



Intel® Server Board SE7520AF2

Technical Product Specification

Intel order number C77866-003

Revision 1.2

March, 2005

Enterprise Platforms and Services Marketing

Revision History

Date	Revision Number	Modifications
01/2004	0.5	First release
05/2004	0.97	Second release
05/2004	0.99	Added integrated Intel® RAID Controller SROMBU42E Chapter
08/2004	1.0	First public release
01/2005	1.1	Q1'2005 Update (rolled Spec Update documentation changes and updated BIOS Setup screen changes)
03/2005	1.2	Updated mBMC and Intel Management Module platform sensors tables

Disclaimers

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel® Server Board SE7520AF2 may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

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

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

Copyright © Intel Corporation 2005.

< This page intentionally left blank. >

Table of Contents

1. Introduction	21
1.1 Chapter Outline.....	21
1.2 Server Board Use Disclaimer	22
2. Server Board Overview.....	23
2.1 Intel® Server Board SE7520AF2 SKU Options.....	23
2.2 Intel® Server Board SE7520AF2 Feature Set.....	23
2.3 Server Board Layout.....	25
2.3.1 Component Placement	25
2.3.2 Mechanical Drawing	26
2.3.3 ATX I/O Layout	27
3. Functional Architecture	29
3.1 Intel® E7520 Chipset.....	30
3.1.1 Memory Controller Hub (MCH)	30
3.1.2 PCI-X Hub (PXH).....	34
3.1.3 IOP332 I/O Processor.....	35
3.1.4 I/O Controller Hub (ICH5-R)	35
3.2 Processor Sub-system.....	39
3.2.1 Processor VRD	40
3.2.2 Reset Configuration Logic	40
3.2.3 Processor Module Presence Detection	40
3.2.4 GTL2006.....	41
3.2.5 Common Enabling Kit (CEK) Design Support.....	41
3.3 Memory Sub-System	41
3.3.1 Supported Memory	42
3.3.2 DIMM Population Rules and Supported Configurations	42
3.3.3 Single Channel Operation.....	44
3.3.4 I ² C Bus.....	44
3.3.5 Memory Voltage Margining	44
3.3.6 Memory RASUM Features.....	44
3.4 I/O Sub-System	51
3.4.1 PCI Subsystem	51
3.4.2 Interrupt Routing	55

3.4.3	SCSI Support.....	60
3.4.4	RAID Functionality	62
3.4.5	Parallel ATA (PATA) Support	63
3.4.6	Serial ATA (SATA) Support	64
3.4.7	Video Controller	65
3.4.8	Network Interface Controller (NIC)	67
3.4.9	USB 2.0 Support.....	68
3.4.10	Super I/O	68
3.4.11	BIOS Flash	70
3.5	Configuration and Initialization.....	70
3.5.1	Memory Space.....	70
3.5.2	I/O Map	76
3.5.3	Accessing Configuration Space.....	78
3.5.4	Hardware Initialization	80
3.6	Clock Generation and Distribution	81
4.	Integrated Intel® RAID Controller SROMBU42E.....	83
4.1	Overview.....	83
4.2	Primary Integrated Intel RAID Controller SROMBU42E Features	83
4.2.1	Configuration on Disk	85
4.2.2	Array Performance Features	85
4.2.3	Fault Tolerance Capabilities	85
4.3	RAID Software Stack	86
4.3.1	General RAID Features	86
4.3.2	RAID Level Migration	87
4.3.3	Fault Tolerant Features	87
4.3.4	Cache Options and Settings	87
4.3.5	Background Tasks	88
4.3.6	Error Handling.....	88
4.3.7	Decoding an Audible Alarm	89
4.4	Levels of RAID	89
4.4.1	RAID 0 - Data Striping	89
4.4.2	RAID 1 - Disk Mirroring/Disk Duplexing	90
4.4.3	RAID 5 - Data Striping with Striped Parity	90
4.4.4	RAID 10 - Combination of RAID 1 and RAID 0.....	91
4.4.5	RAID 50 - Combination of RAID 5 and RAID 0.....	91

4.4.6	Levels of Drive Hierarchy within the Intel® Integrated RAID Firmware	92
4.5	Intel® RAID Controller Drivers	93
4.6	Intel® RAID Web Console	93
4.7	Intel® RAID BIOS Console Configuration Utility	94
4.7.1	Quick Configuration Steps	94
4.7.2	Starting the BIOS Console Utility on the Host Computer	94
4.7.3	Screen and Option Descriptions	95
4.7.4	Configuration Mismatch Screen	98
4.7.5	Detailed Configuration Instructions	98
4.7.6	Finish Configuration	102
4.7.7	Initialize the Drive	102
5.	System BIOS	105
5.1	BIOS Identification String	105
5.2	Supported BIOS Features	105
5.3	Processor Initialization	106
5.3.1	Multiple Processor Initialization	106
5.3.2	Mixed Processor Steppings	107
5.3.3	Mixed Processor Models	107
5.3.4	Mixed Processor Families	107
5.3.5	Mixed Processor Cache Sizes	107
5.3.6	Jumperless Processor Speed Settings	107
5.3.7	Microcode	108
5.3.8	Processor Cache	108
5.3.9	Hyper-Threading Technology	108
5.3.10	Intel SpeedStep® Technology	108
5.3.11	Intel® Extended Memory 64 Technology (Intel® EM64T)	108
5.4	Memory Initialization	108
5.4.1	Memory Sizing	109
5.4.2	Disabling DIMMs	109
5.4.3	Memory On-Line Sparing and Mirroring	109
5.4.4	ECC Memory Initialization	110
5.4.5	Memory Population	110
5.4.6	DIMM Module Capacity	110
5.4.7	Memory Error Handling	111
5.4.8	Memory Test	111

5.5	PCI Initialization	112
5.5.1	Scan Order	112
5.5.2	Resource Assignment.....	112
5.5.3	Automatic IRQ Assignment.....	112
5.5.4	Option ROMs	112
5.5.5	PCI APIs	113
5.5.6	Split Option ROM	113
5.5.7	Dual Video	113
5.5.8	PCI Hot Plug Initialization	113
5.5.9	PCI Express* Initialization.....	116
5.5.10	PCI Express* Enhanced Configuration Mechanisms.....	116
5.5.11	Legacy Universal Serial Bus (USB) Initialization	116
5.5.12	IDE Initialization	116
5.5.13	Removable Media Initialization	117
5.6	Flash ROM.....	117
5.7	BIOS User Interface.....	117
5.7.1	System State Window.....	118
5.7.2	Logo/Diagnostic Window	118
5.7.3	Current Activity Window.....	118
5.8	System Diagnostic Screen.....	118
5.8.1	Static Information Display	118
5.9	Quiet Boot / OEM Splash Screen	119
5.10	BIOS Boot Popup Menu	119
5.11	BIOS Setup Utility	119
5.11.1	Localization.....	120
5.11.2	Console Redirection	120
5.11.3	Configuration Reset	120
5.11.4	Keyboard Commands	120
5.11.5	Entering BIOS Setup	121
5.12	Flash Architecture and Flash Update Utility.....	137
5.13	Rolling BIOS and On-line Updates	138
5.14	Flash Update Utility.....	138
5.14.1	Flash BIOS	139
5.14.2	User Binary Area	140
5.14.3	Recovery Mode.....	140

5.14.4	Update OEM Logo	141
5.15	OEM Binary	142
5.15.1	Scan Point Definitions.....	144
5.15.2	Format of the OEM Binary Structure.....	145
5.16	Quiet Boot / OEM Splash Screen	145
5.17	Operating System Boot, Sleep, and Wake	146
5.17.1	Microsoft* Windows* Compatibility	146
5.17.2	Advanced Configuration and Power Interface (ACPI)	146
5.17.3	Sleep and Wake Functionality	147
5.17.4	On to Off (operating system-Present) or operating system to Sleep	147
5.17.5	Sleep to On (ACPI)	148
5.17.6	System Sleep States	148
6.	Platform Management.....	149
6.1	Management Architecture Overview.....	149
6.1.1	Tiered Server Management Model	149
6.1.2	5V Standby	153
6.1.3	IPMI Messaging, Commands, and Abstractions	154
6.1.4	IPMI 'Sensor Model'	154
6.1.5	Private Management Bus	155
6.1.6	Management Controllers	156
6.2	Onboard Platform Instrumentation Management Features and Functionality	158
6.2.1	mBMC Self-test.....	159
6.2.2	SMBus Interfaces	159
6.2.3	External Interface to mBMC.....	159
6.2.4	Messaging Interfaces.....	160
6.2.5	Direct Platform Control (IPMI over LAN).....	162
6.2.6	Wake On LAN / Power On LAN and Magic Packet Support.....	164
6.2.7	Watchdog Timer	165
6.2.8	System Event Log (SEL)	165
6.2.9	Sensor Data Record (SDR) Repository	166
6.2.10	Event Message Reception	166
6.2.11	Event Filtering and Alerting.....	166
6.2.12	NMI Generation	169
6.3	Platform Management Interconnects	170
6.3.1	Power Supply Interface Signals	170

6.3.2	System Reset Control	171
6.3.3	Temperature-based Fan Speed Control	172
6.3.4	Front Panel Control.....	172
6.3.5	FRU Information	176
6.4	Sensors.....	176
6.4.1	Sensor Type Codes	176
7.	Error Reporting and Handling	189
7.1	Fault Resilient Booting (FRB)	189
7.1.1	FRB-3 – BSP Reset Failures	189
7.1.2	FRB-2 – BSP POST Failures.....	189
7.1.3	FRB-1 – BSP Self-Test Failures	190
7.1.4	OS Watchdog Timer: Operating System Load Failures.....	190
7.1.5	AP Failures	191
7.1.6	Treatment of Failed Processors.....	191
7.2	Error Logging	192
7.2.1	Error Sources and Types	192
7.2.2	SMI Handler	192
7.2.3	BIOS Generated IPMI Events	197
7.2.4	Single Bit ECC Error Throttling Prevention	198
7.3	Error Messages and Error Codes	199
7.3.1	POST Progress Codes and Messages	199
7.3.2	BIOS Messages.....	206
7.3.3	POST Error Messages and Handling	212
7.3.4	Boot Block Error Beep Codes.....	215
7.3.5	POST Error Beep Codes	215
7.3.6	"POST Error Pause" option.....	216
8.	Connectors, Headers and Jumpers	217
8.1	Board Connector Information.....	217
8.2	Main Power Connector	218
8.3	Main Memory Module Connector.....	219
8.4	RAID Memory Module Connector	221
8.5	Processor Socket.....	222
8.6	System Management Headers	225
8.6.1	Intel Management Module Connector.....	225
8.6.2	ICMB Header	228

8.6.3	IPMB Header	229
8.6.4	OEM RMC Header	229
8.6.5	HSBP Header	229
8.7	PCI Slot Connector	229
8.8	Front Panel Connectors	233
8.9	VGA Connector	234
8.10	SCSI Connector	234
8.11	NIC Connectors	235
8.12	ATA Connector	236
8.13	USB Connector	236
8.14	Floppy Connector	237
8.15	Serial Port Connector	238
8.16	Keyboard and Mouse Connector	238
8.17	Fan Headers	239
9.	Configuration Jumpers	241
9.1	System Recovery and Update Jumpers	241
9.2	Rolling BIOS Bank Selection Jumper	242
10.	General Specifications	243
10.1	Absolute Maximum Ratings	243
10.2	Processor Power Support	243
10.3	SE7520AF2 Power Budget	244
10.4	Power Supply Specifications	244
10.4.1	Power Timing	244
10.4.2	Voltage Recovery Timing Specifications	246
11.	Product Regulatory Compliance	249
11.1.1	Product Safety Compliance	249
11.1.2	Product EMC Compliance	249
11.1.3	Product Regulatory Compliance Markings	249
11.2	Electromagnetic Compatibility Notices	250
11.2.1	FCC (USA)	250
11.2.2	Industry Canada (ICES-003)	250
11.2.3	Europe (CE Declaration of Conformity)	250
11.2.4	Australian Communications Authority (ACA) (C-Tick Declaration of Conformity)	251
11.2.5	Ministry of Economic Development (New Zealand) Declaration of Conformity ...	251
11.2.6	Korean RRL Compliance	251

11.2.7 BSMI (Taiwan) 251

11.3 Replacing the Back-Up Battery 251

Integration and Usage Tips I

Glossary II

Reference Specifications and Documents V

List of Figures

Figure 1. Board Layout	25
Figure 2. Board Dimensions	26
Figure 3. ATX Rear IO Connectors	27
Figure 4. Server Board Block Diagram	29
Figure 5. CEK 'Passive' Component Stackup	41
Figure 6. DIMM Socket Configuration	43
Figure 7. Four DIMM Memory Mirror Configuration	48
Figure 8. Six DIMM Memory Mirror Configuration	49
Figure 9. Eight DIMM Memory Mirror Configuration	50
Figure 10. Interrupt Routing Diagram (ICH5-R Internal)	58
Figure 11. Interrupt Routing Diagram	59
Figure 12. Intel® Portable Cache Module	63
Figure 13. Video Controller PCI Bus Interface	67
Figure 14. Intel® Xeon Processor Memory Address Space	71
Figure 15. DOS Compatibility Region	72
Figure 16. Extended Memory Map	74
Figure 17. CONFIG_ADDRES Register	79
Figure 18. Intel® RAID Activation Key	83
Figure 19. RAID 0	90
Figure 20. RAID 1	90
Figure 21. RAID 5	91
Figure 22. RAID 5	91
Figure 23. RAID 50	92
Figure 24. Inte® RAID BIOS Console Screen	95
Figure 25. Adapter Selection Screen	98
Figure 26. Configuration Wizard Screen	99
Figure 27. Array Definition Screen	100
Figure 28. Logical Drive Definition	100
Figure 29. Configuration Wizard Preview Screen	102
Figure 30. Logic Drives Screen	103
Figure 31. SE7520AF2 BIOS Identification String	105
Figure 32. Enhanced Configuration Memory Address Map	116

Figure 33. BIOS Screen Display Areas.....	118
Figure 34. Sample Pop-up Window	119
Figure 35. Block Diagram of Platform Managment Architecture	153
Figure 36. mBMC in a Server Management System.....	159
Figure 37. External Interfaces to mBMC.....	160
Figure 38. IPMI-over-LAN	163
Figure 39. Power Supply Control Signals	170
Figure 40. Location of Diagnostic LEDs on Baseboard	200
Figure 41. SE7520AF2 Configuration Jumpers (J1D1)	241
Figure 42. SE7520AF2 BIOS Bank Jumper (J2J6).....	242
Figure 43. Output Voltage Timing	245
Figure 44. Turn on / off Timing.....	246

List of Tables

Table 1. PCI Express* Hot-Plug LED Function Table.....	33
Table 2. PCI-X Hot-Plug LED Function Table.....	35
Table 3. ICH5-R GPIO Assignment	38
Table 4. Processor Support Matrix	40
Table 5. SE7520AF2 Supported DIMM Modules.....	42
Table 6. I ² C Addresses for Memory Module SMB	44
Table 7. PCI Bus Segment Characteristics.....	52
Table 8. ICH5-R P32-A Configuration IDs	52
Table 9. ICH5-R P32-A Arbitration Connections.....	52
Table 10. PXH P64-A Configuration IDs.....	53
Table 11. PXH P64-B Configuration IDs.....	53
Table 12. PXH P64-A Arbitration Connections	54
Table 13. PHX P64-B Arbitration Connections	54
Table 14. IOP332 P64-A Configuration IDs	54
Table 15. IOP332 P64-B Configuration IDs	55
Table 16. IOP332 P64-B Arbitration Connections	55
Table 17. IOP332 P64-B Arbitration Connections	55
Table 18. PCI Interrupt Routing/Sharing.....	56
Table 19. Interrupt Definitions.....	57
Table 20. Intel® Server Board SE7520AF2 Video Modes	65
Table 21. Video Memory Interface.....	66
Table 22. NIC2 Status LED.....	68
Table 23. Super I/O GPIO Usage Table	68
Table 24. Serial A Header Pin-out	69
Table 25. SMM Space Transaction Handling	76
Table 26. SMM Space Table	76
Table 27. Intel® Server Board SE7520AF2 I/O Map	76
Table 28. PCI Configuration IDs and Device Numbers.....	80
Table 29. Clock Generation and Distribution	81
Table 30. Integrated Intel® RAID Controller SROMBU42E Features.....	84
Table 31. RAID Controller SROMBU42E Array Performance.....	85
Table 32. RAID Controller SROMBU42E Fault Tolerance.....	85

Table 33. RAID Level Migration Matrix	87
Table 34. Adapter Properties Menu Options.....	96
Table 35. BIOS Features	105
Table 36. DIMM Population	110
Table 37. Module Capacity	110
Table 38. PCI Express* Hot-Plug LED Function Table.....	115
Table 39. BIOS Setup Keyboard Command Bar Options	120
Table 40. BIOS Setup Main Menu Options	121
Table 41. BIOS Setup Advance Menu Options.....	122
Table 42. BIOS Setup CPU Configuration Menu Options.....	123
Table 43. BIOS Setup IDE Configuration Menu Options	124
Table 44. BIOS Setup, IDE Device Configuration Sub-menu Selections	125
Table 45. BIOS Setup floppy Configuration Sub-menu Selections.....	126
Table 46. BIOS Setup SuperIO Configuration Sub-menu.....	127
Table 47. BIOS Setup USB Configuration Sub-menu Selections	127
Table 48. BIOS Setup, USB Mass Storage Device Configuration Sub-menu Selections	128
Table 49. BIOS Setup PCI Configuration Sub-menu Selections	128
Table 50. BIOS Setup, Memory Configuration Sub-menu Selections.....	129
Table 51. BIOS Setup, Boot Menu Selections	131
Table 52. BIOS Setup, Boot Settings Configuration Sub-menu Selections	131
Table 53. BIOS Setup, Boot Device Priority Sub-menu Selections	132
Table 54. BIOS Setup, Hard Disk Drive Sub-Menu Selections.....	132
Table 55. BIOS Setup, Removable Drives Sub-menu Selections.....	132
Table 56. BIOS Setup, ATAPI CDROM Drives Sub-menu Selections.....	133
Table 57. BIOS Setup, Security Menu Options.....	133
Table 58. BIOS Setup, Server Menu Selections.....	134
Table 59. BIOS Setup, System Management Sub-menu Selections	135
Table 60. BIOS Setup Serial Console Features Sub-menu Selections	136
Table 61. BIOS Setup, Event Log Configuration Sub-menu Selections	136
Table 62. BIOS Setup, Exit Menu Selections	137
Table 63. Scan Point Definitions	144
Table 64. Format of OEM Binary Structure.....	145
Table 65. Supported Wake Events	148
Table 66. Tiered Platform Management Feature Overview	150
Table 67. Power and Reset Control.....	151

Table 68. Secure Mode Button Actions	151
Table 69. Memory RAS Feature Support by SM Tier	151
Table 70. Supported Channel Assignments	161
Table 71. LAN Channel Capacity.....	162
Table 72. LAN Channel Specifications	164
Table 73. PEF Action Priorities	167
Table 74. mBMC Factory Default Event Filters.....	167
Table 75. Power Control Initiators.....	171
Table 76. System Reset Sources and Actions.....	172
Table 77. Chassis ID LEDs.....	174
Table 78. Fault/Status LED.....	174
Table 79. mBMC Built-in Sensors.....	177
Table 80. Intel® Server Board SE7520AF2 Sensors for OB Platform Instrumentation Management	178
Table 81. Platform Sensors for the Intel® Management Module	180
Table 82. Memory Error Events	194
Table 83. Examples of Event Data Field Contents for Memory Errors	195
Table 84. PCI error events.....	195
Table 85. Event Data Field Contents for PCI Errors	196
Table 86. FRB-2 Error Events.....	196
Table 87. Event Data Field Contents for FRB-2 Errors.....	197
Table 88. BIOS Generated IPMI Events	197
Table 89. POST Progress Code LED Example	199
Table 90. POST Code Checkpoints.....	200
Table 91. Bootblock Initialization Code Checkpoints.....	202
Table 92. Bootblock Recovery Code Checkpoints.....	203
Table 93. DIM Code Checkpoints	204
Table 94. ACPI Runtime Checkpoints	205
Table 95. Memory BIOS Messages	206
Table 96. Boot BIOS Messages.....	206
Table 97. Storage Device BIOS Messages	207
Table 98. Virus Related BIOS Messages	209
Table 99. System Configuration BIOS Messages.....	210
Table 100. CMOS BIOS Messages	211
Table 101. Miscellaneous BIOS Messages	211

Table 102. USB BIOS Error Messages.....	211
Table 103. SMBIOS BIOS Error Messages	212
Table 104. POST Error Messages and Handling.....	212
Table 105. Management Module POST Error Codes and Messages.....	214
Table 106. Bootblock Error Beep Codes	215
Table 107. POST Error Beep Codes	215
Table 108. Troubleshooting BIOS Beep Codes.....	216
Table 109. Board Connector Matrix	217
Table 110. Power Connector Pin-out (J9C1).....	218
Table 111. Power Supply Signal Connector (J9B1).....	218
Table 112. 12V Power Connector (J9E1)	218
Table 113. DIMM Connectors (J9D1, J9D2, J8D2, J8D3, J7D3, J8D1, J7D1, J7D2).....	219
Table 114. RAID DIMM Connector (J1D5)	221
Table 115. Socket 604 Processor Socket Pinout (J8H1, J6H1)	222
Table 116. IMM Connector Pin-out (J3J1).....	225
Table 117. ICMB Header Pin-out (J1A1)	228
Table 118. IPMB Header Pin-out (J1G1).....	229
Table 119. IPMB Header Pin-out (J1B2)	229
Table 120. HSBP Header Pin-out (J1K3, J1K2)	229
Table 121. PCI Slot Characteristics	230
Table 122. Slot 1 and 5 PCI-X 64bit 3.3V Pin-out (J1D4, J4D1)	230
Table 123. Slot 3 and 4 PCI Express* Pin-out (J2C1, J3C1).....	231
Table 124. Slot 6 PCI-X 64-bit 3.3V Pin-out (J4D2, Riser Capable).....	232
Table 125. Front Panel 34-Pin Header Pin-out (J1E1)	233
Table 126. VGA Connector Pin-out (J8A1).....	234
Table 127. 68-pin SCSI Connector Pin-out (J1H1, J1F1).....	234
Table 128. NIC1 and NIC2 1.0Gb RJ45 Connector Pin-out (J6A1).....	235
Table 129. ATA-100 40-pin Connector Pin-out (J3K1)	236
Table 130. USB Connectors Pin-out (J9A2)	236
Table 131. Optional USB Connection Header Pin-out (J1K4)	237
Table 132. Legacy 34-pin Floppy Connector Pin-out (J3K2).....	237
Table 133. Rear DB-9 Serial A Port Pin-out (J8A1).....	238
Table 134. 9-pin Header Serial B Port Pin-out (J1B1).....	238
Table 135. Keyboard and Mouse PS2 Connector Pin-out (J9A1)	238
Table 136. 3-pin CPU Fan Headers Pin-out (J6F1, J5F1).....	239

Table 137. 3 pin System Fan Headers Pin-out (J5A2, J5A1)	239
Table 138. 3+2 pin System Fan Headers Pin-out (J2J3+J2K3, J2J2+J2K1).....	239
Table 139. 6-pin System Fan Headers Pin-out (J2K4, J2K2).....	240
Table 140. Configuration Jumper Options	241
Table 141. BIOS Bank Jumper Option.....	242
Table 142. Absolute Maximum Ratings	243
Table 143. Intel® Xeon™ processor DP TDP Guidelines.....	243
Table 144. SE7520AF2 Power Budget.....	244
Table 145. SE7520AF2 Power Supply Voltage Specification.....	244
Table 146. Voltage Timing Parameters	245
Table 147. Turn On / Off Timing	245
Table 148. Transient Load Requirements.....	247

< This page intentionally left blank. >

1. Introduction

This Technical Product Specification (TPS) provides detail on the architecture and feature set of the Intel® Server Board SE7520AF2.

The target audience for this document is anyone wishing to obtain more in depth detail of the server board than what is generally made available in the board's Users Guide. It is a technical document meant to assist the audience that need higher level of understanding and technical details about the specific features of the board.

This is one of several technical documents available for this server board. All of the functional sub-systems that make up the board are described in this document. However, low-level details of specific sub-systems are not included. Design level information for specific sub-systems can be obtained by ordering the External Product Specification (EPS) or External Architecture Specification (EAS) for a given sub-system. The EPS documents available for this server board include the following:

- Intel® Server Board SE7520AF2 BIOS EPS
- Mini Baseboard Management Controller (mBMC) EPS
- Intel® Server Board SE7520AF2 Baseboard Management Controller (BMC) EPS
- Server Management EAS
- Hot Swap Controller Interface EPS

These documents are not made publicly available and must be ordered by your local Intel representative.

1.1 Chapter Outline

This document is divided into the following chapters

- Chapter 1 – Introduction
- Chapter 2 – Server Platform Overview
- Chapter 3 – Functional Architecture
- Chapter 4 – Integrated Intel® RAID Controller SROMBU42E
- Chapter 5 – System BIOS Architecture
- Chapter 6 – Platform Management Architecture
- Chapter 7 – Error Reporting and Handling
- Chapter 8 – Connectors and Headers
- Chapter 9 – Configurations Jumpers
- Chapter 10 – General Specifications
- Chapter 11 – Product Regulatory Compliance
- Appendix A – Integration and Usage Tips

1.2 Server Board Use Disclaimer

Intel Corporation server boards contain a number of high-density VLSI and power delivery components that need adequate airflow for cooling. Intel ensures through its own chassis development and testing that when Intel server building blocks are used together, the fully integrated platform will meet the intended thermal requirements of these components. It is the responsibility of the platform 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. Server Board Overview

The Intel® Server Board SE7520AF2 is a monolithic printed circuit board with features designed to meet the high-performance, high-reliability, availability and serviceability pedestal and rack server markets.

2.1 Intel® Server Board SE7520AF2 SKU Options

In this document, the product name Intel® Server Board SE7520AF2 is used to describe the family of boards that will be made available under a common marketing name. The core features for each board are common; however each board will have the following distinctions:

- Board SKU 1 – SE7520AF2: Non Hot-Plug PCI SKU
- Board SKU 2 – SE7520HPAF2: Hot-Plug PCI SKU

Throughout this document, all references to the Intel® Server Board SE7520AF2 refer to all board SKUs unless specifically noted otherwise. The board you select to use may or may not include all features described based on the listed board differences.

2.2 Intel® Server Board SE7520AF2 Feature Set

- IA32 server platform with support for dual Intel® Xeon™ processors in the FC-mPGA4 package using the Socket 604 and 800 MT/s Front Side Bus
- Intel® E7520 Memory Controller Hub (MCH)
- Intel® 6700PXH PCI-X Controller Hub (PXH)
- Intel® 80332 I/O Processor with Intel® XScale Technology (IOP332)
- Intel® 82801ER I/O Controller Hub5 (ICH5-R)
- Support for eight DDR2-400MHz compliant registered ECC DIMMs providing up to 16GB of memory in a dual channel architecture
- Five full-length full-height PCI expansion slots:
 - PCI Slot 1: PCI-X 64-bit/133 MHz (Hot plug capable on the SE7520HPAF2 SKU)
 - PCI Slot 3: PCI Express* x4 (Hot plug capable on the SE7520HPAF2 SKU)
 - PCI Slot 4: PCI Express* x8 (Hot plug capable on the SE7520HPAF2 SKU)
 - PCI Slot 5: PCI-X 64-bit/133 MHz (Hot plug capable on the SE7520HPAF2 SKU)
 - PCI Slot 6: PCI-X 64-bit/100 MHz (One slot and two slot riser capable)
- 2D/3D graphics controller: ATI Rage* XL video controller with 8 MB of SDRAM
- Intel® 10/100/1000 82546GB Dual Gigabit Ethernet controller available via two RJ-45 stacked connectors
- Dual-channel LSI Logic* 53C1030 Ultra-320 wide SCSI controller supporting entry level RAID functionality (LSI Logic* Integrated Mirroring* and Integrated Striping*)
- RAID On Motherboard (ROMB)¹ via the IOP332 with support for a maximum of 1 GB of unbuffered ECC DDR 333MHz RAID cache
- Intel® Portable Cache Module for ROMB battery backup support (optional accessory kit)

¹ Requires optional RAID enabling accessory

- Two Serial ATA 150 ports from the ICH5-R supporting entry level RAID functionality (Integrated Mirroring and Integrated Striping)
- LPC (Low Pin Count) bus segment with two embedded devices:
 - Mini-Baseboard Management Controller (mBMC) providing monitoring, alerting, and logging of critical platform information obtained from embedded sensors on server board
 - Super I/O controller chip enabling all PC-compatible I/O (floppy, serial, parallel, keyboard, mouse)
- Intel® Management Module (IMM) Connector supporting platform instrumentation upgrades to either the Professional or Advanced Management Modules (optional accessory kits)
- Three external Universal Serial Bus (USB) 2.0 ports with an additional internal header providing two optional USB ports for front panel support
- Two stacked PS/2* ports for keyboard and mouse (interchangeable)
- Two serial ports: One external serial port (Serial A) on the rear I/O area of the board and one internal header is also available providing an optional port (Serial B)
- One IDE connector, supporting up to two ATA-100 compatible devices
- One standard floppy drive interface
- Six multi-speed system fan headers and two single speed CPU fan headers
- Multiple server management headers providing on-board interconnects to the board's server management features
- SSI-EEB3.5 compliant board form factor (12 x 13 inches)
- SSI-compliant connectors for SSI interface supporting the 34-pin front panel, floppy, ATA-100 and power connectors
- Custom LCD connector enabling support for the Intel® Local Control Panel (optional accessory)
- Intel® Light Guided Diagnostics on critical FRU devices (processors, memory, power)
- Port-80 Diagnostic LEDs to display Port-80 HEX codes

<The rest of this page intentionally left blank>

2.3 Server Board Layout

2.3.1 Component Placement

The following figure shows the board layout of the server board.

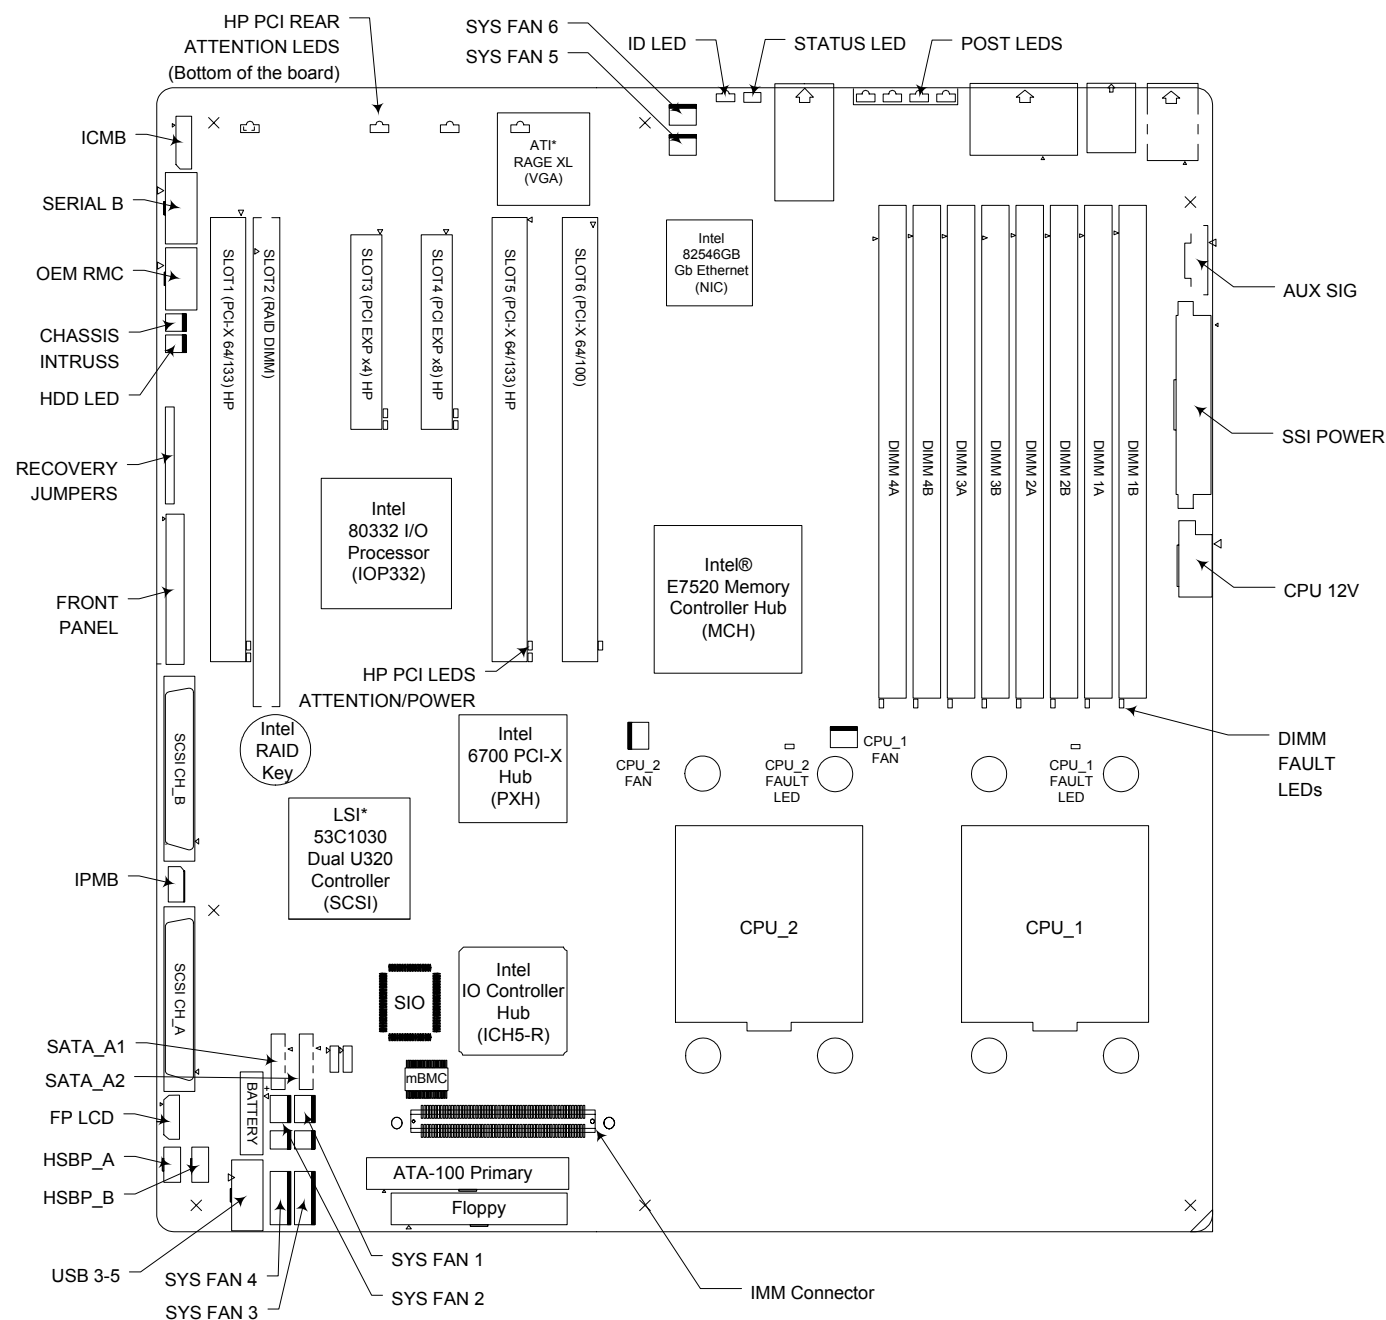


Figure 1. Board Layout

2.3.2 Mechanical Drawing

The following mechanical drawing shows the physical dimensions of the server board.

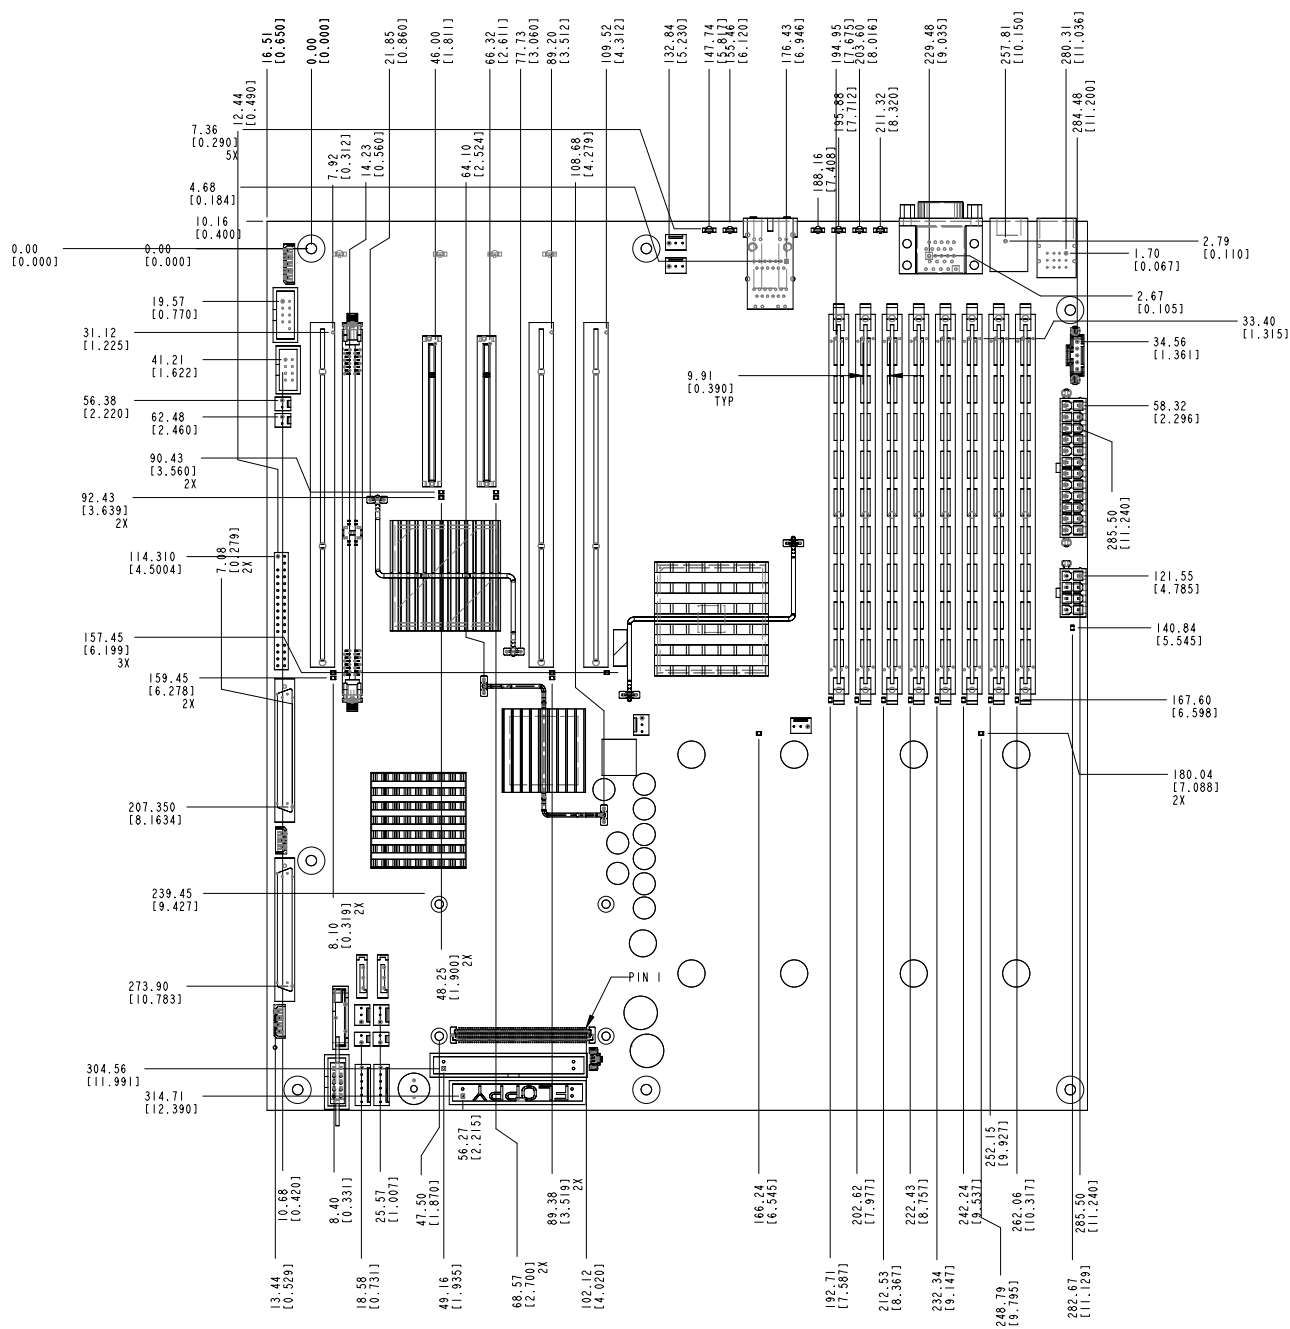


Figure 2. Board Dimensions

2.3.3 ATX I/O Layout

The following figure shows the ATX rear I/O layout of the Server Board SE7520AF2

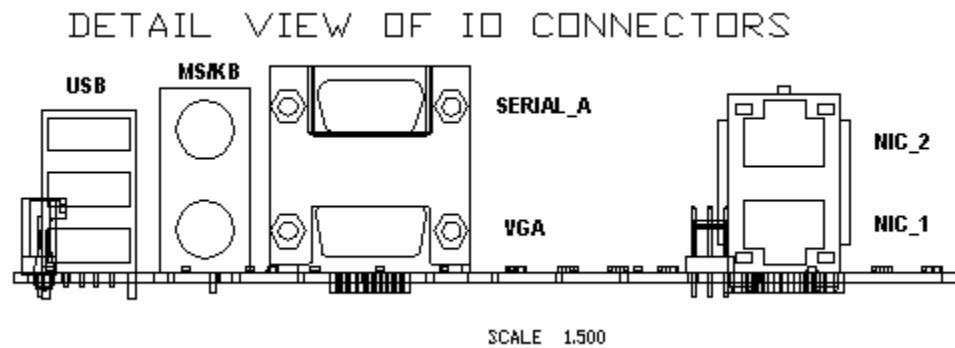


Figure 3. ATX Rear IO Connectors

< This page intentionally left blank. >

3. Functional Architecture

This chapter provides a high-level description of the functionality associated with the architectural blocks that comprise the Intel® Server Board SE7520AF2.

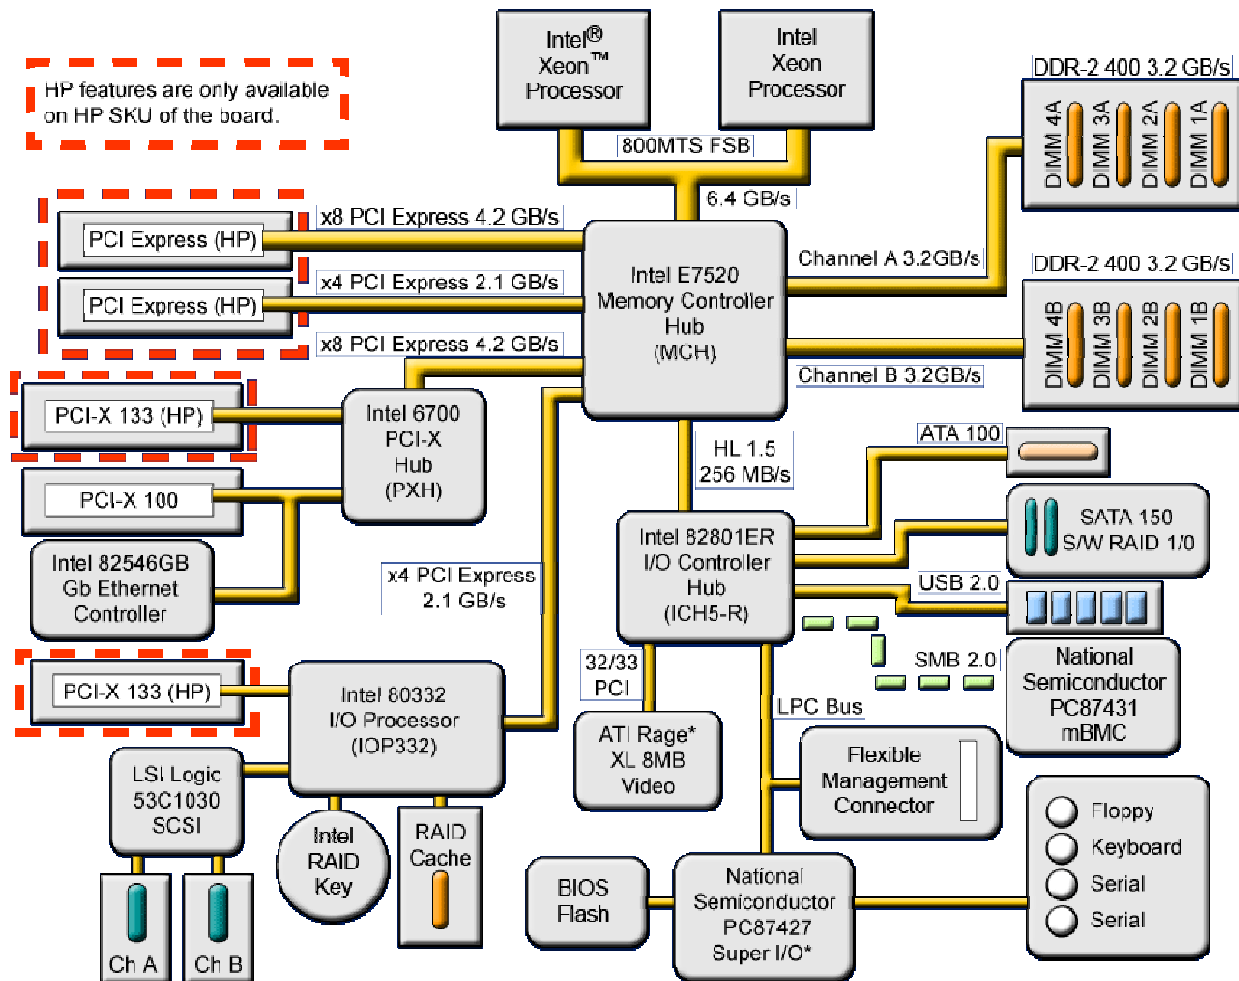


Figure 4. Server Board Block Diagram

3.1 Intel® E7520 Chipset

The architecture of the Intel® Server Board SE7520AF2 is designed around the Intel® E7520 chipset. The chipset used on this board consist of four components which together are responsible for providing the interface between all major sub-systems found on the baseboard including the processor, memory, and I/O sub-systems. These components are:

- Intel® E7520 Memory Controller Hub (MCH)
- Intel® 82801ER I/O Controller Hub (ICH5-R)
- Intel® 6700PXH PCI-X Hub (PXH)
- Intel® 80332 I/O Processor with Intel® XScale Technology (IOP332)

The following sub-sections provide an overview of the primary functions and supported features of each chipset component used on the Intel® Server Board SE7520AF2. Later sections in this chapter provide more detail on how each sub-system is implemented.

Note: The Intel® 80332 I/O processor emulates a second PXH on the platform. When the Intel® RAID Activation Key accessory kit and RAID memory are installed, the RAID On Motherboard (ROMB) function of the I/O processor is enabled to work with the LSI Logic* 53C1030 SCSI controller to provide RAID functionality.

3.1.1 Memory Controller Hub (MCH)

The MCH is a 1077-ball FC-BGA package which supports the following interfaces:

- CPU Front-Side Bus at 200MHz operation using AGTL+ (Assisted Gunner Transceiver Logic) signaling, 4x 64-bit data bus at 6.4GB/s.
- Dual memory channels supporting registered 64-bit data ECC DDR2 400MHz DIMMs with bandwidth of 3.2GB/s per channel giving a total of 6.4GB/s for both channels.
- Three PCI Express* x8 interfaces with aggregate bandwidth of 4.2GB/s each. These interfaces are distributed as follows:
 - One x8 interface to the PXH
 - One x8 interface enabling PCI Slot 4
 - Third x8 PCI Express* interface is configured as two independent x4 interfaces, which connect to the IOP332 and PCI Slot 3¹
- Hub Interface 1.5, with 8-bit data bus running at 66MHz, resulting in a 266MB/s interface to the ICH5-R.
- RASUM support through memory features including Memory Mirroring and Memory Sparing.

The following sub-sections will provide an overview describing the primary functions and supported features of each chipset component as used on the Intel® Server Board SE7520AF2. Later sections in this chapter provide more detail on how each sub-system is implemented.

¹ PCI Slot 3 has a physical x8 PCI Express connector, but it is plumbed with a x4 interface.

3.1.1.1 Front Side Bus (FSB)

The Intel® E7520 MCH supports either single or dual processor configurations using the Intel® Xeon™ processor with 1 MB L2 cache or the Intel® Xeon™ processor with 2 MB L2 cache. The MCH supports a base system bus frequency of 200 MHz. The address and request interface is double pumped to 400 MHz while the 64-bit data interface (+ parity) is quad pumped to 800 MHz. This provides a matched system bus address and data bandwidths of 6.4 GB/s.

3.1.1.2 MCH Memory Sub-System Overview

The MCH provides an integrated memory controller for direct connection to two channels of registered DDR2 400 MHz memory (stacked or unstacked). Peak theoretical memory data bandwidth using DDR2 400 MHz technology is 6.4 GB/s.

When both memory channels are populated and operating, they function in lock-step mode. On the Intel® Server Board SE7520AF2, the maximum supported DDR2 400 MHz memory configuration is 16 GB.

The Intel® E7520 MCH memory interface provides several RASUM (Reliability, Availability, Serviceability, Usability and Manageability) features, including:

- Memory mirroring that allows for two copies of all data in the memory subsystem (one on each channel) to be maintained.
- Memory sparing that allows for one DIMM per channel to be held in reserve and brought on-line if another DIMM in the channel becomes defective. DIMM sparing and memory mirroring are mutually exclusive of one another.
- Hardware periodic memory scrubbing, including demand scrub support.
- Retry on uncorrectable memory errors.
- Intel® x4 Single Device Data Correction (SDDC) for memory error detection and correction of any number of bit failures in a single x4 memory device.

3.1.1.3 PCI Express* Interface

The Intel® E7520 MCH is the first Intel chipset to support the new PCI Express* high-speed serial I/O interface for superior I/O bandwidth. The scalable PCI Express* interface of the MCH complies with the *PCI Express Interface Specification, Rev 1.0a*.

The MCH provides three x8 PCI Express* interfaces, each with a max theoretical bandwidth of 4.2GB/s. Each of these x8 PCI Express* interfaces may alternatively be configured as two independent x4 PCI Express* interfaces. A PCI Express* interface/port is defined as a collection of lanes. Each lane (x1) consists of two striped differential pairs in each direction (transmit and receive). The raw bit-rate on the data pins of 2.5 Gb/s, results in a real bandwidth per pair of 250MB/s given the 8/10 bit encoding used to transmit data across this interface.

The Intel® E7520 MCH is a root class component as defined in the *PCI Express Interface Specification*. The PCI Express* interfaces of the MCH support connection to a variety of bridges and devices compliant with the same revision of the specification.

3.1.1.3.1 *PCI Express* Training*

To establish a connection between PCI Express* endpoints, the endpoints participate in a sequence of steps known as training. This sequence establishes the operational width of the link and adjusts skews of the various lanes within a link so that the data sample points can correctly take a data sample from the link.

In the case of a x8 port, the x4 link pairs first attempt to train independently, and will collapse to a single link at the x8 width upon detection of a single device returning link ID information upstream. Once the number of links has been established, they negotiate to train at the highest common width, and step down in its supported link widths in order to succeed in training. The result may be that the link has trained as a x1 link.

Although the bandwidth of this link size is substantially lower than a x8 link or even a x4 link, it allows communication between the two devices. Software can then interrogate the device at the other end of the link to determine why it failed to train at a higher width (something that would not be possible without support for the x1 link width).

Width negotiation is done only during training or retraining, not during recovery.

3.1.1.3.2 *PCI Express* Retry*

The PCI Express* interface incorporates a link level retry mechanism. The hardware detects when a transmission packet is corrupted and performs a retry of that packet and all following packets. Although this causes a temporary interruption in the delivery of packets, the retry helps to maintain the link integrity.

3.1.1.3.3 *PCI Express* Link Recovery*

If excessive errors occur, the hardware may determine that the quality of the connection is in question, and the end points can enter a quick training sequence known as recovery. The width of the connection will not be renegotiated, but the adjustment of skew between lanes of the link may occur. This occurs without any software intervention, but the software may be notified.

3.1.1.3.4 *PCI Express* Data Protection*

The PCI Express* high-speed serial interface uses traditional CRC protection. The data packets use a 32-bit CRC protection scheme, the same CRC-32 used by Ethernet. The smaller link packets use a 16-bit CRC scheme. Since packets utilize 8B/10B encoding, and not all encodings are used; this provides further data protection, as illegal codes can be detected. If errors are detected on the reception of data packets due to various transients, these data packets can be retransmitted. Hardware logic supports this link-level retry without software intervention.

3.1.1.3.5 *PCI Express* Retrain*

If the hardware is unable to perform a successful recovery, then the link automatically reverts to the polling state, and initiates a full retraining sequence. This is a drastic event with an implicit reset to the downstream device and all subordinate devices, and is logged by the MCH as a "Link Down" error. If escalation of this event is enabled, software is notified of the link DL_DOWN condition. If software is involved, then data is likely lost, and processes need to be

restarted. This is preferred over the necessity of taking down the system, or going offline for an extended period of time.

3.1.1.3.6 *PCI Express* Hot-Plug Support*

Hot-plug support is available on the PCI Express* expansion slots (Slot 3, Slot 4) on the Intel® Server Board SE7520AF2 (SE7520HPAF2 SKU only). The PCI Express* hot-plug controller follows the requirements and recommendations from PCI Express* Base Specification, Rev 1.0.

Three hot-plug LEDs are associated with each PCI Express* slot:

- One power LED (green) on the top, adjacent to each PCI Express* slot
- One attention LED (amber) on the top, adjacent to each PCI Express* slot
- One attention LED (amber) on the bottom of the board, visible from the rear of the Server Chassis SC5300¹.

The hot-plug LEDs support three states: Off, On, or Blinking. The following table lists the description for each state of the power and attention LEDs.

Table 1. PCI Express* Hot-Plug LED Function Table

LED	Off	On (Solid)	Blinking
Power LED	All main voltage and data rails are removed from slot. PCI adapter may be added or removed	Normal power to slot. PCI adapter may NOT be added or removed	Slot power in transition (on-off or off-on). PCI adapter may NOT be added or removed
Attention LED	Normal operation	Slot on attention; power fault or operational problem present on the slot	Slot being identified at the user's request to locate the slot via the Hot-Plug Interface Driver

3.1.1.3.7 *PCI Express* Hot-Plug Controller*

The PCI Express* hot-plug solution is supported via the PCA9555, a 16-bit I2C I/O expander. It interfaces between the PCI Express* slots, hot-plug power controller, hot-plug indicators and the MCH. The PCA9555 facilitates the hot-plug events controlled by the MCH. This includes monitoring presence indicators from the slots, attention indicators, and fault and power good signals. It also drives slot enables and LED indicators during hot-plug events.

3.1.1.3.8 *PCI Express* Hot-Plug Power Controller*

The PCI Express* hot-plug power control is supported via the MIC2591 hot-plug power controller. The MIC2591 controller is a dual PCI Express* hot-plug power controller that provides power control support for 12V, 3.3V and 3.3V aux. This support includes current limiting, voltage supervision, fault indication and independent slot control. The MIC2591 also has a programmable circuit breaker that protects against excessive loads, such as short circuit, for each supply.

¹ Requires installation of light pipe included on the Hot Plug Accessory Kit

3.1.1.4 Hub Link 1.5 Interface

The MCH interfaces with the Intel® 82801ER I/O Controller Hub 5-R (ICH5-R) via a dedicated Hublink Interface supporting a peak bandwidth of 266MB/s using a x4 base clock of 66 MHz.

3.1.2 PCI-X Hub (PXH)

The PXH provides the data interface between the MCH and two PCI-X bus segments over a high-speed PCI Express* x8 link. Each PCI segment in the PXH is individually controlled to operate in PCI/PCI-X mode.

The PCI-X interfaces of the PXH are compliant with the *PCI-X Addendum to the PCI Local Bus Specification Revision 1.0b* and the Mode 1 sections of the *PCI-X Electrical and Mechanical Addendum to the PCI Local Bus Specification Revision 2.0a* and the *PCI-X Protocol Addendum to the PCI Local Bus Specification Revision 2.0a*. The PXH supports PCI bus frequencies of 66 MHz, 100 MHz, and 133 MHz.

On the Server Board SE7520AF2, the PXH is configured to support the following interfaces:

- PCI Segment A: One PCI-X slot (Slot 5 PCI-X 64/133) with maximum speed of 133MHz. Slot also enables PCI hot-plug support in the SE7520HPAF2 SKU.
- PCI Segment B
 - One PCI-X slot (Slot 6 PCI-X 64/100) with maximum speed of 100MHz with all loads present on the segment. Slot also enables optional support for third-party one or two slot riser.
 - Intel® 10/100/1000 82546GB Dual Gigabit Ethernet controller.

3.1.2.1 PCI-X Hot-Plug Option

Hot-plug support is available on the PCI-X expansion Slot 5 and Slot 1 from the IOP332 I/O processor, discussed in a later section, on the SE7520HPAF2 SKU. The PCI-X hot-plug controller follows the requirements and recommendations from *PCI Hot-Plug Specification, Rev 1.1* and the *PCI Standard Hot-plug Controller and Subsystem Specification, Rev 1.0*.

The PCI-X hot-plug control does not need quick switches because the PXH can control bus connection to the slot when only one device is on the bus. A HIP1011 PCI power controller provides power control support for +3.3V, +5V, +12V and -12V to the slots during power-on, power-down, and normal operation. The +3.3V_STBY is controlled by a FET circuitry.

Three hot-plug LEDs are associated with PCI-X Slot 1 and Slot 5:

- One power LED (green) on the top side adjacent to each PCI slot
- One attention LED (amber) on the top side adjacent to each PCI slot
- One attention LED (amber) on the bottom side of the board, visible from the rear of the system¹

The hot-plug LEDs support three states: Off, On, or Blinking. The following table lists the description for each state of the power and attention LEDs.

¹ Requires installation of light pipe included on the Hot Plug Accessory Kit

Table 2. PCI-X Hot-Plug LED Function Table

LED	Off	On (Solid)	Blinking
Power LED	All main voltage and data rails are removed from slot. PCI adapter may be added or removed	Normal power to slot. PCI adapter may NOT be added or removed	Slot power in transition (on-off or off-on). PCI adapter may NOT be added or removed
Attention LED	Normal operation	Slot on attention; power fault or operational problem present on the slot	Slot being identified at the user's request to locate the slot via the Hot-Plug Interface Driver

Note: PCI Slot 6 does not provide hot plug support. A power LED (green) is available adjacent to this slot and is always On to indicate the slot has power and that an adapter should not be hot-added or hot-removed.

3.1.3 IOP332 I/O Processor

The IOP332 I/O processor is a 500MHz multi-function device that integrates the Intel® XScale™ core (ARM* architecture compliant) with intelligent peripherals and dual PCI Express*-to-PCI bridges.

The IOP332 PCI Express* upstream link is capable of x8 lane width at 2.5GHz operation as defined by the PCI Express* Specification, Rev 1.0. This I/O processor integrates dual PCI Express*-to-PCI bridges with the Address Translation Unit (ATU) as an integrated secondary PCI device. The upstream PCI Express* port implements the PCI-to-PCI bridge programming model according to the *PCI Express Specification, Rev 1.0*. The primary ATU is compliant with the *Addendum to the PCI Local Bus Specification, Rev 1.0a* definitions of an application bridge. The IOP332 I/O processor is fully compliant with the *PCI Local Bus Specification, Rev 2.3* and the *PCI Express* Specification, Rev 1.0*.

The IOP332 I/O processor on the Intel® Server Board SE7520AF2 is enabled via a x4 PCI Express* port from the MCH. This I/O processor enables RAID On Mother Board (ROMB) support with the optional Intel® RAID Activation Key accessory and presence of RAID memory. For more details on the ROMB functionality see section 3.4.4 and Chapter 4.

In the absence of the Intel® RAID Activation Key, this I/O processor emulates the functionality of a PXH device. In the Intel® SE7520AF2 Server Board, the IOP332 is configured to support the following interfaces:

- IOP Segment A: LSI Logic* 53C1030 Dual Channel U320 SCSI controller enabling entry level RAID support (LSI Logic* Integrated Mirroring and Integrated Striping).
- IOP Segment B: One PCI-X slot (Slot 1 PCI-X 64/133) with maximum speed of 133MHz. Slot also enables PCI hot-plug support in the SE7520HPAF2 SKU.

3.1.4 I/O Controller Hub (ICH5-R)

The ICH5-R is a multi-function device providing an upstream hub interface for access to several embedded I/O functions and features including:

- *PCI Local Bus Specification, Revision 2.3* with support for 33 MHz PCI operations

- ACPI 2.0 power management logic support
- Enhanced DMA controller, interrupt controller, and timer functions
- Integrated IDE controller with support for Ultra ATA100/66/33
- Integrated SATA 150 controller
- USB host interface with support for eight USB 2.0 ports; via four UHCI host controllers; and one EHCI high-speed host controller
- System Management Bus (SMBus) Specification, Version 2.0 with additional support for I²C devices
- Low Pin Count (LPC) interface

Each function within the ICH5-R has its own set of configuration registers. Once configured, each appears to the system as a distinct hardware controller sharing the same PCI bus interface.

3.1.4.1 PCI Interface

The ICH5-R PCI interface supports a 33 MHz, Revision 2.3 compliant implementation. All PCI signals are 5-V tolerant, except PME#. The ICH5 integrates a PCI arbiter that supports up to six external PCI bus masters in addition to the internal ICH5 requests. On the Intel® Server Board SE7520AF2 this PCI interface is used to support one on-board PCI device: ATI* Rage XL video controller.

3.1.4.2 IDE Interface (Bus Master Capability and Synchronous DMA Mode)

The ICH5-R has an integrated IDE controller with two independent IDE signal channels supporting up to four IDE devices although only one is used on the Intel® Server Board SE7520AF2. This integrated functionality provides the interface for IDE hard disks and ATAPI devices. Each IDE device can have independent timings. The IDE interface supports PIO IDE transfers up to 16 MB/sec and Ultra ATA transfers up to 100 MB/sec. The IDE interface integrates 16x32-bit buffers for optimal transfers and does not consume any ISA DMA resources. The ICH5-R's IDE signal channels can be configured to the standard primary and secondary channels.

On the Intel® Server Board SE7520AF2 the primary bus is available via the one on-board legacy IDE connector.

3.1.4.3 Serial ATA (SATA) Host Controller

The SATA Host controller supports two SATA devices providing an interface for SATA hard disks and ATAPI devices. The SATA interface supports PIO IDE transfers up to 16 MB/s and Serial ATA transfers up to 1.5 Gb/s (150 MB/s). The SATA system for the ICH5-R contains two independent SATA signal ports that can be independently electrically isolated. Each SATA device can have independent timings. They can be configured to the standard primary and secondary channels. In addition, the ICH5-R offers the Intel® Embedded Server RAID Technology enabling data striping (RAID Level 0) for higher performance or data mirroring (RAID Level 1) for fault-tolerance between the two SATA drives alleviating disk bottlenecks by taking advantage of the dual independent SATA controllers integrated in the ICH5-R.

The Intel® Server Board SE7520AF2 supports both SATA 150 ports via two on-board connectors for internal hard drives. For details on SATA RAID see section 3.4.6.

3.1.4.4 Low Pin Count (LPC) Interface

The ICH5-R implements an LPC Interface as described in the Low Pin Count Interface Specification, Revision 1.1. The Low Pin Count (LPC) Bridge function of the ICH5-R resides in PCI Device 31: Function 0. In addition to the LPC bridge interface function, D31:F0 contains other functional units including DMA, interrupt controllers, timers, power management, system management, GPIO, and RTC.

On the Intel® Server Board SE7520AF2, the LPC bus is shared between the Super I/O SIO3 (National Semiconductor* PC87427), the mBMC and the IMM connector.

3.1.4.5 Compatibility Modules (DMA Controller, Timer/Counters, Interrupt Controller)

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

The ICH5-R supports two types of DMA: LPC and PC/PCI. LPC DMA and PC/PCI DMA use the ICH5-R's DMA controller. The PC/PCI protocol allows PCI-based peripherals to initiate DMA cycles by encoding requests and grants via two PC/PC REQ#/GNT# pairs. LPC DMA is handled through the use of the LDRQ# lines from peripherals and special encoding on LAD[3:0] from the host. Single, Demand, Verify, and Increment modes are supported on the LPC interface. Channels 0–3 are 8 bit channels. Channels 5–7 are 16-bit channels. Channel 4 is reserved as a generic bus master request.

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

The ICH5-R provides an ISA-compatible Programmable Interrupt Controller (PIC) that incorporates the functionality of two 82C59 interrupt controllers. The two interrupt controllers are cascaded so that 14 external and two internal interrupts are possible. In addition, the ICH5-R supports a serial interrupt scheme. All of the registers in these modules can be read and restored. This is required to save and restore system state after power has been removed and restored to the platform.

3.1.4.6 Advanced Programmable Interrupt Controller (APIC)

In addition to the standard ISA-compatible PIC described in the previous section, the ICH5-R incorporates the Advanced Programmable Interrupt Controller (APIC).

3.1.4.7 Universal Serial Bus (USB) Controller

The ICH5-R contains an Enhanced Host Controller Interface that supports 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 ICH5-R also contains four Universal Host Controller Interface (UHCI) controllers that support USB full-speed and low-speed signaling.

The Intel® Server Board SE7520AF2 makes use of five of the six USB 2.0 ports from the ICH5-R: three available in the rear I/O area and two available via a DH10 header on the board that

can be routed to the front of the server. All five ports are high-speed, full-speed, and low-speed capable.

3.1.4.8 Real-time Clock (RTC)

The ICH5-R contains a Motorola* MC146818A-compatible real-time clock with 256 bytes of battery backed RAM. The real-time clock performs two key functions: keeping track of the time of day and storing system data, even when the system is powered down. The RTC operates on a 32.768 KHz crystal and a separate 3 V lithium battery. The RTC also supports two lockable memory ranges. By setting bits in the configuration space, two 8-byte ranges can be locked to read and write accesses. This prevents unauthorized reading of passwords or other system security information.

