screen] — [Advanced Screen]

12.2.2.3 Memory Configuration

The Memory Configuration screen allows the user to view details about the DDR3 DIMMs that are installed as system memory, and alter BIOS Memory Configuration settings where appropriate.

For S1400, S1600, S2400, S2600, and S4600 series boards this screen shows memory system information, has options to select, and allows the user to select the “Configure Memory RAS and Performance” screen for further system memory information and configuration.

This screen differs somewhat between different boards which have different memory configurations. Some boards have one processor socket and fewer DIMMs, while other boards have two sockets or four sockets, more DIMMs, and the boards may have RAS and Performance options if configured for them

To access this screen from the **Main** screen, select **Advanced > Memory Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

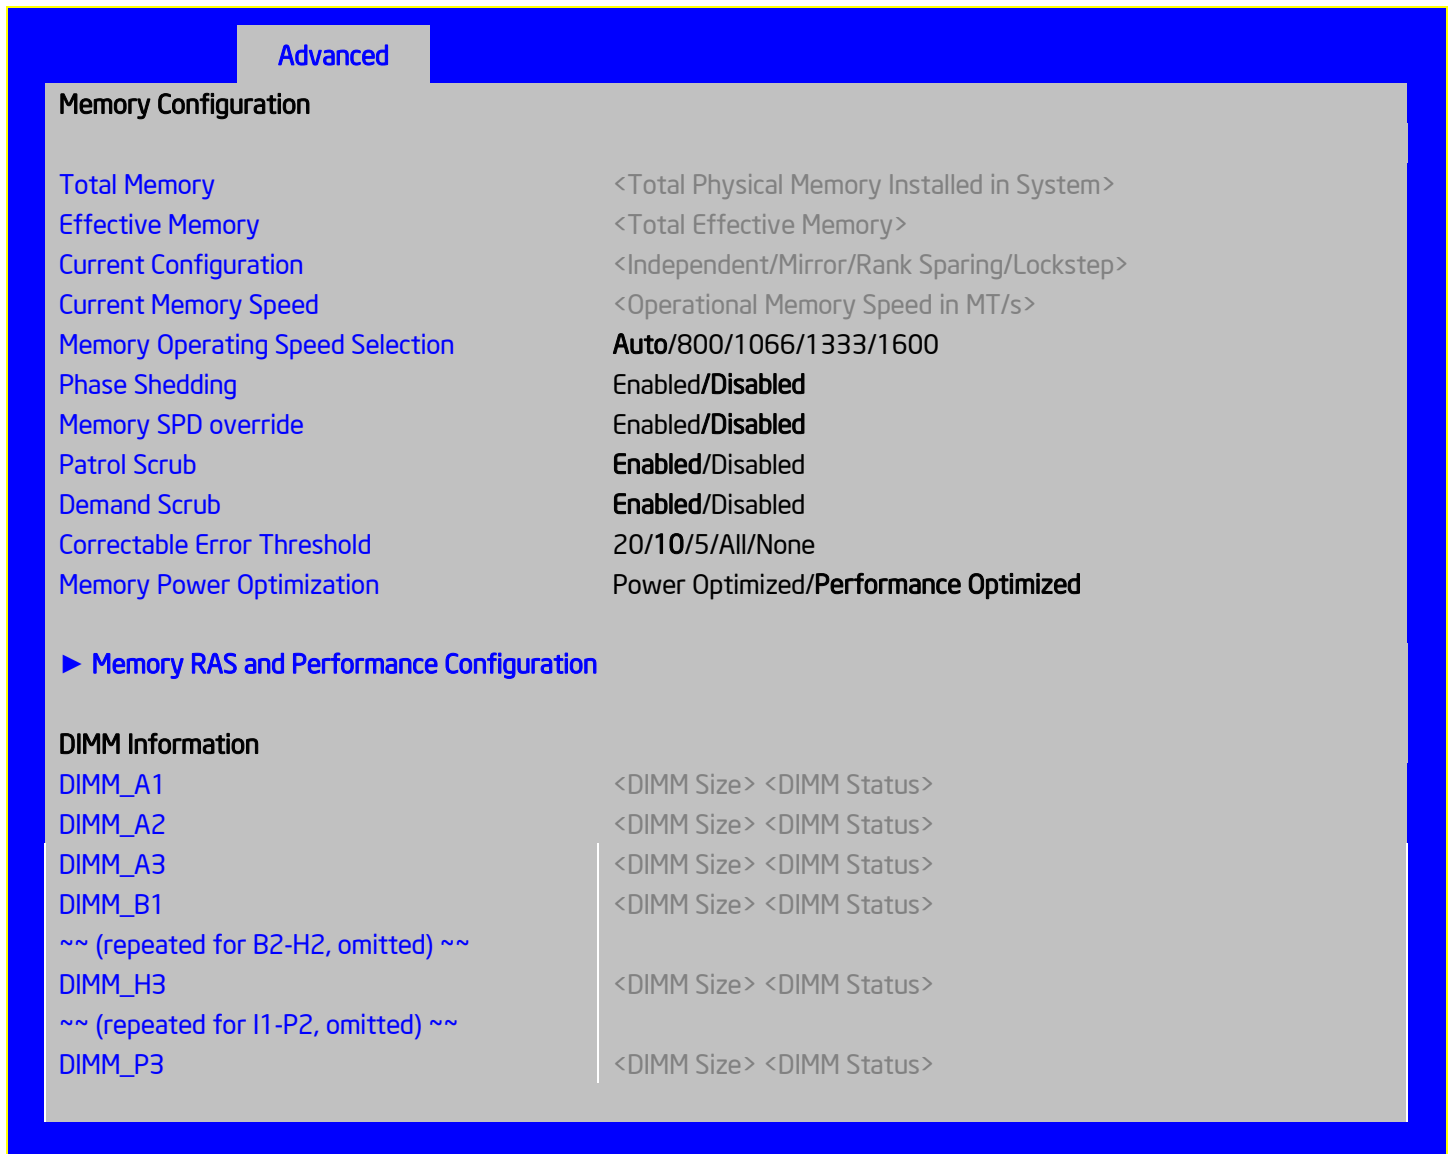


Figure 45. Memory Configuration Screen

Screen Field Descriptions:

1. Total Memory

Option Values: <Total Physical Memory Installed in System>

Help Text: <None>

Comments: Information only. Displays the amount of memory available in the system in the form of installed DDR3 DIMMs, in units of GB.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

2. Effective Memory

Option Values: <Total Effective Memory>

Help Text: <None>

Comments: Information only. Displays the amount of memory available to the OS in MB or GB.

The Effective Memory is the Total Physical Memory minus the sum of all memory reserved for internal usage, RAS redundancy and SMRAM.

Note: Some server operating systems do not display the total physical memory installed.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

3. Current ConfigurationOption Values: *Independent Channel**Mirror**Rank Sparing**Lockstep*

Help Text: <None>

Comments: Information only: Displays one of the following:

- **Independent Channel** – DIMMs are operating in Independent Channel Mode, the default configuration when there is no RAS Mode configured.
- **Mirror** – Mirroring RAS Mode has been configured and is operational.
- **Rank Sparing** – Rank Sparing RAS Mode has been configured and is operational
- **Lockstep** – Lockstep RAS Mode has been configured and is operational

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

4. Current Memory Speed

Option Values: <Operational Memory Speed in MT/s>

Help Text: <None>

Comments: Information only. Displays the speed in MT/s at which the memory is currently running.

The supported memory speeds are 800 MT/s, 1066 MT/s, 1333 MT/s, and 1600 MT/s. The actual memory speed capability depends on the memory configuration.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

5. Memory Operating Speed SelectionOption Values: **Auto**

800

1066

1333

1600

Help Text: *Force specific Memory Operating Speed or use Auto setting.*

Comments: Allows the user to select a specific speed at which memory will operate. Only speeds that are legitimate are available, that is, the user can only specify speeds less than or equal to the auto-selected Memory Operating Speed. The default **Auto** setting will select the highest achievable Memory Operating Speed consistent with the DIMMs and processors installed.

1600 MT/s memory speed is available only on certain models.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

6. Phase Shedding

Option Values: Enabled
 Disabled

Help Text:

Enable/disable DDR3 VR Static Phase Shedding. When enabled, DDR3 VR can be automatically adjusted by typical load.

Comments: This feature helps DDR3 VR optimize for high loading.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

7. Memory SPD Override

Option Values: Enabled
 Disabled

Help Text:

When enabled, it extends the platform's capability to support higher Memory Frequency than the planned settings with Intel® S2600 family processors; disabling it keeps the planned settings.

Comments: When enabled, it extends the platform's capability to support higher Memory Frequency.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

8. Patrol Scrub

Option Values: **Enabled**
 Disabled

Help Text:

When enabled, performs periodic checks on memory cells and proactively walks through populated memory space, to seek and correct soft ECC errors.

Comments: When enabled, Patrol Scrub is initialized to read through all of memory in a 24-hour period, correcting any Correctable ECC Errors it encounters by writing back the corrected data to memory.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

9. Demand Scrub

Option Values: **Enabled**
 Disabled

Help Text:

When enabled, executes when an ECC error is encountered during a normal read/write of data and corrects that data.

Comments: When enabled, Demand Scrub automatically corrects a Correctable ECC Error encountered during a fetch from memory by writing back the corrected data to memory.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

10. Correctable Error Threshold

Option Values: 20

10
5
All
None

Help Text:

Threshold value for logging Correctable Errors (CE) – Threshold of 10 (default) logs 10th CE, "All" logs every CE and "None" means no CE logging. All and None are not valid with Rank Sparing.

Comments: Specifies how many Correctable Errors must occur before triggering the logging of a SEL Correctable Error Event. Only the first threshold crossing is logged, unless "All" is selected. "All" causes every CE that occurs to be logged. "None" suppresses CE logging completely.

When Rank Sparing RAS Mode is configured, "All" and "None" are not valid, so they will not be presented as choices.

This threshold is applied on a per-rank basis. The Correctable Error occurrences are counted for each memory rank. When any one rank accumulates a CE count equal to the CE Threshold, then a single CE SEL Event is logged, and all further CE logging is suppressed.

Note that the CE counts are subject to a "leaky bucket" mechanism that reduces the count as a function of time, to keep from accumulating counts unnecessarily over the term of a long operational run.

This is also the Correctable Error threshold used when Rank Sparing RAS Mode is configured. When a CE threshold crossing occurs in Rank Sparing Mode on a channel which is in Redundant state, it causes a Sparing Fail Over (SFO) event to occur. That threshold crossing will also be logged as a Correctable Error event if it is the first to occur on the system.

An SFO event causes the rank with the error to be replaced by the spare rank for that channel., and the channel goes to a non-redundant state (with a "Redundancy Degraded" SEL Event logged). There may be an SFO for each channel in the system, although only the first one can be logged as a CE event.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

11. Memory Power OptimizationOption Values: Power Optimized
Performance Optimized

Help Text:

Power Optimized enables memory power limiting, Performance Optimized disables it for maximum performance.

Comments: When enabled, the system is configured to allow memory power management by the Node Manager (NM) and Management Engine (ME).

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12. Memory RAS and Performance Configuration

Option Values: <None>

Help Text:

Configure memory RAS (Reliability, Availability, and Serviceability) and view current memory performance information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **[Memory RAS and Performance Configuration](#)** group of configuration settings.

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

13. DIMM_A1

14. DIMM_A2

15. DIMM_A3

16. DIMM_B1

(DIMM_B2 through DIMM_H2 omitted)

17. DIMM_H3

(DIMM_I1 through DIMM_P2 omitted)

18. DIMM_P3

Option Values: <DIMM Size> <DIMM Status>
 Where DIMM Size is:
 Size of DIMM in GB
 Where DIMM Status is:
 Installed&Operational
 Not Installed
 Failed/Disabled

Help Text: <None>

Comments: *Information only*: Displays the status of each DIMM socket present on the board. There is one line for each DIMM socket present on the board.

For each DIMM socket, the DIMM Status reflects one of the following three possible states:

- ***Installed&Operational*** – There is a DDR3 DIMM installed and operational in this slot.
- ***Not Installed*** – There is no DDR3 DIMM installed in this slot.
- ***Failed/Disabled*** – The DIMM installed in this slot has failed during initialization and/or was disabled during initialization.

For each DIMM that is in the ***Installed&Operational*** state, the DIMM Size in GB of that DIMM is displayed. This is the physical size of the DIMM, regardless of how it is counted in the Effective Memory size.

Note: In “DIMM_XY”, X denotes the Channel Identifier A - P, and Y denotes the DIMM Slot identifier 1 - 3 within the Channel. DIMM_A2 is the DIMM socket on Channel A, Slot 2. Not all boards have the same number of channels and slots – this is dependent on the board features.

- **S1400 boards** can have DIMMs A1, A2 to C1, C2 (max 3 channels/2 DPC)
- **S1600 boards** can have DIMMs A1, A2, A3 to D1, D2, D3 (max 4 channels/3 DPC)
- **S2400 boards** can have DIMMs A1, A2 to F1, F2 (max 2 CPU/3 channels/2 DPC)
- **S2600 boards** can have DIMMs A1, A2, A3 to H1, H2, H3 (max 2 CPU/4 chan/3 DPC)
- **S4600 boards** can have DIMMs A1, A2, A3 to P1, P2, P3 (max 4 CPU/4 chan/3 DPC)

[Back to \[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12.2.2.4 Memory RAS and Performance Configuration

The Memory RAS and Performance Configuration screen allows the user to customize several memory configuration options, such as whether to use Memory Mirroring or Memory Sparring.

To access this screen from the **Main** screen, select **Advanced > Memory Configuration > Memory RAS and Performance Configuration**. To move to another screen, press the <Esc> key to return to the **Memory Configuration** screen, if necessary press the <Esc> key again to return to the **Advanced** screen, then select the desired screen.

Advanced

Memory RAS and Performance Configuration

Capabilities

Memory Mirroring Possible	Yes/No
Memory Rank Sparing Possible	Yes/No
Memory Lockstep Possible	Yes/No
Select Memory RAS Configuration	Maximum Performance/Mirroring/Rank Sparing/Lockstep
NUMA Optimized	Enabled/Disabled

Figure 46. Memory RAS and Performance Configuration Screen

Screen Field Descriptions:

1. Memory Mirroring Possible

Option Values: Yes
 No

Help Text: <None>

Comments: Information only. Displays whether the current DIMM configuration is capable of Memory Mirroring. For Memory Mirroring to be possible, DIMM configurations on all paired channels must be identical between the channel pair (Mirroring Domain).

Back to [[Memory RAS and Performance Configuration Screen](#)] — [[Memory Configuration Screen](#)] — [[Advanced Screen](#)]

2. Memory Rank Sparing Possible

Option Values: Yes
 No

Help Text: <None>

Comments: Information only. Displays whether the current DIMM configuration is capable of Rank Sparing. For Rank Sparing to be possible, DIMM configurations on all channels must be capable of supporting Rank Sparing.

Note: The Correctable Error Threshold value is also the Sparing Fail Over threshold value. Threshold values of "All" or "None" are not valid for Rank Sparing. If the Correctable Error Threshold is set to either of those values, Rank Spring will not be possible. See Memory Configuration Setup screen.

Back to [[Memory RAS and Performance Configuration Screen](#)] — [[Memory Configuration Screen](#)] — [[Advanced Screen](#)]

3. Memory Lockstep Possible

Option Values: Yes
 No

Help Text: <None>

Comments: *Information only.* Displays whether the current DIMM configuration is capable of Memory Lockstep. For Memory Lockstep to be possible, DIMM configurations on all paired channels must be identical between the channel pair.

[Back to \[Memory RAS and Performance Configuration Screen\]](#) — [\[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

4. Select Memory RAS Configuration

Option Values: **Maximum Performance**

Mirroring
Rank Sparing
Lockstep

Help Text:

Allows the user to select the memory RAS Configuration to be applied for the next boot.

Comments: Available modes depend on the current memory population. Modes which are not listed as “possible” should not be available as choices. If the only valid choice is “Maximum Performance”, then this option should be grayed out and unavailable.

Maximum Performance – (default) no RAS, but best memory utilization since full amount of memory is available, operating in Independent Channel Mode.

Mirroring - most reliability by using half of memory as a mirror image, can survive an Uncorrectable ECC Error.

Rank Sparing – offers reliability by reserving spare ranks to replace failing ranks which are generating excessive Correctable ECC Errors.

Lockstep – allows SDDC capability with x8 DIMMs installed. No memory size impact, but does have a performance and power penalty.

Note: since only RAS Modes which are listed as “possible” are available for selection, it is not possible to select a RAS Mode without first installing an appropriate DIMM configuration.

[Back to \[Memory RAS and Performance Configuration Screen\]](#) — [\[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

5. NUMA Optimized

Option Values: **Enabled**

Disabled

Help Text:

If enabled, BIOS includes ACPI tables that are required for NUMA-aware Operating Systems.

Comments: This option is only present for boards which have two or more processor sockets. When a multi-socket board has only a single processor installed, this option is grayed out and set as Disabled.

When enabled, the SRAT and SLIT ACPI tables are provided that show the locality of systems resources, especially memory, which allows a “NUMA Aware” OS to optimize which processor threads are used by processes which can benefit by having the best access to those resources.

[Back to \[Memory RAS and Performance Configuration Screen\]](#) — [\[Memory Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12.2.2.5 Mass Storage Controller Configuration

The Mass Storage Configuration screen allows the user to configure the Mass Storage controllers that are integrated into the server board on which the BIOS is executing. This includes only onboard Mass Storage controllers. Mass Storage controllers on add-in cards are not included in this screen, nor are other storage mechanisms such as USB-attached storage devices or Network Attached Storage.

There are two types of onboard controller configured in this screen, the AHCI SATA controller and the Storage Control Unit (SCU) with SATA or SAS drive support and RAID support. There are also informational displays of AHCI and SCU controller configuration, and SATA Drive Information when applicable. If the presence of an Intel® Storage Module is detected, the type of Storage Module is displayed as information-only.

To access this screen from the **Main** screen, select **Advanced > Mass Storage Controller Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

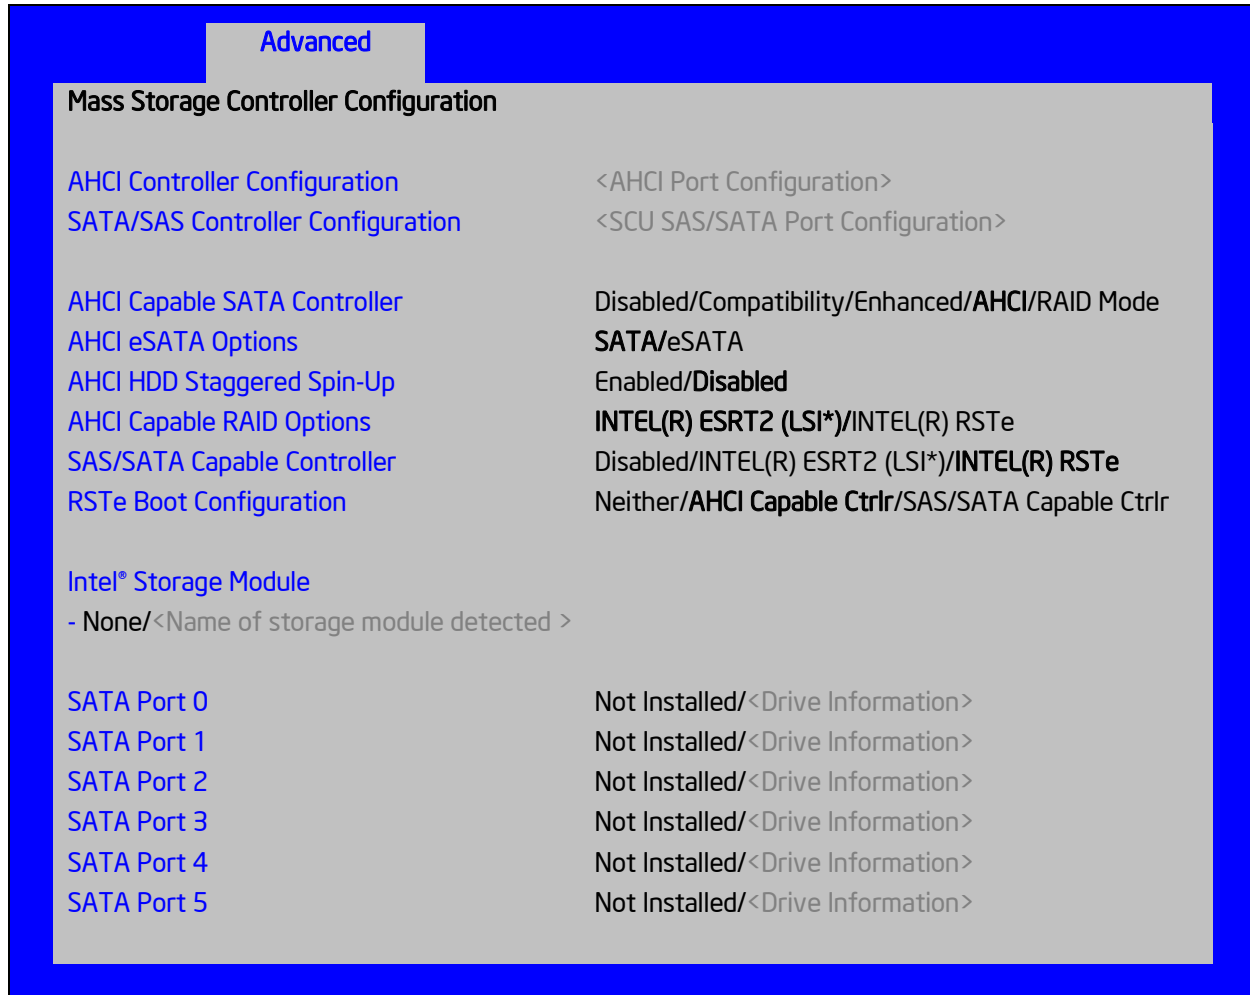


Figure 47. Mass Storage Controller Configuration Screen

Screen Field Descriptions:

AHCI Controller Configuration

Option Values: <AHCI Port Configuration>

One of these strings:

Controller is disabled

2 ports of 6Gb/s SATA

2 ports of 6Gb/s SATA & 4 ports of 3Gb/s SATA

Help Text: <None>

Comments: *Information only.* This is a display showing which ports are available through the onboard AHCI capable SATA controller, if the controller is enabled.

This information is also displayed during POST in the POST Diagnostic Screen.

The number of SATA ports available from the integrated AHCI-capable SATA Controller is dependent on the specific server board installed in the system. Different server board designs expose different SATA port configurations. The Platform ID (Board ID) is displayed in the [Main Screen](#).

[Back to \[Mass Storage Controller Configuration Screen\]](#)

1. SATA/SAS Controller Configuration

Option Values: <SCU SAS/SATA Port Configuration>

One of these strings:

*Controller is disabled
4 ports in SATA mode
4 ports in SAS mode
8 ports in SATA mode
8 ports in SAS mode*

Help Text: <None>

Comments: *Information only.* This is a display showing the number of ports which are available through the SCU controller, and whether they are configured for SATA or SAS.

Various SATA/SAS Capable Controller configurations require the installation of Intel® RAID C600 Upgrade Keys:

4 port SATA requires no key or AXXRKSATA4R5 key
4 port SAS requires AXXRKSAS4 or AXXRKSAS4R5 key
8 port SATA requires AXXRKSATA8 or AXXRKSATA8R5 key
8 port SAS requires AXXRKSAS8 or AXXRKSAS8R5 key

[Back to \[Mass Storage Controller Configuration Screen\]](#)

2. AHCI Capable SATA Controller

Option Values: Disabled
Compatibility
Enhanced
AHCI
RAID Mode

Help Text:

- *Compatibility provides PATA emulation on the SATA device*
- *Enhanced provides Native SATA support*
- *AHCI enables the Advanced Host Controller Interface, which provides Enhanced SATA functionality*
- *RAID Mode provides host based RAID support on the onboard SATA ports*

Comments: This option configures the onboard AHCI-capable SATA controller, which is distinct from the SCU. The number and type of ports it controls differs between board series.

If the SATA Controller is Disabled, the SATA Ports will not operate. and any installed SATA devices will be unavailable.

Compatibility provides PATA emulation on the SATA device, allowing the use of legacy IDE/PATA drivers. Enhanced provides Native SATA support., using native SATA drivers included with the vast majority of current OSes. AHCI enables the Advanced Host Controller Interface, which provides Enhanced SATA functionality, plus possible additional functionality (Native Command Queuing, Hot Plug, Staggered Spin Up). It uses AHCI drivers available for the majority of current OSes.

