



Intel® Server System R2000WF Product Family

Technical Product Specification

An overview of product features, functions, architecture, and support specifications

Revision 1.0

July 2017

<Blank page>

Document Revision History

Date	Revision	Modifications
July 2017	1.0	Production Release

Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at Intel.com, or from the OEM or retailer.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described 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 disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

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

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

Copyright © Intel Corporation. All rights reserved.

Table of Contents

1. Introduction	15
1.1 Document Outline.....	16
1.2 Intel® Server Board Use Disclaimer.....	16
1.3 Product Errata.....	16
2. Server System Family Overview.....	17
2.1 System Features Overview.....	21
2.2 Server Board Architecture	22
2.3 Server Board Features Overview	23
2.4 Back Panel Features.....	26
2.5 Front Control Panel Options	26
2.6 Front Drive Bay Options.....	27
2.7 Locking Front Bezel Support.....	28
2.8 System Dimensional Data.....	29
2.8.1 Chassis Dimensions, 2.5" Drive Bay.....	29
2.8.2 Label Emboss Dimensions	30
2.8.3 Pull-Out Tab Label Emboss Dimensions.....	31
2.9 System Cable Routing Channels.....	32
2.10 Available Rack and Cabinet Mounting Kit Options	33
2.11 System Level Environmental Limits	34
2.12 System Packaging	35
3. System Power.....	36
3.1 Power Supply Configurations	36
3.2 Power Supply Module Options	37
3.2.1 Power Supply Module Efficiency	37
3.2.2 Power Supply Module Mechanical Overview	38
3.2.3 Power Cord Specification Requirements	40
3.3 AC Power Supply Input Specifications.....	40
3.3.1 Power Factor	41
3.3.2 AC Input Voltage Specification.....	41
3.3.3 AC Line Isolation Requirements.....	41
3.3.4 AC Line Dropout/Holdup	41
3.3.5 AC Line Fuse.....	42
3.3.6 AC Inrush	42
3.3.7 AC Line Transient Specification	42
3.3.8 Susceptibility Requirements.....	43
3.3.9 Electrostatic Discharge Susceptibility	43
3.3.10 Fast Transient/Burst	43
3.3.11 Radiated Immunity	43
3.3.12 Surge Immunity.....	43
3.3.13 Power Recovery	43

3.3.14	Voltage Interruptions	44
3.3.15	Protection Circuits	44
3.3.16	Power Supply Status LED	45
3.4	DC Power Supply Input Specifications (iPC – AXX750DCCRPS)	45
3.4.1	DC Input Voltage	45
3.4.2	DC Input Fuse.....	45
3.4.3	DC Inrush Current	45
3.4.4	DC Input under Voltage.....	46
3.4.5	DC Holdup Time and Dropout.....	46
3.4.6	DC Line Surge Voltages (Line Transients).....	46
3.4.7	Susceptibility Requirements.....	46
3.4.8	Protection Circuits	47
3.5	Cold Redundancy Support.....	48
3.5.1	Powering on Cold Standby Supplies to Maintain Best Efficiency.....	48
3.5.2	Powering on Cold Standby Supplies during a Fault or Over Current Condition.....	49
3.5.3	BMC Requirements.....	49
3.5.4	Power Supply Turn on Function	49
3.6	Closed Loop System Throttling (CLST).....	49
3.7	Smart Ride Through (SmaRT)	49
3.8	Server Board Power Connectors	49
3.8.1	Power Supply Module Card Edge Connector.....	50
3.8.2	Optional 12 V Power Connectors for High Power Add-in Cards	50
3.8.3	Hot Swap Backplane Power Connector	51
3.8.4	Peripheral Power Connector	51
4.	Thermal Management	52
4.1	Thermal Operation and Configuration Requirements.....	54
4.1.1	Memory Slot Population Requirements	55
4.2	Thermal Management Overview.....	56
4.2.1	Fan Speed Control.....	57
4.3	System Fans	60
4.4	Power Supply Module Fans	62
4.5	FRUSDR Utility	62
5.	Intel® Xeon Phi™ Coprocessor and Non-Intel GPGPU Add-in Card Support.....	63
5.1	Support Criteria	63
5.2	Intel® Xeon Phi™ Coprocessor Card – System Configuration Requirements.....	66
6.	System Storage and Peripheral Drive Bay Overview.....	67
6.1	Front Mount – Drive Bay Support.....	67
6.2	Hot Swap Drive Carriers	69
6.3	Peripheral Power Sources	71
6.4	Hot Swap Backplane Support.....	71
6.4.1	SGPIO Functionality	73
6.4.2	I2C Functionality.....	73

6.5	8 x 2.5" Drive SATA/SAS/NVMe* Combo Backplane.....	73
6.6	8 x 2.5" Drive Dual Port SAS Backplane.....	75
6.7	8 x 3.5" Drive Hot-Swap Backplane Overview.....	77
6.8	12 x 3.5" Drive Hot-Swap Backplane Overview.....	78
6.9	2 x 2.5" Hot Swap Drive Bay Accessory Kit.....	79
6.10	Internal Fixed Mount Solid State Drive (SSD) Support.....	83
7.	Storage Controller Options Overview	84
7.1	Onboard SATA Support	84
7.1.1	Staggered Disk Spin-Up	86
7.2	Onboard SATA Software RAID Support	86
7.2.1	Intel® Rapid Storage Technology Enterprise (Intel® RSTe) 5.0.....	87
7.2.2	Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 for SATA.....	87
7.3	M.2 SSD Support	89
7.3.1	Embedded RAID Support	89
7.4	PCIe NVMe* Drive Support.....	90
7.4.1	Onboard PCIe* OCuLink Connectors.....	90
7.4.2	Intel® Volume Management Device (Intel® VMD) for NVMe*	90
7.4.3	Intel® Virtual RAID on Chip (Intel® VROC) for NVMe*	94
7.4.4	NVMe* Drive Population Rules for Intel® VROC	95
7.5	Intel® Integrated RAID Module Support.....	99
7.5.1	RAID Maintenance Free Backup Unit (AXXRMFBUX) Support	100
7.6	Intel® RAID Expander Card Support	101
8.	Front Control Panel and I/O Panel Overview.....	105
8.1	I/O Panel Features	105
8.2	Control Panel Features	106
9.	PCIe* Riser Card Support.....	109
9.1	Riser Card Assembly	110
9.2	Riser Slot #1 and Riser Slot #2 Riser Card Options	112
9.2.1	Three-Slot PCIe Riser Card (iPC – A2UL8RISER2)	112
9.2.2	Two-Slot PCIe Riser Card (iPC – A2UL16RISER2).....	112
9.3	Riser Slot #3 Riser Card Option (iPC – A2UX8X4RISER).....	113
10.	OCP* Compatible Intel® Ethernet Network Adapter Support.....	114
11.	Basic and Advanced Server Management Features.....	115
11.1	Dedicated Management Port	116
11.2	Embedded Web Server.....	116
11.3	Advanced Management Feature Support (Intel® RMM4 Lite).....	118
11.3.1	Keyboard, Video, Mouse (KVM) Redirection	118
11.3.2	Remote Console	119
11.3.3	Performance.....	119
11.3.4	Availability	120
11.3.5	Security.....	120
11.3.6	Usage.....	120

11.3.7	Force-enter BIOS Setup	120
11.3.8	Media Redirection	120
Appendix A.	Integration and Usage Tips.....	122
Appendix B.	POST Code Diagnostic LED Decoder.....	123
B.1.	Early POST Memory Initialization MRC Diagnostic Codes.....	124
B.2.	BIOS POST Progress Codes.....	126
Appendix C.	POST Code Errors	129
C.1.	POST Error Beep Codes	135
Appendix D.	System Configuration Table for Thermal Compatibility.....	136
Appendix E.	System Cable Routing Diagrams	153
Appendix F.	Statement of Volatility.....	158
Glossary		161

List of Figures

Figure 1. System components overview	21
Figure 2. 2U standard top cover	21
Figure 3. Server Board Architecture	22
Figure 4. Server board component/feature identification.....	23
Figure 5. Intel® Light Guided Diagnostics – LED identification.....	24
Figure 6. Intel® Light Guided Diagnostics - DIMM fault LEDs.....	25
Figure 7. System configuration and recovery jumpers.....	25
Figure 8. Back panel feature identification.....	26
Figure 9. Front control panel options.....	26
Figure 10. 3.5" x 8 front drive bay configuration – R2308WFxxxx	27
Figure 11. 3.5" x 12 front drive bay configuration – R2312WFxxxx.....	27
Figure 12. 2.5" x 8 front drive bay configuration – R2208WFxxxx	27
Figure 13. 2.5" x 16 front drive bay configuration – R2208WFxxxx + 8 Drive Accessory Kit	27
Figure 14. 2.5" x 24 front drive bay configuration – R2224WFxxxx.....	27
Figure 15. Concept Reference Design – Front Bezel Installation.....	28
Figure 16. Chassis dimensions, 2.5" drive bay	29
Figure 17. Label emboss dimensions	30
Figure 18. Pull-out tab label emboss dimensions, non-storage models.....	31
Figure 19. Pull-out tab label area dimensions, storage models.....	31
Figure 20. System cable routing channels.....	32
Figure 21. Power supply module identification	36
Figure 22. Power supply module mechanical overview	38
Figure 23. 1100W AC power supply module mechanical drawing	39
Figure 24. 1300W AC power supply module mechanical drawing	39
Figure 25. AC power cord specification	40
Figure 26. DC power cord specification.....	40
Figure 27. Power cable for the OPT_12V.....	51
Figure 28. System Airflow and Fan Identification.....	52
Figure 29. System DIMM/DIMM blanks for 8x2.5, 16x2.5, and 8x3.5 drive configurations	55
Figure 30. System DIMM/DIMM blanks for 24x2.5 and 12x3.5 drive configurations.....	55
Figure 31. High-level fan speed control model.....	60
Figure 32. Server board 2U system fan connector locations	61
Figure 33. System fan assembly	62
Figure 34. 2-slot PCIe* riser card.....	63
Figure 35. Auxiliary 12 V power cable (iPC – AXXGPGPUCABLE)	64
Figure 36. High airflow air duct included in accessory kit AWFCOPRODUCTAD	64
Figure 37. Shipping bracket.....	65
Figure 38. Shipping bracket placement.....	66
Figure 39. 2.5" front mount drive bay, 8 drive (R2208WFxxx) and 16 drive configuration option.....	67
Figure 40. 2.5" front mount drive bay, 24 drive (R2224WFxxx) configuration.....	67

Figure 41. 3.5" drive bay, 8 drive configuration	68
Figure 42. 3.5" drive bay, 12 drive configuration.....	68
Figure 43. Drive carrier removal.....	69
Figure 44. 2.5" SSD mounted to 3.5" drive carrier.....	69
Figure 45. Drive tray LED identification	70
Figure 46. Server board peripheral power connectors	71
Figure 47. 3.5" backplane placement	72
Figure 48. 2.5" drive bay module placement options.....	72
Figure 49. 8 x 2.5" SAS/SATA/NVMe* hot swap backplane, front side	73
Figure 50. 8 x 2.5" SAS/SATA/NVMe* hot swap backplane, back side.....	74
Figure 51. 8 x 2.5" dual port SAS backplane, back side	74
Figure 52. 8 x 2.5" dual port SAS backplane, front side.....	75
Figure 53. 8 x 2.5" dual port SAS backplane, Back side.....	76
Figure 54. 8 x 3.5" HSBP Connector Identification.....	77
Figure 55. 12 x 3.5" hot swap backplane, front side.....	78
Figure 56. 12 x 3.5" HSBP connector identification.....	79
Figure 57. 2 x 2.5" rear mount backplane kit placement.....	80
Figure 58. 2 x 2.5" hot swap backplane	81
Figure 59. 2.5" Solid state drive (SSD) mounting option	83
Figure 60. Onboard storage support features	84
Figure 61. BIOS setup Mass Storage Controller Configuration screen.....	86
Figure 62. ESRT2 SATA RAID-5 upgrade key	88
Figure 63. M.2 module connector locations	89
Figure 64. Onboard OCuLink connectors	90
Figure 65. NVMe* storage bus event/error handling	91
Figure 66. VMD support disabled in BIOS setup.....	93
Figure 67. Intel® VMD support enabled in BIOS setup	93
Figure 68. Intel® VROC basic architecture overview.....	94
Figure 69. Intel® VROC upgrade key.....	95
Figure 70. Backplane cabling from two PCIe sources.....	96
Figure 71. Intel® Integrated RAID Module.....	100
Figure 72. Triple RMFBU (default) mounting bracket installation	101
Figure 73. Dual RMFBU mounting bracket installation	101
Figure 74. SAS expander card installation.....	102
Figure 75. Intel® RAID Expander RES3FV288 connector identification.....	103
Figure 76. 12 GB internal 36-Port Intel® RAID Expander Card RES3TV360 connector identification	104
Figure 77. Front I/O panel features	105
Figure 78. Front control panel features	106
Figure 79. PCIe* add-in card support	109
Figure 80. Riser card #1 bracket	110
Figure 81. Left and right views of riser card #2 and #3 bracket.....	111
Figure 82. Full height/full length add-in card support.....	111

Figure 83. Three-slot PCIe* riser card (iPC – A2UL8RISER2)	112
Figure 84. Two-slot PCIe riser card (iPC – A2UL16RISER2).....	112
Figure 85. Low profile riser card (iPC – A2UX8X4RISER).....	113
Figure 86. OCP module placement.....	114
Figure 87. Intel® RMM4 Lite activation key placement.....	116
Figure 88. Dedicated management port.....	116
Figure 89. On-board POST Diagnostic LEDs	123
Figure 90. R2308WFxxx system cable routing diagram	153
Figure 91. R2208WFxxx system cable routing diagram	154
Figure 92. R2312WFxxx system cable routing diagram	155
Figure 93. R2208WFxxx + 8 drives system cable routing diagram	156
Figure 94. R2224WFxxx system cable routing diagram	157

List of Tables

Table 1. Reference Documents	15
Table 2. Intel® Server Board S2600WF product family features.....	17
Table 3. Intel® Server System R2000WF product family configurations.....	19
Table 4. System environmental limits summary.....	34
Table 5. Product weight information	35
Table 6. 1100 Watt AC power supply efficiency (Platinum).....	37
Table 7. 1300 Watt AC power supply efficiency (Titanium).....	37
Table 8. 750 Watt DC power supply efficiency (Gold).....	37
Table 9. AC power cord specifications	40
Table 10. DC power cable connector pin-out.....	40
Table 11. AC Power Factor – 1100W Power Supply	41
Table 12 . AC Power Factor - 1300W Power Supply.....	41
Table 13. AC input voltage range, 1100 W power supply	41
Table 14. AC input voltage range, 1300 W power supply	41
Table 15. AC line holdup time, 1100 W power supply.....	42
Table 16. AC line holdup time, 1300 W power supply.....	42
Table 17. AC line sag transient performance (10 sec interval between each sagging)	42
Table 18. AC line surge transient performance.....	43
Table 19. Performance criteria	43
Table 20. Over current protection, 1300 W power supply.....	44
Table 21. Over current protection, 1100 W power supply.....	44
Table 22. Over voltage protection (OVP) limits, 1300 W power supply	44
Table 23. Over voltage protection (OVP) limits, 1100 W power supply	44
Table 24. LED indicators	45
Table 25. DC input rating.....	45
Table 26. DC holdup time.....	46
Table 27. Line voltage transient limits	46
Table 28. Performance criteria.....	46
Table 29. Over current protection	47
Table 30. Over voltage protection limits.....	48
Table 31. Example load share threshold for activating supplies.....	48
Table 32. Power supply module output power connector pin-out	50
Table 33. Riser slot power pin-out (OPT_12V_PWR_#).....	50
Table 34. Hot swap backplane power connector pin-out ("HSBP PWR").....	51
Table 35. Peripheral drive power connector pin-out ("Peripheral PWR").....	51
Table 36. System volumetric airflow, Intel® Server System R2308WFxxx.....	52
Table 37 . System volumetric airflow, Intel® Server System R2308WFxxx w/Intel® Xeon Phi™ (passive).....	52
Table 38. System volumetric airflow, Intel® Server System R2312WFxxx.....	53
Table 39. System volumetric airflow, Intel® Server System R2208WFxxx.....	53
Table 40 . System volumetric airflow, Intel® Server System R2208WFxxx w/Intel® Xeon Phi™ (passive).....	53

Table 41. System volumetric airflow, Intel® Server System R2208WFxxx + 8 Drive Acc (16x2.5")	53
Table 42 . System volumetric airflow, Intel® Server System R2208WFxxx + 8 Drive Acc (16x2.5") w/Intel® Xeon Phi™ (passive).....	53
Table 43. System volumetric airflow, Intel® Server System R2224WFxxx.....	53
Table 44. PCIe* add-in card airflow (LFM) support limits – R2312WFxxxx and R2224WFxxxx	54
Table 45. PCIe* add-in card airflow (LFM) support limits – R2208WFxxxx, R2208WFxxxx + 8 drive accessory kit, and R2308WFxxxx	55
Table 46. System fan connector pin-out.....	61
Table 47. 2-slot PCIe* riser card description and routing.....	63
Table 48. Drive status LED states	70
Table 49. Drive activity LED states.....	70
Table 50. PCIe* SSD drive status LED states	70
Table 51. I2C cable connector pin-out	74
Table 52. Power harness connector pin-out.....	75
Table 53. Power harness connector pin-out.....	76
Table 54. I2C cable connector pin-out	77
Table 55. 7-pin SATA cable connector pin-out.....	81
Table 56. Power connector pin-out	82
Table 57. I2C connector pin-out.....	82
Table 58. SGPIO connector pin-out	82
Table 59. SATA and sSATA controller feature support	85
Table 60. SATA and sSATA controller BIOS setup utility options.....	85
Table 61. CPU - PCIe* Port Routing.....	92
Table 62. Intel® VROC upgrade key options	95
Table 63. Intel® SAS RAID expander card support.....	103
Table 64. System status LED state definitions.....	107
Table 65. Power/sleep LED functional states	108
Table 66. Riser slot #1 PCIe* root port mapping	110
Table 67. Riser slot #2 PCIe* root port mapping.....	110
Table 68. Riser slot #3 PCIe* root port mapping.....	110
Table 69. Three-slot PCIe* riser card slot description	112
Table 70. Two-slot PCIe riser card slot description	112
Table 71. Low profile riser card slot description	113
Table 72. OCP* Compatible Intel® Ethernet Network Adapters	114
Table 73. Intel® Remote Management Module 4 (RMM4) options.....	115
Table 74. Basic and advanced server management features overview	115
Table 75. POST progress code LED example.....	123
Table 76. MRC progress codes	124
Table 77. MRC fatal error codes.....	125
Table 78. POST progress codes.....	126
Table 79. POST error messages and handling	130
Table 80. POST error beep codes.....	135

Table 81. Integrated BMC beep codes	135
Table 82. Thermal configuration table for system in “normal” operating mode	137
Table 83. Thermal configuration table for system in “fan fail” operating mode.....	145
Table 84. Intel® Server Board S2600WF (iPN – H48104-xxx).....	159
Table 85. 2U 3-Slot PCIe* Riser Card (iPN – H20087-xxx)	159
Table 86. 2U 2-Slot PCIe* Riser Card (iPN – H20078-xxx)	159
Table 87. Riser Slot #3 Low Profile PCIe* Riser Card (iPN – G94347-xxx)	159
Table 88. Common Front Panel Board (iPN – H29366-xxx).....	159
Table 89. Storage System Model Mini- Front Panel Board (iPN – G28538-xxx).....	159
Table 90. 2U 8 x 3.5” Hot Swap Back Plane (iPN – G97160-xxx)	159
Table 91. 2U 12 x 3.5” Hot Swap Back Plane (iPN – H88392-xxx)	159
Table 92. 2U 8 x 2.5” SAS Hot Swap Back Plane (iPN –H95877-xxx).....	160
Table 93. Intel® 2U 8 x 2.5” Dual Port Hot Swap Back Plane Accessory Kit (iPC – A2U8X25S3DPDK) 2U 8 x 2.5” Dual Port Hot Swap Back Plane (iPN – G97166-xxx)	160
Table 94. Intel® 2U 2 x 2.5” Rear Hot Swap Back Plane Accessory Kit (iPC – A2UREARHSDK) 2U 2 x 2.5” Rear Hot Swap Back Plane (iPN – G94339-xxx).....	160
Table 95. Intel® Remote Management Module Lite Accessory Option (iPC – AXXRMM4LITE2).....	160

1. Introduction

This Technical Product Specification (TPS) provides system level information for the Intel® Server System R2000WF product family.

This document describes the embedded functionality and available features of the integrated server system which includes: the chassis layout, system boards, power subsystem, cooling subsystem, storage subsystem options, and available installable options. Note that some system features are provided as configurable options and may not be included standard in every system configuration offered. Please reference the *Intel® Server Board S2600WF Product Family Configuration Guide* for a list of configured options for all system models made available.

For more additional product information, refer to the documents listed in Table 1.

Table 1. Reference Documents

Document Title	Document Classification
Intel® Server Board S2600WF Product Family Technical Product Specification	Public
Intel® Server S2600WF Product Family Configuration Guide	Public
Intel® Server System R1000WF Product Family System Integration and Service Guide	Public
Intel® Server S2600WF Product Family Power Budget & Thermal Configuration Tool	Public
Intel® Servers System BMC Firmware EPS for Intel® Xeon® processor Scalable Family	Intel Confidential
Intel® Server System BIOS EPS for Intel® Xeon® processor Scalable Family	Intel Confidential
Intel® Chipset C62X Product Family External Design Specification	Intel Confidential
Intel® Ethernet Connection X557-AT2 Product Brief	Public
Advanced Configuration and Power Interface Specification, Revision 3.0, http://www.acpi.info/ .	Public
Intelligent Platform Management Interface Specification, Version 2.0. 2004.	Public
Intelligent Platform Management Bus Communications Protocol Specification, Version 1.0. 1998	Public
Platform Support for Serial-over-LAN (SOL), TMode, and Terminal Mode External Architecture Specification, Version 1.1, 02/01/02	Public
Intel® Remote Management Module User's Guide, Intel Corporation.	Public
Alert Standard Format (ASF) Specification, Version 2.0, 23 April 2003, ©2000-2003, Distributed Management Task Force, Inc., http://www.dmtf.org .	Public
SmaRT & CLST Architecture on Intel Systems and Power Supplies Specification	Public
Intel® Remote Management Module 4 Technical Product Specification	Public
Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide	Public

EPS and EDS documents are made available under NDA with Intel and must be ordered through an Intel representative.

1.1 Document Outline

This document is divided into the following chapters:

- Chapter 1 – Introduction
- Chapter 2 – Product Family Overview
- Chapter 3 – System Power
- Chapter 4 – Thermal Management
- Chapter 5 – System Storage and Peripherals Drive Bay Overview
- Chapter 6 – Storage Controller Options Overview
- Chapter 7 – Front Control Panel and I/O Panel Overview
- Chapter 8 – PCIe* Riser Card Support
- Chapter 9 – Intel® I/O Module Support
- Chapter 10 – Basic and Advanced Server Management Features
- Appendix A – Integration and Usage Tips
- Appendix B – POST Code Diagnostic LED Decoder
- Appendix C – Post Code Errors
- Appendix D – System Configuration Tables for Thermal Compatibility
- Appendix E – System Cable Routing Diagrams
- Glossary

1.2 Intel® Server Board Use Disclaimer

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

1.3 Product Errata

Product that is currently shipping may have features or functionality that deviate from published specifications. These deviations are generally discovered after the product has gone into formal production. Intel terms these deviations as product errata. Known product errata are published in the Monthly Specification Update for the given product family which can be downloaded from <http://www.intel.com/support>.

2. Server System Family Overview

The 2U server platforms within the Intel® Server System R2000WF product family offer a variety of system options to meet varied configuration requirements of high-density, high-performance computing environments.

This chapter provides a high-level overview of the system features and available options supported in different system models within this product family. Greater detail for each major sub-system, feature or option is provided in following chapters.

The following table identifies the feature set of each supported server board.

Table 2. Intel® Server Board S2600WF product family features

	Intel® Server Board S2600WF Product Family	
Feature	S2600WFT	S2600WF0
Processor Support	- Two LGA3647-0 (Socket P) processor sockets - Support for one or two Intel® Xeon® processor Scalable family (Platinum, Gold, Silver, and Bronze) - Previous generation Intel® Xeon® processors are not supported - Maximum supported Thermal Design Power (TDP) of up to 205W (Board Only) Note: Intel Server Systems based on this server board family may support a lower maximum Thermal Design Power (TDP). See appropriate Intel System TPS for max supported TDP	
Memory	24 Total DIMM slots (12 DIMMs per processor) 6 Memory Channels per processor / 2 DIMMs per Channel - Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) - Memory Capacity - Up to 1.5TB for Gold and Platinum CPUs; Up to 768GB for Silver and Bronze CPUs - Memory data transfer rates: - Up to 2666 MT/s @ 1 DPC and 2 DPC (Processor SKU Dependent) DPC = DIMMs Per Channel - DDR4 standard voltage of 1.2V	
Intel® Chipset	Intel® C624 Chipset	Intel® C624 Chipset
Intel® Quick Assist Technology (QAT)	No	No
Intel® Omni-Path Fabric Support	Yes	Yes
On-board LAN	Dual Port RJ45 10GbE	No
OCP Module Support iPC = Intel Product Code	- iPC 557T2OCPG1P5 – Dual Port 10Gb RJ45 - iPC 527DA2OCPG1P5 – Dual Port SFP+	- iPC I357T4OCPG1P5 – Quad Port 1Gb RJ45 - iPC X527DA4OCPG1P5 – Quad Port SFP+ - iPC X557T2OCPG1P5 – Dual Port 10Gb RJ45 - iPC X527DA2OCPG1P5 – Dual Port SFP+
Intel® Integrated SAS Module Support	Yes	Yes
Onboard PCIe* NVMe Support	4 – PCIe OCuLink Connectors Intel® VMD Support Intel® RSTe VROC Support – Acc. Option	4 – PCIe OCuLink Connectors Intel® VMD Support Intel® RSTe VROC Support – Acc. Option
Onboard SATA Support	12 x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) - Two single port 7-pin SATA connectors - Two M.2 connectors – SATA / PCIe* - Two 4-port mini-SAS HD (SFF-8643) connectors - Embedded SATA Software RAID - Intel® Rapid Storage RAID Technology (RSTe) 5.0 - Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 with optional RAID 5 key support Note: ESRT2 is only supported on S2600WFT and S2600WF0 boards	

	Intel® Server Board S2600WF Product Family	
Feature	S2600WFT	S2600WF0
Riser Card Support	Concurrent support for up to three riser cards • Riser #1 – PCIe* 3.0 x24 (CPU1 x16, CPU2 x8) – 2 and 3 slot riser card options available • Riser #2 – PCIe* 3.0 x24 (CPU2 x24) – 2 and 3 slot riser card options available • Riser #3 (2U systems only) – PCIe* 3.0 (CPU 2 x12) – 2 slot riser card available	
Video	• Integrated 2D Video Controller • 16MB of DDR4 Video Memory • One DB-15 External Connector • One 14-Pin Internal connector for optional Front Panel Video support	
USB Support	• Three external USB 3.0 ports • One internal Type-A USB 2.0 port • One internal 20-pin connector for optional 2x USB 3.0 port Front Panel support • One Internal 10-pin connector for optional 2x USB 2.0 port Front Panel support	
Serial Port Support	• One external RJ-45 Serial-A port connector • One internal DH-10 Serial-B port header for optional front or rear serial port support	
Server Management	• Integrated Baseboard Management Controller, IPMI 2.0 compliant • Support for Intel® Server Management Software • On-board dedicated RJ45 management port • Support for Advanced Server Management features via an Intel® Remote Management Module 4 Lite Accessory Option (iPC – AXXRMM4LITE2)	
Security	• Intel® Trusted Platform Module 2.0 (Rest of World) – iPC- AXXTPMENC8 (Accessory Option) • Intel® Trusted Platform Module 2.0 (China Version) – iPC- AXXTPMCHNE8 (Accessory Option)	
System Fan Support	• Six System fans supported in two different connector formats hot swap (2U) and cabled (1U) - o Six 10-pin managed system fan headers (Sys_Fan 1-6) – Used for 1U system configuration - o Six 6-pin hot swap capable managed system fan connectors (Sys_Fan 1-6) – Used for 2U system Configuration	

The following table describes the features of the server system configurations. Storage system configurations include R2312WFxxxx and R2224WFxxxx. Non-storage system configurations include R2308WFxxxx, R2208WFxxxx, and R2208WFxxxx +8 (base 8 drive system with 8 drive add in module).