3.1.4.9 General-purpose Input/Output (GPIO)

General-purpose inputs and outputs are provided for custom system design. The number of inputs and outputs depends on the ICH5-R configuration. All unused GPI pins must be pulled high or low, so that they are at a predefined level and do not cause undue side effects and must meet the following guidelines:

- GPIO 0:15 sticky bits on input, level triggered, 61µs min time for latch
- GPI's only: 0:15, 40-47 (note: 42-47 unimplemented)
- GPO's only: 16-23, 48-55 (note: 49-55 unimplemented)
- GPI or GPO: 24-39 (note: 35-39 unimplemented, GPIO[33] is changed to SATA LED and this GPIO is NOT available)
- GPIO resume (3.3VSTBY) power well: 8-15, 24-25, 27-28
- GPIO core (3.3V) power well: 0-7, 16-23, 32-34, 40-41, 48

Table 3. ICH5-R GPIO Assignment

ICH5-R Signal	Type	Pin	PWR Well	Tolerance	Description
INTRUDER DETECT	Input	Y12	Core	3.3V	TP
GPI0/REQA	Input	A5	Core	5V	Board SKU 0
GPI1/REQB/REQ5#	Input	E7	Core	5V	Board SKU 1
GPI6/AGPBUSY#	Input	R5	Core	5V	Bios Recover Boot
GPI7	Input	U3	Core	5V	MCH PME
GPI8	Input	Y2	Resume	3.3V	WAKE#/PCI PME#
GPI9/OC[4]#	Input	B14	Resume	3.3V	Reserved for USB OC4 (3 in back, 2 in front)
GPI10/OC[5]#	Input	A14	Resume	3.3V	MCH GPE
GPI11/SMBALERT#	Input	AC3	Resume	3.3V	Board ID 0
GPI12	Input	W4	Resume	3.3V	SIO SMI
GPI13	Input	W5	Resume	3.3V	BMC IRQ SMI
GPI14/OC[6]#	Input	D13	Resume	3.3V	SIO > ICH5R PME
GPI15/OC[7]#	Input	C13	Resume	3.3V	Password clear
GPO16/GNTA#	Output	E8	Core	3.3V	TP
GPO17/GNTB#/GNT[5]#	Output	B4	Core	3.3V	TP

ICH5-R Signal	Type	Pin	PWR Well	Tolerance	Description
GPO18/STP_PCI#	Output	U21	Core	3.3V	TP
GPO19/SLP_S1#	Output	T20	Core	3.3V	TP
GPO20/STP_CPU#	Output	U22	Core	3.3V	Video Disable
GPO21/C3_STAT#	Output	R1	Core	3.3V	SCSI Disable
GPO22/CPUPERF#	Output	U20	Core	3.3V	TP
GPO23/SSMUXSEL#	Output	F22	Core	3.3V	OEM RMC Connector POST complete Indicator
GPIO24/CLKRUN#	I/O	AC1	Resume	3.3V	Board ID 1
GPIO25	I/O	W3	Resume	3.3V	Board ID 2
GPIO27	I/O	V3	Resume	3.3V	Source Clock Enable 0 (HP PCI-E)
GPIO28	I/O	W2	Resume	3.3V	Source Clock Enable 1 (HP PCI-E)
GPIO32	I/O	T1	Core	3.3V	TP
GPIO34	I/O	F21	Core	3.3V	IDE Primary Cable Sense
GPI40/REQ4#	Input	C6	Core	3.3V	CMOS Clear
GPI41/LDRQ1#	Input	R2	Core	3.3V	Emergency Bank Select
GPO48/GNT4#	Output	A4	Core	3.3V	FRB Timer Halt

3.1.4.10 System Management Bus (SMBus 2.0)

The ICH5-R contains an SMBus host interface that allows the processor to communicate with SMBus slaves. This interface is compatible with most I²C devices and special I²C commands are implemented. The SMBus host controller for the ICH5-R provides a mechanism for the processor to initiate communications with SMBus peripherals (slaves).

The ICH5-R supports slave functionality, including the Host Notify protocol. Hence, the host controller supports eight command protocols of the SMBus interface: Quick Command, Send Byte, Receive Byte, Write Byte/Word, Read Byte/Word, Process Call, Block Read/Write, and Host Notify. See the System Management Bus (SMBus) Specification, Version 2.0.

3.2 Processor Sub-system

The Intel® Server Board SE7520AF2 supports either one or two Intel® Xeon™ processors with 1MB L2 cache or Intel® Xeon™ processors with 2MB L2 cache with frequencies starting at 2.8 GHz using 90 nanometer technology and utilizing an 800MHz front side bus. Previous generations of the Intel Xeon processor are not supported. When two processors are installed, both must be of identical revision, core voltage, cache size, and bus/core speed. When only one processor is installed, it must be populated into the socket labeled “CPU_1” with the other CPU socket remains unpopulated. The support circuitry for the processor sub-system consists of the following:

- Dual 604-pin zero insertion force (ZIF) processor sockets
- Processor host bus AGTL+ support circuitry

- Reset configuration logic
- Processor module presence detection logic
- BSEL detection capabilities
- CPU signal level translation
- CEK CPU retention support

Table 4. Processor Support Matrix

Processor Family	Package Type	Technology	Frequency	Cache Size	Support
Intel® Xeon™	FC-mPGA4	90 nM	2.8 GHz	1024KB	Yes
Intel® Xeon™	FC-mPGA4	90 nM	3.0 GHz	1024KB	Yes
Intel® Xeon™	FC-mPGA4	90 nM	3.2 GHz	1024KB	Yes
Intel® Xeon™	FC-mPGA4	90 nM	3.4 GHz	1024KB	Yes
Intel® Xeon™	FC-mPGA4	90 nM	3.6 GHz	1024KB	Yes
Intel® Xeon™	FC-mPGA4	90 nM	3.8 GHz	1024KB	Yes

The Intel® Server Board SE7520AF2 is designed to provide up to 120A per processors. Processors with higher current requirements are not supported.

3.2.1 Processor VRD

The baseboard has two VRDs (Voltage Regulator Down) providing the appropriate voltages to the installed processors. Each VRD is compliant with the VRD 10.1 specification and are designed to support current and next generation Intel Xeon processors that require up to a sustained maximum current of 105 Amps and peak support of 120 Amps.

The baseboard supports the Flexible Mother Board (FMB) specification for all “Nocona”/“Irwindale” processors with respect to current requirements and processor speed requirements. FMB is an estimation of the maximum values the “Nocona”/“Irwindale” versions of the Intel® Xeon™ processors will have over their lifetime. The value is an estimate; specifications for future processors may differ. At present the current demand per FMB is a sustained maximum of a 105 Amps and peak support of 120 Amps.

3.2.2 Reset Configuration Logic

The BIOS determines the processor stepping, cache size, etc through the CPUID instruction. All processors in the system must operate at the same frequency; have the same cache sizes; and same VID (voltage identification). No mixing of product families is supported. Processors run at a fixed speed and cannot be programmed to operate at a lower or higher speed.

3.2.3 Processor Module Presence Detection

Logic is provided on the baseboard to detect the presence and identity of installed processors. In dual processor configurations, the on-board mini Baseboard Management Controller (mBMC) must read the processor VID bits for each processor before turning on the VRD. If the VIDs of the two processors are not identical, then the mBMC will not turn on the VRD. Before enabling the embedded VRD, circuitry on the baseboard ensures that the following criteria are met:

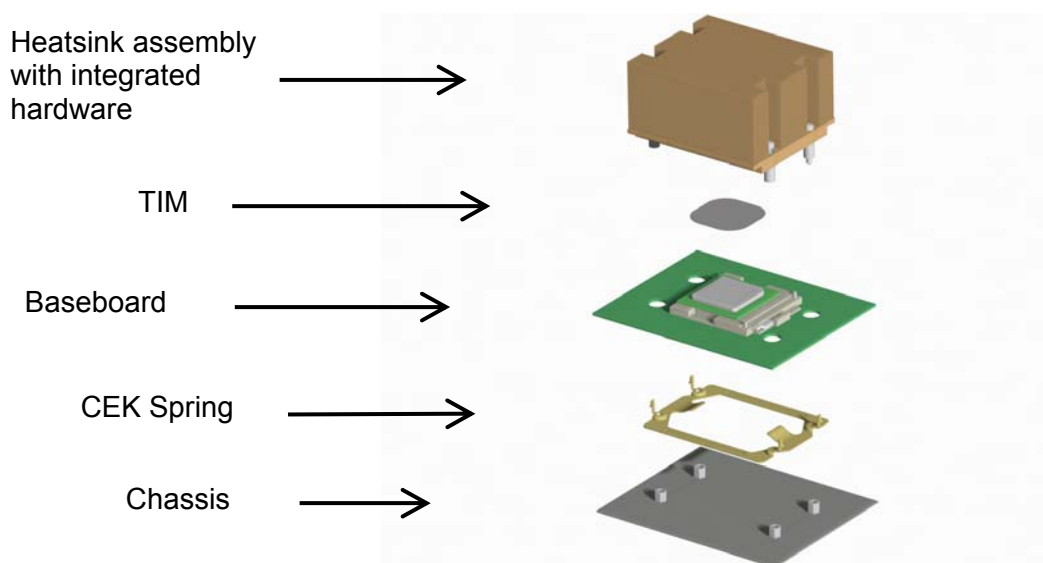
- In a uni-processor configuration, the processor is installed in the CPU1 socket
- Only supported processors are installed to prevent damage to the MCH
- In dual processor configurations, both processors support the same FSB frequency

3.2.4 GTL2006

The GTL2006 is a 13-bit translator designed for 3.3V to GTL/GTL+ translations to the system bus. The translator incorporates all the level shifting and logic functions required to interface between the processor subsystem and the rest of the system.

3.2.5 Common Enabling Kit (CEK) Design Support

The baseboard has been designed to comply with the Intel® Common Enabling Kit (CEK) processor mounting and thermal solution. The baseboard ships from Intel's factory with a CEK spring snapped onto the underside of the board, beneath each processor socket. The CEK spring is removable to allow the use of non-Intel heat sink retention solutions.



Note: When installing the server board into an Intel® Server Chassis SC5300, the passive heatsink solution (no fan) must be used.

Figure 5. CEK 'Passive' Component Stackup

3.3 Memory Sub-System

The Intel® Server Board SE7520AF2 memory sub-system is designed to support Double Data Rate2 (DDR2) Synchronous Dynamic Random Access Memory (SDRAM). The MCH provides two independent DDR channels (A and B), which support 400 MT/s with DDR2 SDRAM DIMMs. The peak bandwidth of each DDR branch channel is 3.2 GB/s (64 bits x 400MT/s). The effective overall peak bandwidth of the DDR memory subsystem is 6.4GB/s.

In the non-mirrored operation, the two DDR channels from the MCH operate in lock step. In memory mirroring operation, the channels operate in lock-step under normal conditions, but independently under failure and recovery conditions.

The DDR data bus is a source synchronous bus where the data signals are latched with the differential strobe signals. Since the strobes are routed with their appropriate data group, the flight times will be virtually the same between strobe and data.

3.3.1 Supported Memory

The Intel® Server Board SE7520AF2 supports four single-ranked registered DDR2-400 DIMMs per channel, which provides a total of eight DIMMs. The maximum memory capacity supported is 16 GB using eight modules of 1 Gbit DRAM technology devices.

Table 5. SE7520AF2 Supported DIMM Modules

Technology	Organization	DRAM Components / DIMM	Row / Column Address Bits
128Mb	4M X 8 X 4bks	8	12/10
	8M X 4 X 4bks	16	12/11
256Mb	8M X 8 X 4bks	8	13/10
	16M X 4 X 4bks	16	13/11
512Mb	16M X 8 X 4bks	8	14/10
	32M X 4 X 4bks	16	14/11
1Gb	32M X 8 X 8bks	8	14/10
	64M X 4 X 8bks	16	14/11

3.3.2 DIMM Population Rules and Supported Configurations

The server board contains four banks of DIMMs: memory Bank 1 through memory Bank 4. Memory Bank 1, with DIMM sockets 1B and 1A, is located closest to the edge of the board. Moving in from the edge on the board are sockets for banks 2, 3, and 4 in order. Sockets associated with any given bank are located next to each other.

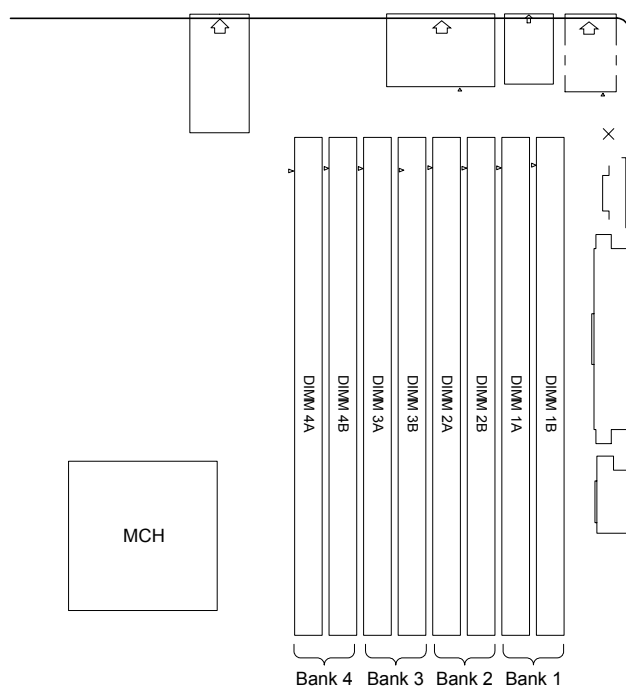


Figure 6. DIMM Socket Configuration

DIMM and memory configurations must adhere to the following:

- DDR2-400 MHz registered ECC DIMM modules
- DIMM organization: x72 ECC
- Pin count: 240 pins
- DIMM capacity: 256 MB, 512 MB, 1 GB, 2 GB, 4 GB DIMMs using current SDRAM densities
- Support up to 16 GB of system memory capacity
- Support a maximum of four ranks per channel¹
- Serial PD: JEDEC Rev 2.0
- Voltage options: 1.8 V (VDD/VDDQ)
- Interface: SSTL2
- Two DIMMs must be populated in a bank for a x144 wide memory data path
- DIMMs must be populated beginning with the socket furthest away from the MCH (i.e. DIMM 1B and DIMM 1A)
- DIMMs must be identical within the same bank. For example, the DIMM in socket DIMM 1A must be identical to the DIMM in socket DIMM 1B)

¹ Note that in order to populate 8 DIMM modules on the board, all DIMMs must be single rank. Refer to the DIMM manufacturer specification to determine the number of ranks on the specific DIMM intended to be used.

When mixing single-ranked and dual-ranked DIMMs, dual-ranked DIMMs must be in lower bank (i.e. Bank 1), followed by single-ranked DIMMs in subsequent banks (i.e. Bank 2). Refer to Table 36. DIMM Population on section 5.4.5 Memory Population.

3.3.3 Single Channel Operation

The Intel® Server Board SE7520AF2 supports single-channel DIMM operation in which only one DIMM is present and is installed in DIMM Bank 1 (DIMM socket 1B or DIMM socket 1A). Population in other DIMM banks is not supported for single-channel operation. The minimal size of memory considered for testing is 256 MB.

Note: All Intel memory qualification is done by testing with complete memory banks of identical memory modules in all DIMM sockets. Memory qualification does not include testing of single channel memory mode, mixed DIMM type and/or vendors.

3.3.4 I²C Bus

To boot the system, the system BIOS on the Intel® Server Board SE7520AF2 uses a dedicated I²C bus to retrieve DIMM information needed to program the MCH memory registers. The following table provides the I²C addresses for each DIMM slot.

Table 6. I²C Addresses for Memory Module SMB

Device	Address
DIMM 1A	0xA6
DIMM 1B	0xAE
DIMM 2A	0xA4
DIMM 2B	0xAC
DIMM 3A	0xA2
DIMM 3B	0xAA
DIMM 4A	0xA0
DIMM 4B	0xA8

3.3.5 Memory Voltage Margining

The supply voltage for memory devices is 1.8V. The Intel® Server Board SE7520AF2 utilizes a HIP6311, which is a multiphase Pulse Width Modulator (PWM) controller, to set this output voltage. This device has Voltage Identification Inputs (VID) on five pins (VID0-VID4.) These VID bits respond to 3.3V TTL logic signals. The HIP6311 decodes these VID bits to determine the output voltage.

3.3.6 Memory RASUM Features

The Intel® E7520 MCH supports several memory RASUM (Reliability, Availability, Serviceability, Usability, and Manageability) features that are traditionally found only on high end server systems. These features include the Intel® x4 Single Device Data Correction (x4 SDDC) for memory error detection and correction, memory Scrubbing, Retry on Correctable Errors, Integrated memory Initialization, DIMM Sparing, and memory Mirroring. The following sections describe how each is supported on the Intel® Server Board SE7520AF2.

Note: Event logging and annunciation is not supported for all RASUM events with the onboard platform instrumentation. Event logging and annunciation for these advanced RASUM memory features requires the presence of an Intel® Management Module – Professional Edition or Advanced Edition.

3.3.6.1 DRAM ECC – Intel® x4 Single Device Data Correction (x4 SDDC)

The DRAM interface uses two different ECC algorithms. The first is a standard SEC/DED ECC across a 64-bit data quantity. The second ECC method is a distributed, 144-bit S4EC-D4ED mechanism, which provides x4 SDDC protection for DIMMS that utilize x4 devices. Bits from x4 parts are presented in an interleaved fashion such that each bit from a particular part is represented in a different ECC word. DIMMs that use x8 devices, can use the same algorithm but do not have x4 SDDC protection, since at most only four bits can be corrected with this method. The algorithm does provide enhanced protection for the x8 parts over a standard SEC-DED implementation. With two memory channels, either ECC method can be utilized with equal performance, although single-channel mode only supports standard SEC/DED.

When memory mirroring is enabled, x4 SDDC ECC is supported in single channel mode when the second channel has been disabled during a fail-down phase. The x4 SDDC ECC is not supported during single-channel operation outside of DIMM mirroring fail-down as it does have significant performance impacts in that environment.

3.3.6.2 Integrated Memory Scrub Engine

The Intel® E7520 MCH includes an integrated engine to walk the populated memory space proactively seeking out soft errors in the memory subsystem. In the case of a single bit correctable error, this hardware detects, logs, and corrects the data except when an incoming write to the same memory address is detected.

For an uncorrectable error, the scrub engine logs the failure. Both types of errors may be reported via multiple alternate mechanisms under configuration control. The scrub hardware executes “demand scrub” writes when correctable errors are encountered during normal operation (on demand reads, rather than scrub-initiated reads). This functionality provides incremental protection against time-based deterioration of soft memory errors from correctable to uncorrectable.

Using this method, a 16 GB system can be completely scrubbed in less than one day. The effect of these scrub-writes does not cause any noticeable degradation to memory bandwidth, although they cause a greater latency if a read that is delayed due to the scrub write cycle. This is a very infrequent occurrence.

Note: An uncorrectable error encountered by the memory scrub engine is a “speculative error.” This designation is applied because no system agent has specifically requested use of the corrupt data, and no real error condition exists in the system until that occurs. It is possible that the error resides in an unmodified page of memory that will be dropped on a swap-back to disk. If dropped the speculative error vanishes from the system, undetected and without adverse consequences.

3.3.6.3 Retry on Uncorrectable Error

The Intel® E7520 MCH includes specialized hardware to resubmit a memory read request upon detection of an uncorrectable error. When a demand fetch (as opposed to a scrub) of memory encounters an uncorrectable error as determined by the enabled ECC algorithm, the memory control hardware causes a (single) full resubmission of the entire cache line request from memory to verify the existence of corrupt data. This feature is expected to greatly reduce or eliminate the reporting of false or transient uncorrectable errors in the DRAM array.

Note: Any given read request is retried only once on behalf of this error detection mechanism. If the uncorrectable error is repeated, it is logged and escalated as directed by device configuration. In the memory mirror mode, the retry on an uncorrectable error is issued to the mirror copy of the target data, rather than back to the devices responsible for the initial error detection. This has the added benefit of making uncorrectable errors in DRAM fully correctable unless the same location in both primary and mirror happens to be corrupt. This RASUM feature may be enabled and disabled via configuration.

3.3.6.4 Integrated Memory Initialization Engine

The Intel® E7520 MCH provides hardware-managed ECC auto-initialization of all populated DRAM space under software control. After the internal configuration has been updated to reflect the types and sizes of populated DIMM devices, the MCH traverses the populated address space initializing all locations with good ECC. This speeds up the mandatory memory initialization step and frees the processor to pursue other machine initialization and configuration tasks.

Additional features have been added to the initialization engine to support high-speed population and verification of a programmable memory range with one of four known data patterns (0/F, A/5, 3/C, and 6/9). This function facilitates a limited, very high-speed memory test and provides a BIOS-accessible memory zeroing capability for use by the operating system.

3.3.6.5 DIMM Sparing Function

To improve fault-tolerance, the Intel® E7520 MCH includes specialized hardware to support fail-over to a spare DIMM device in case a primary DIMM exceeds a specified threshold of runtime errors. One of the DIMMs installed per channel, greater than or equal in size than all installed, is not used but is instead kept in reserve. If a primary DIMM experiences significant failures, the failing DIMM and its corresponding partner in the other channel (if applicable), will, over time have its data copied over to the spare DIMM(s) held in reserve. When the data has all been copied, the reserve DIMM(s) is put into service and the failing DIMM is removed from service. Only one sparing cycle is supported. If this feature is not enabled, then all DIMMs are visible in normal address space.

Note: DIMM Sparing requires that the spare DIMM be at least the size of the largest primary DIMM in use.

Hardware additions for this feature include the implementation of tracking register per DIMM to maintain a history of error occurrence, and a programmable register to hold the fail-over error threshold level. The operational model is as follows: if the fail-over threshold register is set to a non-zero value, the feature is enabled. If the count of errors on any DIMM exceeds the register value, fail-over begins.

The tracking registers are implemented as “leaky buckets,” such that they do not contain an absolute cumulative count of all errors since power-on. Instead, they contain an aggregate count of the number of errors received over a running time period. The “drip rate” of the bucket is selectable by software. It is possible to set the threshold to a value that will never be reached by a “healthy” memory subsystem experiencing the rate of errors expected for the size and type of memory devices in use.

The fail-over mechanism is slightly more complex. Once fail-over has been initiated, the MCH must execute each write twice; once to the primary DIMM, and once to the spare. The MCH begins tracking the progress of its built-in memory scrub engine. Once the scrub engine has covered every location in the primary DIMM, the duplicate write function will have copied every data location to the spare. At that point, the MCH can switch the spare into primary use, and take the failing DIMM off-line.

Until the threshold detection has been triggered to request a data copy, this mechanism requires no software support once it has been programmed and enabled. Hardware detects the threshold initiating the fail-over and escalates the occurrence of the event as directed (signals an SMI, generates an interrupt, or waits to be discovered via polling).

The chosen software routine responds to the threshold detection. Software must select a victim DIMM if multiple DIMMs have crossed the threshold prior to sparing invocation and software must initiate the memory copy.

Hardware automatically isolates the “failed” DIMM after the copy has completed. The data copy is accomplished by address aliasing within the DDR control interface, thus it does not require reprogramming of the DRAM row boundary (DRB) registers, nor does it require notification to the operating system that anything untoward has occurred in memory.

The memory mirroring feature and DIMM sparing are exclusive of each other; only one may be activated during initialization. The feature can be enabled in system BIOS setup; the selected feature must remain enabled until the next power-cycle. Hardware does not include a provision to switch from one feature to the other without rebooting and there is no provision to back out of an enabled feature without a full reboot.

3.3.6.6 Memory Mirroring

The memory mirroring feature provides a way for hardware to maintain two copies of all data in the memory subsystem, such that a hardware failure or uncorrectable error is no longer fatal to the system. When an uncorrectable error is encountered during normal operation, hardware retrieves the mirrored copy of the corrupted data. A system failure does not occur unless both primary and mirrored copies of the data are corrupt at the same time.

Mirroring is supported on dual-channel DIMM populations symmetric both across channels and within each channel. As a result, the Intel® Server Board SE7520AF2 provides three supported configurations for memory mirroring:

- A four-DIMM population uses two identical devices per channel. DIMMs labeled 1A, 2A, 1B, and 2B must all be identical. Refer to the figure below.

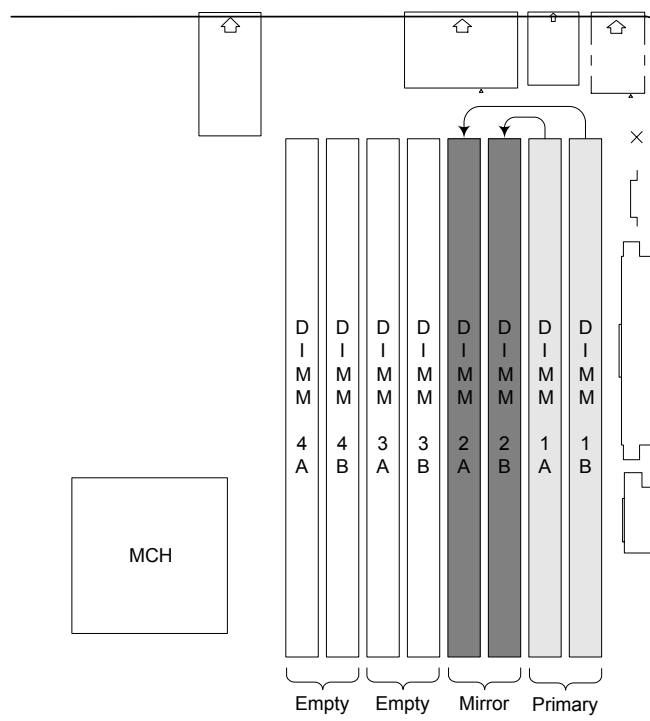


Figure 7. Four DIMM Memory Mirror Configuration

- A six-DIMM population uses identical devices within DIMM Bank 1, and between DIMM Banks 2 and 3. Referring to Figure 8, DIMMs labeled 3A, 3B, 2A, and 2B must be single-rank and identical; and those labeled 1A and 1B must be dual-ranked and identical. The first group does not have to be identical to those in the second group.

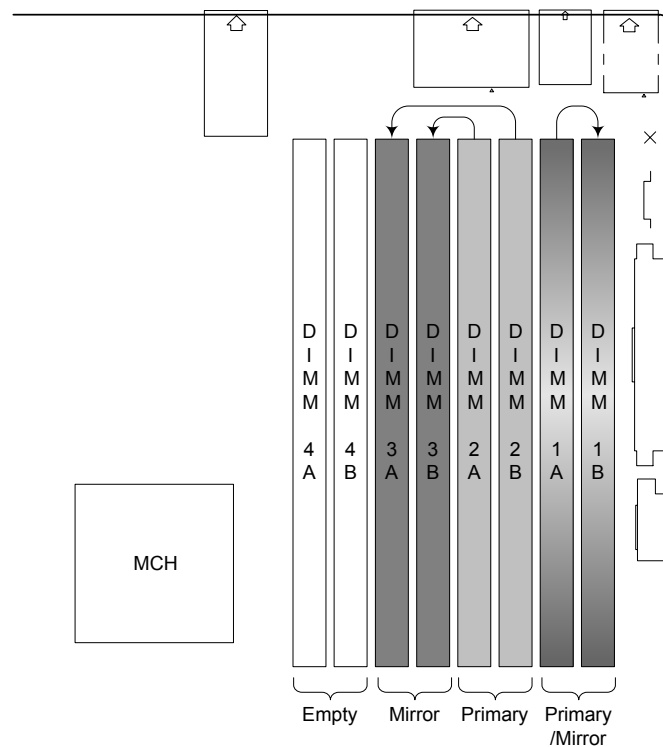


Figure 8. Six DIMM Memory Mirror Configuration

- An eight-DIMM population uses identical devices in DIMM Banks 1 and 2, and in DIMM Banks 3 and 4. Referring to Figure 9, DIMMs labeled 1A, 1B, 2A, and 2B must be identical and those labeled 3A, 3B, 4A, and 4B must be identical. The first group does not need to be identical to those in the second group.

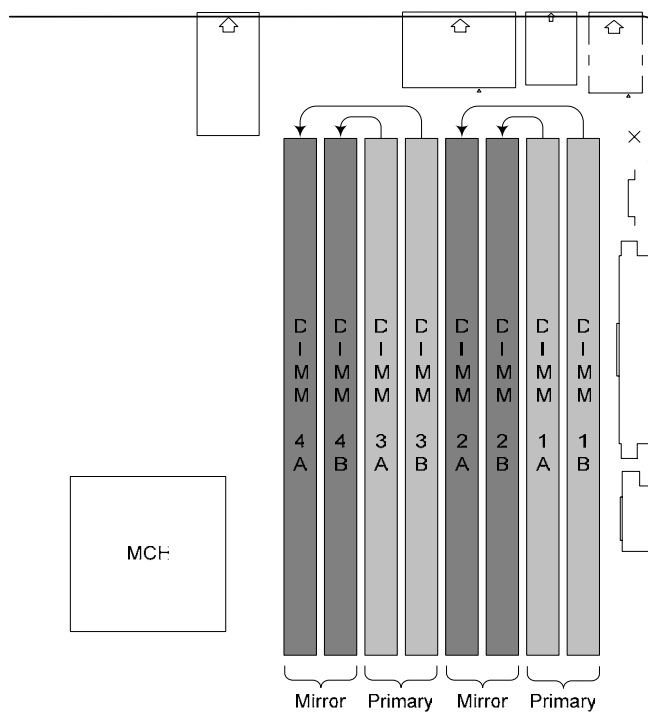


Figure 9. Eight DIMM Memory Mirror Configuration

The symmetry requirements are a result of the hardware mechanism for maintaining two copies of all main memory data while ensuring that each channel has a full copy of all data in preparation for fail-down to single-channel operation. Every write to memory is issued twice, once to the “primary” location, and again to the “mirror” location. The data interleaved across the channel pair are swapped for the second write (1A is a copy of 2B, 1B is a copy of 2A etc.). The resulting memory image has two full copies of all data, and a complete copy is available on each channel.

Hardware in the MCH tracks which DIMM slots are primaries, and which are mirrors, such that data may be internally realigned to correctly reassemble cache lines regardless of which copy is retrieved. There are four distinct cases for retrieval of the “even” and “odd” chunks of a cache-line of data:

- Interleaved dual-channel read to the primary DIMM with “even” data on channel A
- Interleaved dual-channel read to the mirror DIMM with “even” data on channel B
- Non-interleaved single-channel read pair to channel A with “even” data on the primary DIMM
- Non-interleaved single-channel read pair to channel B with “even” data on the mirror DIMM

When mirroring is enabled via the MCH configuration, the memory subsystem maintains two copies of all data, as described above, and retrieves requested data from either the primary DIMM or the mirrored DIMM, based on the state of system address bit 15 (SA[15]).

Software toggles the SA[15] polarity via a configuration register bit setting. SA[15] was chosen because it is the lowest system address bit that is always used to select the memory row address across all DRAM densities and technologies supported by the Intel® E7520 MCH. The toggling of the primary read location based on an address bit distributes request traffic across the primary and mirror DIMMs, thereby distributing the thermal image of the workload across all populated DIMM slots and reducing the chance of thermal-based memory traffic throttling.

In the mirrored operating state, the occurrence of correctable and uncorrectable ECC errors are tracked and logged normally by the MCH and escalated to system interrupt events as specified by the configuration register settings associated with errors on the memory subsystem. Counters implementing the “leaky bucket” function for on-line DIMM sparing track the aggregate count of single-bit and multiple-bit errors on a per DIMM basis.

The memory mirroring feature and DIMM sparing are exclusive of each other, only one may be activated during initialization. The feature can be enabled in system BIOS setup. The selected feature must remain enabled until the next power-cycle.

3.4 I/O Sub-System

The I/O sub-system is comprised of several components:

- The MCH provides the PCI Express* interface to PCI Express* slots 3 and 4
- The PXH provides the PCI-X interface for the dual gigabit Ethernet controller and PCI-X slots 5 and 6
- The IOP332 provides the PCI Express* to PCI bridge to interface the on-board SCSI controller and PCI-X Slot 1
- The ICH5-R provides the interface for the onboard video controller, Super IO chip, and Management Sub-system

This section describes the function of each I/O interface and how they operate on the Intel® Server Board SE7520AF2.

3.4.1 PCI Subsystem

The primary I/O interface for the Intel® Server Board SE7520AF2 is PCI, with six independent PCI bus segments:

- A PCI 32-bit 33 MHz bus segment (P32-A) enabled through the ICH5-R.
- Two PCI-X 64-bit bus segments (P64-A 133MHz and P64-B 100MHz) enabled through the PXH PCI Bridge.
- Two PCI-X 64-bit bus segments (P64-A 133MHz and P64-B 133MHz) enabled through the IOP332 I/O processor.
- Two PCI Express* bus segments (PEXP-B x8 and PEXP-C x4) enabled through the MCH.

The table below lists the characteristics of the PCI bus segments.

Table 7. PCI Bus Segment Characteristics

PCI Bus Segment	Voltage	Width	Speed	Type	PCI I/O Card Slots
-----------------	---------	-------	-------	------	--------------------

ICH5-R P32-A	5 V	32-bits	33 MHz	PCI	ATI* Rage XL
PXH P64-A	3.3 V	64-bits	133 MHz	PCI-X	Slot 5 (PCI-X 64/133)
PXH P64-B	3.3 V	64-bits	100 MHz	PCI-X	Slot 6 (PCI-X 64/100) Intel® 82546GB Dual Gb NIC
IOP332 P64-A	3.3 V	64-bits	133 MHz	PCI-X	LSI Logic* 53C1030 U320 SCSI controller
IOP332 P64-B	3.3 V	64-bits	133 MHz	PCI-X	Slot 1 (PCI-X 64/133)
MCH PExp-B	Diff.	Serial	2.5GHz	PCI Exp	Slot 4 (PCI Exp x8)
MCH PExp-C	Diff.	Serial	2.5GHz	PCI Exp	Slot 3 (PCI Exp x4)

3.4.1.1 ICH5-R P32-A: 32-bit, 33-MHz PCI Subsystem

The 32-bit, 33-MHz PCI segment enabled by the ICH5-R is known as the ICH5-R P32-A segment. This segment supports the on-board 2D/3D Graphics ATI* Rage XL video controller.

3.4.1.1.1 Device IDs (IDSEL)

The device under the PCI hub bridge has its IDSEL signal connected to one bit of AD[31:16], which acts as a chip select on the PCI bus segment in configuration cycles. This determines a unique PCI device ID value for use in configuration cycles. The following table shows the bit to which each IDSEL signal is attached for the 32-bit on-board device.

Table 8. ICH5-R P32-A Configuration IDs

IDSEL Value	Device
28	ATI* Rage XL video controller

3.4.1.1.2 ICH5-R P32-A Arbitration

The ICH5-R P32-A segment only supports the ATI Rage XL video controller which must arbitrate for PCI access, using resources supplied by the ICH5-R. The host bridge PCI interface (ICH5-R) arbitration lines REQx* and GNTx* are a special case in that they are internal to the host bridge. The following table defines the arbitration connections.

Table 9. ICH5-R P32-A Arbitration Connections

Baseboard Signals	Device
PCI_REQ0*/PCI_GNT0*	ATI* Rage XL video controller

3.4.1.2 PXH P64-A and P64-B: 64-bit, 100/133-MHz PCI Subsystem

The two peer 64-bit PCI-X bus segments are directed through the PXH PCI Bridge. The first PCI-X segment, P64-A, supports the interface for the PCI expansion Slot 5 (PCI-X 64/133).

The second PCI-X segment, P64-B, supports the interface to the on-board Intel® 82546GB 10/100/1000 Dual Gigabit Ethernet controller and the PCI expansion Slot 6 (PCI-X 64/100) or Slot 6 and Slot 7 for a two-slot riser.

3.4.1.2.1 *PCI Riser Support*

PCI Slot 6 enables support for third-party add-in riser cards. The slot can be populated with three types of devices:

- Standard PCI-X 1.0 compatible riser card
- 1U/1-slot PCI-X riser card
- 2U/2-slot PCI-X riser card

Note: Intel does not provide a PCI riser card for this server platform. However, in this document, Intel has provided information to assist developers with their PCI riser design.

3.4.1.2.2 *Device IDs (IDSEL)*

Each device under the PCI hub bridge has its IDSEL signal connected to one bit of AD[31:16], which acts as a chip select on the PCI bus segment in configuration cycles. This determines a unique PCI device ID value for use in configuration cycles. The following table shows the bit to which each IDSEL signal is attached for PXH P64-A/P64-B devices and corresponding device description.

Table 10. PXH P64-A Configuration IDs

IDSEL Value	Device
17	Slot 5 (PCI-X 64/133)

Table 11. PXH P64-B Configuration IDs

IDSEL Value	Device
17	Slot 6 (PCI-X 64/100)
19	Slot 6 (Upper slot of optional 2-slot riser)
18	Slot 7 (Lower slot of optional 2-slot riser)
20	Intel® 10/100/1000 82546GB Dual Gigabit Ethernet controller

3.4.1.2.3 *PXH P64-A Arbitration*

The PXH P64-A bus segment supports only the PCI expansion Slot 5 (PCI-X 64/133), which must arbitrate for PCI access using resources supplied by the PXH. The host bridge PCI interface arbitration lines PAREQx* and PAGNTx* are a special case in that they are internal to the host bridge. The following table defines the arbitration connections.

Table 12. PXH P64-A Arbitration Connections

Baseboard Signals	Device
PXH_PAREQ0*/PXH_PAGNT0*	Slot 5 (PCI-X 64/133)

3.4.1.2.4 *PXH P64-B Arbitration*

The PXH P64-B bus segment supports the on-board Intel® 82546GB 10/100/1000 Dual Gigabit Ethernet controller and the PCI expansion Slot 6 (PCI-X 64/100), or Slot 6 and Slot 7 in the event of a two-slot riser. All PCI masters must arbitrate for PCI access, using resources supplied by the PXH. The host bridge PCI interface arbitration lines PBREQx* and PBGNTx* are a special case in that they are internal to the host bridge. The following table defines the arbitration connections.

Table 13. PHX P64-B Arbitration Connections

Baseboard Signals	Device
PXH_PBREQ0*/PXH_PBGNT0*	Intel® 10/100/1000 82546GB Dual Gigabit Ethernet controller
PXH_PBREQ1*/PXH_PBGNT1*	Slot 6 (PCI-X 64/100)
PXH_PBREQ1*/PXH_PBGNT1*	Slot 6 (Upper slot of optional 2-slot riser)
PXH_PBREQ2*/PXH_PBGNT2*	Slot 7 (Lower slot of optional 2-slot riser)

3.4.1.3 **IOP332 P64-A and P64-B: 64-bit, 133-MHz PCI Subsystem**

The two peer 64-bit PCI-X bus segments are directed through the IOP332 I/O bridge. The first PCI-X segment, P64-A, supports the interface to the on-board LSI Logic* 53C1030 U320 Dual Channel SCSI controller. The second PCI-X segment, P64-B, supports the interface for the PCI expansion Slot 1 (PCI-X 64/133).

3.4.1.3.1 *Device IDs (IDSEL)*

Each device under the IOP332 has its IDSEL signal connected to one bit of AD[31:16], which acts as a chip select on the PCI bus segment in configuration cycles. This determines a unique PCI device ID value for use in configuration cycles. The following tables show the bit to which each IDSEL signal is attached for IOP332 P64-A/P64-B devices and corresponding device description.

Table 14. IOP332 P64-A Configuration IDs

IDSEL Value	Device
21	LSI Logic* 53C1030 U320 SCSI controller

Table 15. IOP332 P64-B Configuration IDs

IDSEL Value	Device
17	Slot 1 (PCI-X 64/133)

3.4.1.3.2 IOP332 P64-A Arbitration

The IOP332 P64-A bus segment supports the on-board LSI Logic* 53C1030 U320 Dual Channel SCSI controller which must arbitrate for PCI access using resources supplied by the IOP332. The I/O bridge PCI interface arbitration lines PAREQx* and PAGNTx* are a special case in that they are internal to the host bridge. The following table defines the arbitration connections.

Table 16. IOP332 P64-B Arbitration Connections

Baseboard Signals	Device
IOP_PAREQ0*/IOP_PAGNT0*	LSI Logic* 53C1030 U320 SCSI controller

3.4.1.3.3 IOP332 P64-B Arbitration

The IOP332 P64-B bus segment only supports the PCI expansion Slot 1 (PCI-X 64/133) which must arbitrate for PCI access using resources supplied by the IOP332. The I/O bridge PCI interface arbitration lines PBREQx* and PBGNTx* are a special case in that they are internal to the host bridge. The following table defines the arbitration connections.

Table 17. IOP332 P64-B Arbitration Connections

Baseboard Signals	Device
IOP_PBREQ0*/IOP_PBGNT0*	Slot 1 (PCI-X 64/133)

3.4.1.4 MCH PExp-B and PExp-C: x8/x4 PCI Express* PCI Subsystem

The PCI Express* PExp-B and PExp-C segments are enabled from the MCH. The first segment, PExp-B, supports a x8 interface for the PCI Express* expansion Slot 4 (PCI EXP x8). The second segment, PExp-C, supports a x4 interface for the PCI Express* expansion Slot 3 (PCI EXP x4). Both expansion slots are enabled with x8 physical connectors and offer hot-plug support on the hot-plug Server Board SE7520HPAF2 SKU.

Given the serial point-to-point nature of the PCI Express* architecture, arbitration does not apply; for details on the MCH PCI Express* implementation, refer to Section 3.1.1.3.

3.4.2 Interrupt Routing

The Intel® Server Board SE7520AF2 interrupt architecture accommodates both PC-compatible PIC mode and APIC mode interrupts through use of the integrated I/O APICs in the ICH5-R.

3.4.2.1 Legacy Interrupt Routing

For PC-compatible mode, the ICH5-R provides two 82C59-compatible interrupt controllers. The two controllers are cascaded with interrupt levels 8-15 entering on level 2 of the primary interrupt controller (standard PC configuration). A single interrupt signal is presented to the processors, to which only one processor will respond for servicing. The ICH5-R contains configuration registers that define which interrupt source logically maps to I/O APIC INTx pins.

Interrupts, both PCI and IRQ types, are handled by the ICH5-R. The ICH5-R then translates these to the APIC bus. The numbers in the table below indicate the ICH5-R PCI interrupt input pin to which the associated device interrupt (INTA, INTB, INTC, and INTD) is connected. The ICH5-R' I/O APIC exists on the I/O APIC bus with the processors.

Table 18. PCI Interrupt Routing/Sharing

Interrupt	INT A	INT B	INT C	INT D
USB Controller 1 and 4	ICH5R_PIRQA			
USB Controller 1	ICH5R_PIRQH			
USB Controller 2	ICH5R_PIRQD			
USB Controller 3, IDE, SATA	ICH5R_PIRQC			
Video	ICH5R_PIRQB			
SIO	ICH5R_SERIRQ			
Legacy IDE	ICH5R_IRQ14			
82546GB 1	PXH_PBIRQ4			
82546GB 2	PXH_PBIRQ5			
SCSI Controller 1	IOP_XINT0			
SCSI Controller 2	IOP_XINT1			
Slot 1 (PCI-X 64/133)	IOP_XINT4	IOP_XINT5	IOP_XINT6	IOP_XINT7
Slot 3 (PCI EXP x4)	NA	NA	NA	NA
Slot 4 (PCI EXP x8)	NA	NA	NA	NA
Slot 5 (PCI-X 64/133)	PXH_PAIRQ0	PXH_PAIRQ1	PXH_PAIRQ2	PXH_PAIRQ3
Slot 6 (PCI-X 64/100)	PXH_PBIRQ0	PXH_PBIRQ1	PXH_PBIRQ2	PXH_PBIRQ3
Slot 6 (Upper slot of optional 2-slot riser)	PXH_PBIRQ0	PXH_PBIRQ1	PXH_PBIRQ2	PXH_PBIRQ3
Slot 7 (Lower slot of optional 2-slot riser)	PXH_PBIRQ0	PXH_PBIRQ1	PXH_PBIRQ2	PXH_PBIRQ3

3.4.2.2 APIC Interrupt Routing

For APIC mode, the Intel® Server Board SE7520AF2 interrupt architecture incorporates three Intel I/O APIC devices to manage and broadcast interrupts to local APICs in each processor. The Intel I/O APICs monitor each interrupt on each PCI device including PCI slots in addition to the ISA compatibility interrupts IRQ(0-15). When an interrupt occurs, a message corresponding to the interrupt is sent across a three-wire serial interface to the local APICs. The APIC bus minimizes interrupt latency time for compatibility interrupt sources. The I/O APICs can also supply greater than 16 interrupt levels to the processor(s). This APIC bus consists of an APIC clock and two bidirectional data lines.

3.4.2.3 Legacy Interrupt Sources

The table below recommends the logical interrupt mapping of interrupt sources on the Intel® Server Board SE7520AF2. The actual interrupt map is defined using configuration registers in the ICH5-R.

Table 19. Interrupt Definitions

ISA Interrupt	Description
IRQ0	Timer/counter, HPET #0 in legacy replacement Mode. In APIC mode, cascade from 8259 controller 1
IRQ1	Keyboard
IRQ2	Slave controller INTR output.. In APIC mode Timer/counter, HPET#0
IRQ3	Serial port A
IRQ4	Serial port B
IRQ5	Parallel Port
IRQ6	Floppy
IRQ8	RTC/HPET#1 in legacy replacement mode
IRQ9	Generic, Option for SCI
IRQ10	Generic, Option for SCI
IRQ11	HPET#2, option for SCSI, TCO*
IRQ12	PS2 Mouse
IRQ13	FERR
IRQ14	Primary ATA, legacy mode
PIRQA	USB 1.1 controller 1 and 4
PIRQB	Video
PIRQC	USB 1.1 controller 3, Native IDE, SATA
PIRQD	USB 1.1 controller 2
PIRQE	Option for SCI, TCO, HPET#0,1,2
PIRQF	Option for SCI, TCO, HPET#0,1,2
PIRQG	Option for SCI, TCO, HPET#0,1,2
PIRQH	USB 2.0 EHCI controller 1, Option for SCI, TCO, HPET#0,1,2
Ser IRQ	SIO3

3.4.2.4 Serialized IRQ Support

The Intel® Server Board SE7520AF2 supports a serialized interrupt delivery mechanism. Serialized Interrupt Requests (SERIRQ) consists of a start frame, a minimum of 17 IRQ / data channels, and a stop frame. Any slave device in the quiet mode may initiate the start frame. While in continuous mode, the start frame is initiated by the host controller.

3.4.2.5 IRQ Scan for PCIIRQ

The IRQ / data frame structure includes the ability to handle up to 32 sampling channels with the standard implementation using the minimum 17 sampling channels. The Intel® Server Board SE7520AF2 has an external PCI interrupt serializer for PCIIRQ scan mechanism of ICH5-R to support 16 PCIIRQs.

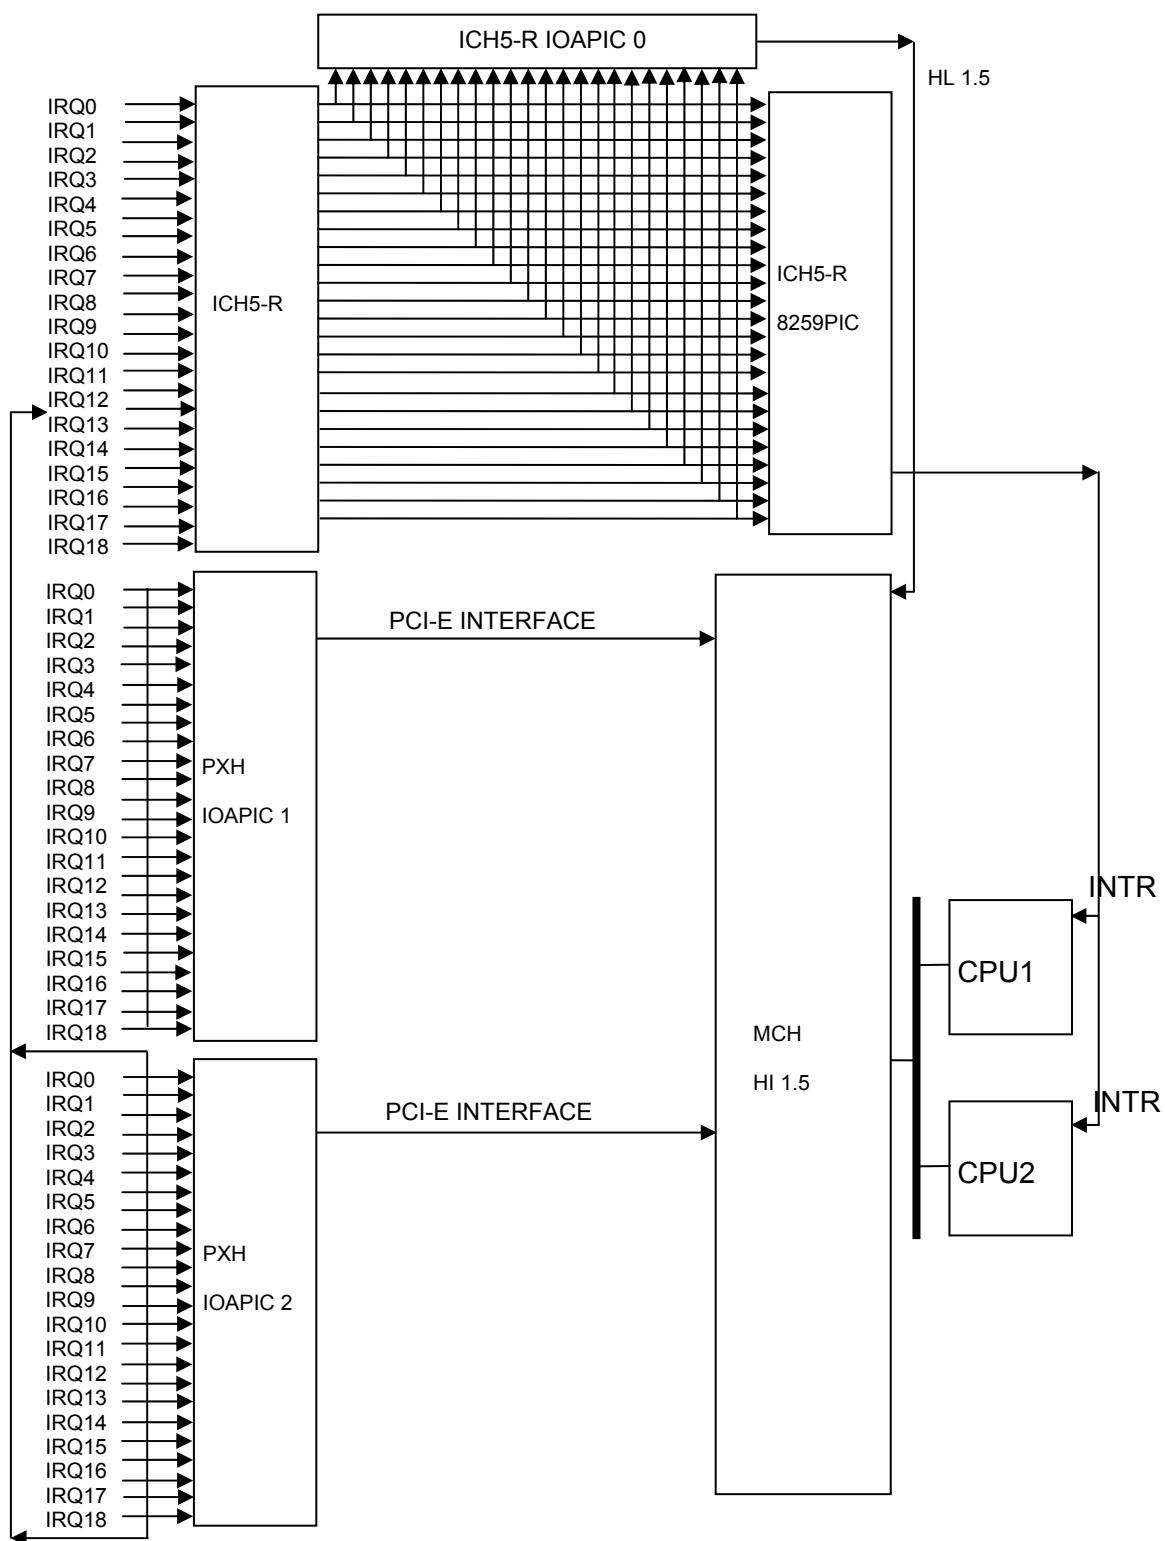


Figure 10. Interrupt Routing Diagram (ICH5-R Internal)

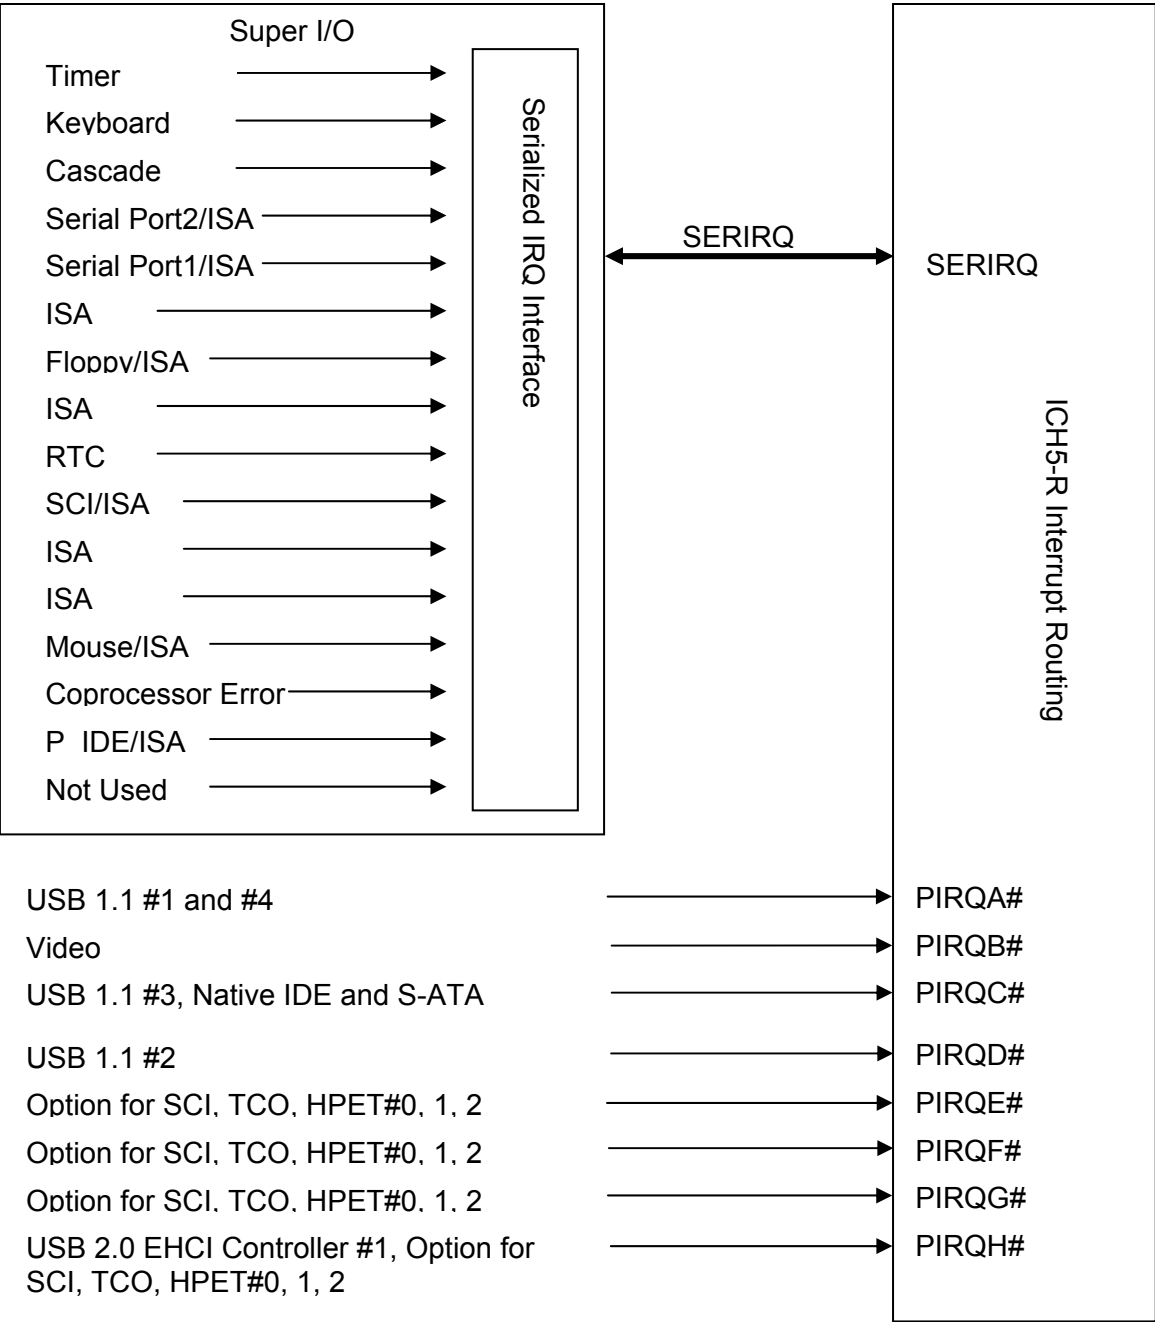


Figure 11. Interrupt Routing Diagram

3.4.3 SCSI Support

The SCSI sub-system on the Intel® Server Board SE7520AF2 is enabled via the LSI Logic* 53C1030 Dual Channel Ultra 320 SCSI controller, two internal 80-pin connector (SCSI Channel A and Channel B), and on-board termination for both SCSI channels.

3.4.3.1 LSI Logic* 53C1030 Dual Channel Ultra 320 SCSI Controller

The LSI Logic* 53C1030 is a PCI-X to Dual Channel Ultra320 SCSI Multifunction controller that supports the *PCI Local Bus Specification, Revision 2.2*, and the *PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a.1*.

The LSI Logic* 53C1030 supports up to a 64-bit, 133 MHz PCI-X bus. The Ultra320 SCSI features for the LSI Logic* 53C1030 include: double transition (DT) clocking, packetized protocol, paced transfers, quick arbitrate and select (QAS), skew compensation, inter-symbol interference (ISI) compensation, cyclic redundancy check (CRC), and domain validation technology. These features comply with the American National Standard Institute (ANSI) T10 SCSI Parallel Interface-4 (SPI-4) draft specification.

DT clocking enables the LSI Logic* 53C1030 to achieve data transfer rates of up to 320 megabytes per second (MB/s) on each SCSI channel, for a total bandwidth of 640 MB/s on both SCSI channels. Domain Validation detects the SCSI bus configuration and adjusts the SCSI transfer rate to optimize bus interoperability and SCSI data transfer rates. Domain Validation provides three levels of domain validation, assuring robust system operation.

The LSI Logic* 53C1030 integrates two high-performance SCSI Ultra320 cores and a 64-bit, 133 MHz PCI-X bus master DMA core. The LSI Logic* 53C1030 employs three ARM966E-S processors to meet the data transfer flexibility requirements of the Ultra320 SCSI, PCI, and PCI-X specifications. Separate ARM* processors support each SCSI channel and the PCI/PCI-X interface. These processors implement the LSI Logic* Fusion-MPT* architecture, a multithreaded I/O algorithm that supports data transfers between the host system and SCSI devices with minimal host processor intervention. Fusion-MPT technology provides an efficient architecture that solves the protocol overhead problems of previous intelligent and non-intelligent adapter designs.

3.4.3.1.1 LSI Logic* 53C1030 Integrated Mirroring and Integrated Striping

The LSI Logic* 53C1030 supports the LSI Logic* Integrated Mirroring Enhanced (IME) and Integrated Striping (IS) technology, which provides physical mirroring or striping of the boot volume through LSI Logic* 53C1030 firmware. These features provide extra reliability and performance for the system's boot volume without burdening the host CPU.

The use of a second disk as a mirror requires the LSI Logic* Fusion-MPT* firmware, which performs writes to both the boot drive and the mirrored drive. Runtime mirroring or striping of the boot drive is transparent to the BIOS, drivers, and operating system.

The IME and IS firmware requires a configuration mechanism, which enables configuration of the mirroring or striping attributes during initial setup or reconfiguration after hardware failures or changes in the system environment. To configure the IME or IS attributes the LSI Logic* BIOS Configuration Utility should be used; this can be accessed by pressing <CTRL+C> during POST. Using the LSI Logic* BIOS and drivers adds support of physical device recognition for

the purpose of Domain Validation and Ultra320 SCSI expander configuration. Host based status software monitors the state of the mirrored drives and reports error conditions as they arise.

Note: Use <CTRL+C> during POST to enter the LSI Logic* BIOS Configuration Utility to configure IME/IS attributes.

For further details on setup and general information on the LSI Logic* Integrated Mirroring Enhanced (IME) and Integrated Striping (IS) technology, refer to LSI Logic Integrated RAID User's Guide.

3.4.3.1.2 53C1030 Summary of Features

The LSI Logic* 53C1030 contains the following SCSI performance features:

- Supports Ultra320 SCSI
- Paced transfers using a free running clock
- Data transfer rate of 320 MB/s on each SCSI channel
- Mandatory packetized protocol
- Quick arbitrate and select (QAS)
- Skew compensation with bus training
- Transmitter pre-compensation to overcome ISI effects for SCSI
- Data signals
- Retained training information (RTI)
- Performance-optimized architecture
- Three ARM966E-S processors provide high performance with low latency
- Two independent Ultra320 SCSI channels
- Designed for optimal packetized performance
- Uses proven integrated LVDlink transceivers for direct attach to either LVD or SE SCSI buses with precision-controlled slew rates
- Supports expanded communication protocol (ECP)
- Supports Integrated Mirroring (IM) which provides features of RAID 1 and RAID 1E for reliability. This configuration allows the user to create one (1) RAID array in Channel A or Channel B (not both) of two to six mirrored disks.
- Supports Integrated Striping (IS) which provides features of RAID 0 for performance and greater capacity. This configuration allows the user to create one (1) RAID array in Channel A or Channel B (not both) of two to six striped disks.
- Uses the Fusion-MPT (Message Passing Technology) drivers to provide support for all Intel® Server Board SE7520AF2 supported operating systems. Refer to the *Intel® Server Board SE7520AF2 Tested Hardware and OS List* for a comprehensive list of supported operating systems.

The LSI Logic* 53C1030 supports these PCI features:

- Operates at up to 133 MHz PCI-X
- Supports 32-bit or 64-bit data
- Supports 32-bit or 64-bit addressing through Dual Address Cycles (DAC)

- Provides a theoretical 1066 Mbytes/s zero wait state transfer rate
- Complies with the PCI Local Bus Specification, Revision 2.2
- Complies with the PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0a
- Complies with PCI Power Management Interface Specification, Revision 1.1
- Complies with the PC2001 System Design Guide
- Offers unmatched performance through the Fusion-MPT architecture
- Provides high throughput and low CPU utilization to off load the host processor
- Presents a single electrical load to the PCI Bus (True PCI Multifunction Device)
- Uses SCSI Interrupt Steering Logic (SISL) to provide alternate interrupt routing for RAID applications
- Reduces Interrupt Service Routine (ISR) overhead with interrupt coalescing
- Supports 32-bit or 64-bit data bursts with variable burst lengths
- Supports the PCI Cache Line Size register
- Supports the PCI Memory Write and Invalidate, Memory Read Line, and Memory Read Multiple commands
- Supports the PCI-X Memory Read DWord, Split Completion, Memory Read Block, and Memory Write Block commands
- Supports up to 8 PCI-X outstanding split transactions
- Supports Message Signaled Interrupts (MSI)

3.4.4 RAID Functionality

The Intel® Server Board SE7520AF2 enables RAID on MotherBoard (ROMB) through the integrated Intel® RAID Controller SROMBU42E. This solution is comprised of the Intel® 80332 I/O Processor with XScale Technology in conjunction with the onboard LSI Logic* 53C1030 SCSI controller and resident RAID firmware on the server board.

The RAID functionality is enabled by using the optional Intel® RAID Activation Key (1-wire serial security EEPROM) and dedicated RAID memory installed on Slot 2 (RAID DIMM). The system BIOS automatically detects the presence of the Intel RAID Activation Key and enables the ROMB subsystem ("IOP ROMB (SROMBU42E) Installed" is displayed in BIOS Setup).



Figure 12. Intel® RAID Activation Key

Note: The Onboard SCSI, Onboard SCSI Option ROM Mode and IOP ROMB options in BIOS Setup are grayed out when an Intel RAID Activation Key is present and compatible RAID memory are present (change of these settings is not allowed). When the Intel RAID Activation Key is present, the user can use <Ctrl+ G> to enter the RAID BIOS Console.

If IOP ROMB is enabled, the system boots using the ROMB option ROM. In this case, the SCSI channels are used by the ROMB solution to enable RAID 5, 10, and 50 functionality. The RAID array configuration is done via the LSI Logic* MegaRAID* setup screens, activated by

<Ctrl+G>. If the Intel RAID Activation Key or RAID DIMM is not present, the system boots on SCSI mode (LSI Logic* 53C1030) and the IOP ROMB option in system BIOS will be grayed out. In both cases, the system BIOS automatically configures and programs the IOP332.

Note: LSI Logic* Integrated Striping/Mirroring and the integrated Intel® RAID Controller SROMBU42E are mutually exclusive solutions. Only one of the two may be used at a time. A RAID array (using one or more drives) transfer from one mode to the other is not supported.

For additional details on the integrated Intel® RAID Controller SROMBU42E solution, see Chapter 4.

3.4.4.1 RAID Memory Subsystem

The Intel® Server Board SE7520AF2 includes a DDR DIMM socket (Slot 2 in the PCI slot area). This DIMM is not shared with main memory and is for exclusive use as RAID cache for ROMB functionality. The ROMB subsystem supports one 128 MB, 256 MB or 512 MB DDR333 unbuffered ECC DIMM, enabling higher performance operation (write back cache). For enhanced data protection and maximum performance, the ROMB functionality on the Intel® Server Board SE7520AF2 supports the Intel® Portable Cache Module accessory, which provides up to 72 hours of battery back-up. The Intel Portable Cache Module accessory is shown below.

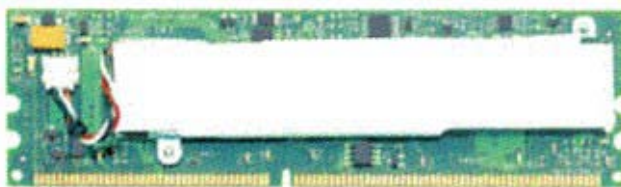


Figure 13. Intel® Portable Cache Module

3.4.5 Parallel ATA (PATA) Support

The Intel® Server Board SE7520AF2 uses one of the integrated IDE controllers of the ICH5-R to provide one parallel ATA channel via a 40-pin IDE connector on the board. It supports up to two drives. The ATA channel can be enabled/disabled and configured in the BIOS Setup Utility.

3.4.5.1 Ultra ATA/100

The IDE interfaces of the ICH5R DMA protocol redefine signals on the IDE cable to allow both host and target throttling of data and transfer rates of up to 100 MB/s.