RAID Mode provides host based RAID support on the onboard SATA ports. RAID levels supported and required drivers depend on the RAID stack selected

[Back to \[Mass Storage Controller Configuration Screen\]](#)

AHCI eSATA Options

Option Values: **SATA**
 eSATA

Help Text:

- *SATA mode enables the switchable internal AHCI SATA (port 1)*
- *eSATA mode enables the switchable external AHCI eSATA (port 1)*
- *These modes are mutually exclusive, so SATA port 1 will only be active on one connector, not both*

Comments: In order to use the external eSATA connection, this option must be set to eSATA. When the external eSATA connector is selected, it disables the corresponding internal SATA port 1 connector. When set to SATA, the internal connector for SATA port 1 is active, and the external eSATA connector is disabled.

This option setting only appears when the SATA Controller is enabled, and only for platforms which support eSATA.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

3. AHCI HDD Staggered Spin-Up

Option Values: Enabled
 Disabled

Help Text:

If enabled for the AHCI Capable SATA controller, Staggered Spin-Up will be performed on drives attached to it. Otherwise these drives will all spin up at boot.

Comments: This option enables or disables Staggered Spin-up only for disk drives attached to ports on the AHCI Capable SATA Controller. Disk drives attached to SATA/SAS ports on the Storage Control Unit are controlled by a different method for Staggered Spin-Up and this option does not affect them.

This option is only visible when the SATA Controller is enabled and AHCI or RAID has been selected as the operational SATA Mode.

Staggered Spin-Up is needed when there are enough HDDs attached to the system to cause a marked startup power demand surge when all drives start spin-up together. Since the power demand is greatest just as the drive spinning is started, the overall startup power demand can be leveled off by starting up each drive at a slightly different time, so the power demand surges for multiple drives do not coincide and cause too great a power draw.

When Staggered Spin-Up is enabled, it does have a possibility of increasing boot time if there are many HDDs attached, because of the interval between starting drives spinning. However, that is exactly the scenario in which Staggered Spin-Up is most needed, because the more disk drives attached, the greater the startup demand surge.

Setting the external eSATA connector Enabled (when available) does not invalidate the Staggered Spin-Up option, although there may be less need for Staggered Spin-Up in a system configured for eSATA use.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

AHCI Capable RAID Options

Option Values: **Intel® ESRT2 (LSI*)**
 Intel® RSTe

Help Text:

- Intel® ESRT2 (Powered By LSI*): Supports RAID 0/1/10 and optional RAID 5 with Intel® RAID C600 Upgrade Keys. Uses Intel ESRT2 drivers (based on LSI* MegaSR).
- Intel® RSTe: Provides pass-through drive support. Also provides host based RAID 0/1/10/5 support. Uses Intel® RSTe iastor drivers.

Comments: This option only appears when the SATA Controller is enabled, and RAID Mode has been selected as the operational SATA Mode. This setting selects the RAID stack to be used for SATA RAID with the onboard AHCI SATA controller.

If a RAID Volume has not previously been created that is compatible with the RAID stack selected, it will be necessary to Save and Exit and reboot in order to create a RAID Volume.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

4. SAS/SATA Capable Controller

Option Values: Disabled
 Intel® ESRT2 (LSI*)
 Intel® RSTe

Help Text:

- Intel® ESRT2: Provides host based RAID 0/1/10 and optional RAID 5. Uses Intel® ESRT2 drivers (based on LSI* MegaSR).
- Intel® RSTe: Provides pass-through drive support. Also provides host based RAID 0/1/10 support, and RAID 5 (in SATA mode only). Uses Intel® RSTe iastor drivers.

Comments: This option selects the RAID stack to be used with the SCU. If Disabled is selected, any drives connected to the SCU will not be usable.

Intel® ESRT2 provides host based RAID 0/1/10 and optional RAID 5. For a RAID 5 configuration, this requires one of the Intel® RAID C600 Upgrade Keys AXXRKSATA4R5, AXXRKSATA8R5, AXXRKSAS4R5, or AXXRKSAS8R5. Uses Intel® ESRT2 drivers (based on LSI* MegaSR), and is also supported by Linux MDRAID.

Intel® RSTe provides pass-through drive support and provides host based RAID 0/1/10 support, and RAID 5 (in SATA mode only). Uses Intel RSTe drivers in Windows, and MDRAID stack in Linux. The Intel® RSTe RAID stack is required if it is necessary to provide pass-through support for non-RAID drives, or if support is needed for more than 8 drives.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

5. RSTe Boot Configuration

Option Values: Neither
 AHCI Capable Ctrlr
 SAS/SATA Capable Ctrlr

Help Text:

This selects the device that will support Bootable Drives, whether they are in RAID arrays or individual pass-through SAS/SATA drives. Once selected and set up (if necessary), individual bootable devices will be listed in the Bootable Devices menu display.

Comments: This option appears only when Intel® RSTe has been selected as the operational mode on both the AHCI and SCU controllers. In that case there is a conflict and only one controller can be selected as having bootable drives attached.

Once selected and set up (if necessary), individual bootable logical or physical drives available on the selected controller will be listed in the Bootable Devices menu display.

If only one device selects RSTe, it will be available as a boot device along with any other devices – this option is only necessary to distinguish between which RSTe device runs the Option ROM instance.

BIOS is required to designate the OPROM for the boot device selected here. Two iterations of the OPROM cannot fully load simultaneously, and the version fully loaded will only show devices connected to the given controller, so the OPROM load order is based on BIOS selecting the correct device.

Note: If RSTe is selected, then only one CONTROLLER can be bootable, so there will be situations where the boot drive *OR* an optical device will be bootable, but not both.

Please also see the product System Guide for restrictions on expander boot support.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

Intel® Storage Module

Option Values: **None**
<Name of Storage Module detected>

Names of Storage Modules supported at this time are:

Intel® Integrated RAID Module
 Intel® Integrated RAID Module RMS25PB040
 Intel® Integrated RAID Module RMT3PB080
 Intel® Integrated RAID Module RMS25CB080
 Intel® Integrated RAID Module RMS25CB040
 Intel® Integrated RAID Module RMT3CB080
 Intel® Integrated RAID Module RMS25JB080
 Intel® Integrated RAID Module RMS25JB040
 Intel® Integrated RAID Module RMS25KB080
 Intel® Integrated RAID Module RMS25KB040

Help Text: <None>

Comments: Information only. If no Intel® Storage Module is detected, then **None** is displayed. This shows the customer the product name of the module installed, which helps in identifying drivers, support, documentation, and so on

[Back to \[Mass Storage Controller Configuration Screen\]](#)

SATA Port

(For Port numbers 0-6)

Option Values: **Not Installed**
<Drive Information>

Help Text: <None>

Comments: Information only. The Drive Information, when present, will typically consist of the drive model identification and size for the disk drive installed on a particular port.

This Drive Information line is repeated for all 6 SATA Port for the onboard AHCI capable SATA Controller. However, for any given board, only the ports which are physically populated on the board are shown. That is, a board which only implements the two 6 GB/s ports 0 and 1 will only show those two ports in this Drive Information list.

This section for Drive Information does not appear at all when the SCU is set to *Disabled* or the SATA operational mode is *RAID Mode*, nor for any drives attached to the SCU SATA or SAS ports. In these cases the BBS information is not available to display.

[Back to \[Mass Storage Controller Configuration Screen\]](#)

12.2.2.6 PCI Configuration

The PCI Configuration screen allows the user to configure the PCI memory space used for onboard and add-in adapters, configure video options, and configure onboard adapter options.

It also includes a selection option to go to the NIC Configuration screen.

To access this screen from the **Main** screen, select **Advanced > PCI Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

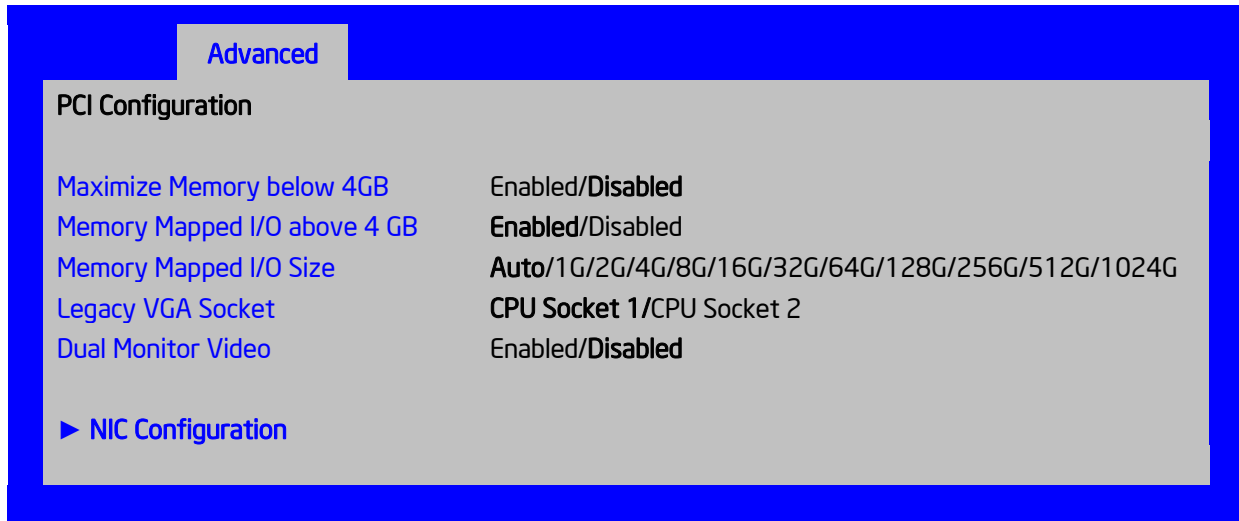


Figure 48. PCI Configuration Screen

Screen Field Descriptions:

1. Maximize Memory below 4GB

Option Values: Enabled
 Disabled

Help Text:

BIOS maximizes memory usage below 4GB for an OS without PAE support, depending on the system configuration. Only enable for an OS without PAE support.

Comments: When this option is enabled, BIOS makes as much memory available as possible in the 32-bit (4GB) address space, by limiting the amount of PCI/PCIe Memory Address Space and PCIe Extended Configuration Space. This option should only be enabled for a 32-bit OS without PAE capability or without PAE enabled.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

2. Memory Mapped I/O above 4 GB

Option Values: Enabled
 Disabled

Help Text:

Enable or disable memory mapped I/O of 64-bit PCI devices to 4 GB or greater address space.

Comments: When enabled, PCI/PCIe Memory Mapped I/O for devices capable of 64-bit addressing is allocated to address space above 4GB, in order to allow larger allocations and avoid impacting address space below 4G.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

3. Memory Mapped I/O Size

Option Values: **Auto**
 1G/2G/4G/8G/16G/32G/64G/128G/256G/512G/1024G

Help Text: Sets MMIO Size: Auto.

Comments: When Memory Mapped I/O above 4GB option enabled, this option sets the preserved MMIO size as PCI/PCIe Memory Mapped I/O for devices capable of 64-bit addressing. This option is grayed out when Memory Mapped I/O above 4GB option is disabled.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

4. Onboard Video

Option Values: **Enabled**
 Disabled

Help Text:

On-board video controller.

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

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

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

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

Note: This option does not appear on some models.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

5. Legacy VGA Socket

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

Help Text:

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

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

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

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

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

6. Dual Monitor Video

Option Values: Enabled
 Disabled

Help Text:

If enabled, both the on-board video controller and an add-in video adapter are enabled for system video. The on-board video controller becomes the primary video device.

Comments: This option must be enabled to use an add-in card as a secondary POST Legacy Video device while also displaying on the Onboard Video device.

If there is no add-in video card in any PCIe slot connected to CPU Socket 1, this options is set to Disabled and grayed out and unavailable.

If the Legacy VGA Socket option is set to CPU Socket 2, this option is grayed out and unavailable, with a value set to Disabled. When Legacy VGA Socket is set back to CPU Socket 1, this option is set to its default value of Disabled, and may become available depending on add-in video card configuration,

Note: This option does not appear on some models.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

7. NIC Configuration

Option Values: <None>

Help Text:

View/Configure NIC information and settings.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **NIC Configuration** group of configuration settings.

[Back to \[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12.2.2.7 NIC Configuration

The NIC Configuration screen allows the user to configure the NIC controller options for BIOS POST. It also displays the NIC MAC Addresses currently in use. This NIC Configuration screen handles network controllers built in on the baseboard (“onboard”) or installed as an IO Module (IOM). It does not configure or report anything having to do with add-in network adapter cards.

To access this screen from the **Main** screen, select **Advanced > PCI Configuration > NIC Configuration**. To move to another screen, press the <Esc> key to return to the **PCI Configuration** screen, if necessary press the <Esc> key again to return to the **Advanced** screen, then select the desired screen.

There is usually one Onboard NIC built into the baseboard, although in some cases there are two Onboard NICs. There are several possible types of NICs which are incorporated into different boards. When an Infiniband* controller is on the baseboard, it appears as an Onboard NIC.

Most boards in this family also can have an IO Module that installs on the board in a specialized connector. There are boards which can have two IO Modules installed.

The descriptive names of the Onboard NIC types are:

1. Intel® 82574 Single-Port Gigabit Ethernet Controller
2. Intel® I350 Dual-Port Gigabit Ethernet Controller
3. Intel® I350 Quad-Port Gigabit Ethernet Controller
4. Intel® I540 Dual-Port X540 10 Gigabit RJ-45 Controller

5. Mellanox* ConnectX-3* Single-Port Infiniband* FD14 Controller

For boards with only one Onboard NIC, the “Onboard NIC2” entries are not present on the screen. The number of “Port” options which are displayed for each NIC will match the number of ports the Onboard NIC presents.

The IO Modules currently available are:

1. Intel® I350 Quad-Port Gigabit Ethernet Module
2. Intel® I540 Dual-Port X540 10 Gigabit RJ-45 Module
3. Intel® 82599 Dual-Port 10 Gigabit SFP+ Module
4. Mellanox* ConnectX-3* Single-Port Infiniband* FD14 Module

For the IO Module entries on the NIC Configuration screen, only entries for modules which are currently installed will appear, and only ports which exist on those IO Modules will appear.

If an IO Module which had been installed is no longer installed when the system is booted, all NIC Configuration entries which are specific to that IO Module will be reset to their default values and hidden. If a different IO Module is installed than had been previously installed, the module-specific settings will still be returned to defaults, but not hidden. This will not necessarily affect the Option ROM settings, which depend on the aggregate capabilities of all installed Onboard and IO Module NICs.

For each NIC port which is present on an Onboard NIC or IO Module other than Infiniband* controllers, there will be a port-specific PXE Boot enabling option and a MAC Address display. Onboard NICs and NIC ports also have enable/disable options. IO Modules and the ports on them cannot be disabled by BIOS.

Infiniband* controllers which appear as Onboard NICs or as IO Modules have a slightly different format. They do not have enable/disable options, but they do have a choice of whether to enable loading and executing the embedded Option ROM for the controller, which will cause it to become bootable. For Infiniband*, both a GUID and a MAC Address are displayed. The GUID is used for Infiniband* Fabric operations, and the MAC Address is used when the controller is attached as an Ethernet device.

For non-Infiniband* NICs, there are different OPRoMs for different protocols, which are also differentiated by speed, 1 Gb or 10 Gb. For a given protocol/speed, all Ethernet controllers of the same speed use the same Option ROM.

PXE – there are two separate PXE Option ROMs, one for 1 Gb NICs and another for 10 Gb NICs. The two are independent of each other, but each must be the only Option ROM enabled in its speed class. If 1 GbE PXE is enabled, then the discs OPRoM cannot be enabled. If 10 GbE PXE is enabled, then neither discs nor 10 GbE FCoE may be enabled.

discs – there is only one discs Option ROM for both 1 GbE and 10 GbE NICs. If discs is enabled, then neither PXE nor FCoE OPRoMs may be enabled for the 1 GbE or 10 GbE NICs.

FCoE – there is a 10 GbE FCoE Option ROM that supports the Intel® 82599 NIC. When it is enabled, the discs OPRoM and the 10 GbE PXE OPRoM must be disabled

Note: These Option ROMs are only in support of onboard NICs and installed IO Modules. They do not support NICs on add-in network cards, even if the NIC on an add-in card is the same type of device as an onboard NIC or IO Module controller.

Only the Option ROMs for which controller capabilities are present are shown in the screen for selection. For example, if there are no 10 GbE NICs installed, then the 10 GbE OPRoMs will not appear for selection. If controller capabilities are present, but all controllers with those capabilities are disabled, then the relevant OPRoM options will appear, but will be disabled and grayed out and not changeable.

Similarly, when the PXE OPROM of a given speed is disabled, all PXE port enable/disable options using that OPROM will be disabled and grayed out. Conversely, if all ports are disabled for PXE, the PXE OPROM will be disabled and grayed out.

When a NIC Port is disabled, the PXE enable/disable option for it will be disabled and grayed out, and the MAC Address will be blank. When a NIC controller is disabled, all Ports and PXE options for that controller will become disabled and grayed out and all MAC Addresses for those ports will be blank. Conversely, if all ports for a given controller are disabled, the controller itself will appear as disabled.

Advanced	
NIC Configuration	
Wake on LAN (PME)	Enabled/Disabled
PXE 1GbE Option ROM	Enabled/Disabled
PXE 10GbE Option ROM	Enabled/Disabled
FCoE 10GbE Option ROM	Enabled/Disabled
discs 1GbE/10GbE Option ROM	Enabled/Disabled
Onboard NIC1 Type	<Onboard NIC Description - Non-InfiniBand>
NIC1 Controller	Enabled/Disabled
NIC1 Port1	Enabled/Disabled
NIC1 Port2	Enabled/Disabled
NIC1 Port3	Enabled/Disabled
NIC1 Port4	Enabled/Disabled
NIC1 Port1 PXE	Enabled/Disabled
NIC1 Port2 PXE	Enabled/Disabled
NIC1 Port3 PXE	Enabled/Disabled
NIC1 Port4 PXE	Enabled/Disabled
NIC1 Port1 MAC Address	<MAC Address display>
NIC1 Port2 MAC Address	<MAC Address display >
NIC1 Port3 MAC Address	<MAC Address display >
NIC1 Port4 MAC Address	<MAC Address display >
Onboard NIC2 Type	<Onboard NIC Description - InfiniBand Only>
NIC2 InfiniBand Option ROM	<u>Enabled/Disabled</u>
NIC2 Port1 GUID	<GUID Display>
NIC2 Port1 MAC Address	<MAC Address display >
Onboard NIC2 Type	<Onboard NIC Description - Non-InfiniBand>
NIC2 Controller	Enabled/Disabled
NIC2 Port1	Enabled/Disabled
NIC2 Port2	Enabled/Disabled
NIC2 Port3	Enabled/Disabled
NIC2 Port4	Enabled/Disabled
NIC2 Port1 PXE	Enabled/Disabled
NIC2 Port2 PXE	Enabled/Disabled
NIC2 Port3 PXE	Enabled/Disabled

NIC2 Port4 PXE	Enabled/Disabled
NIC2 Port1 MAC Address	<MAC Address display >
NIC2 Port2 MAC Address	<MAC Address display >
NIC2 Port3 MAC Address	<MAC Address display >
NIC2 Port4 MAC Address	<MAC Address display >
IO Module 1 Type	<IO Module Description - Non-InfiniBand>
IOM1 Port1 PXE	Enabled/Disabled
IOM1 Port2 PXE	Enabled/Disabled
IOM1 Port3 PXE	Enabled/Disabled
IOM1 Port4 PXE	Enabled/Disabled
IOM1 Port1 MAC Address	<MAC Address display >
IOM1 Port2 MAC Address	<MAC Address display >
IOM1 Port3 MAC Address	<MAC Address display >
IOM1 Port4 MAC Address	<MAC Address display >
IO Module 1 Type	<IO Module Description - InfiniBand Only>
IOM1 InfiniBand Option ROM	Enabled/Disabled
IOM1 Port1 GUID	<GUID Display>
IOM1 Port1 MAC Address	<MAC Address display >
IO Module 2 Type	<IO Module Description - Non-InfiniBand>
IOM2 Port1 PXE	Enabled/Disabled
IOM2 Port2 PXE	Enabled/Disabled
IOM2 Port3 PXE	Enabled/Disabled
IOM2 Port4 PXE	Enabled/Disabled
IOM2 Port1 MAC Address	<MAC Address display >
IOM2 Port2 MAC Address	<MAC Address display >
IOM2 Port3 MAC Address	<MAC Address display >
IOM2 Port4 MAC Address	<MAC Address display >

Figure 49. NIC Configuration Screen

Screen Field Descriptions:

1. Wake on LAN (PME)

Option Values: **Enabled**
Disabled

Help Text:

Enables or disables PCI PME function for Wake on LAN capability from LAN adapters.

Comments: Enables/disables PCI/PCIe PME# signal to generate Power Management Events (PME) and ACPI Table entries required for Wake on LAN (WOL). However, note that this will enable WOL only with an ACPI-capable Operating System which has the WOL function enabled.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

2. PXE 1GbE Option ROM

Option Values: **Enabled**
 Disabled

Help Text:

Enable/Disable Onboard/IOM NIC PXE Option ROM Load.

Comments: This selection is to enable/disable the 1GbE PXE Option ROM that is used by all Onboard and IO Module 1 GbE controllers.

This option is grayed out and not accessible if the discs Option ROM is enabled. It can co-exist with the 10 GbE PXE Option ROM, the 10 GbE FCoE Option ROM, or with an Infiniband* controller Option ROM.

If the 1GbE PXE Option ROM is disabled, and no other Option ROM is enabled, the system cannot perform a Network Boot and cannot respond for Wake-on-LAN.

This 1GbE PXE option does not appear unless there is a 1 GbE NIC installed in the system as an Onboard or IO Module NIC.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

3. PXE 10GbE Option ROM

Option Values: **Enabled**
 Disabled

Help Text:

Enable/Disable Onboard/IOM NIC PXE Option ROM Load.

Comments: This selection is to enable/disable the 10GbE PXE Option ROM that is used by all Onboard and IO Module 10 GbE controllers.

This option is grayed out and not accessible if the discs Option ROM is enabled or the 10 GbE FCoE Option ROM is enabled. It can co-exist with the 1 GbE PXE Option ROM or with an Infiniband* controller Option ROM.

If the 10GbE PXE Option ROM is disabled, and no other Option ROM is enabled, the system cannot perform a Network Boot and cannot respond for Wake-on-LAN.

This 10GbE PXE option does not appear unless there is a 10 GbE NIC installed in the system as an Onboard or IO Module NIC.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

4. FCoE 10GbE Option ROM

Option Values: Enabled
 Disabled

Help Text:

Enable/Disable Onboard/IOM NIC FCoE Option ROM Load.

Comments: This selection is to enable/disable the 10GbE FCoE Option ROM that is used by all Onboard and IO Module 10 GbE controllers capable of FCoE support. At the present time, only the Intel® 82599 10 Gigabit SFP+ NIC supports FCoE for this family of server boards.

This option is grayed out and not accessible if the 10GbE PXE Option ROM is enabled or if the discs Option ROM is enabled. It can co-exist with the 1GbE PXE Option ROM or with an Infiniband* controller Option ROM.

If the FCoE Option ROM is disabled, and no other Option ROM is enabled, the system cannot perform a Network Boot and cannot respond for Wake-on-LAN.

This FCoE option does not appear unless there is a FCoE-capable 10GbE NIC installed in the system as an Onboard or IO Module NIC.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

5. discs 1GbE/10GbE Option ROM

Option Values: Enabled
Disabled

Help Text:

Enable/Disable Onboard/IOM NIC discs Option ROM Load.

Comments: This selection is to enable/disable the discs Option ROM that is used by all Onboard and IO Module 1 GbE and 10 GbE controllers.

This option is grayed out and not accessible if the 1 GbE or 10GbE PXE Option ROM is enabled or if the 10 GbE FCoE Option ROM is enabled. It can co-exist with an Infiniband* controller Option ROM.

If the discs Option ROM is disabled, and no other Option ROM is enabled, the system cannot perform a Network Boot and cannot respond for Wake-on-LAN.

