Intel® Server System R1000WF Product Family

Technical Product Specification

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

Revision 1.0

July 2017

Intel® Server Products and Solutions

<Blank Page>

Revision History

Date	Revision Number	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, Xeon, and Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Table of Contents

List of Figures

List of Tables

1. Introduction

This Technical Product Specification (TPS) provides system level information for the Intel® Server System R1000WF 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 SKUs made available.

For more additional product information, the documents listed in Table 1 should also be referenced.

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 your local 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 Server Board Use Disclaimer

Intel Corporation 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 their published operating or non-operating limits.

1.3 Product Errata

Shipping product may have features or functionality that may 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 will be published in the Monthly Specification Update for the given product family which can be downloaded from the following Intel web site:

http://www.intel.com/support

2. Server System Family Overview

The 1U server platforms within the Intel® Server System R1000WF 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.

Each building block option or integrated system within this product family is configured around the following Intel server board: Intel® Server Board S2600WF product family.

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

Note: Not all server board features identified in Table 2 will be supported in the 1U chassis. See Table 2 for specific 1U system features.

Table 2. Intel® Server Board S2600WF Product Family Feature Set

	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) o 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) o 6 Memory Channels per processor / 2 DIMMs per Channel • Registered DDR4 (RDIMM), Load Reduced DDR4 (LRDIMM) • Memory Capacity o Up to 1.5TB for Gold and Platinum CPUs; Up to 768GB for Silver and Bronze CPUs • Memory data transfer rates: o 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
Onboard 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

	Intel® Server Board S2600WF Product Family	
Feature	**S2600WFT**	**S2600WF0**
Onboard SATA Support	• 12 x SATA 6Gbps ports (6Gb/s, 3 Gb/s and 1.5Gb/s transfer rates are supported) o Two single port 7-pin SATA connectors o Two M.2 connectors – SATA / PCIe* o Two 4-port mini-SAS HD (SFF-8643) connectors • Embedded SATA Software RAID o Intel® Rapid Storage RAID Technology (RSTe) 5.0 o Intel® Embedded Server RAID Technology 2 (ESRT2) 1.60 with optional RAID 5 key support o **NOTE:** ESRT2 is only supported on S2600WFT and S2600WF0 boards	
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 • Onboard 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 identifies system level features associated with the Intel® Server System R1000WF product family.

Table 3. Intel® Server System R1000WF Product Family Feature Set

Feature	Description
Chassis Type	1U Rack Mount Chassis
Server Board	Intel® Server Board S2600WF product family
Maximum Supported Processor Thermal Design Power (TDP)	Up to 165 Watts
External I/O Connections	• DB-15 Video connectors o Front and Back • RJ-45 Serial Port A connector on back panel • Dual RJ-45 Network Interface connectors on back panel – (S2600WFT based systems only) • Dedicated RJ-45 server management port on back panel • Three USB 3.0 connectors on back panel • Two USB 3.0 connectors on front panel
Internal I/O Connectors / Headers	One Type-A USB 2.0 connector One DH-10 Serial Port B connector
System Fans	Six managed 40mm dual rotor system fans One power supply fan for each installed power supply module
Riser Card Support	Support for two riser cards: • Riser #1 – PCIe* 3.0 x24 • Riser #2 – PCIe* 3.0 x24 With two riser cards installed, up to 2 possible add-in cards can be supported: One x16 PCIe* 3.0 Add-in card slot per riser card 2 Full Height / Half Length add-in cards via Risers #1 and #2
Power Supply	Power Supply Module options: • AC 1100W Platinum • DC 750W Gold The server system can support 1 or 2 installed power supply modules, with support for the following power configurations: 1+0 Non-Redundant, 1+1 Redundant, and 2+0 Combined Power
Drive Support	**R1304WFxxx** – 4 x 3.5" hot swap drive bays + SAS/SATA backplane **R1208WFxxx** – 8 x 2.5" hot swap drive bays + SAS/SATA/NVMe combo backplane
Supported Rack Mount Kit Accessory Options	**A1UFULLRAIL** – Tool-less rack mount rail kit – 780mm max travel length **A1USHRTRAIL** – Tool-less rack mount rail kit – 780mm max travel length – No CMA Support **AXXELVRAIL** – Enhanced value rack mount rail kit – 424mm max travel length **AXX1U2UCMA** – Cable Management Arm – (*supported with A1UFULLRAIL only) **AXX2POSTBRCKT** – 2-post fixed mount bracket kit

2.1 System Features Overview

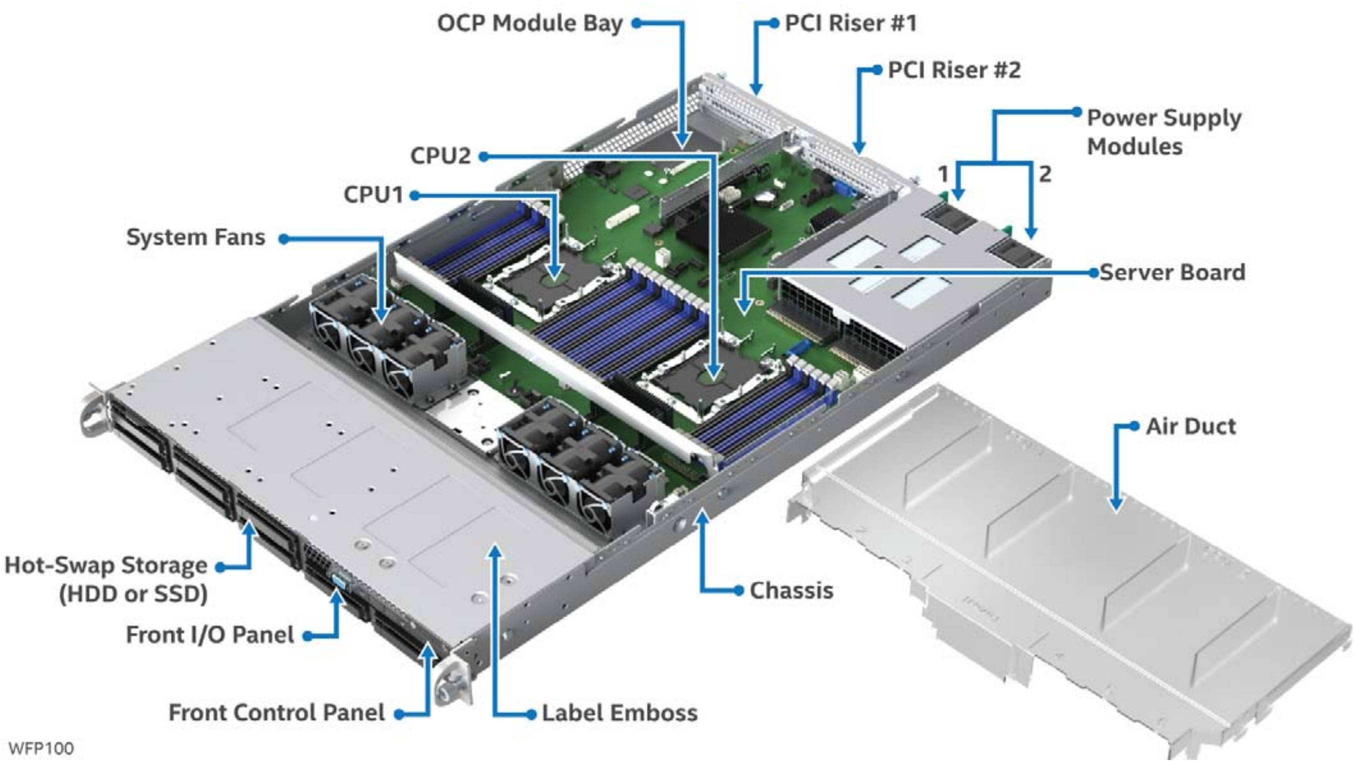

Figure 1. System Components Overview

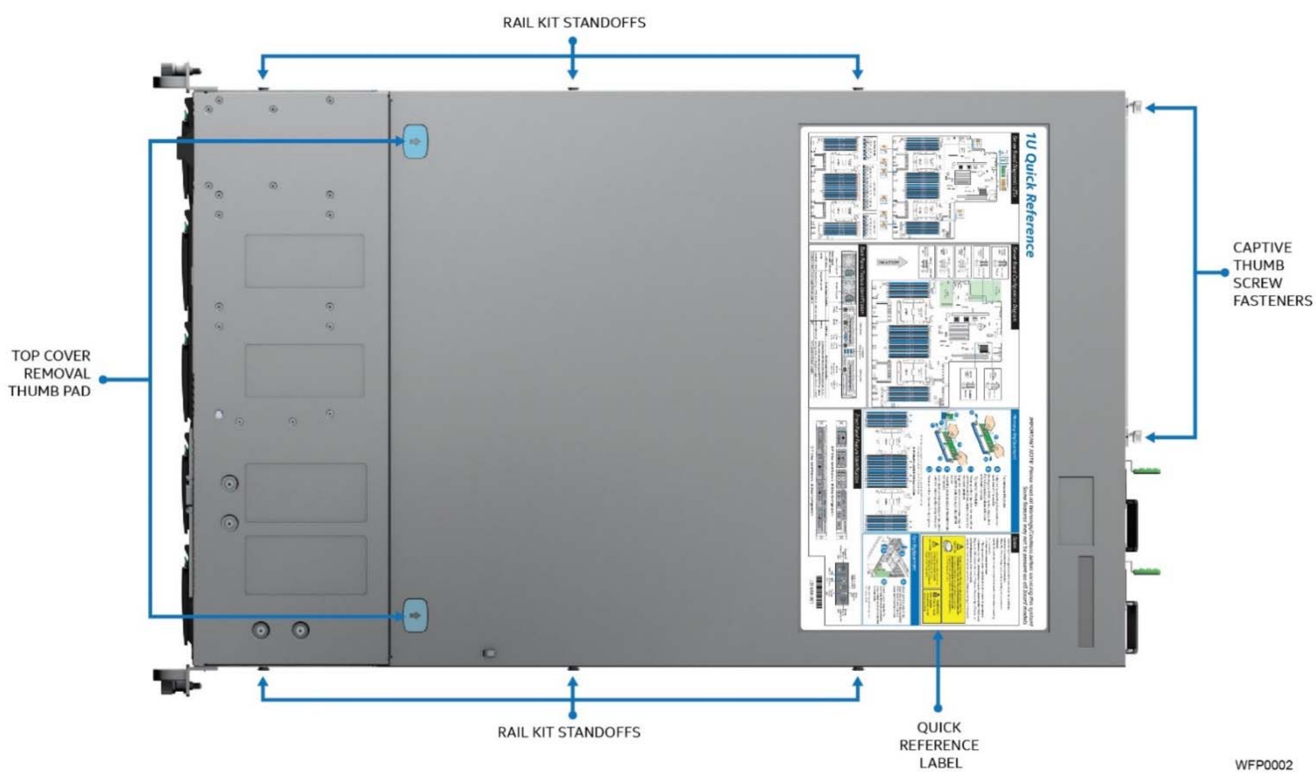

Figure 2. Top Cover Features

2.2 Server Board Architecture

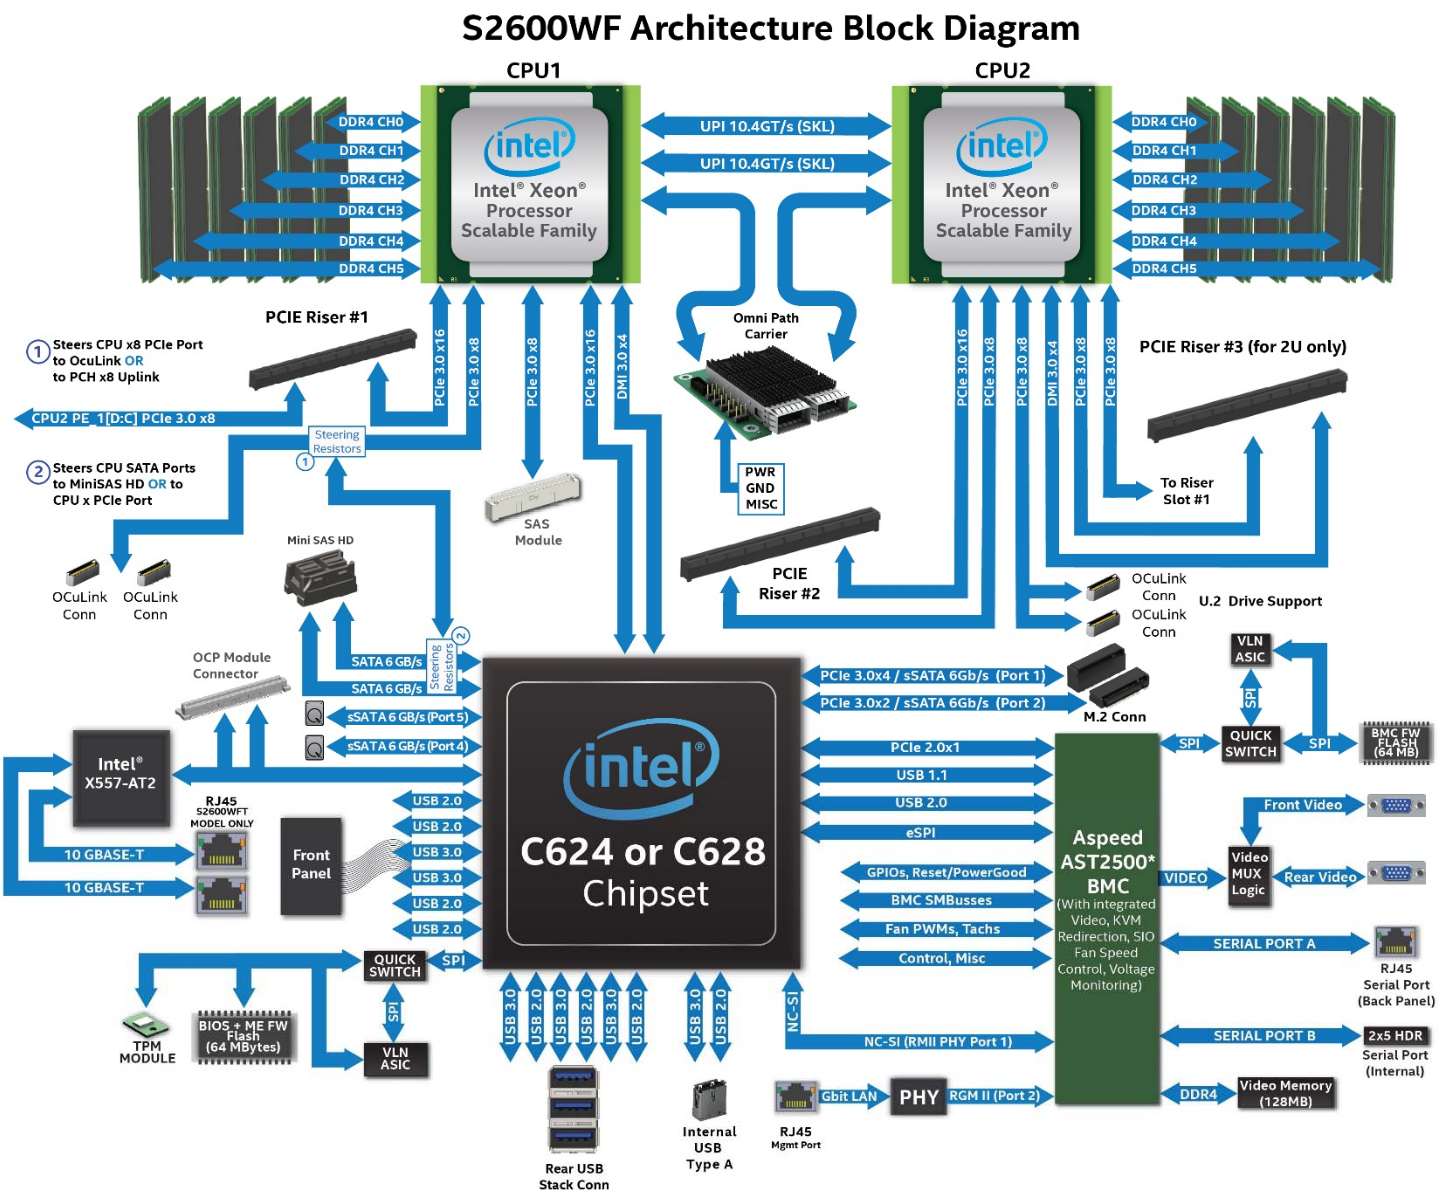

Figure 3 . Server Board Architecture

2.3 Server Board Features Overview

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

Note: Intel® Server Board S2600WFT shown. Some features may not be present on Intel® Server Boards S2600WF0 and/or S2600WFQ.