Table 3. Intel® Server System R2000WF product family configurations

Feature	Description
Chassis Type	2U rack mount chassis
Server Board	Intel® Server Board S2600WF product family
Maximum Supported Processor Thermal Design Power (TDP)	Up to 140 W in all Intel system configuration options 165 W processors supported on Intel® Server Chassis R2208WFxxxx and R2308WFxxxx configurations only See Thermal Config Tables at end of this document for more information
External I/O Connections	- DB-15 video connectors - Front and back (non-storage system configurations only) - RJ-45 serial port A connector - Dual RJ-45 network interface connectors (S2600WFT-based systems only) - Dedicated RJ-45 server management NIC (3) – USB 3.0 connectors on back panel (2) – USB 3.0 connectors on front panel (non-storage system configurations only) (1) – USB 2.0 connector on rack handle (storage configurations only)
Internal I/O Connectors/Headers	(1) – Type-A USB 2.0 connector (1) – DH-10 serial port B connector
System Fans	(6) – managed 60 mm hot swap capable system fans - Integrated fans included with each installed power supply module
Riser Card Support	Support for three riser cards: - Riser #1 – PCIe* 3.0 x24 - up to 3 PCIe slots - Riser #2 – PCIe* 3.0 x24 - up to 3 PCIe slots - Riser #3 – PCIe* 3.0 x16 – up to 2 PCIe slots – Low profile cards only (optional) With three riser cards installed, up to eight possible add-in cards can be supported: (4) full height/half-length + (2) full height/half-length add-in cards via Risers #1 and #2 (2) low profile add in cards via riser #3 (option)
Power Supply Options	The server system can have up to two power supply modules installed, providing support for the following power configurations: 1+0, 1+1 redundant power, and 2+0 combined power. (3) power supply options: - AC 1100W Platinum - AC 1300W Titanium - DC 750W Gold

Feature	Description
Drive Bay Options	Hot Swap Backplane Options: Note: All available backplane options have support for SAS 3.0 (12 Gb/sec) • 8 x 3.5" SAS/ SATA • 8 x 2.5" combo backplane – SAS/SATA/NVMe • 8 x 2.5" Dual Port SAS • 12 x 3.5" SAS/ SATA (supports up to 2 NVMe drives) Storage Bay Options: • 8 x 3.5" SAS/SATA hot swap drive bays • 12 x 3.5" SAS/SATA hot swap drive bays (supports up to 2 NVMe drives) • 8 x 2.5" SAS/SATA/NVMe hot swap drive bays • 16 x 2.5" SAS/SATA/NVMe hot swap drive bays • 24 x 2.5" SAS/SATA/NVMe swap drive bays • 2 x 2.5" SATA SSD Back of Chassis Hot Swap Drive Bays (accessory Option) • 2 x internal fixed mount 2.5" SSDs (all SYSTEM MODELS)
Supported Rack Mount Kit Accessory Options	AXXELVRAIL – Enhanced value rack mount rail kit - 424mm max travel length, 59kgs (130lbs.) max supported weight AXXFULLRAIL – 2U Premium rail with CMA support – 800mm max travel length, 45kgs (99lbs.) max supported weight AXXSHRTRAIL – 2U Premium rail without CMA support - 788mm max travel length, 45kgs (99lbs.) max supported weight AXX2POSTBRCKT – Two post fixed mount bracket kit AXXCMA2 – Cable Management Arm – (*supported with AXXFULLRAIL only)

2.1 System Features Overview




Figure 1. System components overview

Most 2U systems within this product family include the standard top cover shown in Figure 2.




Figure 2. 2U standard top cover

2.2 Server Board Architecture




Figure 3. Server Board Architecture

2.3 Server Board Features Overview

The illustration in Figure 4 provides a general overview of the server board, identifying key feature and component locations. For more information, refer to *Intel® Server Board S2600WF Technical Product Specification*.

Note: Intel® Server Board S2600WFT shown. Some features may not be present on Intel® Server Board S2600WF0.




Figure 4. Server board component/feature identification

WFP0017

The server board includes a number of LEDs to identify system status and/or indicate a component fault. Figure 5 and Figure 6 identify each diagnostic LED and their location.





Figure 5. Intel® Light Guided Diagnostics – LED identification



2UWFP083

Figure 6. Intel® Light Guided Diagnostics - DIMM fault LEDs




Figure 7. System configuration and recovery jumpers

2.4 Back Panel Features




Figure 8. Back panel feature identification

2.5 Front Control Panel Options




Figure 9. Front control panel options

2.6 Front Drive Bay Options




Figure 10. 3.5" x 8 front drive bay configuration – R2308WFxxxx



Figure 11. 3.5" x 12 front drive bay configuration – R2312WFxxxx



Figure 12. 2.5" x 8 front drive bay configuration – R2208WFxxxx



Figure 13. 2.5" x 16 front drive bay configuration – R2208WFxxxx + 8 Drive Accessory Kit



Figure 14. 2.5" x 24 front drive bay configuration – R2224WFxxxx

2.7 Locking Front Bezel Support

The Intel 2U chassis includes features designed into the rack handles and front drive bay by to support a locking front bezel.

Note: Intel will not offer a front bezel accessory option. OEMs looking to develop a locking front bezel can obtain necessary CAD files of the chassis from Intel to aid with front bezel development. Contact your local Intel representative for additional information.




Figure 15. Concept Reference Design – Front Bezel Installation

2.8 System Dimensional Data

2.8.1 Chassis Dimensions, 2.5" Drive Bay




Figure 16. Chassis dimensions, 2.5" drive bay

2.8.2 Label Emboss Dimensions



Figure 17. Label emboss dimensions

2.8.3 Pull-Out Tab Label Emboss Dimensions




Figure 18. Pull-out tab label emboss dimensions, non-storage models




Figure 19. Pull-out tab label area dimensions, storage models

2.9 System Cable Routing Channels

The 2U system provides cable routing channels along each chassis sidewall. No cables should be routed directly in front of the system fans or through the center of the server board between the memory slots and CPU sockets. The system fan assembly must be removed before routing cables.

See Appendix E, System Cable Routing Diagrams for additional information on supported cable routing.




Figure 20. System cable routing channels

2.10 Available Rack and Cabinet Mounting Kit Options

Advisory Note: Available rack and cabinet mounting kits are not designed to support shipment of the server system while installed in a rack. If you chose to do so, Intel advises you verify your shipping configuration with appropriate shock and vibration testing, before shipment. Intel does not perform shipping tests which cover the complex combination of unique rack offerings and custom packaging options.

Caution: Exceeding the specified maximum weight limit of a given rail kit or misalignment of the server in the rack may result in failure of the rack rails, causing damage to the system or personal injury. Using two people or the use of a mechanical assist tool to install and align the server into the rack is highly recommended.

- **AXXELVRAIL** – Enhanced Value Rack Mount Rail Kit
 - o 1U, 2U, 4U compatible
 - o Tool-less chassis attach
 - o Tools required to attach to rails to rack
 - o Rack installation length adjustment from 610mm ~ 765mm
 - o 424mm travel distance (2/3 extension from rack)
 - o 130 lbs. (59 kg) max support weight
- **AXXFULLRAIL** – Premium Rail Kit with Cable Management Arm (CMA) support
 - o 2U compatible
 - o Tool-less chassis attach
 - o Full extension from rack
 - o 800mm travel distance
 - o 99 lbs. (45 Kgs) max supported weight
- **AXXSHRTRAIL** - Premium quality rails with no CMA support
 - o 2U Compatible
 - o Tool-less installation
 - o Full extension from rack
 - o 788mm travel distance
 - o 99 lbs. (45 Kgs) max supported weight
- **AXX2POSTBRCKT** – 2-Post Fixed Mount Bracket Kit
 - o 1U and 2U compatible
 - o Tools required to attach components to rack
- **AXXCMA2** – Cable Management Arm (supports AXXFULLRAIL only)

2.11 System Level Environmental Limits

The following table defines the system level operating and non-operating environmental limits.

Table 4. System environmental limits summary

Parameter		Limits
Temperature	Operating	ASHRAE Class A2 – Continuous Operation. 10 °C to 35 °C (50 °F to 95 °F) with the maximum rate of change not to exceed 10 °C per hour.
		ASHRAE Class A3 – Includes operation up to 40 °C for up to 900 hrs per year.
		ASHRAE Class A4 – Includes operation up to 45 °C for up to 90 hrs per year.
	Shipping	-40 °C to 70 °C (-40 °F to 158 °F)
Altitude	Operating	Support operation up to 3050 m with ASHRAE class de-ratings.
Humidity	Shipping	50% to 90%, non-condensing with a maximum wet bulb of 28 °C (at temperatures from 25 °C to 35 °C)
Shock	Operating	Half sine, 2 g, 11 msec
	Unpackaged	Trapezoidal, 25 g, velocity change is based on packaged weight
	Packaged	ISTA (International Safe Transit Association) Test Procedure 3A 2008
Vibration	Unpackaged	5 Hz to 500 Hz, 2.20 g RMS random
	Packaged	ISTA (International Safe Transit Association) Test Procedure 3A 2008
AC-DC	Voltage	90 Hz to 132 V and 180 V to 264 V
	Frequency	47 Hz to 63 Hz
	Source Interrupt	No loss of data for power line drop-out of 12 msec
	Surge Non-operating and operating	Unidirectional
	Line to earth Only	AC Leads 2.0 kV I/O Leads 1.0 kV DC Leads 0.5 kV
ESD	Air Discharged	12.0 kV
	Contact Discharge	8.0 kV
Acoustics Sound Power Measured	Power	<300 W ≥300 W ≥600 W ≥1000 W
	Servers/Rack Mount Sound Power Level	7.0 7.0 7.0 7.0 Note: System configurations that include one or more Intel® Xeon Phi™ co-processor cards may exhibit sound power levels of up to 8.0 BA.

For system configuration requirements and limitations, refer to Appendix D in this document or the Intel® S2600WF Product Family Power Budget and Thermal Configuration Tool.

2.12 System Packaging

The original Intel packaging is designed to provide protection to a fully configured system and tested to meet International Safe Transit Association (ISTA) Test Procedure 3A (2008). The packaging is also designed to be re-used for shipment after system integration has been completed.

The original packaging includes two layers of boxes – an inner box and the outer shipping box – and various protective inner packaging components. The boxes and packaging components are designed to function together as a protective packaging system. When reused, all of the original packaging material must be used, including both boxes and each inner packaging component. In addition, all inner packaging components must be reinstalled in the proper location to ensure adequate protection of the system for subsequent shipment.

Note: The design of the inner packaging components does not prevent improper placement within the packaging assembly. There is only one correct packaging assembly that allows the package to meet the ISTA Test Procedure 3A (2008) limits. See the *Intel® Server System R2000WF Product Family System Integration and Service Guide* for complete packaging assembly instructions.

Failure to follow the specified packaging assembly instructions may result in damage to the system during shipment.

- **Outer shipping box external dimensions**
 - o Length: 983 mm
 - o Breadth: 577 mm
 - o Height: 260 mm
- **Inner box internal dimension**
 - o Length: 956 mm
 - o Breadth: 550 mm
 - o Height: 202 mm

Table 5. Product weight information

iPC	Product Type	Packaged Gross Weight (kg)	Packaged Gross Weight (lbs)	Un-packaged Net Weight (kg)	Un-packaged Net Weight (lbs)
R2000WFXXX	Chassis	13.08	28.78	6.52	14.34
R2312WFXXX	Chassis	13.06	28.73	6.495	14.29
R2308WFxxx	L6 System	22.88	50.34	16.32	35.9
R2312WFxxx	L6 System	23.41	51.5	16.845	37.06
R2208WFxxx	L6 System	22.35	49.17	15.785	34.73
R2224WFxxx	L6 System	24.13	53.09	17.565	38.64

Note: An L6 system does not include processors, memory, drives, or add-in cards. It is the system configuration shipped from Intel. Weights of integrated system (system configurations that include the items above) vary depending on the final system configuration. For the 2U product family, a fully integrated un-packaged system can weigh upwards of 65 lbs (29.5 kg).

3. System Power

This chapter provides a high level overview of the features and functions related to system power.

3.1 Power Supply Configurations

The server system can have up to two power supply modules installed and can support the following power supply configurations: 1+0 (single power supply), 1+1 redundant power, and 2+0 combined power (non-redundant). 1+1 redundant power and 2+0 combined power configurations are automatically configured depending on the total power draw of the system. If the total system power draw exceeds the power capacity of a single power supply module, then power from the second power supply module is utilized. Should this occur, power redundancy is lost. In a 2+0 power configuration, total power available maybe less than twice the rated power of the installed power supply modules due to the amount of heat produced with both supplies providing peak power. Should system thermal levels exceed programmed limits, platform management attempts to keep the system operational. For details, see section 3.6 and chapter 4.

Caution: Installing two power supply units with different wattage ratings in a system is not supported. Doing so does not provide power supply redundancy and results in multiple errors being logged by the system.

The power supplies are modular, allowing for tool-less insertion and extraction from a bay in the back of the chassis. When inserted, the card edge connector of the power supply mates blindly to a matching slot connector on the server board.




Figure 21. Power supply module identification

In the event of a power supply failure, redundant 1+1 power supply configurations have support for hot-swap extraction and insertion. The AC input is auto-ranging and power factor corrected.

3.2 Power Supply Module Options

There are three power supply options available for this server product family:

- 1100W AC 80 Plus Platinum – iPC **AXX1100CRPS**
- 1300W AC 80 Plus Titanium – iPC **AXX1300TCRPS**
- 750W DC 80 Plus Gold – iPC **AXX750DCCRPS**

iPC= Intel product code

3.2.1 Power Supply Module Efficiency

The following tables provide the required minimum efficiency level at various loading conditions. These are provided at three different load levels: 100%, 50% and 20%.

The AC power supply efficiency is tested over an AC input voltage range of 115 VAC to 220 VAC.

Table 6. 1100 Watt AC power supply efficiency (80 Plus Platinum)

	Loading	100% of maximum	50% of maximum	20% of maximum	10% of maximum
	Minimum Efficiency	91%	94%	90%	82%

Table 7. 1300 Watt AC power supply efficiency (80 Plus Titanium)

	Loading	100% of maximum	50% of maximum	20% of maximum	10% of maximum
	Minimum Efficiency	91%	96%	94%	90%

The DC power supply efficiency is tested with a 53 V DC input.

Table 8. 750 Watt DC power supply efficiency (80 Plus Gold)

	Loading	100% of maximum	50% of maximum	20% of maximum	10% of maximum
	Minimum Efficiency	88%	92%	88%	80%

3.2.2 Power Supply Module Mechanical Overview



1100W AC Power Supply
Module with Dual Fan



1300W AC Power Supply Module



750W DC Power Supply Module



AC Power Cable Connectors



DC Power Cable Connectors

Figure 22. Power supply module mechanical overview

The physical size of the 1100 W AC power supply enclosure is 39 mm x 74 mm x 185 mm. The power supply includes dual in-line 40 mm fans, with one mounted inside the enclosure, and the other extending outside the enclosure.

Note: The second fan protrudes from the back of the power supply bay, making the total length 213 mm front to back.

The power supply has a card edge output that interfaces with a 2x25 card edge connector in the system.




Figure 23. 1100W AC power supply module mechanical drawing




Figure 24. 1300W AC power supply module mechanical drawing

3.2.3 Power Cord Specification Requirements

The AC power cord used must meet the specification requirements listed in the following table.

Table 9. AC power cord specifications

Cable Type	SJT
Wire Size	16 AWG
Temperature Rating	105°C
Amperage Rating	13 A
Voltage Rating	125 V




Figure 25. AC power cord specification




Figure 26. DC power cord specification

Table 10. DC power cable connector pin-out

Pin	Definition
1	+ Return
2	Safety Ground
3	- 48V

3.3 AC Power Supply Input Specifications

The following sections provide the AC input specifications for systems configured with AC power supply modules.

The power supply meets the power factor requirements stated in the ENERGY STAR® Program Requirements for Computer Servers; these requirements are stated below.

3.3.1 Power Factor

Table 11. AC Power Factor – 1100W Power Supply

Output Power	10% load	20% load	50% load	100% load
Power Factor	> 0.65	> 0.80	> 0.90	> 0.95

Tested at 230 VAC, 50 Hz and 60 Hz and 115 VAC, 60 Hz.

Table 12 . AC Power Factor - 1300W Power Supply

Output power	10% load	20% load	50% load	100% load
Power factor	> 0.90	> 0.96	> 0.98	> 0.99
Current THD	< 20%	< 10%	< 8%	< 5%

Tested at 230Vac, 50Hz

3.3.2 AC Input Voltage Specification

The power supply operates within all specified limits over the following input voltage range. Harmonic distortion of up to 10% of the rated line voltage does not cause the power supply to go out of specified limits. Application of an input voltage below 85 VAC does not cause damage to the power supply, including a blown fuse.

Table 13. AC input voltage range, 1100 W power supply

Parameter	MIN	Rated	VMAX	Startup VAC	Power Off VAC
Voltage (110)	90 Vrms	100-127 Vrms	140 Vrms	85 VAC +/-4 VAC	74 VAC +/-5 VAC
Voltage (220)	180 Vrms	200-240 Vrms	264 Vrms		
Frequency	47 Hz	50/60	63 Hz		

Table 14. AC input voltage range, 1300 W power supply

Parameter	MIN	Rated	VMAX	Startup VAC	Power Off VAC
Voltage (110)	90 Vrms	100-127 Vrms	140 Vrms	85VAC +/-4VAC	74VAC +/-5VAC
Voltage (220)	180 Vrms	200-240 Vrms	264 Vrms		
Frequency	47 Hz	50/60	63 Hz		

Maximum input current at low input voltage range is measured at 90 VAC, at max load.

Maximum input current at high input voltage range is measured at 180 VAC, at max load.

This requirement is not to be used for determining agency input current markings.

3.3.3 AC Line Isolation Requirements

The power supply meets all safety agency requirements for dielectric strength. Transformer isolation between primary and secondary windings must comply with the 3000 VAC (4242 VDC) dielectric strength criteria. If the working voltage between primary and secondary dictates a higher dielectric strength test voltage the highest test voltage should be used. In addition, the insulation system must comply with reinforced insulation per safety standard IEC 950. Separation between the primary and secondary circuits, and primary to ground circuits, must comply with the IEC 950 spacing requirements.

3.3.4 AC Line Dropout/Holdup

An AC line dropout is defined to be when the AC input drops to 0VAC at any phase of the AC line for any length of time. During an AC dropout the power supply meets dynamic voltage regulation requirements. An

AC line dropout of any duration does not cause tripping of control signals or protection circuits. If the AC dropout lasts longer than the holdup time, the power supply recovers and meets all turn on requirements. The power supply meets the AC dropout requirement over rated AC voltages and frequencies. A dropout of the AC line for any duration does not cause damage to the power supply.

Table 15. AC line holdup time, 1100 W power supply

Loading	Holdup time
70%	10msec

Table 16. AC line holdup time, 1300 W power supply

Loading during AC dropout / holdup	Holdup time / Dropout duration
0% to 70% of rated load	10msec

3.3.4.1 AC Line 12 VSB Holdup

The 12 VSB output voltage stays in regulation under its full load (static or dynamic) during an AC dropout of 70 msec minimum (12 VSB holdup time) whether the power supply is in an ON or OFF state (PSON asserted or de-asserted).

3.3.5 AC Line Fuse

The power supply has one line fused in the single line fuse on the line (hot) wire of the AC input. The line fusing is acceptable for all safety agency requirements. The input fuse shall be a slow blow type. AC inrush current does not cause the AC line fuse to blow under any conditions. All protection circuits in the power supply do not cause the AC fuse to blow unless a component in the power supply has failed. This includes DC output load short conditions.

3.3.6 AC Inrush

AC line inrush current does not exceed (55 A peak – 1100 W PSU, 35 A peak – 1300 W PSU), for up to one-quarter of the AC cycle, after which the input current is no more than the specified maximum input current. The peak inrush current is less than the ratings of its critical components (including input fuse, bulk rectifiers, and surge limiting device).

The power supply meets the inrush requirements for any rated AC voltage during turn on at any phase of AC voltage, during a single cycle AC dropout condition, upon recovery after AC dropout of any duration, and over the specified temperature range (T_{op}).

3.3.7 AC Line Transient Specification

AC line transient conditions are defined as “sag” and “surge” conditions. “Sag” conditions, also commonly referred to as “brownout,” are defined as the AC line voltage dropping below nominal voltage conditions. “Surge” is defined as conditions when the AC line voltage rises above nominal voltage.

The power supply meets the requirements under the following AC line sag and surge conditions.

Table 17. AC line sag transient performance (10 sec interval between each sagging)

Duration	Sag	Operating AC Voltage	Line Frequency	Performance Criteria
0 to 1/2 AC cycle	95%	Nominal AC Voltage ranges	50/60Hz	No loss of function or performance
> 1 AC cycle	>30%	Nominal AC Voltage ranges	50/60Hz	Loss of function acceptable, self-recoverable

Table 18. AC line surge transient performance

Duration	Surge	Operating AC Voltage	Line Frequency	Performance Criteria
Continuous	10%	Nominal AC Voltages	50/60 Hz	No loss of function or performance
0 to ½ AC cycle	30%	Mid-point of nominal AC Voltages	50/60 Hz	No loss of function or performance

3.3.8 Susceptibility Requirements

The power supply meets the following electrical immunity requirements when connected to a cage with an external EMI filter which meets the criteria defined in the SSI document *EPS Power Supply Specification*. For further information on Intel standards, please request a copy of the *Intel Environmental Standards Handbook*.

Table 19. Performance criteria

Level	Description
A	The apparatus continues to operate as intended. No degradation of performance.
B	The apparatus shall continue to operate as intended. No degradation of performance beyond spec limits.
C	Temporary loss of function is allowed provided the function is self-recoverable or can be restored by the operation of the controls.

3.3.9 Electrostatic Discharge Susceptibility

The power supply complies with the limits defined in EN 55024: 1998/A1: 2001/A2: 2003 using the IEC 61000-4-2: Edition 1.2: 2001-04 test standard and performance criteria B defined in Annex B of CISPR 24.

3.3.10 Fast Transient/Burst

The power supply complies with the limits defined in EN55024: 1998/A1: 2001/A2: 2003 using the IEC 61000-4-4: Second edition: 2004-07 test standard and performance criteria B defined in Annex B of CISPR 24.

3.3.11 Radiated Immunity

The power supply complies with the limits defined in EN55024: 1998/A1: 2001/A2: 2003 using the IEC 61000-4-3: Edition 2.1: 2002-09 test standard and performance criteria A defined in Annex B of CISPR 24.

3.3.12 Surge Immunity

The power supply is tested with the system for immunity to the following for each power supply option:

- **1100W Power Supply** – AC unidirectional wave; 2.6 kV line to ground and 1.3 kV line to line, per EN 55024: 1998/A1: 2001/A2: 2003, EN 61000-4-5: Edition 1.1:2001-04.
- **1300W Power Supply** – AC Unidirectional wave; 2kV line to ground and 1kV line to line, per EN 55024: 1998/A1: 2001/A2: 2003, EN 61000-4-5: Edition 1.1:2001-04 .

The pass criteria include the following: No unsafe operation is allowed under any condition; all power supply output voltage levels to stay within proper specification levels; No change in operating state or loss of data during and after the test profile; No component damage under any condition.

The power supply complies with the limits defined in EN55024: 1998/A1: 2001/A2: 2003 using the IEC 61000-4-5: Edition 1.1:2001-04 test standard and performance criteria B defined in Annex B of CISPR 24.

3.3.13 Power Recovery

The power supply recovers automatically after an AC power failure. AC power failure is defined to be any loss of AC power that exceeds the dropout criteria.

3.3.14 Voltage Interruptions

The power supply complies with the limits defined in EN55024: 1998/A1: 2001/A2: 2003 using the IEC 61000-4-11: Second Edition: 2004-03 test standard and performance criteria C defined in Annex B of CISPR 24.

3.3.15 Protection Circuits

Protection circuits inside the power supply cause only the power supply's main outputs to shut down. If the power supply latches off due to a protection circuit tripping, an AC cycle OFF for 15 seconds and a PS_ON# cycle HIGH for one second reset the power supply.

3.3.15.1 Over-Current Protection

The power supply has current limit to prevent the outputs from exceeding the values shown in Table 20. If the current limits are exceeded, the power supply shuts down and latches off. The latch is cleared by toggling the PS_ON# signal or by an AC power interruption. The power supply is not damaged from repeated power cycling in this condition. 12 VSB is auto-recovered after removing the over current protection limit.

Table 20. Over current protection, 1300 W power supply

Output VOLTAGE	Input voltage range	OVER CURRENT LIMITS	OCP DELAY
+12V	180 – 264VAC	132A min / 138A max	50msec min / 200msec max
		152A min / 160A max	5msec min / 20msec max
+12V	90 – 140VAC	72A min / 77A max	50msec min / 200msec max
	90 – 140VAC	103A min; 107A max	5msec min / 20msec
12VSB	90 – 264VAC	2.5A min; 3.5A max	5msec min / 20msec

Table 21. Over current protection, 1100 W power supply

Output Voltage	Input Voltage Range	Over Current Limits
+12 V	90 – 264 VAC	120 A min/132 A max
	Trip delay	50 msec min
12 VSB	90 – 264 VAC	2.5 A min/3.5 A max

3.3.15.2 Over Voltage Protection (OVP)

The power supply over voltage protection is locally sensed. The power supply shuts down and latches off after an over voltage condition occurs. This latch is cleared by an AC power interruption. The values are measured at the output of the power supply's connectors. The voltage never exceeds the maximum levels when measured at the power connectors of the power supply connector during any single point of fail. The voltage never trips any lower than the minimum levels when measured at the power connector. 12 VSB is auto-recovered after removing OVP limit.

Table 22. Over voltage protection (OVP) limits, 1300 W power supply

Output Voltage	Minimum (V)	Maximum (V)
+12 V	13.5	14.5
+12 VSB	13.5	14.5

Table 23. Over voltage protection (OVP) limits, 1100 W power supply

Output Voltage	Minimum (V)	Maximum (V)
+12V	13.5	14.5
+12VSB	13.5	14.5

3.3.15.3 Over Temperature Protection (OTP)

The power supply is protected against over temperature conditions caused by loss of fan cooling or excessive ambient temperature. In an OTP condition, the PSU shuts down. OT warning SMBAlert# assertion always precedes the OTP shutdown. When the power supply temperature drops to within specified limits, the power supply restores power automatically, while the 12 VSB remains always on. The OTP circuit has built-in margin such that the power supply does not oscillate on and off due to temperature recovering condition. The OTP trip temperature level is at least 5 °C higher than SMBAlert over temperature warning threshold level.

3.3.16 Power Supply Status LED

There is a single bi-color LED to indicate power supply status. The LED operation is defined in Table 24.

Table 24. LED indicators

Power Supply Condition	LED State
Output ON and OK	Solid green
No AC power to all power supplies	Off
AC present/Only 12 VSB on (PS off) or PS in cold redundant state	1 Hz blinking green
AC cord unplugged or AC power lost; with a second power supply in parallel still with AC input power	Solid amber
Power supply warning events where the power supply continues to operate; high temp, high power, high current, slow fan	1 Hz blinking amber
Power supply critical event causing a shutdown; failure, OCP, OVP, fan fail	Solid amber
Power supply FW updating	2 Hz blinking green

3.4 DC Power Supply Input Specifications (iPC – AXX750DCCRPS)

The following sections provide the DC input specifications for systems configured with DC power supply modules.

Note: Product safety regulations pertaining to the use of DC power supplies require that chassis grounding studs be used for all DC power supply configurations. In the event that chassis grounding studs are not available on a given server chassis, systems must be configured with two DC power supplies, each connected to separate ground wires while the system is operational.

3.4.1 DC Input Voltage

The power supply must operate within all specified limits over the following input voltage range.

Table 25. DC input rating

Parameter	Minimum	Rated	Maximum
DC Voltage	-40.5 VDC	-48 VDC/-60 VDC	-75 VDC
Input Current	24bA		12.5 A

3.4.2 DC Input Fuse

The power supply should have the -48 VDC input fused. The fusing shall be acceptable for all safety agency requirements. DC inrush current shall not cause the fuse to blow under any conditions. No protection circuits in the power supply shall cause the DC fuse to blow unless a component in the power supply has failed. This includes DC output load short conditions.

3.4.3 DC Inrush Current

Maximum inrush current from power-on shall be limited to a level below the surge rating of the input line cable; input diodes, fuse, and EMI filter components. To allow multiple power cycling events and DC line

transient conditions max I^2t value shall not exceed 20% of the fuse max rating. Repetitive ON/OFF cycling of the DC input line voltage should not damage the power supply or cause the input fuse to blow.

3.4.4 DC Input under Voltage

The power supply shall contain protection circuitry (under-voltage lock-out) such that the application of an input voltage below the specified minimum specified, shall not cause damage (overstress) to the power supply unit (due to over-heating or otherwise).

3.4.5 DC Holdup Time and Dropout

Table 26. DC holdup time

Loading	Holdup Time
750 W (100%)	0.2 msec

During a DC dropout of 0.2 msec or less the power supply must meet dynamic voltage regulation requirements for every rated load condition. A DC line dropout of 0.2 msec or less shall not cause tripping of control signals or protection circuits. Repeated every ten seconds starting at the minimum input voltage, DC line dropout shall not damage the power supply under any specified load conditions. The PWOK signal shall not go to a low state under these conditions. DC dropout transients in excess of 0.2 msec may cause shutdown of the power supply or out of regulation conditions, but shall not damage the power supply. The power supply should recover and meet all turn on requirements for DC dropouts that last longer than 0.2 msec. The power supply must meet the DC dropout requirement over rated DC voltages and output loading conditions.

3.4.6 DC Line Surge Voltages (Line Transients)

The power supply should demonstrate tolerance for transients in the input DC power line caused by switching or lightning. The power supply shall be primarily tested and must be compliant with the requirements of EN61000-4-5: "Electrical Fast transients / Burst Requirements and Surge Immunity Requirements" for surge withstand capability. The test voltage surge levels are to be 500Vpk for each Line to Primary Earth Ground test (none required between the L1 and L2). The exact description can be found in *Intel Environmental Standards Handbook 2001*.

Table 27. Line voltage transient limits

Duration	Slope/Rate	Output	Performance Criteria
200 μ sec max	-48 V $\rightarrow$ -30 V w/ +2 V/ μ s	Rated DC voltages	No loss of function or performance
	-30 V $\rightarrow$ -48 V w/ -2 V/ μ s	Rated DC voltages	No loss of function or performance

3.4.7 Susceptibility Requirements

The power supply shall meet the following electrical immunity requirements when connected to a cage with an external EMI filter which meets the criteria defined in the SSI document *EPS Power Supply Specification*. For further information on Intel standards, please request a copy of the *Intel Environmental Standards Handbook*.

Table 28. Performance criteria

Level	Description
A	The apparatus shall continue to operate as intended. No degradation of performance.
B	The apparatus shall continue to operate as intended. No degradation of performance beyond spec limits.
C	Temporary loss of function is allowed provided the function is self-recoverable or can be restored by the operation of the controls.

3.4.7.1 Electrostatic Discharge Susceptibility

The power supply shall comply with the limits defined in EN 55024: 1998 using the IEC 61000-4-2:1995 test standard and performance criteria B defined in Annex B of CISPR 24. Limits shall comply with those specified in the *Intel Environmental Standards Handbook*.