This discs option does not appear unless there is an discs-capable NIC installed in the system as an Onboard or IO Module NIC.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

6. Onboard NIC1 Type

Onboard NIC2 Type

Option Values: <Onboard NIC Description>

One of these strings:

*Intel® 82574 Single-Port Gigabit Ethernet Controller
Intel® I350 Dual-Port Gigabit Ethernet Controller
Intel® I350 Quad-Port Gigabit Ethernet Controller
Intel® I540 Dual-Port X540 10 Gigabit RJ-45 Controller
Mellanox* ConnectX-3* Single-Port Infiniband* FD14 Controller*

Help Text: <None>

Comments: Information only. This is a display showing which NICs are available as Network Controllers integrated into the baseboard. Each of these Onboard NICs will be followed by a section including a group of options that are specific to the type of NIC, either as an Ethernet controller or an Infiniband* controller.

If a board only has one onboard NIC, the second NIC Type and following options section will not appear. If there is an Infiniband* controller integrated onboard, it will appear as NIC2.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

7. IO Module 1 Type

IO Module 2 Type

Option Values: <IO Module Description>

One of these strings:

*Intel® I350 Quad-Port Gigabit Ethernet Module
Intel® I540 Dual-Port X540 10 Gigabit RJ-45 Module
Intel® 82599 Dual-Port 10 Gigabit SFP+ Module
Mellanox* ConnectX-3* Single-Port Infiniband* FD14 Module*

Help Text: <None>

Comments: *Information only.* This is a display showing which Network Controllers on IO Modules are installed on the baseboard. Each of these IO Module NICs will be followed by a section including a group of options that are specific to the type of NIC, either as an Ethernet controller or an Infiniband* controller.

This descriptive screen image shows an example of an InfiniBand* controller as IOM1. In a system with two IO Module connectors, an Infiniband* IO Module might be installed as either IOM1 or IOM2.

Most boards have only one IO Module connector. In any case, an IO Module Type and following options section will only appear when an IO Module is installed, and a second IO Module Type and options section will only appear if there are two IO Modules installed.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

8. NIC1 Controller NIC2 Controller

Option Values: **Enabled**
 Disabled

Help Text:

Enable/Disable Onboard Network Controller.

Comments: This will completely disable Onboard Network Controller NIC1 or NIC2, along with all included NIC Ports and their associated options. That controller's NIC Ports, Port PXE options, and Port MAC Address displays will not appear.

This option only appears for onboard Ethernet controllers. It does not appear for onboard Infiniband* controllers.

Ethernet controllers on IO Modules do not have a disabling function that can be controlled by BIOS, so there is no corresponding controller enable/disable option for an IOM Ethernet controller.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

9. NIC2 Infiniband* Option ROM 10. IOM1 Infiniband* Option ROM IOM2 Infiniband* Option ROM

Option Values: Enabled
 Disabled

Help Text:

Enable/Disable Infiniband Controller Option ROM and FlexBoot.*

Comments: This option will control whether the associated Infiniband* Controller Option ROM is executed by BIOS during POST. This will also control whether the Infiniband* controller FlexBoot program appears in the list of bootable devices.

This option only appears for Onboard or IO Module Infiniband* controllers. It does not appear for Ethernet controllers.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

11. NIC2 Port1 GUID 12. IOM1 Port1 GUID IOM2 Port1 GUID

Option Values: <GUID Display>

Help Text: <None>

Comments: *Information only.* 16 hex digits of the Port1 GUID of the Infiniband* controller for NIC2, IOM1, or IOM2.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

- 13. NIC1 Port1
- 14. NIC1 Port2
- 15. NIC1 Port3
- NIC1 Port4
- 16. NIC2 Port1
- 17. NIC2 Port2
- 18. NIC2 Port3
- NIC2 Port4

Option Values: **Enabled**
 Disabled

Help Text:

Enable/Disable Onboard NIC<n> Port<x>.

Comments: This will enable or disable Port<x, x = 1-4> of Onboard Network Controller<n, n = 1-2>, including associated Port PXE options. The NIC<n> Port<x> PXE option and MAC Address display will not appear when that port is disabled.

The associated port enable/disable options will not appear when NIC<n> is disabled.

Only ports which actually exist for a particular NIC will appear in this section. That is, Port1-Port4 will appear for a quad-port NIC, Port1-Port2 will appear for a dual-port NIC, and only Port1 will appear for a single-port NIC.

Network controllers installed on an IO Module do not have a port disabling function that is controlled by BIOS, so there are no corresponding options for IO Module NICs.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

- 19. NIC1 Port1 PXE
- 20. NIC1 Port2 PXE
- 21. NIC1 Port3 PXE
- NIC1 Port4 PXE
- 22. NIC2 Port1 PXE
- 23. NIC2 Port2 PXE
- 24. NIC2 Port3 PXE
- NIC2 Port4 PXE
- 25. IOM1 Port1 PXE
- 26. IOM1 Port2 PXE
- 27. IOM1 Port3 PXE
- IOM1 Port4 PXE
- 28. IOM2 Port1 PXE
- 29. IOM2 Port2 PXE
- 30. IOM2 Port3 PXE
- IOM2 Port4 PXE

Option Values: **Enabled**
 Disabled

Help Text:

Enable/Disable Onboard/IOM NIC Port PXE Boot

Comments: This will enable or disable PXE Boot capability for Port<x, x = 1-4> of Onboard NIC<n, n = 1-2> or IO Module<n, n = 1-2>.

This option will not appear for ports on a NIC which is disabled, or for individual ports when the corresponding NIC Port is disabled.

Only ports which actually exist for a particular NIC or IOM will appear in this section. That is, Port1-Port4 will appear for a quad-port NIC, Port1-Port2 will appear for a dual-port NIC, and only Port1 will appear for a single-port NIC.

The default state of each Port PXE Boot option is *Enabled*, if the corresponding PXE Boot OPROM of the same speed is *Enabled*. If a PXE Boot OPROM for 1 GbE or 10 GbE changes from *Disabled* to *Enabled*, then the Port PXE Boot option becomes *Enabled* for all ports of that speed

If the PXE Boot OPROM for 1 GbE NICs or 10 GbE NICs is disabled, PXE Boot will be disabled and grayed out as unchangeable for all ports on NICs or IO Modules of that same speed.

Conversely, if PXE Boot is disabled for all ports of a given speed, the corresponding PXE Option ROM will be disabled, but not grayed out since it could be selected.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

31. NIC1 Port1 MAC Address

32. NIC1 Port2 MAC Address

33. NIC1 Port3 MAC Address

NIC1 Port4 MAC Address

34. NIC2 Port1 MAC Address

35. NIC2 Port2 MAC Address

36. NIC2 Port3 MAC Address

NIC2 Port4 MAC Address

37. IOM1 Port1 MAC Address

38. IOM1 Port2 MAC Address

39. IOM1 Port3 MAC Address

IOM1 Port4 MAC Address

40. IOM2 Port1 MAC Address

41. IOM2 Port2 MAC Address

42. IOM2 Port3 MAC Address

IOM2 Port4 MAC Address

Option Values: <Mac Address Display>

Help Text: <None>

Comments: *Information only.* 12 hex digits of the MAC address of Port1- Port4 of the Network Controller corresponding to NIC1, NIC2, IOM1, or IOM2.

This display will appear only for ports which actually exist on the corresponding Network Controller. If the Network Controller or port is disabled, the port MAC Address will not appear.

[Back to \[NIC Configuration Screen\]](#) — [\[PCI Configuration Screen\]](#) — [\[Advanced Screen\]](#)

12.2.2.8 Serial Port Configuration

The Serial Port Configuration screen allows the user to configure the Serial A and Serial B ports. In Legacy ISA nomenclature, these are ports COM1 and COM2 respectively.

To access this screen from the **Main** screen, select **Advanced > Serial Port Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

The primary usage for these serial ports is to enable Serial Console Redirection and Serial Over LAN (SOL) capabilities. Either port can be used for Serial Console Redirection, but SOL is only supported on Serial A. See [Figure 55](#) for Console Redirection Configuration.

The exception to this is the W2600CR Workstation, which does not provide a Serial A port. With W2600CR, Serial A will not appear for configuration here, and Serial B will support SOL functionality if required.

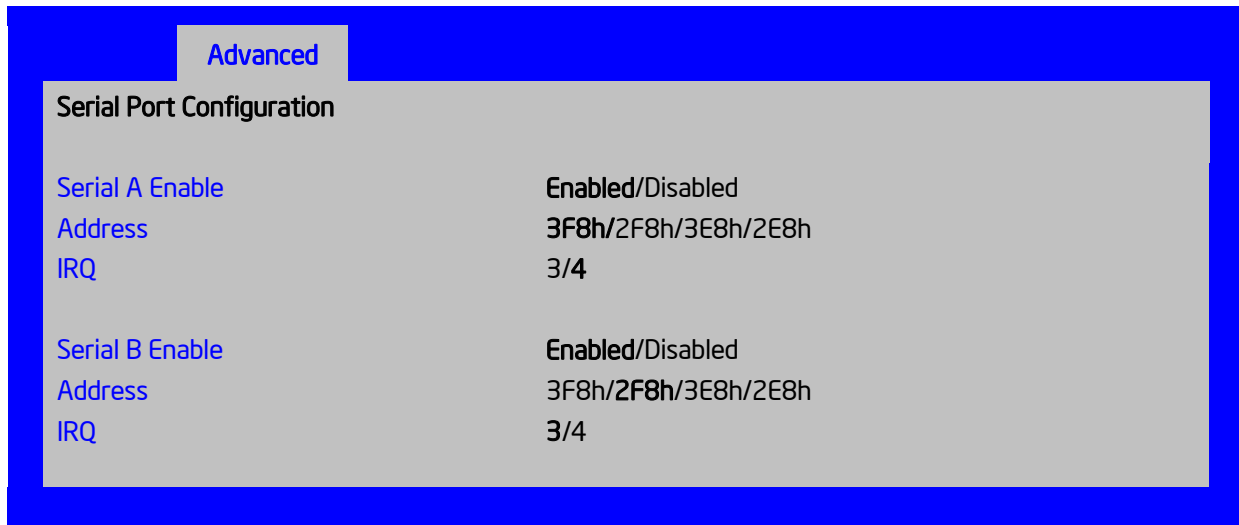


Figure 50. Serial Port Configuration Screen

Screen Field Descriptions:

1. Serial A Enable

Option Values: **Enabled**
 Disabled

Help Text:

Enable or Disable Serial port A.

Comments: Serial Port A can be used for either Serial Over LAN or Serial Console Redirection.

This Setup option should not appear on W2600CR, which does not provide a Serial A port.

[Back to \[Serial Port Configuration Screen\]](#)

2. Address

Option Values: **3F8h**
 2F8h
 3E8h
 2E8h

Help Text:

Select Serial port A base I/O address.

Comments: Legacy I/O port address. This field should not appear when Serial A port enable/disable does not appear.

[Back to \[Serial Port Configuration Screen\]](#)

3. IRQ

Option Values: 3
 4

Help Text:

Select Serial port A interrupt request (IRQ) line.

Comments: Legacy IRQ. This field should not appear when Serial A port enable/disable does not appear.

[Back to \[Serial Port Configuration Screen\]](#)

4. Serial B Enable

Option Values: **Enabled**
 Disabled

Help Text:

Enable or Disable Serial port B.

Comments: Serial Port B can be used for Serial Console Redirection. SOL cannot be routed to Serial B except on W2600CR boards, which do not have a Serial A port.

Back to [Serial Port Configuration Screen]

5. Address

Option Values: 3F8h
 2F8h
 3E8h
 2E8h

Help Text:

Select Serial port B base I/O address.

Comments: Legacy I/O port address.

Back to [Serial Port Configuration Screen]

6. IRQ

Option Values: **3**
 4

Help Text:

Select Serial port B interrupt request (IRQ) line.

Comments: Legacy IRQ

Back to [Serial Port Configuration Screen]

12.2.2.9 USB Configuration

The USB Configuration screen allows the user to configure the available USB controller options.

To access this screen from the **Main** screen, select **Advanced > USB Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

This screen should display all USB Mass Storage devices which have been detected in the system. These include USB-attached Hard Disk Drives (HDDs), Floppy Disk Drives (FDDs), CDROM and DVDROM drives, and USB Flash Memory devices (USB Key, Keyfob, and so on).

Each USB Mass Storage device may be set to allow the media emulation for which it is formatted, or an emulation may be specified. For USB Flash Memory devices in particular, there are some restrictions:

- A USB Key formatted as a CDROM drive will be recognized as an HDD.
- A USB Key formatted without a Partition Table will be forced to FDD emulation.
- A USB Key formatted with one Partition Table, and less than 528 MB in size, will be forced to FDD emulation – otherwise if it is 528 MB or greater in size, it will be forced to HDD emulation.

Note: USB devices can be “hotplugged” during POST, and will be detected and “beeped”. They will be enumerated and displayed on this screen, though they may not be enumerated as bootable devices.

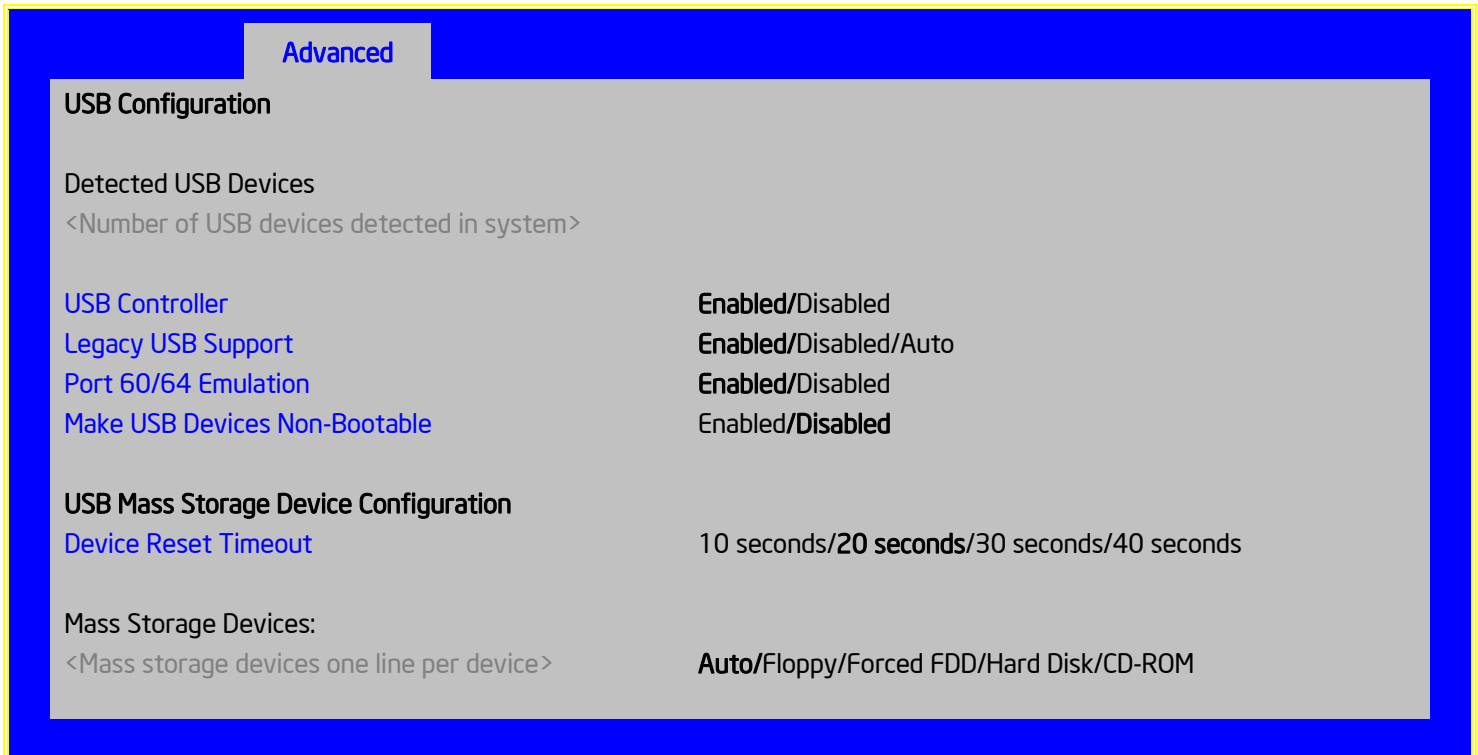


Figure 51. USB Configuration Screen

Screen Field Descriptions:

1. Detected USB Devices

Option Values: <Number of USB devices detected in system>

Help Text: <None>

Comments: Information only. Displays the total number of USB devices of all types which have been detected in POST.

[Back to \[USB Configuration Screen\]](#)

2. USB Controller

Option Values: **Enabled**
Disabled

Help Text:

[Enabled] - All on-board USB controllers are turned on and accessible by the OS.

[Disabled] - All on-board USB controllers are turned off and inaccessible by the OS.

Comments: When the USB controllers are Disabled, there is no USB IO available for either POST or the OS. In that case, all following fields on this screen are grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

3. Legacy USB Support

Option Values: **Enabled**
 Disabled
 Auto

Help Text:

Enables Legacy USB support. AUTO option disables legacy support if no USB devices are connected. Disable option will only keep USB Keyboard devices available for EFI applications.

Comments: When Legacy USB Support is Disabled, USB devices are available only through OS drivers.

If the USB controller setting is Disabled, this field is grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

4. Port 60/64 Emulation

Option Values: **Enabled**
 Disabled

Help Text:

*Enables I/O port 60h/64h emulation support.
This may be needed for legacy USB keyboard support when using an OS that is USB unaware.*

Comments: If the USB controller setting is Disabled, this field is grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

5. Make USB Devices Non-Bootable

Option Values: **Enabled**
 Disabled

Help Text:

*Exclude USB in Boot Table.
[Enabled]- This will remove all USB Mass Storage devices as Boot options.
[Disabled] - This will allow all USB Mass Storage devices as Boot options.*

Comments: This is a security option. When Disabled, the system cannot be booted directly to a USB device of any kind. USB Mass Storage devices may still be used for data storage.

If the USB controller setting is Disabled, this field is grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

6. Device Reset Timeout

Option Values: 10 seconds
 20 seconds
 30 seconds
 40 seconds

Help Text:

*USB Mass Storage device Start Unit command timeout.
Setting to a larger value provides more time for a mass storage device to be ready, if needed.*

Comments: If the USB controller setting is Disabled, this field is grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

7. Mass Storage Devices:

Option Values: **Auto**
 Floppy
 Forced FDD
 Hard Disk
 CD-ROM

Help Text:

[Auto] - USB devices less than 530 MB are emulated as floppies.

[Forced FDD] - HDD formatted drive is emulated as an FDD (for example, ZIP drive).

Comments: This field is hidden if no USB Mass Storage devices are detected.

This setup screen can show a maximum of eight USB Mass Storage devices on the screen. If more than eight devices are installed in the system, the 'USB Devices Enabled' displays the correct count, but only the first eight devices discovered are displayed in this list.

If the USB controller setting is Disabled, this field is grayed out and inactive.

[Back to \[USB Configuration Screen\]](#)

12.2.2.10 System Acoustic and Performance Configuration

The System Acoustic and Performance Configuration screen allows the user to configure the thermal control behavior of the system with respect to what parameters are used in the system's Fan Speed Control algorithms.

To access this screen from the **Main** screen, select **Advanced > System Acoustic and Performance Configuration**. To move to another screen, press the <Esc> key to return to the **Advanced** screen, then select the desired screen.

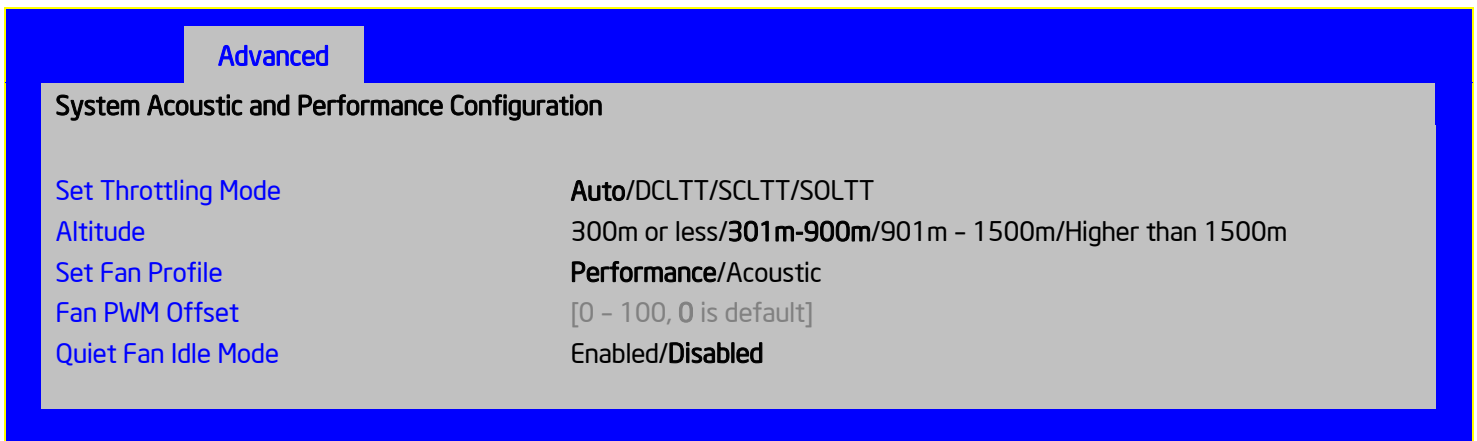


Figure 52. System Acoustic and Performance Configuration

Screen Field Descriptions:

Set Throttling Mode

Option Values: **Auto**
 DCLTT
 SCLTT
 SOLTT

Help Text:

Sets Thermal Throttling mode for memory, to control fans and DRAM power as needed to control DIMM temperatures.

[Auto] – BIOS selects mode. 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.

[SCLTT] – Static Closed Loop Thermal Throttling.

[SOLTT] – Static Open Loop Thermal Throttling.

Comments: The Thermal Throttling Mode chosen reflects whether the DIMMs have Temperature Sensors (TSOD), and whether the chassis is an Intel chassis for which thermal data are available. Note that this is for thermal throttling only, independent of any controls imposed for the purpose of power limiting.

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.

- ▶ DCLTT is the expected mode for a board in an Intel chassis with inlet and outlet air temperature sensors and TSOD. The firmware can update the offset registers for closed loop during runtime, as BIOS sends the dynamic CLTT offset temperature data.
- ▶ SCLTT would be used with an OEM chassis and DIMMs with TSOD. The firmware does not change the offset registers for closed loop during runtime, although the Management Engine can do so.

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.

- ▶ SOLTT is intended for a system with UDIMMs which do not have TSOD. The thermal control registers are configured during POST, and the firmware does not change them.

[Back to \[System Acoustic and Performance Configuration\]](#)

1. Altitude

Option Values: 300m or less

301m-900m

901m-1500m

Higher than 1500m

Help Text:

[300m or less](980ft or less) Optimal near sea level.

[301m-900m](980ft-2950ft) Optimal performance setting at moderate elevation.

[901m-1500m](2950ft-4920ft) Optimal performance setting at high elevation.

[Above 1500m](above 4920ft) Optimal performance setting at the highest elevations.

Comments: This option sets an altitude value in order to choose a Fan Profile that is optimized for the air density at the current altitude at which the system is installed.

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.

[Back to \[System Acoustic and Performance Configuration\]](#)

2. Set Fan Profile

Option Values: **Performance**
 Acoustic

Help Text:

[Performance] - Fan control provides primary system cooling before attempting to throttle memory.

[Acoustic] - The system will favor using throttling of memory over boosting fans to cool the system if thermal thresholds are met.

Comments: This option allows the user to choose a Fan Profile that is optimized for maximizing performance or for minimizing acoustic noise.

When Performance is selected, the thermal conditions in the system are controlled by raising fan speed when necessary to raise cooling performance. This provides cooling without impacting system performance, but may impact system acoustic performance – fans running faster are typically louder. The Performance mode is designed to provide sufficient cooling capability covering all kinds of add-in cards on the market.

When Acoustic is selected, then rather than increasing fan speed for additional cooling, the system will attempt first to control thermal conditions by throttling memory to reduce heat production. This regulates the system's thermal condition without changing the acoustic performance, but throttling memory may impact system performance. The Acoustic mode offers best acoustic experience and appropriate cooling capability covering mainstream and majority of the add-in cards.

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.

[Back to \[System Acoustic and Performance Configuration\]](#)

3. Fan PWM Offset

Option Values: *[Entry Field 0 – 100, 0 is default]*

Help Text:

Valid Offset 0 - 100. This number is added to the calculated PWM value to increase Fan Speed.

Comments: This is a percentage by which the calculated fan speed will be increased. The user can apply positive offsets that result in increasing the minimum fan speeds. This feature is valid when Quiet Fan Idle Mode is at Enabled state.