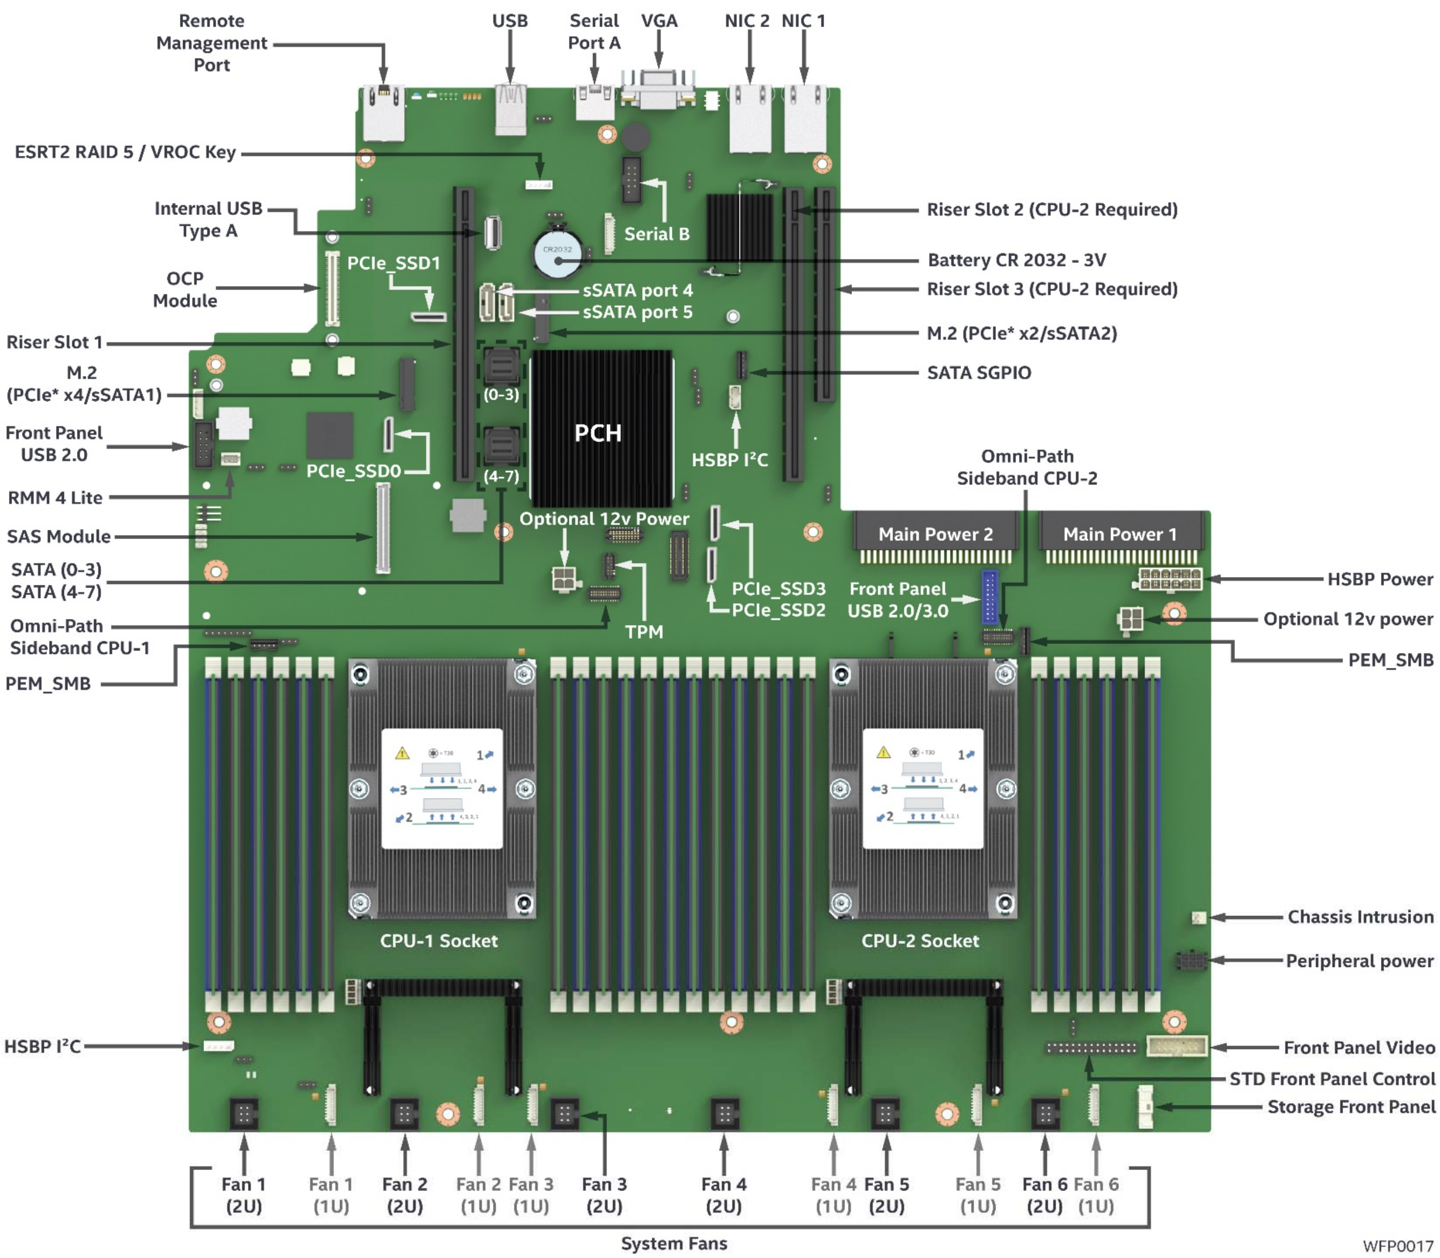

Figure 4. Server Board Features

The server board includes several LEDs to identify system status and / or indicate a component fault. The following illustrations define each Diagnostic LED and identify their location.