3.4.6 Serial ATA (SATA) Support

The Intel® Server Board SE7520AF2 uses the integrated Serial ATA (SATA) controller of the ICH5-R to provide two SATA ports on the baseboard. The SATA ports can be enabled/disabled and/or configured in legacy mode by accessing the BIOS Setup Utility during POST.

The SATA function in the ICH5-R has dual modes of operation to support different operating system conditions. In the case of Native IDE enabled operating systems, the ICH5-R has separate PCI functions for serial and parallel ATA. To support legacy operating systems, a single PCI function is available for both the serial and parallel ATA ports. The MAP register provides the ability to share PCI functions. When sharing is enabled, all decode of I/O is done through the SATA registers. Device 31, Function 1 (IDE controller) is hidden by software writing to the Function Disable Register (D31, F0, offset F2h, bit 1), and its configuration registers are not used. The SATA Capability Pointer Register (offset 34h) will change to indicate that MSI is not supported in combined mode.

The ICH5-R SATA controller features two sets of interface signals that can be independently enabled or disabled. Each interface is supported by an independent DMA controller. The ICH5-R SATA controller interacts with an attached mass storage device through a register interface that is equivalent to that presented by a traditional IDE host adapter. The host software follows existing standards and conventions when accessing the register interface and follows standard command protocol conventions.

SATA interface transfer rates are independent of UDMA mode settings. SATA interface transfer rates operate at the maximum speed for the bus, regardless of the UDMA mode reported by the SATA device or the system BIOS. The SATA interface from the ICH5-R supports data transfer rates up to 1.5 Gb/s (150 MB/s).

3.4.6.1 Intel® Embedded Server RAID Technology

The Intel® Server Board SE7520AF2 enables SATA RAID known as the Intel® Embedded Server RAID Technology. This solution is available via the 82801ER ICH5 R (ICH5R), and offers data stripping for higher performance (RAID Level 0), alleviating disk bottlenecks by taking advantage of the dual independent SATA controllers integrated in the ICH5R. In addition, it offers data mirroring for fault tolerance (RAID Level 1). This solution does not result in the loss of PCI resources (request/grant pair) or add-in card slot.

Intel RAID Technology functionality requires the following items:

- ICH5-R
- Intel® RAID Technology Option ROM (enabled in BIOS Setup **F2**)
- Intel® Application Accelerator RAID Edition drivers, most recent revision
- Two SATA hard disk drives

Note: The Intel® Embedded Server RAID Technology can be enabled in system BIOS Setup under IDE Configuration Sub-menu (Configure S-ATA as RAID = Enabled). Once S-ATA is configured as RAID and S-ATA drives are present, the user can access the S-ATA RAID BIOS configuration to set/change attributes by pressing **<Ctrl+E>** during POST. Intel RAID Technology is not available when the SATA controller is in Legacy Mode.

3.4.6.2 SATA RAID Option ROM

The Intel® RAID Technology for SATA Option ROM provides a pre-operating system user interface for the Intel RAID Technology implementation and provides the ability for an Intel® RAID Technology volume to be used as a boot disk as well as to detect any faults in the Intel® RAID Technology volume(s) attached to the Intel RAID controller.

3.4.7 Video Controller

The Intel® Server Board SE7520AF2 provides an ATI* Rage XL PCI graphics accelerator with 8 MB of video SDRAM and support circuitry for an embedded SVGA video subsystem. The ATI* Rage XL chip contains a SVGA video controller, clock generator, 2D and 3D engine, and RAMDAC in a 272-pin PBGA. One 2Mx32 SDRAM chip provides 8 MB of video memory.

The SVGA subsystem supports a variety of modes, up to 1600 x 1200 resolution in 8/16/24/32 bpp modes under 2D, and up to 1024 x 768 resolution in 8/16/24/32 bpp modes under 3D. It also supports both CRT and LCD monitors with up to 100 Hz vertical refresh rate.

Video is accessed using a standard 15-pin VGA connector found on the back edge of the server board. On-board video can be disabled through the BIOS Setup Utility or when an add-in video card is installed in any of the PCI slots.

3.4.7.1 Video Modes

The ATI* Rage XL chip supports all standard IBM VGA modes. The following table shows the 2D/3D modes supported for both CRT and LCD.

Table 20. Intel® Server Board SE7520AF2 Video Modes

2D Mode	Refresh Rate (Hz)	2D Video Mode Support			
		8 bpp	16 bpp	24 bpp	32 bpp
640x480	60, 72, 75, 90, 100	Supported	Supported	Supported	Supported
800x600	60, 70, 75, 90, 100	Supported	Supported	Supported	Supported
1024x768	60, 72, 75, 90, 100	Supported	Supported	Supported	Supported
1280x1024	43, 60	Supported	Supported	Supported	Supported
1280x1024	70, 72	Supported	–	Supported	Supported
1600x1200	60, 66	Supported	Supported	Supported	Supported
1600x1200	76, 85	Supported	Supported	Supported	–
3D Video Mode Support with Z Buffer Enabled					
3D Mode	Refresh Rate (Hz)	3D Video Mode Support with Z Buffer Enabled			
640x480	60,72,75,90,100	Supported	Supported	Supported	Supported
800x600	60,70,75,90,100	Supported	Supported	Supported	Supported
1024x768	60,72,75,90,100	Supported	Supported	Supported	Supported
1280x1024	43,60,70,72	Supported	Supported	–	–
1600x1200	60,66,76,85	Supported	–	–	–
3D Video Mode Support with Z Buffer Disabled					
3D Mode	Refresh Rate (Hz)	3D Video Mode Support with Z Buffer Disabled			
640x480	60,72,75,90,100	Supported	Supported	Supported	Supported
800x600	60,70,75,90,100	Supported	Supported	Supported	Supported
1024x768	60,72,75,90,100	Supported	Supported	Supported	Supported
1280x1024	43,60,70,72	Supported	Supported	Supported	–
1600x1200	60,66,76,85	Supported	Supported	–	–

3.4.7.2 Video Memory Interface

The memory controller subsystem of the ATI* Rage XL arbitrates requests from direct memory interface, the VGA graphics controller, the drawing coprocessor, the display controller, the video scalar, and hardware cursor. Requests are serviced in a manner that ensures display integrity and maximum CPU/coprocessor drawing performance.

The Intel® Server Board SE7520AF2 supports an 8MB (512Kx32bitx4 banks) SDRAM device for video memory. The following table shows the video memory interface signals:

Table 21. Video Memory Interface

Signal Name	I/O Type	Description
CAS#	O	Column Address Select
CKE	O	Clock Enable for memory
CS#[1..0]	O	Chip Select for memory
DQM[7..0]	O	Memory Data Byte Mask
DSF	O	Memory Special Function Enable
HCLK	O	Memory Clock
[11..0]	O	Memory Address Bus
MD[31..0]	I/O	Memory Data Bus
RAS#	O	Row Address Select
WE#	O	Write Enable

3.4.7.3 Host Bus Interface

The ATI* Rage XL supports a PCI 33 MHz bus. The following diagram shows the signals for the PCI interface:

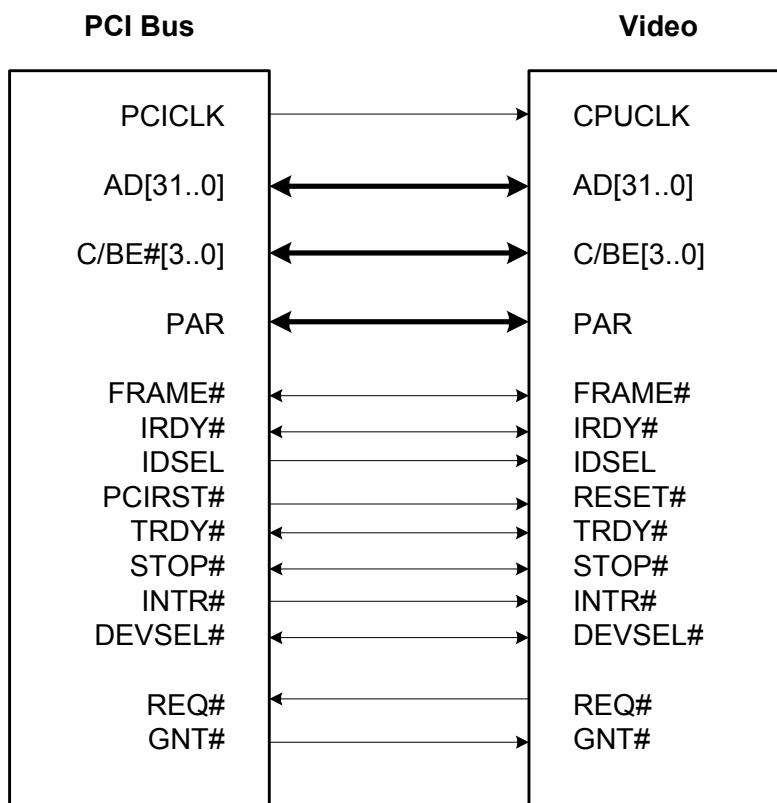


Figure 14. Video Controller PCI Bus Interface

3.4.8 Network Interface Controller (NIC)

The Intel® Server Board SE7520AF2 supports a dual gigabit network interface controller, based on the Intel® 82546GB. The 82546GB is a highly integrated PCI LAN controller in a 21 mm² PBGA package. Each channel supports 10/100/1000 operation and alert-on-LAN functionality. Both channels can be disabled through the BIOS Setup Utility. The 82546EB supports the following features:

- 64-bit PCI-X Rev. 1.0 master interface
- Integrated IEEE 802.3 10Base-T, 100Base-TX and 1000Base-TX compatible PHY
- IEEE 802.3ab auto-negotiation support
- Full duplex support at 10 Mbps, 100Mbps, and 1000 Mbps operation
- Integrated UNDI ROM support
- MDI/MDI-X and HWI support
- Low power +3.3 V device

3.4.8.1 NIC Connector and Status LEDs

The 82546 GB drives two LEDs located on each network interface connector. The link/activity LED to the left of the connector indicates network connection when solidly lit, and Transmit/Receive activity when blinking. The speed LED to the right of the connector indicates 1000-Mbps operations when amber, 100 Mbps operation when green, and 10 Mbps when off.

Table 22. NIC2 Status LED

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

3.4.9 USB 2.0 Support

The Intel® Server Board SE7520AF2 uses the USB 2.0 functionality integrated into ICH5-R. This supports up to six USB ports. The server board utilizes five of these ports: three ports are located on the back edge of the board, and two are available internally via a DH-10 connector that can be routed to the front of the chassis. The front USB ports can be disabled through the BIOS Setup Utility.

3.4.10 Super I/O

Legacy I/O support is provided by using a National Semiconductor* PC87427 Super I/O device. This chip contains the necessary circuitry to control two serial ports, a floppy disk drive, and PS/2-compatible keyboard and mouse. The Intel® Server Board SE7520AF2 supports the following features:

- General Purpose Input/Output (GPIO)
- Two serial ports
- Floppy controller
- Keyboard and mouse controller
- Wake up control

3.4.10.1 General Purpose Input/Output (GPIO)

The National Semiconductor* PC87427 Super I/O provides general-purpose input/output pins that the Intel® Server Board SE7520AF2 utilizes. The following table identifies the pins and the signal names used in the schematic:

Table 23. Super I/O GPIO Usage Table

Pin	Name	IO/GPIO	Intel® Server Board SE7520AF2 Use
50	LED1	Output	FP Power LED
82	LMPCIF	Input/Output	Single wire temperature sensor for PWM2
54	WDO_N	Output	Rolling BIOS Reset
124	GPIO00	Input/Output	Manufacturing Detect Mode
10	GPIO6	Input/Output	BMC SYS IRQ
13	GPIO7/HFCKOUT	Input/Output	Available
37	GPIO27/SIOSMI	Input/Output	SIO SMI
21	GPIOE40	Input/Output	Riser Detect 1

Pin	Name	IO/GPIO	Intel® Server Board SE7520AF2 Use
22	GPIOE41	Input/Output	Riser Detect 0
35	GPIOE42	Input/Output	MROMB
49	GPIOE43	Input/Output	SIO > ICH5R PME
38	GPIOE44/SCI	Input/Output	IMM SCI
51	GPIOE45/LED	Input/Output	BMC (IMM) Card Presence Detect
48	GPO61/SMBSA	Output	LAN A Disable
45	GPO62/LFCLK	Output	LAN B Disable
80	GPEXC	Output	Port 80 Register Clock
79	GPEXD	Input/Output	Serial Data
81	GPEXC2	Output	FRU Register Clock

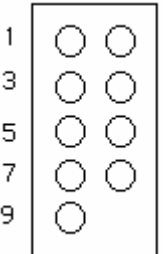
3.4.10.2 Serial Ports

The Intel® Server Board SE7520AF2 supports two serial ports:

- One DB-9 connector located on the back I/O area of the baseboard enabling Serial Port A
- One 9-pin DH-10 header on the server board enables an optional Serial Port B

Serial Port B is an optional port, accessed through a on-board 9-pin internal DH-10 header. A standard DH-10 to DB9 cable can be used to direct Serial Port A out the back or front of a chassis.

Table 24. Serial A Header Pin-out

Pin	Signal Name	Serial Port A Header Pin-out
1	DCD	
2	DSR	
3	RX	
4	RTS	
5	TX	
6	CTS	
7	DTR	
8	RI	
9	GND	

3.4.10.3 Floppy Disk Controller

The floppy disk controller (FDC) in the SIO is functionally compatible with floppy disk controllers in the DP8473 and N844077. All FDC functions are integrated into the SIO including analog data separator and 16-byte FIFO. The Intel® Server Board SE7520AF2 provides access to the floppy interface via an SSI-compliant 36-pin connector.

3.4.10.4 Keyboard and Mouse

Dual stacked PS/2 ports, located on the back edge of the baseboard, are provided for keyboard and mouse support. Either port can support a mouse or keyboard. Neither port will support hot plugging, or connector insertion, while the system is turned on.

3.4.10.5 Wake-up Control

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

3.4.11 BIOS Flash

An Intel® 3 Volt Advanced+ Boot Block 28F320C3 Flash memory component is used as the BIOS flash device. The 28F320C3 is a high-performance 32-megabit memory component that provides 2048Kb x 16 of BIOS and non-volatile storage space. The flash device is connected through the X-bus from the SIO.

3.5 Configuration and Initialization

This section describes the initial programming environment including address maps for memory and I/O, techniques and considerations for programming ASIC registers, and hardware option configuration.

3.5.1 Memory Space

At the highest level, the Intel “Nocona” processor address space is divided into four regions, as shown in the figure below. Each region contains subregions, as described in following sections. Attributes can be independently assigned to regions and subregions using the Intel® Server Board SE7520AF2 registers. The Intel E7520 chipset supports 64 GB of host-addressable memory space and 64KB+3 of host-addressable I/O space. The server board supports only the main memory up to 16GB with DDR2-400.

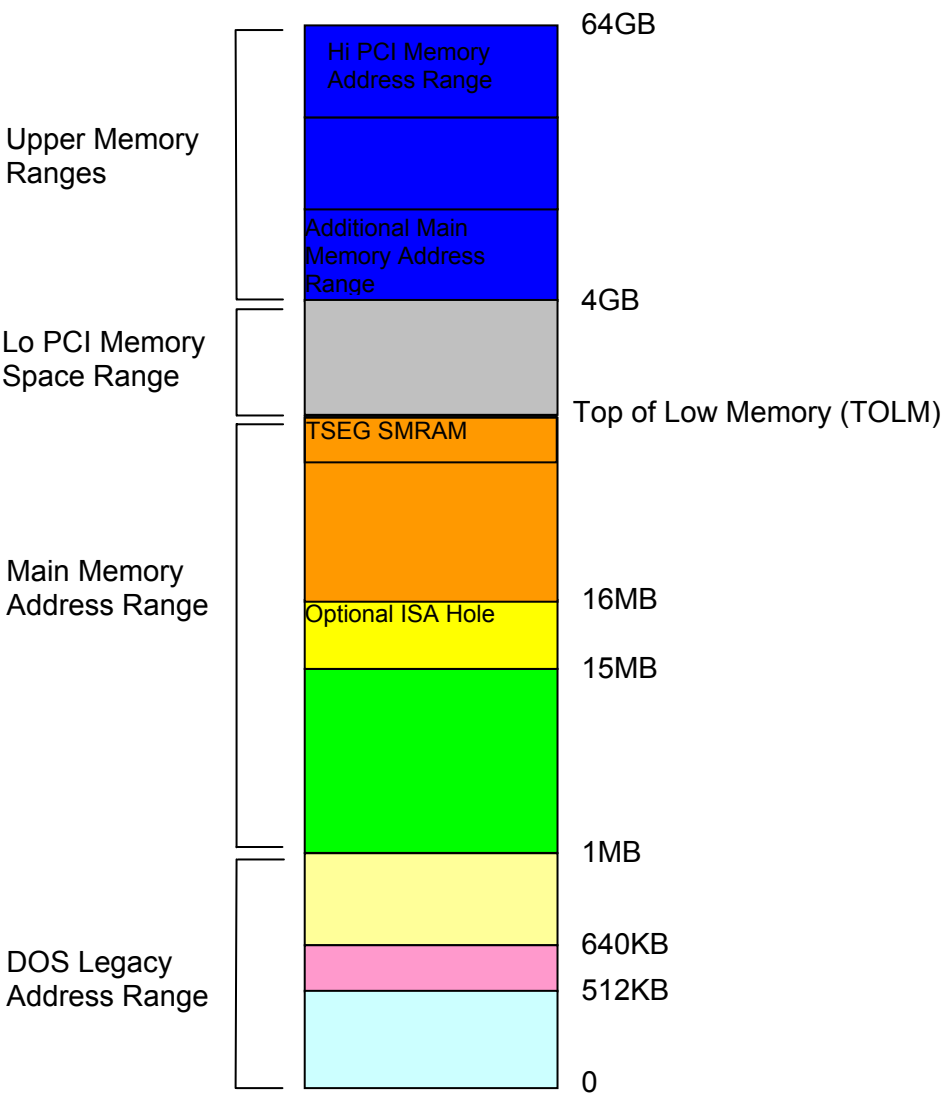


Figure 15. Intel® Xeon Processor Memory Address Space

3.5.1.1 DOS Compatibility Region

The first region of memory below 1 MB was defined for early PCs, and must be maintained for compatibility. The region is divided into subregions as shown in the following figure.

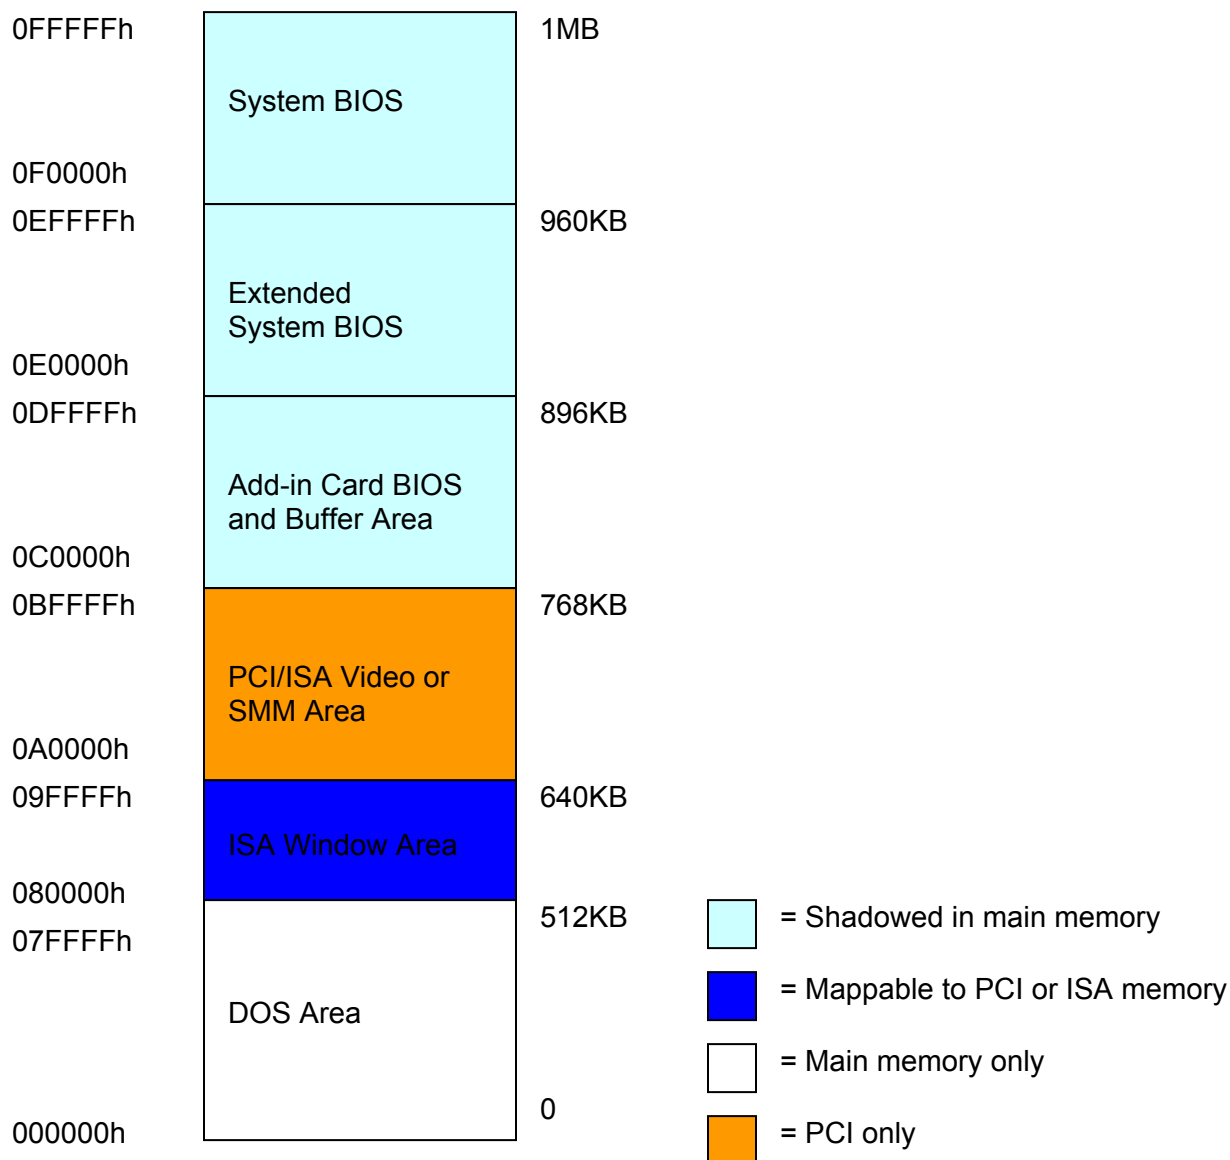


Figure 16. DOS Compatibility Region

3.5.1.1.1 DOS Area

The DOS region is 512 KB in the address range 0 to 07FFFFh. This region is fixed and all accesses go to main memory.

3.5.1.1.2 ISA Window Memory

The ISA window memory is 128 KB, between the address of 080000h to 09FFFFh. This area can be mapped to the PCI bus or main memory.

3.5.1.1.3 *Video or SMM Memory*

The 128 KB graphics adapter memory region, at 0A0000h to 0BFFFFh, is normally mapped to the VGA controller on the PCI bus. This region is also the default region for SMM space.

3.5.1.1.4 *Add-in Card BIOS and Buffer Area*

The 128 KB region, between addresses 0C0000h to 0DFFFFh, is divided into eight segments of 16 KB segments mapped to ISA memory space, each with programmable attributes, for expansion cards buffers. Historically, the 32 KB region from 0C0000h to 0C7FFFh has contained the video BIOS location on the video card

3.5.1.1.5 *Extended System BIOS*

This 64 KB region from 0E0000h to 0EFFFFh is divided into four blocks of 16 KB each, and may be mapped with programmable attributes to map to either main memory or to the PCI bus. Typically this area is used for RAM or ROM. This region can also be used extended SMM space.

3.5.1.1.6 *System BIOS*

The 64 KB region from 0F0000h to 0FFFFFFh is treated as a single block. By default, this area is normally read/write disabled with accesses forwarded to the PCI bus. By manipulating read/write attributes, this region can be shadowed into main memory.

3.5.1.2 *Extended Memory*

Extended memory on Intel® Server Board SE7520AF2 is defined as all address space greater than 1MB. Extended memory region covers 8GB maximum of address space from addresses 0100000h to FFFFFFFFh, as shown in the following figure. PCI memory space can be remapped to top of memory (TOM).

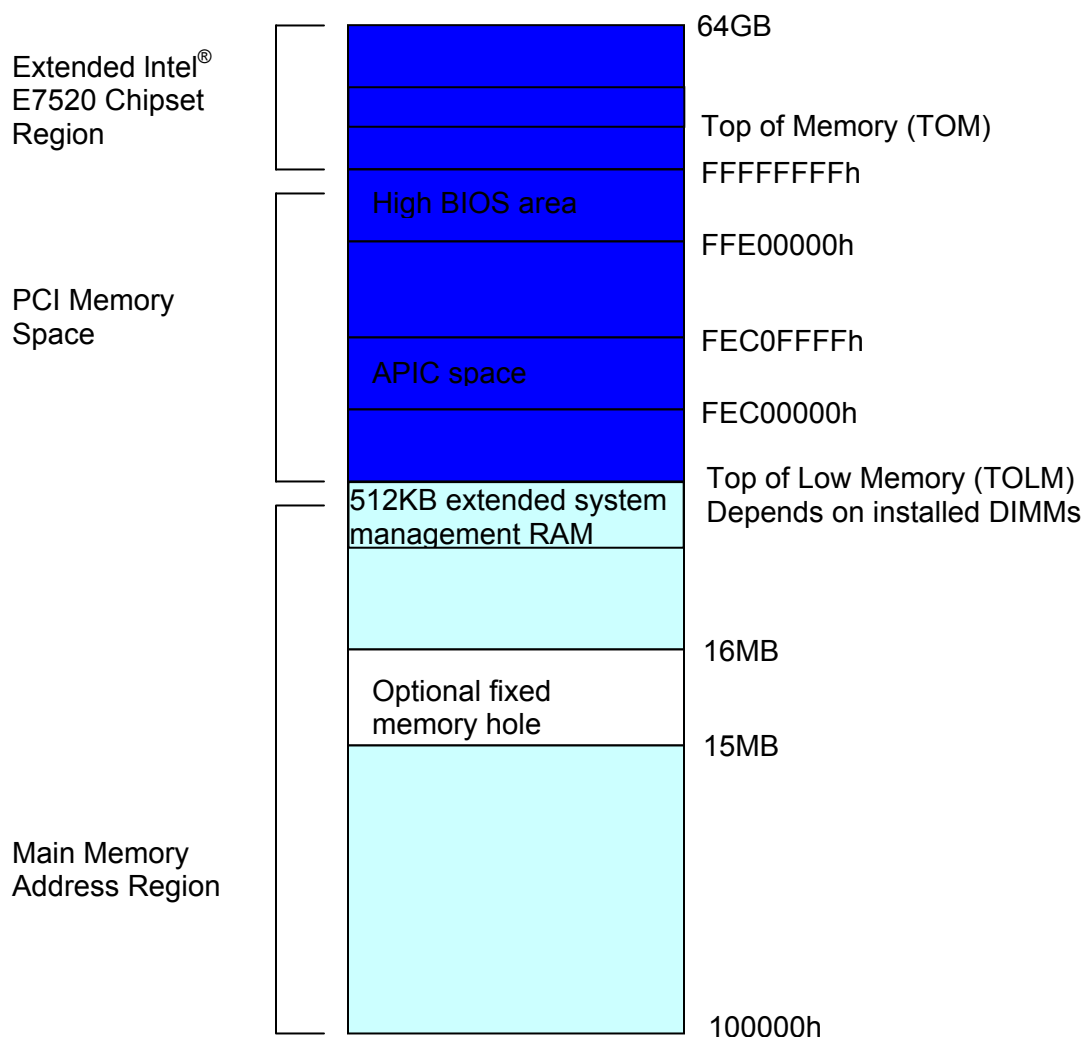


Figure 17. Extended Memory Map

3.5.1.2.1 Main Memory

All installed memory greater than 1 MB is mapped to local main memory, up to 8 GB of physical memory. Memory between 1 MB and 15 MB is considered standard ISA extended memory. 1 MB of memory starting at 15 MB can be optionally mapped to the PCI bus memory space.

The remainder of this space, up to 8 GB, is always mapped to main memory, unless TBSG SMM is used just under TOLM. The range can be from 128KB up to 1MB. 1MB depends on the BIOS setting which limits the top of memory to 256MB. The BIOS occupies 512KB for 32-bit SMI handler.

3.5.1.2.2 *PCI Memory Space*

Memory addresses below 4 GB range are mapped to the PCI bus. This region is divided into three sections: High BIOS, APIC Configuration Space, and General-purpose PCI memory. The General-purpose PCI memory area is typically used memory-mapped I/O to PCI devices. The memory address space for each device is set using PCI configuration registers.

3.5.1.2.3 *High BIOS*

The top 1 MB of Extended memory under 4GB is reserved for the system BIOS, extended BIOS for PCI devices, and A20 aliasing by the system BIOS. The “Nocona” and “Irwindale” processor begins executing from the high BIOS region after reset.

3.5.1.2.4 *I/O APIC Configuration Space*

A 64 KB block located 20 MB below 4 GB (0FEC00000 to 0FEC0FFFFh) is reserved for the I/O APIC configuration space.

The first I/O APIC is located at FEC00000h. The second I/O APIC is located at FEC80000h. The third I/O APIC is located at FEC80100h.

3.5.1.2.5 *Extended “Nocona” and “Irwindale” Processor Region (above 4GB)*

The Intel(r) “Nocona” and “Irwindale” processor -based system can have up to 64 GB of addressable memory. However, the Intel® E7520 chipset supports only 16 GB of addressable memory. The BIOS uses the extended addressing mechanism to use the address ranges.

3.5.1.3 *Memory Shadowing*

The system BIOS and option ROM can be shadowed in main memory. This is done to allow ROM code to execute more rapidly out of RAM. ROM is designated read-only during the copy process while RAM at the same address is designated write-only. After copying, the RAM is designated read-only. After the BIOS is shadowed, the attributes for that memory area are set to read-only so that all writes are forwarded to the expansion bus.

3.5.1.4 *System Management Mode Handling*

The Intel® E7520 chipset supports System Management Mode (SMM) operation in one of three modes. System Management RAM (SMRAM) provides code and data storage space for the SMI_L handler code, and is made visible to the processor only on entry to SMM, or other conditions, which can be configured using Intel® E7520 chipset. The MCH supports three SMM options:

- Compatible SMRAM (C_SMRAM)
- High Segment (HSEG)
- Top of memory Segment (TSEG)

Three abbreviations are used later in the table that describes SMM Space Transaction Handling.

Table 25. SMM Space Transaction Handling

SMM Space Enabled	Transaction Address Space (Adr)	DRAM Space (DRAM)
Compatible (C)	A0000h to BFFFFh	A0000h to BFFFFh
High (H)	0FEDA0000h TO 0FEDBFFFFh	A0000h to BFFFFh
TSEG (T)	(TOLM-TSEG_SZ) to TOLM	(TOLM-TSEG_SZ) to TOLM

Note: High SMM is different than in previous chipsets. In previous chipsets the high segment was the 384KB region from A_0000h to F_FFFFh. However, C_0000h to F_FFFFh was not useful so it is deleted in MCH. TSEG SMM is also different from previous chipsets. In previous chipsets, the TSEG address space was offset by 256MB to allow for simpler decoding and the TSEG was remapped to directly under the TOLM. In the MCH, the TSEG region is not offset by 256MB and it is not remapped.

Table 26. SMM Space Table

Global Enable G_SMRAME	High Enable H_SMRAME	TSEG Enable TSEG_EN	Compatible (C) Range	High (H) Range	TSEG (T) Range
0	X	X	Disable	Disable	Disable
1	0	0	Enable	Disable	Disable
1	0	1	Enable	Disable	Enable
1	1	0	Disable	Enable	Disable
1	1	1	Disable	Enable	Enable

3.5.2 I/O Map

The Intel® Server Board SE7520AF2 allows I/O addresses to be mapped to the processor bus or through designated bridges in a multi-bridge system. Other PCI devices, including the ICH5-R, have built-in features that support PC-compatible I/O devices and functions, which are mapped to specific addresses in I/O space. On the Intel® Server Board SE7520AF2, the ICH5-R provides the bridge to ISA functions.

The I/O map in the following table shows the location of I/O space of all directly I/O-accessible registers. PCI configuration space registers for each device control mapping in I/O and memory spaces, and other features that may affect the global I/O map. All configuration space registers for PCI devices are described in Chapter 6. The sIO chip contains configuration registers that are accessed through an index and data port mechanism; these are also described in Chapter 6.

Table 27. Intel® Server Board SE7520AF2 I/O Map

Address (es)	Resource	Notes
0000h – 000Fh	DMA Controller 1	
0010h – 001Fh	DMA Controller 2	Aliased from 0000h – 000Fh
0020h – 0021h	Interrupt Controller 1	
0022h – 0023h		
0024h – 0025h	Interrupt Controller 1	Aliased from 0020 – 0021h
0026h – 0027h		

Address (es)	Resource	Notes
0028h – 0029h	Interrupt Controller 1	Aliased from 0020h – 0021h
002Ah – 002Bh		
002Ch – 002Dh	Interrupt Controller 1	Aliased from 0020h – 0021h
002Eh – 002Fh	Super I/O (sIO) index and Data ports	
0030h – 0031h	Interrupt Controller 1	Aliased from 0020h – 0021h
0032h – 0033h		
0034h – 0035h	Interrupt Controller 1	Aliased from 0020h – 0021h
0036h – 0037h		
0038h – 0039h	Interrupt Controller 1	Aliased from 0020h – 0021h
003Ah – 003Bh		
003Ch – 003Dh	Interrupt Controller 1	Aliased from 0020h – 0021h
003Eh – 003Fh		
0040h – 0043h	Programmable Timers	
0044h – 004Fh		
0050h – 0053F	Programmable Timers	
0054h – 005Fh		
0060h, 0064h	Keyboard Controller	Keyboard chip select from 87427
0061h	NMI Status and Control Register	
0063h	NMI Status and Control Register	Aliased
0065h	NMI Status and Control Register	Aliased
0067h	NMI Status and Control Register	Aliased
0070h	NMI Mask (bit 7) and RTC address (bits 6::0)	
0072h	NMI Mask (bit 7) and RTC address (bits 6::0)	Aliased from 0070h
0074h	NMI Mask (bit 7) and RTC address (bits 6::0)	Aliased from 0070h
0076h	NMI Mask (bit 7) and RTC address (bits 6::0)	Aliased from 0070h
0071h	RTC Data	
0073h	RTC Data	Aliased from 0071h
0075h	RTC Data	Aliased from 0071h
0077h	RTC Data	Aliased from 0071h
0080h – 0081h	BIOS Timer	
0080h – 008F	DMA Low Page Register	
0090h – 0091h	DMA Low Page Register (aliased)	
0092h	System Control Port A (PC-AT control Port) (this port not aliased in DMA range)	
0093h – 009Fh	DMA Low Page Register (aliased)	
0094h	Video Display Controller	
00A0h – 00A1h	Interrupt Controller 2	
00A4h – 00A5h	Interrupt Controller 2 (aliased)	
00A8h – 00A9h	Interrupt Controller 2 (aliased)	
00ACh – 00ADh	Interrupt Controller 2 (aliased)	
00B0h – 00B1h	Interrupt Controller 2 (aliased)	
00B4h – 00B5h	Interrupt Controller 2 (aliased)	
00B8h – 00B9h	Interrupt Controller 2 (aliased)	
00BCh – 00BDh	Interrupt Controller 2 (aliased)	

Address (es)	Resource	Notes
00C0h – 00DFh	DMA Controller 2	
00F0h	Clear NPX error	Resets IRQ13
00F8h – 00FFh	X87 Numeric Coprocessor	
0102h	Video Display Controller	
0170h – 0177h	Secondary Fixed Disk Controller (IDE)	
01F0h – 01F7h	Primary Fixed Disk Controller (IDE)	
0200h – 0207h	Game I/O Port	
0220h – 022Fh	Serial Port A	
0238h – 023Fh	Serial Port B	
0278h – 027Fh	Parallel Port 3	
0290h – 0298h	NS HW monitor	
02E8h – 02EFh	Serial Port B	
02F8h – 02FFh	Serial Port B	
0338h – 033Fh	Serial Port B	
0370h – 0375h	Secondary Floppy	
0376h	Secondary IDE	
0377h	Secondary IDE/Floppy	
0378h – 037Fh	Parallel Port 2	
03B4h – 03Bah	Monochrome Display Port	
03BCh – 03BFh	Parallel Port 1 (Primary)	
03C0h – 03CFh	Video Display Controller	
03D4h – 03Dah	Color Graphics Controller	
03E8h – 03Efh	Serial Port A	
03F0h – 03F5h	Floppy Disk Controller	
03F6h – 03F7h	Primary IDE – Sec Floppy	
03F8h – 03FFh	Serial Port A (primary)	
0400h – 043Fh	DMA Controller 1, Extended Mode Registers	
0461h	Extended NMI / Reset Control	
0480h – 048Fh	DMA High Page Register	
04C0h – 04CFh	DMA Controller 2, High Base Register	
04D0h – 04D1h	Interrupt Controllers 1 and 2 Control Register	
04D4h – 04D7h	DMA Controller 2, Extended Mode Register	
04D8h – 04DFh	Reserved	
04E0h – 04FFh	DMA Channel Stop Registers	
051Ch	Software NMI (051Ch)	
0CF8h	PCI CONFIG_ADDRESS Register	
0CF9h	Intel® Server Board SE7520AF2 Turbo and Reset Control	
0CFCh	PCI CONFIG_DATA Register	

3.5.3 Accessing Configuration Space

All PCI devices contain PCI configuration space, accessed using mechanism 1 defined in the PCI Local Bus Specification. If dual processors are used, only the processor designated as the

boot-strap processor (BSP) should perform PCI configuration space accesses. Precautions should be taken to guarantee that only one processor performs system configuration.

Two Dword I/O registers in the Intel® E7520 chipset are used for the configuration space register access: CONFIG_ADDRESS (I/O address 0CF8h), CONFIG_DATA (I/O address 0CFCh).

When CONFIG_ADDRESS is written to with a 32-bit value selecting the bus number, device on the bus, and specific configuration register in the device, a subsequent read or write of CONFIG_DATA initiates the data transfer to/from the selected configuration register. Byte enables are valid during accesses to CONFIG_DATA; they determine whether the configuration register is being accessed or not. Only full Dword reads and writes to CONFIG_ADDRESS are recognized as a configuration access by the Intel® E7520 chipset. All other I/O accesses to CONFIG_ADDRESS are treated as normal I/O transactions.

3.5.3.1 CONFIG_ADDRESS Register

CONFIG_ADDRESS is 32 bits wide and contains the field format shown in the following figure. Bits [23::16] choose a specific bus in the system. Bits [15::11] choose a specific device on the selected bus. Bits [10:8] choose a specific function in a multi-function device. Bit [8::2] select a specific register in the configuration space of the selected device or function on the bus.

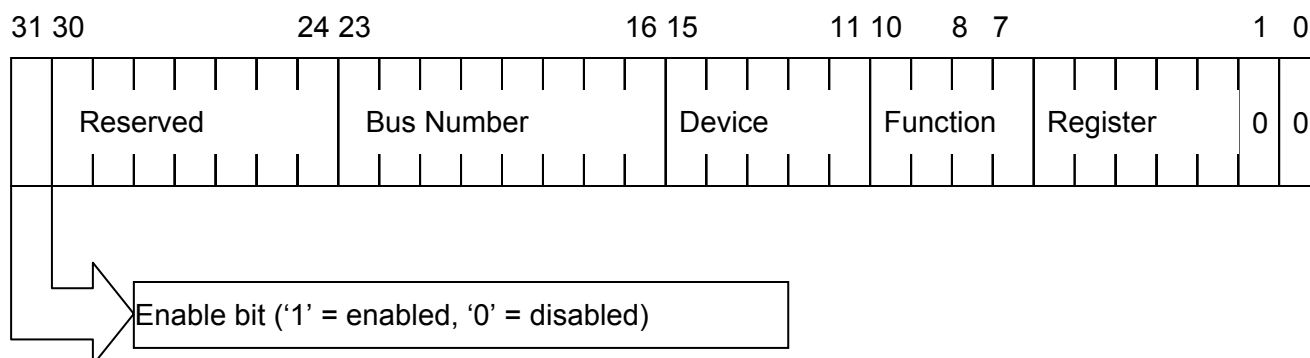


Figure 18. CONFIG_ADDRES Register

3.5.3.1.1 *Bus Number*

PCI configuration space protocol requires that all PCI buses in a system be assigned a bus number. Bus numbers must be assigned in ascending order within hierarchical buses. Each PCI bridge has registers containing its PCI Bus Number and subordinate PCI Bus Number, which must be loaded by POST code. The Subordinate PCI Bus Number is the bus number of the last hierarchical PCI bus under the current bridge. The PCI Bus Number and the Subordinate PCI Bus Number are the same in the last hierarchical bridge.

3.5.3.1.2 Device Number and IDSEL Mapping

Each device under a PCI bridge has its IDSEL input connected to one bit out of the PCI bus address/data signals AD[31::11] for the PCI bus. Each IDSEL-mapped AD bit acts as a chip select for each device on PCI. The host bridge responds to a unique PCI device ID value, that along with the bus number, cause the assertion of IDSEL for a particular device during

configuration cycles. The following table shows the correspondence between IDSEL values and PCI device numbers for the PCI bus. The lower 5-bits of the device number are used in CONFIG_ADDRESS bits [15::11].

Table 28. PCI Configuration IDs and Device Numbers

PCI Device	IDSEL	Bus # / Device # / Function #
MCH host-HI bridge/DRAM controller		00 / 00 / 0,1
MCH EXP Bridge A0		00 / 02 / 00
MCH EXP Bridge A1		00 / 03 / 00
MCH EXP Bridge B0		00 / 04 / 00
MCH EXP Bridge B1		00 / 05 / 00
MCH EXP Bridge C0		00 / 06 / 00
MCH EXP Bridge C1		00 / 07 / 00
ICH5R Hub interface to PCI bridge		00 / 30 / 00
ICH5R PCI to LPC interface		00 / 31 / 00
ICH5R IDE controller		00 / 31 / 01
ICH5R Serial ATA		00 / 31 / 02
ICH5R SMBus controller		00 / 31 / 03
ICH5R USB UHCI controller 1		00 / 29 / 00
ICH5R USB UHCI controller 2		00 / 29 / 01
ICH5R USB UHCI controller 3		00 / 29 / 02
ICH5R USB 2.0 EHCI controller		00 / 29 / 07
Slot 1 (PCI-X 64/133)	AD21	01 / 01 /
Slot 3 (PCI EXP x4)		01 / 04 /
Slot 4 (PCI EXP x8)		02 / 06 /
Slot 5 (PCI-X 64/133)	AD17	00 / 01 /
Slot 6 (PCI-X 64/100)	AD17	01 / 01 /
Slot 6 (Upper Slot of optional 2-slot riser)	AD17	01 / 01 /
Slot 7 (Lower Slot of optional 2-slot riser)	AD18	01 / 02 /
Intel 82546GB Dual Gb NIC	P1B_AD20	/ 04 / 0,1
LSI Logic* 53C1030 Ultra 320 SCSI w/ dual channel	P1A_AD21	/ 05 / 0,1
ATI Rage XL (PCI VGA)	PC_AD28	/ 12 / 0

3.5.4 Hardware Initialization

A system based on the “Nocona” / “Irwindale” processor and the Intel® E7520 chipset is initialized the following manner.

System power is applied. The power-supply provides resets using the PS_GOOD_H signal. The ICH5-R asserts PCI_RST_L to reset the MCH and the PXH. The MCH asserts CPURST_L to reset the processor(s).

The “Nocona” / “Irwindale” processor is initialized, with its internal registers set to default values. Before CPURST_L is de-asserted, the processor asserts BREQ0_L. Processor(s) in the system

determine which host bus agents they are, Agent 0 or Agent 6, according whether their BREQ0_L or BREQ1_L is asserted. This determines bus arbitration priority and order.

The processor(s) in the system determines which processor will be the BSP by using Bootstrap Inter Processor Interrupts (BIPI) on the APIC data bus. The non-BSP processor becomes an application processor and idles, waiting for a Startup Inter Processor Interrupt (SIPI). The BSP begins by fetching the first instruction from the reset vector FFFFFFFF0h. The Intel® E7520 chipset registers are updated to reflect memory configuration, all SDRAM is sized and initialized. All PCI and ISA I/O subsystems are initialized and prepared for booting.

3.6 Clock Generation and Distribution

The main clock source is the CK409B synthesizer/driver component. This device generates majority of the clocks in the design, including the serial reference clock source provided to the DB800 differential buffer. Individual serial reference clocks required for PCI Express* devices/slots are then generated by the DB800. All buses on the Intel® Server Board SE7520AF2 operate using synchronous clocks. Clock synthesizer/driver circuitry on the baseboard generates clock frequencies and voltage levels as detailed in the following table.

Table 29. Clock Generation and Distribution

CK409B	Clk	Pin	Device
	14MHz	REF1	ICH5R
		REF0	Video
	33MHz	PCIF0	ICH5R
		PCIF1	SIO3
		PCIF2	IMM
		PCIO	Video – ATI Rage XL
	48MHz	USB_48	ICH5R
		DOT_48	SIO3
	66MHz	3V66_0	MCH
		3V66_1	ICH5R
	100 MHz	SRC_P	DB800
		SRC_N	DB800
	166MHz	CPU0	MCH
		CPU0_N	MCH
		CPU1	CPU2
		CPU1_N	CPU2
		CPU2	CPU1
		CPU2_N	CPU1
		CPU3	XDP
		CPU3_N	XDP
DB800	Clk	Pin	Device
	100MHz	DIF_0_P	ICH5R
		DIF_0_N	ICH5R
		DIF_1_P	MCH
		DIF_1_N	MCH

		DIF_2_P	IOP332
		DIF_2_N	IOP332
		DIF_3_P	PXH
		DIF_3_N	PXH
		DIF_4_P	HP PCI Express* Slot 3
		DIF_4_N	HP PCI Express* Slot 3
		DIF_5_P	HP PCI Express* Slot 4
		DIF_5_N	HP PCI Express* Slot 4
		DIF_6_P	Pull-down
		DIF_6_N	Pull-down
		DIF_7_N	Pull-down
		DIF_7_N	Pull-down
MCH	Clk	Pin	Device
	200MHz	CMDCLK0_A	DIMM1 (4A)
		CMDCLK0_A_L	DIMM1 (4A)
		CMDCLK1_A	DIMM3 (3A)
		CMDCLK1_A_L	DIMM3 (3A)
		CMDCLK2_A	DIMM5 (2A)
		CMDCLK2_A_L	DIMM5 (2A)
		CMDCLK3_A	DIMM7 (1A)
		CMDCLK3_A_L	DIMM7 (1A)
		CMDCLK0_B	DIMM2 (4B)
		CMDCLK0_B_L	DIMM2 (4B)
		CMDCLK1_B	DIMM4 (3B)
		CMDCLK1_B_L	DIMM4 (3B)
		CMDCLK2_B	DIMM6 (2B)
		CMDCLK2_B_L	DIMM6 (2B)
		CMDCLK3_B	DIMM8 (1B)
		CMDCLK3_B_L	DIMM8 (1B)
PXH	Clk	Pin	Device
	133MHz	PACLKO(0)	Slot 5 (PCI-X 64/133)
	100 MHz	PBCLKO(0)	Slot 6 (PCI-X 64/100)
		PBCLKO(2)	Slot 6 (Upper slot of optional 2-slot riser)
		PBCLKO(1)	Slot 7 (Lower slot of optional 2-slot riser)
		PBCLKO(3)	Intel® 10/100/1000 82546GB Dual Gb Ethernet Controller
IOP332	Clk	Pin	Device
	133 MHz	PACLKO(0)	SCSI
	133 MHz	PBCLKO(0)	Slot 1 (PCI-X 64/133)
ICH5R	Clk	Pin	Device
	SUS CLK	SUSCLK	SIO3, IMM

4. Integrated Intel® RAID Controller SROMBU42E

The Intel® Server Board SE7520AF2 supports Raid-On-MotherBoard (ROMB) through the Intel® 80332 I/O processor (IOP332) with Intel® XScale Technology, the LSI Logic* 53C1030 SCSI controller, resident RAID firmware on the server board and RAID operating system drivers. This chapter provides a general overview of the integrated Intel® RAID Controller SROMBU42E functionality of the Intel® Server Board SE7520AF2.

4.1 Overview

The Intel® RAID Controller SROMBU42E is a high-performance intelligent solution with Redundant Array of Independent Disks (RAID) control capabilities. The controller provides reliability, high performance, and fault-tolerant disk subsystem management. It is an ideal RAID solution for the internal storage of workgroup, departmental, and enterprise systems. The controller offers a cost-effective way to implement RAID on the Intel® Server Board SE7520AF2.

The Intel® RAID Controller SROMBU42E solution (dual-channel) uses the Intel® 80332 I/O processor with Intel® XScale Technology, resident RAID firmware and the onboard LSI Logic* 53C1030 U320 SCSI controller to offer two SCSI RAID channels via the two onboard VHDCI 68-pin SCSI connectors. The integrated controller supports a low voltage differential (LVD) SCSI bus with throughput on each SCSI channel up to 320 Mbytes/s.

The Intel® RAID Controller SROMBU42E functionality can be enabled on the Intel® Server Board SE7520AF2 by using the optional Intel® RAID Activation Key (1-wire serial security EEPROM), using dedicated RAID memory installed on Slot 2 (RAID DIMM) and by enabling the IOP ROMB setup option in BIOS Setup.



Figure 19. Intel® RAID Activation Key

Note: The IOP ROMB option in BIOS Setup is grayed out unless a valid Intel® Activation Raid Key and compatible RAID memory are present.

4.2 Primary Integrated Intel RAID Controller SROMBU42E Features

Features of the integrated Intel® RAID Controller SROMBU42E include:

- Support for DDR333 unbuffered ECC SDRAM memory
- Support for hard disk drives with capacities greater than 8 GB
- Online RAID level migration
- Support for up to 40 logical drives
- Support for up to 32 physical drives per logical array
- RAID remapping
- No reboot necessary after expansion

- Hot-swappable drive support
- Fault-tolerant disk subsystem management
- Variable stripe size
- Instant logical drive availability
- Auto resume during array reconstruction
- Background initialization for quick RAID 5 setup
- Support for non-hard disk drive devices
- SAF-TE intelligent enclosure support
- Ease of management and monitoring using Intel® RAID management utilities
- More than 200 Qtags per array
- User-specified rebuild rate
- Wide Ultra320 LVD SCSI performance up to 320 Mbytes/s
- Support for up to 1 Gbyte of double data rate (DDR) 333MHz unbuffered ECC SDRAM. One 32-, 64-, 128-, 256-, 512-, or 1-Gbyte DIMM can be installed in Slot 2 (RAID DIMM).
- Support for RAID levels 0 (striping), 1 (mirroring), 5 (striping and parity data across all drives), 10 (mirroring and striping), and 50 (RAID 5 and striping)
- Advanced array configuration and management utilities
- Battery backup for up to 72 hours (requires use of the Intel® Portable Cache Module Accessory)
- Support for up to 14 SCSI drives per channel on storage system with SAF-TE enclosures (SCSI accessed fault-tolerant enclosures): 15 SCSI drives per channel for other configurations
- 32 Kbyte x 8 NVRAM for storing RAID system configuration information; the firmware is stored in flash ROM for easy upgrade

Table 30. Integrated Intel® RAID Controller SROMBU42E Features

Feature	Description
RAID Levels	0, 1, 5, 10, 50
SCSI Device Types	Synchronous or Asynchronous
Devices per SCSI Channel	Up to 15 Wide SCSI devices
SCSI Channels	2
SCSI Data Transfer Rate	Up to 320 Mbytes/s per channel
SCSI Bus	LVD only
Cache Function	Write-back (requires Battery Backup Unit), Write-through, Adaptive
	Read Ahead, Non Read Ahead, Read Ahead, Cache I/O, Direct I/O
Multiple Logical Drives/Arrays per Controller	Up to 40 logical drives per controller
Online Capacity Expansion	Yes
Dedicated and Pool Hot Spare	Yes
Hot Swap Devices Supported	Yes
Non-Disk Devices Supported	Yes
Mixed Capacity Hard Disk Drives	Yes
Cluster Support	No

Feature	Description
Hardware Exclusive OR (XOR)	Yes
Direct I/O	Yes
Architecture	Fusion-MPT*

4.2.1 Configuration on Disk

Configuration on Disk saves configuration information both in NVRAM on the Intel® RAID Controller SROMBU42E and on the disk drives attached to the controller. If the Intel® Server Board SE7520AF2 is replaced, the new board detects the RAID configuration from the configuration information on the drives. This maintains the integrity of the data on each drive, even if the drives have changed their target ID.

Note: Configuration on Disk does not work if you change both the SROMBU42E controller and the SCSI connectors to different connectors on the new SROMBU42E controller. It works only if you make one change at a time.

4.2.2 Array Performance Features

Table 31. RAID Controller SROMBU42E Array Performance

SROMBU42E Features	
Drive Data Transfer Rate	320 Mbytes/s
Maximum Scatter/Gathers	26 elements
Maximum Size of I/O Requests	6.4 Mbytes in 64 Kbyte stripes
Maximum Queue Tags per Drive	As many as the drive can accept
Stripe Sizes	2, 4, 8, 16, 32, 64, or 128 Kbyte
Maximum Number of Concurrent Commands	255
Support for Multiple Initiators	Yes

4.2.3 Fault Tolerance Capabilities

Table 32. RAID Controller SROMBU42E Fault Tolerance

SROMBU42E Performance	
Support for SMART ¹	Yes
Optional Battery Backup for Cache memory	Yes, via the Intel Portable Cache Module. Up to 72 hours data retention
Drive Failure Detection	Automatic
Drive Rebuild Using Hot Spares	Automatic
Parity Generation and Checking	Yes

Note:

1. The Self Monitoring Analysis and Reporting Technology (SMART) detects up to 70 percent of all predictable disk drive failures. SMART also monitors the internal performance of all motors, heads, and drive electronics.

4.3 RAID Software Stack

The software described in this document is relevant to the integrated Intel® RAID Controller SROMBU42E functionality of the Intel® Server Board SE7520AF2, and is designated as Software Stack 2. This section describes the RAID controller, drive and array configuration, management utility operation, and drive installation. This software stack is not backward compatible with the previous software stack, Software Stack 1. Software Stack 2 is also not compatible with controllers designed for use with Software Stack 1, including the Intel® RAID Controller SRCMR, SRCZCR, SRCU31, SRCU31L, SRCU32, SRCU42L, and SRCS14L. It is not possible to “Roam” drives between these two software stacks.

4.3.1 General RAID Features

- The Intel® SROMBU42E RAID functionality supports online capacity expansion (OCE). This added capacity can be added to the logical drive and may be presented to the operating system as additional space for the operating system to partition as an additional drive, or may be added to an operating system drive, depending upon the capability of the operating system.
- Online RAID level migration allows for the upgrading of a RAID level (for example, RAID 1 to RAID5), but may require online capacity expansion as part of the process.

Note: This functionality cannot migrate or perform OCE on a spanned RAID, or migrate to a smaller capacity configuration.

- Current firmware and utilities do not allow Random Deletion of logical drives. Drives must be removed in reverse order of creation.
- The total capacity of a configured drive array must be used by one or more logical drives before additional drive arrays can be configured.
- Smart Initialization automatically checks consistency of logical drives when five or more disks are configured in a RAID 5 logical drive. This allows performance optimization by enabling read modify write mode of operation with five or more disks in a RAID 5 array. Peer read mode of operation is used when the RAID 5 array contains three or four physical drives.
- The initialization or rebuild process will automatically resume on the next boot if the system shuts down. Auto resume must be enabled prior to logical drive creation.
- Stripe size is user definable on a per drive basis and can be 2, 4, 8, 16, 32, 64, or 128 kb in size. The default is 64 kb, which has been found to be optimal for many data access types.
- Hot spares can be set as global or dedicated. A global hot spare will automatically come online to replace the first drive to fail on any array on the controller. A dedicated hot spare is assigned to a specific array and will only come online to rebuild a failed drive in that array. A hot spare will only come online if it is the same size or larger than the failing drive (see drive coercion below), and if a drive has been marked as failed. A hot spare can be used across channels on the same RAID controllers. If a drive is removed (and marked as failed) within a logical drive, the hot spare will automatically come online. However, there must be disk activity (I/O to the drive) in order for a missing drive to be marked as failed.

- Drive coercion refers to the ability of the controller to recognize the size of the physical drives that are connected and then force the larger drives to use only the amount of space available on the smallest drive. Drive coercion as implemented in these RAID controllers also allows an option to map out a reserved space to compensate for slightly smaller drive sizes that may be added later. This option allows 16 MB, 256 MB, and 1 GB of drive capacity to remain unused.
- A pulled drive will go to a not responding state if there is no I/O. Not responding is a virtual state, not a persistent state. It is the state that applications will report when management I/O to the drive fails.

4.3.2 RAID Level Migration

Table 33. RAID Level Migration Matrix

	0	1	5
0 ->	N/A	Yes, through OCE	Yes, through OCE
1 ->	Yes	N/A	Yes, through OCE
5 ->	Yes	No	N/A

4.3.3 Fault Tolerant Features

Array configuration information is stored both on the hard drive (COD) and in NVRAM. This helps protect against loss of the configuration due to adapter and/or drive failure. Failed drives are automatically detected and a transparent rebuild of the failed drive automatically occurs using hot spare drives.

Support for SAF-TE enabled enclosures allows enhanced drive failure and rebuild reporting via enclosure LEDs; support also includes hot swapping of hard drives.

Battery backup for cache memory is available as an option for some controllers. RAID controller firmware automatically checks for the presence of the battery module, and if found, allows the write back cache option. The adapter continuously tracks the battery voltage and reports if the battery is low. When low, the battery is first given a fast charge to replenish the charge and is then given a trickle charge to keep the battery at an optimal power level. Adapters that support the battery module include a “dirty cache” LED; when power is lost to the system and data remains in the cache memory that has not been written to disk, this LED signals that this operation needs to be completed. Upon reboot, the data in memory can then be written out to disk.

SMART technology is also supported, which provides a higher level of predictive failure analysis of the hard disk drives by the RAID controller.

4.3.4 Cache Options and Settings

Cache options and settings can be unique for each logical drive.

Cache Write Policy

- Write Through - I/O completion is signaled only after the data is written to hard disk.

- Write Back - I/O completion is signaled when data is transferred to cache.

Cache Policy

- Direct I/O - When possible, no cache is involved for both reads and writes. The data transfers will be directly from host to disk and from disk to host.
- Cached I/O - All reads will first look at cache. If a cache hit occurs, the data will be read from cache; if not, the data will be read from disk and the read data will be buffered into cache. All writes to drive are also written to cache.

Read Policy

- No Read Ahead - Provides no read ahead for the logical drive having this option turned on.
- Read Ahead -Additional consecutive stripes/lines are read and buffered into cache.
- Adaptive - The read-ahead will be automatically turned on and off depending upon whether the disk is accessed for sequential reads or random reads.

4.3.5 Background Tasks

Rebuilding a failed drive is performed in the background. The rebuild rate is tunable from 0-100%.

- The rebuild rate controls the amount of system resources allocated to the rebuild process.
- It is not recommended to increase the rebuild rate over 50% as a higher rebuild rate can result in the operating system requests not being serviced in a timely fashion and causing an operating system panic.

A consistency check validates parity or mirror integrity and can fix a parity or mirror error, if necessary. Consistency checks use the same rate as a rebuild.

Background initialization is a background check of consistency. It has the same functionality as the check consistency option but is automatic and can be canceled only temporarily. If it is canceled, it will start again in five minutes or less. Background initialization is only performed on redundant volumes.

RAID level migration and online capacity expansion are also completed in the background.

4.3.6 Error Handling

Most commands are retried four or more times. The firmware is programmed to provide the best effort to recognize an error and recover from it if possible. Any failures are logged and stored in NVRAM. Operating system-based errors are viewable from the event viewer in Web Console.

RAID-related errors can be reported by the hard drive firmware, SAF-TE controller, or the RAID controller firmware. These errors may be reported to the operating system through RAID management software, through SMART monitoring, or through CIM management. Some errors may also be reported by the SAF-TE controller and logged in the system event log (SEL). In addition, access errors may be reported by the operating system. Depending on the RAID

controller and drive enclosure, the error may be evident via the color of LEDs, the flashing of LEDs, or audible alarms.

4.3.7 Decoding an Audible Alarm

The following list of beep tones is used on Intel RAID controllers using Software Stack 2. These beeps usually indicate that a drive has failed.

- Degraded Array - Short tone, 1 second on, 1 second off
- Failed Array - Long tone, 3 seconds on, 1 second off
- Hot Spare Commissioned - Short tone, 1 second on, 3 seconds off

The tone alarm will stay on during rebuild. After rebuild completes, an alarm with a different tone will sound.

The disable alarm option in either the BIOS Console or Web Console management utilities will hold the alarm disabled after a power cycle. The enable alarm option must be used to re-enable the alarm.

The silence alarm option in either the BIOS Console or Web Console management utilities will silence the alarm until a power cycle or another event occurs.

4.4 Levels of RAID

4.4.1 RAID 0 - Data Striping

In RAID 0, data blocks are split into stripes based on the adjusted stripe size (for example, 128 KB) and the number of hard disks. Each stripe is stored on a separate hard disk. Significant improvement of the data throughput is achieved using this RAID level, especially with sequential read and write. RAID 0 includes no redundancy. When one hard disk fails, all data is lost. A single drive can be chosen as a RAID 0 drive as a method to pass a single drive through to the operating system; however, RAID 0 usually denotes two or more drives.

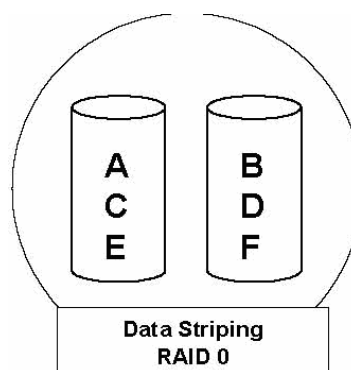


Figure 20. RAID 0

4.4.2 RAID 1 - Disk Mirroring/Disk Duplexing

In RAID 1, all data is stored twice, once each on two identical hard disks making one drive a “mirror” image of the other. When one hard disk fails, all data is immediately available on the other without any impact on performance and data integrity.

With Disk Mirroring, two hard disks are mirrored on one I/O channel. If each hard disk is connected to a separate I/O channel, it is called Disk Duplexing.

RAID 1 represents an easy and highly efficient solution for data security and system availability. It is especially suitable for installations that are not too large (the available capacity is only half of the installed capacity). Disk mirroring requires two drives and each mirrored set is limited to two drives.

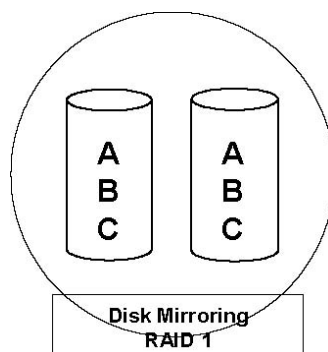


Figure 21. RAID 1

4.4.3 RAID 5 - Data Striping with Striped Parity

RAID 5 works in the same way as RAID 0. The data is striped across the hard disks and the controller calculates redundancy data (parity information) that is striped across all hard disks. Should one hard disk fail, all data remains fully available. Missing data is recalculated from existing data and parity information. The RAID 5 disk array delivers a balanced throughput. Even with small data blocks, which are very likely in a multi-tasking and multi-user environment, the response time is very good. RAID 5 is particularly suitable for systems with medium to large capacity requirements, due to the efficient ratio of installed and available capacity. RAID 5 requires a minimum of three drives in the configuration but can expand to the physical drive capacity of the controller.

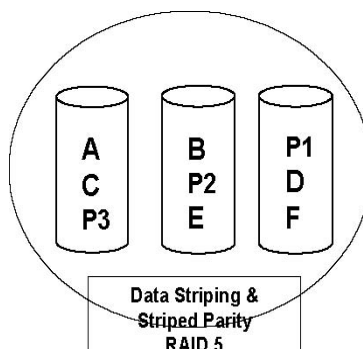


Figure 22. RAID 5

4.4.4 RAID 10 - Combination of RAID 1 and RAID 0

RAID 10 is a combination of RAID 0 (performance) and RAID 1 (data security). Unlike RAID 5, there is no need to calculate parity information. RAID 10 disk arrays offer good performance and data security. As in RAID 0, optimum performance is achieved in highly sequential load situations. Identical to RAID 1, 50% of the installed capacity is lost through redundancy. RAID 10 tolerates a drive failure of one drive per stripe. RAID 10 also requires a minimum of four drives and is set up by spanning two or more RAID 1 arrays. Up to eight RAID 1 arrays can be spanned in a RAID 10 configuration.

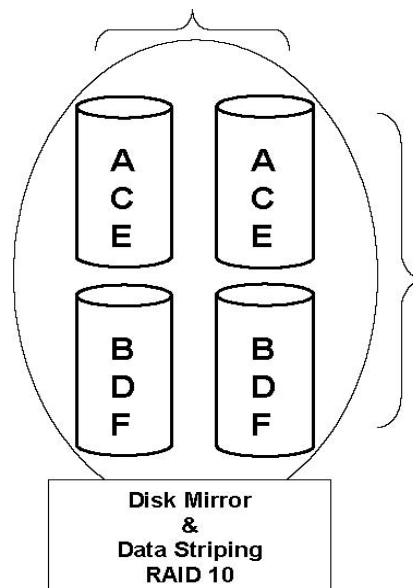


Figure 23. RAID 5

4.4.5 RAID 50 - Combination of RAID 5 and RAID 0

Like RAID 10, RAID 50 is created by first creating multiple RAID 5 arrays and spanning across them. RAID 50 arrays provide the excellent performance of RAID 5 with added data security. RAID 50 requires a minimum of six drives and can span a maximum of eight RAID 5 arrays. RAID 50 tolerates a single drive failure per stripe.

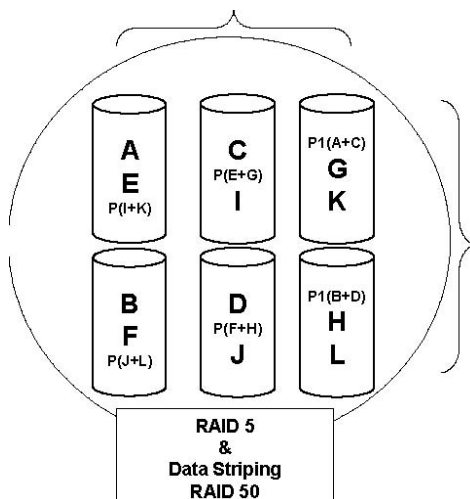


Figure 24. RAID 50

4.4.6 Levels of Drive Hierarchy within the Intel® Integrated RAID Firmware

The Intel Integrated RAID firmware is based on three fundamental levels of hierarchy, as outlined below. Logical drives are created from drive arrays that are in turn created from physical drives.