[Back to \[System Acoustic and Performance Configuration\]](#)

4. Quiet Fan Idle Mode

Option Values: Enabled
 Disabled

Help Text:

Enabling this option allows the system fans to operate in Quiet 'Fan off' mode while still maintaining sufficient system cooling. In this mode, fan sensors become unavailable and cannot be monitored. There will be limited fan related event generation.

Comments: When enabled, this option allows fans to idle or turn off when sufficient thermal margin is available, decreasing the acoustic noise produced by the system and decreasing system power consumption. Fans will run as needed to maintain thermal control. The actual decrease in fan speed depends on the system thermal loading, which in turn depends on system configuration and workload.

While Quiet Fan Idle Mode is engaged, fan sensors become unavailable and are not monitored by the BMC.

Quiet Fan Idle Mode does not conflict with Fan PWM Offset (above) – they work in concert, with Fan PWM Offset applied to fans in Quiet Fan Idle Mode just as when the fans are operating in “normal mode”. A Fan PWM Offset of zero is necessary for fans to actually stop turning.

[Back to \[System Acoustic and Performance Configuration\]](#)

12.2.3 Security Screen (Tab)

The Security screen allows the user to enable and set the Administrator and User passwords and to lock out the front panel buttons so they cannot be used. This screen also allows the user to enable and activate the Trusted Platform Module (TPM) security settings on those boards that support TPM.

Note that it is necessary to activate the TPM in order to be able to enable Intel® Trusted Execution Technology (TXT) on boards that support it. Changing the TPM state in Setup will require a Hard Reset for the new state to become effective. For enabling Intel® TXT, see the Processor Configuration screen.

This BIOS supports (but does not require) “Strong Passwords” for security. The “Strong Password” criteria for both Administrator and User passwords require that passwords be between 8 and 14 characters in length, and a password must contain at least one case-sensitive alphabetic character, one numeric character, and one special character. A warning is given when a password is set which does not meet the Strong Password criteria, but the password is accepted.

For further security, the BIOS optionally may require a Power on Password to be entered in early POST in order to boot the system. When Power On Password is enabled, POST is halted soon after power on while the BIOS queries for a Power On Password. Either the Administrator or the User password may be entered for a Power on Password.

To access this screen from the **Main** screen or other top-level “Tab” screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Security** screen is selected.

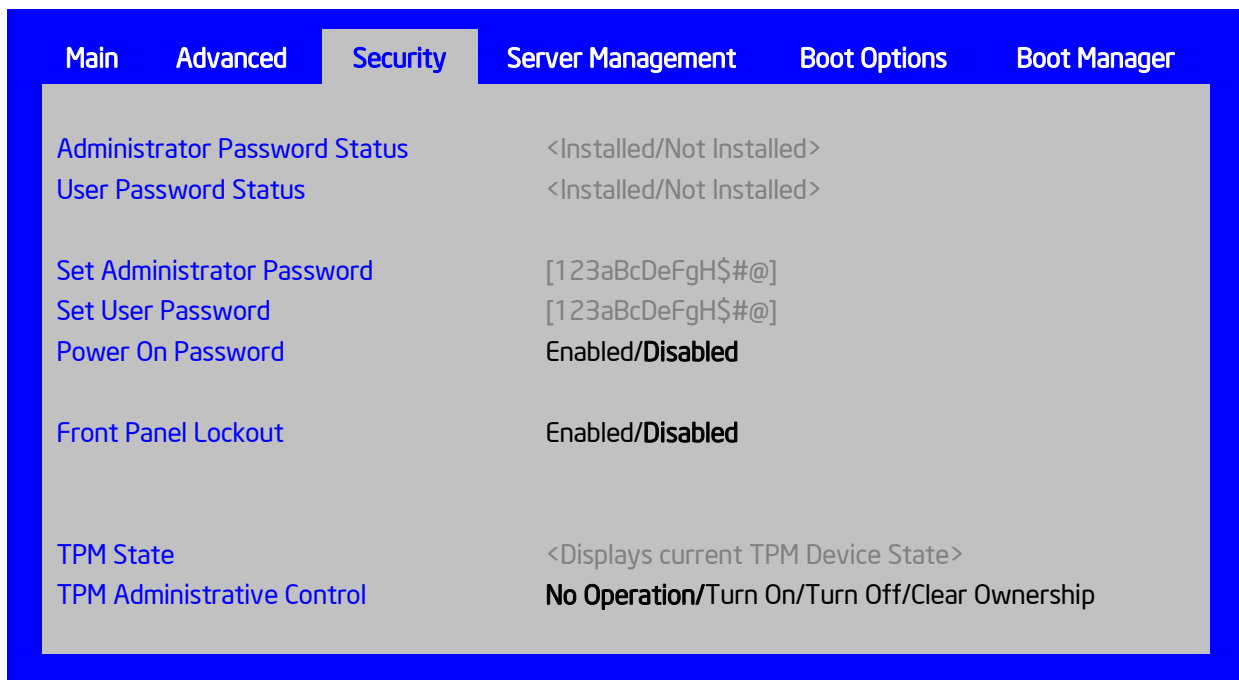


Figure 53. Security Screen

Screen Field Descriptions:**1. Administrator Password Status**

Option Values: *Installed*
 Not Installed

Help Text: <None>

Comments: Information only. Indicates the status of the Administrator Password.

[Back to \[Security Screen\]](#)

2. User Password Status

Option Values: *Installed*
 Not Installed

Help Text: <None>

Comments: Information only. Indicates the status of the User Password.

[Back to \[Security Screen\]](#)

3. Set Administrator Password

Option Values: *[Entry Field – 0-14 characters]*

Help Text:

Administrator password is used if Power On Password is enabled and to control change access in BIOS Setup. Length is 1-14 characters. Case sensitive alphabetic, numeric and special characters !@#\$\$%^&()-_+=? are allowed.*

Note: *Administrator password must be set in order to use the User account.*

Comments: This password controls “change” access to Setup. The Administrator has full access to change settings for any Setup options, including setting the Administrator and User passwords.

When Power On Password protection is enabled, the Administrator password may be used to allow the BIOS to complete POST and boot the system.

Deleting all characters in the password entry field removes a password previously set. Clearing the Administrator Password also clears the User Password.

If invalid characters are present in the password entered, it will not be accepted, and there will be popup error message:

Password entered is not valid. Only case sensitive alphabetic, numeric and special characters !@#\$\$%^&()-_+=? are allowed.*

The Administrator and User passwords must be different. If the password entered is the same as the User password, it will not be accepted, and there will be popup error message:

Password entered is not valid. Administrator and User passwords must be different.

Strong passwords are encouraged, although not mandatory. If a password is entered which does not meet the “Strong Password” criteria, there will be a popup warning message:

Warning – a Strong Password should include at least one each case sensitive alphabetic, numeric, and special character. Length should be 8 to 14 characters.

[Back to \[Security Screen\]](#)

4. Set User Password

Option Values: *[Entry Field – 0-14 characters]*

Help Text:

User password is used if Power On Password is enabled and to allow restricted access to BIOS Setup. Length is 1-14 characters. Case sensitive alphabetic, numeric and special characters !@#\$%^&()-_+=? are allowed.*

Note: *Removing the administrator password also removes the user password.*

Comments: The User password is available only if the Administrator Password has been installed. This option protects Setup settings as well as boot choices. The User Password only allows limited access to the Setup options, and no choice of boot devices.

When Power On Password protection is enabled, the User password may be used to allow the BIOS to complete POST and boot the system.

The password format and entry rules and popup error and warning message are the same for the User password as for the Administrator password (see above).

[Back to \[Security Screen\]](#)

Power On Password

Option Values: Enabled
Disabled

Help Text:

Enable Power On Password support. If enabled, password entry is required in order to boot the system.

Comments: When Power On Password security is enabled, the system will halt soon after power on and the BIOS will ask for a password before continuing POST and booting. Either the Administrator or User password may be used.

If an Administrator password has not been set, this option will be grayed out and unavailable. If this option is enabled and the Administrator password is removed, that will also disable this option.

[Back to \[Security Screen\]](#)

5. Front Panel Lockout

Option Values: Enabled
Disabled

Help Text:

If enabled, locks the power button OFF function and the reset and NMI Diagnostic Interrupt buttons on the system's front panel. If [Enabled] is selected, power off and reset must be controlled from a system management interface, and the NMI Diagnostic Interrupt is not available.

Comments: **Note:** This option does not appear on all boards.

[Back to \[Security Screen\]](#)

6. TPM State

Option Values: <Displays current TPM Device State>

May be:
Enabled & Activated
Enabled & Deactivated
Disabled & Activated
Disabled & Deactivated

Help Text: <None>

Comments: Information only. Shows the current TPM device state.

A **Disabled** TPM device does not execute commands that use the TPM functions and TPM security operations are not available.

An **Enabled & Deactivated** TPM is in the same state as a disabled TPM, except that setting of the TPM ownership is allowed if it is not present already.

An **Enabled & Activated** TPM executes all commands that use the TPM functions and TPM security operations are also available.

Note: This option appears only on boards equipped with a TPM.

Back to [Security Screen]

7. TPM Administrative Control

Option Values: **No Operation**

Turn On
 Turn Off
 Clear Ownership

Help Text:

[No Operation] - No changes to current state.

[Turn On] - Enables and activates TPM.

[Turn Off] - Disables and deactivates TPM.

[Clear Ownership] - Removes TPM ownership & returns TPM to factory default state.

Note: Setting returns to *[No Operation]* on every boot.

Comments: Any Administrative Control operation selected will require the system to perform a Hard Reset in order to become effective.

Note: This option appears only on boards equipped with a TPM.

Back to [Security Screen]

12.2.4 Server Management Screen (Tab)

The Server Management screen allows the user to configure several server management features. This screen also provides an access point to the screens for configuring Console Redirection, displaying system information, and controlling the BMC LAN configuration.

To access this screen from the **Main** screen or other top-level "Tab" screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Server Management** screen is selected.



Figure 54. Server Management Screen

Screen Field Descriptions:

1. Assert NMI on SERR

Option Values: **Enabled**
 Disabled

Help Text:

On SERR, generate an NMI and log an error.

Note: *[Enabled] must be selected for the Assert NMI on PERR setup option to be visible.*

Comments: This option allows the system to generate an NMI when an SERR occurs, which is a method Legacy Operating System error handlers may use instead of processing a Machine Check.

[Back to \[Server Management Screen\]](#)

2. Assert NMI on PERR

Option Values: **Enabled**
 Disabled

Help Text:

On PERR, generate an NMI and log an error.

Note: *This option is only active if the Assert NMI on SERR option has [Enabled] selected.*

Comments: This option allows the system to generate an NMI when a PERR occurs, which is a method Legacy Operating System error handlers may use instead of processing a Machine Check.

[Back to \[Server Management Screen\]](#)

3. PCIe AER Support

Option Values: **Enabled**
 Disabled

Help Text:

[Enabled] – PCIe AER (Advanced Error Reporting) is enabled.

[Disabled] – PCIe AER is disabled. All PCIe AER errors will be masked once PCIe AER is disabled.

Comments: This option allows the system to monitor and handle PCIe AER errors on PCIe devices with PCIe AER support. But as PCIe AER is described in PCI Express Base Specification, any third-party software or OS could override this BIOS policy and take ownership of PCIe AER handling after BIOS POST.

[Back to \[Server Management Screen\]](#)

4. Log Correctable Errors

Option Values: Enabled
 Disabled

Help Text:

[Enabled] – Processor & PCH PCIe correctable error logging is enabled.

[Disabled]– Processor & PCH PCIe correctable error logging is disabled.

Comments: This option allows the system to monitor and handle PCIe correctable errors on PCIe devices behind Processor and PCH.

[Back to \[Server Management Screen\]](#)

5. Reset on CATERR

Option Values: **Enabled**
 Disabled

Help Text:

When enabled system gets reset upon encountering Catastrophic Error (CATERR); when disabled system does not get reset on CATERR.

Comments: This option controls whether the system will be reset when the “Catastrophic Error” CATERR# signal is held asserted, rather than just pulsed to generate an SMI. This indicates that the processor has encountered a fatal hardware error.

Note: If “Reset on CATERR” is *Disabled*, this can result in a system hang for certain error conditions, possibly with the system unable to update the System Status LED or log an error to the SEL before hanging.

[Back to \[Server Management Screen\]](#)

6. Reset on ERR2

Option Values: **Enabled**
 Disabled

Help Text:

When enabled system gets reset upon encountering ERR2 (Fatal error); when disabled system does not get reset on ERR2

Comments: This option controls whether the system will be reset if the BMC's ERR2 Monitor times out, that is, the ERR2 signal has been continuously asserted long enough to indicate that the SMI Handler is not able to service the condition

Note: If "Reset on ERR2" is *Disabled*, this can result in a system hang for certain error conditions, possibly with the system unable to update the System Status LED or log an error to the SEL before hanging.

[Back to \[Server Management Screen\]](#)

7. Resume on AC Power Loss

Option Values: **Stay Off**
 Last State
 Power On

Help Text:

*System action to take on AC power loss recovery.
 [Stay Off] - System stays off.
 [Last State] - System returns to the same state before the AC power loss.
 [Power On] - System powers on.*

Comments: This option controls the policy that the BMC will follow when AC power is restored after an unexpected power outage. The BMC will either hold DC power off or always turn it on to boot the system, depending on this setting – and in the case of **Last State**, depending on whether the power was on and the system was running before the AC power went off.

When this setting is changed in Setup, the new setting will be sent to the BMC. However, the BMC maintains ("owns") this Power Restore Policy setting, and it can be changed independently with an IPMI command to the BMC. BIOS gets this setting from the BMC early in POST, and also for the Setup Server Management screen.

[Back to \[Server Management Screen\]](#)

8. Power Restore Delay

Option Values: **Disabled**
 Auto
 Fixed

Help Text:

Allows a delay in powering up after a power failure, to reduce peak power requirements. The delay can be fixed or automatic between 25-300 seconds.

Comments: When the AC power resume policy (above) is either **Power On** or **Last State**, this option allows a delay to be taken after AC power is restored before the system actually begins to power up. This delay can be either a fixed time or an "automatic" time, where "automatic" means that the BIOS will select a randomized delay time of 25-300 seconds when it sends the Power Restore Delay setting to the BMC.

This option will be grayed out and unavailable when the AC power resume policy is **Stay Off**.

The Power Restore Delay setting is maintained by BIOS. This setting does not take effect until a reboot is done. Early in POST, the Power Restore Policy is read from the BMC, and if the policy is **Power On** or **Last State**, the delay settings are sent to the BMC.

Bear in mind that even if the Power Restore Delay is **Disabled**, there will still be a delay of about 20 seconds while the BMC itself boots up after AC power is restored.

Note: This Power Restore Delay option applies only to powering on when AC is applied. It has no effect on powering the system up using the Power Button on the Front Panel. A DC power on using the Power Button is not delayed.

The purpose of this delay is to avoid having all systems draw “startup surge” power at the same time. Different systems or racks of systems can be set to different delay times to spread out the startup power draws. Alternatively, all systems can be set to Automatic, and then each system will wait for a random period before powering up.

[Back to \[Server Management Screen\]](#)

9. Power Restore Delay Value

Option Values: [Entry Field 25 – 300, **25** is default]

Help Text:

Fixed time period 25-300 seconds for Power Restore Delay.

Comments: When the power restore policy is **Power On** or **Last State**, and the Power Restore Delay selection is **Fixed**, this field allows for specifying how long in seconds that fixed delay will be.

When the Power Restore Delay is **Disabled** or **Auto**, this field will be grayed out and unavailable.

The **Power Restore Delay Value** setting is maintained by BIOS. This setting does not take effect until a reboot is done. Early in POST, the Power Restore Policy is read from the BMC, and if the policy is **Power On** or **Last State**, the delay settings are sent to the BMC. When the **Power Restore Delay** setting is **Fixed**, this delay value is used to provide the length of the delay.

[Back to \[Server Management Screen\]](#)

10. Clear System Event Log

Option Values: Enabled
Disabled

Help Text:

If enabled, clears the System Event Log. All current entries will be lost.

Note: This option is reset to [Disabled] after a reboot.

Comments: This option sends a message to the BMC to request it to clear the System Event Log. The log will be cleared, and then the “Clear” action itself will be logged as an event. This give the user a time/date for when the log was cleared.

[Back to \[Server Management Screen\]](#)

11. FRB-2 Enable

Option Values: **Enabled**
Disabled

Help Text:

Fault Resilient Boot (FRB).

BIOS programs the BMC watchdog timer for approximately 6 minutes. If BIOS does not complete POST before the timer expires, the BMC will reset the system.

Comments: This option controls whether the system will be reset if the BMC Watchdog Timer detects what appears to be a hang during POST. When the BMC Watchdog Timer is purposed as an FRB-2 timer, it is initially set to allow 6 minutes for POST to complete.

However, the FRB-2 Timer is suspended during times when some lengthy operations are in progress, like executing Option ROMs, during Setup, and when BIOS is waiting for a password. or for input to the F6 BBS Boot Menu. The FRB-2 Timer is also suspended while POST is paused with the <Pause> key.

[Back to \[Server Management Screen\]](#)

12. OS Boot Watchdog Timer

Option Values: Enabled
 Disabled

Help Text:

*BIOS programs the watchdog timer with the timeout value selected. If the OS does not complete booting before the timer expires, the BMC will reset the system and an error will be logged.
Requires OS support or Intel Management Software Support.*

Comments: This option controls whether the system will set the BMC Watchdog to detect an apparent to be a hang during OS booting. BIOS sets the timer before starting the OS bootstrap load procedure. If the OS Load Watchdog Timer times out, then presumably the OS failed to boot properly.

If the OS does boot up successfully, it must be aware of the OS Load Watchdog Timer and immediately turn it off before it expires. The OS may turn off the timer, or more often the timer may be repurposed as an OS Watchdog Timer to protect against runtime OS hangs.

Unless the OS does have timer-aware software to support the OS Load Watchdog Timer, the system will be unable to boot successfully with the OS Load Watchdog Timer enabled. When the timer expires without having been reset or turned off, the system will either reset or power off repeatedly.

[Back to \[Server Management Screen\]](#)

13. OS Boot Watchdog Timer Policy

Option Values: **Power off**
 Reset

Help Text:

*If the OS watchdog timer is enabled, this is the system action taken if the watchdog timer expires.
[Reset] - System performs a reset.
[Power Off] - System powers off.*

Comments: This option is grayed out and unavailable when the O/S Boot Watchdog Timer is disabled.

[Back to \[Server Management Screen\]](#)

14. OS Boot Watchdog Timer Timeout

Option Values: 5 minutes
 10 minutes
 15 minutes
 20 minutes

Help Text:

If the OS watchdog timer is enabled, this is the timeout value BIOS will use to configure the watchdog timer.

Comments: This option is grayed out and unavailable when the O/S Boot Watchdog Timer is disabled.

[Back to \[Server Management Screen\]](#)

15. Plug & Play BMC Detection

Option Values: Enabled
 Disabled

Help Text:

If enabled, the BMC will be detectable by OSes which support plug and play loading of an IPMI driver. Do not enable this option if your OS does not support this driver.

Comments: This option controls whether the OS Server Management Software will be able to find the BMC and automatically load the correct IPMI support software for it. If your OS does not support Plug & Play for the BMC, you will not have the correct IPMI driver software loaded.

[Back to \[Server Management Screen\]](#)

EuP LOT6 Off-Mode

Option Values: Enabled
 Disabled

Help Text:

Enable/disable Ecodesign EuP LOT6 "Deep Sleep" Off-Mode for near-zero energy use when powered off.

Comments: This option controls whether the system goes into "Deep Sleep" or more conventional S5 "Soft-Off" when powered off. "Deep Sleep" state uses less energy than S5, but S5 can start up faster and can allow a Wake on LAN action (which cannot be done from a Deep Sleep state).

This option will not appear on platforms which do not support EuP LOT6 Off-Mode.

[Back to \[Server Management Screen\]](#)

16. Shutdown Policy

Option Values: Enabled
 Disabled

Help Text: *Enable/Disable Shutdown Policy.*

Comments: This option is designed for multiple-node system and to control the policy that BMC should shutdown one node if it detected over-current or over-temperature condition. The BIOS and the BMC will synchronize the policy during the BIOS POST and current value of the BMC will be displayed in BIOS Setup.

This option is only displayed when the BMC support this feature on the node.

[Back to \[Server Management Screen\]](#)

17. Console Redirection

Option Values: <None>

Help Text: *View/Configure Console Redirection information and settings.*

Comments: Selection only. Position to this line and press the <Enter> key to go to the **[Console Redirection](#)** group of configuration settings.

[Back to \[Server Management Screen\]](#)

18. System Information

Option Values: <None>

Help Text: *View System Information.*Comments: Selection only. Position to this line and press the <Enter> key to go to the **System Information** group of configuration settings.**[Back to \[Server Management Screen\]](#)****19. BMC LAN Configuration**

Option Values: <None>

Help Text: *View/Configure BMC LAN and user settings.*Comments: Selection only. Position to this line and press the <Enter> key to go to the **BMC LAN Configuration** group of configuration settings.**[Back to \[Server Management Screen\]](#)****12.2.4.1 Console Redirection**

The Console Redirection screen allows the user to enable or disable Console Redirection for Remote System Management, and to configure the connection options for this feature.

To access this screen from the **Main** screen, select **Server Management > Console Redirection**. To move to another screen, press the <Esc> key to return to the **Server Management** screen, then select the desired screen.

When Console Redirection is active, all POST and Setup displays are in Text Mode. The Quiet Boot setting is disregarded, and the Text Mode POST Diagnostic Screen will be displayed regardless of the Quiet Boot setting. This is due to the limitations of Console Redirection, which is based on data terminal emulation using a serial data interface to transfer character data.

Console Redirection can use either of the two Serial Ports provided by the SuperIO in the BMC. However, if Console Redirection is to be coordinated with Serial Over LAN, the user should be aware that SOL is only supported through Serial Port A (except for W200CR, which only has Serial B and supports SOL on Serial B).

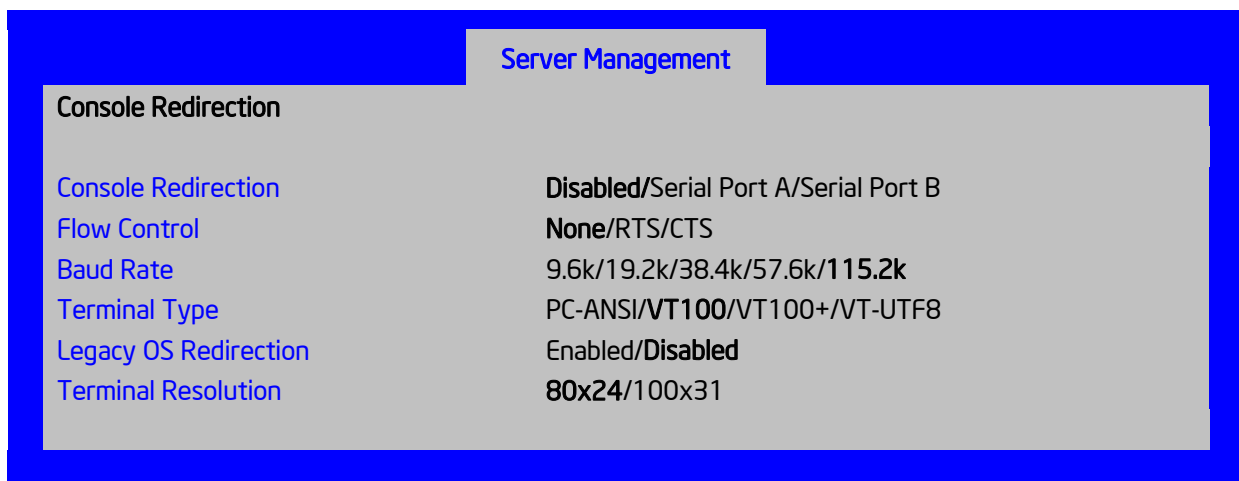


Figure 55. Console Redirection Screen

Screen Field Descriptions:**1. Console Redirection**

Option Values: **Disabled**
 Serial Port A
 Serial Port B

Help Text:

Console redirection allows a serial port to be used for server management tasks.

[Disabled] - No console redirection.

[Serial Port A] - Configure serial port A for console redirection.

Enabling this option will disable display of the Quiet Boot logo screen during POST.

Comments: Serial Console Redirection can use either Serial Port A or Serial Port B. If SOL is also going to be configured, note that SOL is only supported through Serial Port A (with the exception of W2600CR, which only has Serial B so supports SOL on Serial B).