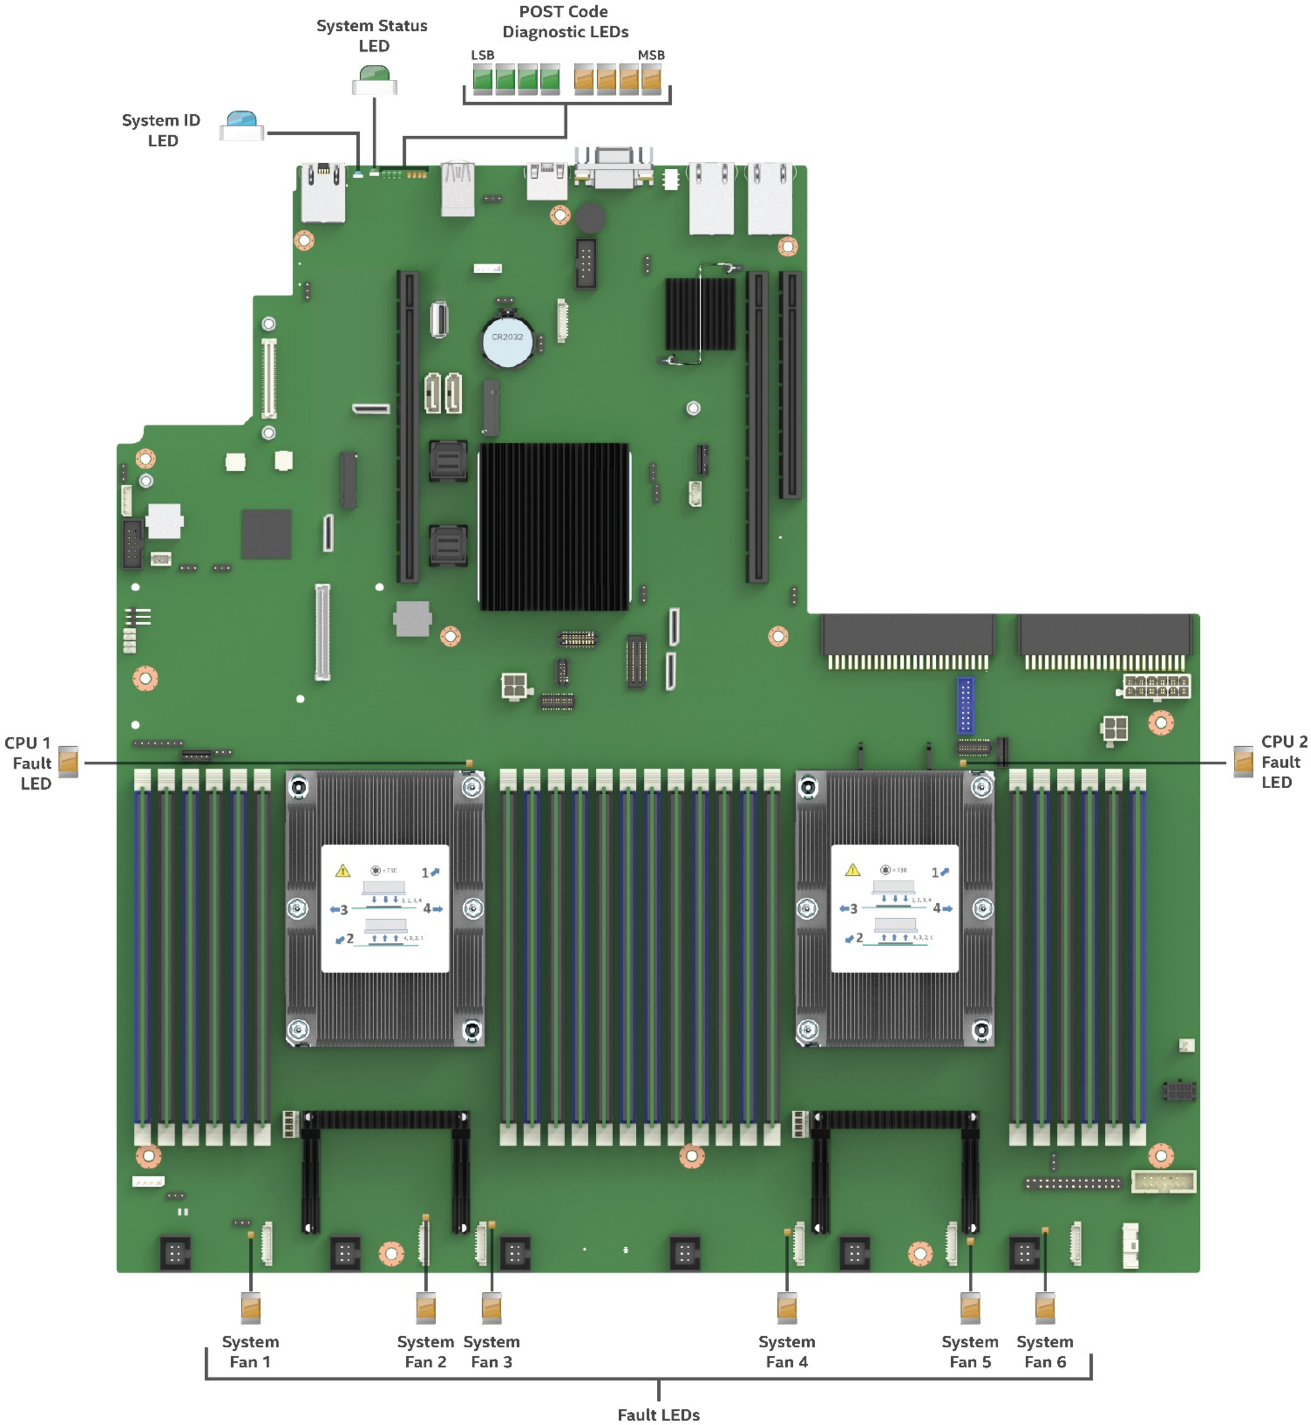

Figure 5. Intel® Light Guided Diagnostic LED Identification

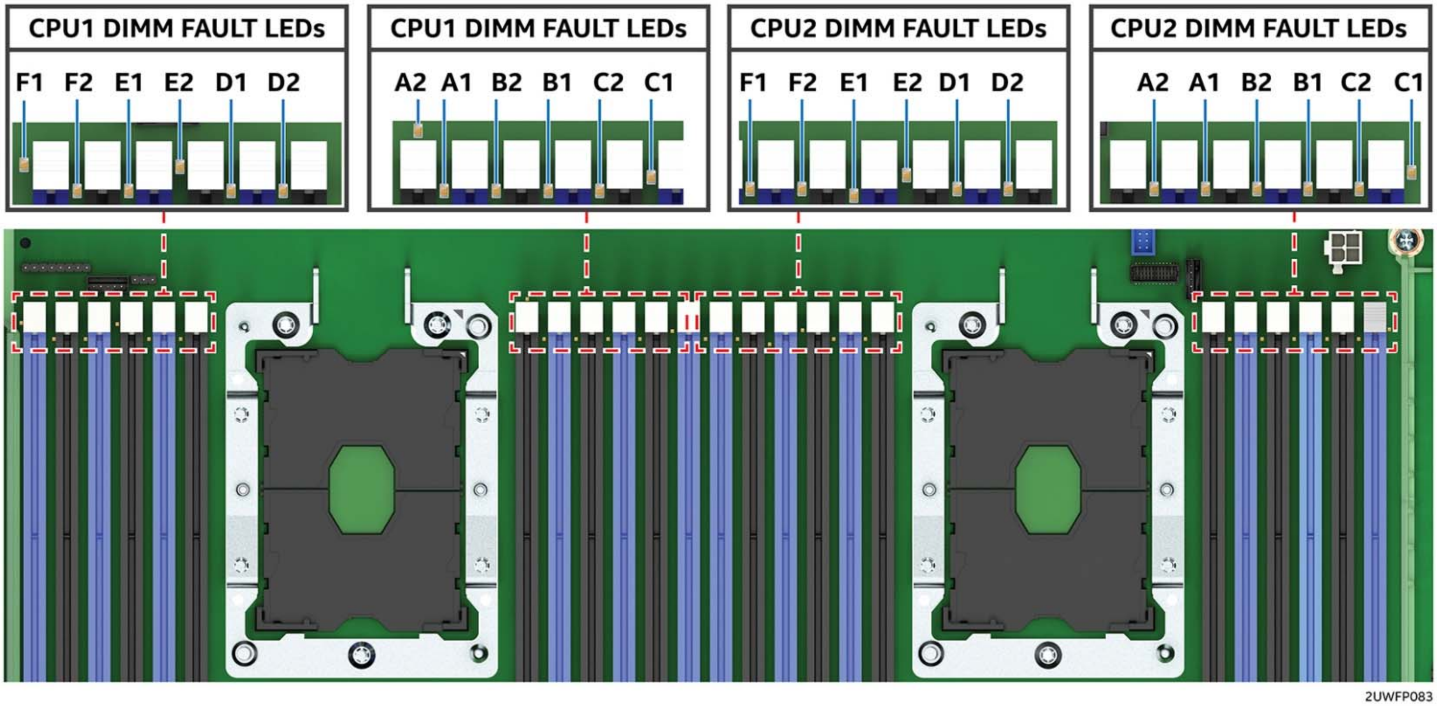

Figure 6. Intel® Light Guided Diagnostics – DIMM Fault LEDs

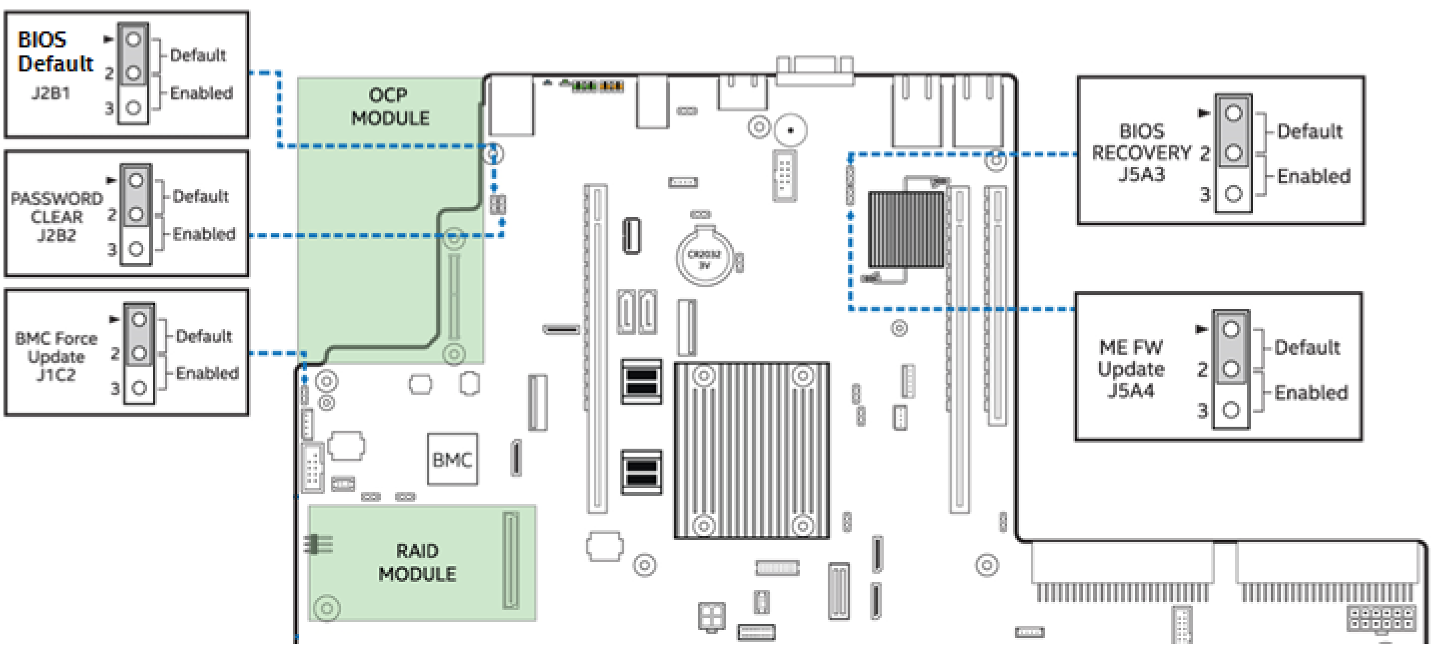

Figure 7. System Reset and Configuration Jumpers

2.4 Back Panel Features

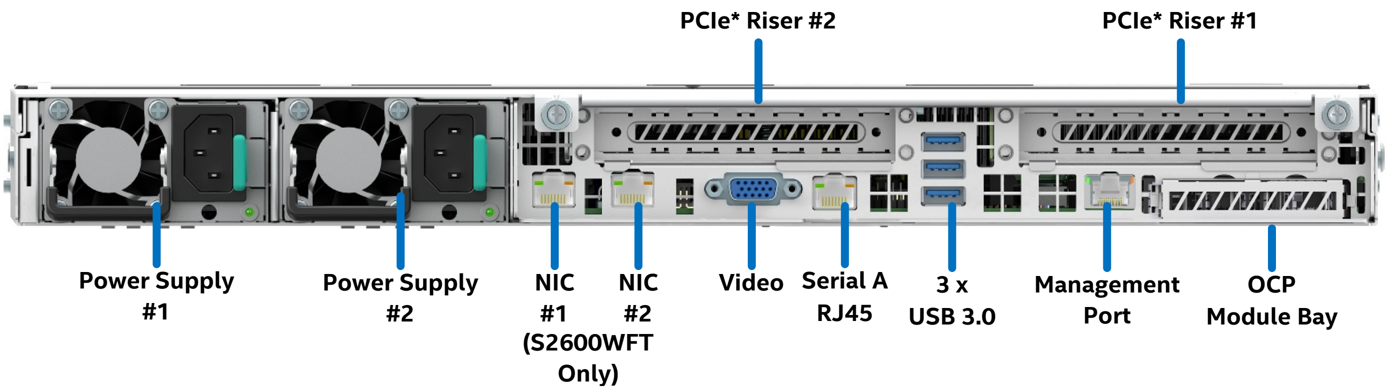

Figure 8. Back Panel Feature Identification

2.5 Front Control Panel

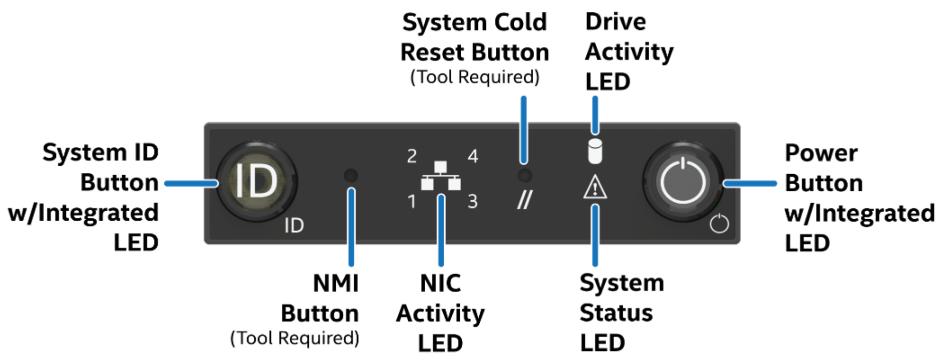

Figure 9. Front Control Panel Options

2.6 Front Drive Bay Options

Figure 10. 3.5" Drive Bay – 4 Drive Configuration (Model R1304WFxxx)

Figure 11. 2.5" Drive Bay – 8 Drive Configuration (Model R1208WFxxx)

2.7 Locking Front Bezel Support

The Intel 1U 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.

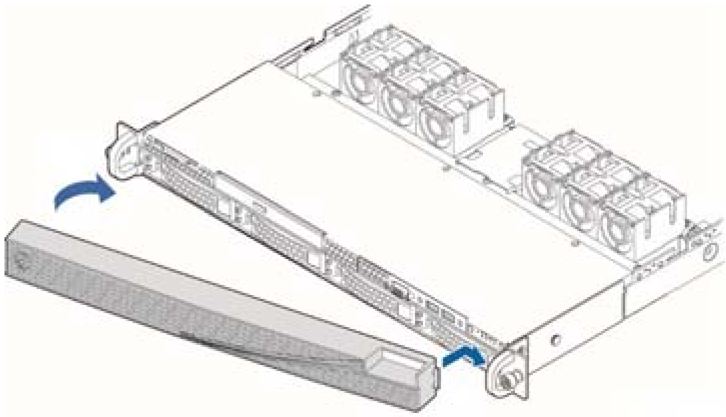

Figure 12. Concept Reference Design - Front Bezel Installation

2.8 System Dimensions

2.8.1 Chassis Dimensions

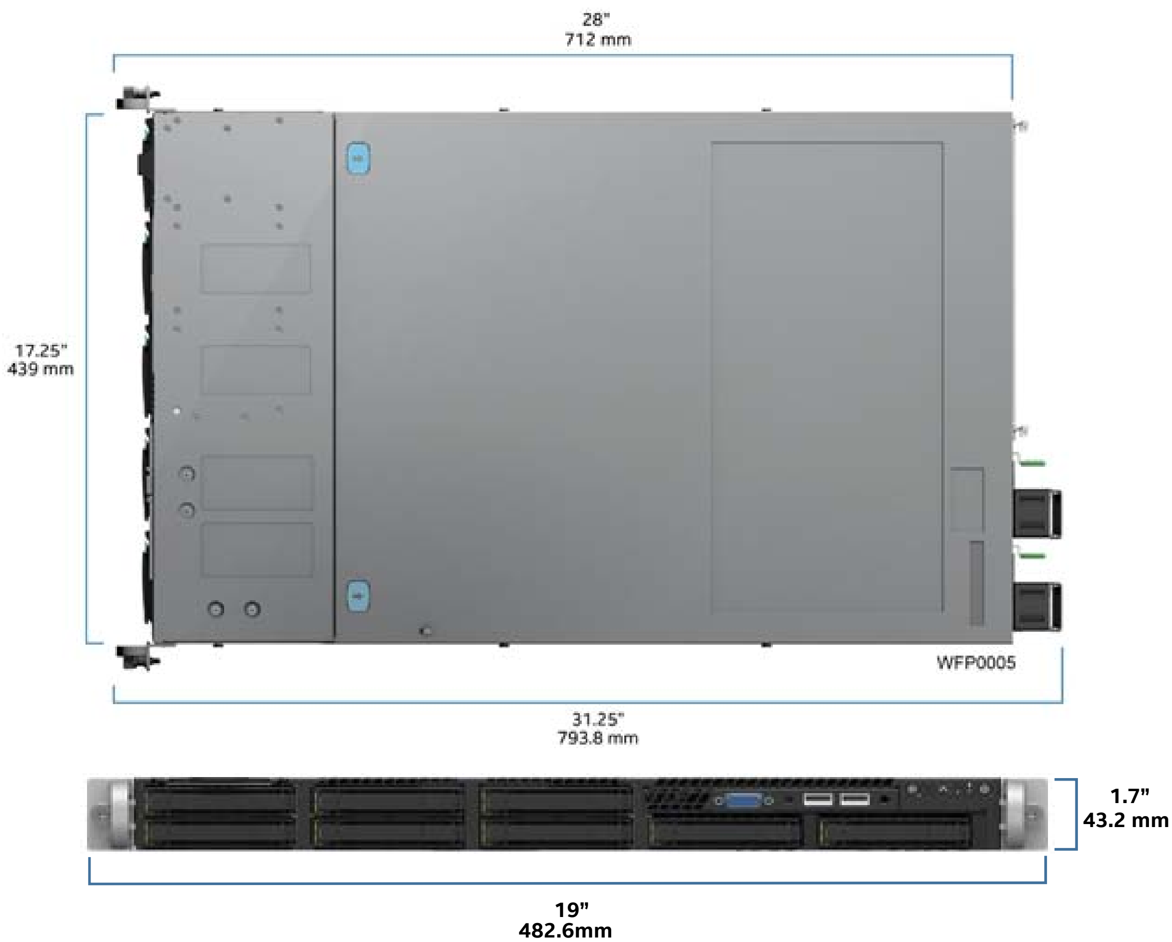

Figure 13. Chassis Dimensions

2.8.2 Label Emboss Dimensions

Figure 14. Label Emboss Dimensions

2.8.3 Pull-out Tab Label Emboss Dimensions

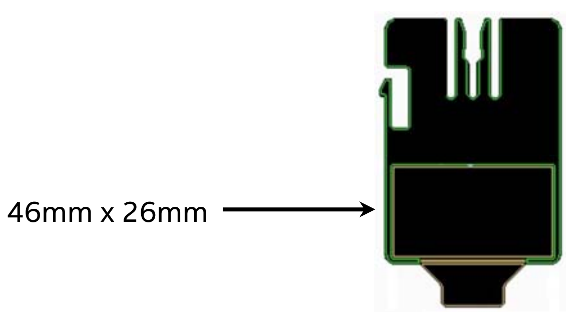

Figure 15. Pull-out Tab Label Emboss Dimensions

2.9 System Cable Routing Channels

The 1U 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.

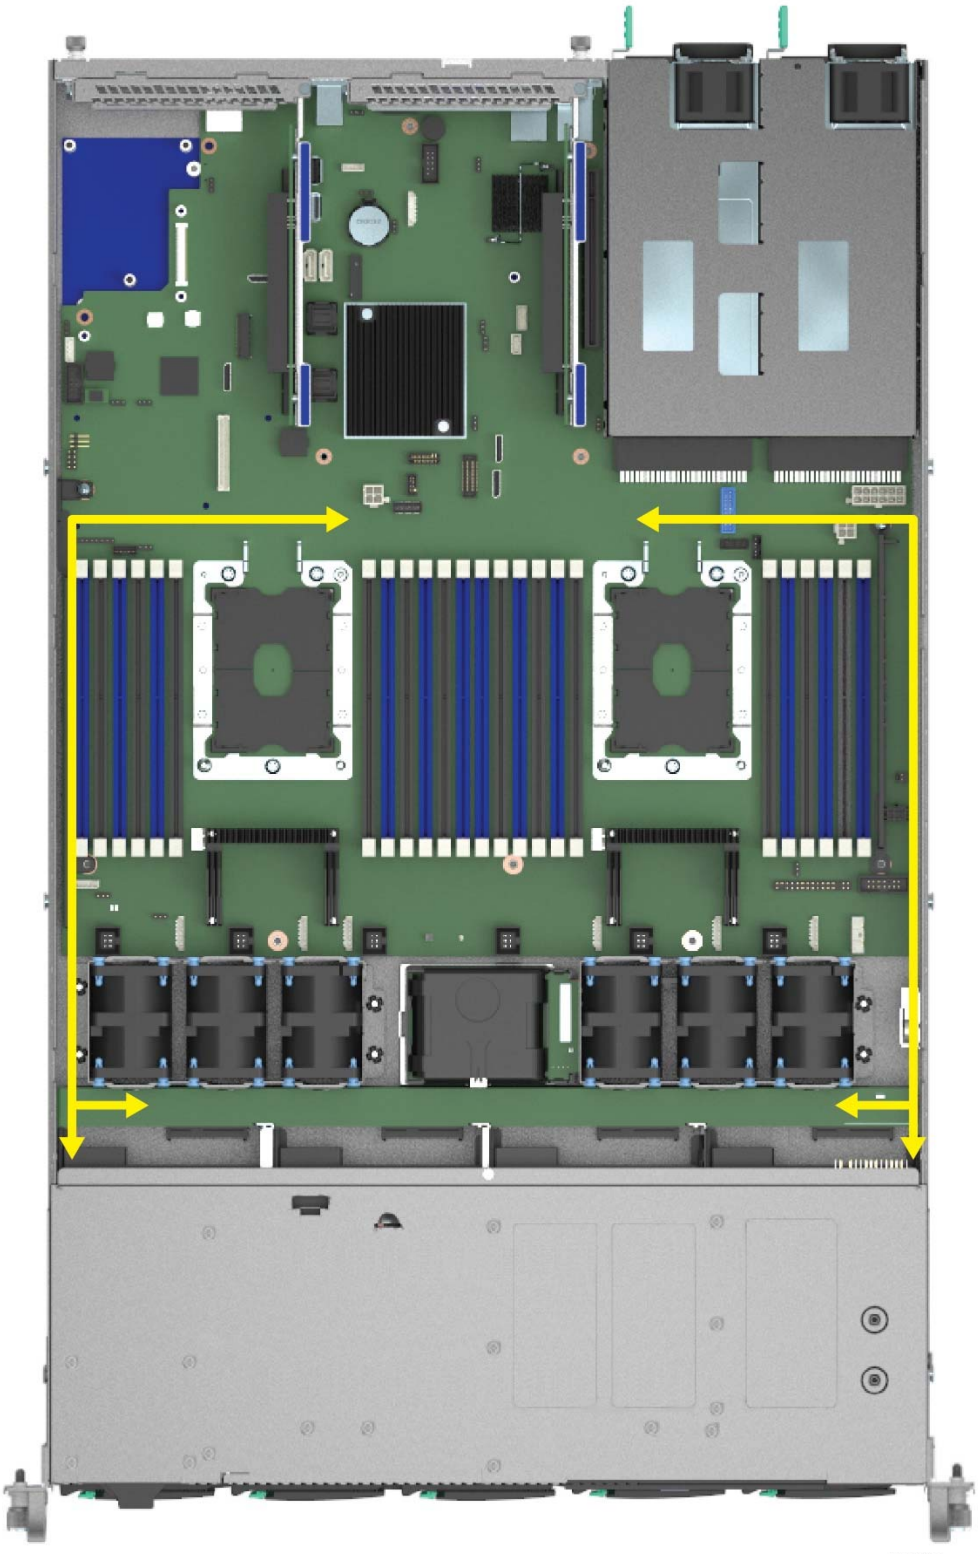

Figure 16. System Cable Routing Channels

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 choose 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 rail kit's specified maximum weight limit or misalignment of the server in the rack may result in failure of the rack rails, resulting in damage to the system or personal injury. 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
  - 1U, 2U, 4U compatible
  - Tool-less chassis attach
  - Tools required to attach rails to rack
  - Rack installation length adjustment: 610mm ~ 765mm
  - 424mm travel distance (2/3 system extension from rack)
  - 130 lbs. (59 Kgs) maximum supported weight