3.4.7.2 Fast Transient/Burst

The power supply shall comply with the limits defined in EN55024: 1998 using the IEC 61000-4-4:1995 test standard and performance criteria B defined in Annex B of CISPR 24. Limits shall comply with those specified in the *Intel Environmental Standards Handbook*.

3.4.7.3 Radiated Immunity

The power supply shall comply with the limits defined in EN55024: 1998 using the IEC 61000-4-3:1995 test standard and performance criteria A defined in Annex B of CISPR 24. Limits shall comply with those specified in the *Intel Environmental Standards Handbook*. Additionally, the power supply must comply with field strength requirements specified in GR 1089 (10 V/m).

3.4.7.4 Surge Immunity

The power supply shall be tested with the system for surge immunity, per EN 55024:1998, EN 61000-4-5:1995 and ANSI C62.45: 1992.

The pass criteria include the following: No unsafe operation is allowed under any condition; all power supply output voltage levels to stay within proper spec levels; No change in operating state or loss of data during and after the test profile; No component damage under any condition.

The power supply shall comply with the limits defined in EN55024: 1998 using the IEC 61000-4-5:1995 test standard and performance criteria B defined in Annex B of CISPR 24. Limits shall comply with those specified in the *Intel Environmental Standards Handbook*.

3.4.8 Protection Circuits

Protection circuits inside the power supply shall cause only the power supply's main outputs to shut down. If the power supply latches off due to a protection circuit tripping, a DC cycle OFF for 15 sec and a PSON# cycle HIGH for 1 sec shall be able to reset the power supply.

3.4.8.1 Over Current Protections (OCP)

The power supply shall have current limit to prevent the outputs from exceeding the values shown in Table 29. If the current limits are exceeded, the power supply shall shutdown and latch off. The latch will be cleared by toggling the PSON# signal or by a DC power interruption. The power supply shall not be damaged from repeated power cycling in this condition. 12VSB will be auto-recovered after removing OCP limit.

Table 29. Over current protection

Output Voltage	Input Voltage Range	Over Current Limits
+12V		72 A min/78 A max
12VSB		2.5 A min/3.5 A max

3.4.8.2 Over Voltage Protection (OVP)

The power supply over voltage protection shall be locally sensed. The power supply shall shutdown and latch off after an over voltage condition occurs. This latch shall be cleared by toggling the PSON# signal or by a DC power interruption. The values are measured at the output of the power supply's connectors. The voltage shall never exceed the maximum levels when measured at the power supply connector during any single point of fail. The voltage shall never trip any lower than the minimum levels when measured at the power connector. 12 VSB will be auto-recovered after removing OVP limit.

Table 30. Over voltage protection limits

Output Voltage	Minimum (V)	Maximum (V)
+12 V	13.3	14.5
+12 VSB	13.3	14.5

3.4.8.3 Over Temperature Protection (OTP)

The power supply will be protected against over temperature conditions caused by loss of fan cooling or excessive ambient temperature. In an OTP condition the PSU will shut down. When the power supply temperature drops to within specified limits, the power supply shall restore power automatically, while the 12 VSB remains always on. The OTP circuit must have built in margin such that the power supply will not oscillate on and off during a temperature recovering condition. The OTP trip level shall have a minimum of 4 °C of ambient temperature margin.

3.5 Cold Redundancy Support

The power supplies support cold redundancy allowing them to go into a low-power state (that is, a cold redundant state) in order to provide increased power usage efficiency when system loads are such that both power supplies are not needed. When the power subsystem is in cold redundant mode, only the needed power supply to support the best power delivery efficiency is ON. Any additional power supplies, including the redundant power supply, is in cold standby state.

Each power supply has an additional signal that is dedicated to supporting cold redundancy: CR_BUS. This signal is a common bus between all power supplies in the system. CR_BUS is asserted when there is a fault in any power supply or the power supplies' output voltages fall below the V_{fault} threshold. Asserting the CR_BUS signal causes all power supplies in cold standby state to power ON.

Enabling power supplies to maintain optimum efficiency is achieved by looking at the load share bus voltage and comparing it to a programmed voltage level via a PMBus command.

Whenever there is no active power supply on the cold redundancy bus driving a HIGH level on the bus, all power supplies are ON no matter their defined cold redundant roll (active or cold standby). This guarantees that incorrect programming of the cold redundancy states of the power supply will never cause the power subsystem to shut down or become over loaded. The default state of the power subsystem is all power supplies ON. There needs to be at least one power supply in cold redundant active state or standard redundant state to allow the cold standby state power supplies to go into cold standby state.

3.5.1 Powering on Cold Standby Supplies to Maintain Best Efficiency

Power supplies in cold standby state shall monitor the shared voltage level of the load share signal to sense when it needs to power on. Depending upon which position (1, 2, or 3) the system defines that power supply to be in, the cold standby configuration slightly changes the load share threshold that the power supply shall power on at.

Table 31. Example load share threshold for activating supplies

State	Enable Threshold for VCR_ON_EN	Disable Threshold for VCR_ON_DIS	CR_BUS De-Asserted/ Asserted States
Standard Redundancy	NA; Ignore dc/dc_active# signal; power supply is always ON		OK = High, Fault = Low
Cold Redundant Active	NA; Ignore dc/dc_active# signal; power supply is always ON		OK = High, Fault = Low
Cold Standby 1 (02h)	3.2 V (40% of max)	$3.2 \text{ V} \times 0.5 \times 0.9 = 1.44 \text{ V}$	OK = Open, Fault = Low
Cold Standby 2 (03h)	5.0 V (62% of max)	$5.0 \text{ V} \times 0.67 \times 0.9 = 3.01 \text{ V}$	OK = Open, Fault = Low
Cold Standby 3 (04h)	6.7 V (84% of max)	$6.7 \text{ V} \times 0.75 \times 0.9 = 4.52 \text{ V}$	OK = Open, Fault = Low

Note: Maximum load share voltage is 8.0 V at 100% of rated output power.

Note: Load share bus thresholds are examples. For a given power supply, these will be customized to maintain the best efficiency curve for that specific model.

3.5.2 Powering on Cold Standby Supplies during a Fault or Over Current Condition

When an active power supply asserts its CR_BUS signal (pulling it low), all parallel power supplies in cold standby mode will power on within 100 μ sec.

3.5.3 BMC Requirements

The BMC uses the Cold_Redundancy_Config command to define and configure the power supply's role in cold redundancy and to turn on/off cold redundancy.

To allow for equal loading over the life time of installed power supplies, the BMC shall schedule a rolling re-configuration of installed power supplies so that each alternates between being the "Active" power supply and the "Cold Stby" power supply.

Events that trigger a re-configuration of the power supplies using the Cold_Redundancy_Config command are listed below.

- AC power ON
- PSON power ON
- Power supply failure
- Power supply inserted into system

3.5.4 Power Supply Turn on Function

Powering on and off of the cold standby power supplies is only controlled by each PSU sensing the Vshare bus. Once a power supply turns on after crossing the enable threshold, it lowers itself to the disable threshold. The system defines the position of each power supply in the cold redundant operation. It does this each time the system is powered on, a power supply fails, or a power supply is added to the system.

The system is relied upon to tell each power supply where it resides in the cold redundancy scheme.

3.6 Closed Loop System Throttling (CLST)

The server system has support for Closed Loop System Throttling (CLST) which prevents the system from crashing if a power supply module is overloaded or over heats. Should system power reach a pre-programmed power limit, CLST throttles system memory and/or processors to reduce power. System performance is impacted should this occur.

3.7 Smart Ride Through (SmaRT)

The server system has support for Smart Ride Through Throttling (SmaRT) which increases the reliability for a system operating in a heavy power load condition and to remain operational during an AC line dropout event. For power supply hold up time requirements for AC line dropout events, see section 3.3.4.

When AC voltage is too low, a fast AC loss detection circuit inside each installed power supply asserts an SMBALERT# signal to initiate a throttle condition in the system. System throttling reduces the bandwidth to both system memory and CPUs, which in turn reduces the power load during the AC line drop out event

3.8 Server Board Power Connectors

The server board provides several connectors to provide power to various system options. This section provide the pin-out definition and a brief usage description for each.

3.8.1 Power Supply Module Card Edge Connector

Each power supply module has a single 2x25 card edge output connection that plugs directly into a matching slot connector on the server board. The connector provides both power and communication signals to the server board. The following table defines the connector pin-out.

Table 32. Power supply module output power connector pin-out

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

3.8.2 Optional 12 V Power Connectors for High Power Add-in Cards

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

Each connector is labeled as “OPT_12V_PWR” on the server board. The following table provides the pin-out for both connectors.

Table 33. Riser slot power pin-out (OPT_12V_PWR_#)

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




Figure 27. Power cable for the OPT_12V

The power cable for the OPT_12V_PWR connectors (shown in Figure 27) is included in the two-slot riser card accessory kit (iPC – A2UL16RISER2) and can support both six- and eight-pin 12 V AUX power connectors found on high power add-in cards.

3.8.3 Hot Swap Backplane Power Connector

The server board includes one white 2x6-pin power connector that is used to provide power for the hot swap backplanes mounted in the front of the chassis only. On the server board, this connector is labeled as “HSBP PWR”. The power for this connector is routed from two separate 12 V power rails, each capable of supporting a current limit of 18 A. The following table provides the pin-out for this connector.

Table 34. Hot swap backplane power connector pin-out (“HSBP PWR”)

Signal Description	Pin#	Pin#	Signal Description
GND	1	7	P12V_240VA3
GND	2	8	P12V_240VA3
GND	3	9	P12V_240VA2
GND	4	10	P12V_240VA2
GND	5	11	P12V_240VA1
GND	6	12	P12V_240VA1

3.8.4 Peripheral Power Connector

The server board includes one brown 2x3-pin power connector intended to provide power to various possible peripheral options including: Solid State Devices (SSDs) mounted to the top side of the air duct or to the rear mounted two drive backplane accessory option. On the server board this connector is labeled as “Peripheral PWR”. The following table provides the pin-out for this connector.

Table 35. Peripheral drive power connector pin-out (“Peripheral PWR”)

Pin	Signal Name	Pin	Signal Name
1	P5V	4	P12V
2	P5V	5	P3V3
3	GND	6	GND

4. Thermal Management

The fully integrated system is designed to operate at external ambient temperatures of between 10 °C and 35 °C with limited excursion based operation up to 45 °C, as specified in Table 4. Working with integrated platform management, several features within the system are designed to move air in a front to back direction, through the system and over critical components to prevent them from overheating and allow the system to operate with best performance.




Figure 28. System Airflow and Fan Identification

The following tables provide air flow data associated with the different system models within this 2U product family and are provided for reference purposes only. The data was derived from actual wind tunnel test methods and measurements using fully configured (worst case) system configurations. Lesser system configurations may produce slightly different data results. In addition, the CFM data provided using server management utilities that utilize platform sensor data may vary slightly from the data listed in the tables.

Note: For system BTU data, see the Intel® Server Board S2600WF Product Family Power Budget and Thermal Configuration Tool.

Table 36. System volumetric airflow, Intel® Server System R2308WFxxx

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	136.2
100%	100%	141.5
55%	auto	70.8

Table 37 . System volumetric airflow, Intel® Server System R2308WFxxx w/Intel® Xeon Phi™ (passive)

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	170.6
100%	100%	173.3
55%	auto	86.5

Table 38. System volumetric airflow, Intel® Server System R2312WFxxx

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	77.9
100%	100%	80.0
55%	auto	40.3

Table 39. System volumetric airflow, Intel® Server System R2208WFxxx

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	154.6
100%	100%	159.1
55%	auto	84.0

Table 40 . System volumetric airflow, Intel® Server System R2208WFxxx w/Intel® Xeon Phi™ (passive)

System Airflow – R2208WFxxx w/Intel Xeon Phi (passive)		
System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	189.0
100%	100%	194.2
55%	auto	97.6

Table 41. System volumetric airflow, Intel® Server System R2208WFxxx + 8 Drive Acc (16x2.5")

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	136.2
100%	100%	140.5
55%	auto	69.8

Table 42 . System volumetric airflow, Intel® Server System R2208WFxxx + 8 Drive Acc (16x2.5") w/Intel® Xeon Phi™ (passive)

System Airflow – R2216WFxxx w/Intel Xeon Phi (passive)		
System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	163.4
100%	100%	165.7
55%	auto	87.0

Table 43. System volumetric airflow, Intel® Server System R2224WFxxx

System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	105.5
100%	100%	108.4
55%	auto	48.9

Note: Intel® Xeon Phi™ processor is not supported in Intel® Server Systems R2312WFxxx and R2224WFxxx.

Note: Airflow data is preliminary and is subject to change.

The Intel® Server System R2000WF product family supports short-term, excursion-based, operation up to 45 °C (ASHRAE A4) with limited performance impact. The configuration requirements and limitations are described in the configuration matrix found in Appendix D of this document or in the Intel® S2600WF Product Family Power Budget and Thermal Configuration Tool available as a download online at:

<http://www.intel.com/support>

The installation and functionality of several system components are used to maintain system thermals. They include six managed 60 mm system fans, fans integrated into each installed power supply module, an air duct, populated drive carriers, and installed CPU heatsinks. Drive carriers can be populated with a storage device (SSD or HDD) or supplied drive blank. In addition, it may be necessary to have specific DIMM slots populated with DIMMs or supplied DIMM blanks. System configurations that require population of specific DIMM slots ship from Intel with DIMM blanks pre-installed. Pre-installed DIMM blanks should only be removed when installing a memory module in its place. Chassis only products include DIMM blanks separately in the accessory kit and must be installed in the appropriate DIMM slot locations as defined in the following section.

4.1 Thermal Operation and Configuration Requirements

To keep the system operating within supported maximum thermal limits, the system must meet the following operating and configuration guidelines:

- The system is designed to sustain operations at an ambient temperature of up to 35 °C (ASHRAE Class A2) with short term excursion based operation up to 45 °C (ASHRAE Class A4).
- The system can operate up to 40 °C (ASHRAE Class A3) for up to 900 hours per year.
- The system can operate up to 45 °C (ASHRAE Class A4) for up to 90 hours per year.
- System performance may be impacted when operating within the extended operating temperature range.
- There is no long term system reliability impact when operating at the extended temperature range within the documented limits.

Specific configuration requirements and limitations are documented in the configuration matrix found in Appendix D of this document or in the Intel® Server Board S2600WF product family Power Budget and Thermal Configuration Tool, available as a download online at <http://www.intel.com/support>.

The CPU-1 processor and CPU heat sink must be installed first. The CPU-2 heat sink must be installed at all times, with or without a processor installed.

Thermally, a system supporting fan redundancy can support the following PCI add-in cards when the system is operating at up to a maximum operating ambient temperature of 35 °C (ASHRAE Class 2).

Airflow for add-in cards is measured at local inlet. Add-in card thermal support limits are listed by riser and by riser slot in the following tables. This limitation is driven primarily by fan failure performance.

Table 44. PCIe* add-in card airflow (LFM) support limits – R2312WFxxx and R2224WFxxx

PCIe* Add-In Slot	Riser Slot #1	Riser Slot #2	Riser Slot #3
Top	100 LFM	200 LFM	100 LFM
Middle	100 LFM	200 LFM	100 LFM
Bottom	200 LFM	200 LFM	N/A

Table 45. PCIe* add-in card airflow (LFM) support limits – R2208WFxxx, R2208WFxxx + 8 drive accessory kit, and R2308WFxxx

PCIe* Add-In Slot	Riser Slot #1	Riser Slot #2	Riser Slot #3
Top	200 LFM	200 LFM	200 LFM
Middle	200 LFM	200 LFM	200 LFM
Bottom	300 LFM	300 LFM	N/A

Note: Most PCI cards have cooling requirements of airflow up to 100 LFM (0.5 m/s). Some of the more difficult to cool cards have airflow requirements of up to 200 LFM (1 m/s). The top and middle PCI locations for the first riser have limited thermals during fan fail.

4.1.1 Memory Slot Population Requirements

System thermal requirements dictate that specific airflow be maintained over or between critical system components. To ensure that proper air flow is achieved, specific memory slots must be populated with a DIMM or factory installed DIMM blank while the system is in operation. Figure 29 and Figure 30 identify the memory slots, identified with an arrow, which must be populated in all 2U system configurations.

NOTE: To maintain system thermals while the system is operational, DIMM slots identified with  must be populated with a DIMM or supplied DIMM blank

**FOR SYSTEMS WITH THE FOLLOWING FRONT DRIVE BAY CONFIGURATIONS:
8x2.5", 16x2.5" 8x3.5"**




Figure 29. System DIMM/DIMM blanks for 8x2.5, 16x2.5, and 8x3.5 drive configurations

**FOR SYSTEMS WITH THE FOLLOWING FRONT DRIVE BAY CONFIGURATIONS:
24x2.5", 12x3.5"**




Figure 30. System DIMM/DIMM blanks for 24x2.5 and 12x3.5 drive configurations

The following memory population rules apply when installing DIMMs:

- DIMM population rules require that DIMMs within a channel be populated starting with the BLUE DIMM slot or DIMM farthest from the processor in a "fill-farthest" approach.

- When only one DIMM is used for a given memory channel, it must be populated in the BLUE DIMM slot (furthest from the CPU).
- Mixing of DDR4 DIMM Types (RDIMM, LRDIMM, 3DS RDIMM, 3DS LRDIMM, or NVDIMM) within a channel socket or across sockets produces a Fatal Error Halt during Memory Initialization.
- Mixing DIMMs of different frequencies and latencies is not supported within or across processor sockets. If a mixed configuration is encountered, the BIOS will attempt to operate at the highest common frequency and the lowest latency possible.
- When populating a Quad-rank DIMM with a Single- or Dual-rank DIMM in the same channel, the Quad-rank DIMM must be populated farthest from the processor. Intel MRC will check for correct DIMM placement. A maximum of 8 logical ranks can be used on any one channel, as well as a maximum of 10 physical ranks loaded on a channel.
- In order to install 3 QR LRDIMMs on the same channel, they must be operated with Rank Multiplication as $RM = 2$, this will make each LRDIMM appear as a DR DIMM with ranks twice as large.
- The memory slots associated with a given processor are unavailable if the corresponding processor socket is not populated.
- A processor may be installed without populating the associated memory slots, provided a second processor is installed with associated memory. In this case, the memory is shared by the processors. However, the platform suffers performance degradation and latency due to the remote memory.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as Memory RAS, Error Management,) in the BIOS setup are applied commonly across processor sockets.
- For multiple DIMMs per channel:
 - For RDIMM, LRDIMM, 3DS RDIMM, 3DS LRDIMM; Always populate DIMMs with higher electrical loading in slot1, followed by slot 2.
 - NVM DIMMs must be populated in DIMM slot 1 of a given memory channel when paired with an RDIMM or 3DS RDIMM type

4.2 Thermal Management Overview

In order to maintain the necessary airflow within the system, all of the previously listed components and top cover need to be properly installed. For optimum system performance, the external ambient temperature should remain below 35 °C and all system fans should be operational. System fan redundancy can be supported for system configurations that meet the necessary fan redundancy support limits identified in Appendix D.

For system configurations that support fan redundancy, should a single fan failure occur (system fan or power supply fan), integrated platform management changes the state of the system status LED to blinking green, reports an error to the system event log, and automatically adjusts fan speeds as needed to maintain system temperatures below maximum thermal limits.

Note: All system fans are controlled independently of each other. The fan control system may adjust fan speeds for different fans based on increasing/decreasing temperatures in different thermal zones within the chassis.

In the event that system temperatures should continue to increase with the system fans operating at their maximum speed, platform management may begin to throttle bandwidth of either the memory subsystem,

the processors, or both to keep components from overheating and keep the system operational. Throttling of these sub-systems continues until system temperatures are reduced below preprogrammed limits.

The power supply is protected against over temperature conditions caused by excessive ambient temperature. In an over-temperature protection condition, the power supply module shuts down.

Should system thermals increase to a point beyond the maximum thermal limits, the system shuts down, the system status LED changes to solid amber, and the event is logged to the system event log. Should power supply thermals increase to a point beyond their maximum thermal limits or if a power supply fan should fail, the power supply shuts down.

Note: For proper system thermal management, Sensor Data Records (SDRs) for any given system configuration must be loaded by the system integrator as part of the initial system integration process. SDRs are loaded using the FRUSDR utility which is part of the System Update Package (SUP) or One-boot Firmware Update (OFU) package which can be downloaded from <http://downloadcenter.intel.com>.

4.2.1 Fan Speed Control

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

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

A fan domain has three states:

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

The fan domain state is controlled by several factors. They are listed below in order of precedence, high to low:

- Boost – if any of the below conditions apply, the fans are set to a fixed boost state speed
 - o Associated fan is in a critical state or missing. The SDR describes which fan domains are boosted in response to a fan failure or removal in each domain. If a fan is removed when the system is in 'Fans-off' mode it is not detected and there is not any fan boost until the system comes out of 'Fans-off' mode.
 - o Any associated temperature sensor is in a critical state. The SDR describes which temperature threshold violations cause fan boost for each fan domain.
 - o The BMC is in firmware update mode or the operational firmware is corrupted.
- Nominal
 - o A fan domain's nominal fan speed can be configured as static (fixed value) or controlled by the state of one or more associated temperature sensors.

4.2.1.1 Programmable Fan PWM Offset

The system provides a BIOS setup option to boost the system fan speed by a programmable positive offset or a Max setting. Setting the programmable offset causes the BMC to add the offset to the fan speeds to which it would otherwise be driving the fans. The Max setting causes the BMC to replace the domain minimum speed with alternate domain minimums that also are programmable through SDRs.

This capability is offered to provide system administrators the option to manually configure fan speeds in instances where the fan speed optimized for a given platform may not be sufficient when a high end add-in adapter is configured into the system. This enables easier usage of the fan speed control to support Intel as well as third party chassis and better support of ambient temperatures higher than 35 °C.

4.2.1.2 Hot-Swap Fans

Hot-swap fans are supported. These fans can be removed and replaced while the system is powered on and operating. The BMC implements fan presence sensors for each hot-swappable fan.

When a fan is not present, the associated fan speed sensor is put into the reading/unavailable state and any associated fan domains are put into the boost state. The fans may already be boosted due to a previous fan failure or fan removal.

When a removed fan is inserted, the associated fan speed sensor is rearmed. If there are no other critical conditions causing a fan boost condition, the fan speed returns to the nominal state. Power cycling or resetting the system re-arms the fan speed sensors and clears fan failure conditions. If the failure condition is still present, the boost state returns once the sensor has re-initialized and the threshold violation is detected again.

4.2.1.3 Fan Redundancy Detection

The BMC supports redundant fan monitoring and implements a fan redundancy sensor. A fan redundancy sensor generates events when it's associated set of fans transitions between redundant and non-redundant states, as determined by the number and health of the fans. The definition of fan redundancy is configuration dependent. The BMC allows redundancy to be configured on a per fan redundancy sensor basis through OEM SDR records.

A fan failure or removal of hot-swap fans up to the number of redundant fans specified in the SDR in a fan configuration is a non-critical failure and is reflected in the front panel status. A fan failure or removal that exceeds the number of redundant fans is a non-fatal, insufficient-resources condition and is reflected in the front panel status as a non-fatal error.

Redundancy is checked only when the system is in the DC-on state. Fan redundancy changes that occur when the system is DC-off or when AC is removed will not be logged until the system is turned on.

4.2.1.4 Fan Domains

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

The BMC controls the average duty cycle of each PWM signal through direct manipulation of the integrated PWM control registers.

The same device may drive multiple PWM signals.

4.2.1.5 Nominal Fan Speed

A fan domain's nominal fan speed can be configured as static (fixed value) or controlled by the state of one or more associated temperature sensors.

OEM SDR records are used to configure which temperature sensors are associated with which fan control domains and the algorithmic relationship between the temperature and fan speed. Multiple OEM SDRs can reference or control the same fan control domain and multiple OEM SDRs can reference the same temperature sensors.

The PWM duty cycle value for a domain is computed as a percentage using one or more instances of a stepwise linear algorithm and a clamp algorithm. The transition from one computed nominal fan speed

(PWM value) to another is ramped over time to minimize audible transitions. The ramp rate is configurable by means of the OEM SDR.

Multiple stepwise linear and clamp controls can be defined for each fan domain and used simultaneously. For each domain, the BMC uses the maximum of the domain's stepwise linear control contributions and the sum of the domain's clamp control contributions to compute the domain's PWM value, except that a stepwise linear instance can be configured to provide the domain maximum.

Hysteresis can be specified to minimize fan speed oscillation and to smooth fan speed transitions. If a Tcontrol SDR record does not contain a hysteresis definition (for example, an SDR adhering to a legacy format), the BMC assumes a hysteresis value of zero.

4.2.1.6 Thermal and Acoustic Management

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

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

4.2.1.7 Thermal Sensor Input to Fan Speed Control

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

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

- Front panel temperature sensor ¹
- CPU margin sensors ^{2, 4, 5}
- DIMM thermal margin sensors ^{2, 4}
- Exit air temperature sensor ^{1, 7, 9}
- PCH temperature sensor ^{3, 5}
- Onboard Ethernet controller temperature sensors ^{3, 5}
- Add-in SAS module temperature sensors ^{3, 5}
- PSU thermal sensor ^{3, 8}
- CPU VR temperature sensors ^{3, 6}
- DIMM VR temperature sensors ^{3, 6}
- BMC temperature sensor ^{3, 6}
- Global aggregate thermal margin sensors ⁷
- Hot swap backplane temperature sensors
- Intel® OCP module temperature sensor (with option installed)
- Intel® SAS module (with option installed)
- Riser card temperature sensors (2U systems only)
- Intel® Xeon Phi™ coprocessor (2U system only with option installed)

Notes:

¹ For fan speed control in Intel chassis

² Temperature margin to max junction temp

³ Absolute temperature

⁴ PECI value or margin value

⁵ On-die sensor⁶ Onboard sensor⁷ Virtual sensor⁸ Available only when PSU has PMBus⁹ Calculated estimate

Figure 31 shows a high-level representation of the fan speed control structure that determines fan speed.




Figure 31. High-level fan speed control model

4.3 System Fans

Six (6) 60 x 38 mm system fans and embedded fans for each installed power supply module provide the primary airflow for the system.

The system is designed for fan redundancy when configured with two power supply modules, all system fan rotors are operational, and ambient air remains at or below ASHRAE class 2 limits. Should a single system fan fail, platform management will adjust air flow of the remaining system fans and manage other platform features to maintain system thermals. Fan redundancy is lost if more than one system fan is in a failed state.

The server board is capable of supporting up to a total of six system fans. On the server board, each system fan includes a pair of fan connectors, a 1x10 pin connector to support a dual rotor cabled fan typically used in 1U system configurations, and a 2x3 pin connector to support a single rotor hot swap fan assembly typically used in 2U system configurations. Concurrent use of both fan connector types for any given system fan pair is not supported.





Figure 32. Server board 2U system fan connector locations

Table 46. System fan connector pin-out

Pin	Signal Name	Pin	Signal Name
1	GND	4	FAN PWM
2	P12V FAN	5	SYS FAN PRSNT
3	FAN TACH	6	LED FAN FAULT

All six system fans are mounted within a single fan assembly module, which can be removed for cable routing or to service other components within the chassis. Each system fan includes support for the following:

- Each individual fan is hot-swappable
- Each fan is blind mated to a matching 6-pin connector located on the server board
- Each fan is designed for tool-less insertion and extraction from the fan assembly
- Each fan has a tachometer signal that allows the integrated BMC to monitor its status.
- Fan speed for each fan is controlled by integrated platform management. As system thermals fluctuate high and low, the integrated BMC firmware increases and decreases the speeds to specific fans within the fan assembly to regulate system thermals.
- On top of each fan is an integrated fault LED. Platform management illuminates the fault LED for the failed fan.






Figure 33. System fan assembly

4.4 Power Supply Module Fans

Each installed power supply module includes embedded (non-removable) 40-mm fans. They are responsible for airflow through the power supply module. These fans are managed by the fan control system. Should a fan fail, the power supply shuts down.

4.5 FRUSDR Utility

The purpose of the embedded platform management and fan control systems is to monitor and control various system features, and to maintain an efficient operating environment. Platform management is also used to communicate system health to supported platform management software and support mechanisms. The FRUSDR utility is used to program the server board with platform specific environmental limits, configuration data, and the appropriate sensor data records (SDRs) for use by these management features.

The FRUSDR utility must be run as part of the initial platform integration process before it is deployed into a live operating environment. Once the initial FRU and SDR data is loaded on to the system, all subsequent system configuration changes automatically update SDR data using the BMC auto configuration feature, without having to run the FRUSDR utility again. However, to ensure the latest sensor data is installed, the SDR data should be updated to the latest available as part of a planned system software update.

The FRUSDR utility for the given server platform can be downloaded as part of the System Update Package (SUP) or One-boot Firmware Update (OFU) package from:

<http://downloadcenter.intel.com>

Note: The embedded platform management system may not operate as expected if the platform is not updated with accurate system configuration data. The FRUSDR utility must be run with the system fully configured during the initial system integration process for accurate system monitoring and event reporting.

5. Intel® Xeon Phi™ Coprocessor and Non-Intel GPGPU Add-in Card Support

Some (not all) system configurations of the Intel® Server System R2000WF product family have thermal and mechanical support for Intel® Xeon Phi™ coprocessor cards and non-Intel GPGPU add-in cards. The card and system configurations that are supported depend on power and thermal limits of the system and power and thermal requirements of the add-in card. This chapter provides an overview of the system configuration criteria needed to support these high power cards.

Disclaimer: Intel incorporates necessary features and conducts appropriate thermal testing of its server systems to ensure the thermal requirements of Intel® Xeon Phi™ coprocessor cards can be supported when the documented support criteria are met.