When Console Redirection is set to Disabled, all other options on this screen will be grayed out and unavailable.

Only Serial Ports which are Enabled should be available to choose for Console Redirection. If neither Serial A nor Serial B is set to Enabled, then Console Redirection will be forced to Disabled, and grayed out as inactive. In that case, all other options on this screen will also be grayed

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

2. Flow Control

Option Values: **None**
 RTS/CTS

Help Text:

Flow control is the handshake protocol.

This setting must match the remote terminal application.

[None] - Configure for no flow control.

[RTS/CTS] - Configure for hardware flow control.

Comments: Flow control is necessary only when there is a possibility of data overrun. In that case the Request To Send/Clear to Send (RTS/CTS) hardware handshake is a relatively conservative protocol which can usually be configured at both ends.

When Console Redirection is set to Disabled, this option will be grayed out and unavailable.

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

3. Baud Rate

Option Values: 9.6k
 19.2k
 38.4k
 57.6k
 115.2k

Help Text:

Serial port transmission speed. This setting must match the remote terminal application.

Comments: In most modern Server Management applications, serial data transfer is consolidated over an alternative faster medium like LAN, and 115.2k is the speed of choice.

When Console Redirection is set to *Disabled*, this option will be grayed out and unavailable.

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

4. Terminal Type

Option Values: PC-ANSI

VT100
VT100+
VT-UTF8

Help Text:

Character formatting used for console redirection. This setting must match the remote terminal application.

Comments: The VT100 and VT100+ terminal emulations are essentially the same. VT-UTF8 is a UTF8 encoding of VT100+. PC-ANSI is the native character encoding used by PC-compatible applications and emulators.

When Console Redirection is set to *Disabled*, this option will be grayed out and unavailable.

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

5. Legacy OS Redirection

Option Values: Enabled

Disabled

Help Text:

This option enables legacy OS redirection (that is, DOS) on serial port. If it is enabled, the associated serial port is hidden from the legacy OS.

Comments: Operating Systems which are “redirection-aware” implement their own Console Redirection mechanisms. For a Legacy OS which is not “aware”, this option allows the BIOS to handle redirection.

When Console Redirection is set to *Disabled*, this option will be grayed out and unavailable.

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

6. Terminal Resolution

Option Values: **80x24**

100x31

Help Text:

Remote Terminal Resolution

Comments: This option allows the use of a larger terminal screen area, although it does not change Setup displays to match.

When Console Redirection is set to *Disabled*, this option will be grayed out and unavailable.

[Back to \[Console Redirection Screen\]](#) — [\[Server Management Screen\]](#)

12.2.4.2 System Information

The System Information screen allows the user to view part numbers, serial numbers, and firmware revisions. This is an *Information Only* screen.

To access this screen from the **Main** screen, select **Server Management > System Information**. To move to another screen, press the <Esc> key to return to the **Server Management** screen, then select the desired screen.

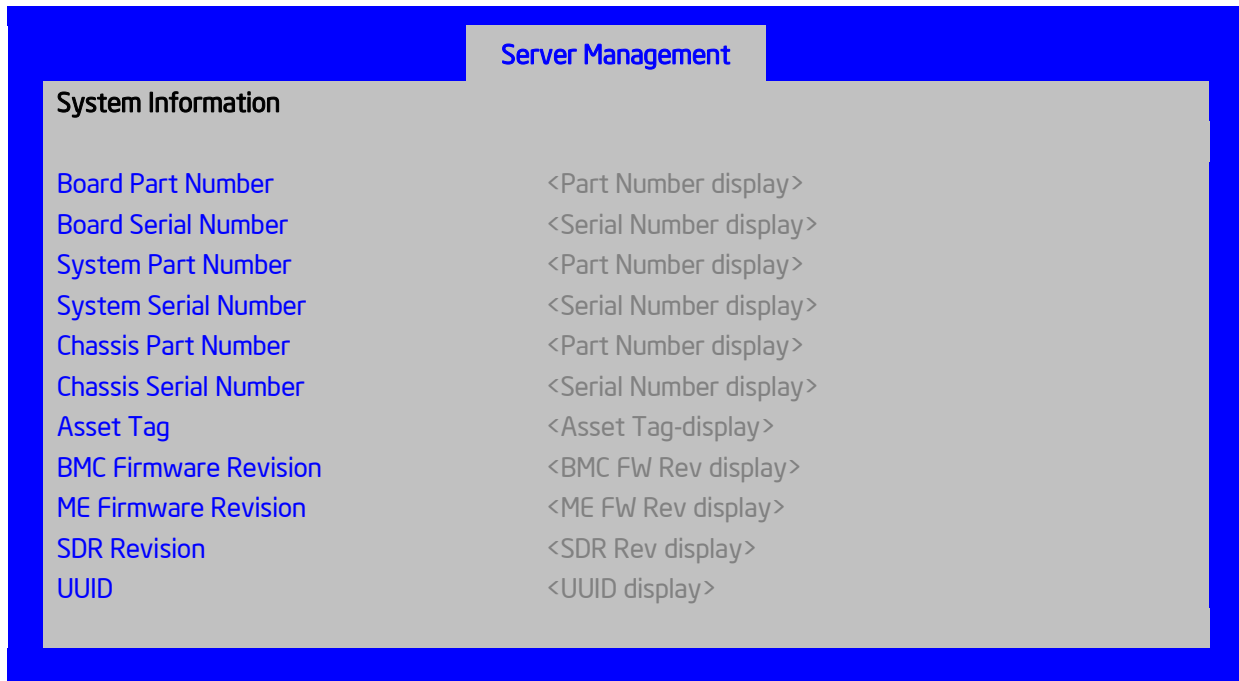


Figure 56. System Information Screen

Screen Field Descriptions:

1. Board Part Number

Option Values: <Part Number display>

Help Text: <None>

Comments: Information only.

[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)

2. Board Serial Number

Option Values: <Serial Number display>

Help Text: <None>

Comments: Information only.

[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)

3. System Part Number

Option Values: <Part Number display>

Help Text: <None>

Comments: Information only.

[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)

4. System Serial Number

Option Values: <Serial Number display>

Help Text: <None>

Comments: Information only.

[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)

5. Chassis Part Number

Option Values: <Part Number display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

6. Chassis Serial Number

Option Values: <Serial Number display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

7. Asset Tag

Option Values: <Asset Tag-display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

8. BMC Firmware Revision

Option Values: <BMC FW Rev display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

9. ME Firmware Revision

Option Values: <ME FW Rev display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

10. SDR Revision

Option Values: <SDR Rev display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

11. UUID

Option Values: <UUID display>

Help Text: <None>

Comments: Information only.**[Back to \[System Information Screen\]](#) — [\[Server Management Screen\]](#)**

12.2.4.3 BMC LAN Configuration

To access this screen from the **Main** screen, select **Server Management > System Information**. To move to another screen, press the <Esc> key to return to the **Server Management** screen, then select the desired screen.

The BMC configuration screen allows the user to configure the BMC Baseboard LAN channel and an Intel® RMM4 LAN channel, and to manage BMC User settings for up to five BMC Users.

An Intel® RMM4 Management Module may be installed in the server system.

If the Management Module is installed, it may also have a Dedicated Server Management NIC Module (DMN) installed with it. In that case, the LAN settings for the Intel®RMM4 with Dedicated Server Management NIC may be configured.

When there is no Management Module installed in the system, or there is an Intel® RMM4-Lite without a DMN installed, the LAN settings specific to the Intel® RMM4 are grayed out and not available.

This screen has a choice of IPv4 or IPv6 addressing. When IPv6 is disabled, only the IPv4 addressing options appear. When IPv6 is enabled, the IPv4 options are grayed out and unavailable, and there is an additional section active for IPv6-addressing. This is true for both the Baseboard LAN configuration and the Intel® RMM4 with Dedicated Server Management NIC Module.

IP addresses for either IPv4 or IPv6 addressing can be assigned by static IP addresses manually typed in, or by dynamic IP addresses supplied by a Dynamic Host Configuration Protocol (DHCP) server. IPv6 addressing can also be provided by “stateless autoconfiguration” which does not require a DHCP server.

The BMC LAN Configuration screen is unusual in that the LAN Configuration parameters are maintained by the BMC itself, so this screen is just a User Interface to the BMC configuration. As such, the initial values of the LAN options shown on the screen are acquired from the BMC when this screen is initially accessed by a user,. Any values changed by the user are communicated back to the BMC when a “Save Changes” or “Save Changes and Exit” action is performed. If a “Discard Changes” or “Discard Changes and Exit” action is performed instead, any accumulated changes from this screen will be disregarded and lost.



Figure 57. BMC LAN Configuration Screen

Screen Field Descriptions:

IP Source

Option Values: Static
 Dynamic

Help Text:

Select BMC IP Source: If [Static], IP parameters may be edited. If [Dynamic], these fields are display-only and IP address is acquired automatically (DHCP).

Comments: This specifies the IP Source for IPv4 addressing for the Baseboard LAN. There is a separate IP Source field for the Intel® RMM4 LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC, and its setting determines whether the other Baseboard LAN IPv4 addressing fields are display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

1. IP Address

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit IP Address. Press <Enter> to edit.

Comments: This specifies the IPv4 Address for the Baseboard LAN. There is a separate IPv4 Address field for the Intel® RMM4 LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

2. Subnet Mask

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit Subnet Mask. Press <Enter> to edit.

Comments: This specifies the IPv4 addressing Subnet Mask for the Baseboard LAN. There is a separate IPv4 Subnet Mask field for the Intel® RMM4 LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

3. Gateway IP

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit Gateway IP. Press <Enter> to edit.

Comments: This specifies the IPv4 addressing Gateway IP for the Baseboard LAN. There is a separate IPv4 Gateway IP field for the Intel® RMM4 LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

4. IPv6

Option Values: Enabled
 Disabled

Help Text:

Option to Enable/Disable IPv6 addressing and any IPv6 network traffic on these channels.

Comments: The initial value for this field is acquired from the BMC. It may be changed in order to switch between IPv4 and IPv6 addressing technologies.

When this option is set to **Disabled**, all other IPv6 fields will not be visible for the Baseboard LAN and Intel® RMM4 DMN (if installed). When IPv6 addressing is **Enabled**, all IPv6 fields for the Baseboard LAN and Intel® RMM4 DMN will become visible, and all IPv4 fields will be grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

5. IPv6 Source

Option Values: Static
 Dynamic
 Auto

Help Text:

Select BMC IPv6 source: If [Static], IPv6 parameters may be edited. If [Dynamic], these fields are display-only and IPv6 address is acquired automatically (DHCP). If [Auto], these fields are display-only and IPv6 address is acquired using ICMPv6 router/neighbor discovery.

Comments: This specifies the IP Source for IPv6 addressing for the Baseboard LAN configuration. There is a separate IPv6 Source field for the Intel® RMM4 LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is **Enabled**, the initial value for this field is acquired from the BMC, and its setting determines whether the other Baseboard LAN IPv6 addressing fields are display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

6. IPv6 Address

Option Values: *[Entry Field 0000.0000.0000.0000.0000.0000.0000.0000,
 0000.0000.0000.0000.0000.0000.0000.0000 is default]*

Help Text:

View/Edit IPv6 address. Press <Enter> to edit. IPv6 addresses consist of 8 hexadecimal 4 digit numbers separated by colons.

Comments: This specifies the IPv6 Address for the Baseboard LAN. There is a separate IPv6 Address field for the Intel® RMM4 LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

7. Gateway IPv6

Option Values: *[Entry Field 0000.0000.0000.0000.0000.0000.0000.0000, 0000.0000.0000.0000.0000.0000.0000.0000 is default]*

Help Text:

View/Edit Gateway IPv6 address. Press <Enter> to edit. Gateway IPv6 addresses consist of 8 hexadecimal 4 digit numbers separated by colons.

Comments: This specifies the Gateway IPv6 Address for the Baseboard LAN. There is a separate Gateway IPv6 Address field for the Intel® RMM4 LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

8. IPv6 Prefix Length

Option Values: *[Entry Field 0 – 128, 64 is default]*

Help Text:

View/Edit IPv6 Prefix Length from zero to 128 (default 64). Press <Enter> to edit.

Comments: This specifies the IPv6 Prefix Length for the Baseboard LAN. There is a separate IPv6 Prefix Length field for the Intel® RMM4 LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

9. Intel® RMM4

Option Values: *Not Present
Intel® RMM4-Lite
Intel® RMM4 + DMN*

Help Text: *<None>*

Comments: *Information only.* Displays whether an Intel® RMM4 component is currently installed. This information may come from querying the BMC.

Intel® RMM4-Lite is the Management Module without the Dedicated Server Management NIC Module. When this is present, or if the Management Module is **Not Present** at all, the fields for Intel® RMM4 LAN Configuration will not be visible.

When an **Intel® RMM4 + DMN** is installed, the options for **Intel® RMM4 LAN Configuration** will be visible. When **IPv6** is **Disabled**, the IPv4 configuration fields will be visible and the IPv6 configuration fields will not be visible. When **IPv6** is **Enabled**, the IPv4 fields will be grayed out and inactive, while the IPv6 Configuration fields will be visible.

In either case, the Intel® RMM4 section **IP Source** or **IPv6 Source** will determine whether the IPv4 or IPv6 address fields are display-only or can be edited.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

10. IP Source

Option Values: Static
 Dynamic

Help Text:

Select RMM4 IP source: If [Static], IP parameters may be edited. If [Dynamic], these fields are display-only and IP address is acquired automatically (DHCP).

Comments: This specifies the IP Source for IPv4 addressing for the Intel® RMM4 DMN LAN connection. There is a separate IP Source field for the Baseboard LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC, and its setting determines whether the other Intel® RMM4 DMN LAN IPv4 addressing fields are display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

11. IP Address

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit IP Address. Press <Enter> to edit.

Comments: This specifies the IPv4 Address for the Intel® RMM4 DMN LAN. There is a separate IPv4 Address field for the Baseboard LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

12. Subnet Mask

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit Subnet Mask. Press <Enter> to edit.

Comments: This specifies the IPv4 addressing Subnet Mask for the Intel® RMM4 DMN LAN. There is a separate IPv4 Subnet Mask field for the Baseboard LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

13. Gateway IP

Option Values: *[Entry Field 0.0.0.0, **0.0.0.0** is default]*

Help Text:

View/Edit Gateway IP. Press <Enter> to edit.

Comments: This specifies the IPv4 addressing Gateway IP for the Intel® RMM4 DMN LAN. There is a separate IPv4 Gateway IP field for the Baseboard LAN configuration.

When IPv4 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IP Source** determines whether this field is display-only (when **Dynamic**) or can be edited (when **Static**).

When IPv6 addressing is enabled, this field is grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

14. IPv6 Source

Option Values: Static
 Dynamic
 Auto

Help Text:

Select Intel® RMM4 IPv6 source: If [Static], IPv6 parameters may be edited. If [Dynamic], these fields are display-only and IPv6 address is acquired automatically (DHCP). If [Auto], these fields are display-only and IPv6 address is acquired using ICMPv6 router/neighbor discovery.

Comments: This specifies the IP Source for IPv6 addressing for the Intel® RMM4 DMN LAN configuration. There is a separate IPv6 Source field for the Baseboard LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is **Enabled**, the initial value for this field is acquired from the BMC, and its setting determines whether the other Intel® RMM4 DMN LAN IPv6 addressing fields are display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

15. IPv6 Address

Option Values: *[Entry Field 0000.0000.0000.0000.0000.0000.0000.0000,
 0000.0000.0000.0000.0000.0000.0000.0000 is default]*

Help Text:

View/Edit IPv6 address. Press <Enter> to edit. IPv6 addresses consist of 8 hexadecimal 4 digit numbers separated by colons.

Comments: This specifies the IPv6 Address for the Intel® RMM4 DMN LAN. There is a separate IPv6 Address field for the Baseboard LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

16. Gateway IPv6

Option Values: *[Entry Field 0000.0000.0000.0000.0000.0000.0000.0000,
 0000.0000.0000.0000.0000.0000.0000.0000 is default]*

Help Text:

View/Edit Gateway IPv6 address. Press <Enter> to edit. Gateway IPv6 addresses consist of 8 hexadecimal 4 digit numbers separated by colons.

Comments: This specifies the Gateway IPv6 Address for the Intel® RMM4 DMN LAN. There is a separate Gateway IPv6 Address field for the Baseboard LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

17. IPv6 Prefix Length

Option Values: *[Entry Field 0 – 128, **64** is default]*

Help Text:

View/Edit IPv6 Prefix Length from zero to 128 (default 64). Press <Enter> to edit.

Comments: This specifies the IPv6 Prefix Length for the Intel® RMM4 DMN LAN. There is a separate IPv6 Prefix Length field for the Baseboard LAN configuration.

This option is only visible when the **IPv6** option is set to **Enabled**.

When IPv6 addressing is used, the initial value for this field is acquired from the BMC. The setting of **IPv6 Source** determines whether this field is display-only (when **Dynamic** or **Auto**) or can be edited (when **Static**).

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

18. BMC DHCP Host Name

Option Values: *[Entry Field, 2-63 characters]*

Help Text:

View/Edit BMC DHCP host name. Press <Enter> to edit. Host name should start with an alphabetic, remaining can be alphanumeric characters. Host name length may be from 2 to 63 characters.

Comments: This field is active and may be edited whenever at least one of the **IP Source** or **IPv6 Source** options is set to **Dynamic**. This is the name of the DHCP Host from which dynamically assigned IPv4 or IPv6 addressing parameters are acquired.

The initial value for this field is supplied from the BMC, if there is a DHCP Host available. The user can edit the existing Ho or enter a different DHCP Host Name.

If none of the **IP/IPv6 Source** fields is set to **Dynamic**, then this **BMC DHCP Host Name** field will be grayed out and inactive.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

19. User ID

Option Values: **anonymous**

root
User3
User4
User5

Help Text:

Select the User ID to configure: User1 (anonymous), User2 (root), and User3/4/5 are supported.

Comments: These 5 User IDs are fixed choices and cannot be changed. The BMC supports 15 User IDs natively, but only the first 5 are supported through this interface.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

20. Privilege

Option Values: Callback
 User
 Operator
 Administrator

Help Text:

View/Select user privilege. User2 (root) privilege is "Administrator" and cannot be changed.

Comments: The level of privilege that is assigned for a User ID affects which functions that user may perform.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

21. User Status

Option Values: Enabled
 Disabled

Help Text:

Enable/Disable LAN access for selected user. Also enables/disables SOL, KVM, and media redirection.

Comments: Note that status setting is **Disabled** by default until set to **Enabled**.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

22. User Name

Option Values: *[Entry Field, 4 - 15 characters]*

Help Text:

Press <Enter> to edit User Name. User Name is a string of 4 to 15 alphanumeric characters, and must begin with an alphabetic character. User Name cannot be changed for User1 (anonymous) and User2 (root).

Comments: User Name can only be edited for users other than "anonymous" and "root". Those two User Names may not be changed.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

23. User Password

Option Values: *[Popup Entry Field, 0 - 15 characters]*

Help Text:

Press <Enter> key to enter password. Maximum length is 15 characters. Any ASCII printable characters can be used: case-sensitive alphabetic, numeric, and special characters.

***Note: Password entered will override any previously set password.*

Comments: This field will not indicate whether there is a password set already. There is no display - just press <Enter> for a popup with an entry field to enter a new password. Any new password entered will override the previous password, if there was one.

[Back to \[BMC LAN Configuration Screen\]](#) — [\[Server Management Screen\]](#)

12.2.5 Boot Options Screen (Tab)

The Boot Options screen displays all bootable media encountered during POST, and allows the user to configure the desired order in which boot devices are to be tried.

To access this screen from the **Main** screen or other top-level "Tab" screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Boot Options** screen is selected.

The first boot device in the specified Boot Order which is present and is bootable during POST will be used to boot the system, and will continue to be used to reboot the system until the boot device configuration has changed (that is, which boot devices are present), or until the system has been powered down and booted in a “cold” power-on boot.

Note: USB devices can be “hotplugged” during POST, and will be detected and “beeped”. They will be enumerated and displayed on the USB Configuration Setup screen. However, they may not be enumerated as bootable devices, depending on when in POST they were hotplugged. If they were recognized before the enumeration of bootable devices, they will appear as Boot Devices if appropriate. If they were recognized after Boot Device enumeration, they will not appear as a bootable device for the Boot Options screen, the Boot Manager screen, or the F6 Boot Menu.

There are two main types of boot order control, Legacy Boot and EFI Optimized boot. These are mutually exclusive – when EFI Optimized Boot is enabled, Legacy Boot (the default) is disabled. Within Legacy Boot operation, there are two further methods of ordering boot devices, Dynamic Boot Order and Static Boot Order.

The default for Boot Order control is Legacy Boot, with Dynamic Boot Order. If all types of bootable devices are installed in the system, then the default Boot Order is as follows :

- CD/DVD-ROM
- Floppy Disk Drive
- Hard Disk Drive
- PXE Network Device
- BEV (Boot Entry Vector) Device

EFI Shell and EFI Boot paths

In this default Boot Order, a USB device may appear in any of several Device Classes, due to the flexibility of USB connections and USB emulation of various types of devices.

Note: A USB Key (USB Flash Drive) can be formatted to emulate either a Floppy Drive or a Hard Drive and will appear in that Boot Device Class. However, although it can be formatted as a CDROM Drive, it will not be detected as such. It will be treated as a Hard Disk and will appear in the list of available Hard Drives.



Figure 58. Boot Options Screen

Screen Field Descriptions:

System Boot Timeout

Option Values: *[Entry Field 0 – 65535, 0 is default]*

Help Text:

The number of seconds BIOS will pause at the end of POST to allow the user to press the [F2] key for entering the BIOS Setup utility.

Valid values are 0-65535. Zero is the default. A value of 65535 causes the system to go to the Boot Manager menu and wait for user input for every system boot.

Comments: After entering the desired timeout, press the <Enter> key to register that timeout value to the system. These settings are in seconds. The timeout value entered will take effect on the next boot.

This timeout value is independent of the FRB2 setting for BIOS boot failure protection. The FBR2 countdown will be suspended during the time that the Boot Timeout countdown is active.

Also, if the <Pause> key is pressed during the time that the Boot Timeout is active, the Boot Timeout countdown will be suspended until the Pause state has been dismissed and normal POST processing has resumed.

[Back to \[Boot Options Screen\]](#)

System Early POST Timeout

Option Values: *[0 – 65535, **0** is default]*

Help Text:

The boot timeout takes effect before loading any third party option ROMs.

Comments: The number of seconds the BIOS will pause at the early stage of POST before the BIOS loads any third party option ROMs.

Valid values are 0-65535. Zero is the default. After entering the desired timeout, press the <Enter> key to register that timeout value to the system. These settings are in seconds. The timeout value entered will take effect on the next boot.

This timeout value is independent of the FRB2 setting for BIOS boot failure protection. The FBR2 countdown will be suspended during the time that the Boot Timeout countdown is active.

Also, neither the <Pause> key nor the <ESC> key works at this moment as Keyboard service is not ready.

Back to [Boot Options Screen]

1. Boot Option #1
2. Boot Option #2

Boot Option <#n>

Option Values: *<Available Boot Device #n>*

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: When the Boot order has been chosen, it will take effect on the next boot. The system will go down the list and boot from the first device on the list which is available and bootable.

This establishes the Boot Order only with respect to the normal boot path. This order has no effect on the Boot Manager selection list or the <F6> BIOS Boot Menu popup, both of which simply list all bootable devices available in the order in which they were detected. Whether or not a potential Boot Device is in this list has no bearing on the presence or order of Boot Devices shown for Boot Manager or the BIOS Boot Menu.

Back to [Boot Options Screen]

3. CDROM Order

Option Values: *<None>*

Help Text:

Set the order of the legacy devices in this group.

Comments: Selection only. Position to this line and press the <Enter> key to go to the ***CDROM Order Screen***.

This option appears when one or more bootable CDROM drives are available in the system. This includes USB CDROM devices, but not USB Keys formatted for CRDOM emulation, which are seen as Hard Disk drives.

Back to [Boot Options Screen]

4. Hard Disk Order

Option Values: <None>

Help Text:

Set the order of the legacy devices in this group.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Hard Disk Order Screen**.

This option appears when one or more bootable Hard Disk drives are available in the system. This includes USB Hard Disk devices and USB Keys formatted for Hard Disk or CRDOM emulation.

Back to [Boot Options Screen]

5. Floppy Order

Option Values: <None>

Help Text:

Set the order of the legacy devices in this group.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Floppy Order Screen**.