- **A1UFULLRAIL** – Premium Rack Mount Rail Kit
  - 1U compatible only
  - Tool-less installation
  - 780mm travel distance (Full system extension from rack)
  - 39 lbs. (18 Kgs) maximum supported weight
  - Support for Cable Management Arm (CMA) AXX1U2UCMA

- **A1USHRTRAIL** – Premium Rack Mount Rail Kit
  - 1U compatible only
  - Tool-less installation
  - 780mm travel distance (Full system extension from rack)
  - 39 lbs. (18 Kgs) maximum supported weight
  - No Cable Management Arm support

- **AXX2POSTBRCKT** – 2-Post Fixed Mount Bracket Kit
  - 1U and 2U compatible
  - Tools required to attach components to rack

- **AXX1U2UCMA** – Cable Management Arm
  - Support for Premium Rack Mount Kit A1UFULLRAIL 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. 10º C to 35º C (50º F to 95º F) with the maximum rate of change not to exceed 10°C per hour
		ASHRAE Class A3 – Includes operation up to 40C for up to 900 hours per year.
		ASHRAE Class A4 – Includes operation up to 45C for up to 90 hours per year.
	Shipping	-40º C to 70º C (-40º F to 158º F)
Altitude	**Operating**	Support operation up to 3050m 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, 2g, 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 in Watts**	<300 W ≥300 W ≥600 W ≥1000 W
	Servers/Rack Mount Sound Power Level (in BA)	7.0 7.0 7.0 7.0

See Appendix D in this document or the *Intel® Server Board S2600WF Product Family Power Budget and Thermal Configuration Tool* for system configuration support limits.

2.12 System Packaging

Intel system packaging is designed to provide a fully configured system the necessary protection when exposed to the rigors of shipment. The packaging was designed and tested to meet International Safe Transit Association (ISTA) Test Procedure 3A (2008). The packaging was 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 will allow the package to meet the ISTA (International Safe Transit Association) Test Procedure 3A (2008) limits. See the ***Intel® Server System R1000WF 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:

Length = 983mm

Breadth = 577mm

Height = 260mm

Inner Box Internal Dimensions:

Length = 956mm

Breadth = 550mm

Height = 202mm

Table 5. Intel Product Weight Information

Product Code	Product Type	Packaged Gross Weight (Kg)	Packaged Gross Weight (Lbs)	Un-packaged Net Weight (Kg)	Un-packaged Net Weight (Lbs)
R1304WFXXX	Chassis	12.57	27.65	4.265	9.38
R1208WFXXX	Chassis	13.17	28.97	4.865	10.7
R1304WFXXX	System	20.6	45.32	12.3	27.1
R1208WFXXX	System	21.42	47.12	13.12	28.86

Note: Integrated system weights will vary depending on the final system configuration. For the 1U product family, a fully integrated un-packaged system can weigh upwards of 40 pounds (18+ Kg).

Reference the ***Intel® Server S2600WF Product Family Configuration Guide*** for product weight information associated with each available product SKU.

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, supporting 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 will be utilized. Should this occur, power redundancy is lost. In a 2+0 power configuration, total power available may be 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 thermals exceed programmed limits, platform management will attempt to keep the system operational. See Closed Loop System Throttling (CLST) later in this chapter, and Chapter 4 Thermal Management, for details.

Caution: Installing two Power Supply Units with different wattage ratings in a system is not supported. Doing so will not provide Power Supply Redundancy and will result 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.

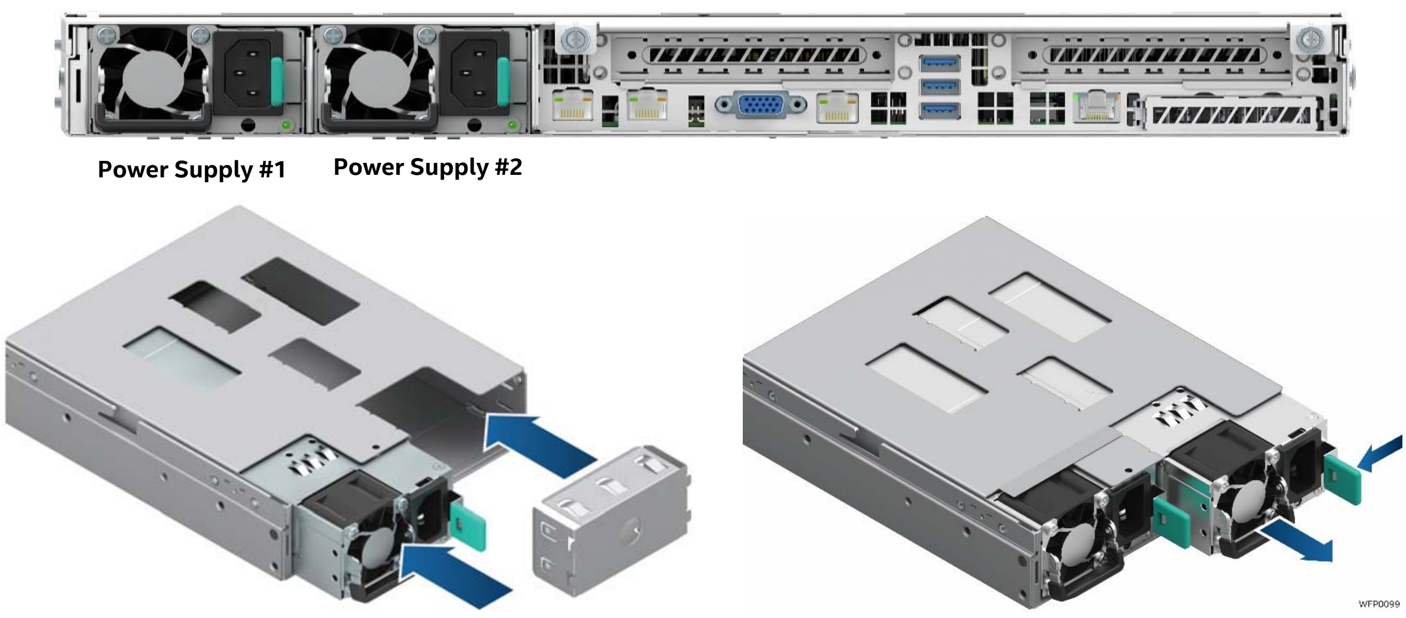

Figure 17. Power Supply Bay

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 two power supply options available for this server product family:

- 1100W AC 80 Plus Platinum – iPC **AXX1100CRPS**
- 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 four different load levels: 100%, 50%, 20% and 10%.

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%

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

Table 7. 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 / Dual Fans

750W DC Power Supply Module

AC Power Cable Connector

DC Power Cable Connector

Figure 18. Power Supply Module Overview

The physical size of the 1100W AC power supply enclosure is 39mm x 74mm x 185mm. The power supply includes dual, in-line, 40mm fans, 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 213mm front to back.

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

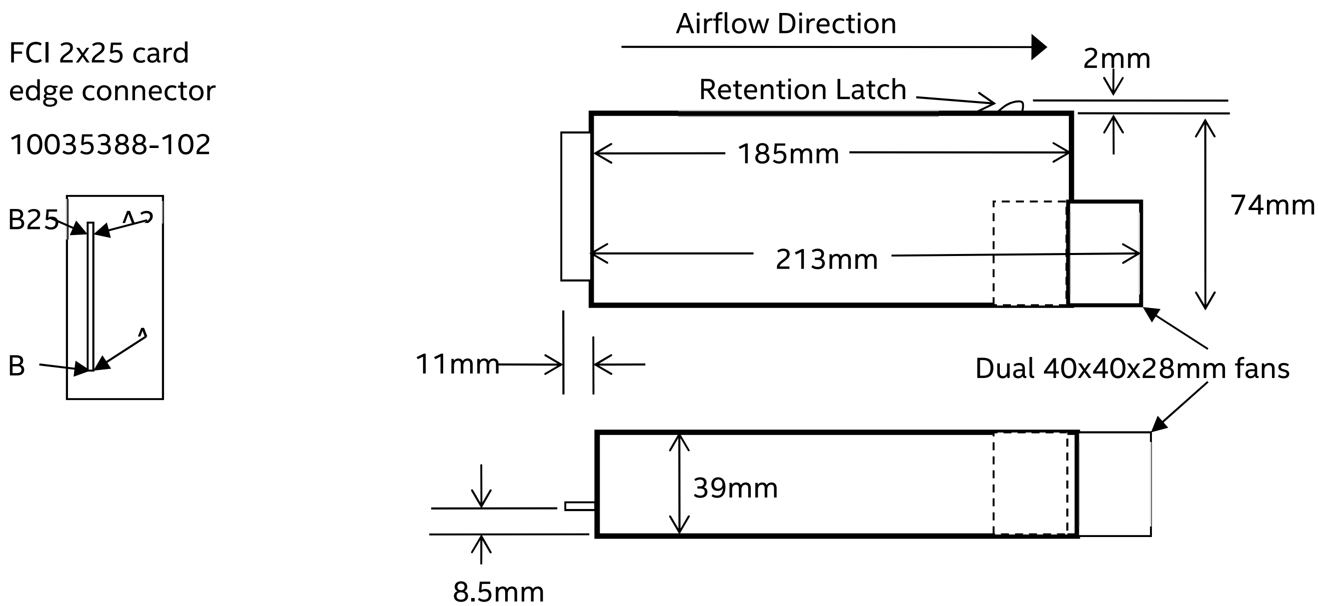

Figure 19. 1100W AC Power Supply Module Mechanical Drawing

3.2.3 Power Cord Specification Requirements

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

Table 8. AC Power Cord Specifications

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

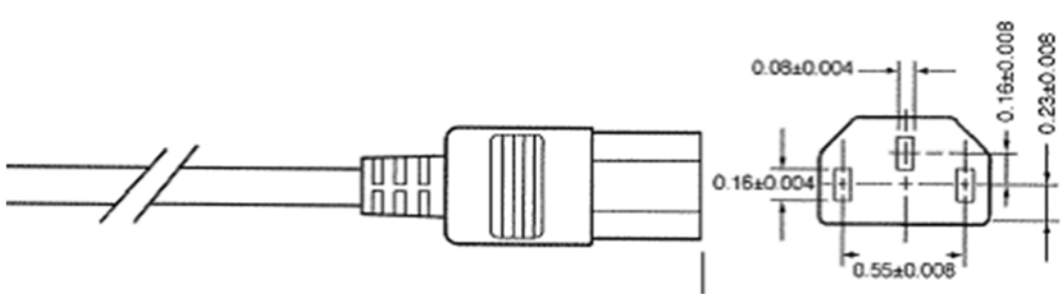

Figure 20. AC Power Cord

The DC power cord used must meet the specification requirements listed in the following tables.

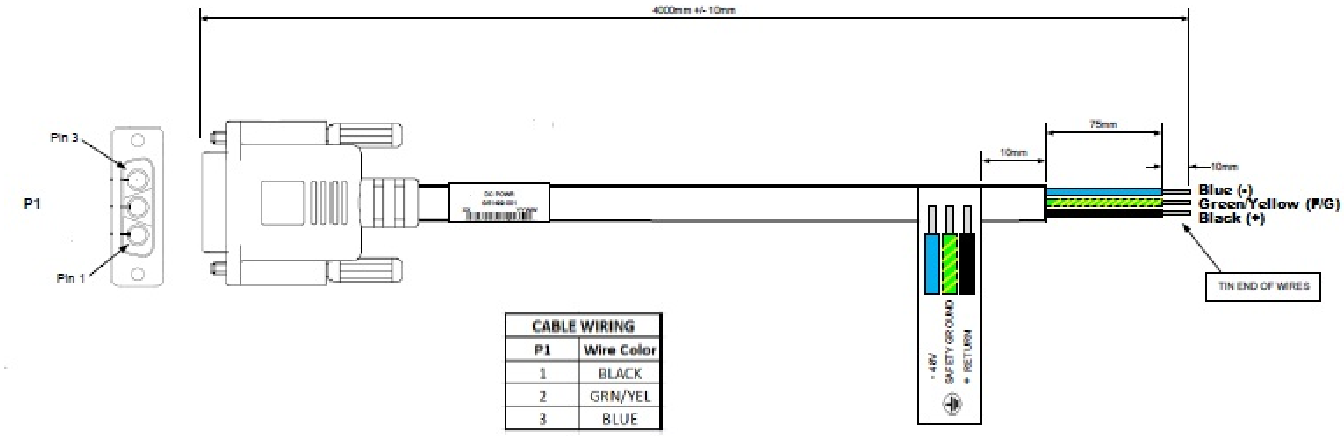

Figure 21. DC Power Cord Specification

Table 9. 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.

3.3.1 Power Factor

The power supply meets the power factor requirements stated in the Energy Star* Program Requirements for Computer Servers. These requirements are stated below.

Table 10. Power Factor Requirements

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

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

3.3.2 AC Input Voltage Specification

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

Table 11. AC Input Voltage Range – 1100W Power Supply

Parameter	MIN	Rated	VMAX	Start-up 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		

1. Maximum input current at low input voltage range shall be measured at 90VAC, at max load.
2. Maximum input current at high input voltage range shall be measured at 180VAC, at max load.
3. 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. Transformers' isolation between primary and secondary windings comply with the 3000Vac (4242Vdc) 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 complies with reinforced insulation per safety standard IEC 950. Separation between the primary and secondary circuits, and primary to ground circuits 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 should recover and meet 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 will not cause damage to the power supply.