Intel can make support recommendations for non-Intel GPGPU cards listed on its Tested Hardware and Support List, based on the same support criteria. However, Intel cannot guarantee its server systems can meet the thermal requirements of these non-Intel high power add-in cards. System integrators should verify non-Intel GPGPU add-in card air flow and thermal requirements from vendor specifications before integrating them into the system. It is the responsibility of the system integrator to conduct appropriate system testing on server systems configured with a non-Intel GPGPU card to ensure reliable operation.

For non-Intel GPGPU cards listed on Intel's tested hardware and OS list, Intel will provide best internal support efforts to duplicate and understand reported issues. However, Intel cannot guarantee an issue will be root caused nor that a fix can be made available to correct it.

5.1 Support Criteria

Due to air flow and thermal limitations, Intel® Xeon Phi™ coprocessor cards and non-Intel GPGPU cards can only be supported in the following base system models: R2308WFxxxx, R2208WFxxxx, and R2208WFxxxx + 8 Drive Accessory Kit (16x2.5" drive configuration).

See Appendix D at the end of this document for additional thermal configuration information.

Each Intel Xeon Phi coprocessor card and non-Intel GPGPU card must be installed into the top slot of a two-slot PCIe* riser card. With a two-slot riser card installed, the embedded fan speed control operates system fans at a higher speed. Three-slot riser cards (as shipped in the standard platform configuration) cannot be used due to airflow and thermal limitations as embedded fan speed control does not drive system fans to higher speeds needed to support the thermal requirements of these cards.




Figure 34. 2-slot PCIe* riser card

Table 47. 2-slot PCIe* riser card description and routing

Slot #	Description	PCIe* Lane Routing Riser Slot #1	PCIe* Lane Routing Riser Slot #2
Slot-1 (top)	PCIe x16 lanes, x16 slot	CPU1	CPU2
Slot-2 (bottom)	PCIe x8 lanes, x8 slot	CPU2	CPU2




Figure 35. Auxiliary 12 V power cable (iPC – AXXGPGPUCABLE)

Riser slots 1 and 2 are capable of supporting up to 100 W maximum power each. Available power to each add-in card slot on a given riser card is defined by the PCIe* specification. Add-in cards with power requirements above 75 W require the additional power to be drawn from either of two 2x2 pin “OPT_12V_PWR” connectors on the server board. A 12 V aux power cable is routed from the “OPT_12V_PWR” on the server board to a six- or eight-pin 12 V AUX power connector located on the add-in cards.

Note: Both the two-slot PCIe riser card and the 12 V auxiliary power cable are included in accessory kit A2UL16RISER2.

Configuration Note: An add-in card that requires additional 12 V power and is installed in Riser Slot #1 must have its AUX 12 V power connector located on the back edge of the card to avoid interference with the chassis side wall.

High power add-in cards with passive cooling solutions (no fan on the heat sink) can only be supported in systems configured with the high air flow air duct included in accessory kits AWFCOPRODUCTAD and AWFCOPRODUCTBKT.



Figure 36. High airflow air duct included in accessory kit AWFCOPRODUCTAD

Air Duct Note: The high air flow air duct included in accessory kit AWFCOPRODUCTAD includes two plastic knock-out panels to allow for additional air flow over a specific riser card area. A knock-out panel is only removed when installing a high power add-in card that is a double wide card with a passive heat sink

solution (heat sink with no fan attached). If the add-in card does not meet these criteria, the given knock-out panel should not be removed.

When installing a double wide add-in card with a passive heat sink solution, remove only the plastic knock-out from the side on which the add-in card is installed. Only remove both plastic knock outs when installing two cards that meet the criteria defined above (one per riser card). Once a knock-out panel is removed, it cannot be re-installed.

High power add-in cards with active cooling (heat sink with mounted fan), are only supported in system configurations installed with the standard shipping air duct (default). High power add-in cards with active cooling (heat sink with mounted fan) are not supported in systems configured with the high air flow air duct included in accessory kit AWFCOPRODUCTAD.

A system configured with the contents of accessory kits AWFCOPRODUCTAD and AWFCOPRODUCTBKT and an Intel® Xeon Phi™ card can also support the following additional PCIe* add-in card configurations.

- PCIe add-in cards with minimum 100 LFM (0.5 m/s) or 200 LFM (1m/s) air flow requirement can be installed in any available add-in card slot in Riser Card #1, Riser Card #2, and Riser Card #3 (if installed)
- PCIe add-in cards with a minimum 300 LFM (1.5 m/s) air flow requirement can be installed in the bottom add-in card slot of Riser Card #1 and Riser Card #2.
- PCIe add-in cards with air flow requirement greater than 300 LFM cannot be supported.

Systems configured with an Intel Xeon Phi card and/or non-Intel GPGPU card must have a shipping bracket installed before the system is exposed to any level of shock or vibration or is transported to the end user location. Failure to install the shipping bracket has the potential to cause serious damage to various components within the system. The shipping bracket is only available as part of accessory kit AWFCOPRODUCTBKT.



Figure 37. Shipping bracket




Figure 38. Shipping bracket placement

The shipping bracket can be used with both the standard air duct and the high airflow air duct included in accessory kit AWFCOPRODUCTAD.

With no shipping bracket installed, the air duct has support for the addition of up to two SSD storage devices mounted to the air duct via fastener screws. Mounting standard hard disk drives to the air duct is not a supported configuration. Mounting hard disk drives to the CPU air duct may cause data loss, performance degradation, and premature drive failures.

High power add-in cards that require up to 300 W can be supported. A total system power budget should be calculated to determine if the power draw of the desired system configuration meets the power limits of the installed power supplies. A power budget tool for the specified system can be downloaded from <http://www.intel.com/support>.

The latest posted system software updates must be installed on the system to ensure proper fan speed control is enabled. The latest system update package can be downloaded from <http://downloadcenter.intel.com>.

5.2 Intel® Xeon Phi™ Coprocessor Card – System Configuration Requirements

For optimum system performance and proper fan speed control, ensure the following BIOS setup utility options are set:

- MMIO over 4 GB = Enabled
- MMIO Size = 256 GB

Intel highly recommends that the latest available system software be installed. The latest system update package can be downloaded from <http://downloadcenter.intel.com>.

6. System Storage and Peripheral Drive Bay Overview

The Intel® Server System R2000WF product family has support for a variety of different storage options including:

- Support for 8, 16, or 24 front mount hot swap 2.5" SAS/SATA/NVMe* drives,
- Support for 8 front mount hot swap 3.5" SAS/SATA drives
- Support for 12 front mount hot swap 3.5" drives – 12 SATA/SAS or up to 2 NVMe + SATA /SAS
- Accessory kit option to support two rear mount 2.5" hot swap SSDs
- Up to two Internal fixed mount SSDs in all system models
- Up to two server board mounted M.2 PCIe* / SATA SSDs

Support for different storage and peripheral options vary depending on the system model and/or available accessory options installed. This section provides an overview of each available option.

6.1 Front Mount – Drive Bay Support

The 2U product family provides options to support 8, 16 and 24 front mounted 2.5" devices. System configurations of 8 or 16 drives also include front panel I/O support.




Figure 39. 2.5" front mount drive bay, 8 drive (R2208WFxxx) and 16 drive configuration option



Figure 40. 2.5" front mount drive bay, 24 drive (R2224WFxxx) configuration

The 2U product family also includes options to support 8 and 12 front mounted 3.5" drives. System configured with eight devices also include front panel I/O.

Note: 3.5" drive trays in the 8 and 12 drive system configurations are capable of supporting 3.5" drives or 2.5" SSDs with the supplied device bracket.




Figure 41. 3.5" drive bay, 8 drive configuration



Figure 42. 3.5" drive bay, 12 drive configuration

Note: Drive numbering in the system illustrations is for general reference only. Actual drive numbering is dependent on SAS / SATA controller configuration and how they are cabled to the backplane. On the backside of each installed backplane is a multi-port mini-SAS HD data connector for each set of four SATA/SAS drives. Backplanes that support PCIe NVMe drives will also include a single PCIe OCuLink connector for each NVMe drive supported on the backplane.

6.2 Hot Swap Drive Carriers

Each SAS/SATA/NVMe* drive that interfaces with a backplane is mounted to a tool-less hot swap drive carrier.



Drive carriers include a latching mechanism used to assist with drive extraction and drive insertion.



Figure 43. Drive carrier removal

Note: To ensure proper system air flow requirements, all front drive bays must be populated with a drive carrier. Drive carriers must be installed with either a drive or supplied drive blank.

There are drive carriers to support 2.5" drives and 3.5" drives. Drive blanks used with the 3.5" drive carrier can also be used to mount a 2.5" SSD.




Figure 44. 2.5" SSD mounted to 3.5" drive carrier

Note: Due to degraded performance and reliability concerns, the use of the 3.5" drive blank as a 2.5" drive bracket is intended to support SSD type storage devices only. Installing a 2.5" hard disk drive into the 3.5" drive blank cannot be supported.

Each drive carrier includes separate LED indicators for drive activity and drive status. Light pipes integrated into the drive carrier assembly direct light emitted from LEDs mounted next to each drive connector on the backplane to the drive carrier faceplate, making them visible from the front of the system.




Figure 45. Drive tray LED identification

Table 48. Drive status LED states

	LED State	Drive Status
Amber	Off	No access and no fault
	Solid on	Hard drive fault has occurred
	1 Hz blinking	RAID rebuild in progress
	2 Hz blinking	Locate (identify)

Table 49. Drive activity LED states

	Condition	Drive Type	LED Behavior
Green	Power on with no drive activity	SAS/NVMe*	LED stays on
		SATA	LED stays off
	Power on with drive activity	SAS/NVMe*	LED blinks off when processing a command
		SATA	LED blinks on when processing a command
	Power on and drive spun down	SAS/NVMe*	LED stays off
		SATA	LED stays off
	Power on and drive spinning up	SAS/NVMe*	LED blinks
		SATA	LED stays off

Table 50. PCIe* SSD drive status LED states

	LED State	Drive Status
Amber	Off	No fault, OK
	4 Hz blinking	Locate (identify)
	Solid on	Fault/fail
	1 Hz blinking	Rebuild

Note: The drive activity LED is driven by signals coming from the drive itself. Drive vendors may choose to operate the activity LED different from what is described in the table above. Should the activity LED on a given drive type behave differently than what is described, customers should reference the drive vendor specifications for the specific drive model to determine the expected drive activity LED operation.

6.3 Peripheral Power Sources

Power for all backplanes and peripheral storage devices is drawn from two power connectors labeled as “HSBP_PWR” and “Peripheral PWR” on the server board as illustrated below.




Figure 46. Server board peripheral power connectors

- **HSBP power** – The hot swap backplane power connector provides power for all front mounted backplane options. Appropriate power cables to support any given backplane option are included with the given system model or given backplane accessory kit. For the HSBP power connector pin-out, see Table 34.
- **Peripheral power** – The peripheral power connector is used to provide power to various add-in options including: internal mounted SSD drives, and the 2 x 2.5” rear mount hot swap backplane accessory kit. Appropriate power cables for supported peripheral options are included with the given system model or given accessory kit option.

6.4 Hot Swap Backplane Support

The 2U system has support for several backplane options.

For 2.5” drives, available backplane options include:

- 8 x 2.5” drive combo backplane with support for SAS/SATA/NVMe*
- 8 x 2.5” drive dual port SAS backplane
- 2 x 2.5” drive rear mount SATA SSD backplane

For 3.5” drives, available options include:

- 8 x 3.5” drive backplane with support for SAS/SATA
- 12 x 3.5” drive backplane with support SAS/SATA and up to two NVMe drives

The 3.5" drive backplanes are fixed mounted to the back of the drive bay within the chassis (as shown in the following diagram).




Figure 47. 3.5" backplane placement

All 2.5" drive backplane options are mounted to a drive bay module that either slide into open bays within the front of the chassis or is mounted to a fixed mount bay in the rear of the chassis as shown below.



Figure 48. 2.5" drive bay module placement options

Backplanes include the following features:

- 12 Gb SAS and 6 Gb SAS/SATA or slower support
- Drive interface connectors
 - o 29-pin SFF-8680 – SATA/SAS only
 - o 68-pin SFF-8639 – SATA/SAS/NVMe – NVMe compatible backplanes only
- Hot swap drive support
- Cable connectors
 - o SFF-8643 Mini-SAS HD – 2.5" and 3.5" backplanes – 12 Gb SAS capable
 - o OCuLink PCIe* interface – NVMe compatible backplanes only
 - o 1x5-pin connector – I2C interface for device status communication to the BMC over slave SMBus
 - o 2x2-pin connector – power
- SGPIO SFF-8485 interface embedded within the sideband of the mini-SAS HD connectors
- HSBP microcontroller – Cypress* CY8C22545-24AXI PSoC* Programmable System-on-Chip device

- LEDs to indicate drive activity and status for each attached device
- Device presence detect inputs to the microcontroller
- 5 V VR for devices
- 3.3 V VR for microcontroller
- In-application microcontroller firmware updateable over the I2C interface
- FRU EEPROM support
- Temperature sensor through the use of a TMP75 (or equivalent) thermistor implementation with the microcontroller

6.4.1 SGPIO Functionality

Backplanes include support for a SFF-8485 compliant SGPIO interface used to activate the status LED. This interface is also monitored by the microcontroller for generating FAULT, IDENTIFY, and REBUILD registers that in turn are monitored by the server board BMC for generating corresponding SEL events.

6.4.2 I2C Functionality

The microcontroller has a master/slave I2C connection to the server board BMC. The microcontroller is not an IPMB compliant device. The BMC generates SEL events by monitoring registers on the HSBP microcontroller for DRIVE PRESENCE, FAULT, and RAID REBUILD in progress.

6.5 8 x 2.5" Drive SATA/SAS/NVMe* Combo Backplane

This section applies to the 8 x 2.5" drive SAS/SATA/NVMe combo backplane (iPC – F2U8X25S3PHS).

All system SKUs capable of supporting 2.5" drives will include one or more eight drive backplanes capable of supporting 12 Gb/sec SAS, 6 Gb/sec SATA drives, and PCIe NVMe* drives.

The front side of the backplane includes 68-pin SFF-8639 drive interface connectors, each capable of supporting SAS, SATA, or NVMe drives. The connectors are numbered 0 through 7.




Figure 49. 8 x 2.5" SAS/SATA/NVMe* hot swap backplane, front side

The backside of the backplane includes two multi-port mini-SAS HD connectors labeled SAS/SATA_0-3 and SAS/SATA_4-7, and eight PCIe OcuLink connectors, each labeled PCIe SSD#, where # = 0 -7, one connector for each installed NVMe drive.




Figure 50. 8 x 2.5" SAS/SATA/NVMe* hot swap backplane, back side

I2C cable connectors – The backplane includes two 1x5 cable connectors (labeled I2C_IN and I2C_OUT) used as a management interface between the server board and the installed backplanes. In systems configured with multiple backplanes, a short jumper cable is attached between backplanes, with connector A used on the first board and connector D used on the second board, extending the SMBus to each installed backplane.

Table 51. I2C cable connector pin-out

Pin	Signal
1	SMB_3V3SB_DAT
2	GND
3	SMB_3V3SB_CLK
4	SMB_ADD0
5	SMB_ADD1




Figure 51. 8 x 2.5" dual port SAS backplane, back side

Multi-port mini-SAS HD cable connectors – The backplane includes two multi-port mini-SAS HD cable connectors (labeled PORT 0-3 and PORT 4-7), each providing SGPIO and I/O signals for up to four SAS/SATA devices installed in the hot swap drive bay. Input cables can be routed from matching connectors on the

server board (on-board SATA only), from installed add-in SAS/SATA RAID cards, or from an optionally installed SAS expander card for drive configurations of greater than eight hard drives.

Power harness connector– The backplane includes a 2x2 connector supplying power to the backplane. Power is routed to each installed backplane via a multi-connector power cable harness from the server board.

Table 52. Power harness connector pin-out

Pin	Signal
1	GND
2	GND
3	P12V
4	P12V

PCIe OCuLink* Connectors The backplane has support for up to eight (8) NVMe SFF SSDs. The backside of the backplane includes eight OCuLink* cable connectors, one for each drive connector on the front side of the backplane. Each installed NVMe drive must have PCIe signals cabled to the appropriate backplane OCuLink connector from any of the following PCIe signal sources:

- Available onboard PCIe* OCuLink connectors on the server board
- Optional PCIe 4 or 8 Port Switch Add-in Card
- Optional Tri-Mode Raid Module or Add-in Card

See section 7.4 for NVMe support information.

6.6 8 x 2.5" Drive Dual Port SAS Backplane

This section applies to the 8 x 2.5" drive dual port SAS backplane (accessory kit iPC – A2U8X25S3DPDK).

With the 8 x drive dual port backplane accessory kit installed, the server can support up to eight dual port SAS drives. The front side of the backplane includes 8 x 29-pin drive interface connectors, each capable of supporting 12 Gb SAS or 6 Gb SAS/SATA. The connectors are numbered 0 through 7. Signals for each set of four drive connectors (0-3 and 4-7), are routed to a pair of mini-SAS HD connectors on the back side of the backplane, one for the primary port and one for the secondary port.





Figure 52. 8 x 2.5" dual port SAS backplane, front side

Note: Letters in parenthesis denote references to connectors on the back side of the backplane as illustrated in Figure 51.



A – I2C-out cable connector for multi-backplane support
 B – Ports 4-7 (secondary) mini-SAS HD cable connector
 C – Ports 4-7 (primary) mini-SAS HD cable connector
 D – Ports 0-3 (secondary) mini-SAS HD cable connector

E – Ports 0-3 (primary) mini-SAS HD cable connector
 F – I2C-in cable connector – from server board or other backplane
 G – Power connector

Figure 53 . 8 x 2.5" dual port SAS backplane, Back side

Multi-port mini-SAS HD cable connectors (connectors B and C) – Primary and secondary port connectors for drives 4-7. Cables can be routed from matching connectors on the server board, from installed add-in SAS/SATA RAID cards, or from optionally installed SAS expander cards for drive configurations of greater than eight hard drives.

Multi-port mini-SAS cable connectors (connectors D and E) – Primary and secondary port connectors for drives 0-3. Cables can be routed from matching connectors on the server board, from installed add-in SAS/SATA RAID cards, or from optionally installed SAS expander cards for drive configurations of greater than eight hard drives.

Power harness connector (connector G) – The backplane includes a 2x2 connector supplying power to the backplane. Power is routed to each installed backplane via a multi-connector power cable harness from the server board.

Table 53. Power harness connector pin-out

Pin	Signal
1	GND
2	GND
3	P12V
4	P12V

I2C cable connectors (connectors A and F) – The backplane includes two 1x5 cable connectors used as a management interface between the server board and the installed backplanes. In systems configured with multiple backplanes, a short jumper cable is attached between backplanes, with connector A used on the first board and connector F used on the second board, extending the SMBus to each installed backplane.

Table 54. I2C cable connector pin-out

Pin	Signal
1	SMB_3V3SB_DAT
2	GND
3	SMB_3V3SB_CLK
4	SMB_ADD0
5	SMB_ADD1

6.7 8 x 3.5" Drive Hot-Swap Backplane Overview

This section applies to the 8 x 3.5" drive hot-swap backplane (iPC – F2U8X35S3HSBP).


2U server systems that support the 8 drive backplane have support for the following drive form factors:

- 3.5" hard drives
- 2.5" SSDs using the included drive blanks as a drive bracket

Note: Due to degraded performance and reliability concerns, the use of the 3.5" drive blank as a 2.5" drive bracket is intended to support SSD type storage devices only. Installing a 2.5" hard disk drive into the 3.5" drive blank cannot be supported.

This backplane has support for up to 6Gb/s SATA drives or up to 12 Gb/s SAS drives

Mounted on the front side of the backplane are eight 29-pin (SFF-8482) drive connectors. On the backside of the backplane are connectors for power, I2C management, and SAS/SATA data.



Note: Letters in parenthesis denote references to connectors on the back side of the backplane.




Figure 54. 8 x 3.5" HSBP Connector Identification

Power harness connector (connector A) – The backplane includes a 2x2 connector supplying power to the backplane. Power is routed to the backplane via a power cable harness from the server board

Multi-port mini-SAS HD cable connectors (connectors B and C) – The 8 x drive backplane includes two multi-port mini-SAS cable connectors, each providing SGPIO and I/O signals for four SAS/SATA hard drives on the backplane. Cables can be routed from matching connectors on the server board, from add-in SAS/SATA RAID cards, or from an optionally installed SAS expander card. Each mini-SAS HD connector includes a silk-screen identifying which drives the connector supports: drives 0-3 and drives 4-7.

I2C cable connector (connector D) – The backplane includes a 1x5 cable connector used as a management interface to the server board

6.8 12 x 3.5" Drive Hot-Swap Backplane Overview

This section applies to the 12 x 3.5" drive hot-swap backplane (iPC – F2U12X35S3PH).

2U server systems that support the 12 drive backplane have support for the following drive form factors:


- 3.5" hard drives
- 2.5" SSDs using the included drive blanks as a drive bracket

Note: Due to degraded performance and reliability concerns, the use of the 3.5" drive blank as a 2.5" drive bracket is intended to support SSD type storage devices only. Installing a 2.5" hard disk drive into the 3.5" drive blank cannot be supported.

Drive technologies supported include:

- 6 Gb/s SATA
- 12 Gb/s SAS
- PCIe* NVMe (up to 2 drives)

Mounted on the front side of the backplane are ten 29-pin (SFF-8482) drive connectors supporting SAS/SATA drives only, and two 68-pin SFF-8639 drive connectors supporting SAS/SATA/NVMe drives.



2UWFP086

Figure 55. 12 x 3.5" hot swap backplane, front side

On the backside of the backplane are connectors for power, I2C management, SAS/SATA data, and PCIe NVMe.

Note: Letters in parenthesis denote references to connectors on the back side of the backplane as illustrated in Figure 56.




Figure 56. 12 x 3.5" HSBP connector identification

Multi-port mini-SAS HD cable connectors (connectors A, B, and C) – The 12 x drive backplane includes three multi-port mini-SAS cable connectors, each providing SGPIO and I/O signals for four SAS/SATA drives on the backplane. Cables can be routed from matching connectors on the server board, from add-in SAS/SATA RAID cards, or from an optionally installed SAS expander card. Each mini-SAS HD connector includes a silk-screen identifying which drives the connector supports: drives 0-3, drives 4-7, and drives 8-11.

Power harness connector (connector D) – The backplane includes a 2x2 connector supplying power to the backplane. Power is routed to the backplane via a power cable harness from the server board

I2C cable connectors (connector E) – The backplane includes a 1x5 cable connector used as a management interface to the server board.

PCIe SSD connector (connectors F and G) – The backplane includes two PCIe OcuLink cable connectors (labeled PCIe SSD0 and PCIe SSD1), each providing support for one NVMe drive.

6.9 2 x 2.5" Hot Swap Drive Bay Accessory Kit

This section applies to the 2 x 2.5" hot swap drive bay accessory kit (iPC – A2UREARHSDK).

The 2U product family provides the option to support two 6 Gb/sec hot swap SATA SSDs installed to a modular drive bay mounted in the back of the system.




Figure 57. 2 x 2.5" rear mount backplane kit placement

Supported SATA SSDs must not exceed the following power and thermal limits:

- One or two SATA SSDs supporting up to 4 W per device with a case temperature rating of 70 °C
- One or two SATA SSDs supporting up to 1.5 W per device with a case temperature rating of 60 °C

Note: The maximum supported SSD power and thermal limits documented above were derived based on thermal testing using a maximum system configuration with fan redundancy support operating at ambient input air temp of 35 °C. The test system was based on the system models utilizing a maximum number of front and rear drive storage devices – 12 x 3.5" or 24 x 2.5" in the front and 2 x 2.5" in the back.

Because of thermal limits in this area of the chassis, the rear hot swap drive bay option cannot support hard disk drives in any system configuration.

By lowering the maximum supported ambient air temperature to 27 °C, and limiting the system configuration to support 8 or 16 devices up front and no storage devices configured on the air duct, supported SATA SSDs must not exceed the following power and thermal limits:

- One or two SATA SSDs supporting up to 6.4 W per device with a case temperature rating of 70 °C
- One or two SATA SSDs supporting up to 3.6 W per device with a case temperature rating of 60 °C

The backplane includes several connectors and a jumper block, as defined in the following illustrations.




Figure 58. 2 x 2.5" hot swap backplane

- **Activity/status LEDs** – The backplane includes LEDs for both activity and status for each drive.
- **SATA drive connectors** – The backplane includes two hot-swap capable SATA drive connectors
- **7-pin SATA cable connectors** – The backplane includes two 7-pin connectors capable of supporting a SATA interface. The following table defines the pin-out for these connectors.

Table 55. 7-pin SATA cable connector pin-out

Pin	Signal
1	GND
2	SAS_SATA_TX_P
3	SAS_SATA_TX_N
4	GND
5	SAS_SATA_RX_N
6	SAS_SATA_RX_P
7	GND

- **Power connector** – The backplane includes a 2x3 pin power connector. The following table defines the pin-out for this connector:

Table 56. Power connector pin-out

Pin	Signal
1	P5V
2	P5V
3	GND
4	P12V
5	P3V3
6	GND

- **I2C connector** – The backplane includes a 1x5 pin I2C connector. This connector is cabled to a matching HSBP I2C connector on the server board and is used as a communication path to the onboard BMC. The following table defines the pin out for this connector.

Table 57. I2C connector pin-out

Pin	Signal
1	SMB_3V3SB_DAT
2	GND
3	SMB_3V3SB_CLK
4	SMB_ADD0
5	SMB_ADD1

- **SGPIO connector** – The backplane includes a 1x5 pin SGPIO connector. When the backplane is cabled to the on-board SATA ports, this connector is cabled to a matching SGPIO connector on the server board and provides support for drive activity and fault LEDs. The following table defines the pin out for this connector.

Table 58. SGPIO connector pin-out

Pin	Signal
1	SGPIO_CLOCK
2	SGPIO_LOAD
3	GND
4	SGPIO_DATAOUT
5	SGPIO_DATAIN

6.10 Internal Fixed Mount Solid State Drive (SSD) Support

The system can support up to two internal fixed mount 2.5" solid state drives (SSDs). Each drive is attached to the top side of the air duct as illustrated below.




Figure 59. 2.5" Solid state drive (SSD) mounting option

SSDs mounted to the air duct must not exceed the following power and thermal limits.

- One or two SATA SSDs supporting up to 6.2 W per device with a case temperature rating of 70 °C
- One or two SATA SSDs supporting up to 2.3 W per device with a case temperature rating of 60 °C

Note: Maximum supported SSD power and thermal limits were derived based on thermal testing using a maximum system configuration with fan redundancy support using a maximum ambient temperature of 35 °C. The test system was based on the system models utilizing a maximum number of front and rear drive storage devices – 12 x 3.5" or 24 x 2.5" in the front and 2 x 2.5" in the back. Higher SSD power and thermal limits maybe possible with system configurations that support 8 or 16 devices up front, no storage devices configured in back, and lowering the maximum operating ambient temperature.

Note: Mounting standard hard disk drives to the CPU air duct is not a supported configuration. Mounting hard disk drives to the CPU air duct may cause data loss, performance degradation, and premature drive failures.

A 2x3 pin power connector on the server board labeled "Peripheral PWR" is designed to provide power to both SSDs or a single SSD and an optical drive. Using the supplied power harness, one power lead is routed to the optical drive bay, and two longer power leads are routed to the SSDs. SATA signals for one SSD and one optical drive or two SSDs are cabled from the two white 7-pin AHCI SATA connectors on the server board.

7. Storage Controller Options Overview

The Intel® Server System R2000WF product family has support for a variety of storage controller and storage device options including:

- Onboard SATA support
- Embedded software RAID support
- M.2 SSD support
- NVMe* SFF SSD support
- Intel® Integrated RAID Modules support

7.1 Onboard SATA Support

The server board utilizes two chipset embedded AHCI SATA controllers, identified as **SATA** and **sSATA**, providing for up to ten 6 Gb/sec Serial ATA (SATA) ports.

The AHCI **sSATA** controller provides support for up to 4 SATA ports on the server board:

- Two ports accessed via two white single port 7-pin connectors labeled “sSATA-4” and “sSATA-5” on the server board
- Two ports (sSATA 1 and sSATA 2) via two M.2 SSD connectors

The AHCI **SATA** controller provides support for up to 8 SATA ports on the server board: (S2600WFT & S2600WF0 boards only)

- Four ports from the Mini-SAS HD (SFF-8643) connector labeled “SATA Ports 0-3” on the server board
- Four ports from the Mini-SAS HD (SFF-8643) connector labeled “SATA Ports 4-7” on the server board.

Note: The onboard SATA controllers are not compatible with and cannot be used with SAS expander cards.




Figure 60. Onboard storage support features

Table 59. SATA and sSATA controller feature support

Feature	Description	AHCI Mode	RAID Mode Intel® RSTe	RAID Mode Intel® ESRT2
Native Command Queuing (NCQ)	Allows the device to reorder commands for more efficient data transfers	Supported	Supported	
Auto Activate for DMA	Collapses a DMA Setup then DMA Activate sequence into a DMA Setup only	Supported	Supported	
Hot Plug Support	Allows for device detection without power being applied and ability to connect and disconnect devices without prior notification to the system	Supported	Supported	
Asynchronous Signal Recovery	Provides a recovery from a loss of signal or establishing communication after hot plug	Supported	Supported	
6 Gb/s Transfer Rate	Capable of data transfers up to 6 Gb/s	Supported	Supported	Supported
ATAPI Asynchronous Notification	A mechanism for a device to send a notification to the host that the device requires attention	Supported	Supported	
Host & Link Initiated Power Management	Capability for the host controller or device to request Partial and Slumber interface power states	Supported	Supported	
Staggered Spin-Up	Enables the host the ability to spin up hard drives sequentially to prevent power load problems on boot	Supported	Supported	Supported
Command Completion Coalescing	Reduces interrupt and completion overhead by allowing a specified number of commands to complete and then generating an interrupt to process the commands	Supported	N/A	

The SATA controller and the sSATA controller can be independently enabled and disabled and configured through the BIOS setup utility under the Mass Storage Controller Configuration menu screen. The following table identifies supported setup options.