4.4.6.1 Level 1

Physical drives are hard drives and removable hard disks. Physical drives are the basic components of logical drives. Before they can be used by the firmware, these hard disks must be added to an array. This process identifies a drive by its physical ID and maps it to a logical address. For reasons of data coherency, this information is extremely important for any logical drive construction consisting of more than one physical drive.

4.4.6.2 Level 2

On this level of hierarchy, the firmware forms the array drives. An array drive can be any of the following:

- Single drives or multiple disks
- Chaining sets (concatenation of several hard disks)
- RAID 0 array drives
- RAID 1 array drives
- RAID 5 array drives
- RAID 10 array drives

4.4.6.3 Level 3

On level 3, the firmware forms logical drives. Only these drives can be accessed by the host operating system of the computer. Drives C, D, etc., under MS-DOS*, etc., are always referred to as logical drives by the firmware. The same applies to Novell* NetWare* and UNIX/Linux drives. The firmware automatically transforms each newly installed drive array into a logical

drive by using one or more drive arrays. A single drive array can be used to create a RAID 0, RAID 1, or RAID 5 logical drive, provided there are enough drives in the array for the desired RAID level. Multiple drive arrays can be spanned to create RAID 10 or RAID 50 logical drives. With this option, the data stripe is written across multiple drive arrays.

4.4.6.4 Non-Direct Access Devices

A device that is not a hard disk or a removable hard disk, or that does not behave like one, is called a Non-Direct Access Device. Such a device (for example, a CD-ROM, DAT or tape) is not configured with the management utilities and does not become part of a drive array or a logical drive. Devices of this kind are either operated through the Advanced SCSI programming interface (MS-DOS* or Novell* NetWare*) or are directly accessed from the operating system (UNIX).

4.5 Intel® RAID Controller Drivers

To use the Intel® RAID Controller SROMBU42E, you must install the software driver that is appropriate for your operating system. Refer to the *Intel SE7520AF2 Tested Hardware and OS List* for the complete list of operating systems supported on the SE7520AF2 server platform.

4.6 Intel® RAID Web Console

Intel® RAID products provide a powerful set of software tools for configuring and managing RAID systems. Intel® RAID Web Console is an object-oriented GUI utility that configures and monitors RAID systems locally or over a network. Web Console runs on Microsoft* Windows* XP, Microsoft* Windows* 2000, Microsoft* Windows* Server 2003, Novell NetWare* 6.x, Red Hat* Linux 8.0 and 9.0, Red Hat* Enterprise Linux AS 2.1 and 3.0, and SuSE* Linux 8.0 and 8.2.

With Web Console, you can perform the same tasks as with the BIOS Console. The Intel® RAID Web Console provides on-the-fly RAID migration, creating almost limitless adaptability and expansion of any logical drive while the system remains operational.

Web Console also allows you to:

- Create and manage logical drives
- Add a drive to a RAID logical drive
- Convert from a RAID 0 configuration to a RAID 1 or 5 configuration by adding a physical drive
- Change a Degraded redundant logical drive to an Optimal RAID 0 logical drive
- Remove physical drives from a logical drive
- Convert a RAID 1 or 5 logical drive to a RAID 0 drive.

4.7 Intel® RAID BIOS Console Configuration Utility

The Intel® RAID BIOS Console Configuration Utility provides full-featured, HTML-based configuration and management of RAID arrays. The BIOS Console resides in the resident firmware and is independent of the operating system. The BIOS Console Configuration Utility lets you:

- Choose a configuration method for physical arrays and logical disks
- Create drive arrays
- Define logical drives
- Initialize logical drives
- Access controller, logical drives, and physical arrays to display their properties
- Create hot spare drives
- Rebuild failed drives
- Verify data redundancy in RAID 1, 5, 10, or 50 logical drives

4.7.1 Quick Configuration Steps

This section provides the steps to configure arrays and logical drives using the BIOS Console utility. The following sections describe how to perform each action using the BIOS Console utility. The steps are:

1. Power-on the system.
2. Start the BIOS Console utility by pressing <Ctrl><G>.
3. Start the Configuration Wizard.
4. Choose a configuration method.
5. Create arrays using the available physical drives.
6. Define the logical drive(s) using the space in the arrays.
7. Initialize the new logical drives.

4.7.2 Starting the BIOS Console Utility on the Host Computer

When the host computer boots, hold the <Ctrl> key and press the <G> key when the following appears:

```
Press <Ctrl>-<G> for BIOS Console
```

After you press <Ctrl><G>, the Adapter Selection screen displays. Select an adapter and press the Start button to begin the configuration.

Note: If there is a configuration mismatch between the disks and the NVRAM, the utility automatically displays the Select Configuration screen. Choose whether the configuration should be read from the RAID array or from NVRAM

4.7.3 Screen and Option Descriptions

This section describes the various BIOS Console screens and options as shown in the following figure.

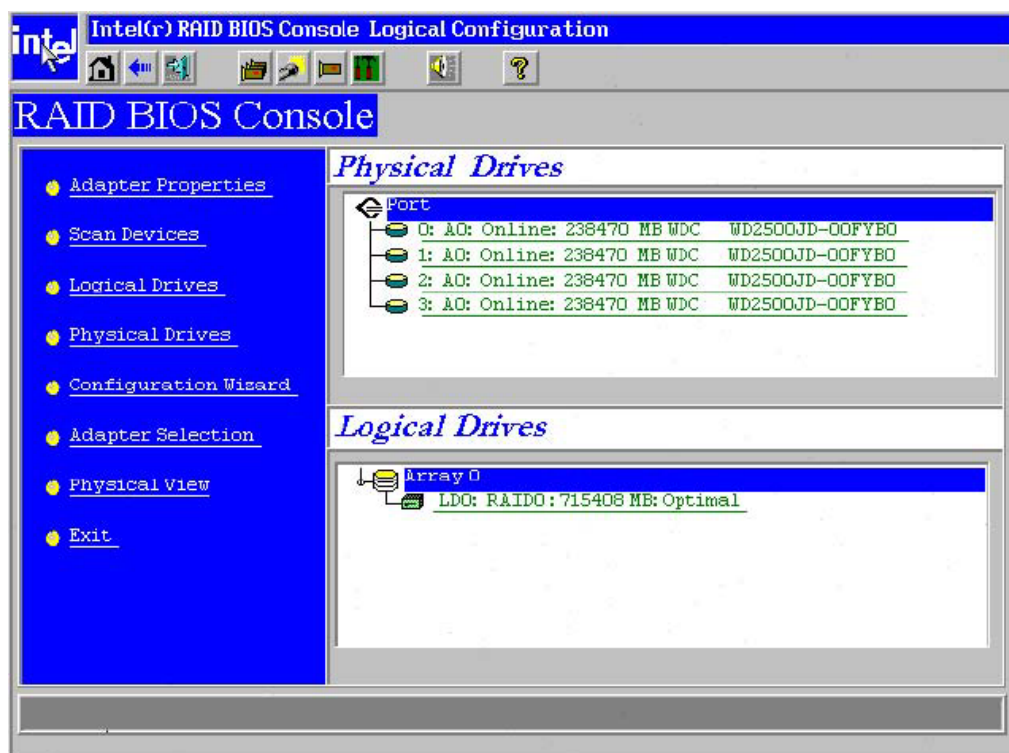


Figure 25. Inte® RAID BIOS Console Screen

4.7.3.1 Main Screen

When you press <Ctrl>-<G> on the host computer, the configuration utility displays the main screen (see previous figure). From the main screen you can scan the devices connected to the controller, and alternate between the physical devices view and the logical devices view. The main screen also provides access to the following screens: Adapter Properties, Physical Devices, Logical Devices, and Configuration Wizard.

4.7.3.2 Adapter Properties Screen

When you select the Adapter Selection option on the main screen, BIOS Console displays a list of the Intel RAID adapters in the system.

The Adapter Properties screen allows you to view and configure the software and hardware of the selected adapter. You access the Adapter Properties screen from the BIOS Console main screen. The following table describes the Adapter Properties menu options.

Table 34. Adapter Properties Menu Options

Option	Description
Firmware Version	This option displays the firmware version number.

Option	Description
BIOS Version	This option displays the BIOS version number.
Battery Backup	This option displays if the optional battery backup unit is installed.
RAM Size	This option displays the amount of RAM on the RAID adapter.
Initiator ID	This option configures the initiator ID for cluster mode.
Rebuild Rate	This option sets the rebuild rate.
Alarm Control	This option disables or silences the alarm.
Coercion Algorithm	This option sets the size of the drive for the coercion operation
BIOS Stops on Error	This option sets if the adapter POST stops if an error occurs.
BIOS Echoes Messages	This option sets if the adapter POST echoes error messages.
BIOS Config Auto Selection	This option sets the action for NVRAM mismatched correction.
Spin Up Parameters	This option sets the hard disk drive spin up delay timing
FlexRAID PowerFail	This option enables the FlexRAID PowerFail feature, which allows drive reconstruction to continue when the system restarts after a power failure. Reserves 2 MB for rebuild and Online Capacity Expansion (OCE).
Fast Initialization	This option sets the span of the logical drive initialization.
PCI Delay Transfer	This option enables PCI delay transfers.
Adapter BIOS	This option enables the adapter BIOS.
Set Factory Defaults	This option loads the default Intel RAID BIOS Console CU settings.
Auto Rebuild	This option automatically rebuilds drives when they fail.
Class Emulation Mode	This option is preset for Mass Storage as class emulation mode.

4.7.3.3 Scan Devices Screen

When you select the Scan Devices option on the Main screen, BIOS Console checks the physical and logical drives for any changes of the drive status. BIOS Console displays the results of the scan in the physical and logical drive descriptions.

4.7.3.3.1 SCSI Channel Properties

When you select this option, properties of the SCSI channel can be viewed. Depending on the I/O bus properties, options available may include bus speed selection and termination configuration.

4.7.3.4 Logical Drives Screen

You can access the Logical Drives screen by clicking on a logical drive in the logical drive list on the main screen. The upper right section of the screen displays the logical drives that currently exist. The Logical Drives screen provides options to:

- Initialize the logical drives
- Check consistency
- Display the logical drive properties
- Boot from a logical drive Press Go to perform the selected action. Press Reset to delete any changes.

Initialization – The Initialize option initializes the selected logical drive by writing zeroes to the entire volume. You should initialize each new logical drive that you configure.

Warning: Initializing a logical drive deletes all information on the physical drives that compose the logical drive.

Check Consistency – This option verifies the correctness of the redundancy data and is available for arrays using RAID 1, 5, 10, or 50. If a difference in the data is found, BIOS Console assumes that the data is accurate and automatically corrects the parity value.

Properties – Through the Properties option you can:

- Display the logical drive properties (such as RAID level, logical drive size, and stripe size)
- Display the read, write, and I/O policies
- Change the read, write, and I/O policies
- Start initialization
- Start a consistency check

4.7.3.5 Physical Drives Screen

This screen displays the physical drives for each channel or port. From this screen, you can rebuild the physical arrays or view the properties for the physical drive you select. Press Reset to return to the configuration that existed before you made any changes.

4.7.3.6 Configuration Wizard Screen

This option enables you to clear a configuration, create a new configuration, or add a configuration. The “Detailed Configuration Instructions” section provides detailed steps for using the Configuration Wizard.

4.7.3.7 Adapter Selection

This option allows you to choose an Intel RAID adapter installed in the system.

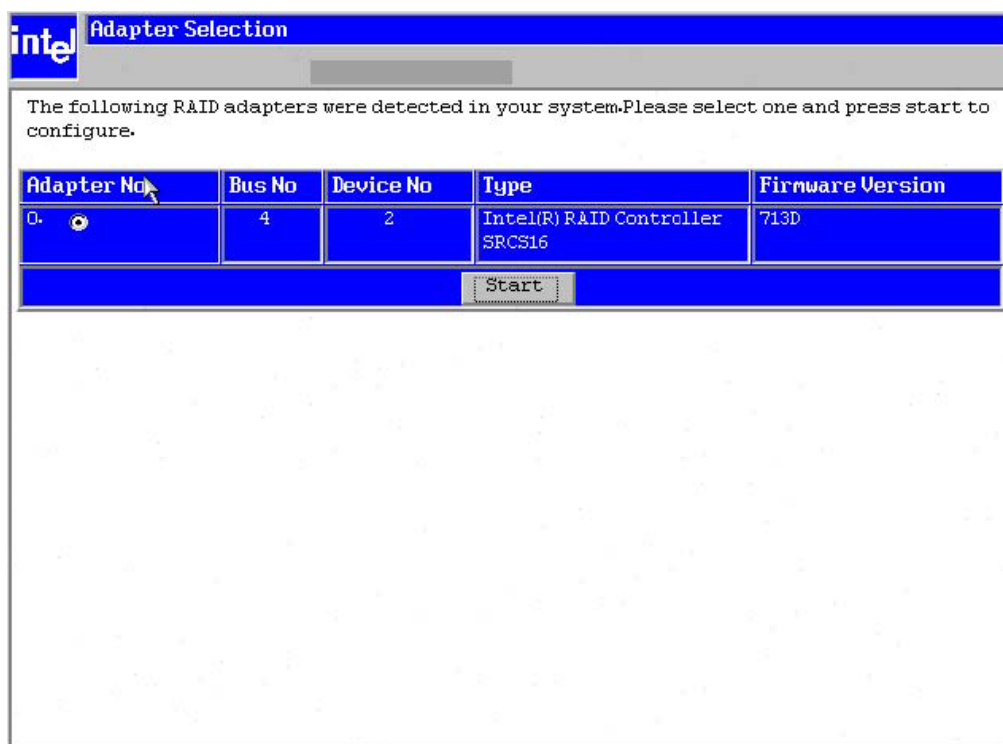


Figure 26. Adapter Selection Screen

4.7.3.8 Physical View/Logical View Option

This option toggles between Physical View and Logical View.

4.7.3.9 Exit

This option allows you to exit and reboot the system.

4.7.4 Configuration Mismatch Screen

A configuration mismatch occurs when the data in the NVRAM and the hard disk drives are different. It will be automatically displayed after POST when a configuration mismatch occurs. The Configuration Mismatch screen allows you to:

1. Select Create New Configuration to delete the previous configuration and create a new configuration
2. Select View Disk Configuration to restore the configuration from the hard disk
3. Select View NVRAM Configuration to restore the configuration from the NVRAM
4. Select the configuration method.

4.7.5 Detailed Configuration Instructions

This section provides detailed steps for using the Configuration Wizard to set up a RAID array.

4.7.5.1 Configuration Wizard

Start the Configuration Wizard by selecting the Configuration Wizard icon on the BIOS Console main screen.

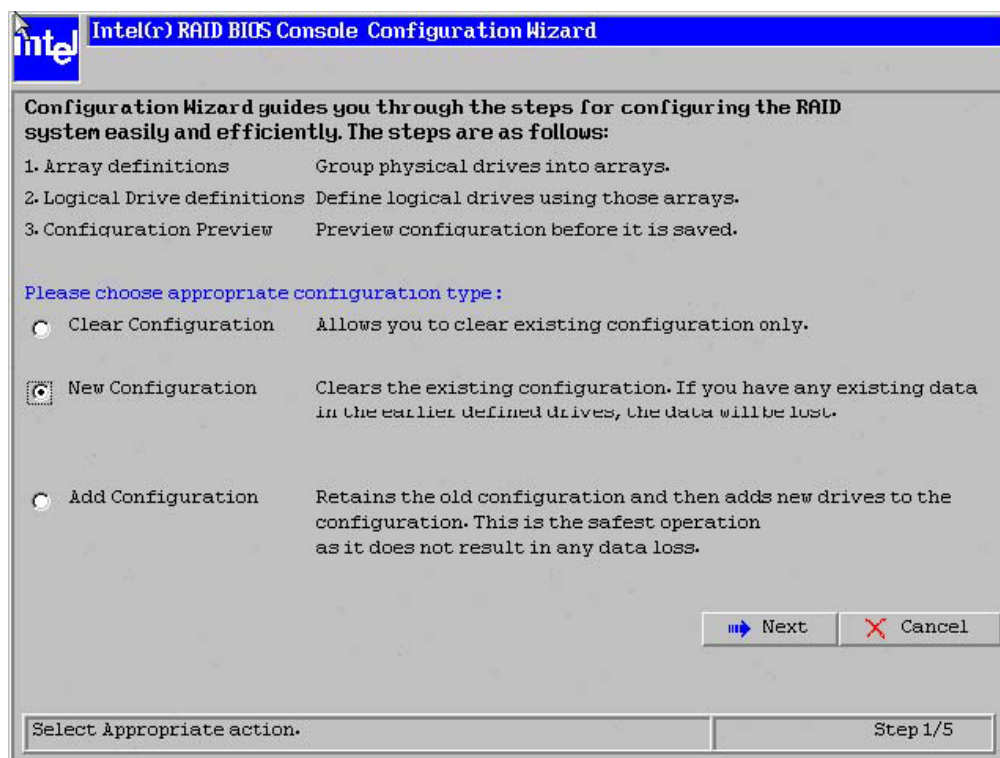


Figure 27. Configuration Wizard Screen

Select Auto configuration with redundancy, Auto configuration without redundancy, or Custom configuration. These options are described below.

4.7.5.2 Auto Configuration with Redundancy

This option configures RAID 1 for systems with 2 drives or RAID 5 for systems with 3 or more drives. All available physical drives will be included in the logical drive using all available capacity on the disks.

Note: Hot spare drives must be designated before starting auto configuration using all available capacity on the disks.

4.7.5.3 Custom Configuration

Configures all available drives as a RAID 0 logical drive. Choose Custom Configuration and select Next.

4.7.5.4 Define the Drive Arrays

At the array definition screen, hold down the CNTL key and click each drive you want included in the array. To undo the changes, press the Reclaim button. Click Accept Array and Next.

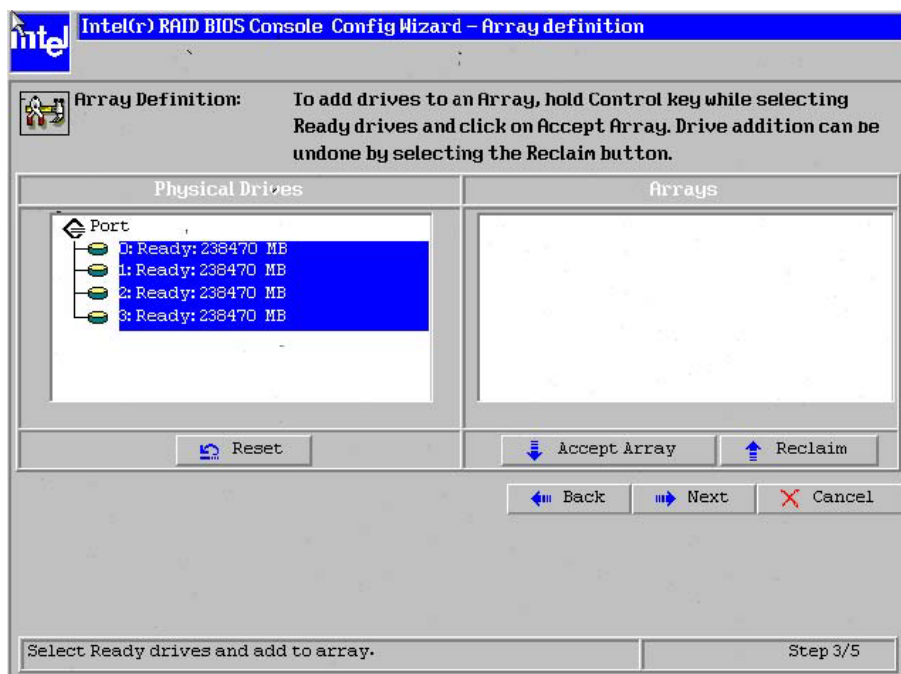


Figure 28. Array Definition Screen

4.7.5.5 Configure the Logical Drive

The logical drive parameters are the RAID level, stripe size, and read-ahead policy.

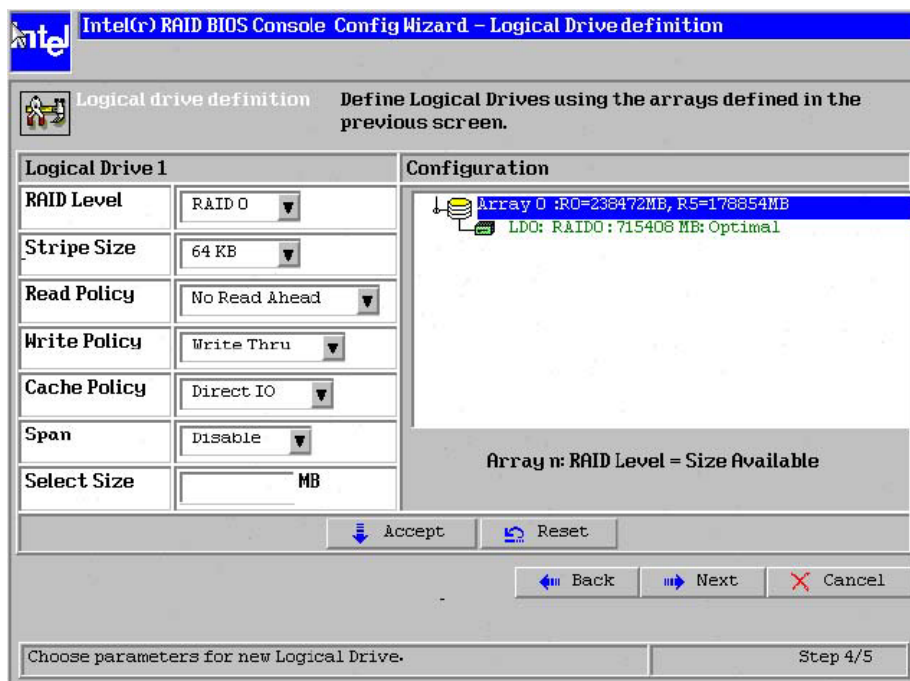


Figure 29. Logical Drive Definition

RAID Level – Choose the RAID level.

Stripe Size – Specify the size of the segment written to each disk in a RAID configuration. You can set the stripe size to 4, 8, 16, 32, 64, or 128 Kbytes. The default is 64 Kbytes.

Read Policy – Enables the read-ahead feature for the logical drive. You can set this parameter to Normal, Read-ahead, or Adaptive. Normal specifies that the controller does not use read-ahead for the current logical drive. Read-ahead specifies that additional consecutive stripes are Read and Buffered into Cache. This option will improve performance for sequential reads. Adaptive specifies that the controller begins using read-ahead if the two most recent disk accesses occurred in sequential sectors. Read-ahead is the default setting.

Write Policy – Use this screen to select the write and cache policies. The write policy parameter specifies the cache write policy. You can set the write policy to Write-back or Write-through. In Write-back caching, the controller sends a data transfer completion signal to the host when the controller cache receives all the data in a transaction. In Write-through caching, the controller sends a data transfer completion signal to the host after the disk subsystem receives all the data in a transaction. Write-through caching is the default setting. Write-through caching has a data security advantage over write-back caching. Write-back caching has a performance advantage over write-through caching, but it should only be enabled when the optional battery backup module is installed.

Caution: Do not use write-back caching for any logical drive in a Novell* NetWare* volume.

Cache Policy – The cache policy applies to I/O on a specific logical drive. It does not affect the read ahead cache. The options are Cached I/O or Direct I/O. Cached I/O buffers all reads in cache memory. Direct I/O does not buffer reads in cache memory. When possible, Direct I/O does not override the cache policy settings. Direct I/O transfers data to cache and the host concurrently. If the same data block is read again, the host reads it from cache memory.

Span – Enable or disable the spanning mode for the current logical drive. If spanning is enabled, the logical drive can occupy space in more than one array. If spanning is disabled, the logical drive can occupy space in only one array. For two arrays to be spannable, they must have the same stripe width and must be consecutively numbered. If these criteria are not met, the utility ignores the Span setting. Spanning allows the logical drive to stripe across multiple arrays. Using this option enables the use of RAID 10 and RAID 50.

RAID 10: To configure a RAID 10 array, create multiple RAID 1 drive arrays (minimum of two) and enable the span option, then choose the size of the logical drive. It is possible to span up to eight arrays in a logical drive. When looking at the logical drive, you will not see a “RAID 10” label, but you will see multiple arrays in the logical drive representing a RAID 10 configuration.

RAID 50: To configure a RAID 50 array, create multiple RAID 5 drive arrays (minimum of two) and enable the span option, then choose the size of the logical drive. It is possible to span up to 8 arrays in a logical drive. When looking at the logical drive, you will not see a “RAID 50” label, but you will see multiple arrays in the logical drive representing a RAID 50 configuration.

Select Size – Set the size of the logical drive in Mbytes. The right pane of the logical drive configuration window will list the maximum capacity that can be selected, depending on the RAID level chosen.

Accept the changes. Click on the Accept button to accept the changes, or click on the Reset button to delete the changes and return to the previous settings. The BIOS Console CU displays a preview of the configuration. Click on Accept to save the configuration, or click on Back to return to the previous screens and change the configuration.

4.7.6 Finish Configuration

Click on the Next button and then the Accept button to complete the selection. Accept to accept the configuration. You are prompted to save the configuration and then to initialize the logical drive. Choose Yes to initialize the local drive.

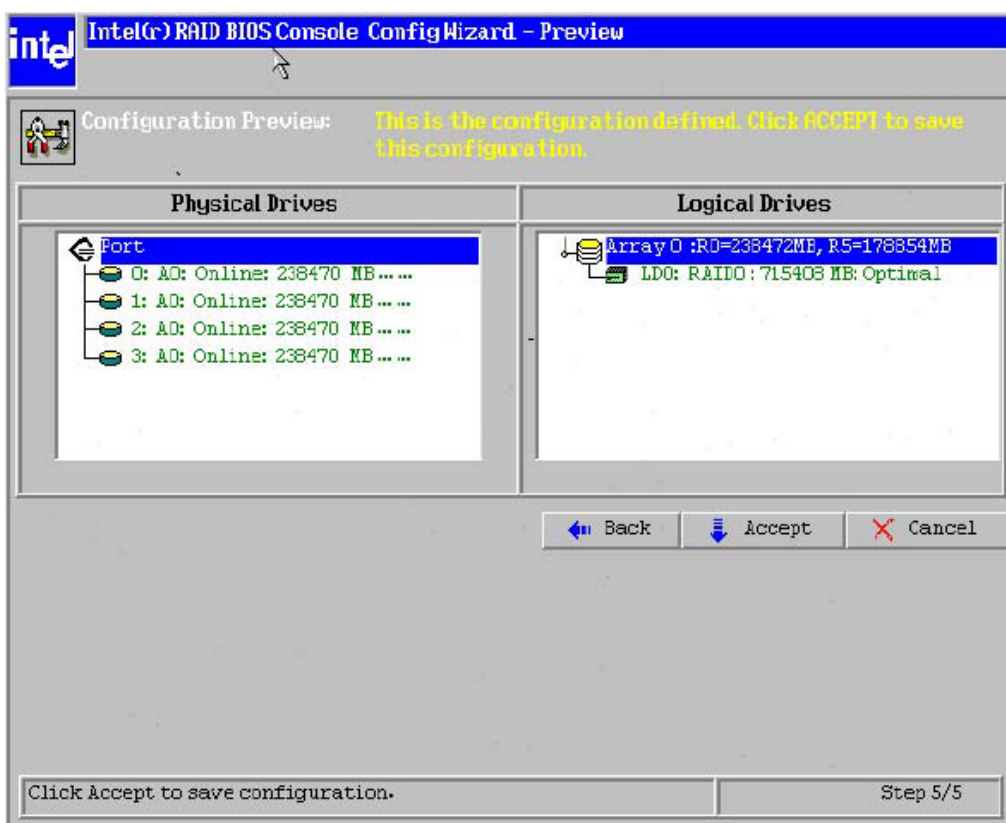


Figure 30. Configuration Wizard Preview Screen

4.7.7 Initialize the Drive

Click Initialize to begin the initialization process. If fast initialization is selected as the default for the adapter properties, a preliminary initialization will complete within seconds and full initialization will run in the background after the operating system is booted. If fast initialization is not selected as the default for the adapter properties, the initialization may take up to two hours to complete.

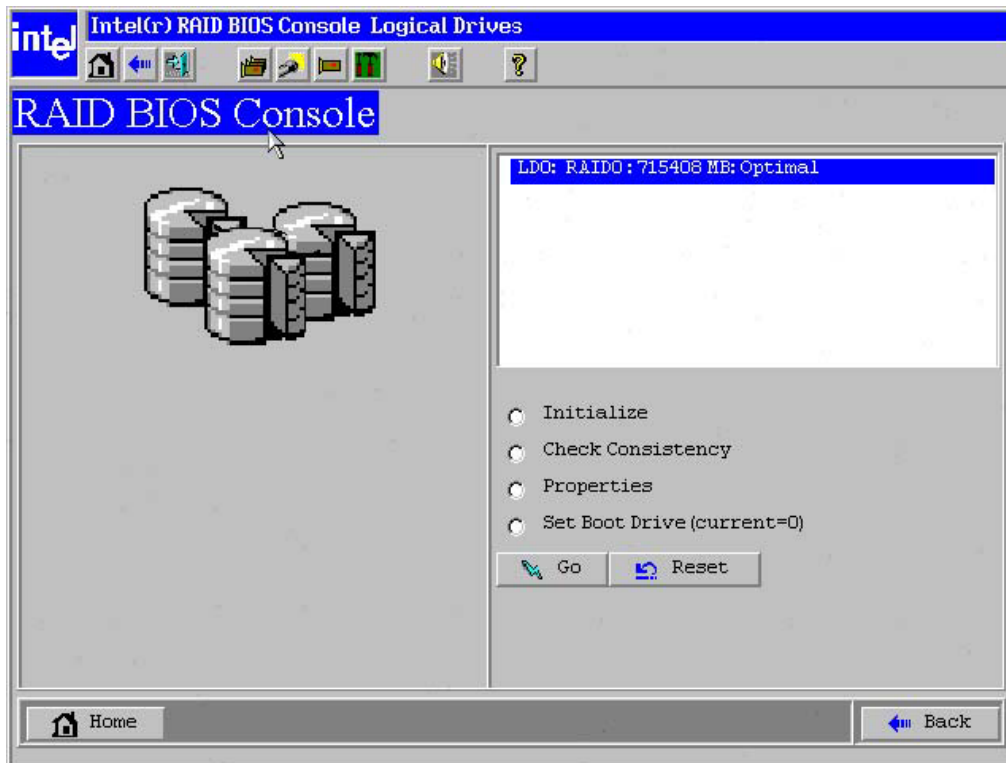


Figure 31. Logic Drives Screen

Click the Home button to return to the main configuration screen. Select an additional logical drive to configure or exit the BIOS Console Configuration Utility and reboot the server system.

< This page intentionally left blank. >

5. System BIOS

The following chapter describes the features supported by the SE7520AF2 BIOS.

5.1 BIOS Identification String

The BIOS identification string is used to uniquely identify the revision of the BIOS being used on the system. The string is formatted as depicted in the following figure:

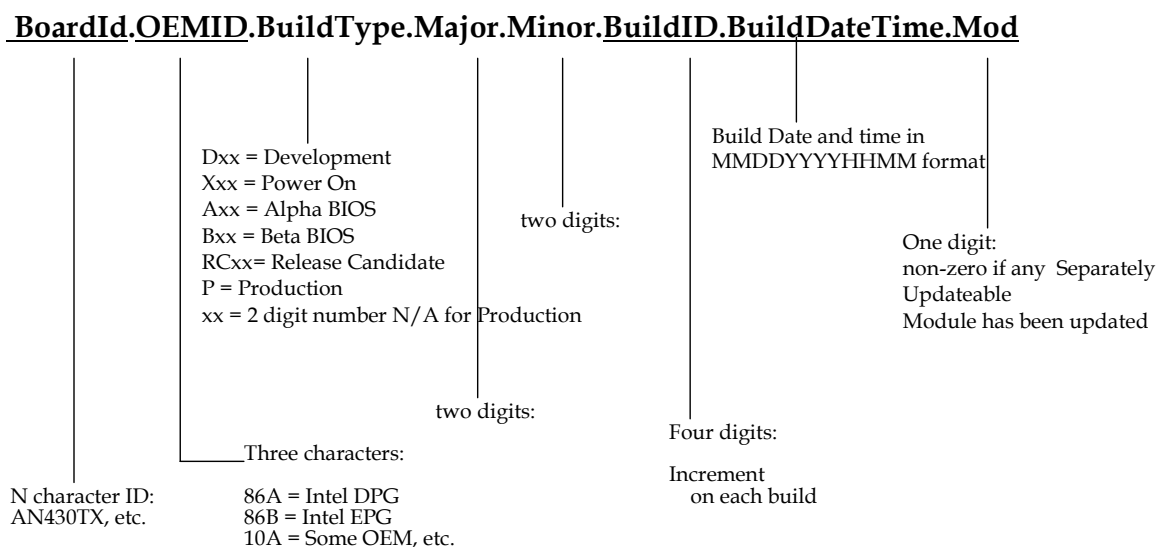


Figure 32. SE7520AF2 BIOS Identification String

The system BIOS has the unique Board ID of “SE7520AF2”. The following is a sample production data string that is displayed during POST:

SE7520AF20.86B.P01.01.00.0052.081320040715

5.2 Supported BIOS Features

Table 35. BIOS Features

Processors	<ul style="list-style-type: none"> Two 800 MHz system bus Dual Processing capable Intel® Xeon™ processors The packaging uses features provided by Intel® Xeon™ processor family to identify processors installed and features supported All processors in the system must operate at the same frequency and have the same cache sizes
Memory	<ul style="list-style-type: none"> Only registered DDR2-400MHz SDRAM will be supported via eight DIMM slots present on the board Memory mirroring supported in Dual channel configuration only Memory sparing supported in Dual and single channel configuration Intel x4 Single Device Data Correction (SDDC) only for memory configurations in dual

	channel mode
Chipset	Intel® E7520 chipset: <ul style="list-style-type: none"> ▪ Memory Controller Hub ▪ PXH PCI Bridge ▪ ICH5-R Controller Hub ▪ IOP332 I/O processor
PCI Slots	Five PCI segments: <ul style="list-style-type: none"> ▪ Two PCI Express* slots with x8 connectors (hot plug support enabled on HP SKU) ▪ Two 64 bit, 133 MHz PCI-X slots (hot plug support enabled on HP SKU) ▪ One 64 bit, 100 MHz PCI-X slot
I/O Controls	<ul style="list-style-type: none"> ▪ National Semiconductor* PC87427Super I/O ▪ LSI Logic* 53C1030 U320 Dual Channel SCSI Controller (with integrated mirroring and striping) ▪ ATI* 2D/3D Rage XL PCI graphics controller with 8 MB SDRAM ▪ Intel® 82546GB Dual Gigabit NIC controller
Server Management	<ul style="list-style-type: none"> ▪ On-board National Semiconductor* PC87431 mini Baseboard Management Controller (mBMC) ▪ Sahalee based BMC via Flexible Management Module upgrade Supporting Emergency Management Port (EMP), SATA and SCSI SAF-TE controller, KVM capture/redirection to LAN or EMP, Blue screen capture/re-direction (as backup), Console re-direct over LAN
BIOS Flash	Intel® TE28F320C3B flash part with 4MB
I/O Ports	<ul style="list-style-type: none"> ▪ Five USB 2.0 ports (two port header, three external ports in back) ▪ Two RS232 serial ports (Serial A via DB9 on I/O panel, Serial B via DH10 header on board) ▪ Single parallel legacy ATA-100 ▪ Native Serial ATA 150 with RAID 0 support ▪ Two interchangeable PS/2 ports for keyboard and mouse (I/O panel) ▪ Two RJ45 connectors for Gbit Ethernet ports ▪ One video connector (I/O panel) ▪ SSI Front Panel header ▪ ICMB/IPMB headers ▪ LCD header ▪ Floppy

5.3 Processor Initialization

The following section describes the memory initialization performed by the Intel® SE7520AF2 Server BIOS.

5.3.1 Multiple Processor Initialization

IA32 processors have a microcode-based BSP-arbitration protocol. On reset, all of the processors compete to become the bootstrap processor (BSP). If a serious error is detected during a Built-in Self-Test (BIST), that processor will not participate in the initialization protocol. The first processor that successfully passes BIST is automatically selected by the hardware as the BSP and starts executing from the reset vector (F000:FFF0h). A processor that does not perform the role of BSP is referred to as an application processor (AP).

The BSP is responsible for executing the BIOS power-on self-test (POST) and preparing the machine to boot the operating system. At boot time, the system is in virtual wire mode and the BSP alone is programmed to accept local interrupts (INTR driven by programmable interrupt controller (PIC) and non-maskable interrupt (NMI)).

As a part of the boot process, the BSP wakes each application processor (AP). When awakened, an AP programs its Memory Type Range Registers (MTRRs) to be identical to those of the BSP. All APs execute a halt instruction with their local interrupts disabled. If the BSP determines that an AP exists that is a lower-featured processor or that has a lower value returned by the CPUID function, the BSP switches to the lowest-featured processor in the system.

5.3.2 Mixed Processor Steppings

For optimum system performance, only identical processors should be installed in a system. Processor steppings can be mixed in a system provided that there is no more than a 1-stepping difference in all processors installed. If the installed processors are more than 1-stepping apart, an error (8080 through 8183) is logged in System Event Log (SEL) and an error (01298000 through 01298003) is reported to the Management Module. Acceptable mixed steppings are not reported as errors.

Note: The error logging format for the Management Module and System Event Logging differ from each other. See the Error Tables to determine which errors are logged in Management Module and which errors are logged as part of System Error Logging.

5.3.3 Mixed Processor Models

Processor models cannot be mixed in a system. If this condition is detected an error is reported and logged in SEL.

5.3.4 Mixed Processor Families

Processor families cannot be mixed in a system. If this condition is detected an error is reported and logged in SEL.

5.3.5 Mixed Processor Cache Sizes

If the installed processors have mixed cache sizes, an error is reported and logged in the SEL. The size of all cache levels must match between all installed processors.

5.3.6 Jumperless Processor Speed Settings

The Intel® Xeon™ processor does not use jumpers or switches to set the processor frequency. The BIOS reads the highest ratio register from all processors in the system. If all processors are the same speed, the Actual Ratio register is programmed with the value read from the High Ratio register. If all processors do not match, the highest common value between High and Low Ratio is determined and programmed to all processors. If there is no value that works for all installed processors, processors not capable of speeds supported by the BSP are disabled and an error is displayed.

5.3.7 Microcode

IA32 processors can correct specific errata through the loading of an Intel-supplied data block (i.e. microcode update). The BIOS is responsible for storing the update in nonvolatile memory and loading it into each processor during POST. The BIOS allows a number of microcode updates to be stored in the flash, limited by the amount of free space available. The BIOS supports variable size microcode updates and it verifies the signature prior to storing the update in the flash. The system BIOS supports the real mode INT15, D042h interface for updating the microcode updates in the flash.

5.3.8 Processor Cache

The BIOS enables all levels of processor cache as early as possible during POST. User options are not available to modify the cache configuration, size or policies. All detected cache sizes are reported in the SMBIOS Type 7 structures. The largest and highest level cache detected is reported in BIOS Setup.

5.3.9 Hyper-Threading Technology

Intel® Xeon™ processors support Hyper-Threading Technology. The BIOS detects processors that support this feature and enables it during POST. BIOS Setup provides an option to selectively enable or disable this feature. The default behavior is enabled.

The BIOS creates additional entries in the ACPI MP tables to describe the virtual processors. The SMBIOS Type 4 structure shows only the physical processors installed. It does not describe the virtual processors.

Because some operating systems are not able to efficiently utilize the Hyper-Threading Technology, the BIOS does not create entries in the MP tables to describe the virtual processors.

5.3.10 Intel SpeedStep® Technology

Intel® Xeon™ processors support the “Geyserville3” (GV3) feature of Intel SpeedStep® Technology. This feature changes the processor operating ratio and voltage similar to the Thermal Monitor 2 (TM2) feature. The E7520 platforms support GV3 feature in conjunction with TM2 feature.

5.3.11 Intel® Extended Memory 64 Technology (Intel® EM64T)

The system BIOS on the Intel® Server Board SE7520AF2 supports the Intel® Extended Memory 64 Technology (Intel® EM64T) capability of the Intel® Xeon™ processors. This capability does not need to be activated or de-activated in BIOS Setup. The system is in IA32 compatibility mode when booting an operating system.

5.4 Memory Initialization

The following section describes the memory initialization performed by the Intel® SE7520AF2 Server BIOS.

5.4.1 Memory Sizing

The Intel® Server Board SE7520AF2 provides eight DIMM sockets for DDR-2 400 MHz registered ECC SDRAM. The maximum local memory capacity is 16GB; memory DIMM sizes supported include: 256MB, 512MB, 1GB and 2GB.

The memory is 2-way interleaved, providing 3.2 GB/sec available peak bandwidth per bus. DIMMs on channel A are paired with DIMMs on channel B to configure 2-way interleaving. The minimum memory configuration is two DIMMs, which requires identical DIMM populated on each channel.

The BIOS reads the Serial Presence Detect (SPD) EEPROMs on each installed memory module to determine the size and timing of the installed memory modules. The memory-sizing algorithm determines the size of each row of DIMMs. The BIOS programs the memory controller in the chipset accordingly. The total amount of configured memory can be found using BIOS Setup.

The DIMM pair, which constitutes interleaving, is referred to as a row. The row can be further divided into two rows, based on single-sided or double-sided DIMMs. If both DIMMs in a pair are single-sided, only one row is said to be present in the system. For double-sided DIMMs, both rows are said to be present.

Memory DIMMs must be populated in pairs. The Intel® Server Board SE7520AF2 has eight DIMM sockets, for four DIMM pairs. Both DIMMs in a pair must be identical (same manufacturer, CAS latency, number of rows, columns and devices, timing parameters etc.). Memory sizing and configuration is guaranteed only for qualified DIMMs approved by Intel.

5.4.2 Disabling DIMMs

The BIOS provides a mechanism to disable a DIMM if it is detected to be faulty. A faulty DIMM is defined to have either multiple correctable errors or a single uncorrectable error. Memory errors are logged during runtime and single bit ECC Errors are counted. Though DIMMs are marked as “Disabled”, they are actually disabled only during the next reboot.

At the next system boot, memory-sizing code reads the recorded state of the DIMMs and skips sizing DIMMs marked as disabled. Because DIMMs are always used in 2-way interleaving, the DIMM pair is disabled. If all DIMMs in a system have been disabled, the BIOS generates beep codes to indicate that the system has no usable memory.

Disabled DIMMs/rows may be re-enabled through a BIOS Setup option. The DIMM slot will no longer be disabled if the system boots without memory in the DIMM slot.

5.4.3 Memory On-Line Sparing and Mirroring

The Intel® E7520 chipset provides RAS features including Memory On-line Sparing and Memory Mirroring. Memory On-line Sparing and Memory Mirroring modes are mutually exclusive. Both modes cannot be supported at the same time.

In Memory On-line Sparing mode, a minimum of four DIMMs are required. Two DIMMs (one DIMM per memory channel constituting a bank) are designated as the spare DIMMs. These spare DIMMs are not included in the system main memory address space. In case of an excessive failure rate on one of the active DIMMs, the content of the failing DIMM and its

corresponding DIMM on other channel is copied to the spare DIMMs, which are put into service. The failing DIMMs are removed from service. The failure rate is preset in the BIOS and cannot be changed. For Memory On-Line Sparing, the spare pair of DIMMs must be equal or larger in size than the other DIMMs in the system.

In Memory Mirroring mode, a minimum of four DIMMs are required to maintain two copies of the memory data at all times. The memory data on one channel is the mirror image of that on the other channel. In this mode, the memory capacity is effectively half of the physical DIMMs being populated.

5.4.4 ECC Memory Initialization

ECC memory must be initialized by the BIOS before it can be used. The BIOS must initialize all memory locations before using them. The BIOS uses the auto-initialize feature of the MCH to initialize ECC. ECC memory initialization cannot be aborted and may result in a noticeable delay in the boot process depending on the amount of memory installed in the system.

5.4.5 Memory Population

The following memory populations are supported:

- Farthest row from MCH first
- Dual rank (d/r) before any single rank (s/r)
- Limited up to four ranks per channel

Supported DIMM populations include:

Table 36. DIMM Population

Bank 4	Bank 3	Bank 2	Bank 1
Empty	Empty	Empty	s/r
Empty	Empty	Empty	d/r
Empty	Empty	s/r	s/r
Empty	Empty	s/r	d/r
Empty	Empty	d/r	d/r
Empty	s/r	s/r	s/r
Empty	s/r	s/r	d/r
s/r	s/r	s/r	s/r

5.4.6 DIMM Module Capacity

Table 37. Module Capacity

Parts	128Mb	256Mb	512Mb	1Gb
X8, single row	128MB	256MB	512MB	1GB
X8, double row	256MB	512MB	1GB	2GB
X4, single row	256MB	512MB	1GB	2GB

Number of DIMMS	Mirror	128Mb	256Mb	512Mb	1Gb
1	No	256MB	512MB	1GB	2GB
2	No	512MB	1GB	2GB	4GB
4	No	1GB	2GB	4GB	8GB
4	Yes	512MB	1GB	2GB	4Gb
6	No	1.5GB	3GB	6GB	12GB
6	Yes	768MB	1.5GB	3GB	6GB
8	No	2GB	4GB	8GB	16GB
8	Yes	1GB	2GB	4GB	7GB

Note: Memory between 4 GB and 4 GB minus 512 MB is not accessible for use by the operating system and may be lost to the user. This area is reserved for BIOS, APIC configuration space, PCI adapter interface, and virtual video memory space. This means that if 4 GB of memory is installed, 3.5 GB of this memory is usable. The chipset should allow the remapping of unused memory above the 4 GB address, but this memory may not be accessible to an operating system that has a 4 GB memory limit.

5.4.7 Memory Error Handling

The chipset detects and corrects single-bit memory errors and detects double-bit memory errors. The chipset supports Intel® x4 single device data correction (x4 SDDC) when in dual channel mode. Both single-bit and double-bit memory errors are reported to baseboard management by the BIOS, which handles SMI events generated by the MCH.

5.4.8 Memory Test

System memory is classified as base and extended memory. Base memory is memory that is required for POST. Extended memory is the remaining memory in the system. Extended memory may be contiguous or may have one or more holes. The BIOS memory test accesses all memory except for memory holes.

Memory test consists of separate base and extended memory tests. The base memory test runs before video is initialized to verify memory required for POST. The BIOS enables video as early as possible during POST to provide a visual indication that the system is functional. After video output has been enabled, the BIOS executes the extended memory test. The status of the extended memory test is displayed on the console.

The extended memory test can be configured through BIOS Setup options. The coverage of the test can be configured to one of the following:

- Test every location (Extensive)
- Test one interleave width per kilo-byte of memory (Sparse)
- Test one interleave width per mega-byte of memory (Quick)

The “interleave width” of a memory subsystem is dependent on the chipset configuration. By default, both the base and extended memory tests are configured to the Disabled setting. The extended memory test can be aborted by pressing the <Space> key during the test.

5.5 PCI Initialization

The following describes the PCI initialization performed by the BIOS.

5.5.1 Scan Order

The BIOS assigns PCI bus numbers in a depth-first hierarchy, in accordance with PCI Specifications. When a bridge device is located, the bus number is incremented in exception of a bridge device in the chipsets. Scanning continues on the secondary side of the bridge until all subordinate busses are assigned numbers. PCI bus number assignments may vary from boot to boot with varying presence of PCI devices with PCI-PCI bridges. If a device with a bridge with a single bus behind it is inserted into a PCI bus, all subsequent PCI bus numbers below the current bus will be increased by one.

The bus assignments occur once, early in the BIOS boot process, and never change during the pre-boot phase.

5.5.2 Resource Assignment

The resource manager assigns the PIC-mode interrupt for the devices that will be accessed by the legacy code. The BIOS configures the PCI Base Address Registers (BAR) and the command register of each device. Code must not make assumptions about the scan order of devices or the order in which resources are allocated to them. The BIOS supports the INT 1Ah PCI BIOS interface calls.

5.5.3 Automatic IRQ Assignment

The BIOS automatically assigns IRQs to devices in the system for legacy compatibility. No method is provided to manually configure the IRQs for devices.

5.5.4 Option ROMs

The BIOS dispatches the option ROMs to available memory space in the address range 0c0000h-0e7fffh. Given the limited space for option ROMs, BIOS allows for disabling of legacy ROM posting via Setup. On-board and per-slot option ROM disable options are available in Setup. The option to disable the onboard video option ROM is not available for obvious reasons.

This option ROM space is also used by the console redirection binary (if enabled) and the user binary (if present and configured for run-time usage).

The BIOS integrate option ROMs for the Intel® 82546GB Gb NIC, the ATI* Rage XL, and the LSI Logic* 53C1030 SCSI controller.

5.5.5 PCI APIs

The system BIOS supports the INT 1Ah, AH = B1h functions as defined in the PCI BIOS Specification. The system BIOS supports the real mode interfaces and does not support the protected mode interfaces.

5.5.6 Split Option ROM

The BIOS supports the split option ROM algorithm according to the *PCI Local Bus Specification 3.0*. Refer to the *PCI Local Bus Specification 3.0* for further information.

5.5.7 Dual Video

The BIOS supports single and dual video modes. The BIOS enables dual video mode by default.

- In single mode (Dual Monitor Video = disabled), the onboard video controller is disabled when an add-in video card is detected.
- In dual mode (Onboard Video = enabled, Dual Monitor Video = enabled), the onboard video controller is enabled and is the primary video device. The external video card is allocated resources and is considered the secondary video device.

5.5.8 PCI Hot Plug Initialization

The following describes the PCI hot plug initialization performed by the BIOS on the SE7520HPAF2 board SKU.

5.5.8.1 PCI Hot-Plug Controller Initialization in Pre-Boot Phase

The hot plug controllers (HPC) come up in an un-initialized state after a power-off or reset. As a result, the hot plug slots come up in an un-powered state. The BIOS initializes the hot plug controllers during the boot process. Hot plug controller initialization involves applying power to the hot plug slots, detecting any error conditions and setting up the speed of the hot plug PCI bus.

The BIOS only initializes root HPCs. A root HPC is an HPC that is reset only when the entire system is reset. Such HPCs can depend upon the BIOS to initialize them.

The PXH functions D0:F0 and D0:F2 implement the PCI Standard Hot Plug Controller Revision 1.0 and includes the PCI-X 2.0 addendum for each PCI Segment. The PXH Hot-plug controller is compatible with the PCI Hot-plug specification, allowing PCI card removal, replacement and addition without powering down the system.

5.5.8.2 PCI Hot-Plug Resource Padding in Pre-Boot Phase

The BIOS over-allocates resources to PCI slots during boot process. This is known as resource padding. The over-allocation is done at the PCI-PCI bridge level. Over-allocation of resources allows a limited number of add-in cards to be hot plugged into a PCI bus without disturbing the allocation to the rest of the buses. The BIOS reserves the following set of resources for unpopulated slots:

- One bus number

- 4 KB I/O
- 8 MB non-prefetchable memory below 4 GB
- 32 MB prefetchable memory below 4 GB

An occupied slot gets only its requested resources.

The direct effect of resource padding is that the total memory allocated by the BIOS to the PCI subsystem increases and main memory available to the operating system in the 0 to 4GB range decreases according to the lower bounds of the memory addresses allocated to PCI.

PCI hot plug specifications define a standard usage model for Hot-plug PCI. The usage model specifies several elements including state indicator LEDs, interlocking switches. The Intel® Server Board SE7520AF2 has state indicator LEDs, but does not have physical interlocking switches. This function is performed by the Hot Plug Driver User Interface.

The BIOS provides ACPI hot-plug methods so a PCI adapter can be added, removed, or replaced without shutting down the system. The hot plug controller follows the Standard Hot-Plug Controller (SHPC) and Subsystem Specification. The user can use hot add, hot remove or hot replace PCI cards via the PCI hot plug driver user interface.

The hardware provides green and amber lights to indicate the status of each PCI slot. The green light indicates that power is applied to that PCI slot. Adapters must not be added or removed from a slot while the green light is on. The amber light is an attention indicator. When the amber light is on, the system has detected a power fault or other error condition on that slot. The ACPI hot plug methods check the frequency and PCI/PCI-X capabilities of the adapter. The ACPI hot plug methods do not allow less capable adapters to be connected to a bus segment that is currently operating in a more capable mode. In this situation, the ACPI methods turn off power to the adapter, turn on the amber light, and notify the operating system of a device malfunction.

The Intel® Server Board SE7520AF2 provides support for PCI Hot-Plug via the prescribed ACPI mechanism, through the use of SCIs and associated control methods via the PCI hot plug driver and GUI for insertion/removal/hot add request events.

The hot-plug feature supports several scenarios:

- Hot Replace: the user can replace a PCI card with an identical card. The replacement card uses the same PCI resources assigned to the previous card. The operating system driver is not updated.
- Hot Remove: the user can remove the existing PCI card without installing a replacement.
- Hot Add: the user can add a new, previously uninstalled adapter card into an empty slot.

5.5.8.3 Operating System Interfaces

5.5.8.3.1 ACPI Control Methods

ACPI control methods require an ACPI hot plug aware operating system. The BIOS-provided ACPI methods support the PCI hot plug driver and GUI.

3.3.8.3.2 PCI Hot-Plug Usage Model

This section describes the hot-plug usage model for both the PCI Express* and PCI-X slots. Each of the four hot-plug slots has a set of status LEDs (for power and attention). The PCI hot plug GUI is used to invoke a hot-plug sequence to remove or add or replace an adapter via a software interface. The status LEDs consist of a green power LED and an amber attention LED.

Table 38. PCI Express* Hot-Plug LED Function Table

	Off	On (Solid)	Blinking
Power LED	All main voltage and data rails are removed from slot. PCI adapter may be added or removed	Normal power to slot. PCI adapter may NOT be added or removed	Slot power in transition (on-off or off-on). PCI adapter may NOT be added or removed
Attention LED	Normal operation	Slot on attention; power fault or operational problem present on the slot	Slot being identified at the user's request to locate the slot via the Hot-Plug Interface Driver

5.5.8.4 Hot Removal Procedure

- Open the chassis cover to see the status LEDs and to access the PCI cards.
- Click “Remove,” then “Confirm” in the PCI hot plug GUI. The operating system safely stops any threads that are dependent on the device, flushes any buffers communicating with the device, and removes device drivers specific for the device. After this, the operating system invokes the appropriate BIOS ACPI method to power off the slot and turn off the power LED.
- Wait for power LED to turn off
- Remove the adapter.

5.5.8.5 Hot Insertion Procedure

- Open the chassis cover to see the status LEDs and to access the PCI cards.
- Ensure the power LED on desired slot is off.
- Install the adapter.
- Click “Hot Add” button for that particular slot in the PCI hot plug GUI.
- Confirm the operation by clicking the “Confirm” control button. The GUI generates an interrupt, which is handled by the BIOS ACPI code. The interrupt handler checks the frequency and mode of the bus and card, and powers on the slot.
- Wait for power LED to turn on. If the attention LED continues to blink, a power fault may have occurred. In this case, remove the adapter, wait for the LED to turn off, and repeat the hot insertion process. A lit attention LED can also indicate a bus/card capability mismatch. In this case, this card cannot be hot inserted into this slot.

When a new card is added to the slot, the ACPI BIOS checks the frequency and mode of the corresponding PCI/PCI-X bus and the capabilities of the adapter card. The ACPI BIOS attempts to match the bus mode and frequency to the maximum capabilities of the card by reprogramming and resetting that individual bus.

5.5.9 PCI Express* Initialization

The following describes the PCI Express* initialization performed by the BIOS.

5.5.10 PCI Express* Enhanced Configuration Mechanisms

The PCI Express* extends the configuration space to 4096 bytes per device/function as compared to 256 bytes allowed by PCI 2.3 configuration space. PCI Express* configuration space is divided into a PCI 2.3 compatible region and an extended PCI-E region.

The Enhanced Configuration memory address map is positioned into memory space by use of the PCI Express* Enhanced Configuration Base register, known as EXPECBASE, which corresponds to E000 0000h. Users must access this space in double byte units.

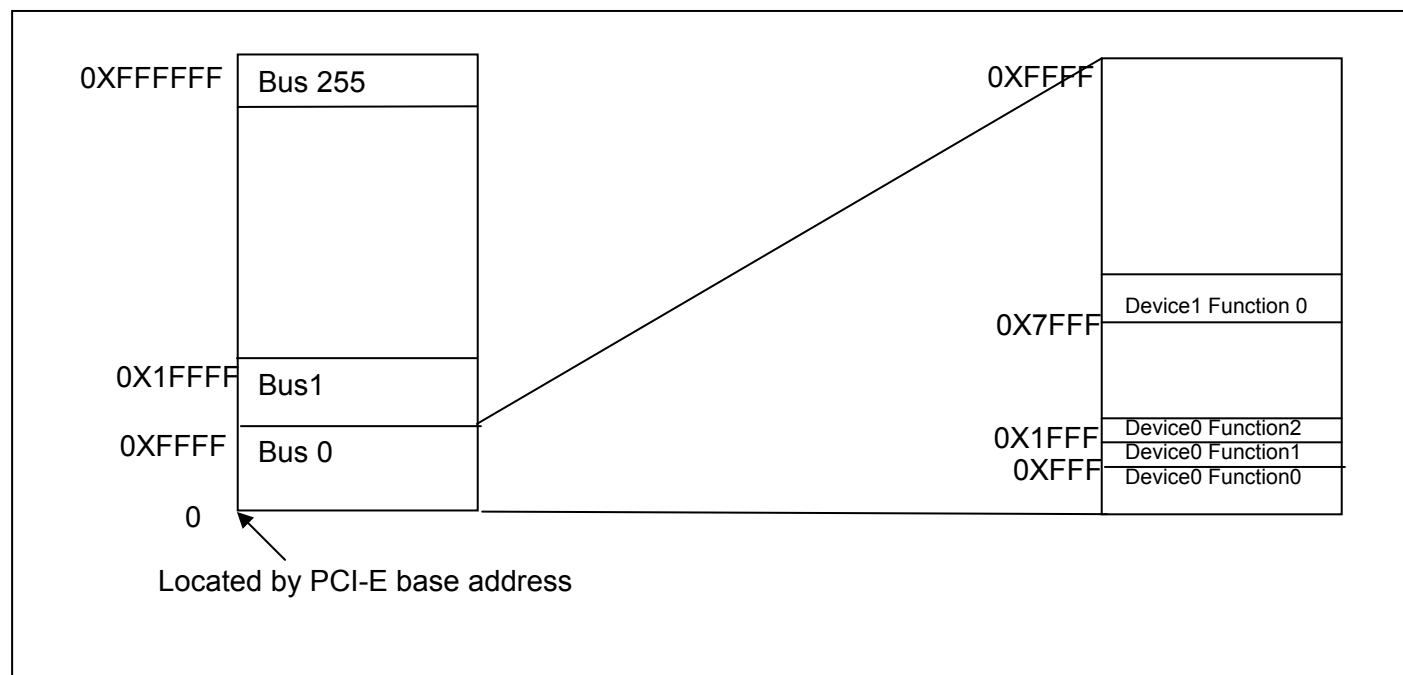


Figure 33. Enhanced Configuration Memory Address Map

5.5.11 Legacy Universal Serial Bus (USB) Initialization

The BIOS supports PS/2* emulation of USB 1.1 keyboards and mice. During POST, the BIOS initializes and configures the root hub ports and then searches for a keyboard and mouse. The USB hub enables them.

5.5.12 IDE Initialization

The BIOS supports the ATA/ATAPI Specification, version 6 or later. The BIOS initializes the embedded IDE controller in the ICH5-R and the IDE devices that are connected to these devices. The BIOS scans the IDE devices and programs the controller and the devices with their optimum timings. The IDE disk read/write services that are provided by the BIOS use PIO mode, but the BIOS programs the necessary Ultra DMA registers in the IDE controller so that the operating system can use the Ultra DMA Modes.

The BIOS initializes and supports ATAPI devices such as LS-120/240, CD-ROM, CD-RW and DVD. The BIOS initializes and supports SATA devices just like P-ATA devices. The BIOS initializes embedded IDE controller in the chipset and SATA devices that are connected to these controllers. From a software standpoint, SATA controllers present the same register interface as the P-ATA controllers. The BIOS utilizes additional SATA capabilities to speed drive detection. Hot plugging of SATA drives during the boot process is not supported by the BIOS and may result in undefined behavior.

5.5.13 Removable Media Initialization

The BIOS supports removable media devices, including 1.44MB floppy removable media devices and optical devices such as a CD-ROM drive or DVD-ROM drive (read only). The BIOS supports booting from USB mass storage devices connected to the chassis USB port, such as a USB key device.

The BIOS supports USB 2.0 media storage devices that are backward compatible to the *Universal Serial Bus Specification, Revision 1.1*.

5.6 Flash ROM

The BIOS supports the Intel® 28F320C3B flash part. The flash part is a 4 MB flash ROM, 2 MB of which is programmable. The flash ROM contains system initialization routines, setup utility, and runtime support routines. The exact layout is subject to change, as determined by Intel. A 128 KB block is available for storing OEM code (user binary) and custom logos.

5.7 BIOS User Interface

Two types of consoles are used to display the user interface: graphical or text based. Graphics consoles are in 640x480 mode; text consoles are 80x25.

Console output is partitioned into three areas: the System Activity/State, Logo/Diagnostic, and Current Activity windows. The System Activity Window displays information about the current state of the system, such as if it is active, hung, or requires user intervention. The Logo/Diagnostic Window displays the OEM splash screen logo or a diagnostic boot screen. The Current Activity Window displays information about the currently executing portion of POST as well as user prompts or status messages.

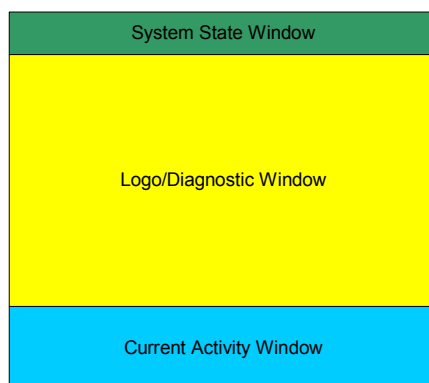


Figure 34. BIOS Screen Display Areas

5.7.1 System State Window

The top row of the screen is reserved for the system state window. On a graphics console, the row is 640x48. On a text console the row is 80x2.

The system state window may be in one of three forms:

- An activity bar that scrolls while the system is busy
- A progress bar that measures percent complete for the current task
- An attention required bar. The attention bar is useful for tasks that require user attention to continue.

5.7.2 Logo/Diagnostic Window

The middle portion of the screen is reserved for the Logo/Diagnostic Window. On a graphics console the screen is 640x384. On a text console the window is 80x20.

The Logo/Diagnostic Window may be in one of two forms: a logo splash screen is displayed in Quiet Boot mode, or a system summary and diagnostic screen is displayed in verbose mode. The default is to display the logo in Quiet Boot mode. If no logo is present in the flash ROM, or if Quiet Boot mode is disabled in the system configuration, the summary and diagnostic screen is displayed. If the user presses <Esc>, the system transfers from the logo screen to the diagnostic screen.

5.7.3 Current Activity Window

The bottom portion of the screen is reserved for the Current Activity Window. On a graphics console the screen is 640x48. On a text console the window is 80x2.

The Current Activity Window is used to display prompts for hot keys and to provide cues to system status.

5.8 System Diagnostic Screen

The diagnostic screen is the console where boot information, options and detected hardware information is displayed.

5.8.1 Static Information Display

This area displays the following information:

- Copyright message
- BIOS ID
- Current processor configuration
- Installed physical memory size

5.9 Quiet Boot / OEM Splash Screen

The BIOS implements Quiet Boot to fulfill the Microsoft* Hardware Design Guide requirement that the BIOS provide minimal startup display during BIOS POST. System start-up must only draw the user's attention in case of errors or when there is a need for user action. By default,

the system must be configured so that the local screen does not display memory counts, device status, and so on. It must present a "clean" BIOS start-up. The only screen display allowed is the OEM splash screen and information like copyright notices.

The Quiet Boot process is controlled by a Setup Quiet-Boot option. If this option is set, the BIOS displays an activity indicator at the top of the screen and a Logo splash screen in the middle section of the screen on the local console. The activity indicator measures POST progress and continues until the Operating System gains control of the system. The splash screen covers up any diagnostic messages in the middle section of the screen. While the logo is displayed on the local console, remote text consoles display diagnostic messages.

Quiet Boot may be disabled by clearing the Setup Quiet-Boot option or by the user pressing the <Esc> key while in Quiet Boot mode. If Quiet Boot is disabled, the BIOS displays diagnostic messages in place of the activity indicator and the splash screen.

The BIOS allows OEMs to override the standard Intel logo with their own logo.

5.10 BIOS Boot Popup Menu

The BIOS Boot Specification (BBS) provides for a Boot Menu Pop-up invoked by pressing the <ESC> key during POST. BBS popup menu can display all available boot devices. The list order in the popup menu is not the same as the boot order in BIOS setup, it simply lists the bootable devices to let the user choose from which to boot.

Please select boot device:
1 st Floppy
Hard Drives
ATAPI CDROM
LAN PXE
EFI Boot Manager
↓ and ↑ to move selection
Enter to select boot device
ESC to boot using defaults

Figure 35. Sample Pop-up Window

5.11 BIOS Setup Utility

The BIOS Setup utility is provided to perform system configuration changes and to display current settings and environment information. The BIOS Setup utility stores configuration settings in system non-volatile storage. Changes affected by BIOS Setup do take effect until the system is rebooted. The BIOS Setup Utility can be accessed through the F2 key when prompted during POST.

5.11.1 Localization

The BIOS Setup utility uses the Unicode standard and is capable of displaying setup forms in English, French, Italian, German, and Spanish. The BIOS supports these languages for console strings as well. The language can be selected using BIOS user interface.

5.11.2 Console Redirection

BIOS Setup is functional via console redirection over various terminal standards emulation. This may limit some functionality for compatibility, such as the use of colors, keys or key sequences, or support of pointing devices.

5.11.3 Configuration Reset

Activating the Clear CMOS jumper produces a “reset system configuration” request. When a request is detected, the BIOS loads the default system configuration values during the next POST.

When the Intel® Management Module is installed it provides two additional methods to issue a “reset system configuration” request. Using the front panel, the user can hold the reset button for 4 seconds and then press the power button while continuing to press the reset button. Alternatively, software can send a specific OEM command to the BMC to indicate the request. In both methods, the BIOS loads the default values during the next POST.