This option appears when one or more bootable Floppy Disk drives are available in the system. This includes USB Floppy Disk devices and USB Keys formatted for Floppy Disk emulation.

Back to [Boot Options Screen]

6. Network Device Order

Option Values: <None>

Help Text:

Set the order of the legacy devices in this group.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **Network Device Order Screen**.

This option appears when one or more bootable Network Devices are available in the system.

Back to [Boot Options Screen]

7. BEV Device Order

Option Values: <None>

Help Text:

Set the order of the legacy devices in this group.

Comments: Selection only. Position to this line and press the <Enter> key to go to the **BEV Device Order Screen**.

This option appears when one or more bootable BEV Devices are available in the system.

Back to [Boot Options Screen]

8. Add EFI Boot Option

Option Values: <None>

Help Text:

Add a new EFI boot option to the boot order.

Comments: Selection only. Position to this line and press the <Enter> key to go to the [Add EFI Boot Option Screen](#).

This option is only displayed if an EFI bootable device is available to the system.

[Back to \[Boot Options Screen\]](#)

9. Delete EFI Boot Option

Option Values: <None>

Help Text:

Remove an EFI boot option from the boot order.

Comments: Selection only. Position to this line and press the <Enter> key to go to the [Delete EFI Boot Option Screen](#).

This option is only displayed if an EFI boot path is included in the Boot Order.

[Back to \[Boot Options Screen\]](#)

10. EFI Optimized BootOption Values: Enabled
Disabled

Help Text:

If enabled, the BIOS only loads modules required for booting EFI-aware Operating Systems.

Comments: If this option is enabled, the system will not boot successfully to a non-EFI-aware OS.

[Back to \[Boot Options Screen\]](#)

11. Use Legacy Video for EFI OSOption Values: Enabled
Disabled

Help Text:

If enabled, the BIOS uses the legacy video ROM instead of the EFI video ROM.

Comments: This option appears only when EFI Optimized Boot is enabled.

[Back to \[Boot Options Screen\]](#)

12. Boot Option RetryOption Values: Enabled
Disabled

Help Text:

If enabled, this continually retries non-EFI-based boot options without waiting for user input.

Comments: This option is intended to keep retrying for cases where the boot devices could possibly be slow to initially respond, for example, if the device were "asleep" and did not wake quickly enough. However, if none of the devices in the Boot Order ever responds, the BIOS will continue to reboot indefinitely.

[Back to \[Boot Options Screen\]](#)

13. USB Boot Priority

Option Values: Enabled
 Disabled

Help Text:

*If enabled, newly discovered USB devices are moved to the top of their boot device category.
If disabled, newly discovered USB devices are moved to the bottom of their boot device category.*

Comments: This option enables or disables the “USB Reorder” functionality. USB Boot Priority, if enabled, is intended for the case where a user wants to be able to plug in a USB device and immediately boot to it, for example in case of a maintenance or System Administration operation. If a User Password is installed, USB Boot Priority action is suspended when a User Password is installed.

[Back to \[Boot Options Screen\]](#)

14. Static Boot Order

Option Values: Enabled
 Disabled

Help Text:

*[Disabled] - Devices removed from the system are deleted from Boot Order Tables.
[Enabled] - Devices removed have positions in Boot Order Tables retained for later reinsertion.*

Comments: When the option changes to “Enabled” from “Disabled”, it will enable Static Boot Order (SBO) from the next boot onward, and also the current Boot Order will be stored as the SBO template.

When the option changes from “Enabled” to “Disabled”, this will disable SBO and the SBO template will be cleared.

Otherwise it will retain the current Enabled/Disabled state.

[Back to \[Boot Options Screen\]](#)

15. Reset Static Boot Order

Option Values: Yes
 No Action

Help Text:

[Yes] Take snapshot of current boot order to save as Static Boot Order Template.

Comments: This option will allow you to save the Boot Order list as the Static Boot Order template without disabling and re-enabling the Static Boot Order option.

Select Yes to snapshot the current Boot Options list into the Static Boot Options list on the next boot. After saving Static Boot Options list, this option will change back to NoAction automatically.

This option is available only when the Static Boot Order option is Enabled. Otherwise it will grayed out and unavailable.

[Back to \[Boot Options Screen\]](#)

12.2.5.1 CDRom Order

The CDRom Order screen allows the user to control the order in which BIOS attempts to boot from the CDRom drives installed in the system. This screen is only available when there is at least one CDRom device available in the system configuration.

Note: A USB attached CDRom device will appear in this section. However, a USB Key formatted as a CRDOM device will not – it will be detected as a Hard Disk device and will be included in the Hard Disk Order Screen.

To access this screen from the **Main** screen, select **Boot Options > CDROM Order**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

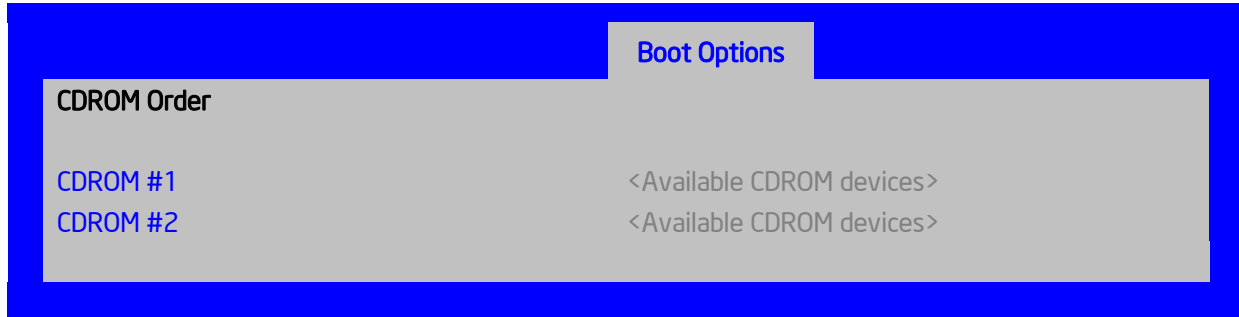


Figure 59. CDROM Order Screen

Screen Field Descriptions:

1. CDROM #1

CDROM #2

Option Values: <Available CDROM devices>

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: Choose the order of booting among CDROM devices by choosing which available CDROM device should be in each position in the order.

[Back to \[CDROM Order Screen\]](#) — [\[Boot Options Screen\]](#)

12.2.5.2 Hard Disk Order

The Hard Disk Order screen allows the user to control the order in which BIOS attempts to boot from the hard disk drives installed in the system. This screen is only available when there is at least one hard disk device available in the system configuration. Note that a USB attached Hard Disk drive or a USB Key device formatted as a hard disk will appear in this section.

To access this screen from the **Main** screen, select **Boot Options > Hard Disk Order**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

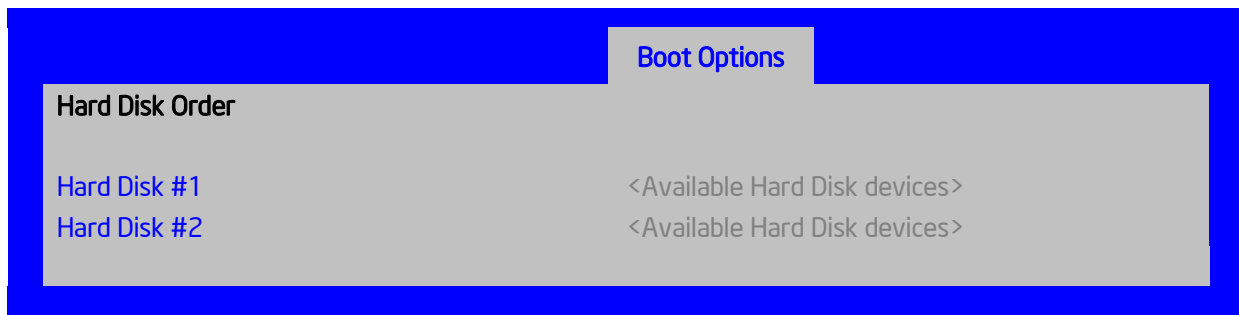


Figure 60. Hard Disk Order Screen

Screen Field Descriptions:

1. Hard Disk #1

Hard Disk #2

Option Values: <Available Hard Disk devices>

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: Choose the order of booting among Hard Disk devices by choosing which available Hard Disk device should be in each position in the order.

[Back to \[Hard Disk Order Screen\]](#) — [\[Boot Options Screen\]](#)

12.2.5.3 Floppy Order

The Floppy Order screen allows the user to control the order in which BIOS attempts to boot from the Floppy Disk drives installed in the system. This screen is only available when there is at least one Floppy Disk (diskette) device available in the system configuration. Note that a USB attached diskette drive or a USB Key device formatted as a diskette drive will appear in this section.

To access this screen from the **Main** screen, select **Boot Options > Floppy Order**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

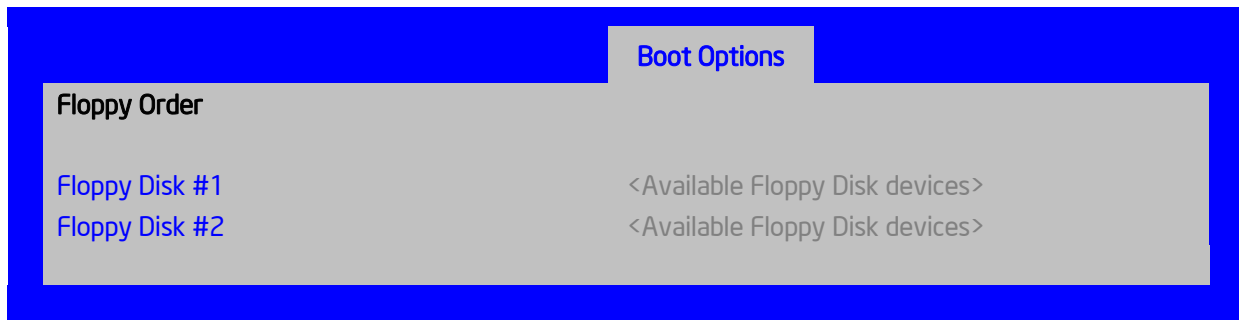


Figure 61. Floppy Order Screen

Screen Field Descriptions:

1. Floppy Disk #1
Floppy Disk #2

Option Values: <Available Floppy Disk devices>

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: Choose the order of booting among Floppy Disk devices by choosing which available Floppy Disk device should be in each position in the order.

[Back to \[Floppy Order Screen\]](#) — [\[Boot Options Screen\]](#)

12.2.5.4 Network Device Order

The Network Device Order screen allows the user to control the order in which BIOS attempts to boot from the network bootable devices installed in the system. This screen is only available when there is at least one network bootable device available in the system configuration.

To access this screen from the **Main** screen, select **Boot Options > Network Device Order**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

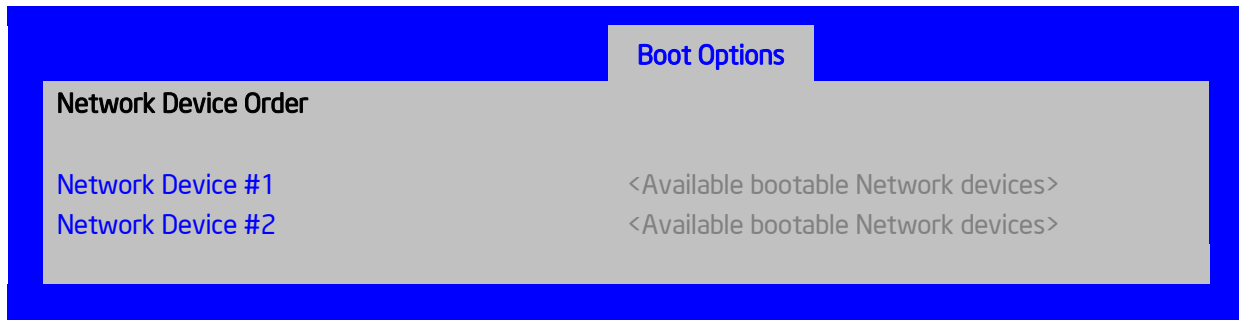


Figure 62. Network Device Order Screen

Screen Field Descriptions:

1. Network Device #1

Network Device #2

Option Values: <Available Network Devices>

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: Choose the order of booting among Network Devices by choosing which available Network Device should be in each position in the order.

[Back to \[Network Device Order Screen\]](#) — [\[Boot Options Screen\]](#)

12.2.5.5 BEV Device Order

The BEV Device Order screen allows the user to control the order in which BIOS attempts to boot from the BEV Devices installed in the system. This screen is only available when there is at least one BEV device available in the system configuration.

To access this screen from the **Main** screen, select **Boot Options > BEV Device Order**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

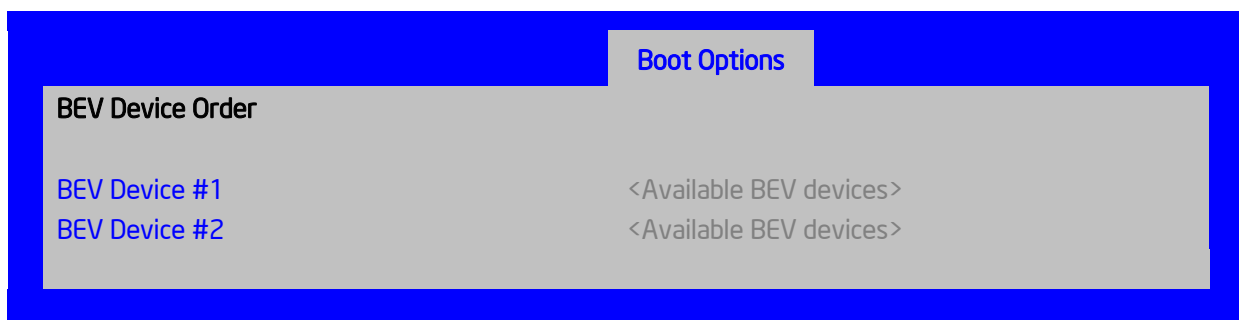


Figure 63. BEV Device Order Screen

Screen Field Descriptions:

1. BEV Device #1

BEV Device #2

Option Values: <Available BEV Devices>

Help Text:

Set system boot order by selecting the boot option for this position.

Comments: Choose the order of booting among BEV Devices by choosing which available BEV Device should be in each position in the order.

[Back to \[BEV Device Order Screen\]](#) — [\[Boot Options Screen\]](#)

12.2.5.6 Add EFI Boot Option

The Add EFI Boot Option screen allows the user to add an EFI boot option to the boot order. This screen is only available when there is at least one EFI bootable device present in the system configuration. The “Internal EFI Shell” Boot Option is permanent and cannot be added or deleted.

To access this screen from the **Main** screen, select **Boot Options > Add EFI Boot Option**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

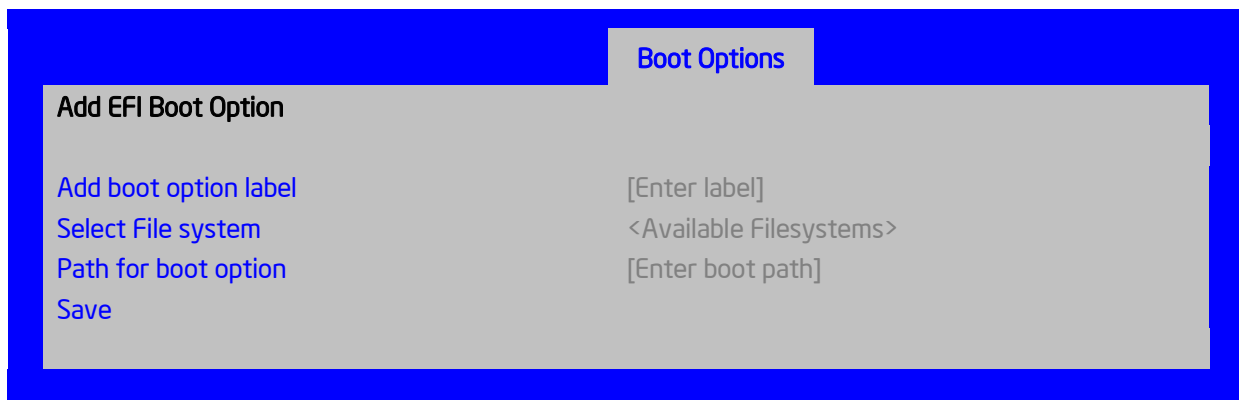


Figure 64. Add EFI Boot Option Screen

Screen Field Descriptions:

1. Add boot option label

Option Values: *[Enter label]*

Help Text:

Create the label for the new boot option.

Comments: This label becomes an abbreviation for this Boot Path.

[Back to \[Add EFI Boot Option Screen\]](#) — [\[Boot Options Screen\]](#)

2. Select File system

Option Values: <Available Filesystems>

Help Text:

Select one filesystem from this list.

Comments: Choose the filesystem on which this boot path resides.

[Back to \[Add EFI Boot Option Screen\]](#) — [\[Boot Options Screen\]](#)

3. Path for boot optionOption Values: *[Enter Boot Path]*

Help Text:

Enter the path to the boot option in the format \path\filename.efi.

Comments: This will be the Boot Path, residing on the filesystem chosen, which will entered into the Boot Order with the Label entered above.

Back to [Add EFI Boot Option Screen] — [Boot Options Screen]

4. SaveOption Values: *<None>*

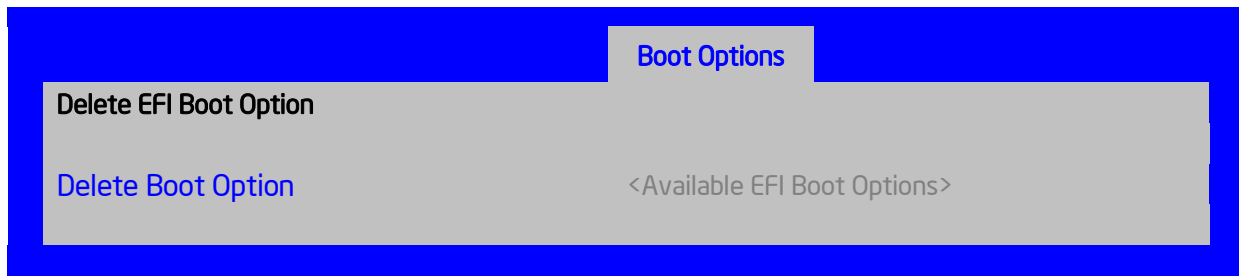
Help Text:

*Save the boot option.*Comments: *Selection only.* This will save the new Boot Option into the Boot Order.***Back to [Add EFI Boot Option Screen] — [Boot Options Screen]***

12.2.5.7 Delete EFI Boot Option

The Delete EFI Boot Option screen allows the user to remove an EFI boot option from the boot order. The “Internal EFI Shell” Boot Option will not be listed, since it is permanent and cannot be added or deleted.

To access this screen from the **Main** screen, select **Boot Options > Delete EFI Boot Option**. To move to another screen, press the <Esc> key to return to the **Boot Options** screen, then select the desired screen.

**Figure 65. Delete EFI Boot Option Screen****Screen Field Descriptions:****1. Delete Boot Option**Option Values: *<Available EFI Boot Options>*

Help Text:

Select one to delete.

Comments: This will not allow a user to delete the EFI Shell.

Back to [Delete EFI Boot Option Screen] — [Boot Options Screen]

12.2.6 Boot Manager Screen (Tab)

The Boot Manager screen allows the user to view a list of devices available for booting, and to select a boot device for immediately booting the system. There is no predetermined order for listing bootable devices. They are simply listed in order of discovery.

Regardless of whether any other bootable devices are available, the “Internal EFI Shell” will always be available,.

Note that this list is ***not*** in order according to the system Boot Option order. Reordering Boot Devices or even removing them from the Boot Order completely has no effect on the Boot Manager.

To access this screen from the ***Main*** screen or other top-level “Tab” screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the ***Boot Manager*** screen is selected.

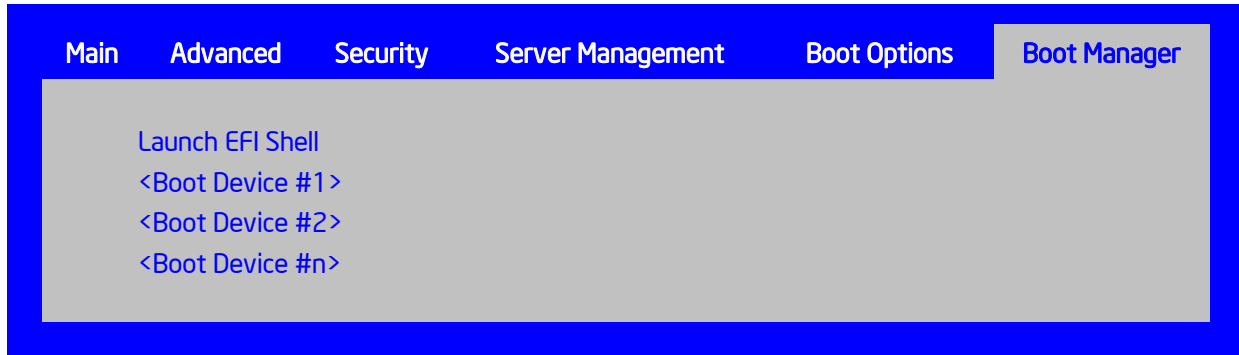


Figure 66. Boot Manager Screen

Screen Field Descriptions:

1. Launch EFI Shell

Option Values: <None>

Help Text:

Select this option to boot now.

Note: *This list is not the system boot option order. Use the Boot Options menu to view and configure the system boot option order.*

Comments: The EFI Shell will always be present in the list of bootable devices.

Back to [Boot Manager Screen]

2. <Boot Device #1>

3. <Boot Device #2>

4. <Boot Device #n>

Option Values: <None>

Help Text:

Select this option to boot now.

Note: *This list is not the system boot option order. Use the Boot Options menu to view and configure the system boot option order.*

Comments: These are names of bootable devices discovered in the system. The system user can choose any of them from which to initiate a one-time boot – that is, booting from any device in this list will not permanently affect the defined system Boot Order.


These bootable devices are not displayed in any specified order, particularly not in the system Boot Order established by the Boot Options screen. This is just a list of bootable devices in the order in which they were enumerated.

Back to [Boot Manager Screen]

12.2.7 Error Manager Screen (Tab)

The Error Manager screen displays any POST Error Codes encountered during BIOS POST, along with an explanation of the meaning of the Error Code in the form of a Help Text. This is an *Information Only* screen.

To access this screen from the **Main** screen or other top-level “Tab” screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Error Manager** screen is selected.



ERROR CODE	SEVERITY	INSTANCE	DESCRIPTION
<Post Code> 5224	<Major/Minor> Major	<Instance #> N/A	<Description> This is an example.

Figure 67. Error Manager Screen

Screen Field Descriptions:

1. ERROR CODE

Option Values: <POST Error Code>

Help Text: <N/A>

Comments: This is a POST Error Code – a BIOS-originated error that occurred during POST initialization.

[Back to \[Error Manager Screen\]](#)

2. SEVERITY

Option Values: Minor
Major
Fatal

Help Text: <N/A>

Comments: Each POST Error Code has a Severity associated with it. refer to the list of POST Error Codes to determine the Severity – Fatal, Major, Minor.

[Back to \[Error Manager Screen\]](#)

3. INSTANCE

Option Values: <Depends on error code>

Help Text: <N/A>

Comments: Where applicable, this field shows a value indicating which one of a group of components was responsible for generating the POST Error Code that is being reported.

[Back to \[Error Manager Screen\]](#)

4. DESCRIPTION

Option Values: <N/A>

Help Text: <Description of POST Error Code>

Comments: This is a description of the meaning of the POST Error Code that is being reported. This text actually appears in the screen space that is usually reserved for “Help” messages.

[Back to \[Error Manager Screen\]](#)

12.2.8 Save & Exit Screen (Tab)

The Save &Exit screen allows the user to choose whether to save or discard the configuration changes made on other Setup screens. It also allows the user to restore the BIOS settings to the factory defaults or to save or restore them to a set of user-defined default values. If Load Default Values is selected, the factory default settings (noted in bold in the Setup screen images) are applied. If Load User Default Values is selected, the system is restored to previously saved User Default Values.

To access this screen from the **Main** screen or other top-level “Tab” screen, press the right or left arrow keys to traverse the tabs at the top of the Setup screen until the **Exit** screen is selected.

Note that there is a Legal Disclaimer footnote at the bottom of the Save & Exit screen:

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

This is reference to any instance in the Setup screens where names belonging to other companies may appear. For example “LSI*” appears in Setup in the context of Mass Storage RAID options.