Table 60. SATA and sSATA controller BIOS setup utility options

SATA Controller	sSATA Controller	Supported
AHCI	AHCI	Yes
AHCI	Disabled	Yes
AHCI	Intel® RSTe	Yes
AHCI	Intel® Embedded Server RAID Technology 2	Microsoft Windows* only
Disabled	AHCI	Yes
Disabled	Disabled	Yes
Disabled	Intel® RSTe	Yes
Disabled	Intel® Embedded Server RAID Technology 2	Yes
Intel® RSTe	AHCI	Yes
Intel® RSTe	Disabled	Yes
Intel® RSTe	Intel® RSTe	Yes
Intel® RSTe	Intel® Embedded Server RAID Technology 2	No
Intel® Embedded Server RAID Technology 2	AHCI	Microsoft Windows only
Intel® Embedded Server RAID Technology 2	Disabled	Yes
Intel® Embedded Server RAID Technology 2	Intel® RSTe	No
Intel® Embedded Server RAID Technology 2	Intel® Embedded Server RAID Technology 2	Yes

7.1.1 Staggered Disk Spin-Up

Because of the high density of disk drives that can be attached to the Intel® C620 onboard AHCI SATA controller and the sSATA controller, the combined startup power demand surge for all drives at once can be much higher than the normal running power requirements and could require a much larger power supply for startup than for normal operations.

In order to mitigate this and lessen the peak power demand during system startup, both the AHCI SATA Controller and the sSATA Controller implement a Staggered Spin-Up capability for the attached drives. This means that the drives are started up separately, with a certain delay between disk drives starting.

For the onboard SATA controller, staggered spin-up is an option – **AHCI HDD Staggered Spin-Up** – in the Mass Storage Controller Configuration screen found in the BIOS setup utility.

7.2 Onboard SATA Software RAID Support

The server board includes support for two embedded SATA RAID options:

- Intel® Rapid Storage Technology enterprise (Intel® RSTe) 5.0
- Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60

By default, onboard RAID options are disabled in BIOS setup. To enable onboard RAID support, access the BIOS setup utility during POST. The onboard RAID options can be found under the **sSATA Controller** or **SATA Controller** options under the following BIOS setup menu:

Advanced -> Mass Storage Controller Configuration



Figure 61. BIOS setup Mass Storage Controller Configuration screen

Note: RAID partitions created using either RSTe or ESRT2 cannot span across the two embedded SATA controllers. Only drives attached to a common SATA controller can be included in a RAID partition.

7.2.1 Intel® Rapid Storage Technology Enterprise (Intel® RSTe) 5.0

Intel® Rapid Storage Technology enterprise (Intel® RSTe) offers several options for RAID to meet the needs of the end user. AHCI support provides higher performance and alleviates disk bottlenecks by taking advantage of the independent DMA engines that each SATA port offers in the chipset. Supported RAID levels include 0, 1, 5, and 10.

- **RAID 0** – Uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.
- **RAID 1** – Uses mirroring so that data written to one disk drive simultaneously writes to another disk drive. This is good for small databases or other applications that require small capacity but complete data redundancy.
- **RAID 5** – Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.
- **RAID 10** – A combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy but uses a larger number of spans.

By using Intel RSTe, there is no loss of PCI resources (request/grant pair) or add-in card slot. Intel RSTe functionality requires the following:

- The embedded RAID option must be enabled in BIOS setup.
- Intel RSTe option must be selected in BIOS setup.
- Intel RSTe drivers must be loaded for the installed operating system.
- At least two SATA drives needed to support RAID levels 0 or 1.
- At least three SATA drives needed to support RAID level 5.
- At least four SATA drives needed to support RAID level 10.
- NVMe SSDs and SATA drives must not be mixed within a single RAID volume

With Intel RSTe software RAID enabled, the following features are made available:

- A boot-time, pre-operating-system environment, text-mode user interface that allows the user to manage the RAID configuration on the system. Its feature set is kept simple to keep size to a minimum, but allows the user to create and delete RAID volumes and select recovery options when problems occur. The user interface can be accessed by pressing **<CTRL-I>** during system POST.
- Boot support when using a RAID volume as a boot disk. It does this by providing Int13 services when a RAID volume needs to be accessed by MS-DOS applications (such as NT loader (NTLDR)) and by exporting the RAID volumes to the system BIOS for selection in the boot order.
- At each boot-up, a status of the RAID volumes provided to the user.

7.2.2 Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 for SATA

Intel® Embedded Server RAID Technology 2 (powered by LSI*) is a driver-based RAID solution for SATA that is compatible with previous generation Intel® server RAID solutions. Intel Embedded Server RAID Technology 2 provides RAID levels 0, 1, and 10, with an optional RAID 5 capability depending on whether a RAID upgrade key is installed.

Note: The embedded Intel Embedded Server RAID Technology 2 option has no RAID support for PCIe* NVMe* SSDs.

Intel Embedded Server RAID Technology 2 is based on LSI* MegaRAID software stack and utilizes the system memory and CPU.

Supported RAID levels include:

- **RAID 0:** Uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.

- **RAID 1:** Uses mirroring so that data written to one disk drive simultaneously writes to another disk drive. This is good for small databases or other applications that require small capacity but complete data redundancy
- **RAID 10:** A combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. It provides high data throughput and complete data redundancy but uses a larger number of spans.
- Optional support for RAID Level 5
 - o Enabled with the addition of an optionally installed ESRT2 SATA RAID 5 Upgrade Key (**IPN - RKSATA4R5**)




Figure 62. ESRT2 SATA RAID-5 upgrade key

- o **RAID 5:** Uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.

Intel Embedded Server RAID Technology 2 on this server board supports a maximum of six drives which is the maximum onboard SATA port support.

The binary driver includes partial source files. The driver is fully open source using an MDRAID layer in Linux*.

Note: RAID configurations cannot span across the two embedded AHCI SATA controllers.

7.3 M.2 SSD Support

The Intel® Server Board S2600WF product family includes two M.2 SSD connectors labeled “M2_x4PCIE/sSATA_1” and “M2_x2PCIE/sSATA_2” on the server board as shown below.




Figure 63. M.2 module connector locations

Each M.2 connector can support PCIe* or SATA modules that conform to a 2280 (80 mm) form factor.

PCIe bus lanes for each connector are routed from the Intel chipset and can be supported in single processor configurations.

The M.2 connector to the left of Riser Slot #1 is supported by **PCIe x4** bus lanes and **sSATA-1** from the chipset embedded sSATA controller. The M.2 connector to the right of Riser Slot #1 is supported by **PCIe x2** bus lanes and **sSATA-2** from the chipset embedded sSATA controller.

7.3.1 Embedded RAID Support

RAID support from embedded RAID options for server board mounted M.2 SSDs is defined as follows:

- Neither Intel® ESRT2 nor Intel® RSTe have RAID support for PCIe* M.2 SSDs when installed to the M.2 connectors on the server board.
 - o **Note:** NVMe RAID support using Intel® RSTe VROC requires that the PCIe bus lanes be routed directly from the CPU. On this server board, the PCIe bus lanes routed to the on-board M.2 connectors are routed from the Intel chipset (PCH).
 - o **Note:** The Intel® ESRT2 onboard RAID option does not support PCIe devices.
- Both Intel® ESRT2 and Intel® RSTe provide RAID support for SATA devices
- Neither embedded RAID option supports mixing of M.2 SATA SSDs and SATA hard drives within a single RAID volume

Note: Storage devices used to create a single RAID volume created using either RSTe or ESRT2, cannot span across the two embedded SATA controllers nor is mixing both SATA and NVMe devices within a single RAID volume supported.

- Open Source Compliance = Binary Driver (includes Partial Source files) or Open Source using MDRAID layer in Linux*

7.4 PCIe NVMe* Drive Support

Several system models within the Intel® Server System R2000WF product family have support for PCIe NVMe drives installed into the front drive bays. For system models supporting 12x3.5" drives, the front drive bay can support up to two (2) NVMe drives. For system models supporting from one up to three 8x2.5" combo backplanes, the front drive bay can support from one (1) up to eight (8) NVMe drives per installed backplane.

Each installed NVMe drive must have its own PCIe interface as cabled to the backplane from one or more of the following PCIe signal sources:

- Onboard PCIe OCUlink connectors on the server board
- 4-port PCIe 3.0 x8 Switch add-in card (accessory kit iPC – AXXP3SWX08040)
- 8-port PCIe 3.0 x8 Switch add-in card (accessory kit iPC – AXXP3SWX08080)
- Intel® Integrated RAID Modules with tri-mode support

7.4.1 Onboard PCIe* OCUlink Connectors

The server board provides four (4) onboard PCIe* OCUlink* connectors, each supporting a x4 PCIe signaling interface. When cabled to the backplane, each connector provides the PCIe interface to a single installed NVMe drive on the backplane.

PCIe* signals for onboard OCUlink connectors "PCIe_SSD0" and "PCIe_SSD1" are routed directly from CPU_1 and PCIe signals for OCUlink connectors "PCIe_SSD2" and "PCIe_SSD3" are directly routed from CPU_2..




Figure 64. Onboard OCUlink connectors

7.4.2 Intel® Volume Management Device (Intel® VMD) for NVMe*

Intel® Volume Management Device (Intel® VMD) is hardware logic inside the processor root complex to help manage PCIe* NVMe* SSDs. It provides robust hot plug support and status LED management. This allows servicing of storage system NVMe SSD media without fear of system crashes or hangs when ejecting or inserting NVMe SSD devices on the PCIe bus.




Figure 65. NVMe* storage bus event/error handling

Intel® VMD handles the physical management of NVMe storage devices as a standalone function but can be enhanced when Intel® VROC support options are enabled to implement RAID based storage systems. See section 7.4.3 for more information.

- Hardware is integrated inside the processor PCIe* root complex.
- Entire PCIe* trees are mapped into their own address spaces (domains).
- Each domain manages x16 PCIe* lanes.
- Can be enabled/disabled in BIOS setup at x4 lane granularity.
- Driver sets up/manages the domain (enumerate, event/error handling)
- May load an additional child device driver that is Intel VMD aware.
- Hot plug support - hot insert array of PCIe* SSDs.
- Support for PCIe* SSDs and switches only (no network interface controllers (NICs), graphics cards, etc.)
- Maximum of 128 PCIe* bus numbers per domain.
- Support for MCTP over SMBus only.
- Support for MMIO only (no port-mapped I/O).
- Does not support NTB, Quick Data Tech, Intel® Omni-Path Architecture, or SR-IOV.
- Correctable errors do not bring down the system.
- Intel® VMD only manages devices on PCIe* lanes routed directly from the processor. Intel® VMD cannot provide device management on PCI lanes routed from the chipset (PCH)
- When Intel VMD is enabled, the BIOS does not enumerate devices that are behind Intel VMD. The Intel VMD-enabled driver is responsible for enumerating these devices and exposing them to the host.
- Intel® VMD supports hot-plug PCIe* SSDs connected to switch downstream ports. Intel® VMD does not support hot-plug of the switch itself.

7.4.2.1 Enabling Intel® VMD support

In order for installed NVMe devices to utilize the Intel® VMD features of the server board, Intel VMD must be enabled on the appropriate CPU PCIe root ports in BIOS setup. By default, Intel VMD support is disabled on all CPU PCIe root ports in BIOS setup.

The following table provides the PCIe root port mapping for all on-board PCIe devices, OCuLink Connectors and Riser Card slots.

Table 61. CPU - PCIe* port routing

CPU 1		CPU 2	
PCI Ports	On-board Device	PCI Ports	On-board Device
Port DMI 3 - x4	Chipset	Port DMI 3 - x4	Riser Slot #3
Port 1A - x4	Riser Slot #1	Port 1A - x4	Riser Slot #2
Port 1B - x4	Riser Slot #1	Port 1B - x4	Riser Slot #2
Port 1C - x4	Riser Slot #1	Port 1C - x4	Riser Slot #1
Port 1D - x4	Riser Slot #1	Port 1D - x4	Riser Slot #1
Port 2A - x4	Chipset (PCH) - uplink	Port 2A - x4	Riser Slot #2
Port 2B - x4	Chipset (PCH) - uplink	Port 2B - x4	Riser Slot #2
Port 2C - x4	Chipset (PCH) - uplink	Port 2C - x4	Riser Slot #2
Port 2D - x4	Chipset (PCH) - uplink	Port 2D - x4	Riser Slot #2
Port 3A - x4	SAS Module	Port 3A - x4	OCuLink PCIe_SSD2
Port 3B - x4	SAS Module	Port 3B - x4	OCuLink PCIe_SSD3
Port 3C - x4	OCuLink PCIe_SSD0	Port 3C - x4	Riser Slot #3
Port 3D -x4	OCuLink PCIe_SSD1	Port 3D -x4	Riser Slot #3

For PCIe add-in card slot root port mapping associated with each riser slot and supported riser cards, See Table 66. Riser slot #1 PCIe* root port mapping, Table 67. Riser slot #2 PCIe* root port mapping, Table 68. Riser slot #3 PCIe* root port mapping.

7.4.3 Intel® Virtual RAID on Chip (Intel® VROC) for NVMe*

Intel® Virtual RAID on Chip (Intel® VROC) enables NVMe* boot on RAID and volume management (Intel® RSTe 5.0 + Intel® VMD)




Figure 68. Intel® VROC basic architecture overview

Intel® VROC supports the following:

- I/O processor with controller (ROC) and DRAM.
- No need for battery backup / RAID maintenance free backup unit.
- Protected write back cache – software and hardware that allows recovery from a double fault.
- Isolated storage devices from OS for error handling.
- Protected R5 data from OS crash.
- Boot from RAID volumes based on NVMe SSDs within a single Intel VMD domain.
- NVMe SSD hot plug and surprise removal on CPU PCIe* lanes.
- LED management for CPU PCIe attached storage.
- RAID / storage management using representational state transfer (RESTful) application programming interfaces (APIs).
- Graphical user interface (GUI) for Linux*.
- 4K native NVMe SSD support.

Enabling Intel VROC support requires installation of an optional upgrade key on to the server board as shown in Figure 69.

Table 62 identifies available Intel VROC upgrade key options.




Figure 69. Intel® VROC upgrade key

Table 62. Intel® VROC upgrade key options

NVMe* RAID Major Features	Standard Intel® VROC (iPC VROCSTANMOD)	Premium Intel® VROC (iPC VROCPREMMOD)
CPU attached NVMe SSD – high perf.	✓	✓
Boot on RAID volume	✓	✓
Third party vendor SSD support	✓	✓
Intel® RSTe 5.0 RAID 0/1/10	✓	✓
Intel® RSTe 5.0 RAID 5	-	✓
RAID write hole closed (RMFBU replacement)	-	✓
Hot plug/ surprise removal (2.5" SSD form factor only)	✓	✓
Enclosure LED management	✓	✓

Note: Intel® VROC Upgrade Keys referenced in

Table 62 are used for PCIe* NVMe* SSDs only. For SATA RAID support, see Section 7.2.

7.4.4 NVMe* Drive Population Rules for Intel® VROC

In order to support NVMe RAID and NVMe Management features, the optional Intel® VROC Key must be installed on to the server board. With the Intel VROC key installed, specific drive population rules exist and must be followed for proper support of the NVMe management features.

The backplane can support PCIe interfaces from the onboard OCuLink connectors on the server board, and/or optional add-in Intel Tri-Mode RAID modules and PCIe Switch cards. When cabling the PCIe interfaces from two different PCIe interface sources to the backplane, the cables from each source must be connected in defined drive sets of four (0,1,2,3) & (4,5,6,7) as shown in the following diagrams.

Note: The use of one or more OCuLink connectors on the server board to the backplane is considered a single source.




Figure 70. Backplane cabling from two PCIe sources

When cabling the backplane from two different PCIe sources, no other drive set combinations beyond those defined above are supported.

Drive population rules will differ depending on the source of the PCIe interface to the backplane. In addition, specific drive population limits exist when populating a backplane with both NVMe and SAS drive types.

The following sections define the drive population rules associated with each of the available PCIe sources to the backplane.


NOTE: When connecting the backplane to two different PCIe sources, the defined population rules for each PCIe source are applied to the drive set connected to it

7.4.4.1 Onboard PCIe OCUlink Connectors and / or Intel Tri-mode RAID module to 8 x 2.5" Combo Backplane


The following information is applicable when PCIe signal to the 8x2.5" combo backplane are cabled from the PCIe OCUlink connectors located on the server board and/or an optionally installed Intel® Tri-mode RAID Module.

- OCUlink connectors on the server board are considered a single PCIe source to the backplane, and therefore can only be connected in defined drive sets **0-3 or 4-7**
- NVMe drive management sideband signals on the backplane are routed between drive connector pairs: **(0,1) (2,3) (4,5) and (6,7)**

- In order to support NVMe drive management within a defined drive pair, an NVMe drive MUST be populated in the first drive connector of the given pair (drives **0, 2, 4, or 6**)
- Combining an NVMe drive with a SAS/SATA drive within a defined drive pair is NOT supported.
Example) With an NVMe drive installed to drive connector 0, drive connector 1 cannot be populated with a SAS/SATA drive. The same rule applies to ALL other drive pairs on the backplane.



The following illustrations identify supported and unsupported drive populations associated with any defined drive pair of the 8x2.5" combo backplane when Intel VROC is used for NVMe drive management.




Where 1st Drive = drive connectors 0, 2, 4, or 6 and 2nd Drive = drive connectors 1, 3, 5, or 7

Note: The NVMe drive population rules defined above are only applicable when the Intel® VROC accessory option is installed and used to provide NVMe drive management.


7.4.4.2 4 port / 8 port PCIe* Switch to 8 x 2.5" Combo Backplane


The following information is applicable when PCIe signal to the 8x2.5" combo backplane are cabled from 4 or 8 port PCIe Switch add-in cards.

- NVMe drive management sideband signals on the backplane are routed between drive connector sets: **(0,1,2,3)** and **(4,5,6,7)**
- In order to support NVMe drive management within a defined drive set, an NVMe drive **MUST** be populated in the first drive connector of the given set (drive connectors **0 or 4**). Additional NVMe drives within the drive set must be populated in sequential order with no gaps between drive connectors.
- Combining NVMe drives and SAS/SATA drives within a defined drive set is **NOT** supported.



The following illustrations identify supported and unsupported drive populations associated with any defined drive set of the 8x2.5" combo backplane when an Intel® VROC key is installed to the server board and the PCIe source to the backplane is from an add-in PCIe Switch card.






Note: The NVMe drive population rules defined above are only applicable when the Intel® VROC accessory option is installed and used to provide NVMe drive management.

7.5 Intel® Integrated RAID Module Support

The system has support for many Intel and third-party PCIe* add-in 12 Gb RAID adapters which can be installed in available PCIe add-in cards slots. For system configurations with limited add-in card slot availability, an optional Intel® Integrated RAID mezzanine module can be installed onto a high density 80-pin connector (labeled “SAS Module”) on the server board.






Figure 71. Intel® Integrated RAID Module

For a list of supported Intel Integrated RAID Module options, please visit the Intel® Server Configurator Tool at <https://serverconfigurator.intel.com>.

Note: Intel Integrated RAID Modules cannot be supported in systems configured with a CPU #1 processor that supports the Integrated® Omni-Path Fabric connector.

7.5.1 RAID Maintenance Free Backup Unit (AXXRMFBUs) Support

The 2U system provides mounting plate options to support the concurrent installation of up to five possible RAID Maintenance Free Backup Units (RMFBUs).

All 2U system/chassis models include a mounting plate capable of supporting up to three AXXRMFBUs or later modules or up to two older AXXRMFBUs plus one AXXRMFBUs or later module. The mounting plate is attached to the air duct using fastener screws and then each RMFBUs module is snapped into place. The triple RMFBUs mounting plate cannot be used if SSDs are mounted to the air duct.




Figure 72. Triple RMFBU (default) mounting bracket installation

Available as an accessory kit option (iPC – AWTAXBBUBKT) is a mounting plate supporting up to two RMFBU modules (AXXRMFBU2 or later). This option is attached to the air duct in two possible mounting locations as shown below.




Figure 73. Dual RMFBU mounting bracket installation

Note: The two RMFBU node mounting plates cannot be installed if a full length PCIe add-in card is installed in the adjacent riser card or if the Intel® Xeon Phi™ shipping bracket is installed in the system.

7.6 Intel® RAID Expander Card Support

System configurations consisting of more than eight SAS drives may require the addition of an SAS expander card. In the 2U 16x2.5" or 24x2.5" system configurations, an SAS expander can be mounted in the area directly behind the drive bays.

Note: The onboard SATA controllers are not compatible with and cannot be used with RAID expander cards.





Figure 74. SAS expander card installation

In the 2U 12x3.5" system configuration, the SAS expander must be installed in any available PCIe add-in slot. Intel makes available the following RAID expander cards.

Table 63. Intel® SAS RAID expander card support

Product Image and iPC	Product Description
iPC – RES3FV288 	Intel® RAID Expander RES3FV288 • SAS 3.0 12 Gb/s expander • 6 Gbps data aggregation for 12 Gbps data transfer with 6 Gb/s devices • Low profile MD2 PCIe* add-in card form factor • 28 internal ports and 8 external ports • Redundant dual-port capable for HA and high performance systems • Power from PCIe x1 • HD mini-SAS 8643 connectors Kit includes: (1) SAS expander card, (2) HD-HD 250 mm expander-to-RAID card cables, PCI brackets for low profile and full height.
iPC – RES3TV360 	Intel® RAID Expander RES3TV360 • SAS 3.0 12 Gb/s expander featuring 6 Gbps data aggregation for 12 Gbps data transfer with 6 Gb/s devices • Internal mount mid-plane form factor • 36 internal ports supporting point-to-point 12, 6, and 3 Gb/s data transfer rates • RA 4-pin power connector • HD mini-SAS 8643 connectors Kit includes: (1) SAS expander card; (1) 130 mm power cable; (1 set) expander-to-backplane cables: (4) HD-HD 165 mm, (1) HD-HD 300 mm, (1) HD-HD 250 mm; (3) rubber pads; mounting screws.

For additional information on Intel RAID Expander Cards, please refer to the appropriate Intel RAID Expander Card User Guide.

**Figure 75. Intel® RAID Expander RES3FV288 connector identification**




Figure 76. 12 GB internal 36-Port Intel® RAID Expander Card RES3TV360 connector identification

See the *Intel® Server S2600WF Product Family Configuration Guide* for SAS cabling requirements.
For system cable routing diagrams, see Appendix A.

8. Front Control Panel and I/O Panel Overview

All system configurations include a control panel on the front of the system providing push button system controls and LED indicators for system features. Systems configured with eight 3.5" hard drive bays or up to sixteen 2.5" hard drive bays will also include an I/O Panel providing additional system I/O features. This section describes the features and functions of both front panel options.

8.1 I/O Panel Features




Figure 77. Front I/O panel features

- **Video connector** – The front I/O panel video connector gives the option of attaching a monitor to the front of the system. When BIOS detects that a monitor is attached to the front video connector, it disables the video signals routed to the on-board video connector on the back of the system. Video resolution from the front video connector may be lower than that of the rear onboard video connector. A short video cable should be used for best resolution. The front video connector is cabled to a 2x7 header on the server board labeled “FP Video”.
- **USB 2.0/3.0 Ports** –The front I/O panel includes two USB 2.0/3.0 ports. The USB ports are cabled to a blue 2x5 connector on the server board labeled “FP_USB”.

Note: Due to signal strength limits associated with USB 3.0 ports cabled to a front panel, some marginally compliant USB 3.0 devices may not be supported from these ports. In addition, server systems based on the Intel® Server Board S2600WF cannot be USB 3.0 certified with USB 3.0 ports cabled to a front panel.

8.2 Control Panel Features

The system includes a control panel that provides push button system controls and LED indicators for several system features. Depending on the hard drive configuration, the front control panel may come in either of two formats; however, both provide the same functionality. This section will provide a description for each front control panel feature.




Figure 78. Front control panel features

- **System ID button w/ integrated LED** – Toggles the integrated ID LED and the blue server board ID LED on and off. The system ID LED is used to identify the system for maintenance when installed in a rack of similar server systems. The system ID LED can also be toggled on and off remotely using the IPMI “Chassis Identify” command which causes the LED to blink for 15 seconds.
- **NMI Button** – When the NMI button is pressed, it puts the server in a halt state and issues a non-maskable interrupt (NMI). This can be useful when performing diagnostics for a given issue where a memory download is necessary to help determine the cause of the problem. To prevent an inadvertent system halt, the actual NMI button is located behind the front control panel faceplate where it is only accessible with the use of a small tipped tool like a pin or paper clip.
- **Network activity LEDs** – The front control panel includes an activity LED indicator for each on-board Network Interface Controller (NIC). When a network link is detected, the LED turns on solid. The LED blinks once network activity occurs at a rate that is consistent with the amount of network activity that is occurring.
- **System cold reset button** – When pressed, this button reboots and re-initializes the system.
- **System status LED** – The system status LED is a bi-color (green/amber) indicator that shows the current health of the server system. The system provides two locations for this feature; one is located on the front control panel and the other is located on the back edge of the server board, viewable from the back of the system. Both LEDs are tied together and show the same state. The system status LED states are driven by the onboard platform management subsystem. The following table provides a description of each supported LED state.

Table 64. System status LED state definitions

LED State	System State	BIOS Status Description
Off	System is not operating.	- System AC power is off. - System is in EuP Lot6 off mode. - System is in S5 soft-off state.
Solid green	System is operating normally.	- System is running (in S0 State) and its status is healthy. The system is not exhibiting any errors. AC power is present and BMC has booted and manageability functionality is up and running. - After a BMC reset, and in conjunction with the chassis ID solid on, the BMC is booting Linux*. Control has been passed from BMC uBoot to BMC Linux* itself. It is in this state for roughly 10-20 seconds.
Blinking green	System is operating in a degraded state although still functioning, or system is operating in a redundant state but with an impending failure warning.	- Redundancy loss such as power-supply or fan. Applies only if the associated platform sub-system has redundancy capabilities. - Fan warning or failure when the number of fully operational fans is less than the minimum number needed to cool the system. - Non-critical threshold crossed – Temperature (including HSBP temp), voltage, input power to power supply, output current for main power rail from power supply and Processor Thermal Control (Therm Ctrl) sensors. - Power supply predictive failure occurred while redundant power supply configuration was present. - Unable to use all of the installed memory (more than 1 DIMM installed). - Correctable Errors over a threshold and migrating to a spare DIMM (memory sparing). This indicates that the system no longer has spared DIMMs (a redundancy lost condition). Corresponding DIMM LED lit. - In mirrored configuration, when memory mirroring takes place and system loses memory redundancy. - Battery failure. - BMC executing in uBoot. (Indicated by Chassis ID blinking at 3Hz). System in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux*. Server will be in this state 6-8 seconds after BMC reset while it pulls the Linux* image into flash. - BMC Watchdog has reset the BMC. - Power Unit sensor offset for configuration error is asserted. - HDD HSC is off-line or degraded.
Blinking amber	System is operating in a degraded state with an impending failure warning, although still functioning. System is likely to fail.	- Critical threshold crossed – Voltage, temperature (including HSBP temp), input power to power supply, output current for main power rail from power supply and PROCHOT (Therm Ctrl) sensors. - VRD Hot asserted. - Minimum number of fans to cool the system not present or failed. - Hard drive fault. - Power Unit Redundancy sensor – Insufficient resources offset (indicates not enough power supplies present). - In non-sparing and non-mirroring mode if the threshold of correctable errors is crossed within the window.

Solid amber	Critical/non-recoverable – system is halted. Fatal alarm – system has failed or shut down.	• CPU CATERR signal asserted. • MSID mismatch detected (CATERR also asserts for this case). • CPU 1 is missing. • CPU Thermal Trip. • No power good – power fault. • DIMM failure when there is only 1 DIMM present and hence no good memory present. • Runtime memory uncorrectable error in non-redundant mode. • DIMM Thermal Trip or equivalent. • SSB Thermal Trip or equivalent. • CPU ERR2 signal asserted. • BMC/Video memory test failed. (Chassis ID shows blue/solid-on for this condition.) • Both uBoot BMC firmware images are bad. (Chassis ID shows blue/solid-on for this condition.) • 240 VA fault. • Fatal Error in processor initialization: - o Processor family not identical - o Processor model not identical - o Processor core/thread counts not identical - o Processor cache size not identical - o Unable to synchronize processor frequency - o Unable to synchronize QPI link frequency • Uncorrectable memory error in a non-redundant mode.
-------------	--	---

- **Power/sleep button** – Toggles the system power on and off. This button also functions as a sleep button if enabled by an ACPI compliant operating system. Pressing this button sends a signal to the integrated BMC, which either powers on or powers off the system. The integrated LED is a single color (green) and is capable of supporting different indicator states as defined in the following table.

Table 65. Power/sleep LED functional states

State	Power Mode	LED	Description
Power-off	Non-ACPI	Off	System power is off, and the BIOS has not initialized the chipset.
Power-on	Non-ACPI	On	System power is on.
S5	ACPI	Off	Mechanical is off and the operating system has not saved any context to the hard disk.
S0	ACPI	On	System and the operating system are up and running.

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

9. PCIe* Riser Card Support

The server board provides three riser card slots identified as: Riser Slot #1, Riser Slot #2, and Riser Slot #3. Per the PCIe* specification, each riser card slot can support a maximum 75 W of power. The PCIe bus interface for each riser card slot is supported by each of the two installed processors. The following tables provide the PCIe bus routing for all supported risers cards.

Note: The riser card slots are specifically designed to support riser cards only. Attempting to install a PCIe add-in card directly into a riser card slot on the server board may damage the server board, the add-in card, or both.

Note: A dual processor configuration is required when using Riser Slot #2 and Riser Slot #3, as well as the bottom add-in card slot for 2U riser cards installed in Riser Slot #1.