5.11.4 Keyboard Commands

The Keyboard Command Bar supports the following:

Table 39. BIOS Setup Keyboard Command Bar Options

Key	Option	Description
Enter	Execute Command	The Enter key is used to activate sub-menus when the selected feature is a sub-menu, or to display a pick list if a selected option has a value field, or to select a sub-field for multi-valued features like time and date. If a pick list is displayed, the Enter key will undo the pick list, and allow another selection in the parent menu.
ESC	Exit	The ESC key provides a mechanism for backing out of any field. This key will undo the pressing of the Enter key. 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 sub-menu, 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 any settings. If “Yes” is selected and the Enter key is pressed, setup is exited and the BIOS continues with POST.
↑	Select Item	The up arrow is used to select the previous value in a pick list, or the previous options 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 affect 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.

Key	Option	Description			
+	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 keyboard, but will have the same effect			
F9	Setup Defaults	<p>Pressing F9 causes the following to appear:</p> <table><tr><td>Setup Confirmation</td></tr><tr><td>Load default configuration now?</td></tr><tr><td>[Yes] [No]</td></tr></table> <p>If “Yes” is selected and the Enter key is pressed, all Setup fields are set to their default values. 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 F9 was pressed without affecting any existing field values</p>	Setup Confirmation	Load default configuration now?	[Yes] [No]
Setup Confirmation					
Load default configuration now?					
[Yes] [No]					
F10	Save and Exit	<p>Pressing F10 causes the following message to appear:</p> <table><tr><td>Setup Confirmation</td></tr><tr><td>Save Configuration changes and exit now?</td></tr><tr><td>[Yes] [No]</td></tr></table> <p>If “Yes” is selected and the Enter key is pressed, all changes are saved and Setup is exited. If “No” is selected and the Enter key 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.</p>	Setup Confirmation	Save Configuration changes and exit now?	[Yes] [No]
Setup Confirmation					
Save Configuration changes and exit now?					
[Yes] [No]					

5.11.5 Entering BIOS Setup

The BIOS Setup utility is accessed by pressing the <F2> hotkey during POST.

5.11.5.1 Main Menu Selection

The Main Menu is the first screen displayed when entering the BIOS Setup Utility. This screen displays the major menu selections available: The following tables describe the available options on the top-level and lower level menus. Default values are in bold type.

Table 40. BIOS Setup Main Menu Options

Feature	Options	Help Text	Description
System Overview			
Server BIOS			
Version	N/A	N/A	BIOS ID string (excluding the build time and date)
Build Date	N/A	N/A	BIOS build date
Processor			
Type	N/A	N/A	Processor brand ID string
Speed	N/A	N/A	Calculated processor speed

Feature	Options	Help Text	Description
Count	N/A	N/A	Detected number of physical processors
Installed Memory			
Size	N/A	N/A	Amount of physical memory detected
Server Board MCH Stepping			
Stepping	N/A	N/A	Version of the Intel® E7520 MCH (C2 or C4)
System Time			
System Time	HH:MM:SS	Use [ENTER], [TAB] or [SHIFT-TAB] to select a field. Use [+] or [-] to configure system Time.	Configures the system time on a 24 hour clock. Default is 00:00:00
System Date			
System Date	DAY MM/DD/YYYY	Use [ENTER], [TAB] or [SHIFT-TAB] to select a field. Use [+] or [-] to configure system Date.	Configures the system date. Default is [Tue 01/01/2002]. Day of the week is automatically calculated.
Language			
Language	English French German Italian Spanish	Select the current default language used by the BIOS.	Select the current default language used by BIOS.

5.11.5.1.1 Advanced Menu Selection

Table 41. BIOS Setup Advance Menu Options

Feature	Options	Help Text	Description
Processor Summary	N/A	Configure processors.	Selects submenu.
IDE Configuration	N/A	Configure the IDE device(s).	Selects submenu.
Floppy Configuration	N/A	Configure the Floppy drive(s).	Selects submenu.
SuperIO Configuration	N/A	Configure the Super I/O Chipset.	Selects submenu.
Memory Configuration	N/A	Configure memory devices.	Selects submenu.
USB Configuration	N/A	Configure the USB support.	Selects submenu.
PCI Configuration	N/A	Configure PCI devices.	Selects submenu.

5.11.5.1.2 Processor Configuration Sub-Menu

Table 42. BIOS Setup CPU Configuration Menu Options

Feature	Options	Help Text	Description
Manufacturer	Intel	N/A	Displays processor manufacturer string

Feature	Options	Help Text	Description
Brand String	N/A	N/A	Displays processor brand ID string
Frequency	N/A	N/A	Displays the calculated processor speed
FSB Speed	N/A	N/A	Displays the processor front-side bus speed.
CPU 1			
CPUID	N/A	N/A	Displays the CPUID of the processor.
Cache L1	N/A	N/A	Displays cache L1 size.
Cache L2	N/A	N/A	Displays cache L2 size.
Cache L3	N/A	N/A	Displays cache L3 size. Visible only if the processor contains an L3 cache.
CPU 2			
CPUID	N/A	N/A	Displays the CPUID of the processor.
Cache L1	N/A	N/A	Displays cache L1 size.
Cache L2	N/A	N/A	Displays cache L2 size.
Cache L3	N/A	N/A	Displays cache L3 size. Visible only if the processor contains an L3 cache.
Max CPUID Value Limit	Enabled Disabled	This should be enabled in order to boot legacy OSes that cannot support processors with extended CPUID functions.	
CPU TM Function	Enabled Disabled	If enabled, Thermal Monitor 2 (TM2) feature of the Intel Xeon processor are turned on (For CPUs 3.6GHz and above).	
Hardware Prefetcher	Enabled Disabled	Enabling this feature can result on higher performance on some applications and operating systems.	
Adjacent Cache Line Prefetch	Enabled Disabled	Enabling this feature can result on higher performance on some applications and operating systems.	
Hyper-Threading Function	Enabled Disabled	Enable Hyper-Threading Technology only if operating system supports it.	Controls Hyper-Threading state. Primarily used to support older Operating Systems that do not support Hyper Threading.
Intel® Speed Step™ Tech	Automatic Disabled	Select disabled for maximum CPU speed. Select enabled to allow the operating system to reduce power consumption.	.

Feature	Options	Help Text	Description
Processor Retest	Enabled Disabled	If enabled, all processors will be activated and retested on the next boot. This option will be automatically reset to disabled on the next boot.	Rearms the processor sensors. Only displayed if the Intel Management Module is present.

5.11.5.1.3 IDE Configuration Sub-Menu

Table 43. BIOS Setup IDE Configuration Menu Options

Feature	Options	Help Text	Description
Onboard P-ATA Channels	Disabled Primary	Disabled: disables the integrated P-ATA controller. Primary: enables only the primary P-ATA controller.	Controls state of integrated P-ATA controller.
Onboard S-ATA Channels	Disabled Enabled	Disabled: disables the integrated SATA controller. Enabled: enables the integrated SATA controller.	Controls state of integrated SATA controller.
Configure S-ATA as RAID	Disabled Enabled	When enabled the SATA channels are reserved to be used as RAID.	When Enabled, S-ATA Mode is grayed out since the system is in Enhanced mode.
S-ATA Mode	Enhanced Legacy	Legacy mode allows installation of legacy operating systems that do not support Enhanced Mode by allowing the user to emulate the ICH5-R S-ATA ports (S-ATA A1/A2) as P-ATA devices.	When in Legacy Mode and P-ATA Channels is set to Primary, the S-ATA A1 emulates Secondary Master, and S-ATA A2 emulates Secondary Slave (A1-Sec M/A2-Sec S) When in Legacy Mode and P-ATA Channels is set to Disabled, the S-ATA A1 emulates Primary Master, and S-ATA A2 emulates Secondary Master (A1-Pri M/A2-Sec M)
Primary IDE Master	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details.
Primary IDE Slave	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details.
Secondary IDE Master	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details. Option displayed grayed out and listing 'Not Detected' when in Enhanced Mode. In Legacy Mode the S-ATA devices will be displayed here.

Secondary IDE Slave	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details. Option displayed grayed out and listing 'Not Detected' when in Enhanced Mode. In Legacy Mode the S-ATA devices will be displayed here.
S-ATA A1	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details. If in Legacy Mode this S-ATA device will be listed in the Secondary IDE fields above.
S-ATA A2	[Device Details]	While entering setup, BIOS auto detects the presence of IDE devices. This displays the status of auto detection of IDE devices.	Selects submenu with additional device details. If in Legacy Mode this S-ATA device will be listed in the Secondary IDE fields above.
Hard Disk Write Protect	Disabled Enabled	Disable/Enable device write protection. This will be effective only if device is accessed through BIOS.	Primarily used to prevent unauthorized writes to hard drives.
IDE Detect Time Out (Sec)	0 5 10 15 20 25 30 35	Select the time out value for detecting ATA/ATAPI device(s).	Primarily used with older IDE devices with longer spin up times.
ATA(Pi) 80Pin Cable Detection	Host & Device Host Device	Select the mechanism for detecting 80Pin ATA(Pi) Cable.	The 80 pin cable is required for UDMA-66 and above. BIOS detects the cable by querying the host and/or device.

Table 44. BIOS Setup, IDE Device Configuration Sub-menu Selections

Feature	Options	Help Text	Description
Device	N/A	N/A	Display detected device info
Vendor	N/A	N/A.	Display IDE device vendor.
Size	N/A	N/A	Display IDE DISK size.
LBA Mode	N/A	N/A	Display LBA Mode
Block Mode	N/A	N/A	Display Block Mode
PiO Mode	N/A	N/A	Display PiO Mode
Async DMA	N/A	N/A	Display Async DMA mode
Ultra DMA	N/A	N/A	Display Ultra DMA mode.
S.M.A.R.T.	N/A	N/A	Display S.M.A.R.T. support.

Feature	Options	Help Text	Description
Type	Not Installed Auto CDROM ARMD	Select the type of device connected to the system.	The Auto setting will work in most cases.
LBA/Large Mode	Disabled Auto	Disabled: Disables LBA Mode. Auto: Enabled LBA Mode if the device supports it and the device is not already formatted with LBA Mode disabled.	The Auto setting will work in most cases.
Block (Multi-Sector Transfer) Mode	Disabled Auto	Disabled: The Data transfer from and to the device occurs one sector at a time. Auto: The data transfer from and to the device occurs multiple sectors at a time if the device supports it.	The Auto setting will work in most cases.
PIO Mode	Auto 0 1 2 3 4	Select PIO Mode.	The Auto setting will work in most cases.
DMA Mode	Auto SWDMA0-2 MWDMA0-2 UWDMA0-5	Select DMA Mode. Auto :Auto detected SWDMA :SinglewordDMA MWDMA :MultiwordDMA UWDMA :UltraDMA	The Auto setting will work in most cases.
S.M.A.R.T.	Auto Disabled Enabled	Self-Monitoring, Analysis and Reporting Technology.	The Auto setting will work in most cases.
32Bit Data Transfer	Disabled Enabled	Enable/Disable 32-bit Data Transfer	

5.11.5.1.4 Floppy Configuration Sub-Menu

Table 45. BIOS Setup floppy Configuration Sub-menu Selections

Feature	Options	Help Text	Description
Floppy A	Disabled 720 KB 3 1/2" 1.44 MB 3 1/2" 2.88 MB 3 1/2"	Select the type of floppy drive connected to the system.	Note: Intel no longer validates 720Kb and 2.88Mb drives.
Onboard Floppy Controller	Disabled Enabled	Allows BIOS to Enable or Disable Floppy controller.	

5.11.5.1.5 SuperIO Configuration Sub-Menu**Table 46. BIOS Setup SuperIO Configuration Sub-menu**

Feature	Options	Help Text	Description
Serial Port 1 Address	Disabled 3F8/IRQ4 2F8/IRQ3 3E8/IRQ4 2E8/IRQ3	Allows BIOS to Select Serial Port A Base Addresses.	Option that is used by other serial port is hidden to prevent conflicting settings.
Serial Port 2 Address	Disabled 3F8/IRQ4 2F8/IRQ3 3E8/IRQ4 2E8/IRQ3	Allows BIOS to Select Serial Port B Base Addresses.	Option that is used by other serial port is hidden to prevent conflicting settings.

5.11.5.1.6 USB Configuration Sub-Menu**Table 47. BIOS Setup USB Configuration Sub-menu Selections**

Feature	Options	Help Text	Description
USB Devices Enabled	N/A	N/A	List of USB devices detected by BIOS.
USB Function	Disabled Enabled	Enables USB HOST controllers.	When set to disabled, other USB options are grayed out.
Legacy USB Support	Disabled Keyboard only Auto Keyboard and Mouse	Enables support for legacy USB. AUTO option disables legacy support if no USB devices are connected. If disabled, USB Legacy Support will not be disabled until booting an operating system.	
Port 60/64 Emulation	Disabled Enabled	Enables I/O port 60/64h emulation support. This should be enabled for the complete USB keyboard legacy support for non-USB aware OSes.	
USB 2.0 Controller	Enabled Disabled	N/A	
USB 2.0 Controller mode	FullSpeed HiSpeed	Configures the USB 2.0 controller in HiSpeed (480Mbps) or FullSpeed (12Mbps).	
USB Mass Storage Device Configuration	N/A	Configure the USB Mass Storage Class Devices.	Selects submenu with USB Device enable.

5.11.5.1.7 USB Mass-Storage Configuration Sub-Menu**Table 48. BIOS Setup, USB Mass Storage Device Configuration Sub-menu Selections**

Feature	Options	Help Text	Description
USB Mass Storage Reset Delay	10 Sec 20 Sec 30 Sec 40 Sec	Number of seconds POST waits for the USB mass storage device after start unit command.	
Device #1	N/A	N/A	Only displayed if a device is detected, includes a DeviceID string returned by the USB device.
Emulation Type	Auto Floppy Forced FDD Hard Disk CDROM	If Auto, USB devices less than 530MB will be emulated as Floppy and remaining as hard drive. Forced FDD option can be used to force a HDD formatted drive to boot as FDD (Ex. ZIP* drive).	
Device #n	N/A	N/A	Only displayed if a device is detected, includes a DeviceID string returned by the USB device.
Emulation Type	Auto Floppy Forced FDD Hard Disk CDROM	If Auto, USB devices less than 530MB will be emulated as Floppy and remaining as hard drive. Forced FDD option can be used to force a HDD formatted drive to boot as FDD (Ex. ZIP drive).	

5.11.5.1.8 PCI Configuration Sub-Menu**Table 49. BIOS Setup PCI Configuration Sub-menu Selections**

Feature	Options	Help Text	Description
Onboard SCSI	Disabled Enabled		
Onboard SCSI ROM Mode	IS (RAID 0) IM/IME (RAID 1/1e)	IM/IME = Integrated Mirroring/Integrated Mirroring Enhanced IS = Integrated Striping Before changing modes, back up array data and delete existing arrays, if any. Otherwise, loss of all data may occur.	After operating system installation with a selected SCSI RAID mode, only change this mode selection if prepared to rebuild RAID array. Changing the mode could damage current operating system installation on RAID volume. Grayed out if device is disabled.

Feature	Options	Help Text	Description
IOP ROMB (SROMBU42E)	Installed Not Installed	Intel RAID On Motherboard, requires the use of the RAID Activation Key and a RAID Memory or Portable Cache Module. Refer to product documentation for RAID configuration details. Before changing modes, back up array data and delete existing arrays, if any. Otherwise, loss of all data may occur.	Display only field, always grayed out. System BIOS will automatically detect presence of Intel RAID Key.
Onboard Video	Disabled Enabled	Enable/Disable on board VGA controller	
Dual Monitor Video	Enabled Disabled	Select which graphics controller to use as the primary boot device. Enabled selects the on board device.	Grayed out if Onboard video is set to "Disabled"
Onboard NIC	Disabled Enabled		
Onboard NIC ROM	Disabled Enabled		Grayed out if device is disabled.
Slot 1 Option ROM	Disabled Enabled	PCI-X 64/133 (Hot-pluggable on BAF2HPBB SKU)	
Slot 3 Option ROM	Disabled Enabled	PCI Express* x4 (Hot-pluggable on BAF2HPBB SKU)	
Slot 4 Option ROM	Disabled Enabled	PCI Express* x8 (Hot-pluggable on BAF2HPBB SKU)	
Slot 5 Option ROM	Disabled Enabled	PCI-X 64/133 (Hot-pluggable on BAF2HPBB SKU)	
Slot 6 Option ROM	Disabled Enabled	PCI-X 64/133 (Hot-pluggable on BAF2HPBB SKU)	

5.11.5.2 Memory Configuration Sub-menu

This sub-menu provides information about the DIMMs detected by BIOS. The DIMM number is printed on the baseboard next to each device.

Table 50. BIOS Setup, Memory Configuration Sub-menu Selections

Feature	Options	Help Text	Description
DIMM 1A	Installed Not Installed Disabled Mirror Spare		Informational display

Feature	Options	Help Text	Description
DIMM 1B	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 2A	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 2B	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 3A	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 3B	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 4A	Installed Not Installed Disabled Mirror Spare		Informational display.
DIMM 4B	Installed Not Installed Disabled Mirror Spare		Informational display.
Extended Memory Test	1 MB 1 KB Every Location Disabled	Settings for extended memory test	
Memory Retest	Enabled Disabled	If "Enabled", BIOS will activate and retest all DIMMs on the next system boot. This option will automatically reset to "Disabled" on the next system boot.	

Feature	Options	Help Text	Description
Memory Remap Feature	Enabled Disabled	Enable: Allow remapping of overlapped PCI memory above the total physical memory. Disable: Do not allow remapping of memory.	
Memory Mirroring / Sparing	Sparing Mirroring Disabled	Disabled provides the most memory space. Sparing reserves memory to replace failures. Mirroring keeps a second copy of memory contents.	Sparing or Mirroring is grayed out if the installed DIMM configuration does not support it.

5.11.5.3 Boot Menu

Table 51. BIOS Setup, Boot Menu Selections

Feature	Options	Help Text	Description
Boot Settings Configuration	N/A	Configure settings during system boot.	Selects submenu.
Boot Device Priority	N/A	Specifies the boot device priority sequence.	Selects submenu.
Hard Disk Drives	N/A	Specifies the boot device priority sequence from available hard drives.	Selects submenu.
Removable Drives	N/A	Specifies the boot device priority sequence from available removable drives.	Selects submenu.
ATAPI CDROM Drives	N/A	Specifies the boot device priority sequence from available ATAPI CD-ROM drives.	Selects submenu.

5.11.5.3.1 Boot Settings Configuration Sub-menu

Table 52. BIOS Setup, Boot Settings Configuration Sub-menu Selections

Feature	Options	Help Text	Description
Quick Boot	Disabled Enabled	Allows BIOS to skip certain tests while booting. This will decrease the time needed to boot the system.	
Quiet Boot	Disabled Enabled	Disabled: Displays normal POST messages. Enabled: Displays OEM Logo instead of POST messages.	
Bootup Num-Lock	Off On	Select power-on state for numlock.	
PS/2 Mouse Support	Disabled Enabled Auto	Select support for PS/2 mouse.	
POST Error Pause	Disabled Enabled	If enabled, the system will wait for user intervention on critical POST errors. If disabled, the system will boot with no intervention, if possible.	
Hit <F2> Message Display	Disabled Enabled	Displays "Press <F2> to run Setup" in POST.	

Scan User Flash Area	Disabled Enabled	Allows BIOS to scan the Flash ROM for user binaries.	
----------------------	----------------------------	--	--

5.11.5.3.2 Boot Device Priority Sub-menu

Table 53. BIOS Setup, Boot Device Priority Sub-menu Selections

Feature	Options	Help Text	Description
1st Boot Device	Varies	Specifies the boot sequence from the available devices. A device enclosed in parenthesis has been disabled in the corresponding type menu.	Number of entries will vary based on system configuration.
nth Boot Device	Varies	Specifies the boot sequence from the available devices. A device enclosed in parenthesis has been disabled in the corresponding type menu.	

5.11.5.3.3 Hard Disk Drive Sub-menu

Table 54. BIOS Setup, Hard Disk Drive Sub-Menu Selections

Feature	Options	Help Text	Description
1st Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.
nth Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.

5.11.5.3.4 Removable drive sub-menu selections

Table 55. BIOS Setup, Removable Drives Sub-menu Selections

Feature	Options	Help Text	Description
1st Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.
nth Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.

5.11.5.3.5 ATAPI CDROM drives sub-menu selections

Table 56. BIOS Setup, ATAPI CDROM Drives Sub-menu Selections

Feature	Options	Help Text	Description
1st Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.

nth Drive	Varies	Specifies the boot sequence from the available devices.	Varies based on system configuration.
-----------	--------	---	---------------------------------------

5.11.5.4 Security menu

Table 57. BIOS Setup, Security Menu Options

Feature	Options	Help Text	Description
Supervisor Password is	N/A	Install / Not installed	Informational display.
User Password is	N/A	Install / Not installed	Informational display.
Change Supervisor Password	N/A	Set or clear Admin password	Set password to null to clear.
Change User Password	N/A	Set or clear User password	This node is grayed out until Admin password is installed. Set password to null to clear.
User Access Level	No Access View Only Limited Full Access	LIMITED: allows only limited fields to be changed such as Date and Time. NO ACCESS: prevents User access to the Setup Utility. VIEW ONLY: allows access to the Setup Utility but the fields can not be changed. FULL: allows any field to be changed except the Supervisor password.	This node is grayed out and becomes active only when Admin password is set.
Clear User Password	N/A	Immediately clears the user password.	Admin uses this option to clear User password (Admin password is used to enter setup is required). This node is graye if Administrator password is not installed.
Fixed disk boot sector protection	Disabled Enabled	Enable/Disable Boot Sector Virus Protection.	
Password Check	Disabled Enabled	If enabled, requires password entry before boot.	This node is grayed out if a user password is not installed.
Secure Mode Timer	1 minute 2 minutes 5 minutes 10 minutes 20 minutes 60 minutes 120 minutes	Period of key/PS/2 mouse inactivity specified for Secure Mode to activate. A password is required for Secure Mode to function. Has no effect unless at least one password is enabled.	This node is grayed out if a user password is not installed.
Secure Mode Hot Key (Ctrl-Alt-)	[Z] [L]	Key assigned to invoke the secure mode feature. Cannot be enabled unless at least one password is enabled. Can be disabled by entering a new key followed by a backspace or by entering delete.	This node is grayed out if a user password is not installed.

Secure Mode Boot	Disabled Enabled	When enabled, allows the host system to complete the boot process without a password. The keyboard will remain locked until a password is entered. A password is required to boot from diskette.	This node is grayed out if a user password is not installed.
Power Switch Inhibit	Disabled Enabled	Disable the Front Panel Power Switch when Secure mode is activated. A password is required to unlock the system.	This node is grayed out if a user password is not installed. This node is hidden if the Intel Management Module is not present.
Reset Switch Inhibit	Disabled Enabled	Disable the Front Panel Reset Switch when Secure mode is activated. A password is required to unlock the system.	This node is grayed out if a user password is not installed. This node is hidden if the Intel Management Module is not present.
NMI Control	Disabled Enabled	Enable / disable NMI control for the front panel NMI button.	

5.11.5.5 Server menu

Table 58. BIOS Setup, Server Menu Selections

Feature	Options	Help Text	Description
System management	N/A	N/A	Selects submenu.
Serial Console Features	N/A	N/A	Selects submenu.
Event Log configuration	N/A	Configures event logging.	Selects submenu.
Assert NMI on SERR	Disabled Enabled	If enabled, NMI is generated on SERR and logged.	
Assert NMI on PERR	Disabled Enabled	If enabled, NMI is generated. SERR option needs to be enabled to activate this option.	Grayed out if “NMI on SERR” is disabled.
AC Link	Stays Off Power On Last State	Determines the mode of operation if a power loss occurs. Stays off, the system will remain off once power is restored. Power On, boots the system after power is restored.	“Last State” is only displayed if the Intel Management Module is present. When displayed, “Last State” is the default. When set to “Stays Off,” “Power Switch Inhibit” is disabled.
Platform Event Filtering	Enabled Disabled	Disable trigger for system sensor events.	
FRB-2 Policy	Disable BSP Do not disable BSP Retry on Next Boot Disable FRB2 Timer	This controls action if the boot processor will be disabled or not.	“Disable BSP” and “Do not disable BSP” are only displayed if the Intel Management Module is present.

Feature	Options	Help Text	Description
Late POST Timeout	Disabled 5 minutes 10 minutes 15 minutes 20 minutes	This controls the time limit for add-in card detection. The system is reset on timeout.	
Hard Disk OS Boot Timeout	Disabled 5 minutes 10 minutes 15 minutes 20 minutes	This controls the time limit allowed for booting an operating system from a Hard disk drive. The action taken on timeout is determined by the OS Watchdog Timer policy setting.	
PXE OS Boot Timeout	Disabled 5 minutes 10 minutes 15 minutes 20 minutes	This controls the time limit allowed for booting an operating system using PXE boot. The action taken on timeout is determined by OS Watchdog Timer policy setting.	
OS Watchdog Timer Policy	Stay On Reset Power Off	Controls the policy upon timeout. Stay on action will take no overt action. Reset will force the system to reset. Power off will force the system to power off.	

5.11.5.5.1 System management sub-menu selections

Table 59. BIOS Setup, System Management Sub-menu Selections

Feature	Options	Help Text	Description
Server Board Part Number	N/A	N/A	Field contents varies
Server Board Serial Number	N/A	N/A	Field contents varies
NIC 1 MAC Address	N/A	N/A	Field contents varies
NIC 2 MAC Address	N/A	N/A	Field contents varies
System Part Number	N/A	N/A	Field contents varies
System Serial Number	N/A	N/A	Field contents varies
Chassis Part Number	N/A	N/A	Field contents varies
Chassis Serial Number	N/A	N/A	Field contents varies
BIOS Version	N/A	N/A	BIOS ID string (excluding the build time and date).
BMC Device ID	N/A	N/A	Field contents varies
BMC Firmware Revision	N/A	N/A	Field contents varies
BMC Device Revision	N/A	N/A	Field contents varies
PIA Revision	N/A	N/A	Field contents varies
FRUSDR Package Revision	N/A	N/A	Field contents varies
Primary HSBP Revision (HSBP A)	N/A	N/A	Firmware revision of the Hot-swap controller. Displays n/a if the controller is not present.

Secondary HSBP Revision (HSBP B)	N/A	N/A	Firmware revision of the Hot-swap controller. Displays n/a if the controller is not present.
----------------------------------	-----	-----	--

5.11.5.5.2 Serial Console features sub-menu selections

Table 60. BIOS Setup Serial Console Features Sub-menu Selections

Feature	Options	Help Text	Description
BIOS Redirection Port	Disabled Serial A Serial B	If enabled, BIOS uses the specified serial port to redirect the console to a remote ANSI terminal. Enabling this option disables Quiet Boot.	When the Intel Management Module is present, the help text directs the user to select Serial B for Serial Over LAN.
		If enabled, BIOS uses the specified serial port to redirect the console to a remote ANSI terminal. Enabling this option disables Quiet Boot. For Serial Over LAN, select Serial B.	
Baud Rate	9600 19.2K 38.4K 57.6K 115.2K	N/A	
Flow Control	No Flow Control CTS/RTS XON/XOFF CTS/RTS + CD	If enabled, it will use the Flow control selected. CTS/RTS = Hardware XON/XOFF = Software CTS/RTS + CD = Hardware + Carrier Detect for modem use.	
Terminal Type	PC-ANSI VT100+ VT-UTF8	VT100+ selection only works for English as the selected language. VT-UTF8 uses Unicode. PC-ANSI is the standard PC-type terminal.	
ACPI Redirection port	Disabled Serial A Serial B	Enable / Disable the ACPI OS Headless Console Redirection.	

5.11.5.5.3 Event Log configuration sub-menu selections

Table 61. BIOS Setup, Event Log Configuration Sub-menu Selections

Feature	Options	Help Text	Description
Clear All Event Logs	Disabled Enabled	Setting this to Enabled will clear the System Event Log during the next boot.	
BIOS Event Logging	Disabled Enabled	Select enabled to allow logging of BIOS events.	Enables BIOS to log events to the SEL. This option controls BIOS events only.

Critical Event Logging	Disabled Enabled	If enabled, BIOS will detect and log events for system critical errors. Critical errors are fatal to system operation. These errors include PERR, SERR, ECC.	Enable SMM handlers to detect and log events to SEL.
ECC Event Logging	Disabled Enabled	Enables or Disables ECC Event Logging.	Grayed out if "Critical Event Logging" option is disabled.
PCI Error Logging	Disabled Enabled	Enables or Disables PCI Error Logging.	Grayed out if "Critical Event Logging" option is disabled.
PCI Express Error Logging	Disabled Enabled	Enables or Disables PCI Express Error Logging.	Grayed out if "Critical Event Logging" option is disabled.
System Bus Event Logging	Disabled Enabled	Enables or Disables System Bus Event Logging.	Grayed out if "Critical Event Logging" option is disabled.
Hub Interface Event Logging	Disabled Enabled	Enables or Disables Hub Interface Event Logging.	Grayed out if "Critical Event Logging" option is disabled.
Memory Buffer Event Logging	Disabled Enabled	Enables or Disables Memory Buffer Event Logging.	Grayed out if "Critical Event Logging" option is disabled.

5.11.5.6 Exit menu

Table 62. BIOS Setup, Exit Menu Selections

Feature	Options	Help Text
Save Changes and Exit	N/A	Exit system setup after saving the changes. F10 key can be used for this operation.
Discard Changes and Exit	N/A	Exit system setup without saving any changes. ESC key can be used for this operation.
Discard Changes	N/A	Discards changes done so far to any of the setup questions. F7 key can be used for this operation.
Load Optimal Defaults	N/A	Load Setup Default values for all the setup questions. F9 key can be used for this operation.
Load Custom Defaults	N/A	Load custom defaults.
Save Custom Defaults	N/A	Save custom defaults

5.12 Flash Architecture and Flash Update Utility

The flash ROM contains system initialization routines, the BIOS Setup Utility, and runtime support routines. The exact layout is subject to change, as determined by Intel. A 64 KB user block is available for user ROM code or custom logos. The flash ROM also contains initialization code in compressed form for on-board peripherals, like SCSI, NIC and video controllers. The flash ROM also contains support for the Rolling Single Boot BIOS update feature.

The complete ROM is visible, starting at physical address 4 GB minus the size of the flash ROM device. The Flash Memory Update utility loads the BIOS image minus the recovery block to the secondary flash partition, and notifies the BIOS that this image should be used on the next

system re-boot. Because of shadowing, none of the flash blocks are visible at the aliased addresses below 1 MB.

A 16 KB parameter block in the flash ROM is dedicated to storing configuration data that controls the system configuration (ESCD). Application software must use standard APIs to access these areas; application software cannot access the data directly.

5.13 Rolling BIOS and On-line Updates

The Online Update nomenclature refers to the ability to update the BIOS while the server is online, in operation, as opposed to having to put the server out of operation while doing a BIOS update. The Rolling BIOS nomenclature refers to the capability for having two copies of BIOS, via the one in use, and the other to which an updated BIOS version can be written. When ready, the system can roll forward to the new BIOS. In case of a failure with the new version, the system can roll back to the previous version.

While the exact nature of hardware changes for the support of on-line update / rolling BIOS are out of scope of this document, BIOS relies on specialized hardware and additional flash space for this. Flash is divided into two partitions, namely primary and secondary. The active partition from which the system boots shall be referred to as primary partition. The AMI FLASH update suite and Intel On-line updates preserve the existing BIOS image on the primary partition. BIOS updates are diverted to the secondary partition. After the update, a notification flag will be set. During the subsequent boot following BIOS update, system will continue to attempt to boot from primary BIOS partition. On determining that a BIOS update occurred in the previous boot, system will attempt to boot the new BIOS. If a failure happens, specialized hardware will switch back to the BIOS on the other partition, thus affecting a “Roll Back”.

The Rolling One Boot Update feature applies to all update mechanisms discussed below.

5.14 Flash Update Utility

Server platforms will support a DOS based update utility. This utility will load a fresh copy of the BIOS into the flash ROM.

The BIOS update may update the following items:

- The system BIOS including the recovery code, setup utility and strings.
- On board video BIOS, SCSI BIOS, and other option ROMS for the devices embedded on the server board.
- OEM Binary Area.
- Microcode Updates.

5.14.1 Flash BIOS

The BIOS flash utility is compatible with DOS, Microsoft Windows (2000/2003), and LINUX operating environments.

The afuXXX AMI Firmware Update Utilities are required for BIOS updates.

5.14.1.1 In DOS

- The flash bootable disk must have ROM image and AFUDOS.EXE.
- Enter in DOS.
- Run AFUDOS /i<ROM filename> [/n] [/p[b][n][c]].

5.14.1.2 In Microsoft* Windows*

- The flash disk must have ROM image, AMIFLDRV.SYS and AFUWIN.EXE.
- Log in to the operating system.
- Run command AFUWIN /i<ROM filename> [/n] [/p[b][n][c]].

5.14.1.3 In LINUX

- The flash disk must have ROM image and AFULNX.
- Log in to the operating system and include floppy device.
- Run command ./afulnx /i<ROM filename> [/n] [/p[b][n][c]].

5.14.1.4 In EFI Shell

- The flash disk must have ROM image and AFUEFI.
- Boot to EFI Shell with the flash disk.
- Do a map -r to obtain the file system on the disk.
- Change drive to the flash disk. E.g., if the flash disk is fs0:, type fs0: at the prompt.
- Run command afuefi [/n] [/p[b][n][c]] <ROM filename> to perform the update.

5.14.1.5 AFU Utility Switches

The AFUxxx utilities must be used on the following format: 'afuXXX /i<ROM filename> [/n] [/p[b][n][c]] [/r<registry_path>] [/s] [/k] [/q] [/h]'; where:

/n - don't check ROM ID
/pbnc -
 b - Program Boot Block
 n - Program NVRAM
 c - Destroy System CMOS
/r - registry path to store result of operation (only for Windows version)
/k - Program non-critical block only
/s - leave signature in BIOS
/q - silent execution
/h - print help

5.14.2 User Binary Area

The baseboard includes an area in flash for implementation-specific OEM add-ons. The OEM binary area can be updated as part of the System BIOS update or it can be updated independently of the System BIOS.

The command line usage for the utility is as follows:

```
UBinD </R> or </I> or </D> [/M<ModID>] /F<RomFileName> /B<NewUserBinaryFileName>  
[ /N<NewRomFileName>] [/O<NCB>]
```

</R> - is used to replace user binary module

</I> - is used to insert user binary module

</D> - is used to delete user binary module from the ROM file.

</?> - is used to display this help.

/M<ModID> - is hexadecimal user binary module ID; Default ModID = 0xF0.

/O<NCB> - is 0-based index of the non-critical block number calculated from the start of the ROM file. Default NCB = 1, used only with insert option. See ROMInfo for reference.

</N<NewRomFileName> - if this option is not included, the ROM is saved with the same name.

5.14.3 Recovery Mode

Three conditions can cause the system to enter recovery mode. Pressing a hot-key, setting the recovery jumper (J1D1, labeled RCVR BOOT) to pins 2-3, and damage to the ROM image will cause the system to enter recovery and update System ROM without the bootblock.

5.14.3.1 BIOS Recovery

The BIOS has its ROM image size of 2 MB. A standard 1.44 MB floppy diskette cannot hold the ROM file due to the larger file size.

The SE7520AF2 BIOS supports Rolling BIOS (see Rolling BIOS and On-line Updates section above for details) and contains a primary and secondary partition. The recovery process performs an update on the secondary partition in the same fashion that the normal flash update process updates the secondary partition. After recovery is complete and the power is cycled to the system, the BIOS partitions will switch and the code executing POST will be the code that was just flashed from the recovery media. The BIOS is made up of a bootblock recovery section, a main BIOS section, an OEM logo/user binary section, and an NVRAM section. The NVRAM section will either be preserved or destroyed based on hot-key press during invocation of the recovery. All the other sections of the secondary BIOS will be updated during the recovery process. If an OEM wishes to preserve the OEM section across an update, it is recommended that the OEM modify the provided AMIBOOT.ROM file with the user binary or OEM logo tools before performing the recovery.

BIOS recovery can be used in one of the following devices: a USB Disk-On-Key, ATAPI CD-ROM/DVD, ATAPI ZIP drives, or a LS-120/LS-240 removable drive. The recovery disk must include the BIOS image file AMIBOOT.ROM.

The recovery mode procedure is as follows:

1. Insert or plug-in a Recovery media with AMIBOOT.ROM
2. Power on system, and when progress code E9 is displayed on port 80h, the system will detect the disk (if there is no image file present, the system will cycle through progress code F1 to EF).
3. When F3 is displayed on port 80h the system will read the BIOS image file.
4. The screen will display flash progress and show if NVRAM and CMOS have been destroyed.
5. When recovery mode is complete, the system will halt and the system can be powered off.

NOTE: One of three different hot-keys that can be invoked:

<Ctrl+Home> - Recovery with CMOS destroyed and NVRAM preserved.

<Ctrl+PageDown> - Recovery with both CMOS and NVRAM preserved.

<Ctrl+PageUp> - Recovery with both CMOS and NVRAM destroyed.

5.14.4 Update OEM Logo

BIOS can change OEMlogo in DOS, Microsoft Windows* (2000/2003). The OEMlogo in the ROM can be changed with utilities tool, then update OEMlogo through flashing the ROM.

5.14.4.1 In DOS

- Boot to the operating system.
- Add OEMLOGOD.exe, Rombuild.exe, RomFile ,and NewOEMlogolImage on harddisk.
- Run “OEMLogoD <RomFileName> <NewOEMImageFileName> [/F or /FN or /N]”

5.14.4.2 In Microsoft Windows*

- Boot to the operating system.
- Add OEMLOGO.exe, Rombuild.exe, RomFile ,and NewOEMlogolImage on harddisk.
- Run command prompt “OEMLogo <RomFileName> <NewOEMImageFileName> [/F or /FN or /N]”

Usage:

OEMLogo <RomFileName> <NewOEMImageFileName> [/F or /FN or /N] or OEMLogo <RomFileName> [/D]

Supported formats (depending on the ROM):

- 16 color BMP, size up to 640x480, even width
- 256 color BMP, 640x480

- JPEG, 640x480, 800x600, or 1024x768
- 256 color PCX, 640x480

[/F] - is used to force replacement of OEM logo even if logo formats do not match.

[/N] - is used to insert 16 color BMP without converting it to AMI Format.

[/FN] - is used to force replacement of OEM logo, without converting 16 color BMP to AMI Format.

[/D] - is used to delete logo module from the ROM file.

NOTE: The Rombuild.exe is different for DOS and Microsoft* Windows*. The user must use the correct Rombuild.exe in DOS or WIN for updating the OEM logo.

5.15 OEM Binary

The Intel® Server Board SE7520AF2 provides users with 16 KB of code and data for use during POST and at run-time. User binary code is executed at several defined hook points within the POST.

The user binary code is stored in the system flash. If no run-time code is added, the BIOS temporarily allocates a code buffer according to [PMM]. If run-time code is present, the BIOS shadows the entire block as though it were an option ROM. The BIOS leaves this region writeable to allow the user binary to update any data structures it defines. System software can locate a run-time user binary by searching for it like an option ROM, checking each 2KB boundary from C0000h to EFFFFh. The system vendor can place a signature within the user binary to distinguish it from other option ROMs.

Intel provides the tools and reference code to help users build their own user binary. The user binary must adhere to the following requirements:

- In order to be recognized by the BIOS and protected from runtime memory managers, the user binary must have an option ROM header (55AA, size).
- The system BIOS performs a scan of the user binary area at predefined points during POST. Mask bits must be set within the user binary to inform the BIOS if an entry point exists for a given time during POST.
- The system state must be preserved by the user binary.
- User binary code must be relocatable. It will be located within the first Megabyte. The user binary code should not make any assumptions about the value of the code segment.
- User binary code will always be executed from RAM and never from flash.
- The code in user binary should not hook critical interrupts, should not reprogram the chipset and should not take any action that affects the correct functioning of the system BIOS.

The BIOS copies the user binary into system memory before the first scan point. If the user binary reports that it does not contain runtime code, it is located in conventional memory (0 - 640 KB).

Reporting that the user binary is POST has only the advantage that it does not use up limited option ROM space, and more option ROMs can be fitted. If user binary code is required at run-time, it is copied to option ROM space. At each scan-point during POST, the system BIOS determines if this scan-point has a corresponding user binary entry point to which it transfers control.

To determine this, the bitmap at byte 4 of the header is tested against the current mask bit that has been determined / defined by the scan point. If the bitmap has the appropriate bit set, the mask is placed in AL and execution is passed to the address computed by (ADR(Byte 5)+5*scan sequence #).

During execution, the user binary may access 11 bytes of Extended BIOS Data Area RAM (EBDA). The segment of the EBDA can be found at address 40:0e. Offset 18 to offset 21h is available for the user binary. The BIOS also reserves eight CMOS bits for the user binary. These bits are in an unchecksummed region of CMOS with default values of zero, and will always be located in the first bank of CMOS. These bits are contiguous, but are not in a fixed location. Upon entry into the user binary, DX contains a 'token' that points to the reserved bits.

This token has the following format:

MSB												LSB			
15				12	11										0
# of bits available -1					Bit offset from start of CMOS of first bit										

The most significant 4 bits are equal to the number of CMOS bits available, minus 1. This field is equal to 7 since 8 CMOS bits are available. The 12 least significant bits define the position of the CMOS bit in RTC. This is a bit address, not a byte address. The CMOS byte location is 1/8th of the 12-bit number, and the remainder is the starting bit position within that byte. For example, if the 12-bit number is 0109h, user binary can use bit 1 of CMOS byte 0108h/8 or 021h.

The following code fragment shows the header and format for a user binary:

```

db    55h, 0AAh, 20h                ; 20h = 8KB USER Area. 40h=16KB.

MyCode    PROC FAR                ; MUST be a FAR procedure
db    CBh                ; Far return instruction
db    04h                ; Bit map to define call points, a 1 in any bit
                        ; specifies that the BIOS is called at that scan
                        ; point in POST

db    CBh                ; First transfer address used to point

```

; to user binary extension structure

```
dw    ?           ; Word Pointer to extension structure
dw    0           ; Reserved
```

; This is a list of 7 transfer addresses, one for each bit in the bitmap.

; 5 Bytes must be used for each.

```
JMP   ErrRet
```

```
JMP   ErrRet
```

```
JMP   Start           ; JMP to maintain proper offset for each entry.
```

; Unused entry JMP's should be filled with 5 byte

; filler or JMP to a RETF

```
JMP   ErrRet
```

```
JMP   ErrRet
```

```
JMP   ErrRet
```

```
JMP   ErrRet
```

Start:

5.15.1 Scan Point Definitions

Table 63. Scan Point Definitions defines the bitmap for each scan point, indicating when the scan point occurs and which resources are available (RAM, stack, binary data area, video, keyboard).

Table 63. Scan Point Definitions

Scan Point	Mask	RAM/ Stack/ BDA	Video/ Keyboard
Near the pointer to the user binary extension structure. The mask bit is 0 if this structure is not present. Instead of a jump instruction, the scan address (offset 5) contains a 0CB followed by a near pointer.	01h	N/A	N/A
Obsolete, no action taken.	02h	N/A	N/A
This scan occurs immediately <u>after</u> video initialization.	04h	Yes	Yes
This scan occurs immediately <u>before</u> video initialization	08h	Yes	No
This scan occurs on POST error. On entry, BX contains the number of the POST error	10h	Yes	Yes
This final scan occurs immediately <u>prior</u> to the INT 19 for normal boot and allows one to completely circumvent the normal INT 19 boot if desired.	20h	Yes	Yes
This scan occurs immediately <u>before</u> the normal option ROM scan.	40h	Yes	Yes
This scan occurs immediately <u>following</u> the option ROM area scan.	80h	Yes	Yes

5.15.2 Format of the OEM Binary Structure

If this structure is not present, that is, if the mask bit 01 is not set, the system BIOS assumes that the user binary is not mandatory, and it is required in runtime.

Table 64: Format of OEM Binary Structure

Offset	Bit Definition
0h	Bit 0 1 if mandatory user binary, 0 if not mandatory. If a user binary is mandatory, it will always be executed. If a platform supports a disabling of the user binary scan through CU, this bit will override CU setting. Bit 1 1 If runtime presence required (other than SMM user binary portion, SMM user binary will always be present in runtime irrespective of setting of this bit) 0, if not required in runtime, and can be discarded at boot time. Bit 7:2 reserved for future expansion
1 - 0Fh	Reserved for future expansion

5.16 Quiet Boot / OEM Splash Screen

The BIOS implements Quiet Boot to fulfill the Hardware Design Guidelines 3.0 requirement that the BIOS provide minimal startup display during BIOS POST. System start-up must only draw the end user's attention in case of errors or when there is a need for user action. By default, the system must be configured so that the local screen does not display memory counts, device status, and so on. It must present a "clean" BIOS start-up. The only screen display allowed is the OEM splash screen and information such as copyright notices.

The Quiet Boot process is controlled by a Setup Quiet-Boot option. If this option is set, BIOS displays an activity indicator at the top of the screen and a Logo splash screen in the middle section of the screen on the local console. The activity indicator measures POST progress and continues until the Operating System gains control of the system. The splash screen covers up any diagnostic messages in the middle section of the screen. While the logo is being displayed on the local console, remote text consoles display diagnostics messages.

Quiet Boot may be disabled by clearing the Setup Quiet-Boot option or by the user pressing the <Esc> key while in Quiet Boot mode. If Quiet Boot is disabled, BIOS displays diagnostic messages in place of the activity indicator and the splash screen.

BIOS allows users to override the standard Intel logo with their own logo.

5.17 Operating System Boot, Sleep, and Wake

5.17.1 Microsoft* Windows* Compatibility

Intel Corporation and Microsoft* Corporation co-author design guides for system designers using Intel® processors and Microsoft* operating systems. These documents are updated yearly to address new requirements and current trends.

PC200x specifications are intended for systems that are designed to work with Windows 2000 and Windows* XP class operating systems. The *Microsoft Hardware Design Guide* for the Windows* XP platform is intended for systems that are designed to work with Windows XP class operating systems. Each specification classifies the systems further and has requirements

based on the intended usage for that system. For example, a server system that will be used in small home/office environments has different requirements than the one used for enterprise applications.

The Intel® Server Board SE7520AF2 supports *Microsoft Hardware Design Guide 3.0*.

5.17.2 Advanced Configuration and Power Interface (ACPI)

The Intel® SE7520AF2 BIOS is ACPI-compliant. The primary role of the BIOS is to provide ACPI tables. The POST creates the ACPI tables and locates them in extended memory (above 1MB). The location of these tables is conveyed to the ACPI-aware operating system through a series of tables located throughout memory. The format and location of these tables is documented in the publicly available ACPI specification.

To prevent conflicts with a non-ACPI-aware operating system, the memory used for the ACPI tables is marked as “reserved” in the INT 15h, function E820h.

As described in the ACPI specification, an ACPI-aware operating system generates an SMI to request that the system be switched into ACPI mode. The BIOS responds by setting up all system (chipset) specific configuration required to support ACPI, issues the appropriate command to the mBMC/BMC to enable ACPI mode and sets the SCI_EN bit as defined by the ACPI specification. The system automatically returns to legacy mode on hard reset or power-on reset.

ACPI has three runtime components:

- **ACPI Tables:** These tables describe the interfaces to the hardware. ACPI tables can make use of a p-code type of language, the interpretation of which is performed by the operating system. The operating system contains and uses an ACPI Machine Language (AML) interpreter that executes procedures encoded in AML and stored in the ACPI tables. AML is a compact, tokenized, abstract machine language. The tables contain information about power management capabilities of the system, APICs, bus structure. The tables also describe control methods that operating system uses to change PCI interrupt routing, control legacy devices in Super I/O, and find out the cause of the wake event, and handle PCI hot plug if applicable.
- **ACPI Registers:** The constrained part of the hardware interface, described (at least in location) by the ACPI tables.
- **ACPI BIOS:** This is the code that boots the machine and implements interfaces for sleep, wake, and some restart operations. The ACPI Description Tables are also provided by the ACPI BIOS.

All IA32 server platforms support S0, S4, and S5 states. In addition to these, the Intel® Server Board SE7520AF2 also supports the S1 state. S1 and S4 are considered sleep states. The ACPI specification defines the sleep states and requires the system to support at least one of them.

While entering the S4 state, the operating system saves the context to the disk and most of the system is powered off. The system can wake on a power button press, or a signal received from a wake-on-LAN compliant LAN card (or on-board LAN), modem ring, PCI power management interrupt, or RTC alarm. The BIOS performs complete POST upon wake up from S4, and initializes the platform.

The system can wake from the S1 state using a PS/2 keyboard, mouse, and USB device in addition to the sources described above.

The wake-up sources are enabled by the ACPI operating systems with co-operation from the drivers; the BIOS has no direct control over the wakeup sources when an ACPI operating system is loaded. The role of the BIOS is limited to describing the wakeup sources to the operating system and controlling secondary control/status bits via DSDT table.

S5 state is equivalent to operating system shutdown. No system context is saved.

5.17.3 Sleep and Wake Functionality

The BIOS supports up to four front panel buttons: the power button, the reset button, the sleep button, and the NMI button¹. The NMI button is a recessed button and may not be accessible on all front panel designs.

The power button is a request that is forwarded by the BMC to the ACPI power state machines in the chipset. It is monitored by the BMC and does not directly control power on the power supply.

The BIOS supports a front panel NMI button. The NMI button is a request that causes the BMC to generate an NMI (non-maskable interrupt). The operating system is responsible for handling the NMI core dump.

The power button behaves differently depending on whether the operating system supports ACPI. If the operating system supports ACPI the power button can be configured as a sleep button. The operating system causes the system to transition to the appropriate system state depending on the user settings.

5.17.4 On to Off (operating system-Present) or operating system to Sleep

If an operating system is loaded, the power button switch generates a request via the system control interrupt (SCI) to the operating system to shutdown the system. The operating system retains control of the system and operating system policy determines if the system transitions into S4/S1 or shuts down.

5.17.5 Sleep to On (ACPI)

If an operating system is loaded, the power button or Wake on LAN (WOL) can generate a wake event to the ACPI chipset and a request (via SCI). If system wakes up from S4/S5, the BIOS POST is completed and then control is given to operating system to wake up the system. If wakeup is from S1, the operating system will wake up the system.

5.17.6 System Sleep States

The platform supports the following ACPI System Sleep States:

- ACPI S0 (working) state
- ACPI S1 (sleep) state

¹ Rack mount configurations may have a fifth button (ID) which is supported by hardware.

- ACPI S4 (suspend to disk) state
- ACPI S5 (soft-off) state

Table 65. Supported Wake Events

Wake Event	Supported via ACPI (by sleep state)	Supported Via Legacy Wake
Power Button	Always wakes system	Always wakes system
Ring indicate from Serial A	Wakes from S1 and S4	Yes
Ring indicate from Serial B	Wakes from S1 and S4. If Serial B (COM2) is used for Emergency Management Port, Serial B wakeup is disabled.	Yes
PME from PCI cards	Wakes from S1 and S4.	Yes
WOL	Wakes from S1 and S4.	No
RTC Alarm	Wakes from S1 and S4.	No
Mouse	Wakes from S1	No
Keyboard	Wakes from S1	No
USB	Wakes from S1	No

6. Platform Management

The Intel® Server Board SE7520AF2 has been designed to support different levels of platform management: Onboard Platform Instrumentation, optional Intel® Management Module – Professional Edition and Intel® Management Module – Advanced Edition.

Integrated onto the baseboard is a “mini” Baseboard Management Controller (mBMC); the mBMC enables entry level platform management features referred to as Onboard Platform Instrumentation. In addition, the baseboard provides an Intel® Management Module Connector (IMM) which offers the option to upgrade to the Professional or Advanced Editions of Server management with one of two optional Intel® Management Modules (IMM).

This chapter will provide an overview of the integrated platform management architecture as well as details of the features and functionality of the onboard platform instrumentation. Refer to the Technical Product Specification for the Intel® Management Modules for a description of their features and functionality.

6.1 Management Architecture Overview

6.1.1 Tiered Server Management Model

To provide the flexibility of different management features, the baseboard supports a three tiered server management model. The features offered with the second tier build upon the features of the first, and the features of the third tier build upon the features of the second.

- **Tier 1 – Onboard Platform Instrumentation (Default).** Essential server management functions are provided by default and are built into the baseboard. This tier is *Intelligent Platform Management Interface Specification, v1.5* compliant, although some functionality may be implemented in a manner different from the Professional and Advanced management models. These functions are provided by a combination of BIOS and the National Semiconductor* PC87431x Mini Baseboard Management Controller (mBMC).
- **Tier 2 – Professional (Optional).** The Professional management model utilizes the feature set of a fully *Intelligent Platform Management Interface Specification, v2.0* compliant Sahalee Baseboard Management Controller (BMC). On the Intel® Server Board SE7520AF2 the Sahalee BMC is located on an optionally installed Intel® Management Module (IMM) that plugs into a dedicated server management connector on the baseboard. When an IMM is installed, the mBMC is automatically converted from an autonomous controller to an I2C based I/O device, allowing the Sahalee BMC to completely manage the system.
- **Tier 3 – Advanced (Optional).** Like the Professional management model, the Advanced management model is implemented using yet an alternate IMM which has a Sahalee BMC integrated onto it. In addition to the features provided with the Professional management module, the Advanced module supports a number of network based Out Of Band (OOB) management capabilities including the availability of a dedicated management NIC for faster network access and support of a full TCP/IP software stack.

The following tables provide an overview of the features supported with each of the three management tiers:

Table 66. Tiered Platform Management Feature Overview

Element	Onboard Platform Instrumentation	Professional	Advanced
IPMI Messaging, Commands, and Abstractions	Yes	Yes	Yes
Baseboard Management Controller (BMC)	Yes	Yes	Yes
Sensors	Limited	Yes	Yes
Sensor Data Records (SDRs) and SDR Repository	Limited	Yes	Yes
FRU Information	Limited	Yes	Yes
Autonomous Event Logging	Yes	Yes	Yes
System Event Log (SEL) Entries	92	3276	3276
BMC Watchdog Timer, covering BIOS and Run-Time software	Limited	Yes	Yes
IPMI Channels, and Sessions	Limited	Yes	Yes
EMP (Emergency Management Port) - IPMI Messaging over Serial/Modem. This feature is also referred to as DPC (Direct Platform Control) over serial/modem.	No	Yes	Yes
Serial/Modem Paging	No	Yes	Yes
Serial/Modem Alerting over PPP using the Platform Event Trap (PET) format	No	Yes	Yes
DPC (Direct Platform Control) - IPMI Messaging over LAN (available via both on-board network controllers)	Yes	Yes	Yes
LAN Alerting using PET	Yes	Yes	Yes
Platform Event Filtering (PEF)	Yes	Yes	Yes
ICMB (Intelligent Chassis Management Bus) - IPMI Messaging between chassis	No	Yes	Yes
PCI SMBus support	No	Yes	Yes
Fault Resilient Booting	Limited	Yes	Yes
Magic Packet and Wake On LAN (WOL) / Power On LAN support	Yes	Yes	Yes
BIOS logging of POST progress and POST errors	Errors Only	Yes	Yes
Integration with BIOS console redirection via IPMI v2.0 Serial Port Sharing	No	Yes	Yes
Wake On Ring (WOR) support	No	Yes	Yes
Access via web browser	No	No	Yes
SNMP access (OOB)	No	No	Yes
Telnet access (OOB)	No	No	Yes
Alerting via Email (OOB operation)	No	No	Yes
Keyboard/Video/Mouse (KVM) redirection via LAN	No	No	Yes
High-speed access to dedicated NIC	No	No	Yes

Table 67. Power and Reset Control

Source	Power Cycle		Power Up		Power Down		Hard Reset	
	OPI	Pro/Adv	OPI	Pro/Adv	OPI	Pro/Adv	OPI	Pro/Adv
DPC (Serial)	No	Yes	No	Yes	No	Yes	No	Yes
DPC (LAN)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AC Power Restore	No	No	Yes	Yes	Yes	Yes	N/A	N/A
IPMB	No	Yes	No	Yes	No	Yes	No	Yes
ICMB	No	Yes	No	Yes	No	Yes	No	Yes
PCI SMBus	No	Yes	No	Yes	No	Yes	No	Yes
System Interface	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Watchdog Timer Expiration	Yes	Yes	N/A	N/A	Yes	Yes	Yes	Yes
PEF Event	Yes	Yes	N/A	N/A	Yes	Yes	Yes	Yes
Front Panel Button	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Wake on LAN1	N/A	N/A	Yes	Yes	N/A	N/A	N/A	N/A
Modem Ring Indicate ¹	N/A	N/A	Yes	Yes	N/A	N/A	N/A	N/A
'Time of Day' request from system RTC ¹	N/A	N/A	Yes	Yes	N/A	N/A	N/A	N/A
ACPI / OS Power Down	N/A	N/A	N/A	N/A	Yes	Yes	N/A	N/A
FRB-3 Timeout	N/A	N/A	N/A	N/A	N/A	N/A	Yes	Yes

Table 68. Secure Mode Button Actions

	ACPI State	Power Switch	Sleep Switch	Reset Switch	NMI Switch	ID Switch
Professional and Advanced	S0 On	Protected – No Action	Protected – No Action	Protected – No Action	Unprotected	Unprotected
	S1 Sleep	Unprotected- Wakes Server	Unprotected	Protected – No Action	Unprotected	Unprotected
	S4/S5 Off	Unprotected – Powers On	Unprotected	Unprotected	Unprotected	Unprotected
Onboard Platform Instrumentation	S0 On	Protected – No Action	Unprotected	Protected – No Action	Unprotected	Unprotected
	S1 Sleep	Protected – No Action	Unprotected	Protected – No Action	Unprotected	Unprotected
	S4/S5 Off	Protected – No Action	Unprotected	Protected – No Action	Unprotected	Unprotected

Table 69. Memory RAS Feature Support by SM Tier

Memory RAS Feature	Onboard Platform	Professional	Advanced
--------------------	------------------	--------------	----------

¹ Via the chip set asserting (for power down) or deasserting (for power up) SLEEP_S5

	Instrumentation		
Inventory	No	Yes	Yes
Correctable Error Reporting	No	Yes	Yes
Uncorrectable Error Reporting	Yes	Yes	Yes
DIMM Sparing	Partial ¹	Yes	Yes
DIMM Mirroring	Partial ¹	Yes	Yes

¹ – No SEL logging.

The following diagram shows a Logical Block Diagram of the platform management architecture implemented on the Server Board SE7520AF2.

Note: The interconnections and blocks shown are to illustrate the functional relationships between the system management elements, and do not map directly to the exact circuit implementation of the architecture.

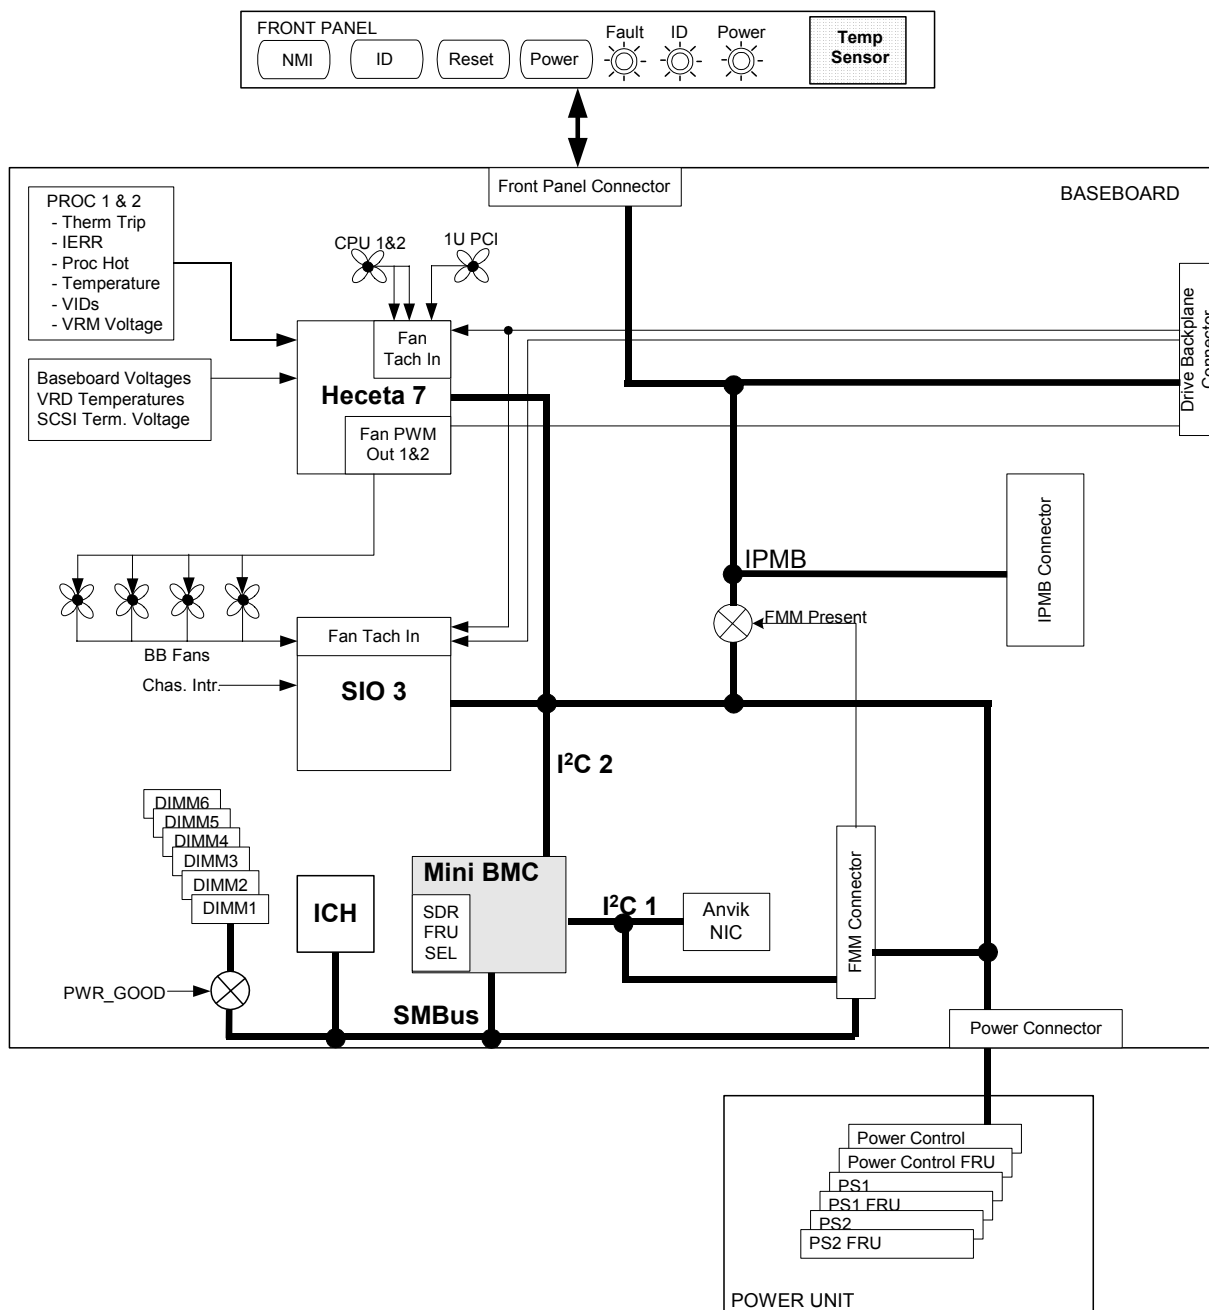


Figure 36. Block Diagram of Platform Management Architecture

6.1.2 5V Standby

The power supply must provide a 5V Standby power source for the platform to provide any management functionality. 5V Standby is a low power 5V supply that is active whenever the system is plugged into AC power. 5V Standby is used by the following onboard management devices:

- Management controller (BMC and/or mBMC) and associated RAM, Flash, and EEPROM which are used to monitor the various system power control sources including the front panel Power Button, the baseboard RTC alarm signal, and power on request messages from the auxiliary IPMB connector and PCI SMBus.
- On-board NICs which support IPMI-over-LAN and LAN Alerting, Wake-On LAN, and Magic Packet operation.
- Emergency management port
- IPMB
- PCI SMBus (e.g. logic and private busses used for power control)
- ICMB Transceiver card (if present)
- IPMB isolation circuit
- System Status LED on the front panel
- System Identify LED

6.1.3 IPMI Messaging, Commands, and Abstractions

The *Intelligent Platform Management Interface Specification* defines a standardized, abstracted, message-based interface between software and the platform management subsystem, and a common set of messages (commands) for performing operations such as accessing temperature, voltage, and fan sensors, setting thresholds, logging events, controlling a watchdog timer, etc.

IPMI also includes a set of records called Sensor Data Records (SDRs) that make the platform management subsystem self-descriptive to system management software. The SDRs include software information such as how many sensors are present, what type they are and what events they generate. The SDRs also include information such as minimum and maximum ranges, sensor type, accuracy and tolerance, etc., that guides software in interpreting and presenting sensor data.

Together, IPMI Messaging and the SDRs provide a self-descriptive, abstracted platform interface that allows management software to automatically configure itself to the number and types of platform management features on the system. In turn, this enables one piece of management software to be used on multiple systems. Since the same IPMI messages are used over the serial/modem and LAN interfaces, a software stack designed for in-band (local) management access can readily be re-used as an out-of-band remote management stack by changing the underlying communications layer for IPMI messaging.

6.1.4 IPMI ‘Sensor Model’

An IPMI-compatible ‘Sensor Model’ is used to unify the way that temperature, voltage, and other platform management status and control is represented and accessed. The implementation of this model is done according to command and data formats defined in the *Intelligent Platform Management Interface Specification*.

The majority of monitored platform elements are accessed as logical ‘Sensors’ under this model. This access is accomplished using an abstracted, message-based interface (IPMI messages). Instead of having system software access the platform monitoring and control hardware registers directly, it sends commands, such as the *Get Sensor Reading* command, for

sensor access. The message-based interface isolates software from the particular hardware implementation.

System Management Software discovers the platform's sensor capabilities by reading the Sensor Data Records from a Sensor Data Record Repository managed by the management controller (i.e. mBMC or IMM). Sensor Data Records provide a list of the sensors, their characteristics, location, type, and associated Sensor Number, for sensors in a particular system. The Sensor Data Records also hold default threshold values (if the sensor has threshold based events), factors for converting a sensor reading into the appropriate units (mV, rpm, degrees Celsius, etc.), and information on the types of events that a sensor can generate.

Sensor Data Records also provide information on where Field Replaceable Unit (FRU) information is located, and information to link sensors with the entity and/or FRU they're associated with.

Information in the SDRs is also used for configuring and restoring sensor thresholds and event generation whenever the system powers up or is reset. This is accomplished via a process called the 'initialization agent'. The management controller reads the SDRs and based on bit settings, writes the threshold data. Then it enables event generation for the various sensors it monitors.

System Management Software uses the data contained in the Sensor Data Record information to locate sensors in order to poll them, interpret, and present their data readings, adjust thresholds, interpret SEL entries, and alter event generation settings.