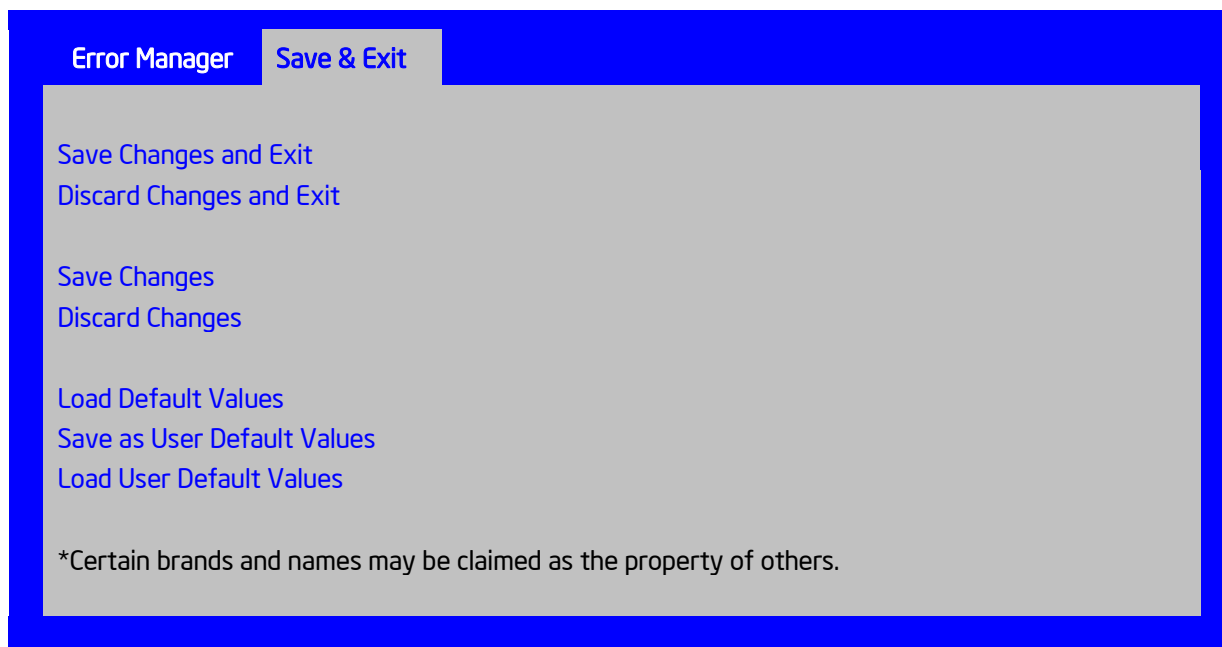


Figure 68. Save & Exit Screen

Screen Field Descriptions:

1. Save Changes and Exit

Option Values: <None>

Help Text:

*Exit BIOS Setup Utility after saving changes. The system will reboot if required.
The [F10] key can also be used.*

Comments: Selection only. Position to this line and press the <Enter> key to exit Setup with any changes in BIOS settings saved. If there have been no changes made in the settings, the BIOS will resume executing POST.

If changes have been made in BIOS settings, a confirmation pop-up will appear. If the “Save Changes & Exit” action is positively confirmed,, any persistent changes will applied and saved to the BIOS settings in NVRAM storage, then the system will reboot if necessary (which is normally the case). If the “Save Changes & Exit” action is not confirmed, BIOS will resume executing Setup.

The <F10 > function key may also be used from anyplace in Setup to initiate a “Save Changes & Exit” action.

Back to [Save & Exit Screen]

2. Discard Changes and Exit

Option Values: <None>

Help Text:

*Exit BIOS Setup Utility without saving changes.
The [Esc] key can also be used.*

Comments: Selection only. Position to this line and press the <Enter> key to exit Setup without saving any changes in BIOS settings. If there have been no changes made in the settings, the BIOS will resume executing POST.

If changes have been made in BIOS settings, a confirmation pop-up will appear. If the “Discard Changes & Exit” action is positively confirmed,, all pending changes will be discarded and BIOS will resume executing POST. If the “Discard Changes & Exit” action is not confirmed, BIOS will resume executing Setup without discarding any changes.

The <Esc > key may also be used in Setup to initiate a “Discard Changes & Exit” action.

Back to [Save & Exit Screen]

3. Save Changes

Option Values: <None>

Help Text:

Save Changes made so far to any of the setup options.

Comments: Selection only. Position to this line and press the <Enter> key to save any pending changes in BIOS settings. If there have been no changes made in the settings,

Also, the user should be aware that most changes require a reboot to become active. If changes have been made and saved, without exiting Setup, the system should be rebooted later even if no additional changes are made.

Back to [Save & Exit Screen]

4. Discard Changes

Option Values: <None>

Help Text:

Discard Changes made so far to any of the setup options.

Comments: Selection only. Position to this line and press the <Enter> key to discard any pending unsaved changes in BIOS settings. If there have been no changes made in the settings, the BIOS will resume executing POST.

If changes have been made in BIOS settings and not yet saved, a confirmation pop-up will appear. If the “Discard Changes” action is positively confirmed, all pending changes will be discarded and BIOS will resume executing POST. If the “Discard Changes” action is not confirmed, BIOS will resume executing Setup without discarding pending changes.

[Back to \[Save & Exit Screen\]](#)

5. Load Default Values

Option Values: <None>

Help Text:

Load Defaults Values for all the setup options.

Comments: Selection only. Position to this line and press the <Enter> key to load default values for all BIOS settings. These are the initial factory settings (“failsafe” settings) for all BIOS parameters.

There will be a confirmation popup to verify that the user really meant to take this action.

After initializing all BIOS settings to default values, the BIOS will resume executing Setup, so the user may made additional changes in the BIOS settings if necessary (for example, Boot Order) before doing a “Save Changes and Exit” with a reboot to make the default settings take effect, including any changes made after loading the defaults.

The <F9> function key may also be used from anyplace in Setup to initiate a “Load Default Values” action.

[Back to \[Save & Exit Screen\]](#)

6. Save as User Default Values

Option Values: <None>

Help Text:

Save the changes made so far as User Default Values.

Comments: Selection only. Position to this line and press the <Enter> key to save the current state of the settings for all BIOS parameters as a customized set of “User Default Values”.

These are a user-determined set of BIOS default settings that can be used as an alternative instead of the initial factory settings (“failsafe” settings) for all BIOS parameters.

By changing the BIOS settings to values that the user prefers to have for defaults, and then using this operation to save them as “User Default Values”, that version of BIOS settings can be restored at any time by using the following “Load User Default Values” operation.

There will be a confirmation popup to verify that the user really intended to take this action.

Loading the “factory default” values with F9 or the “Load Default Values” – or by any other means – does not affect the User Default Values. They remain set to whatever values they were saved as.

[Back to \[Save & Exit Screen\]](#)

7. Load User Default Values

Option Values: <None>

Help Text:

Load the User Default Values to all the setup options.

Comments: Selection only. Position to this line and press the <Enter> key to load User Default Values for all BIOS settings. These are user-customized BIOS default settings for all BIOS parameters, previously established by doing a “Save User Defaults” action (see above).

There will be a confirmation popup to verify that the user really intended to take this action.

Back to [Save & Exit Screen]

Appendix A: Integration and Usage Tips

- When adding or removing components or peripherals from the server board, power cords must be disconnected from the server. With power applied to the server, standby voltages are still present even though the server board is powered off.
- This server board supports the following processor product families:
 - Intel® Xeon® Processor E5-2600 product family with a Thermal Design Power (TDP) of up to and including 135 Watts.
 - Intel® Xeon® Processor E5-2600 v2 product family with a Thermal Design Power (TDP) of up to and including 130 Watts.
 - Previous generations of the Intel® Xeon® processors are not supported.
- Server systems using this server board may or may not meet the TDP design limits of the server board. Validate the TDP limits of the server system before selecting a processor.
- Processors must be installed in order. CPU 1 must be populated for the server board to operate.
- Riser card slot #2 is only functional with two CPUs installed.
- The SCU-1 mini-SAS connector is only enabled when an 8-port Intel® RAID C600 Upgrade Key is installed.
- For the best performance, the number of DDR3 DIMMs installed should be balanced across both processor sockets and memory channels.
- 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.
- The Intel® Remote Management Module 4 (Intel® RMM4) connector is **not** compatible with any previous versions of the Intel® Remote Management Module (Product Order Codes – AXXRMM, AXXRMM2, AXXRMM3).

Appendix B: Integrated BMC Sensor Tables

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

- **Sensor Type**

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

- **Event/Reading Type**

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

- **Event Offset/Triggers**

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

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

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

- **Assertion/De-assertion Enables**

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

- As: Assertions
- De: De-assertion

- **Readable Value/Offsets**

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

- **Event Data**

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

- R: Reading value
- T: Threshold value

- **Rearm Sensors**

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

- 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 can be 1 or 2 (positive or negative hysteresis).

- **Criticality**

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

- **Standby**

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

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

Table 61. BMC Core Sensors

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

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
					06 – Redundant: degraded from fully redundant state.	Degraded					
					07 – Redundant: Transition from non-redundant state.	Degraded					
IPMI Watchdog (<i>IPMI Watchdog</i>)	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 (<i>Physical Scrtcy</i>)	04h	Chassis Intrusion is chassis-specific	Physical Security 05h	Sensor Specific 6Fh	00 - Chassis intrusion	OK	As and De	–	Trig Offset	A	X
					04 - LAN leash lost						
FP Interrupt (<i>FP NMI Diag Int</i>)	05h	Chassis - specific	Critical Interrupt 13h	Sensor Specific 6Fh	00 - Front panel NMI/diagnostic interrupt	OK	As	–	Trig Offset	A	–
SMI Timeout (<i>SMI Timeout</i>)	06h	All	SMI Timeout F3h	Digital Discrete 03h	01 – State asserted	Fatal	As and De	–	Trig Offset	A	–
System Event Log (<i>System Event Log</i>)	07h	All	Event Logging Disabled 10h	Sensor Specific 6Fh	02 - Log area reset/cleared	OK	As	–	Trig Offset	A	X
System Event (<i>System Event</i>)	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	Stand-by
Button Sensor (<i>Button</i>)	09h	All	Button/Switch 14h	Sensor Specific 6Fh	00 – Power Button 02 – Reset Button	OK	AS	–	Trig Offset	A	X
BMC Watchdog	0Ah	All	Mgmt System Health 28h	Digital Discrete 03h	01 – State Asserted	Degraded	As	–	Trig Offset	A	-
Voltage Regulator Watchdog (VR Watchdog)	0Bh	All	Voltage 02h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X
Fan Redundancy ^{Note1} (<i>Fan Redundancy</i>)	0Ch	Chassis-specific	Fan 04h	Generic 0Bh	00 - Fully redundant	OK	As and De	–	Trig Offset	A	–
					01 - Redundancy lost	Degraded					
					02 - Redundancy degraded	Degraded					
					03 - Non-redundant: Sufficient resources. Transition from redundant	Degraded					
					04 - Non-redundant: Sufficient resources. Transition from insufficient.	Degraded					
					05 - Non-redundant: insufficient resources.	Non-Fatal					
					06 – Non-Redundant: degraded from fully redundant.	Degraded					
					07 - Redundant degraded from non-redundant	Degraded					
SSB Thermal Trip (<i>SSB Therm Trip</i>)	0Dh	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/Deassert	Readable Value/Offsets	Event Data	Rearm	Stand-by
IO Module Presence (IO Mod Presence)	0Eh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	X
SAS Module Presence (SAS Mod Presence)	0Fh	Platform-specific	Module/Board 15h	Digital Discrete 08h	01 – Inserted/Present	OK	As and De	–	Trig Offset	M	X
BMC Firmware Health (BMC FW Health)	10h	All	Mgmt Health 28h	Sensor Specific 6Fh	04 – Sensor Failure	Degraded	As	-	Trig Offset	A	X
System Airflow (System Airflow)	11h	All	Other Units 0Bh	Threshold 01h	–	–	–	Analog	–	–	–
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

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
IO Module Temperature (I/O Mod Temp)	26h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 1 Temperature (PCI Riser 1 Temp)	27h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
IO Riser Temperature (IO Riser Temp)	28h	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 1 Temperature (HSBP 1 Temp)	29h	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 2 Temperature (HSBP 2 Temp)	2Ah	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hot-swap Backplane 3 Temperature (HSBP 3 Temp)	2Bh	Chassis-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
PCI Riser 2 Temperature (PCI Riser 2 Temp)	2Ch	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
SAS Module Temperature (SAS Mod Temp)	2Dh	Platform-specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Exit Air Temperature (Exit Air Temp)	2Eh	Chassis & Platform Specific	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Network Interface Controller Temperature (LAN NIC Temp)	2Fh	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
Fan Tachometer Sensors (Chassis specific sensor names)	30h–3Fh	Chassis & Platform Specific	Fan 04h	Threshold 01h	[l] [c,nc]	nc = Degraded c = Non-fatal ^{Note2}	As and De	Analog	R, T	M	-
Fan Present Sensors (Fan x Present)	40h–4Fh	Chassis & 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
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

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
Power Supply 1 Temperature (PS1 Temperature)	5Ch	Chassis-specific	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Power Supply 2 Temperature (PS2 Temperature)	5Dh	Chassis-specific	Temperature	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	X
Hard Disk Drive 16 - 24 Status (HDD 16 - 24 Status)	60h - 68h	Chassis-specific	Drive Slot 0Dh	Sensor Specific 6Fh	00 - Drive Presence	OK	As and De	-	Trig Offset	A	X
					01 - Drive Fault	Degraded					
					07 - Rebuild/Remap in progress	Degraded					
Processor 1 Status (P1 Status)	70h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	-	Trig Offset	M	X
					07 - Presence	OK					
Processor 2 Status (P2 Status)	71h	All	Processor 07h	Sensor Specific 6Fh	01 - Thermal trip	Fatal	As and De	-	Trig Offset	M	X
					07 - Presence	OK					
Processor 1 Thermal Margin (P1 Therm Margin)	74h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Processor 2 Thermal Margin (P2 Therm Margin)	75h	All	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Processor 1 Thermal Control % (P1 Therm Ctrl %)	78h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	-
Processor 2 Thermal Control % (P2 Therm Ctrl %)	79h	All	Temperature 01h	Threshold 01h	[u] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	Trig Offset	A	-

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
Processor 1 ERR2 Timeout (P1 ERR2)	7Ch	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	A	–
Processor 2 ERR2 Timeout (P2 ERR2)	7Dh	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	A	–
Catastrophic Error (CATERR)	80h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor0 MSID Mismatch (P0 MSID Mismatch)	81h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor Population Fault (CPU Missing)	82h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	–	Trig Offset	M	–
Processor1 MSID Mismatch (P1 MSID Mismatch)	87h	All	Processor 07h	Digital Discrete 03h	01 – State Asserted	fatal	As and De	–	Trig Offset	M	–
Processor 1 VRD Temperature (P1 VRD Hot)	90h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	M	–
Processor 2 VRD Temperature (P2 VRD Hot)	91h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	M	–
Processor 1 Memory VRD Hot 0-1 (P1 Mem01 VRD Hot)	94h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–
Processor 1 Memory VRD Hot 2-3 (P1 Mem23 VRD Hot)	95h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	–	Trig Offset	A	–

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
Processor 2 Memory VRD Hot 0-1 (P2 Mem01 VRD Hot)	96h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	-	Trig Offset	A	-
Processor 2 Memory VRD Hot 2-3 (P2 Mem23 VRD Hot)	97h	All	Temperature 01h	Digital Discrete 05h	01 - Limit exceeded	Non-fatal	As and De	-	Trig Offset	A	-
Power Supply 1 Fan Tachometer 1 (PS1 Fan Tach 1)	A0h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 1 Fan Tachometer 2 (PS1 Fan Tach 2)	A1h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 1 (PS2 Fan Tach 1)	A4h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Power Supply 2 Fan Tachometer 2 (PS2 Fan Tach 2)	A5h	Chassis-specific	Fan 04h	Generic – digital discrete	01 – State Asserted	Non-fatal	As and De	-	Trig Offset	M	-
Processor 1 DIMM Aggregate Thermal Margin (P1 DIMM Thrm Mrgn)	B0h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	-
Processor 2 DIMM Aggregate Thermal Margin (P2 DIMM Thrm Mrgn)	B1h	All	Temperature 01h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	-
Processor 1 DIMM Thermal Trip (P1 Mem Thrm Trip)	C0h	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	-	Trig Offset	M	X
Processor 2 DIMM Thermal Trip (P2 Mem Thrm Trip)	C1h	All	Temperature 01h	Digital Discrete 03h	01 – State Asserted	Fatal	As and De	-	Trig Offset	M	X

Full Sensor Name (Sensor name in SDR)	Sensor #	Platform Applicability	Sensor Type	Event/Reading Type	Event Offset Triggers	Contrib. To System Status	Assert/De-assert	Readable Value/Offsets	Event Data	Rearm	Stand-by
Global Aggregate Temperature Margin 1 (Agg Therm Mrgn 1)	C8h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 2 (Agg Therm Mrgn 2)	C9h	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 3 (Agg Therm Mrgn 3)	CAh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 4 (Agg Therm Mrgn 4)	CBh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 5 (Agg Therm Mrgn 5)	CCh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 6 (Agg Therm Mrgn 6)	CDh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 7 (Agg Therm Mrgn 7)	CEh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Global Aggregate Temperature Margin 8 (Agg Therm Mrgn 8)	CFh	Platform Specific	Temperature 01h	Threshold 01h	-	-	-	Analog	R, T	A	-
Baseboard +12V (BB +12.0V)	D0h	All	Voltage 02h	Threshold 01h	[u,l] [c,nc]	nc = Degraded c = Non-fatal	As and De	Analog	R, T	A	-
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

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

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

Appendix C: Management Engine Generated SEL Event Messages

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

Table 62. Server Platform Services Firmware Health Event

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

	<p>Byte 7 – Event Data 3 =<Extended error code. Should be used when reporting an error to the support></p>
--	---

Table 63. Node Manager Health Event

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

Appendix D: 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.




Figure 69. POST Diagnostic LED Location

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 64. 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 65. 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 63									
PEI Phase continued...									
31h	0	0	1	1	0	0	0	1	Memory Installed
32h	0	0	1	1	0	0	1	0	CPU PEIM (Cpu Init)
33h	0	0	1	1	0	0	1	1	CPU PEIM (Cache Init)
34h	0	0	1	1	0	1	0	0	CPU PEIM (BSP Select)
35h	0	0	1	1	0	1	0	1	CPU PEIM (AP Init)
36h	0	0	1	1	0	1	1	0	CPU PEIM (CPU SMM Init)
4Fh	0	1	0	0	1	1	1	1	Dxe IPL started
DXE Phase									
60h	0	1	1	0	0	0	0	0	DXE Core started
61h	0	1	1	0	0	0	0	1	DXE NVRAM Init
62h	0	1	1	0	0	0	1	0	SB RUN Init
63h	0	1	1	0	0	0	1	1	Dxe CPU Init
68h	0	1	1	0	1	0	0	0	DXE PCI Host Bridge Init
69h	0	1	1	0	1	0	0	1	DXE NB Init
6Ah	0	1	1	0	1	0	1	0	DXE NB SMM Init
70h	0	1	1	1	0	0	0	0	DXE SB Init
71h	0	1	1	1	0	0	0	1	DXE SB SMM Init
72h	0	1	1	1	0	0	1	0	DXE SB devices Init
78h	0	1	1	1	1	0	0	0	DXE ACPI Init
79h	0	1	1	1	1	0	0	1	DXE CSM Init
90h	1	0	0	1	0	0	0	0	DXE BDS Started
91h	1	0	0	1	0	0	0	1	DXE BDS connect drivers
92h	1	0	0	1	0	0	1	0	DXE PCI Bus begin
93h	1	0	0	1	0	0	1	1	DXE PCI Bus HPC Init
94h	1	0	0	1	0	1	0	0	DXE PCI Bus enumeration
95h	1	0	0	1	0	1	0	1	DXE PCI Bus resource requested
96h	1	0	0	1	0	1	1	0	DXE PCI Bus assign resource

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	
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
A Eh	1	0	1	0	1	1	1	0	DXE Legacy Boot
A Fh	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	1	0	0	0	0	0	S3 Resume PEIM (S3 started)
E1h	1	1	1	0	0	0	0	1	S3 Resume PEIM (S3 boot script)
E2h	1	1	1	0	0	0	1	0	S3 Resume PEIM (S3 Video Repost)
E3h	1	1	1	0	0	0	1	1	S3 Resume PEIM (S3 OS wake)
BIOS Recovery									
F0h	1	1	1	1	0	0	0	0	PEIM which detected forced Recovery condition
F1h	1	1	1	1	0	0	0	1	PEIM which detected User Recovery condition
F2h	1	1	1	1	0	0	1	0	Recovery PEIM (Recovery started)
F3h	1	1	1	1	0	0	1	1	Recovery PEIM (Capsule found)
F4h	1	1	1	1	0	1	0	0	Recovery PEIM (Capsule loaded)

POST Memory Initialization MRC Diagnostic Codes

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

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

Table 66. 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 67. MRC Fatal Error 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.

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	
									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 E: POST Code Errors

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

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

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

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

Major: The error message is displayed on the Error Manager screen, and an error is logged to the SEL. The POST Error **P**ause option setting in the BIOS setup determines whether the system pauses to the Error Manager for this type of error so the user can take immediate corrective action or the system continues booting.