Figure 79. PCIe* add-in card support

Table 66. Riser slot #1 PCIe* root port mapping

PCIe* Slot	2U – 3-Slot Riser Card iPC – A2UL8RISER2	2U – 2-Slot Riser Card iPC – A2UL16RISER2
Top	CPU #1 – Ports 1A and 1B (x8 elec, x16 mech)	CPU #1 – Ports 1A thru 1D (x16 elec, x16 mech)
Middle	CPU #1 – Ports 1C and 1D (x8 elec, x16 mech)	N/A
Bottom	CPU #2 – Ports 1C and 1D (x8 elec, x8 mech)	CPU #2 – Ports 1C and 1D (x8 elec, x8 mech)

Table 67. Riser slot #2 PCIe* root port mapping

PCIe* Slot	2U – 3-Slot Riser Card iPC – A2UL8RISER2	2U – 2-Slot Riser Card iPC – A2UL16RISER2
Top	CPU #2 – Ports 2A and 2B (x8 elec, x16 mech)	CPU #2 – Ports 2A thru 2D (x16 elec, x16 mech)
Middle	CPU #2 – Ports 2C and 2D (x8 elec, x16 mech)	N/A
Bottom	CPU #2 – Ports 1A and 1B (x8 elec, x8 mech)	CPU #2 – Ports 1A and 1B (x8 elec, x8 mech)


Table 68. Riser slot #3 PCIe* root port mapping

PCIe* Slot	2U – Low Profile Riser Card iPC – A2UX8X4RISER	Notes
Top	CPU #2 – DMI x4 (x4 elec, x8 mech)	Low profile cards only.
Bottom	CPU #2 – Ports 3C and 3D (x8 elec, x8 mech)	Low profile cards only.

Note: Riser Slot# 3 does not support SMBus device aliasing. SMBus aliasing prevents devices with common address requirements from conflicting with each other. Intel recommends that any PCIe add-in card requiring SMBus support be installed into any available add-in card slot on either Riser #1 or Riser #2.

9.1 Riser Card Assembly

The system includes two different riser cards assemblies, one supporting riser slot #1 and one supporting both riser slots #2 and #3 in a back-to-back butterfly configuration.

**Figure 80. Riser card #1 bracket**




Figure 81. Left and right views of riser card #2 and #3 bracket


Depending on the riser card option installed, riser slots #1 and #2 can each support up to two full height, full length add-in cards (top and middle slots on each riser) and one full height, half-length add-in card (bottom slot on each riser).

Note: Add-in cards that exceed the PCI specification for half-length PCI add-in cards (167.65 mm or 6.6 in) may interfere with other installed devices on the server board when installed in the bottom add-in card slot.

Two card guide brackets on the air duct provide support for full height/full length add-in cards when installed in either the middle or top add-in card slots of each installed riser card assembly.

When installed, riser slot #3 can support up to two low profile add-in cards. To avoid possible add-in card bracket interference when installing add-in cards into both riser card #2 and riser card #3, add-in cards in riser #2 should be installed before those to be installed in riser #3.

Installation of each riser card assembly into the chassis is tool-less. Hooks on the back edge of the riser card assembly are aligned with slots on the chassis, then each assembly is pushed down into the respective riser card slots on the server board.



2UWFP047

Figure 82. Full height/full length add-in card support

Note: Different system models within this 2U product family have different PCIe* add-in card population requirements based on supported air flow limits. See section 4.1 for more information. System integrators should identify PCI add-in card air flow requirements from vendor specifications when integrating any add-in card into the system, to ensure the chosen PCIe* add-in card slot can meet the card's air flow requirements.

9.2 Riser Slot #1 and Riser Slot #2 Riser Card Options

Several multi-slot PCI riser card options are available for this server product family. Available riser cards for riser slots #1 and #2 are common between the two slots.

9.2.1 Three-Slot PCIe Riser Card (iPC – A2UL8RISER2)




Figure 83. Three-slot PCIe* riser card (iPC – A2UL8RISER2)

Table 69. Three-slot PCIe* riser card slot description

Slot #	Description
Slot-1 (top)	PCIe x8 elec, x16 mechanical
Slot-2 (middle)	PCIe x8 elec, x16 mechanical
Slot-3 (bottom)	PCIe x8 elec, x8 mechanical

9.2.2 Two-Slot PCIe Riser Card (iPC – A2UL16RISER2)




Figure 84. Two-slot PCIe riser card (iPC – A2UL16RISER2)

Table 70. Two-slot PCIe riser card slot description

Slot #	Description
Slot-1 (top)	PCIe x16 elec, x16 mechanical
Slot-2 (bottom)	PCIe x8 elec, x8 mechanical

9.3 Riser Slot #3 Riser Card Option (iPC – A2UX8X4RISER)

Riser slot #3 is provided to support up to two additional PCIe* add-in card slots for 2U server configurations. The available riser card option is designed to support low profile add-in cards only.




Figure 85. Low profile riser card (iPC – A2UX8X4RISER)

Table 71. Low profile riser card slot description

Slot #	Description
Slot-1 (top)	PCIe x4 elec, x8 mechanical
Slot-2 (bottom)	PCIe x8 elec, x8 mechanical

10. OCP* Compatible Intel® Ethernet Network Adapter Support

To broaden the standard on-board feature set, the server board provides support for one of several available Intel® Ethernet Network Adapters compatible with the Open Compute Project* (OCP*) specification. The OCP compatible modules are mounted to a high density 120-pin connector on the server board labeled OCP_IO_Module.




Figure 86. OCP module placement

Table 72. OCP* Compatible Intel® Ethernet Network Adapters

Description	iPC
Quad Port, 1GB, RJ45	I357T4OCPG1P5
Quad Port, SFP+	X527DA4OCPG1P5
Dual Port, SFP+ (WFT Boards Only)	X527DA2OCPG1P5
Dual Port, 10Gb RJ45 (WFT Boards Only)	X557T2OCPG1P5

Note: Only dual port SFP+ and dual port 10GB RJ45 OCP modules are supported on Intel® Server Board S2600WFT.

11. Basic and Advanced Server Management Features

The integrated BMC has support for basic and advanced server management features. Basic management features are available by default. Advanced management features are enabled with the addition of an optionally installed Intel® Remote Management Module 4 Lite (Intel® RMM4 Lite) key.

Table 73. Intel® Remote Management Module 4 (RMM4) options

iPC	Description	Kit Contents	Benefits
AXXRMM4LITE2	Intel® Remote Management Module 4 Lite	RMM4 Lite 2 Activation Key	Enables KVM & media redirection

When the BMC firmware initializes, it attempts to access the Intel RMM4 Lite. If the attempt to access Intel RMM4 Lite is successful, then the BMC activates the advanced features.

The following table identifies both **Basic** and **Advanced** server management features.

Table 74. Basic and advanced server management features overview

Feature	Basic	Advanced w/RMM4 Lite Key
IPMI 2.0 Feature Support	X	X
In-circuit BMC Firmware Update	X	X
FRB 2	X	X
Chassis Intrusion Detection	X	X
Fan Redundancy Monitoring	X	X
Hot-Swap Fan Support	X	X
Acoustic Management	X	X
Diagnostic Beep Code Support	X	X
Power State Retention	X	X
ARP/DHCP Support	X	X
PECI Thermal Management Support	X	X
E-mail Alerting	X	X
Embedded Web Server	X	X
SSH Support	X	X
Integrated KVM		X
Integrated Remote Media Redirection		X
Lightweight Directory Access Protocol (LDAP)	X	X
Intel® Intelligent Power Node Manager Support	X	X
SMASH CLP	X	X

On the server board, the Intel RMM4 Lite key is installed at the following location.




Figure 87. Intel® RMM4 Lite activation key placement

11.1 Dedicated Management Port

The server board includes a dedicated 1 GbE, RJ45 management port. The management port is active with or without the RMM4 Lite key installed.




Figure 88. Dedicated management port

11.2 Embedded Web Server

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

- Microsoft Internet Explorer*
- Mozilla Firefox*
- Google Chrome*
- Safari*

The embedded web user interface supports strong security (authentication, encryption, and firewall support) since it enables remote server configuration and control. The user interface presented by the embedded web user interface shall authenticate the user before allowing a web session to be initiated. Encryption using 256-bit secure sockets layer (SSL) is supported. User authentication is based on user id and password.

The GUI presented by the embedded web server authenticates the user before allowing a web session to be initiated. It presents all functions to all users but disables those functions that the user does not have privilege to execute. For example, if a user does not have privilege to power control, then the item shall be displayed in a greyed-out font on that user's UI display. The web GUI also provides a launch point for advanced features, KVM and media redirection. These features are also grayed out in the GUI unless the system has been updated to support these advanced features. The embedded web server displays US English or Chinese language output only.

Additional features supported by the web GUI can

- Present all the Basic features to the users
- Power on/Power off/reset the server and view current power state
- Display BIOS, BMC, ME and SDR version information
- Display overall system health.
- Display configuration of various IPMI over LAN parameters for both IPV4 and IPV6
- Display configuration of alerts (SNMP and SMTP)
- Display system asset information for the product, board, and chassis.
- Display BMC-owned sensors (name, status, current reading, enabled thresholds), including color-code status of sensors.
- Provide ability to filter sensors based on sensor type (Voltage, Temperature, Fan and Power supply related)
- Automatically refresh of sensor data with a configurable refresh rate
- Display online help
- Display/clear SEL (display is in easily understandable human readable format)
- Support major industry-standard browsers (Microsoft Internet Explorer* and Mozilla Firefox*)
- Automatically time out the GUI session after a user-configurable inactivity period. By default, this inactivity period is 30 minutes.
- Using the Embedded Platform Debug feature, allow the user to initiate a "debug dump" to a file that can be sent to Intel® for debug purposes
- Employ the Virtual Front Panel to provide the same functionality as the local front panel. The displayed LEDs match the current state of the local panel LEDs. The displayed buttons (for example, power button) can be used in the same manner as the local buttons.
- Display of ME sensor data. Only sensors that have associated SDRs loaded will be displayed.
- Save the SEL to a file
- Force HTTPS connectivity for greater security. This is provided through a configuration option in the UI.
- Display processor and memory information that is available over IPMI over LAN.
- Get and set Node Manager (NM) power policies
- Display power consumed by the server
- View and configure VLAN settings

- Warn users that the reconfiguration of IP address will cause a disconnect
- Block logins for a period of time after several consecutive failed login attempts. The lock-out period and the number of failed logins that initiates the lock-out period are configurable by the user.
- Employ Server Power Control to force boot-up into Setup on a reset
- Report System POST results. The web server provides the system's Power-On Self-Test (POST) sequence for the previous two boot cycles, including timestamps. The timestamps may be displayed as a time relative to the start of POST or the previous POST code.
- Allow customization of ports. The web server provides the ability to customize the port numbers used for SMASH, http, https, KVM, secure KVM, remote media, and secure remote media.

For additional information, refer to the *Intel® Remote Management Module 4 and Integrated BMC Web Console User Guide*.

11.3 Advanced Management Feature Support (Intel® RMM4 Lite)

The integrated baseboard management controller has support for advanced management features which are enabled when an optional Intel® Remote Management Module 4 Lite (RMM4 Lite) is installed. The Intel RMM4-lite option offers convenient, remote KVM access and control through LAN and internet. It captures, digitizes, and compresses video and transmits it with keyboard and mouse signals to and from a remote computer. Remote access and control software runs in the integrated baseboard management controller, utilizing expanded capabilities enabled by the Intel RMM4 hardware.

- Key Features of the RMM4-lite enablement key are:
KVM redirection from either the dedicated management NIC or the server board NICs used for management traffic; up to two KVM sessions. Automatically senses video resolution for best possible screen capture, high performance mouse tracking and synchronization. It allows remote viewing and configuration in pre-boot POST and BIOS setup.
- Media Redirection – The media redirection feature is intended to allow system administrators or users to mount a remote IDE or USB CDROM, floppy drive, or a USB flash disk as a remote device to the server. Once mounted, the remote device appears just like a local device to the server allowing system administrators or users to install software (including operating systems), copy files, update BIOS, or boot the server from this device.

11.3.1 Keyboard, Video, Mouse (KVM) Redirection

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

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

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

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

Other attributes of this feature include:

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

KVM redirection feature supports the following resolutions and refresh rates:

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

11.3.2 Remote Console

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

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

11.3.3 Performance

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

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

For the best possible KVM performance, a 2Mb/sec link or higher is recommended.

The redirection of KVM over IP is performed in parallel with the local KVM without affecting the local KVM operation.

11.3.4 Availability

The remote KVM session is available even when the server is powered off (in stand-by mode). No restart of the remote KVM session shall be required during a server reset or power on/off. A BMC reset (for example, due to a BMC Watchdog initiated reset or BMC reset after BMC FW update) will require the session to be re-established. KVM sessions persist across system reset, but not across an AC power loss.

11.3.5 Security

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

11.3.6 Usage

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

11.3.7 Force-enter BIOS Setup

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

11.3.8 Media Redirection

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

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

The following capabilities are supported:

- The operation of remotely mounted devices is independent of the local devices on the server. Both remote and local devices are usable in parallel.
- Either IDE (CD-ROM, floppy) or USB devices can be mounted as a remote device to the server.
- It is possible to boot all supported operating systems from the remotely mounted device and to boot from disk IMAGE (*.IMG) and CD-ROM or DVD-ROM ISO files. See the Tested/supported Operating System List for more information.
- Media redirection supports redirection for both a virtual CD device and a virtual Floppy/USB device concurrently. The CD device may be either a local CD drive or else an ISO image file; the Floppy/USB device may be either a local Floppy drive, a local USB device, or else a disk image file.
- The media redirection feature supports multiple encryption algorithms, including RC4 and AES. The actual algorithm that is used is negotiated with the client based on the client's capabilities.
- A remote media session is maintained even when the server is powered off (in standby mode). No restart of the remote media session is required during a server reset or power on/off. An BMC reset (for example, due to an BMC reset after BMC FW update) will require the session to be re-established.
- The mounted device is visible to (and usable by) managed system's OS and BIOS in both pre-boot and post-boot states.
- The mounted device shows up in the BIOS boot order and it is possible to change the BIOS boot order to boot from this remote device.

- It is possible to install an operating system on a bare metal server (no OS present) using the remotely mounted device. This may also require the use of KVM-r to configure the OS during install.

USB storage devices will appear as floppy disks over media redirection. This allows for the installation of device drivers during OS installation.

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

11.3.8.1 Availability

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

11.3.8.2 Network Port Usage

The KVM and media redirection features use the following ports:

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

For additional information, reference the Intel® Remote Management Module 4 and Integrated BMC Web Console Users Guide.

Appendix A. Integration and Usage Tips

This section provides a list of useful information that is unique to the Intel® Server System R2000WF product family and should be kept in mind while configuring your server system.

- When adding or removing components or peripherals, power cords must be disconnected from the server. With power applied to the server, standby voltages are still present even though the server board is powered off.
- This server system supports the Intel® Xeon® processor Scalable family with a thermal design power (TDP) of up to and including 145 Watts. Previous generations of the Intel® Xeon® processors are not supported.
- Processors must be installed in order. CPU 1 must be populated for the server board to operate.
- The riser card slots are specifically designed to support riser cards only. Attempting to install a PCIe* add-in card directly into a riser card slot on the server board may damage the server board, the add-in card, or both.
- This server board only supports DDR4 ECC RDIMM – Registered (Buffered) DIMMS and DDR4 ECC LRDIMM – Load Reduced DIMMs.
- For the best performance, the number of DDR4 DIMMs installed should be balanced across both processor sockets and memory channels.
- On the back edge of the server board are eight diagnostic LEDs that display a sequence of amber POST codes during the boot process. If the server board hangs during POST, the LEDs display the last POST event run before the hang.
- The system status LED is set to solid amber for all fatal errors that are detected during processor initialization. A solid amber system status LED indicates that an unrecoverable system failure condition has occurred.
- RAID partitions created using either embedded software RAID option, Intel® RSTe or Intel® Embedded Server RAID Technology 2, cannot span across the two embedded SATA controllers. Only drives attached to a common SATA controller can be included in a RAID partition.
- The FRUSDR utility must be run as part of the initial platform integration process before it is deployed into a live operating environment. Once the initial FRU and SDR data is loaded on to the system, all subsequent system configuration changes automatically update SDR data using the BMC auto configuration feature, without having to run the FRUSDR utility again. However, to ensure the latest sensor data is installed, the SDR data should be updated to the latest available as part of a planned system software update.
- Make sure the latest system software is loaded on the server. This includes system BIOS, BMC firmware, Intel® ME firmware and FRUSDR. The latest system software can be downloaded from <http://downloadcenter.intel.com>.

Appendix B. POST Code Diagnostic LED Decoder

As an aid in troubleshooting a system hang that occurs during a system POST process, the server board includes a bank of eight POST code diagnostic LEDs on the back edge of the server board.

During the system boot process, Memory Reference Code (MRC) and system BIOS execute a number of memory initialization and platform configuration processes, each of which is assigned a hex POST code number.

As each routine is started, the given POST code number is displayed to the POST code diagnostic LEDs on the back edge of the server board.

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

Each POST code is represented by eight LEDs, four green and four amber. The POST codes are divided into two nibbles, an upper nibble and a lower nibble. The upper nibble bits are represented by amber diagnostic LEDs and the lower nibble bits are represented by green diagnostic LEDs. If the bit is set in the upper and lower nibbles, the corresponding LED is lit. If the bit is clear, the corresponding LED is off.




Figure 89. On-board POST Diagnostic LEDs

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

In the following example, the BIOS sends a value of AC to the diagnostic LED decoder. The LEDs are decoded as shown in Table 75, where the upper nibble bits represented by the amber LEDs equal 1010_b or A_h and the lower nibble bits represented by the green LEDs equal 1100_b or C_h . The two are concatenated as AC_h .

Table 75. POST progress code LED example

Upper Nibble						Lower Nibble					
LED 3 (MSB)	LED 2	LED 1	LED 0 (LSB)	Binary Code	Hex Code	LED 3 (MSB)	LED 2	LED 1	LED 0 (LSB)	Binary Code	Hex Code
ON	off	ON	off	1010	A	ON	ON	off	off	1100	C

B.1. Early POST Memory Initialization MRC Diagnostic Codes

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

The MRC progress codes are displayed to the diagnostic LEDs that show the execution point in the MRC operational path at each step.

Table 76. MRC progress codes

Post Code (Hex)	Upper Nibble				Lower Nibble				Description
B0	1	0	1	1	0	0	0	0	Detect DIMM population
B1	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5	1	0	1	1	0	1	0	1	Program registers on the channel level
B6	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9	1	0	1	1	1	0	0	1	Hardware memory test and init
BA	1	0	1	1	1	0	1	0	Execute software memory init
BB	1	0	1	1	1	0	1	1	Program memory map and interleaving
BC	1	0	1	1	1	1	0	0	Program RAS configuration
BF	1	0	1	1	1	1	1	1	MRC is done

Should a major memory initialization error occur, preventing the system from booting with data integrity, a beep code is generated, the MRC displays a fatal error code on the diagnostic LEDs, and a system halt command is executed. Fatal MRC error halts do NOT change the state of the system status LED and they do NOT get logged as SEL events.

The following table lists all MRC fatal errors that are displayed to the diagnostic LEDs.

Note: Fatal MRC errors display POST error codes that may be the same as BIOS POST progress codes displayed later in the POST process. The fatal MRC codes can be distinguished from the BIOS POST progress codes by the accompanying memory failure beep code of three long beeps as identified in Table 80.

Table 77. MRC fatal error codes

Post Code (Hex)	Upper Nibble				Lower Nibble				Description
E8	1	1	1	0	1	0	0	0	No usable memory error 01h = No memory was detected from SPD read, or invalid config that causes no operable memory. 02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error. 03h = No memory installed. All channels are disabled.
E9	1	1	1	0	1	0	0	1	Memory is locked by Intel Trusted Execution Technology and is inaccessible
EA	1	1	1	0	1	0	1	0	DDR4 channel training error 01h = Error on read DQ/DQS (Data/Data Strobe) init 02h = Error on Receive Enable 03h = Error on Write Leveling 04h = Error on write DQ/DQS (Data/Data Strobe)
EB	1	1	1	0	1	0	1	1	Memory test failure 01h = Software memtest failure. 02h = Hardware memtest failed.
ED	1	1	1	0	1	1	0	1	DIMM configuration population error 01h = Different DIMM types (RDIMM, LRDIMM) are detected installed in the system. 02h = Violation of DIMM population rules. 03h = The 3rd DIMM slot cannot be populated when QR DIMMs are installed. 04h = UDIMMs are not supported. 05h = Unsupported DIMM Voltage.
EF	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error

B.2. BIOS POST Progress Codes

The following table provides a list of all POST progress codes.

Table 78. POST progress codes

Post Code (Hex)	Upper Nibble				Lower Nibble				Description
SEC Phase									
01	0	0	0	0	0	0	0	1	First POST code after CPU reset
02	0	0	0	0	0	0	1	0	Microcode load begin
03	0	0	0	0	0	0	1	1	CRAM initialization begin
04	0	0	0	0	0	1	0	0	EI Cache When Disabled
05	0	0	0	0	0	1	0	1	SEC Core at Power on Begin
06	0	0	0	0	0	1	1	0	Early CPU initialization during Sec Phase.
UPI RC (Fully leverage without platform change)									
A1	1	0	1	0	0	0	0	1	Collect info such as SBSP, Boot Mode, Reset type etc
A3	1	0	1	0	0	0	1	1	Setup minimum path between SBSP & other sockets
A7	1	0	1	0	0	1	1	1	Topology discovery and route calculation
A8	1	0	1	0	1	0	0	0	Program final route
A9	1	0	1	0	1	0	0	1	Program final IO SAD setting
AA	1	0	1	0	1	0	1	0	Protocol layer and other uncore settings
AB	1	0	1	0	1	0	1	1	Transition links to full speed operation
AC	1	0	1	0	1	1	0	0	Phy layer setting
AD	1	0	1	0	1	1	0	1	Link layer settings
AE	1	0	1	0	1	1	1	0	Coherency settings
AF	1	0	1	0	1	1	1	1	UPI initialization done
07	0	0	0	0	0	1	1	1	Early SB initialization during Sec Phase.
08	0	0	0	0	1	0	0	0	Early NB initialization during Sec Phase.
09	0	0	0	0	1	0	0	1	End Of Sec Phase.
0E	0	0	0	0	1	1	1	0	Microcode Not Found.
0F	0	0	0	0	1	1	1	1	Microcode Not Loaded.
PEI Phase									
10	0	0	0	1	0	0	0	0	PEI Core
11	0	0	0	1	0	0	0	1	CPU PEIM
15	0	0	0	1	0	1	0	1	NB PEIM
19	0	0	0	1	1	0	0	1	SB PEIM
MRC Progress Codes									
31	0	0	1	1	0	0	0	1	Memory Installed
32	0	0	1	1	0	0	1	0	CPU PEIM (CPU Init)
33	0	0	1	1	0	0	1	1	CPU PEIM (Cache Init)
4F	0	1	0	0	1	1	1	1	DXE IPL started
DXE Phase									
60	0	1	1	0	0	0	0	0	DXE Core started
61	0	1	1	0	0	0	0	1	DXE NVRAM Init
62	0	1	1	0	0	0	1	0	DXE Setup Init
63	0	1	1	0	0	0	1	1	DXE CPU Init
65	0	1	1	0	0	1	0	1	DXE CPU BSP Select
66	0	1	1	0	0	1	1	0	DXE CPU AP Init

Post Code (Hex)	Upper Nibble				Lower Nibble				Description
68	0	1	1	0	1	0	0	0	DXE PCI Host Bridge Init
69	0	1	1	0	1	0	0	1	DXE NB Init
6A	0	1	1	0	1	0	1	0	DXE NB SMM Init
70	0	1	1	1	0	0	0	0	DXE SB Init
71	0	1	1	1	0	0	0	1	DXE SB SMM Init
72	0	1	1	1	0	0	1	0	DXE SB devices Init
78	0	1	1	1	1	0	0	0	DXE ACPI Init
79	0	1	1	1	1	0	0	1	DXE CSM Init
80	1	0	0	0	0	0	0	0	DXE BDS Started
81	1	0	0	0	0	0	0	1	DXE BDS connect drivers
82	1	0	0	0	0	0	1	0	DXE PCI Bus begin
83	1	0	0	0	0	0	1	1	DXE PCI Bus HPC Init
84	1	0	0	0	0	1	0	0	DXE PCI Bus enumeration
85	1	0	0	0	0	1	0	1	DXE PCI Bus resource requested
86	1	0	0	0	0	1	1	0	DXE PCI Bus assign resource
87	1	0	0	0	0	1	1	1	DXE CON_OUT connect
88	1	0	0	0	1	0	0	0	DXE CON_IN connect
89	1	0	0	0	1	0	0	1	DXE SIO Init
8A	1	0	0	0	1	0	1	0	DXE USB start
8B	1	0	0	0	1	0	1	1	DXE USB reset
8C	1	0	0	0	1	1	0	0	DXE USB detect
8D	1	0	0	0	1	1	0	1	DXE USB enable
91	1	0	0	1	0	0	0	1	DXE IDE begin
92	1	0	0	1	0	0	1	0	DXE IDE reset
93	1	0	0	1	0	0	1	1	DXE IDE detect
94	1	0	0	1	0	1	0	0	DXE IDE enable
95	1	0	0	1	0	1	0	1	DXE SCSI begin
96	1	0	0	1	0	1	1	0	DXE SCSI reset
97	1	0	0	1	0	1	1	1	DXE SCSI detect
98	1	0	0	1	1	0	0	0	DXE SCSI enable
99	1	0	0	1	1	0	0	1	DXE verifying SETUP password
9B	1	0	0	1	1	0	1	1	DXE SETUP start
9C	1	0	0	1	1	1	0	0	DXE SETUP input wait
9D	1	0	0	1	1	1	0	1	DXE Ready to Boot
9E	1	0	0	1	1	1	1	0	DXE Legacy Boot
9F	1	0	0	1	1	1	1	1	DXE Exit Boot Services
C0	1	1	0	0	0	0	0	0	RT Set Virtual Address Map Begin
C2	1	1	0	0	0	0	1	0	DXE Legacy Option ROM init
C3	1	1	0	0	0	0	1	1	DXE Reset system
C4	1	1	0	0	0	1	0	0	DXE USB Hot plug
C5	1	1	0	0	0	1	0	1	DXE PCI BUS Hot plug
C6	1	1	0	0	0	1	1	0	DXE NVRAM cleanup
C7	1	1	0	0	0	1	1	1	DXE ACPI Enable
0	0	0	0	0	0	0	0	0	Clear POST Code
S3 Resume									
40	0	1	0	0	0	0	0	0	S3 Resume PEIM (S3 started)

Post Code (Hex)	Upper Nibble				Lower Nibble				Description
41	0	1	0	0	0	0	0	1	S3 Resume PEIM (S3 boot script)
42	0	1	0	0	0	0	1	0	S3 Resume PEIM (S3 Video Repost)
43	0	1	0	0	0	0	1	1	S3 Resume PEIM (S3 OS wake)
BIOS Recovery									
46	0	1	0	0	0	1	1	0	PEIM which detected forced Recovery condition
47	0	1	0	0	0	1	1	1	PEIM which detected User Recovery condition
48	0	1	0	0	1	0	0	0	Recovery PEIM (Recovery started)
49	0	1	0	0	1	0	0	1	Recovery PEIM (Capsule found)
4A	0	1	0	0	1	0	1	0	Recovery PEIM (Capsule loaded)
E8	1	1	1	0	1	0	0	0	No Usable Memory Error:
E9	1	1	1	0	1	0	0	1	Memory is locked by Intel® Trusted Execution Technology and is inaccessible.
EA	1	1	1	0	1	0	1	0	DDR4 Channel Training Error:
EB	1	1	1	0	1	0	1	1	Memory Test Failure
ED	1	1	1	0	1	1	0	1	DIMM Configuration/Population Error
EF	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error
B0	1	0	1	1	0	0	0	0	Detect DIMM population
B1	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5	1	0	1	1	0	1	0	1	Program registers on the channel level
B6	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9	1	0	1	1	1	0	0	1	Hardware memory test and init
Ba	1	0	1	1	1	0	1	0	Execute software memory init
BB	1	0	1	1	1	0	1	1	Program memory map and interleaving
BC	1	0	1	1	1	1	0	0	Program RAS configuration
BF	1	0	1	1	1	1	1	1	MRC is done

Appendix C. POST Code Errors

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

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

Table 79 lists the supported POST error codes. Each error code is assigned an error type which determines the action the BIOS takes when the error is encountered. Error types include minor, major, and fatal. The BIOS action for each is defined as follows:

- **Minor:** An error message may be displayed to the screen or to the BIOS setup error manager and the POST error code is logged to the SEL. The system continues booting in a degraded state. The user may want to replace the erroneous unit. The “POST Error Pause” option setting in the BIOS setup does not have any effect on this error.
- **Major:** An error message is displayed to the error manager screen and an error is logged to the SEL. If the BIOS setup option “Post Error Pause” is enabled, operator intervention is required to continue booting the system. If the BIOS setup option “POST Error Pause” is disabled, the system continues to boot.

Note: For 0048 “Password check failed”, the system halts and then, after the next reset/reboot, displays the error code on the error manager screen.

- **Fatal:** If the system cannot boot, POST halts and display the following message:

Unrecoverable fatal error found. System will not boot until the error is resolved
Press <F2> to enter setup

When the <F2> key on the keyboard is pressed, the error message is displayed on the error manager screen and an error is logged to the system event log (SEL) with the POST error code. The system cannot boot unless the error is resolved. The faulty component must be replaced. The “POST Error Pause” option setting in the BIOS setup does not have any effect on this error.