Table 12. AC Line Holdup Time – 1100W Power Supply

Loading	Holdup time
70%	10msec

3.3.4.1 AC Line 12VSB Holdup

The 12VSB output voltage operates in regulation under its full load (static or dynamic) during an AC dropout of 70ms min (=12VSB 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 single line fuse on the line (Hot) wire of the AC input. The line fusing meets all safety agency requirements. The input fuse is 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 a 55A peak, 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 as well as upon recovery after AC dropout of any duration and over the specified temperature range (Top).

3.3.7 AC Line Transient Specification

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

The power supply meets all requirements under the AC line sag and surge conditions listed in Table 12 above and Table 13 below.

Table 13. AC Line Sag Transient Performance

AC Line Sag (10sec 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 14. AC Line Surge Transient Performance

AC Line Surge				
Duration	**Surge**	**Operating AC Voltage**	**Line Frequency**	**Performance Criteria**
Continuous	10%	Nominal AC Voltages	50/60Hz	No loss of function or performance
0 to ½ AC cycle	30%	Mid-point of nominal AC Voltages	50/60Hz	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 15. 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.

3.3.9 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 was tested with the system for immunity to the following for each power supply option:

- 1100W 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 included:

- No unsafe operation under any condition; all power supply output voltage levels are 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 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 will recover 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 PSON# cycle HIGH for one second reset the power supply.

3.3.15.1 Over-Current Protection

The power supply has a current limit to prevent the outputs from exceeding the values shown in Table 15 below. If the current limits are exceeded the power supply will shut down and latch off. The latch will be cleared by toggling the PSON# signal or by an AC power interruption. The power supply will not be damaged from repeated power cycling in this condition. 12VSB will be auto-recovered after removing Over-Current Protection limit.

Table 16. Over Current Protection – 1100 Watt Power Supply

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

3.3.15.2 Over-Voltage Protection (OVP)

The power supply over voltage protection is locally sensed. The power supply will shut down and latch off after an over voltage condition occurs. This latch will be cleared by an AC power interruption. The values are measured at the output of the power supply's connectors. The voltage will never exceed the maximum levels when measured at the power connectors of the power supply connector during any single point of failure. The voltage does not trip any lower than the minimum levels when measured at the power connector. 12VSB will be auto-recovered after removing OVP limit.

Table 17. Over Voltage Protection (OVP) Limits – 1100 W Power Supply

Output Voltage	MIN (V)	MAX (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 will shut down. When the power supply temperature drops to within specified limits, the power supply will restore power automatically while the 12VSB remains always on. The OTP circuit has a built in margin so that the power supply will not oscillate on and off due to a temperature recovering condition. The OTP trip level has a minimum of 4°C of ambient temperature margin.

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 the following table.

Table 18. LED Indicators

Power Supply Condition	LED State
Output ON and OK	GREEN
No AC power to all power supplies	OFF
AC present / Only 12VSB on (PS off) or PS in Cold redundant state	1Hz Blink, GREEN
AC cord unplugged or AC power lost; with a second power supply in parallel still with AC input power.	AMBER
Power supply warning events where the power supply continues to operate; high temp, high power, high current, slow fan.	1Hz Blink, Amber
Power supply critical event causing a shutdown; failure, OCP, OVP, Fan Fail	AMBER
Power supply FW updating	2Hz Blink, 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, with each connected to separate ground wires while the system is operational.

3.4.1 DC Input Voltage

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

Table 19. DC Input Rating

Parameter	MIN	Rated	MAX
DC Voltage	-40.5 VDC	-48VDC/-60VDC	-75VDC
Input Current	24A		12.5A

3.4.2 DC Input Fuse

The -48VDC power supply input is fused. The fusing meets all safety agency requirements. DC inrush current does not cause the fuse to blow under any conditions nor do the protection circuits in the power supply 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 is 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 does 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 contains protection circuitry (under-voltage lock-out) such that the application of an input voltage below the specified minimum specified, will not cause damage (overstress) to the power supply unit due to over-heating or other condition.

3.4.5 DC Holdup Time and Dropout

Loading	Holdup Time
750W (100%)	0.2msec

During a DC dropout of 0.2ms or less the power supply meets dynamic voltage regulation requirements for every rated load condition. A DC line dropout of 0.2ms or less does not cause tripping of control signals or protection circuits. Repeated every 10 seconds starting at the min input voltage DC line dropout will not damage the power supply under any specified load conditions. The PWOK signal does not go to a low state under these conditions. DC dropout transients in excess of 0.2 milliseconds may cause shutdown of the PSU, but will not damage the power supply. The power supply will recover and meet all turn on requirements for DC dropouts that last longer than 0.2ms. The power supply meets the DC dropout requirement over rated DC voltages and output loading conditions.

3.4.6 DC Line Surge Voltages (Line Transients)

The Power Supply tolerates transients in the input DC power line caused by switching or lightning. The power supply has been tested and is 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 were: 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 20. Line Voltage Transient Limits

Duration	Slope/Rate	Output	Performance criteria
200μs max	-48V → -30V w/ +2V/μs	Rated DC Voltages	No loss of function or performance
	-30V → -48V w/ -2V/μs	Rated DC Voltages	No loss of function or performance

3.4.7 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 21. Susceptibility Requirements

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 complies within 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 comply with those specified in the Intel Environmental Standards Handbook.

3.4.7.2 Fast Transient/Burst

The power supply complies within 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 with those specified in the Intel Environmental Standards Handbook.

3.4.7.3 Radiated Immunity

The power supply complies within 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 comply with those specified in the Intel Environmental Standards Handbook. Additionally, they also comply with field strength requirements specified in GR 1089 (10V/meter).

3.4.7.4 Surge Immunity

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

The pass criteria included:

- 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 complies 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 also comply with those specified in the Intel Environmental Standards Handbook.

3.4.8 Protection Circuits

Protection circuits inside the power supply will only cause 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 15sec and a PSON# cycle HIGH for 1sec shall be able to reset the power supply.

3.4.8.1 Current Limit (OCP)

The power supply has current limits to prevent the outputs from exceeding the values shown in Table 20 below. If the current limits are exceeded the power supply will shut down and latch off. The latch will be cleared by toggling the PSON# signal or by a DC power interruption. The power supply will not be damaged from repeated power cycling in this condition. 12VSB will be auto-recovered after removing OCP limit.

Table 22. Over Current Protection – 750 Watt Power Supply

Output Voltage	Input Voltage Range	Over Current Limits
+12V	90 – 264VAC	72A min; 78A max
12VSB	90 – 264VAC	2.5A min; 3.5A 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 connectors of 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. 12VSBwill be auto-recovered after removing OVP limit.

Table 23. Over Voltage Protection Limits – 750 Watt Power Supply

Output Voltage	MIN (V)	MAX (V)
+12V	13.3	14.5
+12VSB	13.3	14.5

3.4.8.3 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 will shut down. When the power supply temperature drops to within specified limits, the power supply will restore power automatically, while the 12VSB remains always on. The OTP circuit has a built in margin so that the power supply will not oscillate on and off due to a temperature recovering condition. The OTP trip level has 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, 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 power supply needed 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 voltage falls below the Vfault 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 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 the power supply to be in, the cold standby configuration will slightly change the load share threshold that the power supply will power on at.

Table 24. Load Share Threshold for Activating Supplies Example

	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)	3.2V (40% of max)	3.2V x 0.5 x 0.9 = 1.44V	OK = Open Fault = Low
Cold Standby 2 (03h)	5.0V (62% of max)	5.0V x 0.67 x 0.9 = 3.01V	OK = Open Fault = Low
Cold Standby 3 (04h)	6.7V (84% of max)	6.7V x 0.75 x 0.9 = 4.52V	OK = Open Fault = Low

Notes:

1. Maximum load share voltage = 8.0V at 100% of rated output power
2. These are example load share bus thresholds; for a given power supply, these should 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 100µsec.

3.5.3 BMC Requirements

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

The BMC schedules a rolling change for which PSU is the Active, Cold Stby1, Cold Stby 2, and Cold Stby 3 power supply. This allows for equal loading across power supplies over their life.

The list below are events that trigger a re-configuration of the power supplies using the Cold_Redundancy_Config command.

- 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 controlled by each PSU sensing the Vshare bus. Once a power supply turns on after crossing the enable threshold, it lowers its threshold to the disable threshold. The system defines the "position" of each power supply in the Cold Redundant operation. It will do 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). CLST prevents the system from crashing if a power supply module is overloaded or insufficiently cooled. Should system power reach a pre-programmed power limit or power supply thermal sensor hit the threshold, CLST will throttle system memory and/or processors to reduce power. System performance will be impacted should this occur.

3.7 Smart Ride Through (SmaRT)

The server system has support for Smart Ride Through Throttling (SmaRT). This feature increases the reliability for a system operating in a heavy power load condition to remain operational during an AC line dropout event. See Section 3.3.4, AC Line Dropout / Holdup for power supply hold up time requirements for AC Line dropout events.

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 for various system options. The following sub-sections will 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 25. 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	PSON	B21	12V stby
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 Hot Swap Backplane Power Connector

The server board includes one white, 2x6-pin power connector that is cabled to provide power for hot swap backplanes. On the server board, this connector is labeled “HSBP PWR”. The following table provides the pin-out for this connector.

Table 26. Hot Swap Backplane Power Connector Pin-out (“HSBP PWR”)

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

3.8.3 Peripheral Power Connector

The server board includes one brown, 2x3-pin power connector intended to provide power to optionally installed peripheral devices. On the server board this connector is labeled "Peripheral PWR". The following table provides the pin-out for this connector.

Table 27. Peripheral Drive Power Connector Pin-out ("Peripheral PWR")

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

4. Thermal Management

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

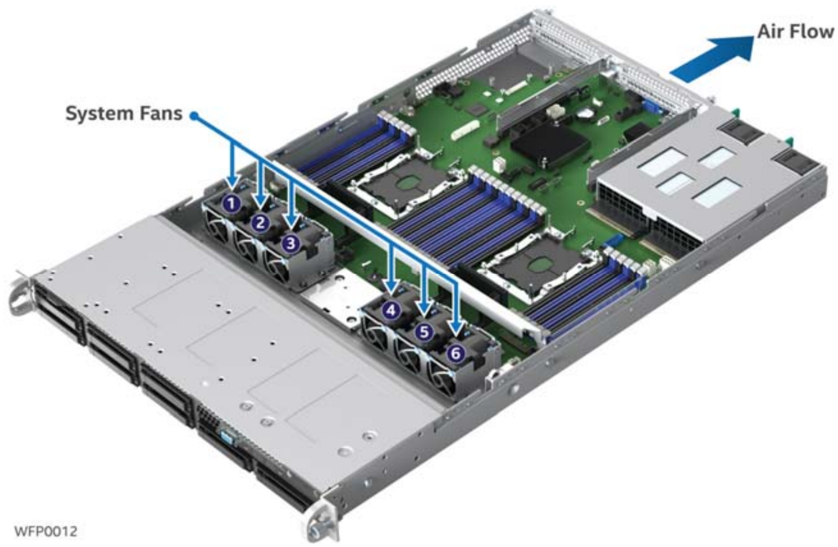

Figure 22. System Air Flow and Fan Identification

The following tables provide air flow data associated with the different system models within this 1U 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 was derived using server management utilities that utilize platform sensor data and may vary slightly from the data listed in the tables.

Table 28. System Volumetric Air Flow

System Airflow – R1304WFxxxx				System Airflow – R1208WFxxxx		
System Fan	PSU Fan	Total Airflow (CFM)		System Fan	PSU Fan	Total Airflow (CFM)
100%	auto	85.6		100%	auto	89.0
80%	auto	67.7		80%	auto	69.6
60%	auto	48.9		60%	auto	50.8
40%	auto	31.6		40%	auto	32.6
20%	auto	13.6		20%	auto	13.8
100%	100%	89.1		100%	100%	92.6

The Intel® Server System R1000WF 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® Server Board 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 physical system components are used to maintain system thermals. They include six managed 40mm dual rotor system fans, fans integrated into each installed power supply module, an air duct, populated drive carriers, populated DIMM slots, and installed CPU heats sinks. Drive carriers can be populated with a storage device (SSD or Hard Disk Drive) or supplied drive blank. In addition, it is necessary to have specific DIMM slots populated with DIMMs or supplied DIMM blanks. See Figure 26. System configurations that require population of specific DIMM slots will ship from Intel with DIMM blanks pre-installed. Pre-installed DIMM blanks should only be removed when installing a memory module in its place.

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 a maximum operating ambient temperature of 35°C (ASHRAE Class 2).
  - Add-in cards with a minimum 300 LFM (1.5 m/s) air flow requirement or lower can be installed in available add-in card slots in Riser Card #1 and Riser Card #2
  - Add-in cards with an air flow requirement greater than 300 LFM cannot be supported in any PCIe slot on any riser
  - **Note:** Most PCI add-in cards have minimum air flow requirements of 100 LFM (0.5m/s). Some high power add-in cards have minimum air flow requirements of 300 LFM (1.5 m/s) or higher. System integrators should verify PCI add-in card air flow requirements from vendor specifications when integrating add-in cards into the system.

4.1.1 Memory Slot Population Requirements

System thermal requirements dictate that a specific airflow be maintained over or between critical system components. To ensure the 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. The following illustration identifies the memory slots which must be populated in all 1U system configurations.

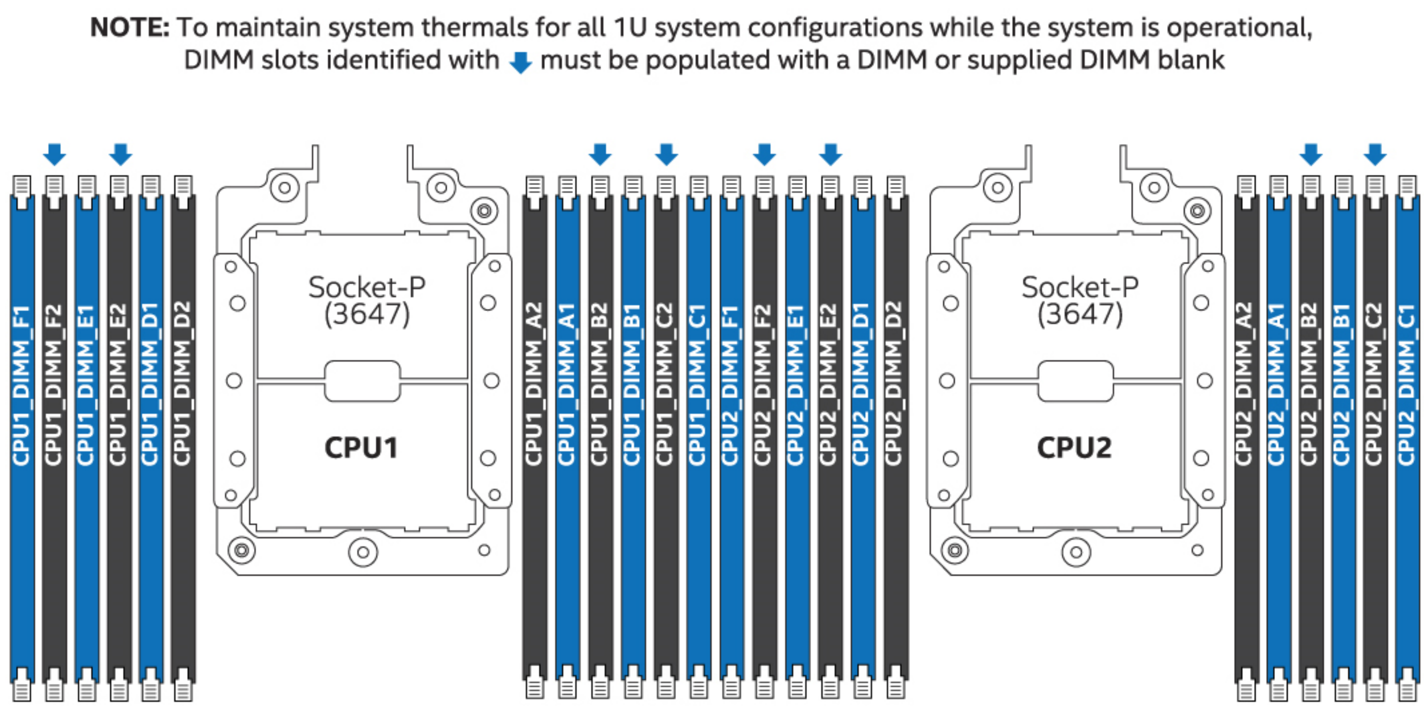

Figure 23. DIMM Population Layout

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:

- o For RDIMM, LRDIMM, 3DS RDIMM, 3DS LRDIMM; Always populate DIMMs with higher electrical loading in slot1, followed by slot 2.