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

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

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

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

Table 68. 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

Error Code	Error Message	Response
0146	PCI out of resources error	Major
0191	Processor core/thread count mismatch detected	Fatal
0192	Processor cache size mismatch detected	Fatal
0194	Processor family mismatch detected	Fatal
0195	Processor Intel® QPI link frequencies unable to synchronize	Fatal
0196	Processor model mismatch detected	Fatal
0197	Processor frequencies unable to synchronize	Fatal
5220	BIOS Settings reset to default settings	Major
5221	Passwords cleared by jumper	Major
5224	Password clear jumper is Set	Major
8130	Processor 01 disabled	Major
8131	Processor 02 disabled	Major
8132	Processor 03 disabled	Major
8133	Processor 04 disabled	Major
8160	Processor 01 unable to apply microcode update	Major
8161	Processor 02 unable to apply microcode update	Major
8162	Processor 03 unable to apply microcode update	Major
8163	Processor 04 unable to apply microcode update	Major
8170	Processor 01 failed Self Test (BIST)	Major
8171	Processor 02 failed Self Test (BIST)	Major
8172	Processor 03 failed Self Test (BIST)	Major
8173	Processor 04 failed Self Test (BIST)	Major
8180	Processor 01 microcode update not found	Minor
8181	Processor 02 microcode update not found	Minor
8182	Processor 03 microcode update not found	Minor
8183	Processor 04 microcode update not found	Minor
8190	Watchdog timer failed on last boot	Major
8198	OS boot watchdog timer failure	Major
8300	Baseboard management controller failed self-test	Major
8305	Hot Swap Controller failure	Major
83A0	Management Engine (ME) failed Selftest	Major
83A1	Management Engine (ME) Failed to respond.	Major
84F2	Baseboard management controller failed to respond	Major
84F3	Baseboard management controller in update mode	Major
84F4	Sensor data record empty	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8501	DIMM Population Error	Major
8520	DIMM_A1 failed test/initialization	Major
8521	DIMM_A2 failed test/initialization	Major
8522	DIMM_A3 failed test/initialization	Major
8523	DIMM_B1 failed test/initialization	Major
8524	DIMM_B2 failed test/initialization	Major
8525	DIMM_B3 failed test/initialization	Major
8526	DIMM_C1 failed test/initialization	Major
8527	DIMM_C2 failed test/initialization	Major
8528	DIMM_C3 failed test/initialization	Major
8529	DIMM_D1 failed test/initialization	Major
852A	DIMM_D2 failed test/initialization	Major
852B	DIMM_D3 failed test/initialization	Major
852C	DIMM_E1 failed test/initialization	Major
852D	DIMM_E2 failed test/initialization	Major
852E	DIMM_E3 failed test/initialization	Major
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

Error Code	Error Message	Response
8534	DIMM_G3 failed test/initialization	Major
8535	DIMM_H1 failed test/initialization	Major
8536	DIMM_H2 failed test/initialization	Major
8537	DIMM_H3 failed test/initialization	Major
8538	DIMM_I1 failed test/initialization	Major
8539	DIMM_I2 failed test/initialization	Major
853A	DIMM_I3 failed test/initialization	Major
853B	DIMM_J1 failed test/initialization	Major
853C	DIMM_J2 failed test/initialization	Major
853D	DIMM_J3 failed test/initialization	Major
853E	DIMM_K1 failed test/initialization	Major
853F (Go to 85C0)	DIMM_K2 failed test/initialization	Major
8540	DIMM_A1 disabled	Major
8541	DIMM_A2 disabled	Major
8542	DIMM_A3 disabled	Major
8543	DIMM_B1 disabled	Major
8544	DIMM_B2 disabled	Major
8545	DIMM_B3 disabled	Major
8546	DIMM_C1 disabled	Major
8547	DIMM_C2 disabled	Major
8548	DIMM_C3 disabled	Major
8549	DIMM_D1 disabled	Major
854A	DIMM_D2 disabled	Major
854B	DIMM_D3 disabled	Major
854C	DIMM_E1 disabled	Major
854D	DIMM_E2 disabled	Major
854E	DIMM_E3 disabled	Major
854F	DIMM_F1 disabled	Major
8550	DIMM_F2 disabled	Major
8551	DIMM_F3 disabled	Major
8552	DIMM_G1 disabled	Major
8553	DIMM_G2 disabled	Major
8554	DIMM_G3 disabled	Major
8555	DIMM_H1 disabled	Major
8556	DIMM_H2 disabled	Major
8557	DIMM_H3 disabled	Major
8558	DIMM_I1 disabled	Major
8559	DIMM_I2 disabled	Major
855A	DIMM_I3 disabled	Major
855B	DIMM_J1 disabled	Major
855C	DIMM_J2 disabled	Major
855D	DIMM_J3 disabled	Major
855E	DIMM_K1 disabled	Major
855F (Go to 85D0)	DIMM_K2 disabled	Major
8560	DIMM_A1 encountered a Serial Presence Detection (SPD) failure	Major
8561	DIMM_A2 encountered a Serial Presence Detection (SPD) failure	Major
8562	DIMM_A3 encountered a Serial Presence Detection (SPD) failure	Major
8563	DIMM_B1 encountered a Serial Presence Detection (SPD) failure	Major
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

Error Code	Error Message	Response
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	DIMM_K2 encountered a Serial Presence Detection (SPD) failure	Major
(Go to 85E0)		
85C0	DIMM_K3 failed test/initialization	Major
85C1	DIMM_L1 failed test/initialization	Major
85C2	DIMM_L2 failed test/initialization	Major
85C3	DIMM_L3 failed test/initialization	Major
85C4	DIMM_M1 failed test/initialization	Major
85C5	DIMM_M2 failed test/initialization	Major
85C6	DIMM_M3 failed test/initialization	Major
85C7	DIMM_N1 failed test/initialization	Major
85C8	DIMM_N2 failed test/initialization	Major
85C9	DIMM_N3 failed test/initialization	Major
85CA	DIMM_O1 failed test/initialization	Major
85CB	DIMM_O2 failed test/initialization	Major
85CC	DIMM_O3 failed test/initialization	Major
85CD	DIMM_P1 failed test/initialization	Major
85CE	DIMM_P2 failed test/initialization	Major
85CF	DIMM_P3 failed test/initialization	Major
85D0	DIMM_K3 disabled	Major
85D1	DIMM_L1 disabled	Major
85D2	DIMM_L2 disabled	Major
85D3	DIMM_L3 disabled	Major
85D4	DIMM_M1 disabled	Major
85D5	DIMM_M2 disabled	Major
85D6	DIMM_M3 disabled	Major
85D7	DIMM_N1 disabled	Major
85D8	DIMM_N2 disabled	Major
85D9	DIMM_N3 disabled	Major
85DA	DIMM_O1 disabled	Major
85DB	DIMM_O2 disabled	Major
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

Error Code	Error Message	Response
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 69. 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 Table 67 and Table 68.	System halted because a fatal error related to the memory was detected.
2	BIOS Recovery started	NA	Recovery boot has been initiated.
4	BIOS Recovery failure	NA	BIOS recovery has failed. This typically happens so quickly after recovery was 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 70. 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 F: Supported Intel® Server Systems

Two Intel® Server System product families integrate the Intel® Server board S2600GZ and Intel® Server S2600GL, they are the 1U rack mount Intel® Server System R1000GZ/GL product family and the 2U rack mount Intel® R2000GZ/GL product family.

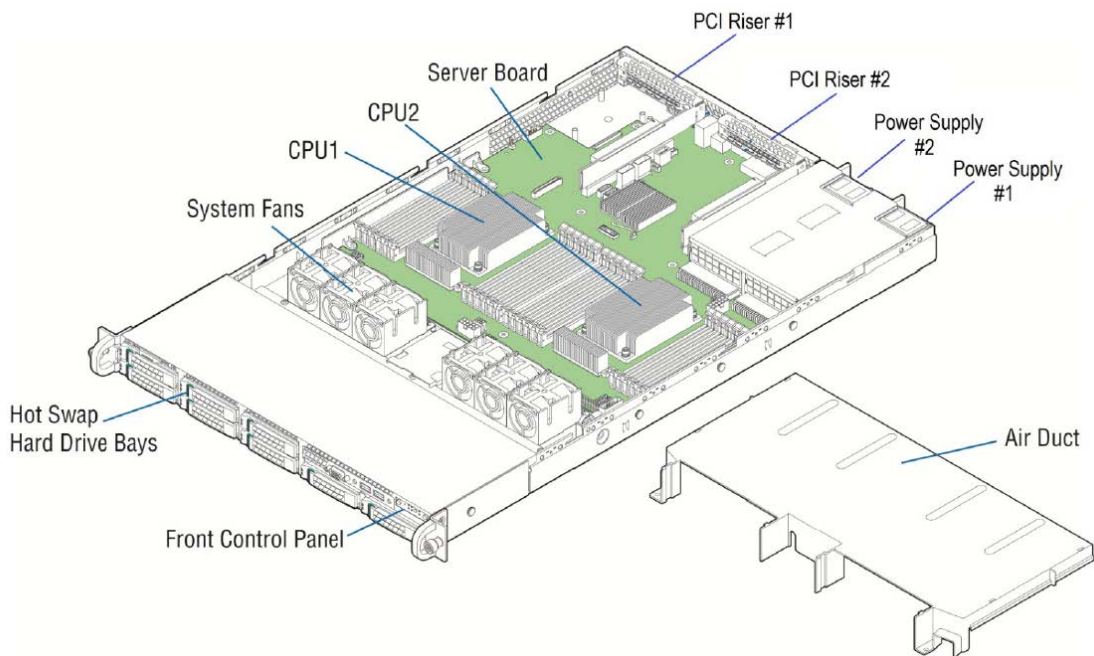


Figure 70. Intel® Server System R1000GZ/GL

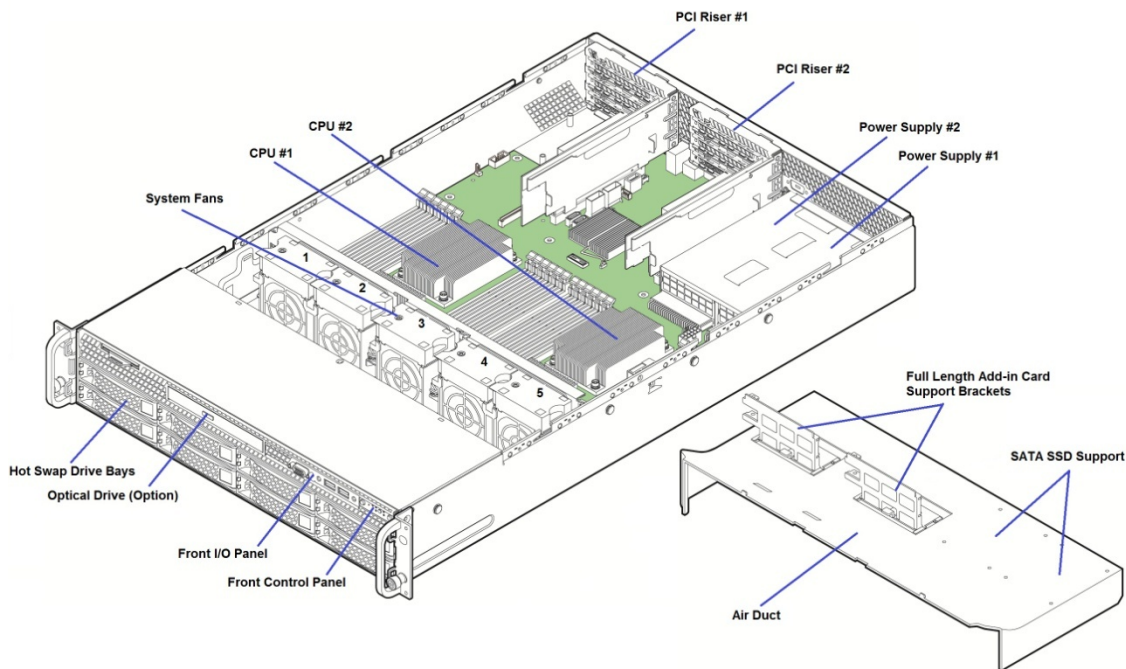


Figure 71. Intel® Server System R2000GZ/GL

Table 71. Intel® Server System R1000GZ/GL Product Family Feature Set

Feature	Description
Processor Support	<ul style="list-style-type: none"> ▪ Support for one or two processors: <ul style="list-style-type: none"> ○ Intel® Xeon® processor E5-2600 product family with TDP support up to 135 W^{1,2} ○ Intel® Xeon® processor E5-2600 v2 product family with TDP support up to 130 W
Memory	<ul style="list-style-type: none"> ▪ S2600GL - 16 DIMM slots – 2 DIMMs/Channel – 4 memory channels per processor ▪ S2600GZ - 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor ▪ Unbuffered DDR3 (UDIMM), registered DDR3 (RDIMM), Load Reduced DDR3 (LRDIMM) ▪ Memory DDR3 data transfer rates of 800, 1066, 1333 MT/s, 1600, 1866³ MT/s ▪ DDR3 standard I/O voltage of 1.5V and DDR3 Low Voltage of 1.35V
Chipset	Intel® C602 chipset with support for optional Storage Option Select keys
External I/O connections	<ul style="list-style-type: none"> ▪ Video (back and front video connectors) ▪ RJ-45 Serial- A Port ▪ Four RJ-45 Network Interface Connectors supporting 10/100/1000Mb ▪ USB 2.0 connectors - 3 on back panel + 2 on front panel
Internal I/O connectors/headers	<ul style="list-style-type: none"> ▪ One Type-A USB 2.0 connector ▪ One DH-10 Serial-B port connector
Optional I/O Module support	<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"> ▪ AXX4P1GBPWL IOM – Quad port 1 GbE based on Intel® Ethernet Controller I350 ▪ AXX10GBTWL IOM – Dual RJ-45 port 10GBase-T I/O Module based on Intel® Ethernet Controller x540 ▪ AXX10GBNIA IOM – Dual SFP+ port 10GbE module based on Intel® 82599 10 GbE controller ▪ AXX1FDRIB IOM – Single Port FDR 56GT/S speed InfiniBand module with QSFP connector ▪ AXX2FDRIB IOM – Dual port FDR 56GT/S speed infiniband module with QSFP connector ▪ AXXQAA IOMOD – Intel® Quick Assist Accelerator Card
System Fans	<ul style="list-style-type: none"> ▪ Six dual rotor managed system fans ▪ One power supply fan for each installed power supply module
Riser Cards	<p>Concurrent support for two PCIe riser cards. Each riser card slot has support for the following riser card options:</p> <ul style="list-style-type: none"> ▪ Single add-in card slot – PCIe x16, x16 mechanical
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"> ▪ One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices ▪ Two 7-pin single port AHCI SATA connectors capable of supporting up to 6 GB/sec ▪ Two SCU 4-port mini-SAS connectors capable of supporting up to 3 GB/sec SAS/SATA ▪ Intel® RAID C600 Upgrade Key support providing optional expanded SCU SATA/SAS RAID capabilities ▪ Intel® Integrated RAID module support (Optional)
Security	Intel® Trusted Platform Module (TPM) - AXXTPE5 (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

Feature	Description
Power Supply Options	<ul style="list-style-type: none"> ▪ The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 Redundant Power, and 2+0 Combined Power ▪ Three power supply options: <ul style="list-style-type: none"> ○ AC 460W Gold ○ AC 750W Platinum ○ DC 750W
Power Supply Options	<ul style="list-style-type: none"> ▪ The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 Redundant Power, and 2+0 Combined Power ▪ Three power supply options: <ul style="list-style-type: none"> ○ AC 460W Gold ○ AC 750W Platinum
Storage Bay Options	<ul style="list-style-type: none"> ▪ 4x – 3.5” SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support ▪ 8x – 2.5” SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support (capable)
Supported Rack Mount Kit Accessory Options	<ul style="list-style-type: none"> ▪ Tool-less rack mount rail kit – Intel Product Code – AXXPRAIL ▪ Value rack mount rail kit – Intel Product Code – AXXVRAIL ▪ Cable Management Arm – Intel Product Code – AXX1U2UCMA (*supported with AXXPRAIL only) ▪ 2-post fixed mount bracket kit – Intel Product Code – AXX2POSTBRCKT

Notes:

1. With a system fan failure, processor throttling may occur.
2. Processor throttling may occur with systems configured using the following processors: E5-2690, E5-2643.
3. Intel® Xeon® processor E5-2600 v2 product family only

Table 72. Intel® Server System R2000GZ/GL Product Family Feature Set

Feature	Description
Processor Support	<ul style="list-style-type: none"> ▪ Support for one or two processors: <ul style="list-style-type: none"> ○ Intel® Xeon® processor E5-2600 product family with TDP support up to 135 W^{1,2} ○ Intel® Xeon® processor E5-2600 v2 product family with TDP support up to 130 W
Memory	<ul style="list-style-type: none"> ▪ S2600GL - 16 DIMM slots – 2 DIMMs/Channel – 4 memory channels per processor ▪ S2600GZ - 24 DIMM slots – 3 DIMMs/Channel – 4 memory channels per processor ▪ Unbuffered DDR3 (UDIMM), Registered DDR3 (RDIMM), and Load Reduced DDR3 (LRDIMM) ▪ Memory DDR3 data transfer rates of 800, 1066, 1333 MT/s, 1600 and 1866³ MT/s ▪ DDR3 standard I/O voltage of 1.5V and DDR3 Low Voltage of 1.35V
Chipset	Intel® C602 chipset with support for optional Storage Option Select keys
External I/O connections	<ul style="list-style-type: none"> ▪ Video – Back Panel + Front Panel on Non-Max Hard drive SKUs ▪ RJ-45 Serial- A Port ▪ Four RJ-45 Network Interface Connectors supporting 10/100/1000Mb ▪ USB 2.0 connectors - 3 on back panel + 2 on front panel on non-max hard drive SKUs
Internal I/O connectors/headers	<ul style="list-style-type: none"> ▪ One Type-A USB 2.0 connector ▪ One DH-10 Serial-B port connector
Optional I/O Module support	<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"> ▪ AXX4P1GBPWL IOM – Quad port 1 GbE based on Intel® Ethernet Controller I350 ▪ AXX10GBTWLIOM – Dual RJ-45 port 10GBase-T I/O Module based on Intel® Ethernet Controller x540 ▪ AXX10GBNIA IOM – Dual SFP+ port 10GbE module based on Intel® 82599 10 GbE controller ▪ AXX1FDRIB IOM – Single Port FDR 56GT/S speed InfiniBand module with QSFP connector ▪ AXX2FDRIB IOM – Dual port FDR 56GT/S speed infiniband module with QSFP connector ▪ AXXQAAIOMOD - Intel® Quick Assist Accelerator Card
System Fans	<ul style="list-style-type: none"> ▪ Five managed system fans ▪ One power supply fan for each installed power supply module
Riser Cards	<p>Support for two riser card slots. Each riser card slot has support for the following riser card options:</p> <ul style="list-style-type: none"> ▪ 3-slot PCIe Riser Card: (Slots 1 & 2) – PCIe x16 slot, x8 lanes, (Slot 3) – PCIe x8 slot, x8 lanes ▪ 2-slot PCIe Riser Card: (Slot 1) – PCIe x16 slot, x16 lanes, (Slot 2) – PCIe x8 slot, x8 lanes ▪ 3-slot PCIx/PCIe Riser Card: (Slots 1 & 2) – PCIx 64-bit, (Slot 3) – PCIe x8 slot, x8 lanes
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"> ▪ One eUSB 2x5 pin connector to support 2mm low-profile eUSB solid state devices ▪ Two 7-pin single port AHCI SATA connectors capable of supporting up to 6 GB/sec ▪ Two SCU 4-port mini-SAS connectors capable of supporting up to 3 GB/sec SAS/SATA ▪ Intel® Integrated RAID module support (Optional) ▪ Intel® RAID C600 Upgrade Key support providing optional expanded SATA/SAS RAID capabilities
Security	Intel® Trusted Platform Module (TPM) - AXXTPE5 (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	<ul style="list-style-type: none"> ▪ The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 Redundant Power, and 2+0 Combined Power ▪ Three power supply options: <ul style="list-style-type: none"> ▪ AC 460W Gold ▪ AC 750W Platinum
Storage Bay Options	<ul style="list-style-type: none"> ▪ 8x – 3.5" SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support ▪ 12x – 3.5" SATA/SAS Hot Swap Hard Drive Bays ▪ 8x – 2.5" SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support ▪ 16x – 2.5" SATA/SAS Hot Swap Hard Drive Bays + Optical Drive support ▪ 24x – 2.5" SATA/SAS Hot Swap Hard Drive Bays
Supported Rack Mount Kit Accessory Options	<ul style="list-style-type: none"> ▪ Tool-less rack mount rail kit – Intel Product Code – AXXPRAIL ▪ Value rack mount rail kit – Intel Product Code – AXXVRAIL ▪ Cable Management Arm – Intel Product Code – AXX1U2UCMA (*supported with AXXPRAIL only) ▪ 2-post fixed mount bracket kit – Intel Product Code – AXX2POSTBRCKT

Notes:

1. With a system fan failure, processor throttling may occur.
2. Processor throttling may occur with systems configured using the following processors: E5-2690, E5-2643.
3. Intel® Xeon® processor E5-2600 v2 product family only

Refer to the Technical Product Specification for each these two Intel® Server Product families for more information.

Glossary

This appendix contains important terms used in this document. For ease of use, numeric entries are listed first (for example, “82460GX”) followed by alpha entries (for example, “AGP 4x”). Acronyms are followed by non-acronyms.

Term	Definition
ACPI	Advanced Configuration and Power Interface
AP	Application Processor
APIC	Advanced Programmable Interrupt Control
ARP	Address Resolution Protocol
ASIC	Application Specific Integrated Circuit
ASMI	Advanced Server Management Interface
BIOS	Basic Input/Output System
BIST	Built-In Self Test
BMC	Baseboard Management Controller
BPP	Bits per pixel
Bridge	Circuitry connecting one computer bus to another, allowing an agent on one to access the other
BSP	Bootstrap Processor
Byte	8-bit quantity
CBC	Chassis Bridge Controller (A microcontroller connected to one or more other CBCs, together they bridge the IPMB buses of multiple chassis.)
CEK	Common Enabling Kit
CHAP	Challenge Handshake Authentication Protocol
CMOS	Complementary Metal-oxide-semiconductor In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory, which normally resides on the server board.
DHCP	Dynamic Host Configuration Protocol
DPC	Direct Platform Control
EEPROM	Electrically Erasable Programmable Read-Only Memory
EHCI	Enhanced Host Controller Interface
EMP	Emergency Management Port
EPS	External Product Specification
ESB2	Enterprise South Bridge 2
FBD	Fully Buffered DIMM
F MB	Flexible Mother Board
FRB	Fault Resilient Booting
FRU	Field Replaceable Unit
FSB	Front Side Bus
GB	1024 MB
GPA	Guest Physical Address
GPIO	General Purpose I/O
GTL	Gunning Transceiver Logic
HPA	Host Physical Address
HSC	Hot-swap Controller
Hz	Hertz (1 cycle/second)

Term	Definition
I ² C	Inter-Integrated Circuit Bus
IA	Intel® Architecture
IBF	Input Buffer
ICH	I/O Controller Hub
ICMB	Intelligent Chassis Management Bus
IERR	Internal Error
IFB	I/O and Firmware Bridge
ILM	Independent Loading Mechanism
IMC	Integrated Memory Controller
INTR	Interrupt
I/OAT	I/O Acceleration Technology
IOH	I/O Hub
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IR	Infrared
ITP	In-Target Probe
KB	1024 bytes
KCS	Keyboard Controller Style
KVM	Keyboard, Video, Mouse
LAN	Local Area Network
LCD	Liquid Crystal Display
LDAP	Local Directory Authentication Protocol
LED	Light Emitting Diode
LPC	Low Pin Count
LUN	Logical Unit Number
MAC	Media Access Control
MB	1024 KB
MCH	Memory Controller Hub
MD2	Message Digest 2 – Hashing Algorithm
MD5	Message Digest 5 – Hashing Algorithm – Higher Security
ME	Management Engine
MMU	Memory Management Unit
ms	Milliseconds
MTTR	Memory Type Range Register
Mux	Multiplexor
NIC	Network Interface Controller
NMI	Nonmaskable Interrupt
OBF	Output Buffer
OEM	Original Equipment Manufacturer
Ohm	Unit of electrical resistance
OVP	Over-voltage Protection
PECI	Platform Environment Control Interface
PEF	Platform Event Filtering
PEP	Platform Event Paging

Term	Definition
PIA	Platform Information Area (This feature configures the firmware for the platform hardware)
PLD	Programmable Logic Device
PMI	Platform Management Interrupt
POST	Power-On Self Test
PSMI	Power Supply Management Interface
PWM	Pulse-Width Modulation
QPI	QuickPath Interconnect
RAM	Random Access Memory
RASUM	Reliability, Availability, Serviceability, Usability, and Manageability
RISC	Reduced Instruction Set Computing
RMII	Reduced Media-Independent Interface
ROM	Read Only Memory
RTC	Real-Time Clock (Component of ICH peripheral chip on the server board)
SDR	Sensor Data Record
SECC	Single Edge Connector Cartridge
SEEPROM	Serial Electrically Erasable Programmable Read-Only Memory
SEL	System Event Log
SIO	Server Input/Output
SMBUS*	System Management BUS
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
SPS	Server Platform Services
SSE2	Streaming SIMD Extensions 2
SSE3	Streaming SIMD Extensions 3
SSE4	Streaming SIMD Extensions 4
TBD	To Be Determined
TDP	Thermal Design Power
TIM	Thermal Interface Material
UART	Universal Asynchronous Receiver/Transmitter
UDP	User Datagram Protocol
UHCI	Universal Host Controller Interface
URS	Unified Retention System
UTC	Universal time coordinare
VID	Voltage Identification
VRD	Voltage Regulator Down
VT	Virtualization Technology
Word	16-bit quantity
WS-MAN	Web Services for Management
ZIF	Zero Insertion Force

Reference Documents

- *Advanced Configuration and Power Interface Specification*, Revision 3.0, <http://www.acpi.info/>.
- *Intelligent Platform Management Bus Communications Protocol Specification*, Version 1.0. 1998. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- *Intelligent Platform Management Interface Specification*, Version 2.0. 2004. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- *Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification*, Version 1.1, 02/01/02, Intel Corporation.
- *Intel® Remote Management Module User's Guide*, Intel Corporation.
- *Alert Standard Format (ASF) Specification, Version 2.0, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., <http://www.dmtf.org>.*
- *BIOS for EPSP Platforms Based on Intel® Xeon Processor E5-4600/2600/2400/1600 Product Families External Product Specification*
- *EPSP Platforms Based On Intel Xeon® Processor E5 4600/2600/2400/1600 Product Families BMC Core Firmware External Product Specification*
- *SmaRT & CLST Architecture on "Romley" Systems and Power Supplies Specification (Doc Reference # 461024)*
- *Intel Integrated RAID Module RMS25PB080, RMS25PB040, RMS25CB080, and RMS25CB040 Hardware Users Guide*
- *Intel® Remote Management Module 4 Technical Product Specification*
- *Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide*
- *Intel® Server System R1000GZ/GL Technical Product Specification*
- *Intel® Server System R2000GZ/GL Technical Product Specification*
- *Intel® Deployment Assistant (IDA) 5.2 Users Guide*
- *One Boot Flash Update Installation and Users Guide*
- *Intel® Active System Control Console User Guide*
- *Intel® Intelligent Power Node Manager White Paper*
- *Intel® System Management Software Installation Guide*

Pin	Name	Pin	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#