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

Table 79. POST error messages and handling

Error Code	Error Message	Action message	Error Type
0012	System RTC date/time not set		Major
0048	Password check failed	Please put right password.	Major
0140	PCI component encountered a PERR error		Major
0141	PCI resource conflict		Major
0146	PCI out of resources error	Please enable Memory Mapped I/O above 4 GB item at SETUP to use 64bit MMIO.	Major
0191	Processor core/thread count mismatch detected	Please use identical CPU type.	Fatal
0192	Processor cache size mismatch detected	Please use identical CPU type.	Fatal
0194	Processor family mismatch detected	Please use identical CPU type.	Fatal
0195	Processor Intel(R) UPI link frequencies unable to synchronize		Fatal
0196	Processor model mismatch detected	Please use identical CPU type.	Fatal
0197	Processor frequencies unable to synchronize	Please use identical CPU type.	Fatal
5220	BIOS Settings reset to default settings		Major
5221	Passwords cleared by jumper		Major
5224	Password clear jumper is Set	Recommend to remind user to install BIOS password as BIOS admin password is the master keys for several BIOS security features.	Major
8130	Processor 01 disabled		Major
8131	Processor 02 disabled		Major
8160	Processor 01 unable to apply microcode update		Major
8161	Processor 02 unable to apply microcode update		Major
8170	Processor 01 failed self-test (BIST)		Major
8171	Processor 02 failed self-test (BIST)		Major
8180	Processor 01 microcode update not found		Minor
8181	Processor 02 microcode update not found		Minor
8190	Watchdog timer failed on last boot		Major
8198	OS boot watchdog timer failure		Major
8300	Baseboard management controller failed self-test		Major
8305	Hot Swap Controller failure		Major
83A0	Management Engine (ME) failed self-test		Major
83A1	Management Engine (ME) Failed to respond		Major
84F2	Baseboard management controller failed to respond		Major
84F3	Baseboard management controller in update mode		Major
84F4	Sensor data record empty	Please update right SDR.	Major
84FF	System event log full	Please clear SEL through EWS or SELVIEW utility.	Minor
8500	Memory component could not be configured in the selected RAS mode		Major
8501	DIMM Population Error	Please plug DIMM at right population.	Major
8520	CPU1_DIMM_A1 failed test/initialization	Please remove the disabled DIMM.	Major
8521	CPU1_DIMM_A2 failed test/initialization	Please remove the disabled DIMM.	Major
8522	CPU1_DIMM_A3 failed test/initialization	Please remove the disabled DIMM.	Major
8523	CPU1_DIMM_B1 failed test/initialization	Please remove the disabled DIMM.	Major
8524	CPU1_DIMM_B2 failed test/initialization	Please remove the disabled DIMM.	Major
8525	CPU1_DIMM_B3 failed test/initialization	Please remove the disabled DIMM.	Major

8526	CPU1_DIMM_C1 failed test/initialization	Please remove the disabled DIMM.	Major
8527	CPU1_DIMM_C2 failed test/initialization	Please remove the disabled DIMM.	Major
8528	CPU1_DIMM_C3 failed test/initialization	Please remove the disabled DIMM.	Major
8529	CPU1_DIMM_D1 failed test/initialization	Please remove the disabled DIMM.	Major
852A	CPU1_DIMM_D2 failed test/initialization	Please remove the disabled DIMM.	Major
852B	CPU1_DIMM_D3 failed test/initialization	Please remove the disabled DIMM.	Major
852C	CPU1_DIMM_E1 failed test/initialization	Please remove the disabled DIMM.	Major
852D	CPU1_DIMM_E2 failed test/initialization	Please remove the disabled DIMM.	Major
852E	CPU1_DIMM_E3 failed test/initialization	Please remove the disabled DIMM.	Major
852F	CPU1_DIMM_F1 failed test/initialization	Please remove the disabled DIMM.	Major
8530	CPU1_DIMM_F2 failed test/initialization	Please remove the disabled DIMM.	Major
8531	CPU1_DIMM_F3 failed test/initialization	Please remove the disabled DIMM.	Major
8532	CPU1_DIMM_G1 failed test/initialization	Please remove the disabled DIMM.	Major
8533	CPU1_DIMM_G2 failed test/initialization	Please remove the disabled DIMM.	Major
854A	CPU1_DIMM_D2 disabled	Please remove the disabled DIMM.	Major
854B	CPU1_DIMM_D3 disabled	Please remove the disabled DIMM.	Major
854C	CPU1_DIMM_E1 disabled	Please remove the disabled DIMM.	Major
854D	CPU1_DIMM_E2 disabled	Please remove the disabled DIMM.	Major
854E	CPU1DIMM_E3 disabled	Please remove the disabled DIMM.	Major
854F	CPU1DIMM_F1 disabled	Please remove the disabled DIMM.	Major
8550	CPU1DIMM_F2 disabled	Please remove the disabled DIMM.	Major
8551	CPU1DIMM_F3 disabled	Please remove the disabled DIMM.	Major
8552	CPU1DIMM_G1 disabled	Please remove the disabled DIMM.	Major
8553	CPU1DIMM_G2 disabled	Please remove the disabled DIMM.	Major
8554	CPU1DIMM_G3 disabled	Please remove the disabled DIMM.	Major
8555	CPU1DIMM_H1 disabled	Please remove the disabled DIMM.	Major
8556	CPU1DIMM_H2 disabled	Please remove the disabled DIMM.	Major
8557	CPU1DIMM_H3 disabled	Please remove the disabled DIMM.	Major
8558	CPU2_DIMM_A1 disabled	Please remove the disabled DIMM.	Major
8559	CPU2_DIMM_A2 disabled	Please remove the disabled DIMM.	Major
855A	CPU2_DIMM_A3 disabled	Please remove the disabled DIMM.	Major
855B	CPU2_DIMM_B1 disabled	Please remove the disabled DIMM.	Major
855C	CPU2_DIMM_B2 disabled	Please remove the disabled DIMM.	Major
855D	CPU2_DIMM_B3 disabled	Please remove the disabled DIMM.	Major
855E	CPU2_DIMM_C1 disabled	Please remove the disabled DIMM.	Major
855F (Go to 85D0)	CPU2_DIMM_C2 disabled	Please remove the disabled DIMM.	Major
8560	CPU1_DIMM_A1 encountered a Serial Presence Detection (SPD) failure		Major
8561	CPU1_DIMM_A2 encountered a Serial Presence Detection (SPD) failure		Major
8562	CPU1_DIMM_A3 encountered a Serial Presence Detection (SPD) failure		Major
8563	CPU1_DIMM_B1 encountered a Serial Presence Detection (SPD) failure		Major
8564	CPU1_DIMM_B2 encountered a Serial Presence Detection (SPD) failure		Major

8565	CPU1_DIMM_B3 encountered a Serial Presence Detection (SPD) failure		Major
8566	CPU1_DIMM_C1 encountered a Serial Presence Detection (SPD) failure		Major
8567	CPU1_DIMM_C2 encountered a Serial Presence Detection (SPD) failure		Major
8568	CPU1_DIMM_C3 encountered a Serial Presence Detection (SPD) failure		Major
8569	CPU1_DIMM_D1 encountered a Serial Presence Detection (SPD) failure		Major
856A	CPU1_DIMM_D2 encountered a Serial Presence Detection (SPD) failure		Major
856B	CPU1_DIMM_D3 encountered a Serial Presence Detection (SPD) failure		Major
856C	CPU1_DIMM_E1 encountered a Serial Presence Detection (SPD) failure		Major
856D	CPU1_DIMM_E2 encountered a Serial Presence Detection (SPD) failure		Major
856E	CPU1_DIMM_E3 encountered a Serial Presence Detection (SPD) failure		Major
856F	CPU1_DIMM_F1 encountered a Serial Presence Detection (SPD) failure		Major
8570	CPU1_DIMM_F2 encountered a Serial Presence Detection (SPD) failure		Major
8571	CPU1_DIMM_F3 encountered a Serial Presence Detection (SPD) failure		Major
8572	CPU1_DIMM_G1 encountered a Serial Presence Detection (SPD) failure		Major
8573	CPU1_DIMM_G2 encountered a Serial Presence Detection (SPD) failure		Major
8574	CPU1_DIMM_G3 encountered a Serial Presence Detection (SPD) failure		Major
8575	CPU1_DIMM_H1 encountered a Serial Presence Detection (SPD) failure		Major
8576	CPU1_DIMM_H2 encountered a Serial Presence Detection (SPD) failure		Major
8577	CPU1_DIMM_H3 encountered a Serial Presence Detection (SPD) failure		Major
8578	CPU2_DIMM_A1 encountered a Serial Presence Detection (SPD) failure		Major
8579	CPU2_DIMM_A2 encountered a Serial Presence Detection (SPD) failure		Major
857A	CPU2_DIMM_A3 encountered a Serial Presence Detection (SPD) failure		Major
857B	CPU2_DIMM_B1 encountered a Serial Presence Detection (SPD) failure		Major
857C	CPU2_DIMM_B2 encountered a Serial Presence Detection (SPD) failure		Major
857D	CPU2_DIMM_B3 encountered a Serial Presence Detection (SPD) failure		Major
857E	CPU2_DIMM_C1 encountered a Serial Presence Detection (SPD) failure		Major

857F (Go to 85E0)	CPU2_DIMM_C2 encountered a Serial Presence Detection (SPD) failure		Major
85C0	CPU2_DIMM_C3 failed test/initialization	Please remove the disabled DIMM.	Major
85C1	CPU2_DIMM_D1 failed test/initialization	Please remove the disabled DIMM.	Major
85C2	CPU2_DIMM_D2 failed test/initialization	Please remove the disabled DIMM.	Major
85C3	CPU2_DIMM_D3 failed test/initialization	Please remove the disabled DIMM.	Major
85C4	CPU2_DIMM_E1 failed test/initialization	Please remove the disabled DIMM.	Major
85C5	CPU2_DIMM_E2 failed test/initialization	Please remove the disabled DIMM.	Major
85C6	CPU2_DIMM_E3 failed test/initialization	Please remove the disabled DIMM.	Major
85C7	CPU2_DIMM_F1 failed test/initialization	Please remove the disabled DIMM.	Major
85C8	CPU2_DIMM_F2 failed test/initialization	Please remove the disabled DIMM.	Major
85C9	CPU2_DIMM_F3 failed test/initialization	Please remove the disabled DIMM.	Major
85CA	CPU2_DIMM_G1 failed test/initialization	Please remove the disabled DIMM.	Major
85CB	CPU2_DIMM_G2 failed test/initialization	Please remove the disabled DIMM.	Major
85CC	CPU2_DIMM_G3 failed test/initialization	Please remove the disabled DIMM.	Major
85CD	CPU2_DIMM_H1 failed test/initialization	Please remove the disabled DIMM.	Major
85CE	CPU2_DIMM_H2 failed test/initialization	Please remove the disabled DIMM.	Major
85CF	CPU2_DIMM_H3 failed test/initialization	Please remove the disabled DIMM.	Major
85D0	CPU2_DIMM_C3 disabled	Please remove the disabled DIMM.	Major
85D1	CPU2_DIMM_D1 disabled	Please remove the disabled DIMM.	Major
85D2	CPU2_DIMM_D2 disabled	Please remove the disabled DIMM.	Major
85D3	CPU2_DIMM_D3 disabled	Please remove the disabled DIMM.	Major
85D4	CPU2_DIMM_E1 disabled	Please remove the disabled DIMM.	Major
85D5	CPU2_DIMM_E2 disabled	Please remove the disabled DIMM.	Major
85D6	CPU2_DIMM_E3 disabled	Please remove the disabled DIMM.	Major
85D7	CPU2_DIMM_F1 disabled	Please remove the disabled DIMM.	Major
85D8	CPU2_DIMM_F2 disabled	Please remove the disabled DIMM.	Major
85D9	CPU2_DIMM_F3 disabled	Please remove the disabled DIMM.	Major
85DA	CPU2_DIMM_G1 disabled	Please remove the disabled DIMM.	Major
85DB	CPU2_DIMM_G2 disabled	Please remove the disabled DIMM.	Major
85DC	CPU2_DIMM_G3 disabled	Please remove the disabled DIMM.	Major
85DD	CPU2_DIMM_H1 disabled	Please remove the disabled DIMM.	Major
85DE	CPU2_DIMM_H2 disabled	Please remove the disabled DIMM.	Major
85DF	CPU2_DIMM_H3 disabled	Please remove the disabled DIMM.	Major
85E0	CPU2_DIMM_C3 encountered a Serial Presence Detection (SPD) failure		Major
85E1	CPU2_DIMM_D1 encountered a Serial Presence Detection (SPD) failure		Major
85E2	CPU2_DIMM_D2 encountered a Serial Presence Detection (SPD) failure		Major
85E3	CPU2_DIMM_D3 encountered a Serial Presence Detection (SPD) failure		Major
85E4	CPU2_DIMM_E1 encountered a Serial Presence Detection (SPD) failure		Major
85E5	CPU2_DIMM_E2 encountered a Serial Presence Detection (SPD) failure		Major
85E6	CPU2_DIMM_E3 encountered a Serial Presence Detection (SPD) failure		Major

85E7	CPU2_DIMM_F1 encountered a Serial Presence Detection (SPD) failure		Major
85E8	CPU2_DIMM_F2 encountered a Serial Presence Detection (SPD) failure		Major
85E9	CPU2_DIMM_F3 encountered a Serial Presence Detection (SPD) failure		Major
85EA	CPU2_DIMM_G1 encountered a Serial Presence Detection (SPD) failure		Major
85EB	CPU2_DIMM_G2 encountered a Serial Presence Detection (SPD) failure		Major
85EC	CPU2_DIMM_G3 encountered a Serial Presence Detection (SPD) failure		Major
85ED	CPU2_DIMM_H1 encountered a Serial Presence Detection (SPD) failure		Major
85EE	CPU2_DIMM_H2 encountered a Serial Presence Detection (SPD) failure		Major
85EF	CPU2_DIMM_H3 encountered a Serial Presence Detection (SPD) failure		Major
8604	POST Reclaim of non-critical NVRAM variables		Minor
8605	BIOS Settings are corrupted		Major
8606	NVRAM variable space was corrupted and has been reinitialized		Major
8607	Recovery boot has been initiated. Note: The Primary BIOS image may be corrupted or the system may hang during POST. A BIOS update is required.		Fatal
92A3	Serial port component was not detected		Major
92A9	Serial port component encountered a resource conflict error		Major
A000	TPM device not detected		Minor
A001	TPM device missing or not responding		Minor
A002	TPM device failure		Minor
A003	TPM device failed self-test		Minor
A100	BIOS ACM Error		Major
A421	PCI component encountered a SERR error		Fatal
A5A0	PCI Express component encountered a PERR error		Minor
A5A1	PCI Express component encountered an SERR error		Fatal
A6A0	DXE Boot Services driver: Not enough memory available to shadow a Legacy Option ROM	Please disable OpRom at SETUP to save runtime memory.	Minor

C.1. POST Error Beep Codes

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

Table 80. POST error beep codes

Beeps	Error Message	POST Progress Code	Description
1 short	USB device action	N/A	Short beep sounded whenever USB device is discovered in POST, or inserted or removed during runtime.
1 long	Intel® TXT security violation	AE, AF	System halted because Intel® Trusted Execution Technology detected a potential violation of system security.
3 short	Memory error	Multiple	System halted because a fatal error related to the memory was detected.
3 long and 1 short	CPU mismatch error	E5, E6	System halted because a fatal error related to the CPU family/core/cache mismatch was detected.
2 short	BIOS recovery started	N/A	Recovery boot has been initiated.
4 short	BIOS recovery failed	N/A	Recovery has failed. This typically happens so quickly after recovery is initiated that it sounds like a 2-4 beep code.

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

Table 81. Integrated BMC beep codes

Code	Error Message	Description
1-5-1-2	VR Watchdog Timer	VR controller DC power on sequence not completed in time.
1-5-1-4	Power Supply Status	The system does not power on or unexpectedly power off and a power supply unit (PSU) is present that is an incompatible model with one or more other PSUs in the system
1-5-2-1	CPU Socket Population Error	CPU1 socket is empty or sockets are populated incorrectly – CPU1 must be populated before CPU2
1-5-2-4	MSID Mismatch	MSID mismatch occurs if a processor is installed into a system board that has incompatible power capabilities.
1-5-4-2	Power Fault	DC power unexpectedly lost (power good dropout) – Power unit sensors report power unit failure offset.
1-5-4-4	Power Control Fault	Power good assertion timeout – Power unit sensors report soft power control failure offset.

ix D. System Configuration Table for Thermal Compatibility

configuration compatibility data based on various supported system operating thermal limits. Two tables are identifies supported system configurations while the system is in “normal” operating mode meaning all systems fans are onal. The second table identifies supported system configurations while the system is in a “fan fail” mode meaning rotor is no longer on-line or operational and fan redundancy is lost.

icate support criteria associated with specific configurations identified in the following tables. Each relevant note to y a reference number in the table. Listed notes that are not specified in the table reflect support criteria for a similar intel® Server S2600WF product family, details of which can be found in the Intel® Server System R1000WF Technical

notes:

ne is limited to elevations of 900m or less. Altitudes higher than 900m need to be de-rated to ASHRAE Class 2

undancy, the system must be configured with two power supplies to maintain sufficient cooling. Concurrent system s is not supported.

ccur which may impact system performance. CPU reliability is not impacted.

Module are only supported in the specified base system model configured.

* slot is limited to add-in cards that have air flow requirements of 100 LFM or less. See add-in card specs for air flow

ass 4 support, the following power supply margining is required to meet thermal specifications:

configurations, the power budget must fit within a single power supply rated load and be installed in a dual

ly configurations, the power budget must be sized with 30% margin to single power supply rated load.

el GPGPU cards may have performance impact during ASHRAE Class 2 and Class 3 ambient air excursions.

an only be supported in the bottom add-in card slot on Riser Slot #1 and Riser Slot #2.

ce Free Backup Unit (AXXRMFBUX) can support a case temperature of up to 45°C with the system operating in

with the system operating in a fan fail mode. The case temperature of Intel® Smart RAID Battery (AXXRSBBUX) can

ormal and fan fail mode. Excursions over these specs may result in a reliability impact.

gured with Intel® accessory kits AWFCOPRODUCTAD, AWFCOPRODUCTBKT and A2UL16RISER2 in order to support

GPGPU add-in cards with passive cooling solutions.

y testing was carried out in environmental lab controlled conditions according to ASHRAE standard.

data center environmental temperature and relative humidity levels controls provided by end user.

he system integrator to both consider the thermal configuration matrix and power budget tool documents in order

16. See TPS for thermal limitations.
17. Intel NVME drives used for thermal testing.
18. PSU 2 supported up to max CPU TDP of 165 W.
19. PSU 2 supported up to max CPU TDP of 140 W.
20. PSU 2 supported up to max CPU TDP of 125 W.
21. PSU 2 supported up to max CPU TDP of 105 W.
22. PSU 2 supported up to max CPU TDP of 85 W.
23. AEP electrically/system not supported on the program.
24. During fan fail mode CPU throttling might occur due power limiting feature for PSU2 Unit.
25. PSU 2 supported up to max CPU TDP of 135 W (Apply for Fabric sku only).
26. PSU 2 supported up to max CPU TDP of 113 W (Apply for Fabric sku only).

Table 82. Thermal configuration table for system in “normal” operating mode

Thermal Configuration Matrix Normal Operating Mode "●" = Full Support without limitation "4,5" (Cell with number) = Conditional support with limitation (See notes above) " " (Blank) = No support			Base System SKUs: R2312 WF STANDARD AIR DUCT				Base System SKUs: R2224 WF SATA / NVME 200 GB - 600 GB STANDARD AIR DUCT				Base System SKUs: R2224 WF NVME 800 GB - 2.0 TB STANDARD AIR DUCT				Base System SKUs R2308 WF R2216 WF STANDARD AIR DUCT			
ASHRAE (See Note 1)		Classifications	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4
		Max Ambient	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45° C
		Altitude	900	900	900	900	900	900	900	900	900	900	900	900	900	900	900	900
PS (See Notes 2,6)		1300W AC	●	20 26			20 26	21			20 26	22			18	19 25	20 26	22
		1100W AC	●	●			●	●			●	●			●	●	●	●
CPU TDP	205 W	Intel Xeon Platinum 8180_28C																
		Intel Xeon Platinum 8168_24C													●			
	200 W	Intel Xeon Gold 6154_18C													●			

Intel® Server System R2000WF Product Family Technical Product Specification

	173 W	Intel® Xeon® Platinum 8176F_28C													•	3		
	165 W	Intel® Xeon® Platinum 8176_28C	3				•								•	3		
		Intel® Xeon® Platinum 8170_26C	3				•								•	3		
		Intel® Xeon® Gold 6150_18C	3				•								•	•		
		Intel® Xeon® Platinum 8170F_26C	3				•								•	•		
	160 W	Intel® Xeon® Platinum 8160F_24C	3				•								•	•		
		Intel® Xeon® Gold 6148F_20C	3				•								•	•		
		Intel® Xeon® Gold 6142F_16C	3				•								•	•		
	150 W	Intel® Xeon® Platinum 8164_26C	•				•	3							•	•		
		Intel® Xeon® Platinum 8160_24C	•				•	3							•	•		
		Intel® Xeon® Gold 6148_20C	•				•	3							•	•		
		Intel® Xeon® Gold 6136_12C	•				•	•							•	•		
		Intel® Xeon® Platinum 8158_12C	•				•	•							•	•		
		Intel® Xeon® Gold 6142_16C	•				•	3							•	•		
	140 W	Intel® Xeon® Gold 6132_14C	•	3			•	3			•				•	•	•	3
		Intel® Xeon® Gold 6152_22C	•	•			•	•			•				•	•	•	•
		Intel® Xeon® Gold 6140_18C	•	•			•	•			•				•	•	•	•
	135 W	Intel® Xeon® Gold 6138F_20C	•	•			•	•			3				•	•	•	•
		Intel® Xeon® Gold 6130F_16C	•	•			•	•			3				•	•	•	•
		Intel® Xeon® Gold 6126F_12C	•	•			•	•			3				•	•	•	•
	130 W	Intel® Xeon® Gold 6134_8C	•	•			•	•			•				•	•	•	•
	125 W	Intel® Xeon® Gold 6138_20C	•	•			•	•			•				•	•	•	•
		Intel® Xeon® Gold 6130_16C	•	•			•	•			•				•	•	•	•
		Intel® Xeon® Platinum 8153_16C	•	•			•	•			•				•	•	•	•
		Intel® Xeon® Gold 6126_12C	•	•			•	•			•				•	•	•	•
	115 W	Intel® Xeon® Gold 6128_6C	•	•			•	•			•				•	•	•	•
	113 W	Intel® Xeon® Gold 5117F_14C	•	•			•	•			•	•			•	•	•	•
	105 W	Intel® Xeon® Gold 5122_4C	•	•			•	•			•	•			•	•	•	•

Intel® Server System R2000WF Product Family Technical Product Specification

		Intel® Xeon® Platinum 8156_4C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Gold 5120_14C	•	•			•	•			•	3			•	•	•	•
		Intel® Xeon® Gold 5118_12C	•	•			•	•			•	3			•	•	•	•
	85 W	Intel® Xeon® Gold 5115_10C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4116_12C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4114_10C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4110_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4108_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Bronze 3106_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Bronze 3104_6C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4112_4C	•	•			•	•			•	•			•	•	•	•
Memory Type (See Note 23)		RDIMM-2Rx8, 1Rx4, 1Rx8	•	•			•	•			•	•			•	•	•	•
		RDIMM-DRx4	•	•			•	•			•	•			•	•	•	•
		LRDIMM-QRx4 DDP	•	•			•	•			•	3			•	•	3	3
Add-in Cards (See Note 5)		Riser #1 - Bottom Slot (1U riser and 2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #1 - Middle Slot (2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #1 - Top Slot (2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #2 - Bottom Slot (1U riser and 2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #2 - Middle Slot (2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #2 - Top Slot (2U riser)	•	•			•	•			•	•			•	•	•	•
3rd PCI Riser		Riser #3 - Bottom Slot	•	•			•	•			•	•			•	•	•	5
		Riser #3 - Top Slot	•	•			•	•			•	•			•	•	•	5
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	•	•			•	•			•	•			•	•	•	•
		16 port tri mode mezzanine	•	•			•	•			•	•			•	•	•	•
		Dual 10GBase-T Taylor Canal	•	•			•	•			•	•			•	•	•	•
		Dual SFP+ Cascade Canal	•	•			•	•			•	•			•	•	•	•
		Quad SFP+ Cascade Canal	•	•			•	•			•	•			•	•	•	•
		1GBase-T Quad Main Canal	•	•			•	•			•	•			•	•	•	•
Battery Backup (See Note 9)		BBU (rated to 45C)					•				•				•	•	•	•
		Supercap (rated to 55C)	•	•			•	•			•	•			•	•	•	•

Intel® Server System R2000WF Product Family Technical Product Specification

		Cache Offload Module (rated to 55C)	•	•			•	•			•	•			•	•	•	•
Internal SSD		Rated to 60C					•				•				•	•	•	•
		Rated to 70C	•	•			•	•			•	•			•	•	•	•
Rear SSD		Rated to 60C	•				•				•				•	•		
		Rated to 70C	•	•			•	•			•	•			•	•	•	•
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 17)		1600GB/2TB									•	•			•	•	•	•
		800GB									•	•			•	•	•	•
		600GB					•	•			•	•			•	•	•	•
		400GB					•	•			•	•			•	•	•	•
		200GB					•	•			•	•			•	•	•	•
SATA HDD		Up to 2 TB	•	•			•	•			•	•			•	•	•	•
PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		1600GB/2TB	•				•								•	8	8	
		800GB	•				•								•	8	8	
		600GB	•	•			•	•			•	•			•	•	8	8
		400GB	•	•			•	•			•	•			•	•	8	8
		200GB	•	•			•	•			•	•			•	•	•	•
Intel® Xeon Phi™ (See Notes 7, 10)		Active Cooling up to 300W													•	•		
		Active Cooling up to 225W													•	•	•	•
		Intel® Xeon Phi™ w/Passive Cooling up to 225W																
		Intel® Xeon Phi™ w/Passive Cooling up to 245W																
		Intel® Xeon Phi™ w/Passive Cooling up to 300W																
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	14	14			14	14			14	14			•	•	14	14
		Non-Intel M.2	16	16			16	16			16	16			•	•	16	16
M.2 Light Workload (See Note 15)		Intel 120G and 80G SATA	•	•			•	14			•	14			•	•	•	14
		Non-Intel M.2	•	•			•	16			•	16			•	•	•	16

Thermal Configuration Matrix Normal Operating Mode "●" = Full Support without limitation "4,5" (Cell with number) = Conditional support with limitation (See notes above) " " (Blank) = No support			Base System SKUs: R2208 WF STANDARD AIRDUCT				Xeon PHI Computing Optimized: R2308 WF R2208 WF +8 AWFCOPRODUCTAD AIRDUCT				Xeon PHI Computing Optimized: R2208WF AWFCOPRODUCTAD AIRDUCT			
ASHRAE (See Note 1)		Classifications	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4
		Max Ambient	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C
		Altitude	900	900	900	900	900	900	900	900	900	900	900	900
PS (See Notes 2,6)		1300W AC	●	18	19 25	20 26	●	19 25	20 26	21	●	19 25	20 26	21
		1100W AC	●	●	●	●	●	●	●	●	●	●	●	●
CPU TDP	205 W	Intel Xeon Platinum 8180_28C	●	●	●									
		Intel Xeon Platinum 8168_24C	●	●	●						●			
	200 W	Intel Xeon Gold 6154_18C	●	●	●						●			
	173 W	Intel® Xeon® Platinum 8176F_28C	●	●	●	●					●	3		
	165 W	Intel® Xeon® Platinum 8176_28C	●	●	●	●	3				●	3		
		Intel® Xeon® Platinum 8170_26C	●	●	●	●	●				●	3		
		Intel® Xeon® Gold 6150_18C	●	●	●	●	●				●	●		
		Intel® Xeon® Platinum 8170F_26C	●	●	●	●	●				●	●		
	160 W	Intel® Xeon® Platinum 8160F_24C	●	●	●	●	●				●	●		
		Intel® Xeon® Gold 6148F_20C	●	●	●	●	●				●	●		
		Intel® Xeon® Gold 6142F_16C	●	●	●	●	●				●	●		
	150 W	Intel® Xeon® Platinum 8164_26C	●	●	●	●	●	3			●	●	3	3
		Intel® Xeon® Platinum 8160_24C	●	●	●	●	●	3			●	●	3	3
		Intel® Xeon® Gold 6148_20C	●	●	●	●	●	3			●	●	3	3

Intel® Server System R2000WF Product Family Technical Product Specification

		Intel® Xeon® Gold 6136_12C	•	•	•	•	•	•			•	•	•	•
		Intel® Xeon® Platinum 8158_12C	•	•	•	•	•	•			•	•	•	•
		Intel® Xeon® Gold 6142_16C	•	•	•	•	•	3			•	•	3	3
	140 W	Intel® Xeon® Gold 6132_14C	•	•	•	•	•	•	3		•	•	•	•
		Intel® Xeon® Gold 6152_22C	•	•	•	•	•	•	3		•	•	•	•
		Intel® Xeon® Gold 6140_18C	•	•	•	•	•	•	•		•	•	•	•
	135 W	Intel® Xeon® Gold 6138F_20C	•	•	•	•	•	•	•		•	•	•	•
		Intel® Xeon® Gold 6130F_16C	•	•	•	•	•	•	•		•	•	•	•
		Intel® Xeon® Gold 6126F_12C	•	•	•	•	•	•	•		•	•	•	•
	130 W	Intel® Xeon® Gold 6134_8C	•	•	•	•	•	•	•		•	•	•	•
	125 W	Intel® Xeon® Gold 6138_20C	•	•	•	•	•	•	•	3	•	•	•	•
		Intel® Xeon® Gold 6130_16C	•	•	•	•	•	•	•	3	•	•	•	•
		Intel® Xeon® Platinum 8153_16C	•	•	•	•	•	•	•	3	•	•	•	•
		Intel® Xeon® Gold 6126_12C	•	•	•	•	•	•	•	•	•	•	•	•
	115 W	Intel® Xeon® Gold 6128_6C	•	•	•	•	•	•	•	•	•	•	•	•
	113 W	Intel® Xeon® Gold 5117F_14C	•	•	•	•	•	•	•	•	•	•	•	•
	105 W	Intel® Xeon® Gold 5122_4C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Platinum 8156_4C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Gold 5120_14C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Gold 5118_12C	•	•	•	•	•	•	•	•	•	•	•	•
	85 W	Intel® Xeon® Gold 5115_10C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Silver 4116_12C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Silver 4114_10C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Silver 4110_8C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Silver 4108_8C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Bronze 3106_8C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Bronze 3104_6C	•	•	•	•	•	•	•	•	•	•	•	•
		Intel® Xeon® Silver 4112_4C	•	•	•	•	•	•	•	•	•	•	•	•