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 (all rotors) should be operational. The system is designed for fan redundancy when the system is configured with two power supplies, all system fans are present and operational, and ambient air remains at or below ASHRAE Class 2 limits. In fan redundancy mode, should a single system fan rotor failure occur, integrated platform management will: change the state of the System Status LED to flashing Green, report an error to the system event log, and automatically adjust remaining 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 or the processors or both, in order to keep components from overheating and keep the system operational. Throttling of these subsystems will continue until system temperatures are reduced below preprogrammed limits.

Should system thermals increase to a point beyond maximum thermal limits, the system will shut down, the System Status LED will change to a solid Amber state, and the event will be logged to the system event log. Should power supply thermals increase to a point beyond its maximum thermal limit or if a power supply fan should fail, the power supply will shut 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 the following Intel website:

http://downloadcenter.intel.com

4.2.1 Fan Speed Control

The baseboard management controller (BMC) controls and monitors the system fans. Each fan is associated with a fan speed sensor that detects fan failure. 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
  - 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 will not be detected and there will not be any fan boost until the system comes out of 'Fans-off; mode.
  - Any associated temperature sensor is in a critical state. The SDR describes which temperature threshold violations cause fan boost for each fan domain.
  - The BMC is in firmware update mode, or the operational firmware is corrupted.
  - If any of the above conditions apply, the fans are set to a fixed boost state speed.
- Nominal
  - 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 value to the fan speeds it would otherwise be driving. 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 for ambient temperatures higher than 35°C.

4.2.1.2 Fan Redundancy Detection

The BMC supports redundant fan monitoring through 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 up to the number of redundant fans specified in the SDR in a fan configuration is a non-critical failure and is reflected at the front panel. A fan failure or removal that exceeds the number of redundant fans is a non-fatal, insufficient-resources condition and is reflected at the front panel 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.

4.2.1.3 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.

4.2.1.4 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. A stepwise linear instance can also 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.

4.2.1.5 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 acoustic standards might require a trade-off between fan speed and system performance parameters that contribute to the cooling requirements and 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.

4.2.1.6 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 and 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]
- 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 system 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 sensor
6. Onboard sensor
7. Virtual sensor
8. Available only when PSU has PMBus
9. Calculated estimate

A simple model is shown in the following figure which gives a high level representation of how the fan speed control structure creates the resulting fan speeds.

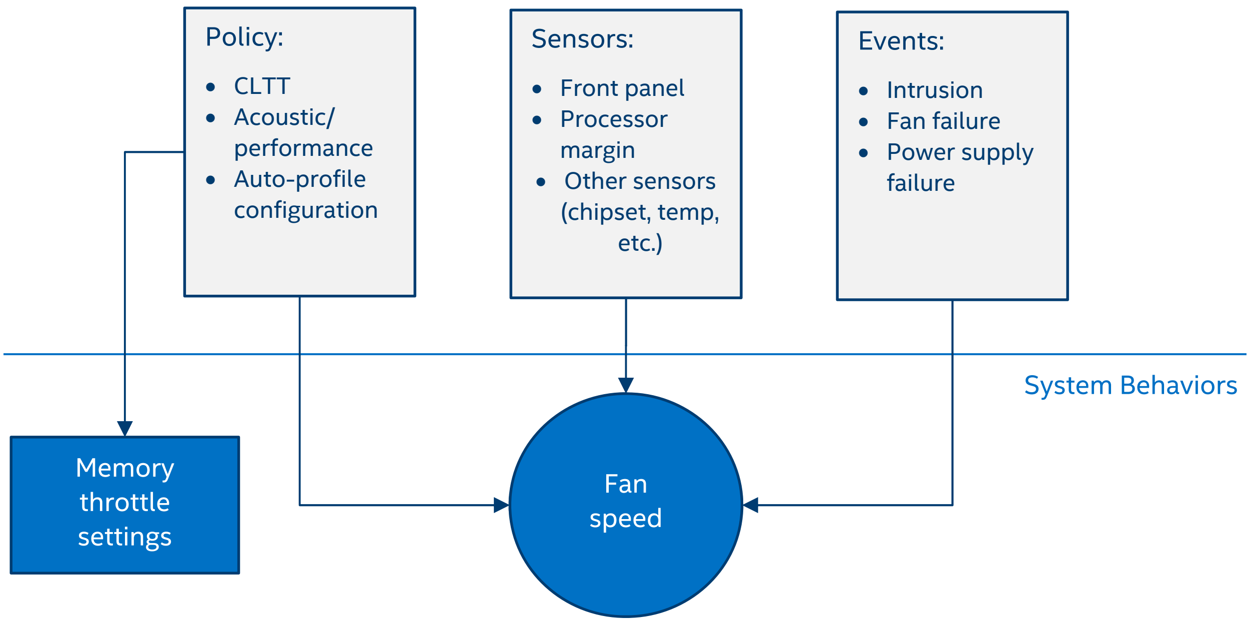

Figure 24. High Level Fan Speed Control Model

4.3 System Fans

Six dual rotor, 40 x 56mm system fans and an embedded fan for each installed power supply 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 rotor 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 rotor is in a failed state.

The system includes two system fan assemblies of three dual rotor fans each. The fan assemblies are held in place by fitting them over mounting pins coming up from the chassis base.

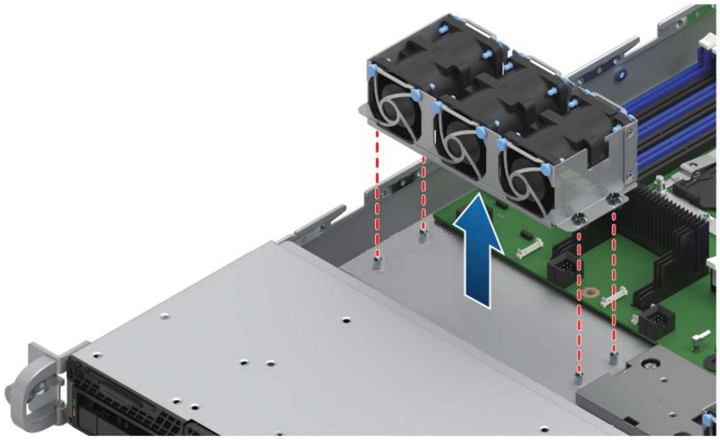

Figure 25. System Fan Assembly

- System fans are NOT hot-swappable
- Each fan and fan assembly is designed for tool-less insertion and extraction from the system. For instructions on fan replacement, see the *Intel® Server System R1000WF System Integration and Service Guide*
- Each fan and fan assembly incorporates vibration dampening features used to minimize fan vibration affects within the chassis
- Fan speed for each fan is controlled by the integrated BMC on the server board. As system thermals fluctuate high and low, the integrated BMC firmware will increase and decrease the speeds to specific fans to regulate system thermals
- Each fan has a tachometer signal for each rotor that allows the Integrated BMC to monitor their status
- Each system fan includes a fault LED located near each system fan connector on the server board
- Each fan has a 10-pin wire harness that connects to a matching connector on the server board

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.

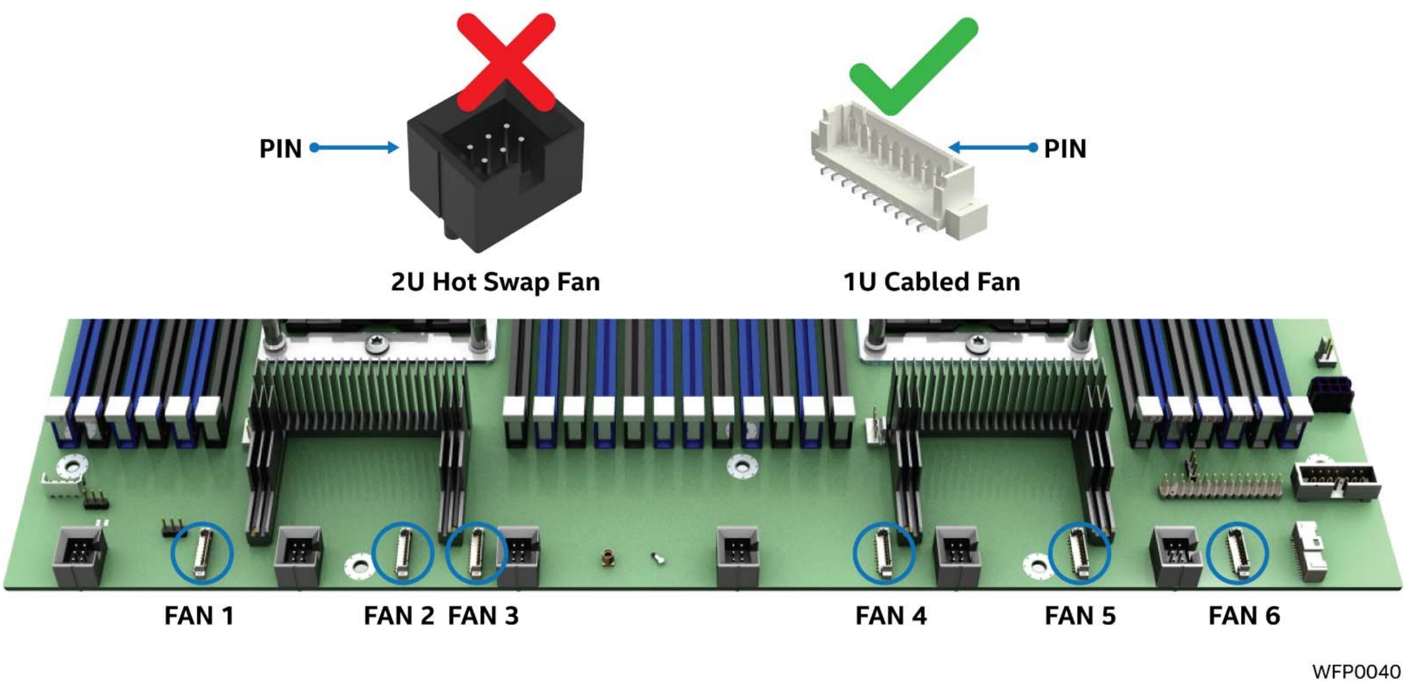

Figure 26. System Fan Connector Locations on a 1U Server Board

Table 29. System Fan Connector Pin-out

SYS_FAN 1		SYS_FAN 2		SYS_FAN 3	
Signal Description	**Pin#**	**Signal Description**	**Pin#**	**Signal Description**	**Pin#**
FAN_TACH1	1	FAN_TACH3	1	FAN_TACH5	1
FAN_PWM0	2	FAN_PWM1	2	FAN_PWM2	2
P12V_CPU_DIMM	3	P12V_CPU_DIMM	3	P12V_CPU_DIMM	3
P12V_CPU_DIMM	4	P12V_CPU_DIMM	4	P12V_CPU_DIMM	4
FAN_TACH0	5	FAN_TACH2	5	FAN_TACH4	5
GROUND	6	GROUND	6	GROUND	6
GROUND	7	GROUND	7	GROUND	7
FM_SYS_FAN0_PRSNT_N	8	FM_SYS_FAN1_PRSNT_N	8	FM_SYS_FAN2_PRSNT_N	8
LED_FAN_FAULT0_R	9	LED_FAN_FAULT1_R	9	LED_FAN_FAULT2_R	9
LED_FAN0	10	LED_FAN1	10	LED_FAN2	10
SYS_FAN 4		SYS_FAN 5		SYS_FAN 6	
Signal Description	Pin#	Signal Description	Pin#	Signal Description	Pin#
FAN_TACH7	1	FAN_TACH9	1	FAN_TACH11	1
FAN_PWM3	2	FAN_PWM4	2	FAN_PWM5	2
P12V_CPU_DIMM	3	P12V_CPU_DIMM	3	P12V_CPU_DIMM	3
P12V_CPU_DIMM	4	P12V_CPU_DIMM	4	P12V_CPU_DIMM	4
FAN_TACH6	5	FAN_TACH8	5	FAN_TACH10	5
GROUND	6	GROUND	6	GROUND	6
GROUND	7	GROUND	7	GROUND	7
FM_SYS_FAN3_PRSNT_N	8	FM_SYS_FAN4_PRSNT_N	8	FM_SYS_FAN5_PRSNT_N	8
LED_FAN_FAULT3_R	9	LED_FAN_FAULT4_R	9	LED_FAN_FAULT5_R	9
LED_FAN3	10	LED_FAN4	10	LED_FAN5	10

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 will shut 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 will 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 revision available as part of a planned system software update.