In Standard and Advanced management models, SDRs also provide a mechanism for extending the baseboard management with additional chassis or OEM 'value-added' monitoring and events. The baseboard monitoring can be extended by implementing an IPMI-compatible management controller, connecting it to the IPMB, and adding new SDRs describing that controller and its sensors to the SDR Repository. System Management Software can then read the SDRs and use them to automatically incorporate the additional sensors.

6.1.5 Private Management Bus

A 'Private Management Bus' is a single-master I²C bus that is controlled by the management controller. Access to any of the devices on the Private Management Bus is accomplished indirectly via commands to the management controller via the IPMB or system interfaces. The Private Management bus is a common mechanism used for accessing temperature sensors, system processor information, and other baseboard monitoring devices that are located in various locations in the system.

The devices on the Private Management Bus are isolated from traffic on the IPMB. Since devices such as temperature sensors are polled by the management controller, this gets the polling traffic off the 'public' IPMB bus. This also increases the reliability of access to the information, since issues with IPMB bus arbitration and message retries are avoided.

Furthermore, placing managed I²C devices on the private management bus frees up the I²C addresses that those devices would have used up on the IPMB.

6.1.6 Management Controllers

At the heart of platform management is a management controller. To support the tiered management model, the Intel® Server Board SE7520AF2 supports two different management controllers. Integrated onto the baseboard is the National Semiconductor* 87431x Mini-BMC (mBMC) to provide the functionality of the onboard platform instrumentation management tier. The Standard and Advanced modules electrically replace the mBMC with the more full featured 'Sahalee' microcontroller. Sahalee is a custom ARM7-TDMI based microcontroller designed for baseboard management applications on Intel Server baseboards.

The management controller is a microcontroller that provides the intelligence at the heart of the Intelligent Platform Management architecture. The primary purpose of the management controller is to autonomously monitor system 'sensors' for system platform management events, such as over-temperature, out-of-range voltages, etc., and log their occurrence in the non-volatile System Event Log (SEL). The management controller also provides the interface to the sensors and SEL so System Management Software can poll and retrieve the present status of the platform. The contents of the log can be retrieved 'post mortem' in order provide failure analysis information to field service personnel. It is also accessible by System Management Software, such as Intel Server Management (ISM), running under the operating system.

The management controller 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.

The management controller includes 'recovery control' functions that allow local or remote software to request actions such as power on/off, power cycle, and system hard resets, plus an IPMI Watchdog Timer that can be used by BIOS and run-time management software as a way to detect software hangs.

The management controller provides 'out-of-band' (OOB) remote management interfaces providing access to the platform health, event log, and recovery control features via LAN (all tiers). Standard and Advanced systems also allow access via serial/modem, IPMB, PCI SMBus, and ICMB interfaces. These interfaces remain active on standby power, providing a mechanism where the SEL, SDR, and recovery control features can be accessed even when the system is powered down.

Because the management controller operates independently from the main processor(s), the management controller monitoring and logging functions, and the out-of-band interfaces can remain operative even under failure conditions that cause the main processors, operating system, or local system software to stop.

The management controller also provides the interface to the non-volatile 'Sensor Data Record (SDR) Repository'. IPMI Sensor Data Records provide a set of information that system management software can use to automatically configure itself for the number and type of IPMI sensors (e.g. temperature sensors, voltage sensors, etc.) in the system. This information allows management software to automatically adapt itself to the particular system, enabling the development of management software that can work on multiple platforms without requiring the software to be modified.

The following is a list of the major functions that are managed by either the mBMC or the BMC in the IMM modules.

- Sensors and Sensor Polling
- FRU Information Access. FRU (Field Replaceable Unit) information is non-volatile storage for serial number, part number, asset tag and other inventory information for the baseboard and chassis. The FRU implementation on SE7520AF2 includes write support for OEM-specific records.
- Autonomous Event Logging. The management controller autonomously polls baseboard sensors and generates IPMI Platform Events, also called Event Messages, when an event condition is detected. The events are automatically logged to the System Event Log (SEL).
- System Event Log (SEL). Non-volatile storage for platform health events. Events can be autonomously logged by the BMC, or by sending Event Messages via the system interface or IPMB to the BMC. This enables BIOS, software, and add-in cards to also log events.
- Sensor Data Record (SDR) Repository. Non-volatile storage holding records describing the number and type of management sensors on the baseboard and in the chassis. Includes write support for OEM-specific records and sensors.
- SDR/SEL Timestamp Clock. A clock internally maintained by the management controller that is used for time-stamping events and recording when SDR and SEL contents have changed.
- Intelligent Platform Management Bus (IPMB). The IPMB is a two-wire, multi-master serial bus that provides a point for extending the baseboard management to include chassis management features, and for enabling add-in cards to access the baseboard management subsystem. (Standard and Advanced systems only.)
- Watchdog Timer with selectable timeout actions (power off, power cycle, reset, or NMI) and automatic logging of timeout event
- Direct Platform Control (DPC) LAN Remote Management Connection
- LAN Alerting via PET (Platform Event Trap) format SNMP trap
- Serial/Modem Remote Management Connection (Standard and Advanced systems only)
- Serial/Modem Event Paging/Alerting (Standard and Advanced systems only)
- Platform Event Filtering (PEF)
- Keyboard Controller Style (KCS) IPMI-System Interface (Standard and Advanced systems only)
- SMBus IPMI-System Interface (Onboard Platform Instrumentation systems only)
- Intelligent Chassis Management Bus (ICMB) support (Standard and Advanced systems only)
- Remote Boot Control
- Local and Remote Power On/Off/Reset Control
- Local and Remote Diagnostic Interrupt (NMI) Control
- Fault-Resilient Booting
- Front Panel LED Control
- Platform Management Interrupt Routing (Standard and Advanced systems only)

- Power Distribution Board (PDB) monitoring (Standard and Advanced systems only)
- Updateable BMC Firmware
- System Management Power Control (including providing Sleep/Wake and power push-button interfaces)
- Platform Event Filtering (PEF)
- Platform Fan Speed Control and Failure Monitoring (mBMC, Heceta*, FRUSDR)
- Speaker ‘Beep’ Capability (used to indicate conditions such as FRB failure) (Standard and Advanced systems only)
- Baseboard FRU Information interface
- Diagnostic Interrupt (Front Panel NMI) Handling
- SMI/NMI status monitor (Standard and Advanced systems only)
- System interface to the IPMB (via System Interface Ports) (Standard and Advanced systems only)
- System interface to the PCI SMBus (via System Interface Ports) (Standard and Advanced systems only)
- Secure Mode Control - front panel lock/unlock initiation.
- *Intelligent Platform Management Interface Specification v1.5* Management Controller Initialization Agent function
- Emergency Management Port (EMP) Serial/Modem platform management interface (Standard and Advanced systems only)
- Dedicated Network Interface controller and full TCP/IP software stack (Advanced only)

6.2 Onboard Platform Instrumentation Management Features and Functionality

The mini Baseboard Management Controller (mBMC) is an Application Specific Integrated Circuit (ASIC) with many peripheral devices embedded into it. The mBMC contains the logic needed for controlling the system, monitoring the sensors, and communicating with other systems and devices via various external interfaces.

The following figure is a block diagram of the mBMC as it is used in a server management system. The external interface blocks to the mBMC are the discrete hardware peripheral device interface modules.

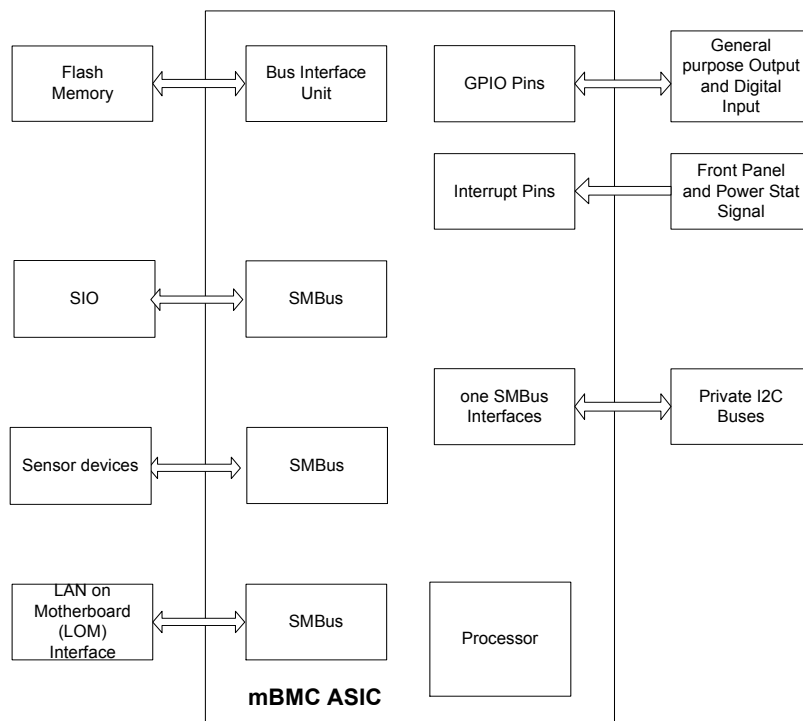


Figure 37. mBMC in a Server Management System

6.2.1 mBMC Self-test

The mBMC performs various tests as part of its initialization. If a failure is determined, the mBMC stores the error internally. A failure may be caused by a corrupt mBMC FRU, SDR, or SEL. The *IPMI 1.5 Get Self Test Results* command can be used to return the first error detected.

Executing the *Get Self Test Results* command causes the mBMC self-test to be run. It is strongly recommended to reset the mBMC via the *Cold Reset* command afterwards.

6.2.2 SMBus Interfaces

The mBMC incorporates one slave and two master-only SMBus interfaces. The mBMC interfaces with the host through the slave SMBus interface. It interfaces with the LAN On Motherboard (LOM) and peripherals through two independent master bus interfaces.

6.2.3 External Interface to mBMC

Figure 38 shows the data/control flow to and within the functional modules of the mBMC. External interfaces from the host system, LOM, and peripherals, interact with the mBMC through the corresponding interface modules as shown.

The mBMC communicates with the internal modules using its private SMBus. External devices and sensors interact with the mBMC using the peripheral SMBus. LOM communicates through the LOM SMBus. GPIO pins are available and are used for general input and output functions. Dedicated LED lines are used for LED/color control.

Also built into the mBMC are the control functions for both the power supply and front panel.

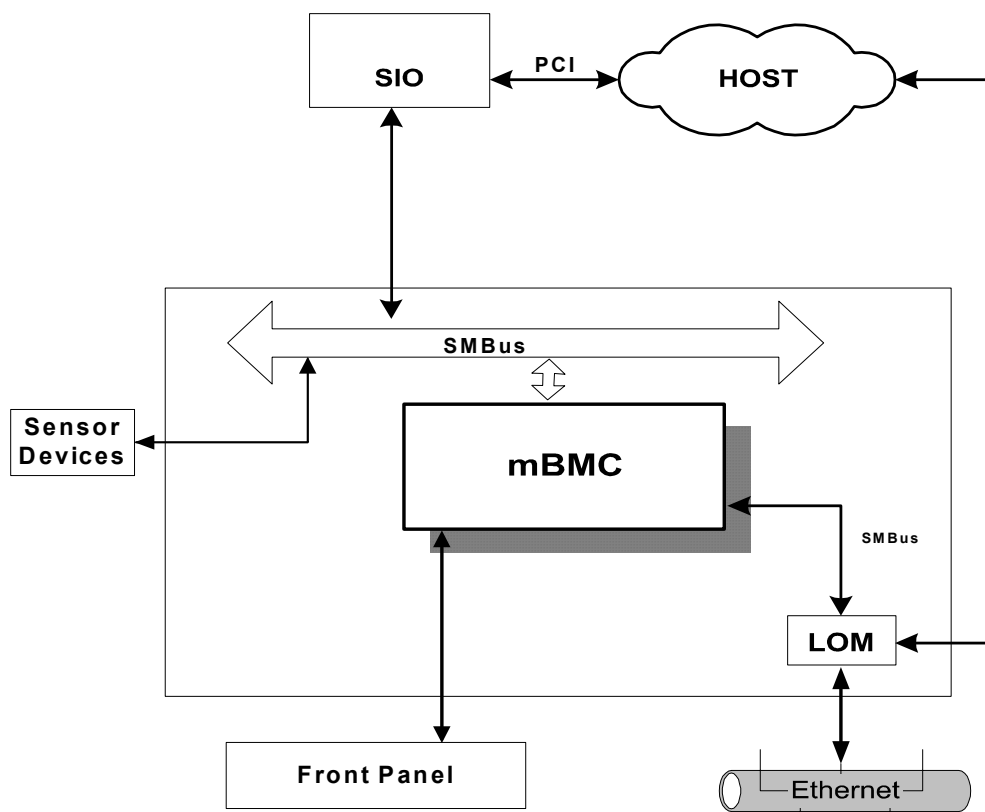


Figure 38. External Interfaces to mBMC

6.2.3.1 Private Management I²C Buses

The mBMC implements a single private management bus. The mBMC is the sole master on this bus. External agents must use the mBMC *Master Write/Read I²C* command if they require direct communication with a device on this bus. In addition, the mBMC provides a *Reserve Device* command that gives an external agent exclusive access to a specific device for a selectable time.

6.2.4 Messaging Interfaces

This section describes the supported mBMC communication interfaces:

- Host SMS interface via SMBus interface
- LAN interface using the LAN On Motherboard SMBus

6.2.4.1 Channel Management

The mBMC supports two channels: System interface, and 802.3 LAN.

Table 70. Supported Channel Assignments

Channel Id	Media type	Interface	Supports Sessions
1	802.3 LAN Interface	IPMI-SMBus	Multi sessions
2	System Interface	IPMI-SMBus	Session-less

6.2.4.2 User Model

The mBMC supports one anonymous user (null user name) with a settable password.

6.2.4.3 Request/Response Protocol

All of the protocols used in the host interface and the LOM interface are Request/Response protocols. A Request Message is issued to an intelligent device, to which the device responds with a separate Response Message.

6.2.4.4 System Communication Interface

The host communicates with the mBMC via the System Management Bus (SMBus). The interface consists of three signals:

- SMBus clock signal (SCLH)
- SMBus data signal (SDAH)
- Optional SMBus alert signal (SMBAH). The signal notifies the host that the PC87431x has data to provide.

The mBMC is a slave device on the bus. The host interface is designed to support polled operations. Host applications can optionally handle an SMBus alert interrupt if the mBMC is unable to respond immediately to a host request. In this case, “Not Ready” is indicated in one of two ways:

- The host interface bandwidth is limited by the bus clock and mBMC latency. To meet the device latency, the mBMC slows down the bus periodically by extending the SMBus clock low interval (SCLH).
- If the mBMC is in the middle of a LAN or peripheral device communication, or if a response to the host request is not yet ready, the mBMC does not acknowledge the device address (“NAK”). This forces the host software to stop and restart the session.

For more information on read-write through SMBus refer the *System Management Bus Specification 2.0*.

6.2.4.5 LAN Interface

The baseboard supports one DPC LAN interface via a UDP port 26Fh. The mBMC supports a maximum of one simultaneous session across all authenticated channels. The baseboard

implements gratuitous ARP support according to the *Intelligent Platform Management Interface Specification v1.5*.

The *Intelligent Platform Management Interface Specification v1.5* defines how IPMI messages, encapsulated in RMCP packet format, can be sent to and from the mBMC. This capability allows a remote console application to access the mBMC and perform the following operations:

- Chassis Control, e.g., get chassis status, reset chassis, power-up chassis, power-down chassis
- Get system sensor readings
- Get and Set system boot options
- Get Field Replaceable Unit (FRU) information
- Get System Event Log (SEL) entries
- Get Sensor Data Records (SDR)
- Set Platform Event Filtering (PEF)
- Set LAN configurations

In addition, the mBMC supports LAN alerting in the form of SNMP traps that conform to the IPMI Platform Event Trap (PET) format.

Table 71. LAN Channel Capacity

LAN CHANNEL Capability	Options
Number of Sessions	1
Number of Users	1
User	Name NULL (anonymous)
User Password	Configurable
Privilege Levels	User, Operator, Administrator
Authentication Types	MD5
Number of LAN Alert Destinations	1
Address Resolution Protocol (ARP)	Gratuitous ARP

6.2.5 Direct Platform Control (IPMI over LAN)

Direct Platform Control provides a mechanism for delivering IPMI Messages directly to the management controllers via a LAN connection. The NICs and the management controllers remain active on standby power, enabling the IPMI Messaging when the system is powered up, powered down, and in a system sleep state. This allows a remote console application to be able to access the management controller capabilities, including:

- Power on/off and reset control with the ability to set BIOS boot flags
- FRU, SDR, and SEL access
- BMC configuration access
- Ability to transfer IPMI messages between the LAN interface and other interfaces. This capability enables messages to be delivered to system management software, and

provides the ability to access sensors and FRU information on other management controllers.

IPMI Messages are encapsulated in a packet format called Remote Management Control Protocol (RMCP). The Distributed Management Task Force (DMTF) has defined RMCP for supporting pre-operating system and operating system-absent management. RMCP is a simple request-response protocol that can be delivered using UDP datagrams. IPMI-over-LAN uses version 1 of the RMCP protocol and packet format.

UDP port 26Fh is a 'well known port' address that is specified to carry RMCP formatted UDP datagrams. The on-board Intel network interface controllers contain circuitry that enables detecting and capturing RMCP packets that are received on Port 26Fh and making them available to the management controller via a 'side-band' interface that is separate from the PCI interface to the NIC. Similarly, the management controller can use the side-band interface to send packets from Port 26Fh, as shown in the following figure.

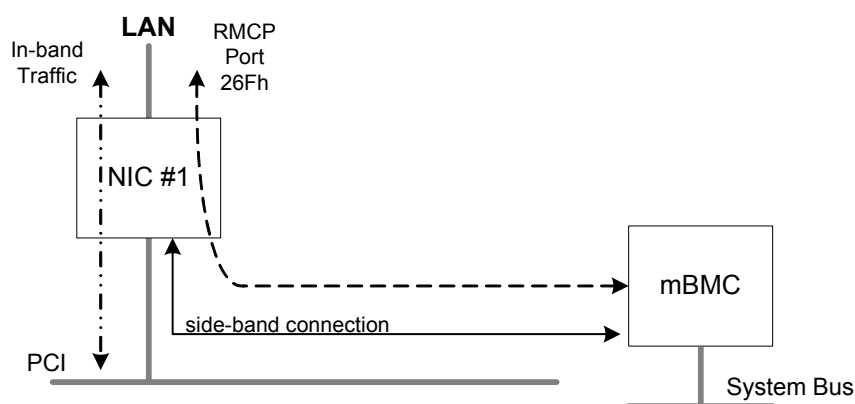


Figure 39. IPMI-over-LAN

RMCP includes a field that indicates the class of messages that can be embedded in an RMCP message packet. For RMCP version 1.0, the defined classes are IPMI, ASF, and OEM. IPMI-over-LAN uses the IPMI class to transfer IPMI Messages encapsulated in RMCP packets. Intelligent Platform Management Interface v1.5 Specification specifies the packet formats and commands used to perform IPMI Messaging on LAN via RMCP.

The management controller transmits to other port addresses as needed. For example, LAN Alerts, which are sent as SNMP Traps, can be transmitted to the SNMP Trap 'well known' port address, 162 (0A2h).

6.2.5.1 LAN Channel Specifications

The following table presents the minimum support that will be provided. Note that system management software and utilities may not use all the available management controller options and capabilities. For detailed technical information on the operation of the LAN channel operation and LAN Alerting, refer to *Intelligent Platform Management Interface Specification v1.5*.

Table 72. LAN Channel Specifications

Configuration Capability	Options	Description/Notes
Channel Access Modes	Always-active, disabled	This option determines when the BMC can be accessed via IPMI Messaging over LAN.
Number of Sessions	1 (OB Platform Instrumentation)	The number of simultaneous sessions that can be supported is shared across the LAN and serial/modem channels.
Number of Users	1 (OB Platform Instrumentation)	User information is a resource that is shared across the LAN and serial/modem channels.
Configurable User Names	No (OB Platform Instrumentation)	User information is a resource that is shared across the LAN and serial/modem channels.
Configurable User Passwords	Yes	
Privilege Levels	User, Operator, Administrator	
IPMI Message Authentication Type Support	MD5	
Number of LAN Alert destinations	1 (OB Platform Instrumentation)	
PET Acknowledge support	Yes	
Gratuitous ARP Support	Yes	

6.2.5.2 LAN Drivers and Setup

The IPMI-over-LAN feature must be used with the appropriate Intel NIC Driver, and the NIC correctly configured in order for DPC LAN operation to occur transparently to the operating system and network applications. If an incorrect driver or NIC configuration is used, it is possible to get driver timeouts when the IPMI-over-LAN feature is enabled.

6.2.5.3 BIOS Boot Flags

A remote console application can use the IPMI *Set System Boot Options* command to configure a set of BIOS boot flags and boot initiator parameters that are held by the management controller. These parameters include information that identifies the party that initiated the boot, plus flags and other information that can be used to direct the way booting proceeds after a system reset or power-up. For example, the system can be configured to boot normally, boot using PXE, etc.

6.2.6 Wake On LAN / Power On LAN and Magic Packet Support

The baseboard supports Wake On LAN / Power On LAN capability using the on-board network interface chips or an add-in network interface card. An add-in network card can deliver the wake signal to the baseboard via the PME signal on the PCI bus. The actual support for Magic Packet and/or packet filtering for Wake On LAN / Power On LAN is provided by the NIC. The baseboard handles the corresponding wake signal.

6.2.6.1 Wake On LAN in S4/S5

A configuration option is provided that allows the on-board NICs to be enabled to wake the system in an S4/S5 state, even if the operating system disabled Wake-On-LAN when it powered down the system. This provides an option for users who want to use standard, but non-secure, WOL capability for operations such as after-hours maintenance. Note that the DPC LAN capability provides a secure system power-up, plus the ability to provide BIOS boot options, by sending authenticated IPMI messages directly to the BMC via the on-board NICs.

6.2.7 Watchdog Timer

The mBMC implements an IPMI 1.5-compatible watchdog timer. See the IPMI specification for details. SMI and NMI pre-timeout actions are supported, as are hard reset, power down, and power cycle timeout actions.

6.2.8 System Event Log (SEL)

The mBMC implements the logical System Event Log device as specified in the *Intelligent Platform Management Interface Specification, v1.5*. The SEL is accessible via all communication transports. In this way, the SEL information can be accessed while the system is down by means of out-of-band interfaces. The maximum SEL size that is supported by mBMC is 92 entries.

Supported commands are:

- Get SEL Info
- Reserve SEL
- Get SEL Entry
- Add SEL Entry
- Clear SEL
- Get SEL Time
- Set SEL Time

6.2.8.1 Timestamp Clock

The management controller (mBMC or BMC) maintains a four-byte internal timestamp clock used by the SEL and SDR subsystems. This clock is incremented once per second. It is read using the *Get SEL Time* command and set using the *Set SEL Time* command. The *Get SDR Time* command can also be used to read the timestamp clock. These commands are specified in the *Intelligent Platform Management Interface Specification, v1.5*.

After a reset or power up, the management controller sets the initial value of the timestamp clock to 0x00000000. It is incremented once per second after that. A SEL event containing a timestamp from 0x00000000 to 0x140000000 has a timestamp value that is relative to management controller initialization.

During POST, the BIOS tells the management controller the current time via the *Set SEL Time* command. The management controller maintains this time, incrementing it once per second, until the controller is reset or the time is changed via another *Set SEL Time* command.

If the RTC changes during system operation, system management software must synchronize the management controller time with the system time. If this is not done, the server should be reset so that BIOS will pass the new time to the mBMC.

6.2.9 Sensor Data Record (SDR) Repository

The mBMC includes built-in Sensor Data Records that provide platform management capabilities (sensor types, locations, event generation and access information). The SDR Repository is accessible via all communication transports. This way, out-of-band interfaces can access the SDR Repository information if the system is down.

The mBMC supports 2176 bytes of storage for SDR records. The SDR defines the type of sensor, thresholds, hysteresis values and event configuration. The mBMC supports up to six threshold values for threshold-based full sensor records, and up to 15 events for non threshold-based full and compact sensor records. It also supports both low-going and high-going sensor devices.

6.2.9.1 Initialization Agent

The mBMC implements the internal sensor initialization agent functionality specified in the *Intelligent Platform Management Interface Specification, v1.5*. When the mBMC initializes, or when the system boots, the initialization agent scans the SDR repository and configures the sensors referenced by the SDRs. This includes setting sensor thresholds, enabling/disabling sensor event message scanning, and enabling/disabling sensor event messages.

6.2.10 Event Message Reception

The mBMC supports externally (e.g., BIOS) generated events via the Platform Event Message command. Events received via this command will be logged to the SEL and processed by Platform Event Filtering (PEF) as described below.

6.2.11 Event Filtering and Alerting

The mBMC implements the following *Intelligent Platform Management Interface Specification v1.5* alerting features:

- PEF
- Alert over LAN

6.2.11.1 Platform Event Filtering (PEF)

The mBMC monitors platform health and logs failure events into the SEL. The Platform Event Filtering feature provides a configurable mechanism to allow events to trigger alert actions. PEF provides a flexible, general mechanism that enables the mBMC to perform selectable actions triggered by a configurable set of platform events. The mBMC supports the following IPMI PEF actions:

- Power-down
- Soft shut-down
- Power cycle
- Reset

- Diagnostic Interrupt
- Alert

The mBMC maintains an Event Filter table with 30 entries that is used to select the actions to perform. Also maintained is a fixed/read-only Alert Policy Table entry. No alert strings are supported.

Note: All Fault/Status LED and ID LED behaviors are driven off of PEF. PEF should not be disabled and the manufacturing shipping configuration should not be modified or those behaviors will be changed.

Each time the PEF module receives either an externally or internally generated event message, it compares the event data against the entries in the event filter table. The mBMC scans all entries in the table and determines a set of actions to be performed. If a combination of actions is identified, such as power down, power cycle, and/or reset actions, the action are performed according to PEF Action Priorities. Action priorities are outlined in the table below.

Note: An action that has changed from delayed to non-delayed, or an action whose delay time has been reduced has a higher priority. Each generated event is logged by SEL.

Table 73. PEF Action Priorities

Action	Priority	Delayed	Type	Note
Power-down	1	Yes	PEF Action	
Soft shut-down	2	Yes	OEM PEF Action	Not executed if a power-down action was also selected.
Power cycle	3	Yes	PEF Action	Not executed if a power-down action was also selected.
Reset	4	Yes	PEF Action	Not executed if a power-down action was also selected.
Diagnostic Interrupt (NMI)	5	No	PEF Action	Not executed if a power-down action was also selected.
PET Alert	6	No	PEF Action	When selected, always occurs immediately after detection of a critical event.
IPMB message event	8	No	OEM PEF Action	When selected, always occurs immediately after detection of a critical event.

Table 74. mBMC Factory Default Event Filters

Event Filter #	Offset Mask	Events
1	Non-critical	Voltage Assert
2	Non-critical	Voltage Deassert
3	Critical	Voltage Assert
4	Critical	Voltage Deassert
5	Critical	PS Soft Fail Assert
6	Critical	PS Soft Fail Deassert
7	Critical	Proc 1-2 Thermal Trip Assert

Event Filter #	Offset Mask	Events
8	Critical	Proc 1-2 Thermal Trip, Config Error and IERR Deassert
9	Degraded	Proc 1-2 FRB3 Assert
10	Degraded	Proc 1-2 FRB3 Deassert
11	Degraded	Proc 1-2 Hot Assert
12	Degraded	Proc 1-2 Hot Deassert
13	Critical	FP NMI Assert
14	Critical	FP NMI Deassert
15	Non Critical	SCSI Terminator Fail Assert
16	Non Critical	SCSI Terminator Fail Deassert
17	N/A	ID Button Assert
18	N/A	ID Button Deassert
19	Critical	Fan Speed Assert
20	Critical	Fan Speed Deassert
21	Non Critical	Fan Speed Assert
22	Non Critical	Fan Speed Deassert
23	Critical	Temperature Assert
24	Critical	Temperature Deassert
25	Non Critical	Temperature Assert
26	Non Critical	Temperature Deassert
27	Critical	Proc 1-2 IERR Assert
28	Critical	CPU Configuration Error
29	N/A	Reserved for ISM
30	N/A	Reserved for ISM

6.2.11.2 Alert over LAN

LAN alerts are sent as SNMP traps in ASF formatted Platform Event Traps to a specified alert destination. The Alert over LAN feature is used to send either Platform Event Trap alerts or directed events to a remote system management application, regardless of the state of the host's operating system. LAN alerts are sent over the dedicated server management NIC (NIC1). LAN alerts can be used by PEF to send out alerts to selected destination when ever an event matches an event filter table entry. For more information on LAN alerts, see the *Intelligent Platform Management Interface Specification v1.5*.

6.2.11.3 System Identification in Alerts

The PET alert format used in PPP and LAN Alerting contains a system GUID field that can be used to uniquely identify the system that raised the alert. In addition, since the PET is carried in a UDP packet, the alerting system's IP Address is also present.

6.2.11.4 Platform Alerting Setup

The management controller provides commands via the System Interface that support setting/retrieving the alerting configuration LAN alerting in mBMC NV storage.

The user does not typically deal with filter contents directly. Instead, the Server Setup Utility provides a user interface that allows the user to select among a fixed set of pre-configured event filters.

The following list presents the type of Alerting configuration options that are provided:

- Enabling/Disabling PEF.
- Configuring Alert actions.
- Selecting which pre-configured events trigger an alert.
- Generating a 'test' event to allow the paging configuration to be checked¹
- Configuring the PPP Accounts for PPP Alerting¹. (PPP Accounts represent the phone number and user login information necessary to connect to a remote system via PPP).

6.2.11.5 Alerting Events while in Power Down

The mBMC is capable of generating alerts while the system is powered down. A watchdog power-down event alert is sent after the power down so that the alert does not delay the power-down action.

6.2.11.6 Alerting On System Reset Events

The alerting process must complete before the system reset is completed. This is done to simplify timing interactions between the mBMC and BIOS initialization after a system reset.

6.2.11.7 Alert-in-Progress Termination

An alert in progress will be terminated by a system reset or power on, or by disabling alerting via commands to the management controller.

6.2.12 NMI Generation

The following may cause the mBMC to generate an NMI pulse:

- Receiving a Chassis Control command issued from one of the command interfaces. Use of this command will not cause an event to be logged in the SEL.
- Detecting that the front panel Diagnostic Interrupt button has been pressed.
- A PEF table entry matching an event where the filter entry has the NMI action indicated.
- A processor IERR or Thermal Trip (if the mBMC is so configured).
- Watchdog timer pre-timeout expiration with NMI pre-timeout action enabled.

The mBMC-generated NMI pulse duration is 200ms. This time is chosen to try to avoid the BIOS missing the NMI if the BIOS is in the SMI Handler and the SMI Handler is masking the NMI.

¹ available only with Professional & Advanced server management

6.3 Platform Management Interconnects

6.3.1 Power Supply Interface Signals

The mBMC supports two power supply control signals: *Power On* and *Power Good*. The *Power On* signal connects to the chassis power subsystem and is used to request power state changes (asserted = request *Power On*). The *Power Good* signal from the chassis power subsystem indicates current the power state (asserted = power is on).

Figure 40 shows the power supply control signals and their sources. To turn the system on, the mBMC asserts the *Power On* signal and waits for the *Power Good* signal to assert in response, indicating that DC power is on.

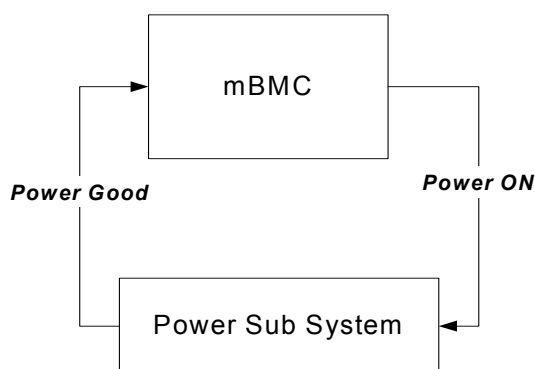


Figure 40. Power Supply Control Signals

The mBMC uses the *Power Good* signal to monitor whether the power supply is on and operational, and to confirm whether the actual system power state matches the intended system on/off power state that was commanded with the *Power On* signal.

De-assertion of the *Power Good* signal generates an interrupt that the mBMC uses to detect either power subsystem failure or loss of AC power. If AC power is suddenly lost, the mBMC logs a *Power Off* event in the SEL.

6.3.1.1 Power-up Sequence

When turning on the system power in response to one of the event occurrences listed in Table 75 below, the mBMC executes the following procedure:

1. The mBMC asserts *Power On* and waits for the power subsystem to assert *Power Good*. The system is held in reset.
2. The mBMC initializes all sensors to their *Power On* initialization state by running the Init Agent.
3. The mBMC attempts to boot the system by running the FRB3 algorithm, if enabled in BIOS setup.

6.3.1.2 Power-down Sequence

To power down the system, the mBMC effectively performs the sequence of power-up steps in reverse order. This operation can be initiated by one of the event occurrences listed in the following table and proceeds as follows:

1. The mBMC asserts system reset (de-asserts *Power Good*).
2. The mBMC de-asserts the *Power On* signal.
3. The power subsystem turns off system power upon de-assertion of the *Power On* signal.

6.3.1.3 Power Control Sources

The sources listed in the following table can initiate power-up and/or power-down activity.

Table 75. Power Control Initiators

#	Source	External Signal Name or Internal Subsystem	Capabilities
1	Power Button	FP Power button	Turns power On or Off
2	mBMC Watchdog Timer	Internal mBMC timer	Turns power Off, or power cycle
3	Platform Event Filtering	PEF	Turns power Off, or power cycle
4	Command	Routed through command processor	Turns power On or Off, or power cycle
5	Power state retention	Implemented via mBMC internal logic	Turns power On when AC power returns
6	Chipset	Sleep S5	Turns power On or Off

6.3.2 System Reset Control

6.3.2.1 Reset Signal Output

The mBMC asserts the *System Reset* signal on the baseboard to perform a system reset. The mBMC asserts the *System Reset* signal before powering the system up. After power is stable (as indicated by the power subsystem *Power Good* signal), the mBMC sets the processor enable state as appropriate and de-asserts the *System Reset* signal, taking the system out of reset.

To reset the system without a power state change, the mBMC:

1. Asserts the *System Reset* signal.
2. Holds this state for as long as the reset button is pushed.
3. De-asserts the *System Reset* signal.

6.3.2.2 Reset Control Sources

The following table shows the reset sources and the actions taken by the system.

Table 76. System Reset Sources and Actions

#	Reset Source	System Reset?	mBMC Reset
1	Standby power comes up	No (no DC power)	Yes
2	Main system power comes up	Yes	No
3	Reset button or in-target probe (ITP) reset	Yes	No
4	Warm boot (example: DOS Ctrl-Alt-Del)	Yes	No
5	Command to reset the system	Yes	No
6	Set Processor State command	Yes	No
7	Watchdog timer configured for reset	Yes	No
8	FRB3 failure	Yes	No
9	PEF action	Optional	No

6.3.3 Temperature-based Fan Speed Control

Baseboard hardware implements an ambient-temperature-based Fan Speed control that is part of *normal system operation*. The feature allows the baseboard to drive different fan speeds based on various temperature measurements in order to lower the acoustic noise of the system.

The ambient-temperature thresholds at which the Fan Speed increases does not correspond to a non-critical (warning) condition for the fan - since the fan's state is still 'OK' from the system point-of-view.

The baseboard has two analog Fan Speed signals that are driven by pulse-width modulator (PWM) circuits by the baseboard hardware. These signals can be driven to several levels according to temperature measurements. Multiple bytes of a Sensor Initialization Table is used to hold parameters that set the temperature thresholds and corresponding PWM duty cycles. This table is loaded as part of the baseboard configuration.

The management controller firmware expects to find an LM30 temperature sensor on the front panel board. Thus, the ambient temperature-based fan speed control capability is not enabled by default for SE7520AF2 as a baseboard-only product, but can be enabled via a management controller configuration change.

6.3.3.1 Fan Kick Start

Some fans may not begin rotating unless started at high speed. To ensure that the fans start, the baseboard hardware will start and run the fans at high speed for a brief interval following system power up.

6.3.4 Front Panel Control

The mBMC provides the main 'front panel control' functions. These include control of the system Power Button, Reset Button, Diagnostic Interrupt (Front Panel NMI) Button, System Identify Button, System ID LED, Status/Fault LED, and Chassis Intrusion Switch. Front panel control also includes the front panel lockout features.

6.3.4.1 Power Button

The front panel *Power Button* signal is routed to the mBMC, after debouncing this signal, the mBMC asserts PWBTOUT signal to the chipset PWRBTN input. The chipset responds by

deasserting SLEEP_S5 to the CPU configuration circuitry. If the configuration is good, PS_PWRON is asserted to the power supply. The supply then asserts POWERGOOD back to the mBMC.

If the system is in Secure Mode or the *Power Button* is forced protected, then when the power switch is pressed, a Platform Security Violation Attempt event message is generated and no power control action is taken.

In the case of simultaneous button presses, the *Power Button* action takes priority over all other buttons. For example, if the sleep button is depressed for one second and then the *Power Button* is pressed and released, the system powers down. Due to the routing of the de-bounced *Power Button* signal to the chipset, the power signal action overrides the action of the other switch signals.

6.3.4.2 Reset Button

The reset button is a momentary contact button on the front panel. Its signal is routed through the front panel connector to the mBMC, which monitors and de-bounces it. The signal must be stable for at least 25ms before a state change is recognized.

An assertion of the front *Panel Reset* signal to the mBMC causes the mBMC to start the reset and reboot process. This action is immediate and without the cooperation of any software or operating system running on the system.

If *Secure Mode* is enabled or the button is forced protected, the reset button does not reset the system, but instead a Platform Security Violation Attempt event message is generated. The reset button is disabled in sleep mode.

6.3.4.3 Diagnostic Interrupt Button (Front Panel NMI)

As stated in the *IPMI 1.5 Specification*, a Diagnostic Interrupt is a non-maskable interrupt or signal for generating diagnostic traces and core dumps from the operating system. The mBMC generates the NMI, which can be used as an OEM-specific diagnostic front panel interface.

The Diagnostic Interrupt button is connected to the mBMC through the front panel connector. A Diagnostic Interrupt button press causes the mBMC to generate a 200 ms system NMI pulse.

This generates an event (NMI button sensor); the NMI is actually generated by a factory defined PEF filter.

6.3.4.4 Chassis ID Button and LED

The front panel interface supports a *Chassis Identify* Button and a corresponding Blue *Chassis Identify* LED (both available only on the rack version of the Server Chassis SC5300). A second Blue Chassis Identify LED is mounted on the back edge of the baseboard where it may be visible when viewed from the back of an integrated system.

The LED can provide a mechanism for identifying one system out of a group of identical systems in a rack environment

The Chassis Identify LED can be turned on either locally via the push-button signal, or by local or remote software using the IPMI *Chassis Identify* command. The following list summarizes the Chassis Identify Push-button and LED operation:

- The Identify signal state is preserved on Standby power across system power-on/off and system hard resets when an Intel Management Module is present. It is not preserved with the on-board platform instrumentation (mBMC and no IMM present) or when A/C power is removed. The initial LED state is OFF when A/C power is applied.
- The IPMI Chassis Identify command can be used to control the LED. If the Chassis Identify command is used to turn the LED On, the command will automatically time out and turn off the LED unless another Chassis Identify command to turn on the LED is received. The default timeout for the command is 15 seconds. The baseboard supports the optional command parameter to allow the timeout to be set anywhere from 1 to 255 seconds.
- The optional timeout parameter in the Chassis Identify command also allows software to tell the LED to go Off immediately.
- The Chassis Identify push button blinks the ID LED for 15 seconds when pressed (on-board platform instrumentation). When an Intel Management Module is present, however, this button works using a toggle “push-on/push-off” operation. Each press of the push-button toggles the LED signal state between On and Off. If the pushbutton is used to turn the LED On, it will stay on indefinitely, until either the button is pressed again or a Chassis Identify LED command causes the LED to go Off.

Table 77. Chassis ID LEDs

Color	Condition	When
Blue	Off	Default
	On/Blink	Identify button pressed or Chassis Identify command executed

6.3.4.5 Status/Fault LED

The following table shows mapping of sensors/faults to the LED state.

Table 78. Fault/Status LED

Color	Condition	When
Green	Solid	System Ready
	Blink	System Ready, but degraded. CPU fault, DIMM killed
Amber	Solid	Critical Failure: critical fan, voltage, temperature state
	Blink	Non-Critical Failure: non-critical fan, voltage, temperature state
Off	Solid	Not Ready. POST error/NMI event/CPU or terminator missing

Critical Condition

Any critical or non-recoverable threshold crossing associated with the following events:

- Temperature, voltage, or fan critical threshold crossing
- Power subsystem failure. The management controller asserts this failure whenever it detects a power control fault (e.g., detects that the system power is remaining on even though the management controller has de-asserted the signal to turn off power to the system).
- “Critical Event Logging” errors, including: System memory Uncorrectable ECC error and Fatal/Uncorrectable Bus errors, such as PCI SERR and PERR

Non-Critical Condition

- Temperature, voltage, or fan non-critical threshold crossing
- Chassis intrusion

Degraded Condition

- One or more processors are disabled by Fault Resilient Boot (FRB) or BIOS
- BIOS has disabled or mapped out some of the system memory

6.3.4.6 Chassis Intrusion Switch

The Intel® SE7520AF2 server platform supports chassis intrusion detection. The mBMC monitors the state of the *Chassis Intrusion* signal and makes the status of the signal available via the *Get Chassis Status* command and *Physical Security* sensor state. If enabled, a chassis intrusion state change causes the mBMC to generate a *Physical Security* sensor event message with a *General Chassis Intrusion* offset.

6.3.4.7 Front Panel Lockout

The management controller monitors a ‘Secure Mode’ signal from the keyboard controller on the baseboard. When the Secure Mode signal is asserted, the management controller locks out the ability to power down or reset the system using the power or reset push buttons, respectively. Secure Mode does not block the ability to initiate a sleep request using the Sleep push-button.

The management controller generates a ‘Secure Mode Violation Attempt’ event message if an attempt is made to power-down, or reset the system using the push buttons while Secure Mode is active.

The set of buttons protected when Secure Mode is active varies depending on the system ACPI power state and whether the Sahalee BMC or the mBMC is in control as shown in table Table 68. Differences are highlighted.

Note: The mBMC will prevent the system from powering up via button press when either secure mode or the front panel lockout I/O signal is asserted.

6.3.5 FRU Information

The platform management architecture supports providing Field Replaceable Unit (FRU) information for the baseboard and major replaceable modules in the chassis. ‘Major Module’ is defined as any circuit board in the system containing active electronic circuitry.

FRU information includes board serial number, part number, name, asset tag, and other information. FRUs that contain a management controller, use the controller to provide access to the FRU information. FRUs that lack a management controller can make their FRU information available via a EEPROM directly connected to the mBMC's sensor device private I²C bus. This allows the system integrator to provide a chassis FRU device without having to implement a management controller. This information can only be accessed via IPMI Master Write-Read commands.

The mBMC implements the interface for logical FRU inventory devices as specified in the *Intelligent Platform Management Interface Specification, v1.5*. This functionality provides commands used for accessing and managing the FRU inventory information associated with the baseboard (FRU ID 0). These commands can be delivered via all interfaces. All other FRUs must be accessed using IPMI Master Write-Read commands.

6.3.5.1 mBMC FRU Inventory Area Format

The mBMC FRU inventory area format follows the Platform Management FRU Information Storage Definition. Refer to *Platform Management FRU Information Storage Definition, v1.0* for details.

The mBMC provides only low-level access to the FRU inventory area storage. It does not validate or interpret the data stored in FRU memory. The baseboard's FRU information is kept in the mBMC internal flash memory.

6.4 Sensors

6.4.1 Sensor Type Codes

The following section list the sensor identification numbers and information regarding the sensor type, name, supported thresholds, assertion and deassertion information, and a brief description of the sensor purpose. Refer to the *Intelligent Platform Management Interface Specification, v1.5*, for sensor and event/reading-type table information.

- **Sensor Type**
The Sensor Type references the values enumerated in the *Sensor Type Codes* table in the IPMI specification. It provides the context in which to interpret the sensor, e.g., the physical entity or characteristic that is represented by this sensor.
- **Event/Reading Type**
The Event/Reading Type references values from the *Event/Reading Type Code Ranges* and *Generic Event/Reading Type Codes* tables in the *Intelligent Platform Management Interface Specification, v1.5*. Note that digital sensors are a specific type of discrete sensors, which have only two states.
- **Event Offset/Triggers**
Event Thresholds are 'supported event generating thresholds' for threshold types of sensors.
 - [u,l][nr,c,nc] upper nonrecoverable, upper critical, upper noncritical, lower nonrecoverable, lower critical, lower noncritical
 - uc, lc upper critical, lower critical

Event Triggers are ‘supported event generating offsets’ for discrete type sensors. 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/Deassertion Enables**

Assertions and Deassertion indicators reveals the type of events the sensor can generate:

- As: Assertions
- De: Deassertion

▪ **Readable Value / Offsets**

- Readable Value indicates the type of value returned for threshold and other non-discrete type sensors.
- Readable Offsets indicates the offsets for discrete sensors that are readable via the *Get Sensor Reading* command. Unless otherwise indicated, all Event Triggers are readable, i.e., *Readable Offsets* consists of the reading type offsets that do not generate events.

▪ **Event Data**

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

- R: Reading value
- T: Threshold value

The following table lists the core sensors located within the mBMC. These sensors are fixed and hard-coded. They cannot be modified by a user.

Table 79. mBMC Built-in Sensors

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value/Offsets	Event Data	PEF Action	SDR Record Type
Physical Security Violation	01	Physical Security 05h	Sensor Specific 6Fh	LAN Leash Lost	As	LAN Leash Lost	Trig Offset		
Platform Security Violation	02	Platform Security Violation Attempt 06h	Sensor Specific 6Fh	Out-of-band access password violation	As	–	Trig Offset		
Power Unit Status	03	Power Unit 09h	Sensor Specific 6Fh	Power On/Off Power cycle AC Lost	As	–	Trig Offset		

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value/Offsets	Event Data	PEF Action	SDR Record Type
Button	04h	Button 14h	Sensor Specific 6Fh	Power Button Reset Button	As	–	Trig Offset		
Watchdog	05h	Wtchdog2 23h	Sensor Specific 6Fh	Timer Expired Hard Reset Power Down Power cycle Timer Interrupt	As	–	Trig Offset		
System Boot	06h	System boot Initiated 1Dh	Sensor Specific 6Fh	Initiated by power up Initiated by hard reset Initiated by warm reset	As	–	Trig Offset		
System PEF Event	07h	System Event 12h	Sensor Specific 6Fh	PEF Action	As	–	Trig Offset		
Platform Allert	08h	Platform Alert 24h	Sensor Specific 6Fh	Platform Event Trap generated	As	–	Trig Offset		
Physical Security Violation	09h	Physical Security 05h	Sensor Specific 6Fh	General Chassis Intrusion	As & De	General Chassis Intrusion	Trig Offset		

The following table shows the platform sensors that are supported by the mBMC.

Table 80. Intel® Server Board SE7520AF2 Sensors for OB Platform Instrumentation Management

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value/Offsets	Event Data	PEF Action	SDR Record Type
Physical Security Violation	0Ah	Physical Security 05h	Sensor Specific 6Fh	General Chassis Intrusion	As	General Chassis Intrusion	Trig Offset	X	02
CPU1 12v	0Bh	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
CPU2 12v	0Ch	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
BB +1.5V	0Dh	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
BB +1.8V	0Eh	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value/Offsets	Event Data	PEF Action	SDR Record Type
BB +3.3V	0Fh	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
BB +5V	10h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
BB +12V	11h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
BB -12V	12h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
FSB Vtt	13h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
MCH Vtt	14h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
SCSI Core(1.8v)	15h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc1 VCCP	16h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc2 VCCP	17h	Voltage 02h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 1	18h	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 2	19h	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 3	1Ah	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 4	1Bh	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 5	1Ch	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Tach Fan 6	1Dh	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc1 IERR	1Eh	Processor 07h	Sensor Specific 6Fh	IERR	As	–	Trig Offset	–	02
Proc2 IERR	1Fh	Processor 07h	Sensor Specific 6Fh	IERR	As	–	Trig Offset	–	02
Proc1 Thermal trip	20h	Processor 07h	Sensor Specific 6Fh	Thermal Trip	As	–	Trig Offset	Fault LED Action	02
Proc2 Thermal trip	21h	Processor 07h	Sensor Specific 6Fh	Thermal Trip	As	–	Trig Offset	Fault LED Action	02
CPU 1 Thermal Throttle	22h	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value/Offsets	Event Data	PEF Action	SDR Record Type
CPU 2 Thermal Throttle	23h	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Diagnostic Interrupt Button	24h	Critical Interrupt 13h	Sensor Specific 6Fh	FP NMI Button	As	–	Trig Offset	NMI Pulse	02
Chassis Identify Button	25h	Button 14h	Generic 03h	Sate Deasserted State Assert	As & De	–	Trig Offset	ID LED Action	02
Proc1 Fan	26h	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc2 Fan	27h	Fan 04h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc1 Core temp	28h	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
Proc2 Core temp	29h	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	R, T	Fault LED Action	01
CPU Configuration Error	2Ah	Processor 07h	Generic 03h	State Asserted	As & De	Discrete	R, T	Fault LED Action	02
Baseboard Temp 1	2Bh	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	Trig Offset	Fault LED Action	01
Baseboard Temp 2	2Ch	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	Trig Offset	Fault LED Action	01
Front Panel Temp	2Dh	Temp 01h	Threshold 01h	[u,l][nr, c,nc]	As & De	Analog	Trig Offset	Fault LED Action	01

The following table shows the platform sensors that are supported by the Intel® Management Module (IMM) Professional and Advanced Editions when installed on the Intel® Server Board SE7520AF2 and updated with SE7520AF2 platform specific firmware.

Table 81. Platform Sensors for the Intel® Management Module

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Power Unit Status	01h	Power Unit 09h	Sensor Specific 6Fh	Power Off Power Cycle A/C Lost Soft Power Control Fault Power Unit Failure Predictive Failure	As	–	Trig Offset	A	X

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Power Unit Redundancy	02h	Power Unit 09h	Generic 0Bh	Redundancy Regained Redundancy lost Redundancy Degraded Non-red:Suff res from redund Non-red:Suff res from insuff res Non-red:Insuff res Redundancy Degraded from full redundancy Redundancy Degraded from non-redundant	As	–	Trig Offset	A	X
Watchdog	03h	Watchdog2 23h	Sensor Specific 6Fh	Timer Expired Hard Reset Power Down Power Cycle Timer Interrupt	As & De	–	Trig Offset	A	X
Platform Security Violation	04h	Platform Security Violation Attempt 06h	Sensor Specific 6Fh	Secure mode violation attempt Out-of-band access password violation	As	–	Trig Offset	A	X
Physical Security Violation	05h	Physical Security 05h	Sensor Specific 6Fh	General Chassis Intrusion LAN Leash Lost	As & De	General Chassis Intrusion LAN Leash Lost	Trig Offset	A	X
POST Error	06h	POST error 0Fh	Sensor Specific 6Fh	POST error	As	–	POST Code	A	–
Critical Interrupt Sensor	07h	Critical Interrupt 13h	Sensor Specific 6Fh	Front Panel NMI Bus Error	As & De	–	Trig Offset	A	–
Memory	08h	Memory 0Ch	Sensor Specific 6Fh	Uncorrectable ECC	As	–	Trig Offset	A	–
Event Logging Disabled	09h	Event Logging Disabled 10h	Sensor Specific 6Fh	Correctable Memory Error Logging Disabled Log Area Reset/Cleared	As	–	Trig Offset	A	X

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Session Audit	0Ah	Session Audit 2Ah	Sensor Specific 6Fh	00: Session Activation 01: Session Deactivation	As	–	As defined by IPMI	A	X
BB +1.2V Vtt	10h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +1.5V	11h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +1.5V PXH Core	12h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	X
BB +1.5V NIC Core	13h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +1.8V	14h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +1.8V SCSI Core	15h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +3.3V	16h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +3.3V Standby	17h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +5V	18h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB +12V	19h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB -12V	1Ah	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB Mem Vtt	1Bh	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
BB Vbat	1Ch	Battery 29h	Sensor Specific 6Fh	Battery Low Battery Failed Battery Presence Detected	As	–	Trig Offset	A	X
BB Temp	30h	Temp 01h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	X
Front Panel Temp	32h	Temp 01h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	X
Tach Fan 1	40h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Tach Fan 2	41h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Tach Fan 3	42h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Tach Fan 4	43h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Tach Fan 5	44h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Tach Fan 6	45h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Fan Redundancy	4Ah	Fan 04h	Discrete 0Bh	Redundancy Regained Redundancy lost Redundancy Degraded Non-red:Suff res from redund Non-red:Suff res from insuff res Non-red:Insuff res Redundancy Degraded from full redundancy Redundancy Degraded from non-redundant	As	–	Trig Offset	A	X
Fan 1 Presence	4Bh	Slot/Connector 21h	Sensor Specific 6Fh	Device installed	As & De	–	Trig Offset	A	–
Fan 2 Presence	4Ch	Slot/Connector 21h	Sensor Specific 6Fh	Device installed	As & De	–	Trig Offset	A	–
Fan 3 Presence	4Dh	Slot/Connector 21h	Sensor Specific 6Fh	Device installed	As & De	–	Trig Offset	A	–
Fan 4 Presence	4Eh	Slot/Connector 21h	Sensor Specific 6Fh	Device installed	As & De	–	Trig Offset	A	–
LVDS SCSI channel 1 terminator fault	60h	Terminator 1Ch	Digital Discrete 06h	Performance Met or Lags	As	–	Trig Offset	A	–
LVDS SCSI channel 2 terminator fault	61h	Terminator 1Ch	Digital Discrete 06h	Performance Met or Lags	As	–	Trig Offset	A	–
Power Supply Status 1	70h	Power Supply 08h	Sensor Specific 6Fh	Presence Failure Predictive Fail A/C Lost	As & De	–	Trig Offset	A	X
Power Supply Status 2	71h	Power Supply 08h	Sensor Specific 6Fh	Presence Failure Predictive Fail A/C Lost	As & De	–	Trig Offset	A	X
Power Nozzle Power Supply 1	78h	Current 03h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Power Nozzle Power Supply 2	79h	Current 03h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Power Gauge V1 rail (+12v) Power Supply 1	7Ah	Current 03h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Power Gauge V1 rail (+12v) Power Supply 2	7Bh	Current 03h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Power Gauge (aggregate power) Power Supply 1	7Ch	Other Units 0Bh	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Power Gauge (aggregate power) Power Supply 2	7Dh	Other Units 0Bh	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor Missing	80h	Module / Board 15h	Digital Discrete 03h	State Asserted	As	–	Trig Offset	A	–
System ACPI Power State	82h	System ACPI Power State 22h	Sensor Specific 6Fh	S0 / G0 S1 S2 S3 S4 S5 / G2 S4/S5 soft-off G3 Mechanical Off G1 S5 Legacy On Legacy Off Unknown	As	–	Trig Offset	A	X
System Event	83h	System Event 12h	Sensor Specific 6Fh	OEM System Boot Event (Hard Reset) PEF Action	As	–	Trig Offset	A	–
Button	84h	Button 14h	Sensor Specific 6Fh	Power Button Sleep Button Reset Button	As	–	Trig Offset	A	X
SMI Timeout	85h	SMI Timeout F3h	Digital Discrete 03h	State Asserted	As	–	Trig Offset	A	–
NMI Signal State	87h	OEM C0h	Digital Discrete 03h	State Asserted	–	–	–	–	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
SMI Signal State	88h	OEM C0h	Digital Discrete 03h	State Asserted	–	–	–	–	–
Processor 1 Status	90h	Processor 07h	Sensor Specific 6Fh	IERR Thermal Trip FRB1, FRB2, FRB3 Config Error Presence Disabled	As & De	–	Trig Offset	M	X
Processor 2 Status	91h	Processor 07h	Sensor Specific 6Fh	IERR Thermal Trip FRB1, FRB2, FRB3 Config Error Presence Disabled	As & De	–	Trig Offset	M	X
Processor 1 Core Temp	98h	Temp 01h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor 2 Core Temp	99h	Temp 01h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor 1 Fan	A8h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Processor 2 Fan	A9h	Fan 04h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	M	–
Processor 1 12v VRM	B8h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor 2 12v VRM	B9h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor 1 Thermal Control	C0h	Temp 01h	Digital Discrete 07h	Transitioned to OK Transitioned to Non-Critical from OK	As & De	–	Trig Offset	M	–
Processor 2 Thermal Control	C1h	Temp 01h	Digital Discrete 07h	Transitioned to OK Transitioned to Non-Critical from OK	As & De	–	Trig Offset	M	–
Processor 1 VRD Over Temp	C8h	Temp 01h	Digital Discrete 07h	Transitioned to OK Transitioned to Non-Critical from OK	As & De	–	Trig Offset	M	–
Processor 2 VRD Over Temp	C9h	Temp 01h	Digital Discrete 07h	Transitioned to OK Transitioned to Non-Critical from OK	As & De	–	Trig Offset	M	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
Processor 1 Vcc	D0h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
Processor 2 Vcc	D1h	Voltage 02h	Threshold 01h	[u,l][nr,c,nc]	As & De	Analog	R, T	A	–
CPU Configuration Error	D8h	Processor 07h	Generic 03h	State Asserted	As & De	Discrete	R, T	A	-
DIMM 1	E0h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 2	E1h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 3	E2h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 4	E3h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 5	E4h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 6	E5h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
DIMM 7	E6h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
DIMM 8	E7h	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Disabled Slot Holds Spare Device	As	–	Trig Offset	A	–
PCI-X Slot 1 Hot-plug Status	ECh	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Device ready for removal Slot power is off	As & De	–	Trig Offset	A	–
PCI-Express Slot 3 Hot-plug Status	EDh	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Device ready for removal Slot power is off	As & De	–	Trig Offset	A	–
PCI-Express Slot 4 Hot-plug Status	EEh	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Device ready for removal Slot power is off	As & De	–	Trig Offset	A	–
PCI-X Slot 5 Hot-plug Status	EFh	Slot Connector 21h	Sensor Specific 6Fh	Fault Status Asserted Device Installed Device ready for removal Slot power is off	As & De	–	Trig Offset	A	–
DIMM Domain 1 Sparing Enabled	F0h	Entity Presence 25h	Sensor Specific 6Fh	Entity Present Entity Absent	As	–	Trig Offset	A	–
DIMM Domain 1 Sparing Redundancy	F1h	Memory 0Ch	Discrete 0Bh	Fully Redundant Non-red:Suff res from redund Non-red:Insuff res	As	–	Trig Offset	A	–
DIMM Domain 2 Sparing Enabled	F2h	Entity Presence 25h	Sensor Specific 6Fh	Entity Present Entity Absent	As	–	Trig Offset	A	–

Sensor Name	Sensor #	Sensor Type	Event / Reading Type	Event Offset Triggers	Assert / Deassert	Readable Value / Offsets	EventData	Rearm	Standby
DIMM Domain 2 Sparing Redundancy	F3h	Memory 0Ch	Discrete 0Bh	Fully Redundant Non-red:Suff res from redund Non-red:Insuff res	As	–	Trig Offset	A	–

7. Error Reporting and Handling

This section covers general information regarding Error reporting, error logging and handling by BIOS and firmware.

The BIOS of the Intel® Server Board SE7520AF2 indicates the current testing phase during POST by writing a hex code to I/O location 80h. If errors are encountered, error messages or codes will either be displayed to the video screen, or if an error has occurred prior to video initialization, errors will be reported through a series of audio beep codes.

The error codes are defined by Intel and whenever possible are backward compatible with error codes used on earlier platforms.

7.1 Fault Resilient Booting (FRB)

It should be noted that processors will not be disabled if the system is using the mBMC, this advanced resilience server feature is only available with an Intel® Management Module installed.

7.1.1 FRB-3 – BSP Reset Failures

The BIOS and firmware provides a feature to guarantee that the system boots, even if one or more processors fail during POST. The BMC contains two watchdog timers that can be configured to reset the system upon time-out. The first timer (FRB-3) starts counting down whenever the system comes out of hard reset. If the Boot Strap Processor (BSP) successfully resets and begins executing, the BIOS disables the FRB-3 timer in the BMC and the system continues executing POST. If the timer expires because of the BSP's failure to fetch or execute BIOS code, the BMC resets the system and disables the failed processor. The BMC continues to change the bootstrap processor until the BIOS successfully disables the FRB-3 timer. The BMC sounds beep codes on the system speaker, if it fails to find a good processor. It will continue to cycle until it finds a good processor. The process of cycling through all the processors is repeated upon system reset or power cycle. Soft resets do not affect the FRB-3 timer. The duration of the FRB3 timer is set by system firmware. The mBMC also supports the algorithm described above, with the exception that it does not disable the processor.

7.1.2 FRB-2 – BSP POST Failures

The second timer (FRB-2) is set to several minutes by BIOS and is designed to guarantee that the system completes POST. The FRB-2 timer is enabled just before the FRB-3 timer is disabled to prevent any "unprotected" window of time. Near the end of POST, the BIOS disables the FRB-2 timer. If the system contains more than 1 GB of memory and the user chooses to test every DWORD of memory, the watchdog timer is extended before the extended memory test starts, because the memory test can exceed the timer duration. The BIOS will also disable watchdog timer before prompting the user for a boot password. If the system hangs during POST, before the BIOS disables the FRB-2 timer, the BMC generates an asynchronous system reset (ASR). The BMC retains status bits that can be read by BIOS later in the POST for the purpose of disabling the previously failing processor, logging the appropriate event into the System Event Log (SEL), and displaying an appropriate error message to the user.

Options are provided by the BIOS to control the policy applied to FRB-2 failures. By default, an FRB-2 failure results in the failing processor being disabled during the next reboot. This policy can be overridden to prevent BSP from ever being disabled due to the FRB-2 failure or a policy resulting in disabling the BSP after three consecutive FRB-2 failures can be selected. These options may be useful in systems that experience fatal errors during POST that are not indicative of a bad processor. Selection of this policy should be considered an advanced feature and should only be modified by a qualified system administrator. If supported by the specific platform, these options can be found in BIOS Setup. The mBMC does not support the option to disable the BSP.

7.1.3 FRB-1 – BSP Self-Test Failures

In addition to FRB-3 and FRB-2 timers, the BIOS provides FRB-1. Early in POST, the BIOS checks the Built-in Self Test (BIST) results of the BSP. If the BSP fails BIST, the BIOS requests the BMC to disable the BSP. The BMC disables the BSP, selects a new BSP and generates a system reset. If there is no alternate processor available, the BMC beeps the system speaker and halts the system. If the Sahalee BMC is not installed, then BIOS can only notify the user that BIST failed, no processors will be disabled.

The BIST failure is indicated to the user by displaying during POST and logging an error in the System Event Log (SEL).

7.1.4 OS Watchdog Timer: Operating System Load Failures

The OS Watchdog Timer feature is designed to allow watchdog timer protection of the operating system load process. This is done in conjunction with an operating system-present device driver or application that will disable the watchdog timer once the operating system has successfully loaded. If the operating system load process fails, the BMC will reset the system.

BIOS shall disable the OS Watchdog Timer before handing control to the OS loader if it is determined to be booting from removable media or the BIOS cannot determine the media type.

If the BIOS is going to boot to a known HDD, it will read a user option for OS Watchdog timer for HDD Boots. If this is disabled, the BIOS will ensure the watchdog timer is disabled and boot. Otherwise the BIOS will read the enabled time value from the option and set the OS Watchdog timer for that value (5, 10, 15, or 20 minutes) before trying to load the operating system. If the OS Watchdog timer is enabled, the timer is repurposed as an OS Watchdog timer and is referred to by that title as well. **WARNING:** The BIOS may incorrectly determine that a removable media is a HDD if the media emulates a HDD. In this case the OS Watchdog timer will not be automatically disabled.

If the BIOS is going to boot to a known PXE compliant device, then the BIOS reads a user option for OS Watchdog timer for PXE Boots and either disables the timer or enables the timer with a value read from the option (5, 10, 15, or 20 minutes). If the OS Watchdog timer is enabled, the timer is repurposed as an OS Watchdog timer and is referred to by that title as well.

If the OS Watchdog timer is enabled and if a boot password is enabled, the BIOS will disable the OS Watchdog timer before prompting the user for a boot password regardless of the OS Watchdog timer option setting. Also if the user has chosen to enter BIOS setup, the timer will be

disabled regardless of option settings. Otherwise, if the system hangs during POST, before the BIOS disables the timer, the BMC generates an asynchronous system reset (ASR). The BMC retains status bits that can be read by BIOS later in the POST for the purpose of disabling the previously failing processor, logging the appropriate event into the SEL, and displaying an appropriate error message to the user. As the timer may be repurposed, the BIOS and BMC will also keep track of which timer expired (early FRB-2, late FRB-2, or OS Watchdog) and display the appropriate error message to the user.

All of the user options are intended to allow a system administrator to setup a system such that during a normal boot no gap exists during POST that is not covered by the watchdog timer. Options are provided by the BIOS to control the policy applied to OS Watchdog timer failures. By default, an OS Watchdog timer failure will not cause any action. Other options provided by the BIOS are for the system to reset or power off watchdog timer failure. However, it should be noted that these failures will NOT result in a processor being disabled (as could happen with an FRB-2 failure).

7.1.5 AP Failures

The BIOS and BMC implement additional safeguards to detect and disable the application processors (AP) in a multiprocessor system. If an AP fails to complete initialization within a certain time, it is assumed to be nonfunctional. If the BIOS detects that an AP has failed BIST or is nonfunctional, it requests the BMC to disable that processor. Processors disabled by the BMC are not available for use by the BIOS or the operating system. Since the processors are unavailable, they are not listed in any configuration tables including SMBIOS tables.

7.1.6 Treatment of Failed Processors

All the failures (FRB-3, FRB-2, FRB-1, and AP failures) including the failing processor are recorded into the system event log. The FRB-3 failure is recorded automatically by the BMC while the FRB-2, FRB-1, and AP failures are logged to the SEL by the BIOS. In the case of an FRB-2 failure, some systems will log additional information into the OEM data byte fields of the SEL entry. This additional data indicates the last POST task that was executed before the FRB-2 timer expired. This information may be useful for failure analysis.

The BMC maintains failure history for each processor in nonvolatile storage. This history is used to store a processor's track record. Once a processor is marked "failed," it remains "failed" until the user forces the system to retest the processor by entering BIOS Setup and selecting the "Processor Retest" option. The BIOS reminds the user about a previous processor failure during each boot cycle until all processors have been retested and successfully pass the FRB tests or AP initialization. If all the processors are bad, the system does not alter the BSP and attempts to boot from the original BSP. Error messages are displayed on the console, and errors are logged in the event log of a processor failure.

If the user replaces a processor that has been marked bad by the system, the system must be informed about this change by running BIOS Setup and selecting that processor to be retested. If a bad processor is removed from the system and is replaced with a terminator module, the BMC automatically detects this condition and clears the status flag for that processor during the next boot.