Memory Type (See Note 23)		RDIMM-2Rx8, 1Rx4, 1Rx8	•	•	•	•	•	•	•	•	•	•	•	•
		RDIMM-DRx4	•	•	•	•	•	•	•	•	•	•	•	•
		LRDIMM-QRx4 DDP	•	•	•	•	•	•	3	3	•	•	•	•
Add-in Cards (See Note 5)		Riser #1 - Bottom Slot (1U riser and 2U riser)	•	•	•	•	•	•	•	•	•	•	•	•
		Riser #1 - Middle Slot (2U riser)	•	•	•	•								
		Riser #1 - Top Slot (2U riser)	•	•	•	•	•	•	•	•	•	•	•	•
		Riser #2 - Bottom Slot (1U riser and 2U riser)	•	•	•	•	•	•	•	•	•	•	•	•
		Riser #2 - Middle Slot (2U riser)	•	•	•	•								
		Riser #2 - Top Slot (2U riser)	•	•	•	•	•	•	•	•	•	•	•	•
3rd PCI Riser		Riser #3 - Bottom Slot	•	•	•	5	•	•	5	5	•	•	5	5
		Riser #3 - Top Slot	•	•	•	5	•	•	5	5	•	•	5	5
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	•	•	•	•	•	•	•	•	•	•	•	•
		16 port tri mode mezzanine	•	•	•	•	•	•	•	•	•	•	•	•
		Dual 10GBase-T Taylor Canal	•	•	•	•	•	•	•	•	•	•	•	
		Dual SFP+ Cascade Canal	•	•	•	•	•	•	•	•	•	•	•	
		Quad SFP+ Cascade Canal	•	•	•	•	•	•	•	•	•	•	•	
		1GBase-T Quad Main Canal	•	•	•	•	•	•	•	•	•	•	•	
Battery Backup (See Note 9)		BBU (rated to 45C)	•	•	•	•	•	•	•	•	•	•		
		Supercap (rated to 55C)	•	•	•	•	•	•	•	•	•	•		
		Cache Offload Module (rated to 55C)	•	•	•	•	•	•	•	•	•	•		
Internal SSD		Rated to 60C	•	•	•	•	•	•	•	•	•	•		
		Rated to 70C	•	•	•	•	•	•	•	•	•	•	•	
Rear SSD		Rated to 60C	•	•			•	•	•	•	•	•		
		Rated to 70C	•	•	•	•	•	•	•	•	•	•	•	
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 17)		1600GB/2TB	•	•	•	•					•			
		800GB	•	•	•	•					•	•	•	•
		600GB	•	•	•	•					•	•	•	•
		400GB	•	•	•	•					•	•	•	•
		200GB	•	•	•	•					•	•	•	•
SATA HDD		Up to 2 TB	•	•	•	•	•	•	•	•	•	•	•	•
		1600GB/2TB	•	8	8						8	8	8	

Intel® Server System R2000WF Product Family Technical Product Specification

PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		800GB	•	8	8						8	8	8	
		600GB	•	•	8	8					•	•	8	8
		400GB	•	•	8	8					•	•	8	8
		200GB	•	•	•	•					•	•	•	•
Intel® Xeon Phi™ (See Notes 7, 10)		Active Cooling up to 300W	•	•										
		Active Cooling up to 225W	•	•	•	•								
		Intel® Xeon Phi™ w/Passive Cooling up to 225W					•	•	•	•	•	•	•	•
		Intel® Xeon Phi™ w/Passive Cooling up to 245W					7	7	7	7	•	7	7	7
		Intel® Xeon Phi™ w/Passive Cooling up to 300W					7	7	7	7	7	7	7	7
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	•	•	14	14	•	14	14	14	•	•	14	14
		Non-Intel M.2	•	•	16	16	•	16	16	16	•	•	16	16
M.2 Light Workload (See Note 15)		Intel 120G and 80G SATA	•	•	•	14	•	•	•	14	•	•	•	14
		Non-Intel M.2	•	•	•	16	•	•	•	16	•	•	•	16

Table 83. Thermal configuration table for system in "fan fail" operating mode

Thermal Configuration Matrix Fan Fail Operating Mode			Base System SKUs: R2312 WF STANDARD AIR DUCT				Base System SKUs: R2224 WF SATA / NVME 200 GB - 600 GB STANDARD AIR DUCT				Base System SKUs: R2224 WF NVME 800 GB - 2.0 TB STANDARD AIR DUCT				Base System SKUs R2308 WF R2208 WF +8 STANDARD AIR DUCT			
"●" = Full Support without limitation "4,5" (Cell with number) = Conditional support with limitation (See notes above) " " (Blank) = No support																		
ASHRAE (See Note 1)		Classifications	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4
		Max Ambient	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45° C
		Altitude	900	900	900	900	900	900	900	900	900	900	900	900	900	900	900	900
PS (See Notes 2,6)		1300W AC	2, 21	2			2, 21	2			2, 22	2			2, 20	2, 22	2	2
		1100W AC	2	2			2	2			2	2			2	2	2	2
CPU TDP	205 W	Intel Xeon Platinum 8180_28C													24			
		Intel Xeon Platinum 8168_24C													24			
	200 W	Intel Xeon Gold 6154_18C													24			
	173 W	Intel® Xeon® Platinum 8176F_28C													24	24		
	165 W	Intel® Xeon® Platinum 8176_28C	24				24								24	24		
		Intel® Xeon® Platinum 8170_26C	24				24								24	24		
		Intel® Xeon® Gold 6150_18C	24				24								24	24		
		Intel® Xeon® Platinum 8170F_26C	24				24								24	24		
	160 W	Intel® Xeon® Platinum 8160F_24C	24				24								24	24		
		Intel® Xeon® Gold 6148F_20C	24				24								24	24		
		Intel® Xeon® Gold 6142F_16C	24				24								24	24		
	150 W	Intel® Xeon® Platinum 8164_26C	24				24	24							24	24		
		Intel® Xeon® Platinum 8160_24C	24				24	24							24	24		

Intel® Server System R2000WF Product Family Technical Product Specification

		Intel® Xeon® Gold 6148_20C	24				24	24							24	24		
		Intel® Xeon® Gold 6136_12C	24				24	24							24	24		
		Intel® Xeon® Platinum 8158_12C	24				24	24							24	24		
		Intel® Xeon® Gold 6142_16C	24				24	24							24	24		
	140 W	Intel® Xeon® Gold 6132_14C	24	24			24	24			24				24	24	24	24
		Intel® Xeon® Gold 6152_22C	24	24			24	24			24				24	24	24	24
		Intel® Xeon® Gold 6140_18C	24	24			24	24			24				24	24	24	24
	135 W	Intel® Xeon® Gold 6138F_20C	24	24			24	24			24				•	24	24	24
		Intel® Xeon® Gold 6130F_16C	24	24			24	24			24				•	24	24	24
		Intel® Xeon® Gold 6126F_12C	24	24			24	24			24				•	24	24	24
	130 W	Intel® Xeon® Gold 6134_8C	•	24			•	24			24				•	24	24	24
	125 W	Intel® Xeon® Gold 6138_20C	•	24			•	24			24				•	24	24	24
		Intel® Xeon® Gold 6130_16C	•	24			•	24			24				•	24	24	24
		Intel® Xeon® Platinum 8153_16C	•	24			•	24			24				•	24	24	24
		Intel® Xeon® Gold 6126_12C	•	24			•	24			24				•	24	24	24
	115 W	Intel® Xeon® Gold 6128_6C	•	24			•	24			24				•	•	24	24
	113 W	Intel® Xeon® Gold 5117F_14C	•	24			•	24			24	24			•	•	24	24
	105 W	Intel® Xeon® Gold 5122_4C	•	•			•	•			•	24			•	•	24	24
		Intel® Xeon® Platinum 8156_4C	•	•			•	•			•	24			•	•	24	24
		Intel® Xeon® Gold 5120_14C	•	3			•	•			•	24			•	•	24	24
		Intel® Xeon® Gold 5118_12C	•	3			•	•			•	24			•	•	24	24
	85 W	Intel® Xeon® Gold 5115_10C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4116_12C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4114_10C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4110_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Silver 4108_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Bronze 3106_8C	•	•			•	•			•	•			•	•	•	•
		Intel® Xeon® Bronze 3104_6C	•	•			•	•			•	•			•	•	•	•

Intel® Server System R2000WF Product Family Technical Product Specification

		Intel® Xeon® Silver 4112_4C	•	•			•	•			•	•			•	•	•	•
Memory Type (See Note 23)		RDIMM-2Rx8, 1Rx4, 1Rx8	•	•			•	•			•	•			•	•	•	•
		RDIMM-DRx4	•	•			•	•			•	•			•	•	•	•
		LRDIMM-QRx4 DDP	3	3			3	3			3	3			•	3	3	3
Add-in Cards (See Note 5)		Riser #1 - Bottom Slot (1U riser and 2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #1 - Middle Slot (2U riser)	•	5			•	5			•	5			•	5	5	5
		Riser #1 - Top Slot (2U riser)	•	5			•	5			•	5			•	5	5	5
		Riser #2 - Bottom Slot (1U riser and 2U riser)	•	•			•	•			•	•			•	•	•	•
		Riser #2 - Middle Slot (2U riser)	•	5			•	5			•	5			•	5	5	5
		Riser #2 - Top Slot (2U riser)	•	5			•	5			•	5			•	5	5	5
3rd PCI Riser		Riser #3 - Bottom Slot	•	5			•	5			•	5			•	5	5	5
		Riser #3 - Top Slot	•	5			•	5			•	5			•	5	5	5
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	•	•			•	•			•	•			•	•	•	•
		16 port tri mode mezzanine	•	4			•	4			•	4			•	•	•	•
		Dual 10GBase-T Taylor Canal	•	•			•	•			•	•			•	•	•	•
		Dual SFP+ Cascade Canal	•	•			•	•			•	•			•	•	•	•
		Quad SFP+ Cascade Canal	•	•			•	•			•	•			•	•	•	•
		1GBase-T Quad Main Canal	•	•			•	•			•	•			•	•	•	•
Battery Backup (See Note 9)		BBU (rated to 45C)					•				•				•	•	•	•
		Supercap (rated to 55C)	•	•			•	•			•	•			•	•	•	•
		Cache Offload Module (rated to 55C)	•	•			•	•			•	•			•	•	•	•
Internal SSD		Rated to 60C					•				•				•			
		Rated to 70C	•	•			•	•			•	•			•	•	•	•
Rear SSD		Rated to 60C	•				•				•				•	•	•	•
		Rated to 70C	•	•			•	•			•	•			•	•	•	•
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 17)		1600GB/2TB									•	•			•	•	•	•
		800GB									•	•			•	•	•	•
		600GB					•	•			•	•			•	•	•	•
		400GB					•	•			•	•			•	•	•	•
		200GB					•	•			•	•			•	•	•	•
SATA HDD		Up to 2 TB	•	•			•	•			•	•			•	•	•	•

PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		1600GB/2TB													8	8	8	
		800GB													8	8	8	
		600GB	•	8			•	8			•	8			•	•	8	8
		400GB	•	8			•	8			•	8			•	•	8	8
		200GB	•	•			•	•			•	•			•	•	•	•
Intel® Xeon Phi™ (See Notes 7, 10)		Active Cooling up to 300W													7	7		
		Active Cooling up to 225W													7	7	7	7
		Intel® Xeon Phi™ w/Passive Cooling up to 225W																
		Intel® Xeon Phi™ w/Passive Cooling up to 245W																
		Intel® Xeon Phi™ w/Passive Cooling up to 300W																
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	14	14			14	14			14	14			14	14	14	14
		Non-Intel M.2	16	16			16	16			16	16			16	16	16	16
M.2 Light Workload (See Note 15)		Intel 120G and 80G SATA	14	14			14	14			14	14			14	14	14	14
		Non-Intel M.2	16	16			16	16			16	16			16	16	16	16

Thermal Configuration Matrix Fan Fail Operating Mode "•" = Full Support without limitation "4,5" (Cell with number) = Conditional support with limitation (See notes above) " " (Blank) = No support			Base System SKUs: R2208 WF STANDARD AIRDUCT				Xeon PHI Computing Optimized: R2308 WF R2208 WF +8 AWFCOPRODUCTAD AIRDUCT				Xeon PHI Computing Optimized: R2208WF AWFCOPRODUCTAD AIRDUCT			
ASHRAE (See Note 1)		Classifications	27C	A2	A3	A4	27C	A2	A3	A4	27C	A2	A3	A4
		Max Ambient	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C	27°C (1)	35°C	40°C	45°C
		Altitude	900	900	900	900	900	900	900	900	900	900	900	900
PS		1300W AC	2, 18	2, 19	2, 20	2, 21	2, 21	2	2	2	2, 20	2, 22	2	2

Intel® Server System R2000WF Product Family Technical Product Specification

(See Notes 2,6)		1100W AC	•	•	•	•	2	2	2	2	2	2	2	2
CPU TDP	205 W	Intel Xeon Platinum 8180_28C	•	3	24						24			
		Intel Xeon Platinum 8168_24C	•	3	24						24			
	200 W	Intel Xeon Gold 6154_18C	•	3	24						24			
	173 W	Intel® Xeon® Platinum 8176F_28C	•	3	24	24					24	24		
	165 W	Intel® Xeon® Platinum 8176_28C	•	•	•	24	24				24	24		
		Intel® Xeon® Platinum 8170_26C	•	•	•	24	24				24	24		
		Intel® Xeon® Gold 6150_18C	•	•	•	24	24				24	24		
		Intel® Xeon® Platinum 8170F_26C	•	•	•	24	24				24	24		
	160 W	Intel® Xeon® Platinum 8160F_24C	•	•	•	24	24	24			24	24		
		Intel® Xeon® Gold 6148F_20C	•	•	•	24	24	24			24	24		
		Intel® Xeon® Gold 6142F_16C	•	•	•	24	24	24			24	24		
	150 W	Intel® Xeon® Platinum 8164_26C	•	•	•	•	24	24			24	24	24	24
		Intel® Xeon® Platinum 8160_24C	•	•	•	•	24	24			24	24	24	24
		Intel® Xeon® Gold 6148_20C	•	•	•	•	24	24			24	24	24	24
		Intel® Xeon® Gold 6136_12C	•	•	•	•	24	24			24	24	24	24
		Intel® Xeon® Platinum 8158_12C	•	•	•	•	24	24			24	24	24	24
		Intel® Xeon® Gold 6142_16C	•	•	•	•	24	24			24	24	24	24
	140 W	Intel® Xeon® Gold 6132_14C	•	•	•	•	24	24	24		•	24	24	24
		Intel® Xeon® Gold 6152_22C	•	•	•	•	24	24	24		•	24	24	24
		Intel® Xeon® Gold 6140_18C	•	•	•	•	24	24	24		•	24	24	24
	135 W	Intel® Xeon® Gold 6138F_20C	•	•	•	•	24	24	24		•	24	24	24
		Intel® Xeon® Gold 6130F_16C	•	•	•	•	24	24	24		•	24	24	24
		Intel® Xeon® Gold 6126F_12C	•	•	•	•	24	24	24		•	24	24	24
	130 W	Intel® Xeon® Gold 6134_8C	•	•	•	•	24	24	24		•	24	24	24
	125 W	Intel® Xeon® Gold 6138_20C	•	•	•	•	•	24	24	24	•	•	24	24
		Intel® Xeon® Gold 6130_16C	•	•	•	•	•	24	24	24	•	•	24	24
		Intel® Xeon® Platinum 8153_16C	•	•	•	•	•	24	24	24	•	•	24	24

Intel® Server System R2000WF Product Family Technical Product Specification

		Intel® Xeon® Gold 6126_12C	•	•	•	•	•	24	24	24	•	•	24	24
	115 W	Intel® Xeon® Gold 6128_6C	•	•	•	•	•	24	24	24	•	•	24	24
	113 W	Intel® Xeon® Gold 5117F_14C	•	•	•	•	•	24	24	24	•	•	24	24
	105 W	Intel® Xeon® Gold 5122_4C	•	•	•	•	•	24	24	24	•	•	•	24
		Intel® Xeon® Platinum 8156_4C	•	•	•	•	•	•	24	24	•	•	•	24
		Intel® Xeon® Gold 5120_14C	•	•	•	•	•	•	24	24	•	•	•	24
		Intel® Xeon® Gold 5118_12C	•	•	•	•	•	•	24	24	•	•	•	24
	85 W	Intel® Xeon® Gold 5115_10C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Silver 4116_12C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Silver 4114_10C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Silver 4110_8C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Silver 4108_8C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Bronze 3106_8C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Bronze 3104_6C	•	•	•	•	•	•	•	24	•	•	•	•
		Intel® Xeon® Silver 4112_4C	•	•	•	•	•	•	•	24	•	•	•	•
Memory Type (See Note 23)		RDIMM-2Rx8,1Rx4, 1Rx8	•	•	•	•	•	•	•	•	•	•	•	•
		RDIMM-DRx4	•	•	•	•	•	•	•	•	•	•	•	•
		LRDIMM-QRx4 DDP	•	•	3	3	•	3	3	3	•	•	3	3
Add-in Cards (See Note 5)		Riser #1 - Bottom Slot (1U riser and 2U riser)	•	•	•	•	•	•	5	5	•	•	5	5
		Riser #1 - Middle Slot (2U riser)	•	•	•	•								
		Riser #1 - Top Slot (2U riser)	•	•	•	•	5	5	5	5	5	5	5	5
		Riser #2 - Bottom Slot (1U riser and 2U riser)	•	•	•	•	•	•	•	•	•	•	•	•
		Riser #2 - Middle Slot (2U riser)	•	•	•	•								
		Riser #2 - Top Slot (2U riser)	•	•	•	•	5	5	5	5	5	5	5	5
3rd PCI Riser		Riser #3 - Bottom Slot	•	•	•	5	5	5	5	5	5	5	5	5
		Riser #3 - Top Slot	•	•	•	5	5	5	5	5	5	5	5	5
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	•	•	•	•	•	•	•	•	•	•	•	•
		16 port tri mode mezzanine	•	•	•	•	•	•			•	•		
		Dual 10GBase-T Taylor Canal	•	•	•	•	•	•	•	•	•	•	•	•

Intel® Server System R2000WF Product Family Technical Product Specification

		Dual SFP+ Cascade Canal	•	•	•	•	•	•			•	•		
		Quad SFP+ Cascade Canal	•	•	•	•	•	•			•	•		
		1GBase-T Quad Main Canal	•	•	•	•	•	•	•	•	•	•	•	•
Battery Backup (See Note 9)		BBU (rated to 45C)	•	•	•	•	•	•			•	•		
		Supercap (rated to 55C)	•	•	•	•	•	•			•	•		
		Cache Offload Module (rated to 55C)	•	•	•	•	•	•			•	•		
Internal SSD		Rated to 60C	•	•	•	•	•	•			•	•		
		Rated to 70C	•	•	•	•	•	•			•	•		
Rear SSD		Rated to 60C	•	•			•	•			•	•		
		Rated to 70C	•	•	•	•	•	•			•	•		
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 17)		1600GB/2TB	•	•	•	•					•			
		800GB	•	•	•	•					•		•	•
		600GB	•	•	•	•					•	•	•	•
		400GB	•	•	•	•					•	•	•	•
		200GB	•	•	•	•					•	•	•	•
SATA HDD		Up to 2 TB	•	•	•	•	•	•	•	•	•	•	•	•
PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		1600GB/2TB	•	8	8						8	8		
		800GB	•	8	8						8	8		
		600GB	•	•	8	8					•	•	8	8
		400GB	•	•	8	8					•	•	8	8
		200GB	•	•	•	•					•	•	8	8
Intel® Xeon Phi™ (See Notes 7, 10)		Active Cooling up to 300W	•	•										
		Active Cooling up to 225W	•	•	•	•								
		Intel® Xeon Phi™ w/Passive Cooling up to 225W					7	7	7	7	7	7	7	7
		Intel® Xeon Phi™ w/Passive Cooling up to 245W					7	7	7	7	7	7	7	7
		Intel® Xeon Phi™ w/Passive Cooling up to 300W					7	7	7	7	7	7	7	7
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	•	14	14	14	14	14	14	14	14	14	14	14
		Non-Intel M.2	16	16	16	16	16	16	16	16	16	16	16	16
M.2 Light Workload (See Note 15)		Intel 120G and 80G SATA	•	14	14	14	14	14	14	14	14	14	14	14
		Non-Intel M.2	16	16	16	16	16	16	16	16	16	16	16	16

Appendix E. System Cable Routing Diagrams




Figure 90. R2308WFxxx system cable routing diagram






Figure 91. R2208WFxxx system cable routing diagram

Intel Confidential



- Power Cable
- SAS/SATA Cable
- I2C Cable
- Front Control Panel and I/O Cable

Figure 92. R2312WFxxx system cable routing diagram



- Power Cable
- SAS/SATA Cable
- 12C Cable
- Front Control Panel and I/O Cable
- OcuLink

2UWFP0156

Figure 93. R2208WFxxx + 8 drives system cable routing diagram




Figure 94. R2224WFxxx system cable routing diagram

Appendix F. Statement of Volatility

The tables in this section are used to identify the volatile and non-volatile memory components for system boards used within the Intel® Server System R2000WF product family.

The tables provide the following data for each identified component.

- **Component Type:** Three types of components are on the server board assembly:
 - o **Non-volatile:** Non-volatile memory is persistent, and is not cleared when power is removed from the system. Non-volatile memory must be erased to clear data. The exact method of clearing these areas varies by the specific component. Some areas are required for normal operation of the server, and clearing these areas may render the server board inoperable
 - o **Volatile:** Volatile memory is cleared automatically when power is removed from the system.
 - o **Battery powered RAM:** Battery powered RAM is similar to volatile memory, but is powered by a battery on the server board. Data in battery powered RAM is persistent until the battery is removed from the server board.
- **Size:** Size of each component in bits, Kbits, Mbits, bytes, kilobytes (KB), or megabytes (MB).
- **Board Location:** Board location is the physical location of each component corresponding to information on the server board silkscreen.
- **User Data:** The flash components on the server boards do not store user data from the operating system. No operating system level data is retained in any listed components after AC power is removed. The persistence of information written to each component is determined by its type as described in the table.

Each component stores data specific to its function. Some components may contain passwords that provide access to that device's configuration or functionality. These passwords are specific to the device and are unique and unrelated to operating system passwords. The specific components that may contain password data are:

- o **BIOS:** The server board BIOS provides the capability to prevent unauthorized users from configuring BIOS settings when a BIOS password is set. This password is stored in BIOS flash, and is only used to set BIOS configuration access restrictions.
- o **BMC:** The server boards support an Intelligent Platform Management Interface (IPMI) 2.0 conformant baseboard management controller (BMC). The BMC provides health monitoring, alerting and remote power control capabilities for the Intel® Server Board. The BMC does not have access to operating system level data.

The BMC supports the capability for remote software to connect over the network and perform health monitoring and power control. This access can be configured to require authentication by a password. If configured, the BMC maintains user passwords to control this access. These passwords are stored in the BMC flash.

Table 84. Intel® Server Board S2600WF (iPN – H48104-xxx)

Component Type	Size	Board location	User Data	Name
Non-Volatile	64MB	U3D1	No(BIOS)	BIOS Flash
Non-Volatile	64MB	U1D1	No(FW)	BMC Flash
Non-Volatile	4MB	U5M1	No	10GB NIC EEPROM (S2600WTTR)
Non-Volatile	N/A	U1E3	No	CPLD
Volatile	4GB	U1D2	No	BMC SDRAM
Non-Volatile	8GB	U8N1	No	BMC eMMC

Table 85. 2U 3-Slot PCIe* Riser Card (iPN – H20087-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	U7B1	Yes	FRU

Table 86. 2U 2-Slot PCIe* Riser Card (iPN – H20078-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	U7B1	Yes	FRU

Table 87. Riser Slot #3 Low Profile PCIe* Riser Card (iPN – G94347-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	U2	Yes	FRU

Table 88. Common Front Panel Board (iPN – H29366-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	U1A1	Yes	PSOC / Microcontroller

Table 89. Storage System Model Mini- Front Panel Board (iPN – G28538-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	U4M1	Yes	FRU

Table 90. 2U 8 x 3.5" Hot Swap Back Plane (iPN – G97160-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	U5A1	Yes	PSOC / Microcontroller / FRU
Non-Volatile	1024x8	U3	No	SAS Re-Driver Settings

Table 91. 2U 12 x 3.5" Hot Swap Back Plane (iPN – H88392-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	U6L3	Yes	PSOC / Microcontroller / FRU
Non-Volatile	16384x8	U4L1	Yes	PSOC / Microcontroller / FRU
Non-Volatile	1024x8	U7N1	No	SAS Re-Driver Settings

Table 92. 2U 8 x 2.5" SAS Hot Swap Back Plane (iPN – H95877-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	U3A1	Yes	PSOC / Microcontroller / FRU

**Table 93. Intel® 2U 8 x 2.5" Dual Port Hot Swap Back Plane Accessory Kit (iPC – A2U8X25S3DPDK)
2U 8 x 2.5" Dual Port Hot Swap Back Plane (iPN – G97166-xxx)**

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	U3A2	Yes	PSOC / Microcontroller / FRU
Non-Volatile	N/A	U3A1	No	CPLD Code

**Table 94. Intel® 2U 2 x 2.5" Rear Hot Swap Back Plane Accessory Kit (iPC – A2UREARHSDK)
2U 2 x 2.5" Rear Hot Swap Back Plane (iPN – G94339-xxx)**

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	U1	Yes	PSOC / Microcontroller / FRU

Table 95. Intel® Remote Management Module Lite Accessory Option (iPC – AXXRMM4LITE2)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	1Mbit	U2B1	No	RMM Programming

Glossary

Term	Definition
BMC	Baseboard Management Controller
BIOS	Basic Input/Output System
CMOS	Complementary Metal-oxide-semiconductor
CPU	Central Processing Unit
DDR4	Double Data Rate 4th edition
DIMM	Dual In-line Memory Module
DPC	DIMMs per Channel
EDS	External Design Specification
EPS	External Product Specification
FP	Front Panel
FRB	Fault Resilient Boot
FRU	Field Replaceable Unit
GPGPU	General Purpose Graphic Processing Unit
I2C	Inter-integrated Circuit bus
iPC	Intel Product Code
LED	Light Emitting Diode
LFM	Linear Feet per Minute – Air Flow measurement
LPC	Low-pin Count
LRDIMM	Load Reduced DIMM
LSB	Least Significant Bit
MSB	Most Significant Bit
MTBF	Mean Time Between Failure
NIC	Network Interface Card
NMI	Non-maskable Interrupt
OCuLink	Optical Copper Link
OVP	Over-voltage Protection
PCI	Peripheral Component Interconnect
PCB	Printed Circuit Board
PCIe*	Peripheral Component Interconnect Express*
PCI-X	Peripheral Component Interconnect Extended
PFC	Power Factor Correction
POST	Power-on Self-Test
PSU	Power Supply Unit
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RDIMM	Registered DIMM
ROC	RAID On Chip
SAS	Serial Attached SCSI
SATA	Serial Advanced Technology Attachment
SCA	Single Connector Attachment
SCSI	Small Computer System Interface
SDR	Sensor Data Record
SSD	Solid State Device
TPS	Technical Product Specification

Intel® TXT	Intel® Trusted Execution Technology for servers
VLSI	Very Large Scale Integration
VSB	Voltage Standby
Intel® VROC	Intel® Virtual RAID on CPU

Mouser Electronics

Authorized Distributor

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

[Intel:](#)

[R2208WF0ZS](#) [R2208WFTZS](#) [R2224WFTZS](#) [R2308WFTZS](#) [R2312WF0NP](#) [R2312WFTZS](#)



Компания «ЭлектроПласт» предлагает заключение долгосрочных отношений при поставках импортных электронных компонентов на взаимовыгодных условиях!

Наши преимущества:

- Оперативные поставки широкого спектра электронных компонентов отечественного и импортного производства напрямую от производителей и с крупнейших мировых складов;
- Поставка более 17-ти миллионов наименований электронных компонентов;
- Поставка сложных, дефицитных, либо снятых с производства позиций;
- Оперативные сроки поставки под заказ (от 5 рабочих дней);
- Экспресс доставка в любую точку России;
- Техническая поддержка проекта, помощь в подборе аналогов, поставка прототипов;
- Система менеджмента качества сертифицирована по Международному стандарту ISO 9001;
- Лицензия ФСБ на осуществление работ с использованием сведений, составляющих государственную тайну;
- Поставка специализированных компонентов (Xilinx, Altera, Analog Devices, Intersil, Interpoint, Microsemi, Aeroflex, Peregrine, Syfer, Eurofarad, Texas Instrument, Miteq, Cobham, E2V, MA-COM, Hittite, Mini-Circuits, General Dynamics и др.);

Помимо этого, одним из направлений компании «ЭлектроПласт» является направление «Источники питания». Мы предлагаем Вам помощь Конструкторского отдела:

- Подбор оптимального решения, техническое обоснование при выборе компонента;
- Подбор аналогов;
- Консультации по применению компонента;
- Поставка образцов и прототипов;
- Техническая поддержка проекта;
- Защита от снятия компонента с производства.



Как с нами связаться

Телефон: 8 (812) 309 58 32 (многоканальный)

Факс: 8 (812) 320-02-42

Электронная почта: org@eplast1.ru

Адрес: 198099, г. Санкт-Петербург, ул. Калинина, дом 2, корпус 4, литера А.