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

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. Drive Bay Options

The Intel® Server System R1000WF product family has support for different drive bay options including:

- Up to 8 x 2.5" hot swap NVMe, SAS, or SATA drives
- Up to 4 x 3.5" hot swap SAS or SATA hard disk drives or 2.5" SSDs

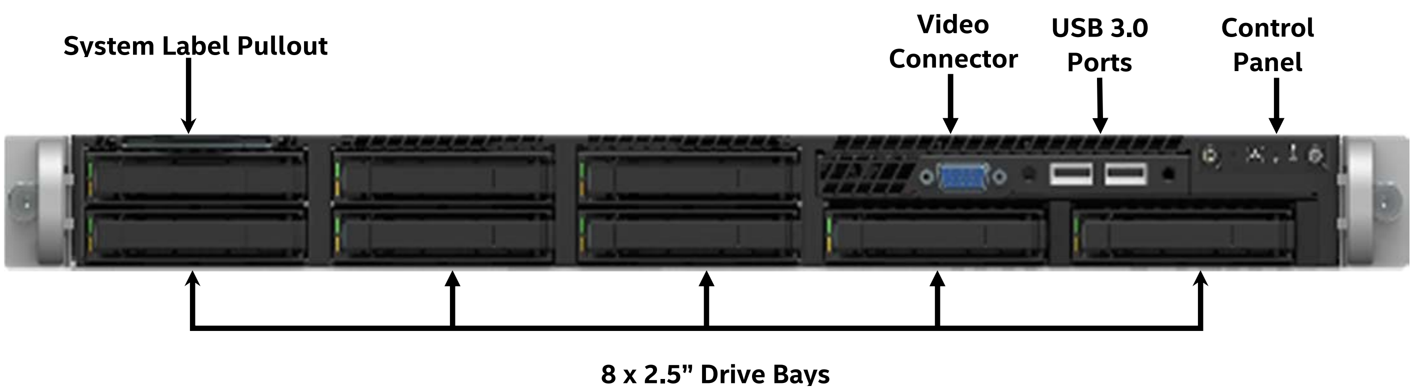

Figure 27. 8x2.5" Drive Bay Configuration (Model R1208WFxxx)

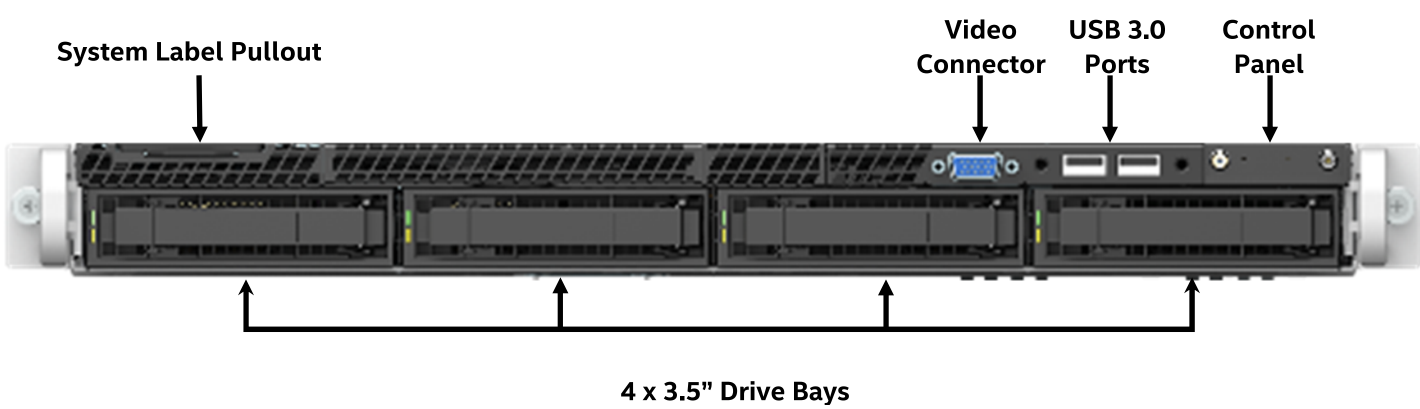

Figure 28. 4x3.5" Drive Bay Configuration (Model R1304WFxxx)

5.1 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 29. Drive Carrier Removal

Note: To ensure proper system air flow requirements, all front drive bays must be populated with a drive tray. Drive trays 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.

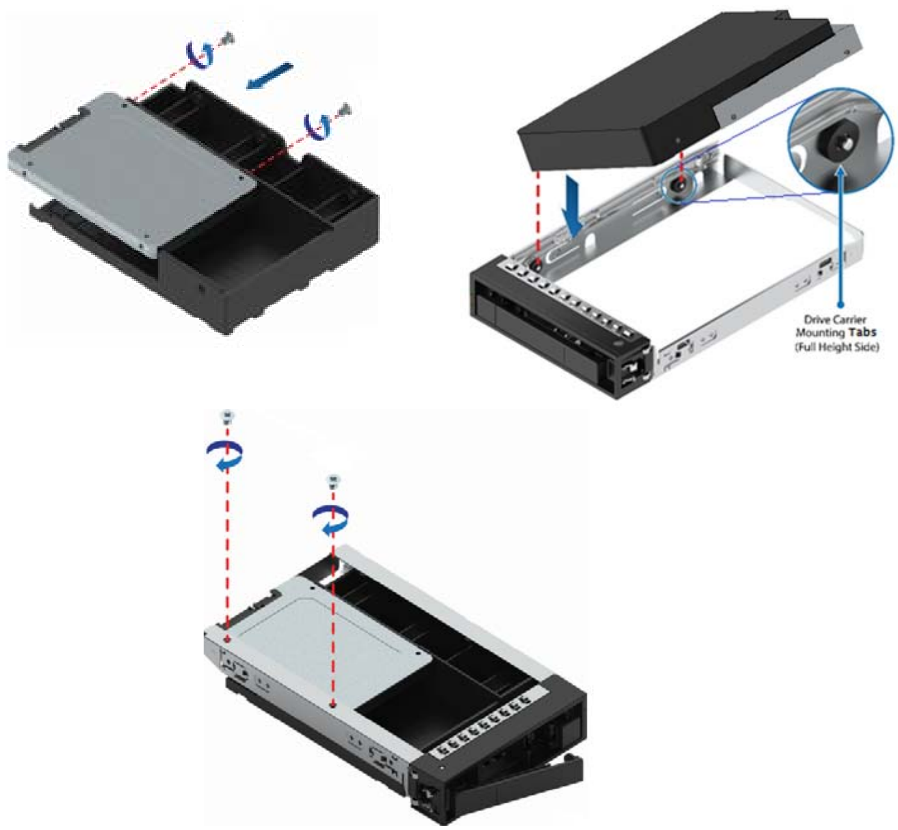

Figure 30. 2.5" SSD Mounted to a 3.5" Drive Tray

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

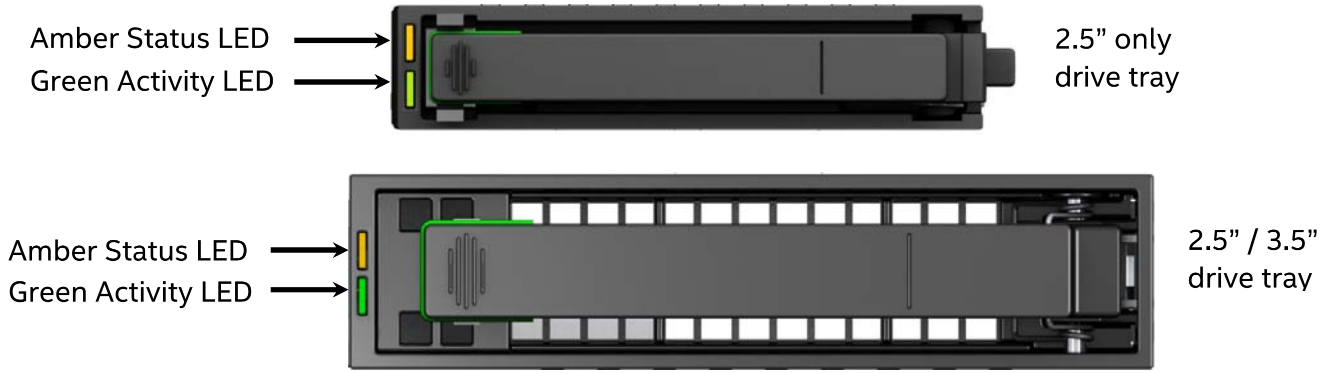

Figure 31. Drive Tray LED Identification

Table 30. Amber Drive Status LED States

Amber	LED State	Drive Status
	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 31. Green Drive Activity LED States

Green	Condition	Drive Type	LED Behavior
	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 32. PCIe SSD Drive Status LED States (VROC)

Amber	LED State	Drive Status
	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.

5.2 Hot Swap Backplane Support

The 1U system has support for two backplane options.

- 8 x 2.5" drive combo backplane with support for SAS/SATA/NVMe drives
- 4 x 3.5" backplane with support for SAS/SATA drives

All available backplane options mount directly to the back of the drive bay as shown in the following illustration.

Figure 32. Backplane Installation

Backplanes include the following features:

- 12 Gb SAS and 6Gb SAS/SATA or slower support
- Drive interface connectors
  - 29-pin SFF-8680 – 3.5" Backplane – SAS/SATA only – 12 Gb rated
  - 68-pin SFF-8639 – 2.5" Backplane – supporting 12 Gb SAS and x4 PCIe* NVMe
- Hot swap drive support
- Cable Connectors
  - SFF-8643 Mini-SAS HD – 2.5" and 3.5" backplanes – 12Gb SAS capable
  - OCuLink* PCIe interface – 2.5" backplane
  - 1x5-pin connector – I2C interface for device status communication to the BMC over slave SMBus
  - 2x2- in connector - Power
- SGPIO SFF-8485 interface embedded within the sideband of the mini-SAS HD connectors
- HSBP microcontroller – Cypress* CY8C22545-24AXI Programmable System-on-Chip (PSoC*) device
- LEDs to indicate drive activity and status for each attached device
- Device presence detect inputs to the microcontroller
- 5V VR for devices
- 3.3V VR for microcontroller
- In-application microcontroller FW updateable over the I2C interface
- FRU EEPROM support
- Temperature sensor through the use of a TMP75 (or equivalent) thermistor implementation with the microcontroller

5.2.1 SGPIO Functionality

Backplanes include support for an 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.

5.2.2 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 will generate SEL events by monitoring registers on the HSBP microcontroller for DRIVE PRESENCE, FAULT, and RAID REBUILD in progress.

5.3 4 x 3.5" Drive SATA/SAS Hot-Swap Backplane

Intel Spare Product Code: FR1304S3HSBP

All 3.5" drive system SKUs within the product family will ship with a 4 x drive backplane capable of supporting 12 Gb/sec SAS and 6 Gb/sec SAS / SATA drives. Both hard disks and Solid State Devices (SSDs) can be supported within a common backplane. Each backplane can support either SATA or SAS devices. However, mixing of SATA and SAS devices within a common hot swap backplane is not supported. Supported devices are dependent on the type of host bus controller driving the backplane: SATA only or SAS.

The front side of the backplane includes, 4 x 29-pin SFF-8680 drive interface connectors, each capable of supporting 12 Gb SAS or 6 Gb SAS/SATA. The connectors are numbered 0 thru 3. Signals for all four drive connectors are routed to a single multi-port, mini-SAS HD SFF-8643 connector on the back side of the backplane.

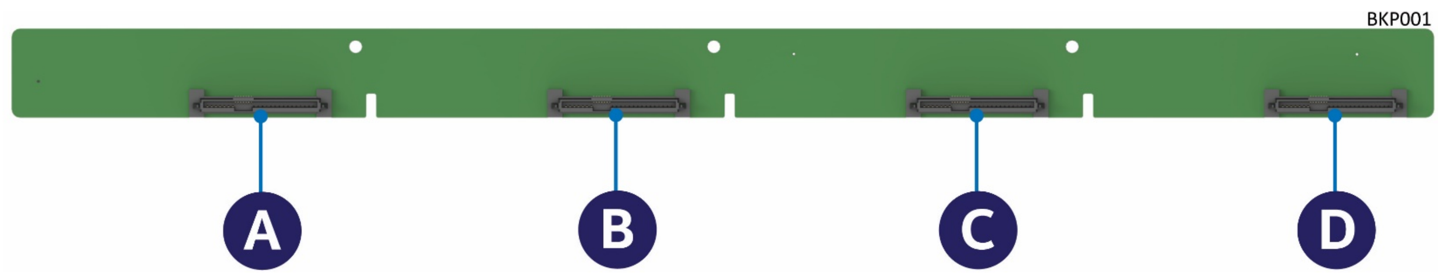

Label	Description
A	HDD_0
B	HDD_1
C	HDD_2
D	HDD_3

Figure 33. 4 x 3.5" Drive Hot-Swap Backplane – Front View

On the backside of the backplane are several connectors. The following illustration identifies each.

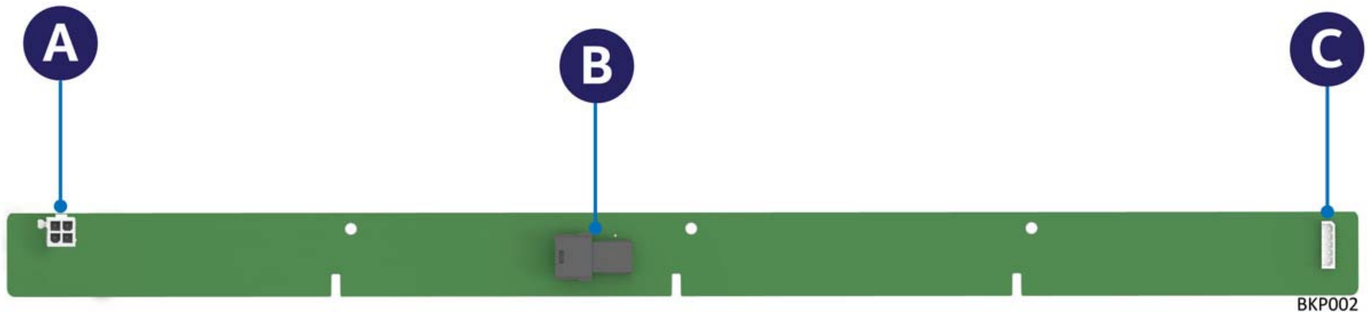

Label	Description
A	Power Harness Connector
B	Multi-port, Mini-SAS Cable Connector
C	I2C Cable Connector

Figure 34. 4 x 3.5" Drive Hot-Swap Backplane – Rear View

A – Power Harness Connector – 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.

B – Multi-port Mini-SAS Cable Connector – The backplane includes one multi-port mini-SAS cable connector providing I/O signals for four SAS/SATA drives on the backplane. A cable can be routed from matching connectors on the server board or add-in SAS/SATA RAID cards.

C – I2C Cable Connector – The backplane includes a 1x5 cable connector used as a management interface to the server board.

5.4 8 x 2.5" Drive SATA / SAS / NVMe* Combo Backplane

Intel Spare Product Code: **F1U8X253PHS**

All R1208…. 2.5" drive systems within the Intel® Server R1000WF product family will include an eight (8) drive combo backplane capable of supporting 12 Gb/sec SAS, 6 Gb/sec SATA and PCIe* NVMe drives.

The 8x2.5" drive combo backplane has support for different drive configurations including SAS or SATA only, NVMe only, or a combination of both SAS and NVMe. Mixing of SATA and SAS devices within a common hot swap backplane is not supported. Hard Disk Drives (HDDs) and Solid State Drives (SSDs) can be supported within a common backplane.

The front side of the backplane includes 8 x 68-pin SFF-8639 drive interface connectors, each capable of supporting SAS, SATA, or PCIe* NVMe drives. The connectors are numbered 0 thru 7.

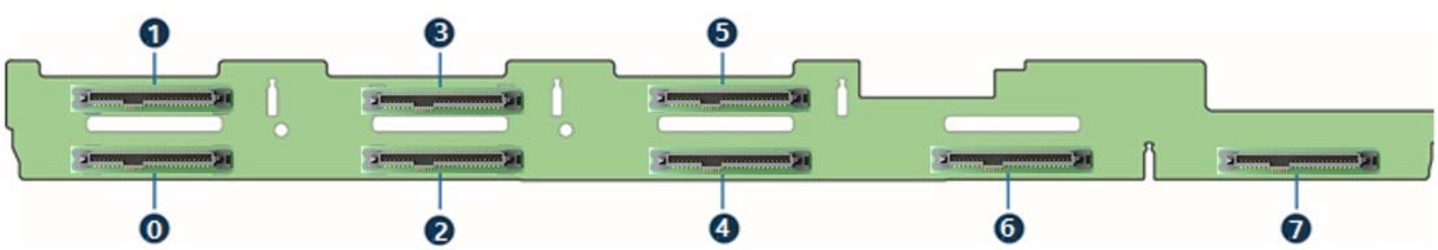

Figure 35. 8 x 2.5" Drive SAS/SATA/NVMe Backplane – Front View

The backplane and backplane management features can support drives installed in any order when populating mixed SATA/NVMe or SAS/NVMe drive configurations. However, when mixing NVMe and SATA/SAS within a common backplane, Intel recommends drive types be populated together and not in an intermixed order. In addition, drive population rules must be followed and drive support limitations exist when the optional Intel® VROC upgrade key is installed to the server board providing support for NVMe RAID and NVMe Management features. See Section 6.4.4.