Three states are possible for each processor slot:

1. Processor installed (status only, indicates processor has passed BIOS POST).
2. Processor failed. The processor may have failed FRB-2, FRB-3, or BIST, and it has been disabled.
3. Processor not installed (status only, indicates the processor slot has no processor in it).

Additional information on FRB may be found in the Sahalee Baseboard Management Controller EPS.

7.2 Error Logging

This section defines how errors are handled by the system BIOS. Also discussed is the role of the BIOS in error handling and the interaction between the BIOS, platform hardware, and server management firmware with regard to error handling. In addition, error-logging techniques are described and beep codes for errors are defined.

7.2.1 Error Sources and Types

One of the major requirements of server management is to correctly and consistently handle system errors. System errors can be categorized as follows:

- PCI bus
- Memory multi-bit errors (single-bit errors are not logged)
- Sensors
- Processor internal errors, bus/address errors, thermal trip errors, temperatures and voltages, and GTL voltage levels
- Errors detected during POST, logged as POST errors

Sensors are managed by the mBMC. The mBMC is capable of receiving event messages from individual sensors and logging system events.

7.2.2 SMI Handler

The SMI handler handles and logs system-level events that are not visible to the server management firmware. If SEL error logging is disabled in the BIOS Setup utility, no SMI signals are generated on system errors. If error logging is enabled, the SMI handler preprocesses all system errors, even those that are normally considered to generate an NMI.

The SMI handler sends a command to the BMC to log the event and provides the data to be logged. For example, The BIOS programs the hardware to generate an SMI on a single-bit memory error and logs the location of the failed DIMM in the system event log. System events that are handled by the BIOS generate SMI.

7.2.2.1 PCI Bus Error

The PCI bus defines two error pins, PERR# and SERR#, for reporting PCI parity errors and system errors, respectively. The BIOS can be instructed to enable or disable reporting PERR#

and SERR# through NMI. Disabling NMI for PERR# and/or SERR# also disables logging of the corresponding event. In the case of PERR#, the PCI bus master has the option to retry the offending transaction, or to report it using SERR#. All other PCI-related errors are reported by SERR#. All the PCI-to-PCI bridges are configured so that they generate SERR# on the primary interface whenever there is SERR# on the secondary side, if SERR# is enabled through Setup. The same is true for PERR#. The format of the data bytes is described above.

7.2.2.2 Processor Bus Error

If the chipset supports ECC on the processor bus, the BIOS will enable the error correction and detection capabilities of the processors by setting appropriate bits in processor model specific register (MSR) and appropriate bits inside the chipset.

In the case of irrecoverable errors on the host processor bus, proper execution of the asynchronous error handler (usually SMI) cannot be guaranteed and the handler cannot be relied upon to log such conditions. The handler will record the error to the system event log only if the system has not experienced a catastrophic failure that compromises the integrity of the handler.

7.2.2.3 Memory Bus Error

The hardware is programmed to generate an SMI on single-bit data errors in the memory array if ECC memory is installed. The SMI handler records the error and the DIMM location to the system event log. Double-bit errors in the memory array are mapped to SMI because the mBMC cannot determine the location of the bad DIMM. The double-bit errors may have corrupted the contents of SMRAM. The SMI handler will log the failing DIMM number to the mBMC if the SMRAM contents are still valid. The ability to isolate the failure down to a single DIMM may not be available on certain platforms, and/or during early POST.

7.2.2.4 System Limit Error

The BMC monitors system operational limits. It manages the A/D converter, defining voltage and temperature limits as well as fan sensors and chassis intrusion. Any sensor values outside of specified limits are fully handled by the BMC. The BIOS does not generate an SMI to the host processor for these types of system events.

Refer to the platform's Server Management External Architecture Specification for details on various sensors and how they are managed.

7.2.2.5 Processor Failure

The BIOS detects processor BIST failure and logs this event. The failed processor can be identified by the first OEM data byte field in the log. For example, if processor 0 fails, the first OEM data byte will be 0. The BIOS will depend upon BMC to log the watchdog timer reset event.

If an operating system device driver is using the watchdog timer to detect software or hardware failures and that timer expires, an Asynchronous Reset (ASR) is generated, which is equivalent to a hard reset. The POST portion of the BIOS can query the BMC for watchdog reset event as the system reboots, and logs this event in the SEL.

7.2.2.6 Boot Event

The BIOS downloads the system date and time to the BMC during POST and logs a boot event. This record does not indicate an error, and software that parses the event log should treat it as such.

7.2.2.7 Logging Format Conventions

The BIOS event log data in the SEL complies with the IPMI specification. IPMI requires use of all but two bytes in each event log entry, called Event Data 2 and Event Data 3. An event generator can specify that these bytes contain OEM-specified values. The system BIOS uses these two bytes to record additional information about the error.

The format of the OEM data bytes (Event Data 2 and Event Data 3) for memory errors, PCI bus errors and FRB2 errors is described here. This format is supported by all platforms that are IPMI version 1.0 (or later) compliant.

Bits 3:1 of the generator ID field define the format revision. The system software ID is a 7-bit quantity. For events covered in this document, the system software ID's will be within the range 0x18-0x1F. System software ID of 0x18 indicates that OEM data byte 2 and 3 are encoded using data format scheme revision 0. Note that the system software IDs in the range 0x10-0x1f are reserved for the SMI handler. The IPMI specification reserves two distinct ranges for BIOS and the SMI handler. Since the distinction between the two is not very important, we use the same values of generator ID's for the BIOS as well as the SMI handler. Technically, the FRB-2 event is not logged by the SMI handler, but it will use the same generator ID range as memory errors.

Table 82. Memory Error Events

Field	IPMI definition	BIOS-Specific Implementation
Generator ID	7:1 System software ID or IPMB slave address. 1=ID is system software ID 0=ID is IPMB slave address	7:4 0x3 for system BIOS 3:1 0 Format revision, Revision of the data format for OEM data bytes 2 and 3, For this revision of the specification, set this field to 0. All other revisions are reserved. 0 1 = ID is system software ID. As a result, the generator ID byte will start from 0x31 and go up to 0x3f, in increments of 2 for events logged by the BIOS.
Sensor Type	See Table 30.3 in the Intelligent Platform Management Specification, Version 1.5.	0xC for memory errors
Sensor Number	Number of sensor that generated this event	Unique value for each type of event because IPMI specification requires it that way. This field has no other significance. Should not be displayed to the end user if the event is logged by BIOS.
Type code	0x6F if event offsets are specific to the sensor	0x6F

Field	IPMI definition	BIOS-Specific Implementation
Event Data 1	7:6 00 = unspecified byte 2 10 = OEM code in byte 2. 5:4 00 = unspecified byte 3 10 = OEM code in byte 3. (The BIOS will not use encodings 01 and 11 for errors covered by this document.). 3:0 Offset from Event Trigger for discrete event state.	Follow IPMI definition. If either of the two data bytes following this do not have any data, that byte should be set to 0xff, and the appropriate field in event data 1 should indicate that that it is unspecified. According to Table 30.3 in the Intelligent Platform Management Specification, Version 1.5, 3:0 is 0 for single bit error and 1 for multi-bit error.
Event Data 2	7:0 OEM code 2 or unspecified.	For format rev 0, if this byte is specified, 7:6 Zero-based memory card number. Matches the number of type 16 entry in SMBIOS table. For example, card 0 corresponds to the first type 16 entry in SMBIOS tables. If all DIMMs are onboard, this field will always be 0.5:0 zero-based DIMM number on the card. DIMM 0 corresponds to the first type 17 record in SMBIOS tables for that memory card.
Event Data 3	7:0 OEM code 3 or unspecified.	If format rev is 0 and if this byte is specified, syndrome byte.

Table 83. Examples of Event Data Field Contents for Memory Errors

Error Type	Event Data 1	Event Data 2	Event Data 3
Single bit error; no information about the error is available.	00	0xFF	0xFF
Multi bit memory error, failed DIMM is the fifth DIMM on the second memory card.	0x81	0x44 (Bits 7:6 = 01 Bits 5:0 = 04)	0xFF
Single bit error. Syndrome is 0x54, DIMM location is not known.	0x20	0xFF	0x54
Multi-bit error, Syndrome byte is 0x1c. DIMMs are onboard, and the second DIMM has failed.	0xA1	0x01 (Bits 7:6 = 00 Bits 5:0 = 01)	0x1C

Table 84. PCI error events

Field	IPMI Definition	BIOS Specific Implementation
Generator ID	7:1 System software ID or IPMB slave address. 1=ID is system software ID; 0=ID is IPMB slave address.	7:4 0x3 for system BIOS 3:1 0 Format revision, Revision of the data format for OEM data bytes 2 and 3, For this revision of the specification, set this field to 0. All other revisions are reserved for now. 0 1=ID is system software ID As a result, the generator ID byte will start from 0x31 and go up to 0x3f, in increments of 2 for events logged by the BIOS.
Sensor Type	See Table 30.3 in <i>Intelligent Platform Management Interface Specification</i>	0x13 for critical interrupt
Sensor number	Number of sensor that generated this event	Unique value for each type of event because IPMI specification requires it that way. This field has no other significance. Should not be displayed to the end user if the event is logged by BIOS.
Type code	0x6F if event offsets are specific to the sensor	0x6F

Event Data 1	7:6 00 = unspecified byte 2; 10 = OEM code in byte 2. 5:4 00 = unspecified byte 3; 10 = OEM code in byte 3 (BIOS will not use encodings 01 and 11 for errors covered by this document.) 3:0 Offset from Event Trigger for discrete event state.	Follow IPMI definition. If either of the two data bytes following this do not have any data, that byte should be set to 0xff, and the appropriate field in event data 1 should indicate that that it is unspecified. According to Table 30.3 in <i>Intelligent Platform Management Interface Specification</i> , 3:0 is 04 for PCI PERR and 05 for PCI SERR.
Event Data 2	7:0 OEM code 2 or unspecified	For format rev 0, if this byte is specified, it contains the PCI bus number on which the failing device resides. If the source of the PCI error cannot be determined, this byte contains 0xff and the event data 1 byte indicates that byte 2 is unspecified.
Event Data 3	7:0 OEM code 3 or unspecified.	For format rev 0, if this byte is specified, it contains the PCI device/function address in the standard format: 7:3 Device number of the failing PCI device 2:0 PCI function number. Will always contain a zero if the device is not a multifunction device. If the source of the PCI error cannot be determined, this byte contains 0xff and the event data 1 byte indicates that byte 3 is unspecified.

7.2.2.8 Examples of Event Data Field Contents for PCI Errors

Table 85. Event Data Field Contents for PCI Errors

Error Type	Event Data 1	Event Data 2	Event Data 3
PCI PERR, failing device is not known	04	0xFF	0xFF
PCI SERR, failing device is not known	05	0xFF	0xFF
PCI PERR, device 3, function 1 on PCI bus 5 reported the error	0xA4	0x05	0x19 (Bits 7:3 = 03 Bits 2:0 = 01)
An unknown device on PCI bus 0 reported the SERR	0x85	0x00	0xFF

7.2.2.9 FRB-2 Error Events

Table 86. FRB-2 Error Events

Field	IPMI Definition	BIOS Specific Implementation
Generator ID	7:1 System software ID or IPMB slave address. 1=ID is system software ID; 0=ID is IPMB slave address.	7:4 0x3 for system BIOS 3:1 0 Format revision, Revision of the data format for OEM data bytes 2 and 3, For this revision of the specification, set this field to 0. All other revisions are reserved for now. 0 1=ID is system software ID As a result, the generator ID byte will start from 0x31 and go up to 0x3f, in increments of 2 for events logged by the BIOS.
Sensor Type	See Table 30.3 in <i>Intelligent Platform Management Interface Specification</i> , v1.5.	0x7 for processor related errors

Sensor number	Number of sensor that generated this event	Unique value for each type of event because IPMI specification requires that. This field has no other significance, and it should not be displayed to the end user if the event is logged by BIOS.
Type code	0x6F if event offsets are specific to the sensor	0x6F
Event Data 1	7:6 00 = unspecified byte 2; 10 = OEM code in byte 2. 5:4 00 = unspecified byte 3; 10 = OEM code in byte 3. (BIOS will not use encodings 01 and 11 for errors covered by this document.) 3:0 Offset from Event Trigger for discrete event state.	If Event data 2 and event data 3 contain OEM codes, bits 7:6 and bits 5:4 contain 10. For platforms that do not include the POST code information with FRB-2 log, both these fields will be 0. BIOS either should specify both bytes or should mark both bytes as unspecified. According to IPMI 1.0 specification, table 30.3, Byte 3:0 is 03 for FRB-2 failure during POST.
Event Data 2	7:0 OEM code 2 or unspecified.	For format rev 0, if this byte is specified, it contains bits 7:0 of the POST code at the time FRB-2 reset occurred (port 80 code)
Event Data 3	7:0 OEM code 3 or unspecified.	For format rev 0, if this byte is specified, it contains bits 15:8 of the POST code at the time FRB-2 reset occurred (port 81 code). If the BIOS only uses one byte POST codes, this byte will always be zero.

7.2.2.10 Examples of Event Data Field Contents for FRB-2 Errors

Table 87. Event Data Field Contents for FRB-2 Errors

Error type	Event Data 1	Event Data 2	Event Data 3
FRB-2 error, failing POST code information not available	0x03	0xFF	0xFF
FRB-2 error, BIOS uses one byte POST codes. The last POST code before FRB-2 reset was 0x60.	0xA3	0x60	0x0
FRB-2 error, BIOS uses one byte POST codes. The last POST code before FRB-2 reset was 0x1942.	0xA3	0x42	0x19

7.2.3 BIOS Generated IPMI Events

The below table details the events that are initiated by the BIOS and the values that these events should use to generate SEL entries.

Table 88. BIOS Generated IPMI Events

Sensor Type	Generator ID	EMV Rev	Sensor Type Code	Sensor number	Type code	Sensor-Specific Offset	Event Data1	Event Data 2,3	Event
Processor	31h	04h	07h	90h	6Fh	03h	A3h	Data2 :port80 Data3 :port81	FRB2/POST Failure
Memory	33h	04h	0Ch	08h	6Fh	00h	A0h	Data2: DIMM location Data3:	Correctable ECC

								Synrome	
Memory	33h	04h	0Ch	08h	6Fh	01h	A1h	Data2: DIMM location Data3: Synrome	Uncorrectabl e ECC
POST Error	31h	04h	0Fh	06h	6Fh	Data2[0: 3]	A0h or Data2	See POST error code table7.3.3. Data2:low byte Data3:hig h byte	POST Error
System Event	03h	04h	12h	83h	6Fh	05h	05h	Data2: 00h(1 st of pair) 80h(2 nd of pair) Data3:0ffh	Timestamp Clock Sync
System Event	03h	04h	12h	83h	6Fh	01h	01h	Data2:0ffh Data3:0ffh	Boot event log
Critical Interrupt	31h	04h	13h	EAh	6Fh	04h	A4h	Data2:bus num Data3:dev and fun	PCI SERR
Critical Interrupt	31h	04h	13h	EBh	6Fh	05h	A5h	Data2:bus num Data3:dev and fun	PCI PERR

7.2.4 Single Bit ECC Error Throttling Prevention

The system detects, corrects, and logs correctable errors. As long as these errors occur infrequently, the system should continue to operate without a problem.

Occasionally, correctable errors are caused by a persistent failure of a single component. For example, a broken data line on a DIMM would exhibit repeated errors until replaced. Although these errors are correctable, continual calls to the error logger can throttle the system, preventing any further useful work.

For this reason, the system counts certain types of correctable errors and disables reporting if they occur too frequently. Correction remains enabled but calls to the error handler are disabled. This allows the system to continue running, despite a persistent correctable failure. The BIOS adds an entry to the event log to indicate that logging for that type of error has been disabled. Such an entry indicates a serious hardware problem that must be repaired at the earliest possible time.

The system BIOS implements this feature for two types of errors, correctable memory errors and correctable bus errors. If ten errors occur in a single wall-clock hour, the corresponding

error handler disables further reporting of that type of error. A unique counter is used for each type of error; i.e., an overrun of memory errors does not affect bus error reporting.

The BIOS re-enables logging and SMLs the next time the system is rebooted.

7.3 Error Messages and Error Codes

The system BIOS displays error messages on the video screen. Prior to video initialization, beep codes inform the user of errors. POST error codes are logged in the event log. The BIOS displays POST error messages on the video monitor. For some errors BIOS may display POST error codes as well.

7.3.1 POST Progress Codes and Messages

7.3.1.1 System ROM BIOS POST Task Test Point (Port 80h Code)

The BIOS sends a 1-byte hex code to port 80 before each task. The port 80 codes provide a troubleshooting method in the event of a system hang during POST. Below is the table of the Port 80 codes and the corresponding task description.

7.3.1.2 Diagnostic LEDs

The value of port 80h will be sent to four tri-color LEDs. The diagnostic LED feature consists of a hardware decoder and four dual color LEDs located on the baseboard. During POST, the LEDs will display all normal POST progress codes representing the progress of the BIOS POST. Each code will be represented by a combination of colors from the four LEDs.

The LEDs are in pairs of green and red. The POST progress codes are divided into two nibbles, an upper nibble and a lower nibble. Each bit in the upper nibble is represented by a red LED and each bit in the lower nibble is represented by a green LED. If both bits are set in the upper nibble and in the lower nibble then both red and green LEDs are lit, resulting in an amber color. If both bits are clear, then both the red and green LEDs are off.

In the below example, BIOS sends a value of ACh to the LEDs. The LEDs are decoded as follows:

- Red bits = 1010b = Ah
- Green bits = 1100b = Ch

Since the red bits correspond to the upper nibble and the green bits correspond to the lower nibble, the two are concatenated to be ACh.

Table 89. POST Progress Code LED Example

LEDs	Red	Green	Red	Green	Red	Green	Red	Green
ACh	1	1	0	1	1	0	0	0
Result	Amber		Green		Red		Off	

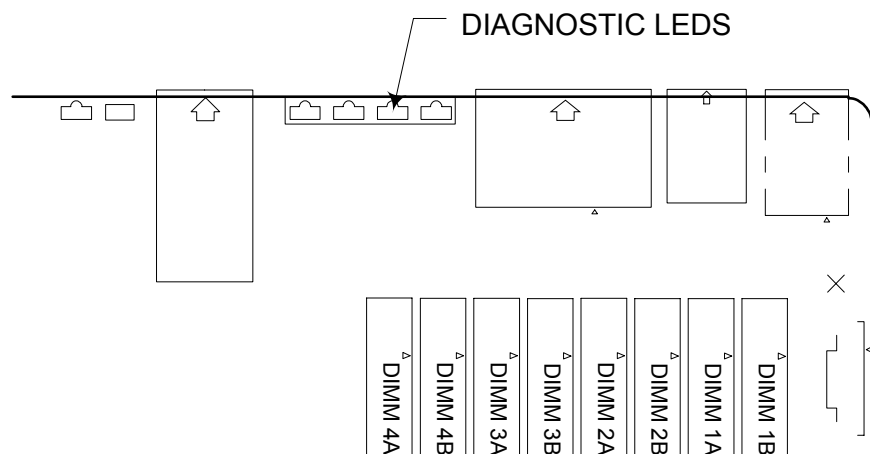


Figure 41. Location of Diagnostic LEDs on Baseboard

7.3.1.3 POST Code Checkpoints

The following table describes the type of checkpoints that may occur during the POST portion of the BIOS.

Table 90. POST Code Checkpoints

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
03	Off	Off	G	G	Disable NMI, Parity, video for EGA, and DMA controllers. Initialize BIOS, POST, Runtime data area. Also initialize BIOS modules on POST entry and GPNV area. Initialize CMOS as mentioned in the Kernel Variable "wCMOSFlags."
04	Off	G	Off	Off	Check CMOS diagnostic byte to determine if battery power is OK and CMOS checksum is OK. Verify CMOS checksum manually by reading storage area. If the CMOS checksum is bad, update CMOS with power-on default values and clear passwords. Initialize status register A. Initializes data variables that are based on CMOS setup questions. Initializes both the 8259 compatible PICs in the system
05	Off	G	Off	G	Initializes the interrupt controlling hardware (generally PIC) and interrupt vector table.
06	Off	G	G	Off	Do R/W test to CH-2 count reg. Initialize CH-0 as system timer. Install the POSTINT1Ch handler. Enable IRQ-0 in PIC for system timer interrupt. Traps INT1Ch vector to "POSTINT1ChHandlerBlock."
08	G	Off	Off	Off	Initializes the CPU. The BAT test is being done on KBC. Program the keyboard controller command byte is being done after Auto detection of KB/MS using AMI KB-5.
C0	R	R	Off	Off	Early CPU Init Start -- Disable Cache - Init Local APIC
C1	R	R	Off	G	Set up boot strap processor Information
C2	R	R	G	Off	Set up boot strap processor for POST
C5	R	A	Off	G	Enumerate and set up application processors

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
C6	R	A	G	Off	Re-enable cache for boot strap processor
C7	R	A	G	G	Early CPU Init Exit
0A	G	Off	G	Off	Initializes the 8042 compatible Key Board controller.
0B	G	Off	G	G	Detects the presence of PS/2 mouse.
0C	G	G	Off	Off	Detects the presence of Keyboard in KBC port.
0E	G	G	G	Off	Testing and initialization of different Input Devices. Also, update the Kernel Variables. Traps the INT09h vector, so that the POST INT09h handler gets control for IRQ1. Uncompress all available language, BIOS logo, and Silent logo modules.
13	Off	Off	G	A	Early POST initialization of chipset registers.
24	Off	G	R	Off	Uncompress and initialize any platform specific BIOS modules.
30	Off	Off	R	R	Initialize System Management Interrupt.
2A	G	Off	A	Off	Initializes different devices through DIM. See <i>DIM Code Checkpoints</i> section of document for more information.
2C	G	G	R	Off	Initializes different devices. Detects and initializes the video adapter installed in the system that have optional ROMs.
2E	G	G	A	Off	Initializes all the output devices.
31	Off	Off	R	A	Allocate memory for ADM module and uncompress it. Give control to ADM module for initialization. Initialize language and font modules for ADM. Activate ADM module.
33	Off	Off	A	A	Initializes the silent boot module. Set the window for displaying text information.
37	Off	G	A	A	Displaying sign-on message, CPU information, setup key message, and any OEM specific information.
38	G	Off	R	R	Initializes different devices through DIM. See <i>DIM Code Checkpoints</i> section of document for more information.
39	G	Off	R	A	Initializes DMAC-1 and DMAC-2.
3A	G	Off	A	R	Initialize RTC date/time.
3B	G	Off	A	A	Test for total memory installed in the system. Also, Check for DEL or ESC keys to limit memory test. Display total memory in the system.
3C	G	G	R	R	Mid POST initialization of chipset registers.
40	Off	R	Off	Off	Detect different devices (Parallel ports, serial ports, and coprocessor in CPU, ... etc.) successfully installed in the system and update the BDA, EBDA...etc.
50	Off	R	Off	R	Programming the memory hole or any kind of implementation that needs an adjustment in system RAM size if needed.
52	Off	R	G	R	Updates CMOS memory size from memory found in memory test. Allocates memory for Extended BIOS Data Area from base memory.
60	Off	R	R	Off	Initializes NUM-LOCK status and programs the KBD typematic rate.
75	Off	A	R	A	Initialize Int-13 and prepare for IPL detection.
78	G	R	R	R	Initializes IPL devices controlled by BIOS and option ROMs.
7A	G	R	A	R	Initializes remaining option ROMs.
7C	G	A	R	R	Generate and write contents of ESCD in NVRam.
84	R	G	Off	Off	Log errors encountered during POST.
85	R	G	Off	G	Display errors to the user and gets the user response for error.
87	R	G	G	G	Execute BIOS setup if needed / requested.

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
8C	A	G	Off	Off	Late POST initialization of chipset registers.
8D	A	G	Off	G	Build ACPI tables (if ACPI is supported)
8E	A	G	G	Off	Program the peripheral parameters. Enable/Disable NMI as selected
90	R	Off	Off	R	Late POST initialization of system management interrupt.
A0	R	Off	R	Off	Check boot password if installed.
A1	R	Off	R	G	Clean-up work needed before booting to operating system.
A2	R	Off	A	Off	Takes care of runtime image preparation for different BIOS modules. Fill the free area in F000h segment with 0FFh. Initializes the Microsoft* IRQ Routing Table. Prepares the runtime language module. Disables the system configuration display if needed.
A4	R	G	R	Off	Initialize runtime language module.
A7	R	G	A	G	Displays the system configuration screen if enabled. Initialize the CPU's before boot, which includes the programming of the MTRR's.
A8	A	Off	R	Off	Prepare CPU for operating system boot including final MTRR values.
A9	A	Off	R	G	Wait for user input at config display if needed.
AA	A	Off	A	Off	Uninstall POST INT1Ch vector and INT09h vector. De-initializes the ADM module.
AB	A	Off	A	G	Prepare BBS for Int 19 boot.
AC	A	G	R	Off	End of POST initialization of chipset registers.
B1	R	Off	R	A	Save system context for ACPI.
00	Off	Off	Off	Off	Passes control to OS Loader (typically INT19h).
61-70	-	-	-	-	OEM POST Error. This range is reserved for chipset vendors and system manufacturers. The error associated with this value may be different from one platform to the next.

7.3.1.4 Bootblock Initialization Code Checkpoints

The Bootblock initialization code sets up the chipset, memory and other components before system memory is available. The following table describes the type of checkpoints that may occur during the bootblock initialization.

Table 91. Bootblock Initialization Code Checkpoints

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
Before D1h	R	R	Off	A	Early chipset initialization is done. Early super I/O initialization is done including RTC and keyboard controller. NMI is disabled.
D1	R	R	Off	A	Perform keyboard controller BAT test. Check if waking up from power management suspend state. Save power-on CPUID value in scratch CMOS.
D0	R	R	Off	R	Go to flat mode with 4GB limit and GA20 enabled. Verify the bootblock checksum.

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
D2	R	R	G	R	Disable CACHE before memory detection. Execute full memory sizing module. Verify that flat mode is enabled.
D3	R	R	G	A	If memory sizing module not executed, start memory refresh and do memory sizing in Bootblock code. Do additional chipset initialization. Re-enable CACHE. Verify that flat mode is enabled.
D4	R	A	Off	R	Test base 512KB memory. Adjust policies and cache first 8MB. Set stack.
D5	R	A	Off	A	Bootblock code is copied from ROM to lower system memory and control is given to it. BIOS now executes out of RAM.
D6	R	A	G	R	Both key sequence and OEM specific method is checked to determine if BIOS recovery is forced. Main BIOS checksum is tested. If BIOS recovery is necessary, control flows to checkpoint E0. See <i>Bootblock Recovery Code Checkpoints</i> section of document for more information.
D7	R	A	G	A	Restore CPUID value back into register. The Bootblock-Runtime interface module is moved to system memory and control is given to it. Determine whether to execute serial flash.
D8	A	R	Off	R	The Runtime module is uncompressed into memory. CPUID information is stored in memory.
D9	A	R	Off	A	Store the Uncompressed pointer for future use in PMM. Copying Main BIOS into memory. Leaves all RAM below 1MB Read-Write including E000 and F000 shadow areas but closing SMRAM.
DA	A	R	G	R	Restore CPUID value back into register. Give control to BIOS POST (ExecutePOSTKernel). See <i>POST Code Checkpoints</i> section of document for more information.
E1-E8 EC-EE	-	-	-	-	OEM memory detection/configuration error. This range is reserved for chipset vendors and system manufacturers. The error associated with this value may be different from one platform to the next.

7.3.1.5 Bootblock Recovery Code Checkpoints

The Bootblock recovery code gets control when the BIOS determines that a BIOS recovery needs to occur because the user has forced the update or the BIOS checksum is corrupt. The following table describes the type of checkpoints that may occur during the Bootblock recovery portion of the BIOS:

Table 92. Bootblock Recovery Code Checkpoints

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
E0	R	R	R	Off	Initialize the floppy controller in the super I/O. Some interrupt vectors are initialized. DMA controller is initialized. 8259 interrupt controller is initialized. L1 cache is enabled.
E9	A	R	R	G	Set up floppy controller and data. Attempt to read from floppy.
EA	A	R	A	Off	Enable ATAPI hardware. Attempt to read from ARMD and ATAPI CDROM.

Check point	Diagnostic LED Decoder G=Green, R=Red, A=Amber				Description
	Hi			Low	
EB	A	R	A	G	Disable ATAPI hardware. Jump back to checkpoint E9.
EF	A	A	A	G	Read error occurred on media. Jump back to checkpoint EB.
F0	R	R	R	R	Search for pre-defined recovery file name in root directory.
F1	R	R	R	A	Recovery file not found.
F2	R	R	A	R	Start reading FAT table and analyze FAT to find the clusters occupied by the recovery file.
F3	R	R	A	A	Start reading the recovery file cluster by cluster.
F5	R	A	R	A	Disable L1 cache.
FA	A	R	A	R	Check the validity of the recovery file configuration to the current configuration of the flash part.
FB	A	R	A	A	Make flash write enabled through chipset and OEM specific method. Detect proper flash part. Verify that the found flash part size equals the recovery file size.
F4	R	A	R	R	The recovery file size does not equal the found flash part size.
FC	A	A	R	R	Erase the flash part.
FD	A	A	R	A	Program the flash part.
FF	A	A	A	A	The flash has been updated successfully. Make flash write disabled. Disable ATAPI hardware. Restore CPUID value back into register. Give control to F000 ROM at F000:FFF0h.

7.3.1.6 DIM Code Checkpoints

The Device Initialization Manager (DIM) gets control at various times during BIOS POST to initialize different system buses. The following table describes the main checkpoints where the DIM module is accessed.

Table 93. DIM Code Checkpoints

Checkpoint	Description
2A	Initialize different buses and perform the following functions: Reset, Detect, and Disable (function 0); Static Device Initialization (function 1); Boot Output Device Initialization (function 2). Function 0 disables all device nodes, PCI devices, and PnP ISA cards. It also assigns PCI bus numbers. Function 1 initializes all static devices that include manual configured onboard peripherals, memory and I/O decode windows in PCI-PCI bridges, and noncompliant PCI devices. Static resources are also reserved. Function 2 searches for and initializes any PnP, PCI, or AGP video devices.
38	Initialize different buses and perform the following functions: Boot Input Device Initialization (function 3); IPL Device Initialization (function 4); General Device Initialization (function 5). Function 3 searches for and configures PCI input devices and detects if system has standard keyboard controller. Function 4 searches for and configures all PnP and PCI boot devices. Function 5 configures all onboard peripherals that are set to an automatic configuration and configures all remaining PnP and PCI devices.

While control is in the different functions, additional checkpoints are output to port 80h as a word value to identify the routines under execution. The low byte value indicates the main POST Code Checkpoint. The high byte is divided into two nibbles and contains two fields. The details of the high byte of these checkpoints are as follows:

HIGH BYTE XY

The upper nibble 'X' indicates the function number that is being executed. 'X' can be from 0 to 7.

- 0 = func#0, disable all devices on the BUS concerned.
- 1 = func#1, static devices initialization on the BUS concerned.
- 2 = func#2, output device initialization on the BUS concerned.
- 3 = func#3, input device initialization on the BUS concerned.
- 4 = func#4, IPL device initialization on the BUS concerned.
- 5 = func#5, general device initialization on the BUS concerned.
- 6 = func#6, error reporting for the BUS concerned.
- 7 = func#7, add-on ROM initialization for all BUSES.
- 8 = func#8, BBS ROM initialization for all BUSES.

The lower nibble 'Y' indicates the BUS on which the different routines are being executed. 'Y' can be from 0 to 5.

- 0 = Generic DIM (Device Initialization Manager).
- 1 = On-board System devices.
- 2 = ISA devices.
- 3 = EISA devices.
- 4 = ISA PnP devices.
- 5 = PCI devices.

7.3.1.7 ACPI Runtime Checkpoints

ACPI checkpoints are displayed when an ACPI capable operating system either enters or leaves a sleep state. The following table describes the type of checkpoints that may occur during ACPI sleep or wake events.

Table 94. ACPI Runtime Checkpoints

Checkpoint	Description
------------	-------------

AC	First ASL check point. Indicates the system is running in ACPI mode.
AA	System is running in APIC mode.
01, 02, 03, 04, 05	Entering sleep state S1, S2, S3, S4, or S5.
10, 20, 30, 40, 50	Waking from sleep state S1, S2, S3, S4, or S5.

7.3.2 BIOS Messages

Table 95. Memory BIOS Messages

Message Displayed	Description
Gate20 Error	The BIOS is unable to properly control the motherboard's Gate A20 function, which controls access of memory over 1 MB. This may indicate a problem with the motherboard.
Multi-Bit ECC Error	This message will only occur on systems using ECC enabled memory modules. ECC memory has the ability to correct single-bit errors that may occur from faulty memory modules. A multiple bit corruption of memory has occurred, and the ECC memory algorithm cannot correct it. This may indicate a defective memory module.
Parity Error	Fatal Memory Parity Error. System halts after displaying this message.

Table 96. Boot BIOS Messages

Message Displayed	Description
Boot Failure ...	This is a generic message indicating the BIOS could not boot from a particular device. This message is usually followed by other information concerning the device.
Invalid Boot Diskette	A diskette was found in the drive, but it is not configured as a bootable diskette.
Drive Not Ready	The BIOS was unable to access the drive because it indicated it was not ready for data transfer. This is often reported by drives when no media is present.
A: Drive Error	The BIOS attempted to configure the A: drive during POST, but was unable to properly configure the device. This may be due to a bad cable or faulty diskette drive.

B: Drive Error	The BIOS attempted to configure the B: drive during POST, but was unable to properly configure the device. This may be due to a bad cable or faulty diskette drive.
Insert BOOT diskette in A:	The BIOS attempted to boot from the A: drive, but could not find a proper boot diskette.
Reboot and Select proper Boot device or Insert Boot Media in selected Boot device	BIOS could not find a bootable device in the system and/or removable media drive does not contain media.
NO ROM BASIC	This message occurs on some systems when no bootable device can be detected.

Table 97. Storage Device BIOS Messages

Message Displayed	Description
Primary Master Hard Disk Error	The IDE/ATAPI device configured as Primary Master could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Primary Slave Hard Disk Error	The IDE/ATAPI device configured as Primary Slave could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Secondary Master Hard Disk Error	The IDE/ATAPI device configured as Secondary Master could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Secondary Slave Hard Disk Error	The IDE/ATAPI device configured as Secondary Slave could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
3 rd Master Hard Disk Error	The IDE/ATAPI device configured as Master in the 3 rd IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
3 rd Slave Hard Disk Error	The IDE/ATAPI device configured as Slave in the 3 rd IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
4 th Master Hard Disk Error	The IDE/ATAPI device configured as Master in the 4th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
4 th Slave Hard Disk Error	The IDE/ATAPI device configured as Slave in the 4th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.

5 th Master Hard Disk Error	The IDE/ATAPI device configured as Master in the 5th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
5 th Slave Hard Disk Error	The IDE/ATAPI device configured as Slave in the 5th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
6 th Master Hard Disk Error	The IDE/ATAPI device configured as Master in the 6th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
6 th Slave Hard Disk Error	The IDE/ATAPI device configured as Slave in the 6th IDE controller could not be properly initialized by the BIOS. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Primary Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Primary Master failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Primary Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Primary Slave failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Secondary Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Secondary Master failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
Secondary Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Secondary Slave failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
3 rd Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Master in the 3 rd IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
3 rd Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Slave in the 3 rd IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
4 th Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Master in the 4 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
4 th Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Slave in the 4 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.

5 th Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Master in the 5 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
5 th Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Slave in the 5 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
6 th Master Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Master in the 6 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
6 th Slave Drive - ATAPI Incompatible	The IDE/ATAPI device configured as Slave in the 6 th IDE controller failed an ATAPI compatibility test. This message is typically displayed when the BIOS is trying to detect and configure IDE/ATAPI devices in POST.
S.M.A.R.T. Capable but Command Failed	The BIOS tried to send a S.M.A.R.T. message to a hard disk, but the command transaction failed. This message can be reported by an ATAPI device using the S.M.A.R.T. error reporting standard. S.M.A.R.T. failure messages may indicate the need to replace the hard disk.
S.M.A.R.T. Command Failed	The BIOS tried to send a S.M.A.R.T. message to a hard disk, but the command transaction failed. This message can be reported by an ATAPI device using the S.M.A.R.T. error reporting standard. S.M.A.R.T. failure messages may indicate the need to replace the hard disk.
S.M.A.R.T. Status BAD, Backup and Replace	A S.M.A.R.T. capable hard disk sends this message when it detects an imminent failure. This message can be reported by an ATAPI device using the S.M.A.R.T. error reporting standard. S.M.A.R.T. failure messages may indicate the need to replace the hard disk.
S.M.A.R.T. Capable and Status BAD	A S.M.A.R.T. capable hard disk sends this message when it detects an imminent failure. This message can be reported by an ATAPI device using the S.M.A.R.T. error reporting standard. S.M.A.R.T. failure messages may indicate the need to replace the hard disk.

Table 98. Virus Related BIOS Messages

Message Displayed	Description
-------------------	-------------

BootSector Write !!	The BIOS has detected software attempting to write to a drive's boot sector. This is flagged as possible virus activity. This message will only be displayed if Virus Detection is enabled in AMIBIOS setup.
VIRUS: Continue (Y/N)?	If the BIOS detects possible virus activity, it will prompt the user. This message will only be displayed if Virus Detection is enabled in AMIBIOS setup.

Table 99. System Configuration BIOS Messages

Message Displayed	Description
DMA-2 Error	Error initializing secondary DMA controller. This is a fatal error, often indication a problem with system hardware.
DMA Controller Error	POST error while trying to initialize the DMA controller. This is a fatal error, often indication a problem with system hardware.
Checking NVRAM..Update Failed	BIOS could not write to the NVRAM block. This message appears when the FLASH part is write-protected or if there is no FLASH part (System uses a PROM or EPROM).
Microcode Error	BIOS could not find or load the CPU Microcode Update to the CPU. This message only applies to INTEL CPUs. The message is most likely to appear when a brand new CPU is installed in a motherboard with an outdated BIOS. In this case, the BIOS must be updated to include the Microcode Update for the new CPU.
NVRAM Checksum Bad, NVRAM Cleared	An error was identified while validating the NVRAM data. This causes POST to clear the NVRAM data.
Resource Conflict	More than one system device is trying to use the same non-shareable resources (memory or I/O).
NVRAM Ignored	The NVRAM data used to store Plug'n'Play (PnP) data was not used for system configuration in POST.
NVRAM Bad	The NVRAM data used to store Plug'n'Play (PnP) data was not used for system configuration in POST due to a data error.
Static Resource Conflict	Two or more Static Devices are trying to use the same resource space (usually memory or I/O).
PCI I/O conflict	A PCI adapter generated an I/O resource conflict when configured by BIOS POST.
PCI ROM conflict	A PCI adapter generated an I/O resource conflict when configured by BIOS POST.
PCI IRQ conflict	A PCI adapter generated an I/O resource conflict when configured by BIOS POST.
PCI IRQ routing table error	BIOS POST (DIM code) found a PCI device in the system but was unable to figure out how to route an IRQ to the device. Usually this error is causing by an incomplete description of the PCI Interrupt Routing of the system.

Timer Error	Indicates an error while programming the count register of channel 2 of the 8254 timer. This may indicate a problem with system hardware.
Interrupt Controller-1 error	BIOS POST could not initialize the Master Interrupt Controller. This may indicate a problem with system hardware.
Interrupt Controller-2 error	BIOS POST could not initialize the Slave Interrupt Controller. This may indicate a problem with system hardware.

Table 100. CMOS BIOS Messages

Message Displayed	Description
CMOS Date/Time Not Set	The CMOS Date and/or Time are invalid. This error can be resolved by readjusting the system time in AMIBIOS Setup.
CMOS Battery Low	CMOS Battery is low. This message usually indicates that the CMOS battery needs to be replaced. It could also appear when the user intentionally discharges the CMOS battery.
CMOS Settings Wrong	CMOS settings are invalid. This error can be resolved by using AMIBIOS Setup.
CMOS Checksum Bad	CMOS contents failed the Checksum check. Indicates that the CMOS data has been changed by a program other than the BIOS or that the CMOS is not retaining its data due to malfunction. This error can typically be resolved by using AMIBIOS Setup.

Table 101. Miscellaneous BIOS Messages

Message Displayed	Description
Keyboard Error	Keyboard is not present or the hardware is not responding when the keyboard controller is initialized.
PS2 Keyboard not found	PS2 Keyboard support is enabled in the BIOS setup but the device is not detected.
PS2 Mouse not found	PS2 Mouse support is enabled in the BIOS setup but the device is not detected.
Keyboard/Interface Error	Keyboard controller failure. This may indicate a problem with system hardware.
Unlock Keyboard	PS2 keyboard is locked. User needs to unlock the keyboard to continue the BIOS POST.
System Halted	The system has been halted. A reset or power cycle is required to reboot the machine. This message appears after a fatal error has been detected.

Table 102. USB BIOS Error Messages

Message Displayed	Description
Warning! Unsupported USB device found and disabled!	This message is displayed when a non-bootable USB device is enumerated and disabled by the BIOS.
Warning! Port 60h/64h emulation is not supported by this USB Host Controller!	This message is displayed to indicate that port 60h/64h emulation mode cannot be enabled for this USB host controller. This condition occurs if USB KBC emulation option is set for non-SMI mode.
Warning! EHCI controller disabled. It requires 64bit data support in the BIOS.	This message is displayed to indicate that EHCI controller is disabled because of incorrect data structure. This condition occur if the USB host controller needs 64-bit data structure while the USB is ported with 32-bit data structure.

Table 103. SMBIOS BIOS Error Messages

Message Displayed	Description
Not enough space in Runtime area!!. SMBIOS data will not be available.	This message is displayed when the size of the SMBIOS data exceeds the available SMBIOS runtime storage size.

7.3.3 POST Error Messages and Handling

Whenever possible the BIOS will output the current boot progress codes on the video screen. Progress codes are 32 bit quantities plus optional data. The 32 bit numbers include Class, subclass and Operation information. Class and subclass point to the type of the hardware that is being initialized, where as the Operation field represents the specific initialization activity. Based upon the data bit availability to display Progress Code, a progress code can be customized to fit the data width. The higher the data bit, higher the granularity of information, which could send on the progress port. The progress codes may be reported by system BIOS or option ROMs.

The Response section in following table is divided in 3 different types

- **Warning** – The message is displayed on screen and error is logged in SEL. System will continue booting with degraded state. User may want to replace erroneous unit.
- **Pause** – The message is displayed on screen, an error is logged in SEL, and user input is required to continue. User can take immediate corrective action or can choose to continue booting.
- **Halt** – The message is displayed on screen, an error is logged in SEL, and the system cannot boot unless the error is resolved. User needs to replace faulty part and restart the system.

Table 104. POST Error Messages and Handling

Error Code	Error Message	Response
0000	Timer Error	Warning
0004	CMOS Settings Wrong	Warning
0005	CMOS Checksum Bad	Warning
0012	CMOS Date/Time Not Set	Warning

Error Code	Error Message	Response
0040	Refresh timer test failed	Halt
0044	DMA Controller Error	Halt
0045	DMA-1 Error	Halt
0046	DMA-2 Error	Halt
0048	Password check failed	Halt
004C	Keyboard/Interface Error	Warning
004D	Primary Master Hard Disk Error	Warning
004E	Primary Slave Hard Disk Error	Warning
0055	Primary Master Drive - ATAPI Incompatible	Warning
0056	Primary Slave Drive - ATAPI Incompatible	Warning
005D	S.M.A.R.T. Status BAD, Backup and Replace	Warning
0120	Thermal Trip Failure	Warning
0146	Insufficient Memory to Shadow PCI ROM	Pause
0150	BSP Processor failed BIST	Warning
0160	Processor missing microcode – P0	Warning
0161	Processor missing microcode – P1	Warning
0180	BIOS does not support current stepping – P0	Pause
0181	BIOS does not support current stepping – P1	Pause
0192	L2 cache size mismatch	Pause
0193	CPUID, Processor stepping are different	Pause
0194	CPUID, Processor family are different	Pause
0195	Front side bus mismatch. System halted.	Pause
0196	CPUID, Processor Model are different	Pause
0197	Processor speeds mismatched	Pause
5120	CMOS Cleared By Jumper	Warning
5121	Password cleared by jumper	Warning
5122	CMOS Cleared By BMC Request	Warning
8103	Warning! Unsupported USB device found and disabled !!!	Warning
8104	Warning! Port 60h/64h emulation is not supported by this USB Host Controller !!!	Warning
8105	Warning! EHCI controller disabled. It requires 64bit data support in the BIOS.	Warning
8110	Processor 1 Internal error (IERR)	Warning
8111	Processor 2 Internal error (IERR)	Warning
8120	Processor 1 Thermal Trip error	Warning
8121	Processor 2 Thermal Trip error	Warning
8130	Processor 1 disabled	Warning
8131	Processor 2 disabled	Warning
8140	Processor 1 failed FRB-3 timer	Warning
8141	Processor 2 failed FRB-3 timer	Warning
8150	Processor 1 failed initialization on last boot.	Warning
8151	Processor 2 failed initialization on last boot.	Warning
8160	Processor 1: unable to apply BIOS update	Pause
8161	Processor 2: unable to apply BIOS update	Pause
8170	Processor 1 failed BIST	Warning
8171	Processor 2 failed BIST	Warning

Error Code	Error Message	Response
8180	BIOS does not support current stepping for Processor 1	Pause
8181	BIOS does not support current stepping for Processor 2	Pause
8190	Watchdog timer failed on last boot	Warning
8198	OS boot watchdog timer failure	Warning
8300	BaseBoard Management Controller failed Self Test	Pause
8301	Front Panel Controller failed to function	Pause
8305	Primary Hot swap Controller failed to function	Warning
8306	Power Share Controller failed to function	Warning
84F1	BIST failed for all available processors	Halt
84F2	BaseBoard Management Controller failed to respond	Pause
84F3	BaseBoard Management Controller in Update Mode	Pause
84F4	Sensor Data Record Empty	Pause
84FF	System Event Log Full	Warning
8500	Bad or missing memory in slot 4A	Pause
8501	Bad or missing memory in slot 3A	Pause
8502	Bad or missing memory in slot 2A	Pause
8503	Bad or missing memory in slot 1A	Pause
8504	Bad or missing memory in slot 4B	Pause
8505	Bad or missing memory in slot 3B	Pause
8506	Bad or missing memory in slot 2B	Pause
8507	Bad or missing memory in slot 1B	Pause
8601	All memory marked as fail. Forcing minimum back online	Pause

The messages in the following table do not appear on the video and nor do they get logged in the SEL. These are error codes that are sent to Management Module for Error Logging as BMC pass through command. The syntax of error logging with management module and SEL is different. Same error may be logged in IPMI format in the SEL.

Table 105. Management Module POST Error Codes and Messages

Error code	Error messages
161	Bad CMOS Battery
301	Keyboard failure
102	System Board failure (Timer tick 2 test failure)
106	Diskette Controller Failure
604	Diskette Drive ? failure
163	Time of the day not set
01298000	The BIOS does not support the current stepping of Processor P0
01298001	The BIOS does not support the current stepping of Processor P1
196	Processor cache mismatch detected.
198	Processor speed mismatch detected.
00019700	Processor P0 failed BIST.
00019701	Processor P1 failed BIST.
00150100	Multi-bit error occurred: forcing NMI DIMM = ??

00150100	Multi-bit error occurred: forcing NMI DIMM = ?? DIMM = ?? (could not isolate)
289	DIMM D?? is Disabled.
00150900	SERR/PERR Detected on PCI bus (no source found)
00151100	MCA: Recoverable Error Detected Proc = ??
00151200	MCA: Unrecoverable Error Detected Proc = ??
00151300	MCA: Excessive Recoverable Errors Proc = ??
00151350	Processor MachineCheck Data a Bank = ?? APIC ID = ?? CR4 = ???? ???? ?
00151351	Processor MachineCheck Data b Address = ???? ???? ???? ???? Time Stamp = ???? ???? ?
00151352	Processor MachineCheck Data b Status = ???? ???? ???? ???? ?
00151500	Excessive Single Bit Errors Detected
00151720	Parity Error Detected on processor bus
00151730	IMB Parity/CRC Error
00151700	Started Hot Spare memory Copy. Failed row/rows = ?? and ?? copied to spare row/rows = ?? and ?? (used on CMIC-HE box)
00151710	Completed Hot Spare memory Copy. Failed row/rows = ?? and ?? copied to spare row/rows = ?? and ?? (used on CMIC-HE box)

7.3.4 Boot Block Error Beep Codes

Table 106. Bootblock Error Beep Codes

Number of Beeps	Description
1	Insert diskette in floppy drive A:
2	'AMIBOOT.ROM' file not found in root directory of diskette in A:
3	Base memory error
4	Flash Programming successful
5	Floppy read error
6	Keyboard controller BAT command failed
7	No Flash EPROM detected
8	Floppy controller failure
9	Boot Block BIOS checksum error
10	Flash Erase error
11	Flash Program error
12	'AMIBOOT.ROM' file size error
13	BIOS ROM image mismatch (file layout does not match image present in flash device)
1 long beep	Insert diskette with AMIBOOT.001 File for Multi-Disk Recovery

7.3.5 POST Error Beep Codes

The following table lists POST error beep codes. Prior to system video initialization, BIOS uses these beep codes to inform users on error conditions.

Table 107. POST Error Beep Codes

Number of Beeps	Description
1	Memory refresh timer error
2	Parity error in base memory (first 64KB block)
3	Base memory read / write test error
4	Motherboard timer not operational
5	Processor error
6	8042 Gate A20 test error (cannot switch to protected mode)
7	General exception error (processor exception error)
8	Display memory error (system video adapter)
9	ROM checksum error
10	CMOS shutdown register read/write error
11	Cache memory test failed

7.3.5.1 Troubleshooting BIOS Beep Codes

Table 108. Troubleshooting BIOS Beep Codes

Number of Beeps	Troubleshooting Action
1, 2 or 3	Reseat the memory, or replace with known good modules.
4-7, 9-11	<p>Fatal error indicating a serious problem with the system. Consult your system manufacturer. Before declaring the motherboard beyond all hope, eliminate the possibility of interference by a malfunctioning add-in card. Remove all expansion cards except the video adapter.</p> <ul style="list-style-type: none"> · If beep codes are generated even when all other expansion cards are absent, consult your system manufacturer's technical support. · If beep codes are not generated when all other expansion cards are absent, one of the add-in cards is causing the malfunction. Insert the cards back into the system one at a time until the problem happens again. This will reveal the malfunctioning add-in card.
8	If the system video adapter is an add-in card, replace or reseat the video adapter. If the video adapter is an integrated part of the system board, the board may be faulty.

7.3.6 "POST Error Pause" option

In case of POST error(s) which occur during system boot-up, BIOS will stop and wait for user to press an appropriate key before booting the O/S or entering BIOS setup.

The user can override this option by setting "POST Error Pause" to "disabled" in BIOS setup Advanced menu page. If "POST Error Pause" option is selected "disabled", system will boot the O/S without user-intervention. Option default value is set to "enabled".

8. Connectors, Headers and Jumpers

8.1 Board Connector Information

The following section provides detailed information regarding all connectors, headers and jumpers on the Server Board SE7520AF2. Table 109 lists all connector types available on the board and corresponding reference designators printed on silkscreen.

Table 109. Board Connector Matrix

Connector	Quantity	Reference Designators	Connector Type	Pin Count
Power supply	3	2X4 = J9E1 2X12 = J9C1 1X5 = J9B1	CPU Power Main Power P/S Aux	8 24 5
CPU	2	J8H1, J6H1	CPU Sockets	604
Main Memory	8	J9D1, J9D2, J8D2, J8D3, J7D3, J8D1, J7D1, J7D2	DIMM Sockets	240
PCI Express	2	J2C1, J3C1	Card Edge	98
PCI-X 266MHz	2	J1D4, J4D1	Card Edge	184
PCI-X 100MHz	1	J4D2	Card Edge	184
RAID Memory	1	J1D5	DIMM Socket	187
RAID Key	1	U1F2	Key Holder	3
IDE	1	J3K1	Shrouded Header	40
SCSI LED	1	J1C2	Header	4
System Fans 3+2 pin	2	J2J3+J2K3, J2J2+J2K1	Header	3+2
System Fans 6 pin	2	J2J4, J2K2	Header	6
System Fans 3 pin	2	J5A2, J5A3	Header	3
CPU Fans	2	J6F1, J5F1	Header	3
Battery	1	BATT1J1	Battery Holder	3
Keyboard/Mouse	1	J9A1	PS2, stacked	12
Rear USB	1	J9A2	External, Stacked	12
Serial Port A	1	J8A1	External, D-Sub	9
Serial Port B	1	J1B1	Header	10
Video connector	1	J8A1	External, D-Sub	15
Dual LAN connector 10/100/1000	1	J6A1	Dual LAN connector with built-in magnetic	38
Floppy drive	1	J3K2	Header	34
Front panel, Main	1	2x17 = J1E1	Header	34
Front panel, USB	1	J1K3	Header	10
Intrusion detect	2	J1C1	Header	2
SCSI connector	2	J1H1, J1F1	Header	80
Serial ATA	2	J2J1, J2J4	Header	7
ICMB	1	J1A1	Header	5
IPMB	1	J1G1	Header	3
OEM RMC	1	J1B2	Header	8

HSBP	2	J1K1, J1K2	Header	4
Recovery Jumpers	1	J1D1	Jumper	9

8.2 Main Power Connector

The main power supply connection is obtained using the 24-pin connector. The following table defines the pin-outs of the connector.

Table 110. Power Connector Pin-out (J9C1)

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

Table 111. Power Supply Signal Connector (J9B1)

Pin	Signal	Color
1	5VSB_SCL	Orange
2	5VSB_SDA	Black
3	PS_ALERT_L	Red
4	GND	Yellow
5	3.3V SENSE(+)	Green

Table 112. 12V Power Connector (J9E1)

Pin	Signal	Color
1	GND	Black
2	GND	Black
3	GND	Black
4	GND	Black
5	+12Vdc	Yellow
6	+12Vdc	Yellow
7	+12Vdc	Yellow

8	+12Vdc	Yellow
---	--------	--------

8.3 Main Memory Module Connector

The Intel® Server Board SE7520AF2 has eight DIMM connectors which support registered ECC DDR-2 400MHz modules. For additional DIMM information, refer to the *DDR2-400Mhz Registered DIMM Specification*.

Table 113. DIMM Connectors (J9D1, J9D2, J8D2, J8D3, J7D3, J8D1, J7D1, J7D2)

Pin	Front	Pin	Back	Pin	Front	Pin	Back
1	VREF	121	VSS	61	A4	181	VDDQ
2	VSS	122	DQ4	62	VDDQ	182	A3
3	DQ0	123	DQ5	63	A2	183	A1
4	DQ1	124	VSS	64	VDD	184	VDD
5	VSS	125	DM0,DQS9	Key		Key	
6	DQS0_L	126	NC,DQS9_L	65	VSS	185	CK0
7	DQS0	127	VSS	66	VSS	186	CK0_L
8	VSS	128	DQ6	67	VDD	187	VDD
9	DQ2	129	DQ7	68	NC	188	A0
10	DQ3	130	VSS	69	VDD	189	VDD
11	VSS	131	DQ12	70	A10/AP	190	BA1
12	DQ8	132	DQ13	71	BA0	191	VDDQ
13	DQ9	133	VSS	72	VDDQ	192	RAS_L
14	VSS	134	DM1,DQS10	73	WE_L	193	S0_L
15	DQS1_L	135	NC,DQS10_L	74	CAS_L	194	VDDQ
16	DQS1	136	VSS	75	VDDQ	195	ODT0
17	VSS	137	CK1	76	S1_L	196	A13
18	RESET_L	138	CK1_L	77	ODT1	197	VDD
19	NC	139	VSS	78	VDDQ	198	VSS
20	VSS	140	DQ14	79	VSS	199	DQ36
21	DQ10	141	DQ15	80	DQ32	200	DQ37
22	DQ11	142	VSS	81	DQ33	201	VSS
23	VSS	143	DQ20	82	VSS	202	DM4,DQS13
24	DQ16	144	DQ21	83	DQS4_L	203	NC,DQS13_L
25	DQ17	145	VSS	84	DQS4	204	VSS
26	VSS	146	DM2,DQS11	85	VSS	205	DQ38
27	DQS2_L	147	NC,DQS11_L	86	DQ34	206	DQ39
28	DQS2	148	VSS	87	DQ35	207	VSS
29	VSS	149	DQ22	88	VSS	208	DQ44
30	DQ18	150	DQ23	89	DQ40	209	DQ45
31	DQ19	151	VSS	90	DQ41	210	VSS
32	VSS	152	DQ28	91	VSS	211	DM5,DQS14
33	DQ24	153	DQ29	92	DQS5_L	212	NC,DQS14_L
34	DQ25	154	VSS	93	DQS5	213	VSS
35	VSS	155	DM3,DQS12	94	VSS	214	DQ46

Pin	Front	Pin	Back	Pin	Front	Pin	Back
36	DQS3_L	156	NC,DQS12_L	95	DQ42	215	DQ47
37	DQS3	157	VSS	96	DQ43	216	VSS
38	VSS	158	DQ30	97	VSS	217	
39	DQ26	159	DQ31	98	DQ48	218	DQ53
40	DQ27	160	VSS	99	DQ49	219	VSS
41	VSS	161	CB4	100	VSS	220	CK2
42	CB0	162	CB5	101	SA2	221	CK2_L
43	CB1	163	VSS	102	NC	222	VSS
44	VSS	164	DM8,DQS17	103	VSS	223	DM6,DQS15
45	DQS8_L	165	NC,DQS17_L	104	DQS6_L	224	NC,DQS15_L
46	DQS8	166	VSS	105	DQS6	225	VSS
47	VSS	167	CB6	106	VSS	226	DQ54
48	CB2	168	CB7	107	DQS5	227	DQ55
49	CB3	169	VSS	108	DQ51	228	VSS
50	VSS	170	VDDQ	109	VSS	229	DQ60
51	VDDQ	171	CKE1	110	DQ56	230	DQ61
52	CKD0	172	VDD	111	DQ57	231	VSS
53	VDD	173	A15	112	VSS	232	DM7,DQS16
54	A16,BA2	174	A14	113	DQS7_L	233	NC,DQS16_L
55	RC02	175	VDDQ	114	DQS7	234	VSS
56	VDDQ	176	A12	115	VSS	235	DQ62
57	A11	177	A9	116	DQ58	236	DQ63
58	A7	178	VDD	117	DQ59	237	VSS
59	VDD	179	A8	118	VSS	238	VDDSPD
60	A5	180	A6	119	SDA	239	SA0
				120	SCL	240	SA1

8.4 RAID Memory Module Connector

The Intel® Server Board SE7520AF2 has one RAID DIMM connector (Slot 2) which supports one unbuffered ECC DDR 333 MHz module. For additional DIMM information, refer to the *DDR 333Mhz Registered DIMM Specification*.

Table 114. RAID DIMM Connector (J1D5)

Pin	Front	Pin	Front	Pin	Front	Pin	Back	Pin	Back	Pin	Back
1	VREF	34	GND	67	DQS5	100	GND	133	DQ31	166	DQ53
2	DQ0	35	DQ25	68	DQ42	101	NC	134	CB4	167	FETEN
3	GND	36	DQS3	69	DQ43	102	NC	135	CB5	168	VDD
4	DQ1	37	A4	70	VDD	103	A13	136	VDDQ	169	DM6
5	DQS0	38	VDD	71	RSVD	104	VDDQ	137	CK0P	170	DQ54
6	DQ2	39	DQ26	72	DQ48	105	DQ12	138	CK0N	171	DQ55
7	VDD	40	DQ27	73	DQ49	106	DQ13	139	GND	172	VDDQ
8	DQ3	41	A2	74	GND	107	DM1	140	DM8	173	NC

Pin	Front	Pin	Front	Pin	Front	Pin	Back	Pin	Back	Pin	Back
9	NC	42	GND	75	RSVD	108	VDD	141	A10	174	DQ60
10	RESET*	43	A1	76	RSVD	109	DQ14	142	CB6	175	DQ61
11	GND	44	CB0	77	VDDQ	110	DQ15	143	VDDQ	176	GND
12	DQ8	45	CB1	78	DQS6	111	CKE1	144	CB7	177	DM7
13	DQ9	46	VDD	79	DQ50	112	VDDQ	145	GND	178	DQ62
14	DQS1	47	DQS8	80	DQ51	113	BA2	146	DQ36	179	DQ63
15	VDDQ	48	A0	81	GND	114	DQ20	147	DQ37	180	VDDQ
16	RSVD	49	CB2	82	VDDID	115	A12	148	VDD	181	SA0
17	RSVD	50	GND	83	DQ56	116	GND	149	DM4	182	SA1
18	GND	51	CB3	84	DQ57	117	DQ21	150	DQ38	183	SA2
19	DQ10	52	BA1	85	VDD	118	A11	151	DQ39	184	VDDSPD
20	DQ11	53	DQ32	86	DQS7	119	DM2	152	GND	185	NC
21	CKE0	54	VDDQ	87	DQ58	120	VDD	153	DQ44	186	NC
22	VDDQ	55	DQ33	88	DQ59	121	DQ22	154	RAS*	187	NC
23	DQ16	56	DQS4	89	GND	122	A8	155	DQ45		
24	DQ17	57	DQ34	90	NC	123	DQ23	156	VDDQ		
25	DQS2	58	GND	91	SDA	124	GND	157	CS0*		
26	GND	59	BA0	92	SCL	125	A6	158	CS1*		
27	A9	60	DQ35	93	GND	126	DQ28	159	DM5		
28	DQ18	61	DQ40	94	DQ4	127	DQ29	160	GND		
29	A7	62	VDDQ	95	DQ5	128	VDDQ	161	DQ46		
30	VDDQ	63	WE*	96	VDDQ	129	DM3	162	DQ47		
31	DQ19	64	DQ41	97	DM0	130	A3	163	RSVD		
32	A5	65	CAS*	98	DQ6	131	DQ30	164	VDDQ		
33	DQ24	66	GND	99	DQ7	132	GND	165	DQ52		

8.5 Processor Socket

The SE7520AF2 has two Socket 604 processor sockets. The following table provides the processor socket pin numbers and pin names:

Table 115. Socket 604 Processor Socket Pinout (J8H1, J6H1)

Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name
A1	Reserved	B6	VCC	C11	A30#	D16	A17#	E21	RS0#
A2	VCC	B7	A31#	C12	A23#	D17	A9#	E22	HIT#
A3	SKTOCC#	B8	A27#	C13	VSS	D18	VCC	E23	VSS
A4	Reserved	B9	VSS	C14	A16#	D19	ADS#	E24	TCK
A5	VSS	B10	A21#	C15	A15#	D20	BR0#	E25	TDO
A6	A32##	B11	A22#	C16	VCC	D21	VSS	E26	VCC
A7	A33#	B12	VCC	C17	A8#	D22	RS1#	E27	FERR#
A8	VCC	B13	A13#	C18	A6#	D23	BPRI#	E28	VCC
A9	A26#	B14	A12#	C19	VSS	D24	VCC	E29	VSS
A10	A20#	B15	VSS	C20	REQ3#	D25	Reserved	E30	VCC
A11	VSS	B16	A11#	C21	REQ2#	D26	VSSENSE	E31	VSS

Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name
A12	A14#	B17	VSS	C22	VCC	D27	VSS	F1	VCC
A13	A10#	B18	A5#	C23	DEFER#	D28	VSS	F2	VSS
A14	VCC	B19	REQ0#	C24	TDI	D29	VCC	F3	VID0
A15	Reserved	B20	VCC	C25	VSS	D30	VSS	F4	VCC
A16	Reserved	B21	REQ1#	C26	IGNNE#	D31	VCC	F5	BPM3#
A17	LOCK#	B22	REQ4#	C27	SMI#	E1	VSS	F6	BPM0#
A18	VCC	B23	VSS	C28	VCC	E2	VCC	F7	VSS
A19	A7#	B24	LINT0	C29	VSS	E3	VID1	F8	VPM1#
A20	A4#	B25	PROCHOT#	C30	VCC	E4	BPM5#	F9	GTLREF
A21	VSS	B26	VCC	C31	VSS	E5	IERR#	F10	VCC
A22	A3#	B27	VCCSENSE	D1	VCC	E6	VCC	F11	VINIT#
A23	HITM#	B28	VSS	D2	VSS	E7	BPM2#	F12	BR1#
A24	VCC	B29	VCC	D3	VID2	E8	BPM4#	F13	VSS
A25	TMS	B30	VSS	D4	STPCLK#	E9	VSS	F14	ADSTB1#
A26	Reserved	B31	VCC	D5	VSS	E10	AP0#	F15	A19#
A27	VSS	C1	VSS	D6	INIT#	E11	BR2# ¹	F16	VCC
A28	VCC	C2	VCC	D7	MCERR#	E12	VCC	F17	ADSTB0#
A29	VSS	C3	VID3	D8	VCC	E13	A28#	F18	DBSY#
A30	VCC	C4	VCC	D9	AP1#	E14	A24#	F19	VSS
A31	VSS	C5	Reserved	D10	BR3# ¹	E15	VSS	F20	BNR#
B1	Reserved	C6	RSP#	D11	VSS	E16	COMP1	F21	RS2#
B2	VSS	C7	VSS	D12	A29#	E17	VSS	F22	VCC
B3	VID4	C8	A35#	D13	A25#	E18	DRDY#	F23	GTLREF
B4	VCC	C9	A34#	D14	VCC	E19	TRDY#	F24	TRST#
B5	OTDEN	C10	VCC	D15	A18#	E20	VCC	F25	VSS
F26	THERMTRIP#	J3	VSS	L24	VCC	P1	VSS	T9	VSS
F27	A20M#	J4	VCC	L25	VSS	P2	VCC	T23	VSS
F28	VSS	J5	VSS	L26	VCC	P3	VSS	T24	VCC
F29	VCC	J6	VCC	L27	VSS	P4	VCC	T25	VSS
F30	VSS	J7	VSS	L28	VCC	P5	VSS	T26	VCC
F31	VCC	J8	VCC	L29	VSS	P6	VCC	T27	VSS
G1	VSS	J9	VSS	L30	VCC	P7	VSS	T28	VCC
G2	VCC	J23	VSS	L31	VSS	P8	VCC	T29	VSS
G3	VSS	J24	VCC	M1	VCC	P9	VSS	T30	VCC
G4	VCC	J25	VSS	M2	VSS	P23	VSS	T31	VSS
G5	VSS	J26	VCC	M3	VCC	P24	VCC	U1	VCC
G6	VCC	J27	VSS	M4	VSS	P25	VSS	U2	VSS
G7	VSS	J28	VCC	M5	VCC	P26	VCC	U3	VCC
G8	VCC	J29	VSS	M6	VSS	P27	VSS	U4	VSS
G9	VSS	J30	VCC	M7	VCC	P28	VCC	U5	VCC
G23	LINT1	J31	VSS	M8	VSS	P29	VSS	U6	VSS
G24	VCC	K1	VCC	M9	VCC	P30	VCC	U7	VCC
G25	VSS	K2	VSS	M23	VCC	P31	VSS	U8	VSS

Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name
G26	VCC	K3	VCC	M24	VSS	R1	VCC	U9	VCC
G27	VSS	K4	VSS	M25	VCC	R2	VSS	U23	VCC
G28	VCC	K5	VCC	M26	VSS	R3	VCC	U24	VSS
G29	VSS	K6	VSS	M27	VCC	R4	VSS	U25	VCC
G30	VCC	K7	VCC	M28	VSS	R5	VCC	U26	VSS
G31	VSS	K8	VSS	M29	VCC	R6	VSS	U27	VCC
H1	VCC	K9	VCC	M30	VSS	R7	VCC	U28	VSS
H2	VSS	K23	VCC	M31	VCC	R8	VSS	U29	VCC
H3	VCC	K24	VSS	N1	VCC	R9	VCC	U30	VSS
H4	VSS	K25	VCC	N2	VSS	R23	VCC	U31	VCC
H5	VCC	K26	VSS	N3	VCC	R24	VSS	V1	VSS
H6	VSS	K27	VCC	N4	VSS	R25	VCC	V2	VCC
H7	VCC	K28	VSS	N5	VCC	R26	VSS	V3	VSS
H8	VSS	K29	VCC	N6	VSS	R27	VCC	V4	VCC
H9	VCC	K30	VSS	N7	VCC	R28	VSS	V5	VSS
H23	VCC	K31	VCC	N8	VSS	R29	VCC	V6	VCC
H24	VSS	L1	VSS	N9	VCC	R30	VSS	V7	VSS
H25	VCC	L2	VCC	N23	VCC	R31	VCC	V8	VCC
H26	VSS	L3	VSS	N24	VSS	T1	VSS	V9	VSS
H27	VCC	L4	VCC	N25	VCC	T2	VCC	V23	VSS
H28	VSS	L5	VSS	N26	VSS	T3	VSS	V24	VCC
H29	VCC	L6	VCC	N27	VCC	T4	VCC	V25	VSS
H30	VSS	L7	VSS	N28	VSS	T5	VSS	V26	VCC
H31	VCC	L8	VCC	N29	VCC	T6	VCC	V27	VSS
J1	VSS	L9	VSS	N30	VSS	T7	VSS	V28	VCC
J2	VCC	L23	VSS	N31	VCC	T8	VCC	V29	VSS
V30	VCC	Y22	VCC	AB1	VSS	AC11	D43#	AD21	D29#
V31	VSS	Y23	D5#	AB2	VCC	AC12	D41#	AD22	DBI1#
W1	VCC	Y24	D2#	AB3	BSEL1 ^b	AC13	VSS	AD23	VSS
W2	VSS	Y25	VSS	AB4	VCCA	AC14	D50#	AD24	D21#
W3	Reserved	Y26	D0#	AB5	VSS	AC15	DP2#	AD25	D18#
W4	VSS	Y27	Reserved	AB6	D63#	AC16	VCC	AD26	VCC
W5	BCLK1	Y28	Reserved	AB7	PWRGOOD	AC17	D34#	AD27	D4#
W6	TESTHI0	Y29	SM_TS1_A1	AB8	VCC	AC18	DP0#	AD28	SM_ALERT#
W7	TESTHI1	Y30	VSS	AB9	DBI3#	AC19	VSS	AD29	SM_WP
W8	TESTHI2	Y31	VSS	AB10	D55#	AC20	D25#	AD30	VCC
W9	GTLREF	AA1	VCC	AB11	VSS	AC21	D26#	AD31	VSS
W23	GTLREF	AA2	VSS	AB12	D51#	AC22	VCC	AE2	VSS
W24	VSS	AA3	BSEL0 ^b	AB13	D52#	AC23	D23#	AE3	VCC
W25	VCC	AA4	VCC	AB14	VCC	AC24	D20#	AE4	Reserved
W26	VSS	AA5	VSSA	AB15	D37#	AC25	VSS	AE5	TEST
W27	VCC	AA6	VCC	AB16	D32#	AC26	D17#	AE6	SLP#
W28	VSS	AA7	TESTHI4	AB17	D31#	AC27	DBI0#	AE7	D58#

Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name	Pin No.	Pin Name
W29	VCC	AA8	D61#	AB18	VCC	AC28	SM_CLK	AE8	VCC
W30	VSS	AA9	VSS	AB19	D14#	AC29	SM_DAT	AE9	D44#
W31	VCC	AA10	D54#	AB20	D12#	AC30	VSS	AE10	D42#
Y1	VSS	AA11	D53#	AB21	VSS	AC31	VCC	AE11	VSS
Y2	VCC	AA12	VCC	AB22	D13#	AD1	Reserved	AE12	DBI2#
Y3	Reserved	AA13	D48#	AB23	D9#	AD2	VCC	AE13	D35#
Y4	BCLK0	AA14	D49#	AB24	VCC	AD3	VSS	AE14	VCC
Y5	VSS	AA15	VSS	AB25	D8#	AD4	VCCIOPLL	AE15	Reserved
Y6	TESTHI3	AA16	D33#	AB26	D7#	AD5	TESTHI5	AE16	Reserved
Y7	VSS	AA17	VSS	AB27	VSS	AD6	VCC	AE17	DP3#
Y8	RESET#	AA18	D24#	AB28	SM_EP_A2	AD7	D57#	AE18	VCC
Y9	D62#	AA19	D15#	AB29	SM_EP_A1	AD8	D46#	AE19	DP1#
Y10	VCC	AA20	VCC	AB30	VCC	AD9	VSS	AE20	D28#
Y11	DSTBP3#	AA21	D11#	AB31	VSS	AD10	D45#	AE21	VSS
Y12	DSTBN3#	AA22	D10#	AC1	Reserved	AD11	D40#	AE22	D27#
Y13	VSS	AA23	VSS	AC2	VSS	AD12	VCC	AE23	D22#
Y14	DSTBP2#	AA24	D6#	AC3	VCC	AD13	D38#	AE24	VCC
Y15	DSTBN2#	AA25	D3#	AC4	VCC	AD14	D39#	AE25	D19#
Y16	VCC	AA26	VCC	AC5	D60#	AD15	VSS	AE26	D16#
Y17	DSTBP1#	AA27	D1#	AC6	D59#	AD16	COMP0	AE27	VSS
Y18	DSTBN1#	AA28	SM_TS1_A0	AC7	VSS	AD17	VSS	AE28	SM_VCC
Y19	VSS	AA29	SM_EP_A0	AC8	D56#	AD18	D36#	AE29	SM_VCC
Y20	DSTBP0#	AA30	VSS	AC9	D47#	AD19	D30#		
Y21	DSTBN0#	AA31	VCC	AC10	VCC	AD20	VCC		

Note:

- These are “Reserved” pins on the Intel® Xeon™ processor. In systems utilizing the Intel® Xeon™ processor, the system designer must terminate these signals to the processor Vcc.
- Base boards treating AA3 and AB3 as Reserved will operate correctly with a bus clock of 100MHz

8.6 System Management Headers

8.6.1 Intel Management Module Connector

A 120-pin IMM Connector (J3J1) is included on the baseboard to support the optionally installed “Standard” or “Advanced” Intel Management Modules.

Table 116. IMM Connector Pin-out (J3J1)

FMC Signal Name	FMC Pin	Description
DVI_TX1M	2	Green TMDS differential DVI output of graphics chip
DVI_TX0M	3	Blue TMDS differential DVI output of graphics chip
DVI_TX1P	4	Green TMDS differential DVI output of graphics chip
DVI_TX0P	5	Blue TMDS differential DVI output of graphics chip
DVI_CLK_TX1CM	8	TMDS differential DVI clock output of graphics chip

FMC Signal Name	FMC Pin	Description
DVI_TX2M	9	Red TMDS differential DVI output of graphics chip
DVI_CLK_TX1CP	10	TMDS differential DVI clock output of graphics chip
DVI_TX2P	11	Red TMDS differential DVI output of graphics chip
SIO_MS_DAT	14	KVM mouse data from SIO
SIO_KB_DAT	15	KVM keyboard data from SIO
SIO_MS_CLK	16	KVM mouse clock from SIO
SIO_KB_CLK	17	KVM keyboard clock from SIO
PS2_MS_DAT	18	KVM passthrough mouse data from PS2 connector
PS2_KB_DAT	19	KVM passthrough keyboard data from PS2 connector
PS2_MS_CLK	20	KVM passthrough mouse clock from PS2 connector
PS2_KB_CLK	21	KVM passthrough keyboard clock from PS2 connector
KM_INHIB_N	22	KVM enable of baseboard Switch for mouse and keyboard
FML_SDA	25	Fast Management Link Data In. This signal is driven by the FML Slave, i.e. NIC controller
FML_MCL_I2CSCL	26	Fast Management Link Clock Out. This signal is driven by the FML Master, i.e. IMM. When not configured as FML, this signal is used as I2C clock.
FML_SINTEX	27	Fast Management Link Slave Interrupt/Clock Extension. This signal is driven by the FML Slave, and has a dual usage: Used as an Alert signal for the slave to notify master that data is ready to be read from slave Used as a clock Extension (Stretching) for the slave to indicate to the master to extend its low period of the clock
FML_MDA_I2CSDA	28	Fast Management Link Data Out. This signal is driven by the FML Master. When not configured as FML, this signal is used as I2C data
ICH_LCLK	31	LPC 33Mhz clock input
USB_M	32	Reserved for future use as USB input. Baseboard can leave as NC
IMM_SYSIRQ	33	KCS interrupt signal from IMM Card.
USB_P	34	Reserved for future use as USB input. Baseboard can leave as NC
ICH_LAD1	35	LPC Address/data bus Bit 1
IMM_RSMRST_N	36	When this signal is asserted, the IMM is held in reset. This is a Standby reset indication, and should be driven by a Standby monitor device such as the Heceta7 or Dallas DS1815
ICH_LFRAME_N	37	LPC Cycle Framing
ICH_LAD0	38	LPC Address/data bus Bit 0
ICH_LAD3	39	LPC Address/data bus Bit 3
ICH_LPCPD_N	40	LPC Power down indication
ICH_LAD2	41	LPC Address/data bus Bit 2
IMM_LPCRST_N	40	LPC bus reset. Must be properly buffered on motherboard to ensure monotonicity
DFP_CLK	46	Serial clock signal for DFP EDID device. Must connect to DFP_CLK pin on the graphics chip.
DFP_DAT	48	Serial data signal for DFP EDID device. Must connect to DFP_DAT pin on the Graphics chip.
IPMB_I2C_5VSB_SDA	49	Connects to IPMB header
IPMB_I2C_5VSB_SCL	50	Connects to IPMB header

FMC Signal Name	FMC Pin	Description
SMB_I2C_3VSB_SDA	51	This bus should connect to the PCI slots, ICH, and mBMC (host I/F). An isolated version of this bus (non-Standby) should connect to the DIMMs, and clock buffer(s)
SMB_I2C_3VSB_SCL	52	This bus should connect to the PCI slots, ICH, and mBMC (host I/F). An isolated version of this bus (non-Standby) should connect to the DIMMs, and clock buffer(s)
PERIPH_I2C_3VSB_SDA	53	This bus should connect to the mBMC (Peripheral I/F), SIO, Heceta, Front panel header. A level shifted version of this bus (5V Standby) should connect to the Power Supply header
PERIPH_I2C_3VSB_SCL	54	This bus should connect to the mBMC (Peripheral I/F), SIO, Heceta, Front panel header. A level shifted version of this bus (5V Standby) should connect to the Power Supply header
MCH_I2C_3V_SDA	55	This bus should connect to the Northbridge and I/O bridge (MCH and PXH respectively in the LH chipset). In a system that supports PCI hot plug, this bus should also connect to the Power control devices if possible (such as the MIC2591 for PCI Express* for example)
MCH_I2C_3V_SCL	56	This bus should connect to the Northbridge and I/O bridge (MCH and PXH respectively in the LH chipset). In a system that supports PCI hot plug, this bus should also connect to the Power control devices if possible (such as the MIC2591 for PCI Express* for example)
LAN_I2C_3VSB_SDA	57	LAN usage
LAN_I2C_3VSB_SCL	58	LAN usage
HDD_FLT_LED_N	64	Drive Fault LED output driven when IMM detects a bad drive from the Hot Swap controller on the Hot Swap disk Drive sub-system.
IMM_PS_PWR_ON_N	65	Power On Request to the system Power Supply
COOL_FLT_LED_N	66	Cool Fault LED output driven when IMM detects a bad Fan if SSI front panel is detect.
IMM_CPU_VRD_EN	67	This signal is driven by the IMM to enable the CPU VRDs and allow the VRD power good chain to complete. This signal can also be used to keep the system in reset for an extended time, beyond what the chipset RST_BTN_N can provide.
IMM_SCI_N	68	SCI event request. If ACPI EC is supported by IMM, this signal is used for ACPI interrupts.
ICH_PWR_BTN_N	69	IMM pass through of Front panel power button to chipset
IMM_SPKR_N	72	IMM uses this to create Beep Codes on the system audible alarm. This signal is configured as an Open Drain buffer in the IMM and must be pulled up to 3.3V Standby on the motherboard
FP_NMI_BTN_N	73	NMI / Diagnostic interrupt from front panel. Actual NMI generated by SMBUS command to mBMC
FP_SLP_BTN_N	74	Front panel Sleep Button input, if used
FP_ID_BTN_N	75	Front panel ID button, will cause the ID light to toggle
SYS_PWR_GD	76	Signal from the end of the baseboard VRD Power good chain. This signal should be the last VRD power good indication generated on the baseboard. Usually this would be the signal feeding the Chipset Power OK input. Used by IMM in conjunction with RST_PWRGD_PS to determine if all critical VRDs have successfully reached their nominal value.
CPU2_SKTOCC_N	80	Indicates that a processor is in the application processor socket
CLK_32K_RTC	81	This signal is used for "Synchronized clock with system RTC". IMM can synchronize own RTC with system RTC. IPMI define synchronized method. clock comes from the Chipset RTC function.

FMC Signal Name	FMC Pin	Description
CPU1_SKTOCC_N	82	Indicates that a processor is in the primary processor socket. If this socket is detected empty and there's an attempt to power up the system, the IMM will output an Error Beep Code and prevent the System from turning on
IMM_SOUT	85	EMP/SOL Serial Data Out. This is the Serial Port data output from IMM and should be connected to the SIN signal in the SIO3 device
IMM_SIN	86	EMP/SOL Serial Data In. This is the Serial Port data input into the IMM and should be connected to the SOUT signal in the SIO3 device
IMM_DCD_N	87	EMP/SOL Data Carrier Detect. This is the Serial Port Data Carrier Detect input into the IMM and should be connected to the DCD signal in the SIO3 device
IMM_RTS_N	88	EMP/SOL Request to Send. This is the Serial Port Request to Send output from IMM and should be connected to the CTS (Clear to Send) signal in the SIO3 device
IMM_DTR_N	89	EMP/SOL Data Terminal Ready. This is the Serial Port Data Terminal Ready output from IMM and should be connected to the DSR (Data Set Ready) signal in the SIO3 device
IMM_CTS_N	90	EMP/SOL Clear to Send. This is the Serial Port Clear to Send input into the IMM and should be connected to the RTS (Ready to Send) signal in the SIO device
ICMB_RX	93	Inter Chassis Communication Management Bus receive data
ICMB_TX	94	Inter Chassis Communication Management Bus transmit data
ICMB_TX_EN	96	Inter Chassis Communication Management Bus transceiver enable
IMM_RI_BUF_N	97	Ring Indicator from the EMP serial port on the baseboard
RST_PWRGD_PS	101	Power good signal from power subsystem. In typical system, this signal is connected to PWR_OK signal on power supply. This signal is monitored by the IMM to detect a Power Supply failure
LAN_SMBALERT_N	102	Alert signal from the motherboard NIC (LOM).
ICH_SLP_S4_N	103	Power Off request from the Chipset
ICH_SMI_BUFF_N	105	SMI signal from Chipset. This signal is monitored by the IMM to detect an "SMI Time-out" condition. If this signal is asserted for longer than a predefined SMI Time-out timer, an event is logged and the IMM interrogates the chipset for further data, such as fatal errors.
CHPSET_ERR_ALERT_N	106	When available from chipset, indicates that a error occurred and IMM will need interrogate Chipset for further data, such as fatal errors. If not available, leave as NC.
FP_RST_BTN_N	109	Front panel Reset Button input.
ICH_RST_BTN_N	110	Passthrough of front panel Reset button to the chipset. IMM chassis control command will also use this.
FP_PWR_BTN_N	113	Front panel power button input.
IMM_IRQ_SMI_N	116	IMM might use this signal to generate an SMI to the system.
IMM_PRES_N	120	When IMM is present, this signal is asserted. This signal can be used to notify BIOS that a module is present (via routing to GPIO), as well as to control any logic which behaves differently when IMM is present, such as the FML mux (if supported), etc

8.6.2 ICMB Header

Table 117. ICMB Header Pin-out (J1A1)

Pin	Signal Name	Type	Description
1	5 V standby	Power	+5 V Standby
2	Receive	Signal	UART signals
3	Transmit Enable	Signal	UART signals
4	Transmit	Signal	UART signals
5	Ground	GND	

8.6.3 IPMB Header

Table 118. IPMB Header Pin-out (J1G1)

Pin	Signal Name	Description
1	Local I2C SDA	BMC IMB 5 V Standby Data Line
2	GND	
3	Local I2C SCL	BMC IMB 5 V Standby Clock Line

Note: IPMB bus is only available when an Intel® Management Module (Professional or Advanced) is present.

8.6.4 OEM RMC Header

Table 119. IPMB Header Pin-out (J1B2)

Pin	Signal Name	Description
1	I2C SDA	Peripheral 3V Standby Data Line
2	GND	
3	I2C SCL	Peripheral 3V Standby Clock Line
4	5 V standby	Power
5	POST Status	ICH5 GPIO23
6	RST#	Reset
7	5V	Power
8	PWR BTN#	

8.6.5 HSBP Header

Table 120. HSBP Header Pin-out (J1K3, J1K2)

Pin	Signal Name	Description
1	5VSB SDA	Data Line
2	GND	
3	5VSB SCL	Clock Line
4	5V – HSBP_A	

	GND – HSBP_B	
--	--------------	--

8.7 PCI Slot Connector

The Intel® SE7520AF2 Server Board enables 5 PCI expansion slots: Slot 1 (PCI-X 64/133), Slot 3 (PCI Exp x4), Slot 4 (PCI Exp x8), Slot 5 (PCI-X 64/133), Slot 6 (PCI-X 64/100). The PCI Express* slots both are enabled with a x8 physical connector, although only Slot 4 has a x8 interface (Slot 3 has a x4 interface). All PCI slots support full-length full-height PCI add-in cards. The tables below list the characteristics for the PCI slots and their pin-out.

Table 121. PCI Slot Characteristics

Slot No.	Mode	Width	Speed ¹	Voltage	Notes
1	PCI-X	64-bit	133/100/66 MHz	3.3V	
3	PCI Express	Serial	2.5GHz x4	Diff	x8 connector
4	PCI Express	Serial	2.5GHz x8	Diff	x8 connector
5	PCI-X	64-bit	133/100/66 MHz	3.3V	
6	PCI-X	64-bit	100/66 MHz	3.3V	2-slot riser support enabled

Table 122. Slot 1 and 5 PCI-X 64bit 3.3V Pin-out (J1D4, J4D1)

Pin	Side B	Side A	Pin	Side B	Side A
1	-12 V	TRST#	49	M66EN	AD[09]
2	TCK	+12 V	50	Ground	Ground
3	Ground	TMS	51	Ground	Ground
4	TDO	TDI	52	AD[08]	C/BE[0]#
5	+5 V	+5 V	53	AD[07]	+3.3 V
6	+5 V	INTA#	54	+3.3 V	AD[06]
7	INTB#	INTC#	55	AD[05]	AD[04]
8	INTD#	+5 V	56	AD[03]	Ground
9	PRSNT1#	RSV	57	Ground	AD[02]
10	RSV	+3.3 V	58	AD[01]	AD[00]
11	PRSNT2#	RSV	59	+3.3 V	+3.3 V
12	Connector Key	Connector Key	60	ACK64#	REQ64#
13	Connector Key	Connector Key	61	+5 V	+5 V
14	RSV	3.3 VAUX	62	+5 V	+5 V
15	Ground	RST#		Connector Key	Connector Key
16	CLK	+3.3 V		Connector Key	Connector Key
17	Ground	GNT#	63	RSV	Ground
18	REQ#	Ground	64	Ground	C/BE[7]#
19	+3.3 V	PME#	65	C/BE[6]#	C/BE[5]#

¹ The actual bus mode/speed is determined by capabilities of the adapter installed on that bus. PCI Express slots will attempt to link train at their supported link and fall-back to a x1 link should the adapter not support the slot's maximum link.

Pin	Side B	Side A	Pin	Side B	Side A
20	AD[31]	AD[30]	66	C/BE[4]#	+3.3 V
21	AD[29]	+3.3 V	67	Ground	PAR64
22	Ground	AD[28]	68	AD[63]	AD[62]
23	AD[27]	AD[26]	69	AD[61]	Ground
24	AD[25]	Ground	70	+3.3 V	AD[60]
25	+3.3 V	AD[24]	71	AD[59]	AD[58]
26	C/BE[3]#	IDSEL	72	AD[57]	Ground
27	AD[23]	+3.3 V	73	Ground	AD[56]
28	Ground	AD[22]	74	AD[55]	AD[54]
29	AD[21]	AD[20]	75	AD[53]	+3.3 V
30	AD[19]	Ground	76	Ground	AD[52]
31	+3.3 V	AD[18]	77	AD[51]	AD[50]
32	AD[17]	AD[16]	78	AD[49]	Ground
33	C/BE[2]#	+3.3V	79	+3.3 V	AD[48]
34	Ground	FRAME#	80	AD[47]	AD[46]
35	IRDY#	Ground	81	AD[45]	Ground
36	+3.3 V	TRDY#	82	Ground	AD[44]
37	DEVSEL#	Ground	83	AD[43]	AD[42]
38	Ground	STOP#	84	AD[41]	+3.3 V
39	LOCK#	+3.3 V	85	Ground	AD[40]
40	PERR#	SMBUS SCL	86	AD[39]	AD[38]
41	+3.3 V	SMBUS SDA	87	AD[37]	Ground
42	SERR#	Ground	88	+3.3 V	AD[36]
43	+3.3 V	PAR	89	AD[35]	AD[34]
44	C/BE[1]#	AD[15]	90	AD[33]	Ground
45	AD[14]	+3.3 V	91	Ground	AD[32]
46	Ground	AD[13]	92	RSV	RSV
47	AD[12]	AD[11]	93	RSV	GND
48	AD[10]	Ground	94	Ground	RSV

Table 123. Slot 3 and 4 PCI Express* Pin-out (J2C1, J3C1)

Pin	Side B	Side A	Pin	Side B	Side A
1	+12 V	PRSNT1#	27	PETp3	Ground
2	+12 V	+12 V	28	PETn3	Ground
3	RSVD	+12 V	29	Ground	PERp3
4	Ground	Ground	30	RSVD	PERn3
5	SMCLK	JTAG2	31	PRSNT2#	Ground
6	SMDAT	JTAG3	32	Ground	RSVD
7	Ground	JTAG4	End of the x4 connector		
8	+3.3 V	JTAG5	33	PETp4	RSVD
9	JTAG1	+3.3 V	34	PETn4	Ground
10	3.3 Vaux	+3.3 V	35	Ground	PERp4
11	Wake#	PERST#	36	Ground	PERn4

Pin	Side B	Side A	Pin	Side B	Side A
	Connector Key	Connector Key	37	PETp5	Ground
12	RSVD	Ground	38	PETn5	Ground
13	Ground	REFCLK+	39	Ground	PERp5
14	PETp0	REFCLK-	40	Ground	PERn5
15	PETn0	Ground	41	PETp6	Ground
16	Ground	PERp0	42	PETn6	Ground
17	PRSNT2#	PERn0	43	Ground	PERp6
18	Ground	Ground	44	Ground	PERn6
End of the x1 connector			45	PETp7	Ground
19	PETp1	RSVD	46	PETn7	Ground
20	PETn1	Ground	47	Ground	PERp7
21	Ground	PERp1	48	PRSNT2#	PERn7
22	Ground	PERn1	49	Ground	Ground
23	PETp2	Ground	End of the x8 connector		
24	PETn2	Ground			
25	Ground	PERp2			
26	Ground	PERn2			

Table 124. Slot 6 PCI-X 64-bit 3.3V Pin-out (J4D2, Riser Capable)

Pin	Side B	Side A	Pin	Side B	Side A
1	-12 V	TRST#	49	M66EN	AD[09]
2	TCK	+12 V	50	Ground	Ground
3	Ground	Lower Slot CLK	51	Ground	Ground
4	TDO	Upper Slot CLK	52	AD[08]	C/BE[0]#
5	+5 V	+5 V	53	AD[07]	+3.3 V
6	+5 V	INTA#	54	+3.3 V	AD[06]
7	INTB#	INTC#	55	AD[05]	AD[04]
8	INTD#	+5 V	56	AD[03]	Ground
9	PRSNT1#	RSVD	57	Ground	AD[02]
10	RSVD	+3.3 V	58	AD[01]	AD[00]
11	PRSNT2#	RSVD	59	+3.3 V	+3.3 V
12	Connector Key	Connector Key	60	ACK64#	REQ64#
13	Connector Key	Connector Key	61	+5 V	+5 V
14	RSVD	3.3 VAUX	62	+5 V	+5 V
15	Ground	RST#		Connector Key	Connector Key
16	CLK	+3.3 V		Connector Key	Connector Key
17	Ground	GNT#	63	RSVD	Ground
18	REQ#	Ground	64	Ground	C/BE[7]#
19	+3.3 V	PME#	65	C/BE[6]#	C/BE[5]#
20	AD[31]	AD[30]	66	C/BE[4]#	+3.3 V
21	AD[29]	+3.3 V	67	Ground	PAR64
22	Ground	AD[28]	68	AD[63]	AD[62]
23	AD[27]	AD[26]	69	AD[61]	Ground

Pin	Side B	Side A	Pin	Side B	Side A
24	AD[25]	Ground	70	+3.3 V	AD[60]
25	+3.3 V	AD[24]	71	AD[59]	AD[58]
26	C/BE[3]#	IDSEL	72	AD[57]	Ground
27	AD[23]	+3.3 V	73	Ground	AD[56]
28	Ground	AD[22]	74	AD[55]	AD[54]
29	AD[21]	AD[20]	75	AD[53]	+3.3 V
30	AD[19]	Ground	76	Ground	AD[52]
31	+3.3 V	AD[18]	77	AD[51]	AD[50]
32	AD[17]	AD[16]	78	AD[49]	Ground
33	C/BE[2]#	+3.3V	79	+3.3 V	AD[48]
34	Ground	FRAME#	80	AD[47]	AD[46]
35	IRDY#	Ground	81	AD[45]	Ground
36	+3.3 V	TRDY#	82	Ground	AD[44]
37	DEVSEL#	Ground	83	AD[43]	AD[42]
38	PCIXCAP	STOP#	84	AD[41]	+3.3 V
39	LOCK#	+3.3 V	85	Ground	AD[40]
40	PERR#	SMBUS SCL	86	AD[39]	AD[38]
41	+3.3 V	SMBUS SDA	87	AD[37]	Ground
42	SERR#	Ground	88	+3.3 V	AD[36]
43	+3.3 V	PAR	89	AD[35]	AD[34]
44	C/BE[1]#	AD[15]	90	AD[33]	Ground
45	AD[14]	+3.3 V	91	Ground	AD[32]
46	Ground	AD[13]	92	Riser Presence1	Riser Presence2
47	AD[12]	AD[11]	93	Upper Slot REQ#	GND
48	AD[10]	Ground	94	Ground	Upper Slot GNT#

8.8 Front Panel Connectors

A standard SSI 34-pin header is provided to support a system front panel. The header contains reset, NMI, power control buttons, and LED indicators. The following tables detail the pin outs of the header.

Table 125. Front Panel 34-Pin Header Pin-out (J1E1)

Pin	Signal Name	Pin	Signal Name
1	Power LED Anode	2	5VSB
3	Key	4	Fan Fail LED Anode
5	Power LED Cathode	6	Fan Fail LED Cathode
7	HDD Activity LED Anode	8	Power Fault LED Anode
9	HDD Activity LED Cathode	10	Power Fault LED Cathode
11	Power Switch	12	NIC1 Activity LED Anode
13	GND (Power Switch)	14	NIC1 Activity LED Cathode

Pin	Signal Name	Pin	Signal Name
15	Reset Switch	16	I2C SDA
17	GND (Reset Switch)	18	I2C SCL
19	ACPI Sleep Switch	20	Chassis Intrusion
21	GND (ACPI Sleep Switch)	22	NIC2 Activity LED Anode
23	NMI to CPU Switch	24	NIC2 Activity LED Cathode
25	KEY	26	Key
27	ID LED Anode	28	System Ready Anode
29	ID LED Cathode	30	System Ready Cathode
31	ID Switch	32	HDD Fault Anode
33	GND (ID Switch)	34	HDD Fault Cathode

8.9 VGA Connector

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

Table 126. VGA Connector Pin-out (J8A1)

Pin	Signal Name
1	Red (analog color signal R)
2	Green (analog color signal G)
3	Blue (analog color signal B)
4	No connection
5	GND
6	GND
7	GND
8	GND
9	No connection
10	GND
11	No connection
12	DDCDAT
13	HSYNC (horizontal sync)
14	VSYSN (vertical sync)
15	DDCCLK

8.10 SCSI Connector

The Intel® Server Board SE7520AF2 provides two internal wide SCSI connectors (Channel A = J1H1, Channel B = J1F1). The following table details the pin-out of the SCSI connectors.

Table 127. 68-pin SCSI Connector Pin-out (J1H1, J1F1)

Connector Contact Number	Signal Name	Signal Name	Connector Contact Number
1	+DB(12)	-DB(12)	35

Connector Contact Number	Signal Name	Signal Name	Connector Contact Number
2	+DB(13)	-DB(13)	36
3	+DB(14)	-DB(14)	37
4	+DB(15)	-DB(15)	38
5	+DB(P1)	-DB(P1)	39
6	+DB(0)	-DB(0)	40
7	+DB(1)	-DB(1)	41
8	+DB(2)	-DB(2)	42
9	+DB(3)	-DB(3)	43
10	+DB(4)	-DB(4)	44
11	+DB(5)	-DB(5)	45
12	+DB(6)	-DB(6)	46
13	+DB(7)	-DB(7)	47
14	+DB(P)	-DB(P)	48
15	GROUND	GROUND	49
16	DIFFSENSE	GROUND	50
17	TERMPWR	TERMPWR	51
18	TERMPWR	TERMPWR	52
19	RESERVED	RESERVED	53
20	GROUND	GROUND	54
21	+ATN	-ATN	55
22	GROUND	GROUND	56
23	+BSY	-BSY	57
24	+ACK	-ACK	58
25	+RST	-RST	59
26	+MSG	-MSG	60
27	+SEL	-SEL	61
28	+C/D	-C/D	62
29	+REQ	-REQ	63
30	+I/O	-I/O	64
31	+DB(8)	-DB(8)	65
32	+DB(9)	-DB(9)	66
33	+DB(10)	-DB(10)	67
34	+DB(11)	-DB(11)	68

8.11 NIC Connectors

The Intel® Server Board SE7520AF2 provides two Gigabit network interfaces using two stacked identical RJ45 connectors. The following table details the pin-out of each connector.

Table 128. NIC1 and NIC2 1.0Gb RJ45 Connector Pin-out (J6A1)

Pin	Signal Name	Pin	Signal Name
1	P2V5	2	MDIA[1]-

Pin	Signal Name	Pin	Signal Name
3	MDIA[3]-	4	MDIA[2]+
5	MDIA[0]+	6	P2V5
7	P2V5	8	MDIA[3]+
9	MDIA[1]+	10	MDIA[0]-
11	MDIA[2]-	12	P2V5
13	LILED#	14	ACTLED#
15	LINKA 100#	16	LINKA 1000#

8.12 ATA Connector

The SE7520AF2 board provides one 40-pin low-density ATA-100 connector. The pin-out for this connector is listed in the following table.

Table 129. ATA-100 40-pin Connector Pin-out (J3K1)

Pin	Signal Name	Pin	Signal Name
1	RESET_L	2	GND
3	DD7	4	IDE_DD8
5	DD6	6	IDE_DD9
7	DD5	8	IDE_DD10
9	DD4	10	IDE_DD11
11	DD3	12	IDE_DD12
13	DD2	14	IDE_DD13
15	DD1	16	IDE_DD14
17	DD0	18	IDE_DD15
19	GND	20	KEY
21	IDE_DMAREQ	22	GND
23	IDE_IOW_L	24	GND
25	IDE_IOR_L	26	GND
27	IDE_IORDY	28	GND
29	IDE_DMAACK_L	30	GND
31	IRQ_IDE	32	Test Point
33	IDE_A1	34	DIAG
35	IDE_A0	36	IDE_A2
37	IDE_DCS0_L	38	IDE_DCS1_L
39	IDE_HD_ACT_L	40	GND

8.13 USB Connector

The Intel® Server Board SE7520AF2 supports three USB connectors, which are stacked in a single housing. The pin-out for each connector is identical and is detailed in the following table.

Table 130. USB Connectors Pin-out (J9A2)

Pin	Signal Name
1	USB_PWR<0> (Fused 5 V)

2	USB_BCK0_L
3	USB_BCK0
4	GND
5	USB_PWR<1> (Fused 5 V)
6	USB_BCK1_L
7	USB_BCK1
8	GND
9	USB_PWR<2> (Fused 5 V)
10	USB_BCK2_L
11	USB_BCK2
12	GND

In addition, a header on the server board (DH10) provides an option to support two additional USB ports. The pin-out of the header is detailed in the following table.

Table 131. Optional USB Connection Header Pin-out (J1K4)

Pin	Signal Name	Description
1	USB_PWR<5>	USB Port 5 Power
2	USB_PWR<4>	USB Port 4 Power
3	USB_BCK5_L	USB Port 5 Negative Signal
4	USB_BCK4_L	USB Port 4 Negative Signal
5	USB_BCK5	USB Port 5 Positive Signal
6	USB_BCK4	USB Port 4 Positive Signal
7	Ground	
8	Ground	
9	No Connect	KEY
10	TP_USB_OVRCUR3_L	Front Panel USB Overcurrent signal. This signal is not used

8.14 Floppy Connector

The Intel® Server Board SE7520AF2 provides a standard 34-pin interface to the floppy drive controller. The following tables detail the pin-out of the 34-pin legacy floppy connector.

Table 132. Legacy 34-pin Floppy Connector Pin-out (J3K2)

Pin	Signal Name	Pin	Signal Name
1	GND	2	FD_DENSEL0
3	GND	4	Test Point
5	KEY	6	FD_DENSEL1
7	GND	8	FD_INDEX_L
9	GND	10	FD_MTR0_L
11	GND	12	FD_DS1_L
13	GND	14	FD_DS0_L
15	GND	16	FD_MTR1_L
17	Test Point	18	FD_DIR_L
19	GND	20	FD_STEP_L

21	GND	22	FD_WDATA_L
23	GND	24	FD_WGATE_L
25	GND	26	FD_TRK0_L
27	Test Point	28	VCC
29	GND	30	FD_RDATA_L
31	GND	32	FD_HDSEL_L
33	GND	34	FD_DSKCHG_L

8.15 Serial Port Connector

The SE7520AF2 supports two serial ports:

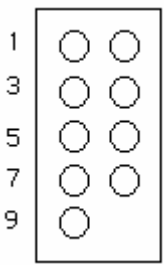
- A DB-9 connector located on the back I/O area of the baseboard enabling Serial Port A
- A 9-pin DH-10 header on the server board (J1B1) enables an optional Serial Port B port

The following tables detail the pin-outs of these two ports:

Table 133. Rear DB-9 Serial A Port Pin-out (J8A1)

Pin	Signal Name	Description
7	RTS	Request To Send
4	DTR	Data Terminal Ready
3	TD	Transmit Data
5	SGND	Signal Ground
9	RI	Ring Indicate
2	RD	Receive Data
1	DCD	Carrier Detect
8	CTS	Clear to send
6	DSR	Data Set Ready

Table 134. 9-pin Header Serial B Port Pin-out (J1B1)

Pin	Signal Name	Description	COM2 Pin-out
1	DCD	Carrier Detect	
2	DSR	Data Set Ready	
3	RD	Receive Data	
4	RTS	Request To Send	
5	TD	Transmit Data	
6	CTS	Clear To Send	
7	DTR	Data Terminal Ready	
8	RI	Ring Indicator	
9	SGND	Signal Ground	

8.16 Keyboard and Mouse Connector

Two stacked PS/2* ports are provided in a single housing for keyboard and mouse support. Although the board set supports swapping of these connections, the top port is color-coded green to designate mouse support, and the bottom port is color-coded violet to designate keyboard support. The following table details the pin-out of the PS/2 connectors.

Table 135. Keyboard and Mouse PS2 Connector Pin-out (J9A1)

Keyboard		Mouse	
Pin	Signal Name	Pin	Signal Name
1	KBDATA	1	MSDATA
2	N/C	2	N/C
3	GND	3	GND
4	Fused 5V	4	Fused 5V
5	KBCLK	5	MSCLK
6	N/C	6	N/C

8.17 Fan Headers

The Intel® Server Board SE7520AF2 provides eight fan headers: two for CPU fans and six for system fans. The CPU fans are labeled “CPU1_FAN” and “CPU2_FAN”. The CPU fan connectors only support steady 12-volt power.

The system fans are labeled “SYS FAN_1” through “SYS FAN_6”. All system fan connectors have variable speed control and are capable of supporting variable speed fans. “SYS FAN_1” and “SYS FAN_2” are enabled via a 3-pin + 2-pin header; “SYS FAN_3” and “SYS FAN_4” are enabled via a 6-pin header; and “SYS FAN_5” and “SYS FAN_6” are enabled via a 3-pin header.

Table 136. 3-pin CPU Fan Headers Pin-out (J6F1, J5F1)

Pin	Signal Name	Type	Description
1	Ground	Power	GROUND is the power supply ground
2	Fan Power	Power	Straight 12V
3	Fan Tach	Out	FAN_TACH signal is connected to the BMC to monitor the FAN speed

Table 137. 3 pin System Fan Headers Pin-out (J5A2, J5A1)

Pin	Signal Name	Type	Description
1	Ground	Power	GROUND is the power supply ground
2	Fan Power	Power	Variable Speed Fan Power
3	Fan Tach	Out	FAN_TACH signal is connected to the BMC to monitor the FAN speed

Table 138. 3+2 pin System Fan Headers Pin-out (J2J3+J2K3, J2J2+J2K1)

Pin	Signal Name	Type	Description
1	Ground	Power	GROUND is the power supply ground
2	Fan Power	Power	Variable Speed Fan Power
3	Fan Tach	Out	FAN_TACH signal is connected to the BMC to monitor the FAN speed
Pin	Signal Name	Type	Description
1	Fan LED	Power	Hot Swap fan LED power
2	Fan Presence	In	Fan Presence detect signal

Table 139. 6-pin System Fan Headers Pin-out (J2K4, J2K2)

Pin	Signal Name	Type	Description
1	Fan LED	Power	Hot Swap fan LED power
2	Fan Presence	In	Fan Presence detect signal
3	PWM	In	
4	Ground	Power	GROUND is the power supply ground
5	Fan Power	Power	Variable Speed Fan Power
6	Fan Tach	Out	FAN_TACH signal is connected to the BMC to monitor the FAN speed

9. Configuration Jumpers

This section describes configuration jumper options on the Server Board SE7520AF2.

9.1 System Recovery and Update Jumpers

The Intel® Server Board SE7520AF2 provides an 11-pin single inline header (J1D1), located on the edge of the baseboard next to the Front Panel connector, this connector provides a total of three 3-pin jumper blocks that are used to configure several system recovery and update options. The figure below shows the factory default locations for each jumper option.

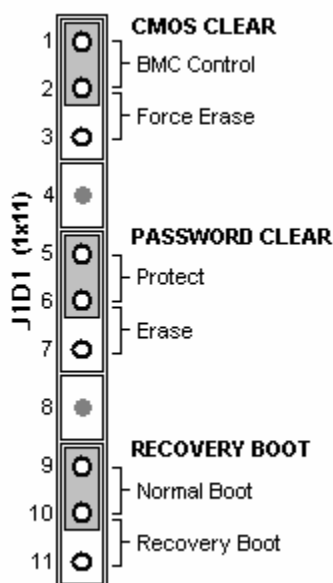


Figure 42. SE7520AF2 Configuration Jumpers (J1D1)

The following table describes each jumper option.

Table 140. Configuration Jumper Options

Option	Description
CMOS Clear	If pins 1 and 2 are jumpered (default), preservation of configuration CMOS through system reset is controlled by the BMC. If pins 2 and 3 are jumpered, CMOS contents are set to manufacturing default during system reset.
Password Clear	If pins 1 and 2 are jumpered (default), the current BIOS Setup Utility passwords are maintained during system reset. If pins 2 and 3 are jumpered, the Administrator and user passwords are cleared on reset.
Recovery Boot	If pins 1 and 2 are jumpered (default) the system will attempt to boot using the BIOS programmed in the Flash memory. If pins 2 and 3 are jumpered, the BIOS will attempt a recovery boot, loading BIOS code from a floppy disk into the Flash device. This is typically used when the BIOS code has been corrupted.

9.2 Rolling BIOS Bank Selection Jumper

The Intel® Server Board SE7520AF2 provides a 3-pin header (J2J6), located next to the SATA headers on the board to provide the option to force the board to boot from Bank 0 as part of the Rolling BIOS feature. The figure below shows the factory default location for the jumper option.

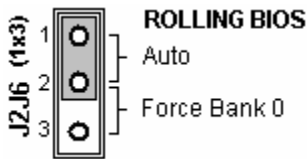


Figure 43. SE7520AF2 BIOS Bank Jumper (J2J6)

The following table describes the jumper option.

Table 141. BIOS Bank Jumper Option

Option	Description
Auto	If pins 1 and 2 are jumpered (default), the platform instrumentation on the board controls which BIOS bank has the BIOS the board is intended to boot from.
Force Bank0	If pins 2 and 3 are jumpered, the server board is forced to boot by using the BIOS code stored in Bank0.

10. General Specifications

10.1 Estimated Mean-Time Between Failures (MTBF)

The estimated Mean-time Between Failures (MTBF) is calculated at 67,640 hours at a maximum operating temperature of 35 C.

10.2 Absolute Maximum Ratings

Operating an SE7520AF2 baseboard at conditions beyond those shown in the following table may cause permanent damage to the system. The table is provided for stress testing purposes only. Exposure to absolute maximum rating conditions for extended periods may affect system reliability.

Table 142. Absolute Maximum Ratings

Operating Temperature	0 degrees C to 55 degrees C
Non-operating Temperature	-40 degrees C to +70 degrees C
Voltage on any signal with respect to ground	-0.3 V to Vdd + 0.3V
3.3 V Supply Voltage with Respect to ground	-0.3 V to 3.63 V
5 V Supply Voltage with Respect to ground	-0.3 V to 5.5 V

Notes:

1. Chassis design must provide proper airflow to avoid exceeding Intel® Xeon processor maximum case temperature.
2. VDD means supply voltage for the device

Note: Intel Corporation server boards contain a number of high-density VLSI and power delivery components which 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 air flow required for their specific application and environmental conditions. Intel Corporation can not 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.

10.3 Processor Power Support

The SE7520AF2 is designed to support the Thermal Design Point (TDP) guideline for Intel® Xeon™ processors. In addition, the Flexible Motherboard Guidelines (FMB) have been followed to help determine the suggested thermal and current design values for anticipating future processor needs. The following table provides maximum values for I_{cc}, TDP power and T_{CASE} for the Intel® Xeon™ processor family

Table 143. Intel® Xeon™ processor DP TDP Guidelines

TDP Power	Max TCASE	Icc MAX
103 W	72° C	92 A

Note: These values are for reference only. The processor EMTS contains the actual specifications for the processor. If the values found in the EMTS are different then those published here, the EMTS values will supersede these, and should be used.

10.4 SE7520AF2 Power Budget

The following section describes the power consumption of the Intel® Server Board SE7520AF2.

Table 144. SE7520AF2 Power Budget

	+3.3V (amps)	+5V (amps)	+12V (amps)	-12V (amps)	+5VSB (amps)	Total System Watts
Server Board	11.68	3.83	1.25	0.003	0.04	72.91
Processors – Qty 2	0.00	0.00	17.17	0.00	0.00	206.00
Memory – Qty 8	0.00	0.00	4.16	0.00	0.00	50.00
PCI Adapters – Qty 5	12.00	4.00	3.00	0.00	0.00	95.60
Hard Drives – Qty 10	0.00	7.50	9.00	0.00	0.00	145.50
Chassis and Peripherals	0.00	3.73	9.40	0.00	0.00	131.45
System Totals	23.68	19.06	43.98	0.003	0.04	701.38

10.5 Power Supply Specifications

This section provides power supply design guidelines for an SE7520AF2-based system, including voltage and current specifications, and power supply on/off sequencing characteristics.

Table 145. SE7520AF2 Power Supply Voltage Specification

Output	Min	Max	Tolerance
+3.3 V	3.14 V	3.46 V	+5 / -5 %
+5 V	4.75 V	5.25 V	+5 / -5 %
+12 V	11.40 V	12.60 V	+5 / -5%
-12 V	-11.40 V	-13.08 V	+5 / -9 %
+5 V SB	4.75 V	5.25 V	+5/ -5%

10.5.1 Power Timing

This section discusses the timing requirements for operation with a single power supply. The output voltages must rise from 10% to within regulation limits ($T_{\text{vout_rise}}$) within 5 ms to 70 ms. The +3.3 V, +5 V and +12 V output voltages start to rise approximately at the same time. All outputs must rise monotonically. The +5 V output must be greater than the +3.3 V output during any point of the voltage rise, however, never by more than 2.25 V. Each output voltage shall

reach regulation within 50 ms ($T_{\text{vout_on}}$) of each other and begin to turn off within 400 ms ($T_{\text{vout_off}}$) of each other. The following figure shows the output voltage timing parameters.

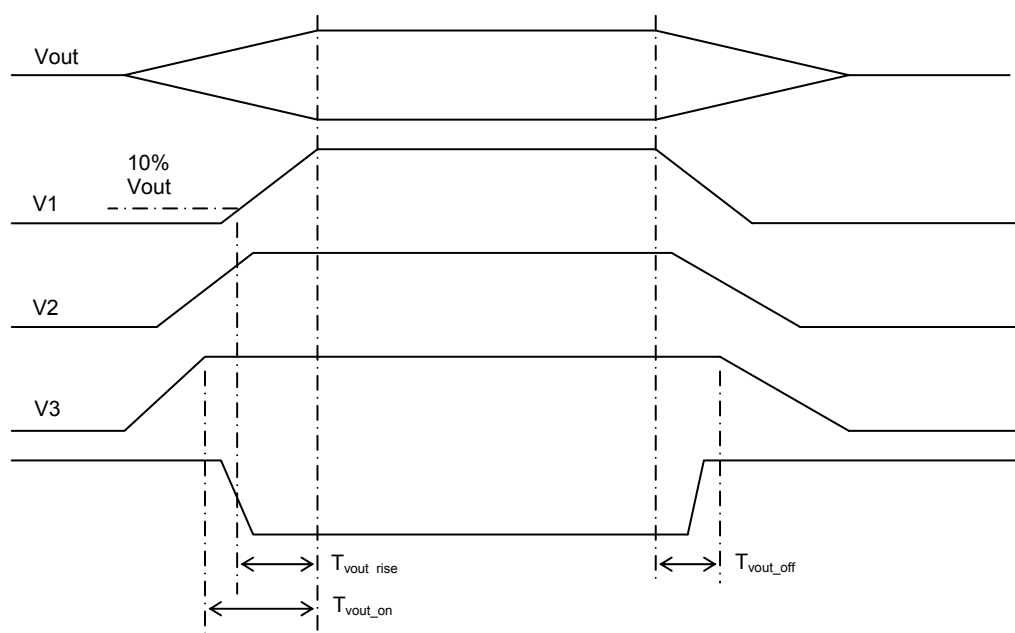


Figure 44. Output Voltage Timing

The following tables show the timing requirements for a single power supply being turned on and off via the AC input, with PSON held low and the PSON signal, with the AC input applied. The ACOK# signal is not being used to enable the turn on timing of the power supply.

Table 146. Voltage Timing Parameters

Item	Description	Min	Max	Units
$T_{\text{vout_rise}}$	Output voltage rise time from each main output.	5	70	msec
$T_{\text{vout_on}}$	All main outputs must be within regulation of each other within this time.		50	msec
$T_{\text{vout_off}}$	All main outputs must leave regulation within this time.		400	msec

Table 147. Turn On / Off Timing

Item	Description	Min	Max	Units
Tsb_on_delay	Delay from AC being applied to 5VSB being within regulation.		1500	msec
T ac_on_delay	Delay from AC being applied to all output voltages being within regulation.		2500	msec
Tvout_holdup	Time all output voltages stay within regulation after loss of AC.	21		msec
Tpwok_holdup	Delay from loss of AC to de-assertion of PWOK	20		msec
Tpson_on_delay	Delay from PSON# active to output voltages within regulation limits.	5	400	msec

T pson_pwok	Delay from PSON# deactive to PWOK being de-asserted.		50	msec
Tpwok_on	Delay from output voltages within regulation limits to PWOK asserted at turn on.	100	1000	msec
T pwok_off	Delay from PWOK de-asserted to output voltages (3.3V, 5V, 12V, -12V) dropping out of regulation limits.	1	200	msec
Tpwok_low	Duration of PWOK being in the de-asserted state during an off/on cycle using AC or the PSON signal.	100		msec
Tsb_vout	Delay from 5 V SB being in regulation to O/Ps being in regulation at AC turn on.	50	1000	msec
T5vsb_holdup	Time the 5 VSB output voltage stays within regulation after AC lost.	70		msec

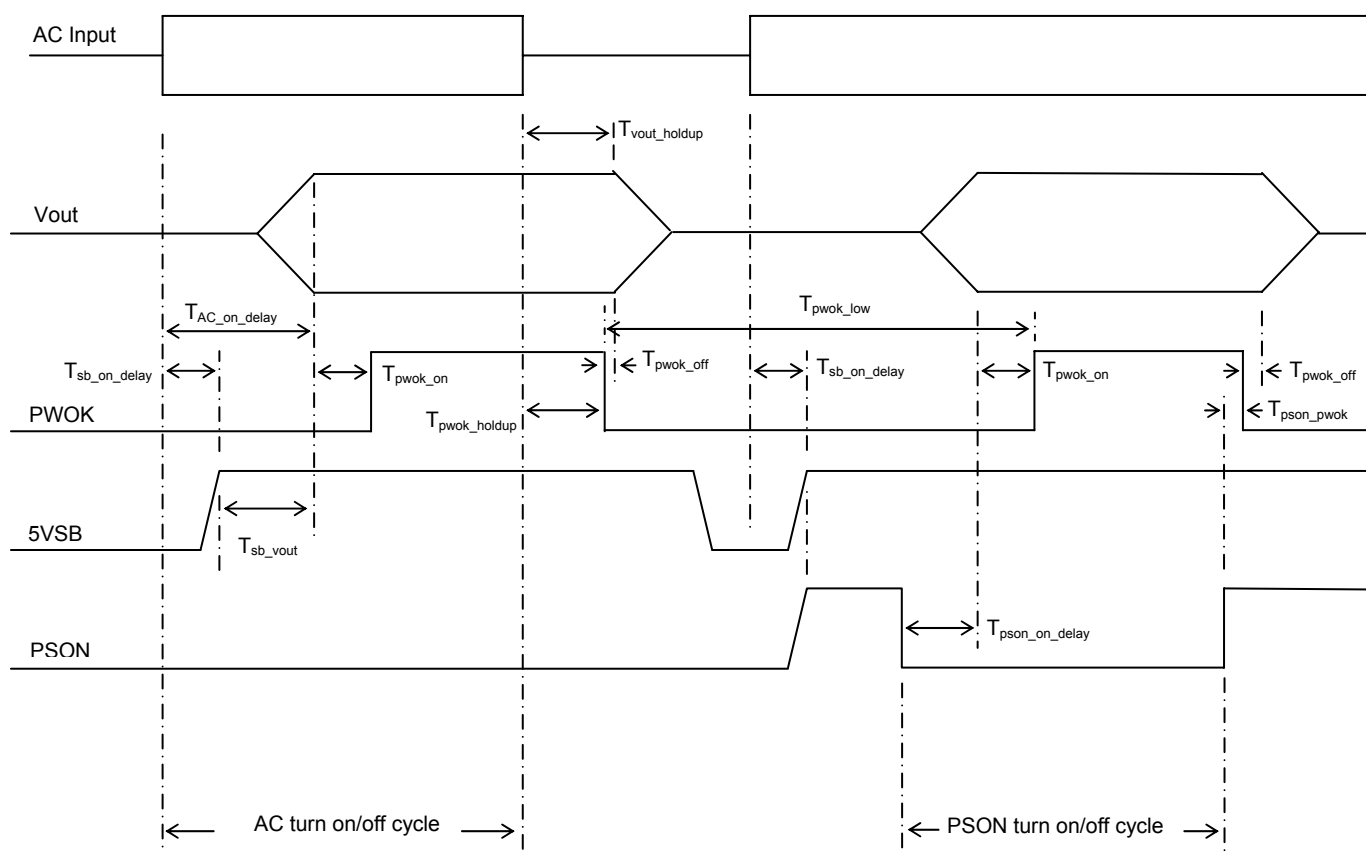


Figure 45. Turn on / off Timing

10.5.2 Voltage Recovery Timing Specifications

The power supply must conform to the following specifications for voltage recovery timing under load changes:

- Voltage shall remain within +/- 5% of the nominal set voltage on the +5 V, +12 V, 3.3 V, -5 V and -12 V output, during instantaneous changes in load shown in the following table.
- Voltage regulation limits shall be maintained over the entire AC input range and any steady state temperature and operating conditions specified.
- Voltages shall be stable as determined by bode plot and transient response. The combined error of peak overshoot, set point, regulation, and undershoot voltage shall be less than or equal to +/-5% of the output voltage setting. The transient response measurements shall be made with a load changing repetition rate of 50 Hz to 5 kHz. The load slew rate shall not be greater than 0.2 A/ μ s.

Table 148. Transient Load Requirements

Output	Step Load Size	Load Slew Rate	Capacity Load
+3.3 V	7.0 A	0.25 A/ μ s	4700 μ F
+5 V	7.0 A	0.25 A/ μ s	1000 μ F
+12 V	6.25 A	0.25 A/ μ s	675 μ F
+5 VSB	500 mA	0.25 A/ μ s	20 μ F

<This page intentionally left blank>

11. Product Regulatory Compliance

11.1.1 Product Safety Compliance

The Intel® Server Board SE7520AF2 complies with the following safety requirements:

- UL 1950 - CSA 950 (US/Canada)
- EN 60 950 (European Union)
- IEC60 950 (International)
- CE – Low Voltage Directive (73/23/EEC) (European Union)
- EMKO-TSE (74-SEC) 207/94 (Nordics)
- GOST R 50377-92 (Russia)

11.1.2 Product EMC Compliance

The SE7520AF2 has been tested and verified to comply with the following electromagnetic compatibility (EMC) regulations when installed in a compatible Intel host system. For information on compatible host system(s), contact your local Intel representative.

- FCC (Class A Verification) – Radiated and Conducted Emissions (USA)
- ICES-003 (Class A) – Radiated and Conducted Emissions (Canada)
- CISPR 22, 3rd Edition (Class A) – Radiated and Conducted Emissions (International)
- EN55022 (Class A) – Radiated and Conducted Emissions (European Union)
- EN55024 (Immunity) (European Union)
- CE – EMC Directive (89/336/EEC) (European Union)
- AS/NZS 3548 (Class A) – Radiated and Conducted Emissions (Australia / New Zealand)
- RRL (Class A) Radiated and Conducted Emissions (Korea)
- BSMI (Class A) Radiated and Conducted Emissions (Taiwan)
- GOST R 29216-91 (Class A) Radiated and Conducted Emissions (Russia)
- GOST R 50628-95 (Immunity) (Russia)

11.1.3 Product Regulatory Compliance Markings

This product is provided with the following product certification markings:

- cURus Recognition Mark
- CE Mark
- Russian GOST Mark
- Australian C-Tick Mark
- Korean RRL MIC Mark
- Taiwan BSMI DOC Mark and BSMI EMC Warning

11.2 Electromagnetic Compatibility Notices

11.2.1 FCC (USA)

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

For questions related to the EMC performance of this product, contact:

Intel Corporation
5200 N.E. Elam Young Parkway
Hillsboro, OR 97124
1-800-628-8686

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and the receiver.
- Connect the equipment to an outlet on a circuit other than the one to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Any changes or modifications not expressly approved by the grantee of this device could void the user's authority to operate the equipment. The customer is responsible for ensuring compliance of the modified product.

Only peripherals (computer input/output devices, terminals, printers, etc.) that comply with FCC Class A or B limits may be attached to this computer product. Operation with noncompliant peripherals is likely to result in interference to radio and TV reception.

All cables used to connect to peripherals must be shielded and grounded. Operation with cables, connected to peripherals, that are not shielded and grounded may result in interference to radio and TV reception.

11.2.2 Industry Canada (ICES-003)

This digital apparatus does not exceed the Class A limits for radio noise emissions from digital apparatus set out in the interference-causing equipment standard entitled: "Digital Apparatus," ICES-003 of the Canadian Department of Communications.

Cet appareil numérique respecte les limites bruits radioélectriques applicables aux appareils numériques de Classe A prescrites dans la norme sur le matériel brouilleur: "Appareils Numériques", NMB-003 édictée par le Ministre Canadien des Communications.

11.2.3 Europe (CE Declaration of Conformity)

This product has been tested in accordance too, and complies with the Low Voltage Directive (73/23/EEC) and EMC Directive (89/336/EEC). The product has been marked with the CE mark to illustrate its compliance.

11.2.4 Australian Communications Authority (ACA) (C-Tick Declaration of Conformity)

This product has been tested to AS/NZS 3548, and complies with ACA emission requirements. The product has been marked with the C-Tick mark to illustrate its compliance.

11.2.5 Ministry of Economic Development (New Zealand) Declaration of Conformity

This product has been tested to AS/NZS 3548, and complies with New Zealand's Ministry of Economic Development emission requirements.

11.2.6 Korean RRL Compliance

This product has been tested and complies with MIC Notices No. 1997-41 and 1997-42. The product has been marked with the MIC logo to illustrate compliance.



1. 기기의 명칭(모델명) :
2. 인증번호 :
3. 인증받은 자의 상호 :
4. 제조년월일:
5. 제조자/제조국가 :

11.2.7 BSMI (Taiwan)

The BSMI DOC Mark is silk screened on the component side of the server board; and the following BSMI EMC warning is marked on the server board.



11.3 Replacing the Back-Up Battery

The lithium battery on the server board powers the RTC for up to 5 years in the absence of power. When the battery starts to weaken, it loses voltage, and the server settings stored in CMOS RAM in the RTC (for example, the date and time) may be wrong. Contact your customer service representative or dealer for a list of approved devices.



WARNING

Danger of explosion if battery is incorrectly replaced. Replace only with the same or equivalent type recommended by the equipment manufacturer. Discard used batteries according to manufacturer's instructions.



ADVARSEL!

Lithiumbatteri - Eksplosionsfare ved fejlagtig håndtering. Udskiftning må kun ske med batteri af samme fabrikat og type. Levér det brugte batteri tilbage til leverandøren.

**ADVARSEL**

Lithiumbatteri - Eksplosjonsfare. Ved utskifting benyttes kun batteri som anbefalt av apparatfabrikanten. Brukt batteri returneres apparatleverandøren.

**VARNING**

Explosionsfara vid felaktigt batteribyte. Använd samma batterityp eller en ekvivalent typ som rekommenderas av apparattillverkaren. Kassera använt batteri enligt fabrikantens instruktion.

**VAROITUS**

Paristo voi räjähtää, jos se on virheellisesti asennettu. Vaihda paristo ainoastaan laitevalmistajan suosittelemaan tyyppiin. Hävitä käytetty paristo valmistajan ohjeiden mukaisesti.

Integration and Usage Tips

This section provides a bullet list of useful information that is unique to the Intel® Server Board SE7520AF2 and should be kept in mind while assembling and configuring your SE7520AF2 based server.

- Only Intel® Xeon™ processors are supported on the Intel® Server Board SE7520AF2
- Processors must be populated in the sequential order; that is, socket CPU1 must be populated before socket CPU 2
- You do not need to populate a terminator in an unused processor socket
- The Intel® Server Board SE7520AF2 supports DDR2 400MHz SDRAM memory only. Memory installation occurs in pairs of contiguous sockets (e.g. DIMM 1A and DIMM 1B). Within each pair, the DIMMs need to be the same size and vendor. DIMM pair 1 is located closest to the edge of board.
- When integrating the Intel® Server Board SE7520AF2 in the SC5300 server chassis, users will be required to install additional standoffs and one bumper in the chassis base plate. Refer to the *Quick Start User's Guide* for further details.
- The Intel® Server Board SE7520AF2 enables six (6) System Fan Headers: Sys Fan 1 through Sys Fan 6. Sys Fan3 and Sys Fan 4 are used when integrating the server board in the SC5300 Base Chassis. Other system fans are available for reference chassis.
- When integrating the Intel® Server Board SE7520AF2 in the Intel® Server Chassis SC5300, the system utilizes the 2U passive (no fan) heatsink solution of the Intel Xeon processor.

Glossary

Term	Definition
ACPI	Advanced Configuration and Power Interface
ANSI	American National Standards Institute
AP	Application processor
ASIC	Application Specific Integrated Circuit
ASR	Asynchronous Reset
BGA	Ball-grid Array
BIOS	Basic input/output system
BIST	Built-in self test
BMC	Server board Management Controller
Bridge	Circuitry connecting one computer bus to another, allowing an agent on one to access the other.
BSP	Bootstrap processor
Byte	8-bit quantity.
CMOS	In terms of this specification, this describes the PC-AT compatible region of battery-backed 128 bytes of memory, which normally resides on the server board.
DDR	Synchronous Dynamic RAM
DMA	Direct Memory Access
DMTF	Distributed Management Task Force
ECC	Error Correcting Code
EMC	Electromagnetic Compatibility
EMP	Emergency management port.
EPS	External Product Specification
ESCD	Extended System Configuration Data
FDC	Floppy Disk Controller
FIFO	First-In, First-Out
FRB	Fault resilient booting
FRU	Field replaceable unit
GB	1024 MB.
GPIO	General purpose I/O
GUID	Globally Unique ID
HDG	Hardware Design Guide
Hz	Hertz (1 cycle/second)
I ² C	Inter-integrated circuit bus
IA	Intel® architecture
ICMB	Intelligent Chassis Management Bus
IERR	Internal error
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMI	Intelligent Platform Management Interface
IRQ	Interrupt Request
ISC	Intel® Server Control
ITP	In-target probe
KB	1024 bytes

Term	Definition
KCS	Keyboard Controller Style
LAN	Local area network
LBA	Logical Block Address
LCD	Liquid crystal display
LPC	Low pin count
LSB	Least Significant Bit
LVD	Low-Voltage Differential
MB	1024 KB
MBE	Multi-Bit Error
ms	milliseconds
MSB	Most Significant Bit
MT/s	Mega transfers per second
MTBF	Mean Time Between Failures
Mux	multiplexor
NIC	Network Interface Card
NMI	Non-maskable Interrupt
OEM	Original equipment manufacturer
Ohm	Unit of electrical resistance
P32-A	32-bit PCI Segment A
P64-B	64-bit PCI Segment B
P64-C	64-bit PCI Segment C
PBGA	Pin Ball Grid Array
PDB	Power Distribution Board
PEF	Platform Event Filtering
PERR	Parity Error
PET	Platform Even Trap
PIO	Programmable I/O
PMB	Private Management Bus
PMC	Platform Management Controller
PME	Power Management Event
PnP	Plug and Play
POST	Power-on Self Test
PWM	Pulse-Width Modulator
RAM	Random Access Memory
RI	Ring Indicate
RISC	Reduced instruction set computing
RMCP	Remote Management Control Protocol
ROM	Read Only Memory
RTC	Real Time Clock
SAF-TE	SCSI Accessed Fault-Tolerant Enclosure Specification
SBE	Single-Bit Error
SCI	System Configuration Interrupt
SDR	Sensor Data Record
SDRAM	Synchronous Dynamic RAM

Term	Definition
SEL	System Event Log
SERIRQ	Serialized Interrupt Requests
SERR	System Error
SM	Server Management
SMI	Server management interrupt. SMI is the highest priority nonmaskable interrupt
SMM	System Management Mode
SMS	System Management Software
SNMP	Simple Network Management Protocol
SPD	Serial Presence Detect
SSI	Server Standards Infrastructure
SSU	Server Setup Utility
TPS	Technical Product Specification
UART	Universal asynchronous receiver and transmitter
USB	Universal Serial Bus
VGA	Video Graphic Adapter
VID	Voltage Identification
VRM	Voltage Regulator Module
Word	16-bit quantity

Reference Specifications and Documents

Refer to the following documents for additional information:

- SE7520AF2 BIOS External Product Specification. Intel Corporation
- SE7520AF2 BMC External Product Specification. Intel Corporation
- SE7520AF2 Server Management External Architecture Specification (EAS).
- Intel® Corporation Server Management Firmware External Architecture Specification (EAS). Intel® Corporation
- Mini BMC Core External Product Specification. Intel Corporation
- Sahalee Baseboard Management Controller Core External Product Specification (Sahalee BMC Core EPS for IPMI 2.0 Systems), Intel Corporation.
- Sahalee Platform Information Area External Product Specification (Sahalee PIA EPS) ver 1.0, Intel Corporation
- Server Chassis SC5300 Technical Product Specification, Intel Corporation.
- Advanced Configuration and Power Interface Specification. Intel Corporation, Microsoft* Corporation, Toshiba Corporation.
- Intelligent Chassis Management Bus (ICMB) Specification, Version 1.0. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- 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. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- Platform Management FRU Information Storage Definition. 1998. Intel Corporation, Hewlett-Packard Company, NEC Corporation, Dell Computer Corporation.
- PCI Local Bus Specification Revision 2.2
- PCI-X Specification 1.0b
- ATI Rage XL Graphics Controller Specifications, Technical Reference Manual, Rev 2.01
- RAID I/O Steering (RAIDIOS) Specification version 1.0
- VRM 10.1 DC-DC Converter Design Guide Line
- Intel® Xeon™ processor Voltage Regulator Down Design Guidelines
- The I2C Bus and How to Use It, January 1992. Phillips Semiconductors.
- Power Supply Management Interface (PSMI), Revision 1.4, 2003. Intel Corporation