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.

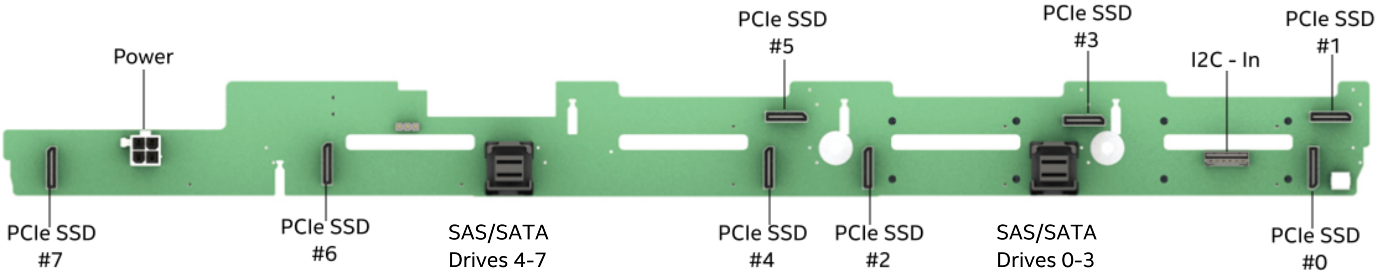

Figure 36. 8 x 2.5" Drive SAS/SATA Backplane – Rear View

Power Connector – 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.

PIN	SIGNAL	SIGNAL	PIN
1	GND	P12V	3
2	GND	P12V	4

I2C Cable Connector – The backplane includes a 1x5 cable connector used as a management interface to the server board.

PIN	SIGNAL
1	SMB_3V3SB_DAT
2	GND
3	SMB_3V3SB_CLK
4	SMB_ADD0
5	SMB_ADD1

Multi-port Mini-SAS Cable Connectors – The backplane includes two, multi-port mini-SAS HD (SFF-8643) cable connectors, each providing I/O signals for four SAS/SATA hard drives on the backplane. Cables can be routed from matching connectors on the server board for SATA only support, or from an add-in SAS/SATA RAID card.

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

See section 6.4 for NVMe support information.

6. Storage Controller Options Overview

The Intel® Server System R1000WF 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

6.1 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

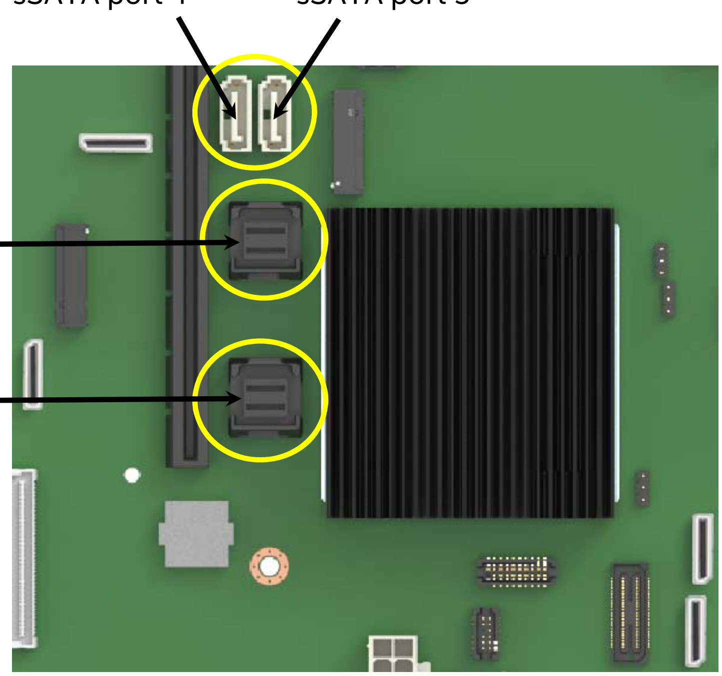

Figure 37. Onboard Storage Support Features

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

Table 33. SATA and sSATA Controller Feature Support

Feature	Description	AHCI Mode	RAID Mode RSTe	RAID Mode 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 <F2> BIOS Setup Utility under the “Mass Storage Controller Configuration” menu screen. The following table identifies supported setup options.

Table 34. SATA and sSATA Controller BIOS Utility Setup 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

6.1.1 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 Setup Mass Storage Controller Configuration screen found in the <F2> BIOS Setup Utility.

6.2 Onboard SATA Software RAID support

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

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

By default, onboard RAID options are set to DISABLED in <F2> BIOS Setup. To enable onboard RAID support, access the <F2> 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

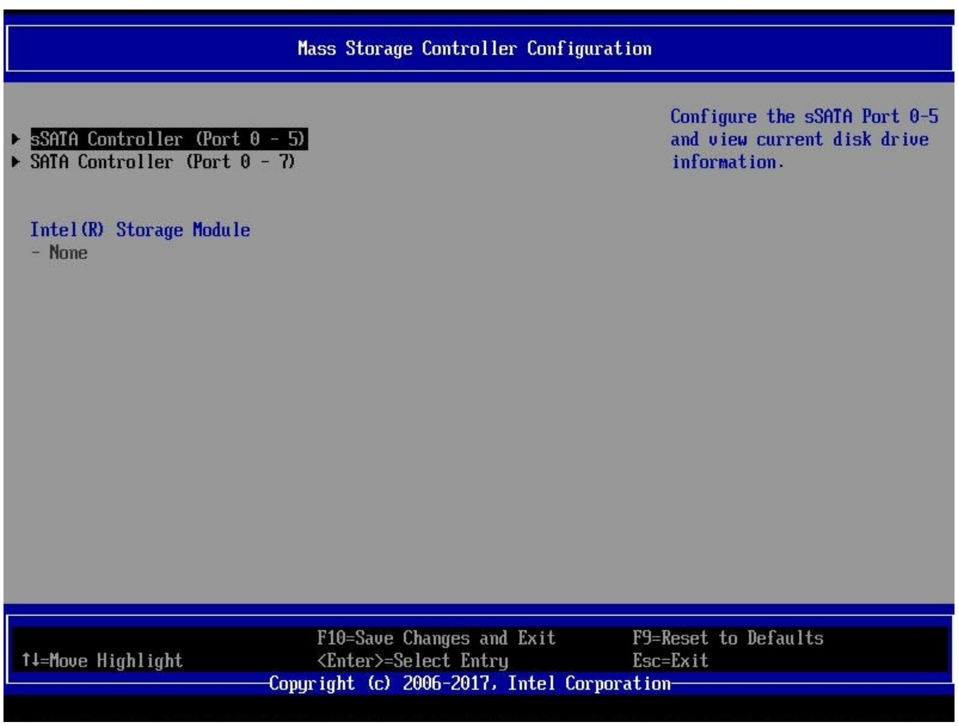

Note: RAID partitions created using either RSTe or ESRT2 cannot span across the two embedded SATA controllers. Only drives attached to a common SATA controller can be included in a RAID partition.

6.2.1 Intel® Rapid Storage Technology (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 SW-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.

6.2.2 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
  - Enabled with the addition of an optionally installed ESRT2 SATA RAID 5 Upgrade Key (**iPN - RKSATA4R5**)

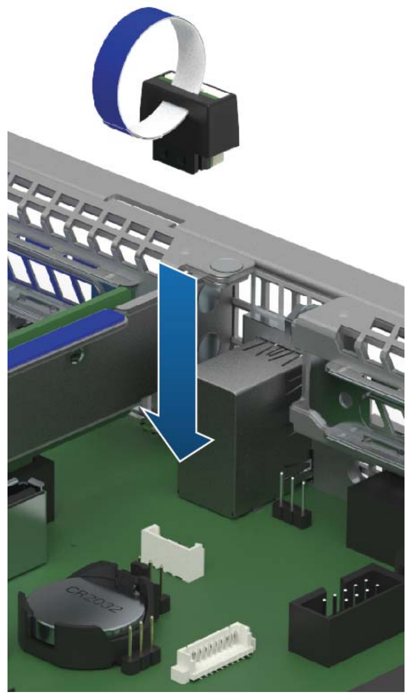

Figure 38. ESRT2 SATA RAID-5 Upgrade Key

  - **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.

6.3 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.

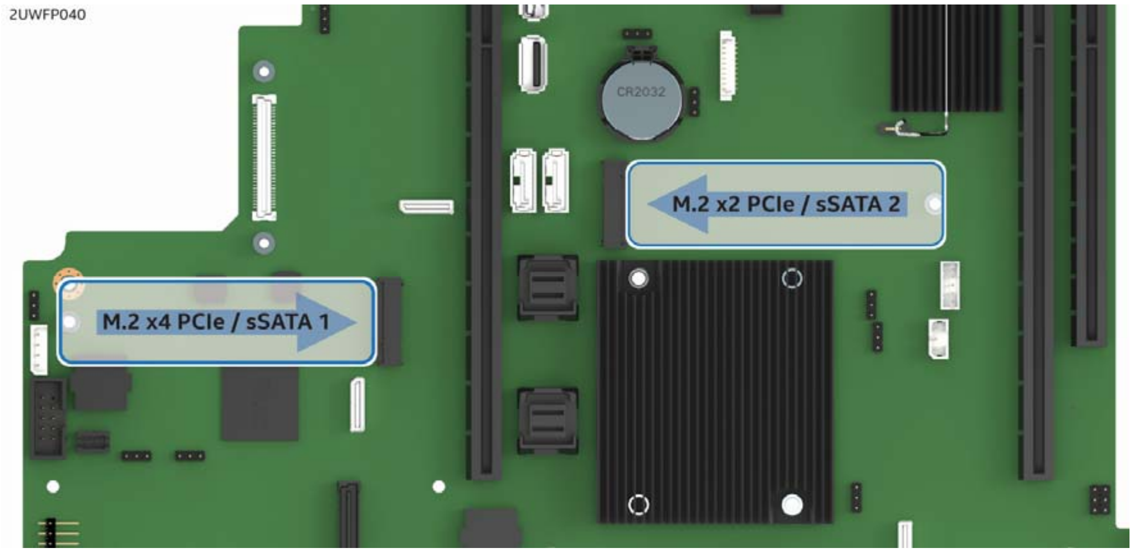

Figure 39. M.2 Module Connector Locations

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 Sot #1is supported by **PCIe x2** bus lanes and **sSATA-2** from the chipset embedded sSATA controller.

6.3.1 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.
    - **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 onboard M.2 connectors are routed from the Intel chipset (PCH).
    - **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*

6.4 PCIe NVMe* Drive Support

The eight drive combo backplane used within all R1208…. system models of the Intel® Server System R1000WF product family has support for up to eight (8) PCIe NVMe 2.5" Small Form Factor (SFF) SSDs. 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)

6.4.1 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.

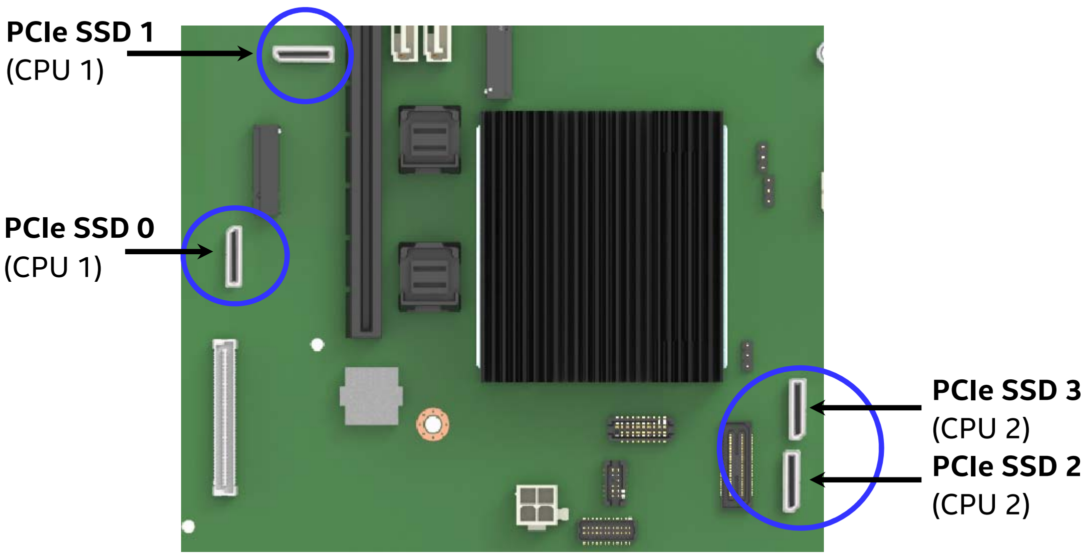

Figure 40. Onboard OCuLink Connectors

6.4.2 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.

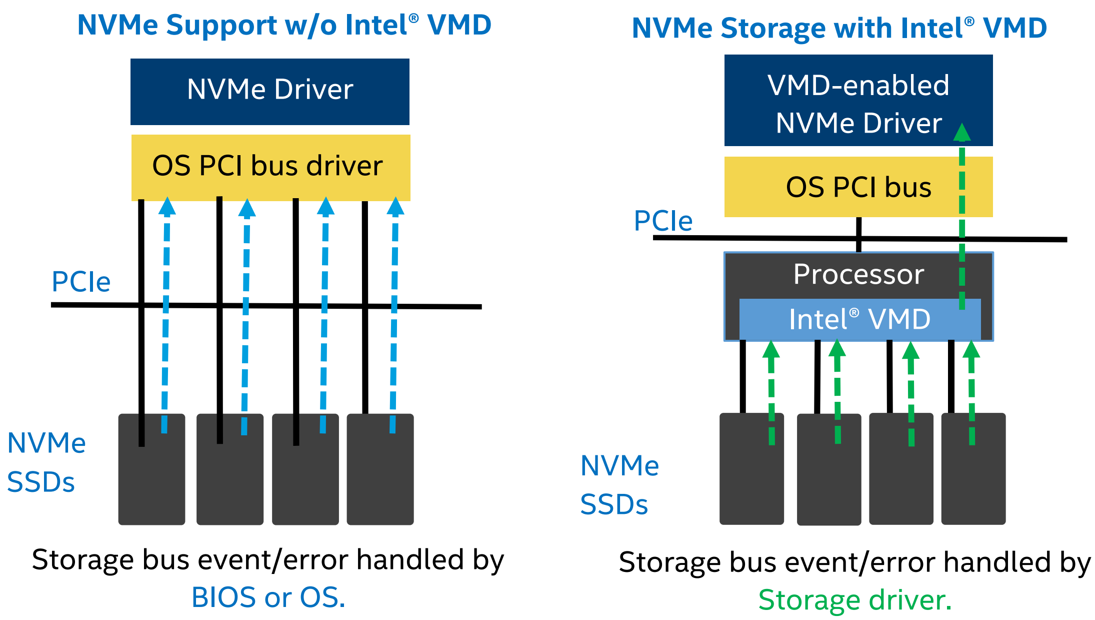

Figure 41. Storage Bus Event/Error Handled by BIOS or OS/ Storage Drive

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 6.4.3 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.

6.4.2.1 Enabling VMD support

In order for installed NVMe devices to utilize the VMD features of the server board, VMD must be ENABLED on the appropriate CPU PCIe* Root Ports in <F2> BIOS Setup. By default, VMD support is DISABLED on all CPU PCIe* root ports in <F2> BIOS Setup.

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

Table 35. CPU - PCIe* Port Routing

CPU 1		CPU 2	
PCI Ports	Onboard Device	PCI Ports	Onboard 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 39. Riser Slot #1 & #2 – PCIe* Root Port Mapping

In <F2> BIOS Setup, the Intel VMD support menu can be found under the following BIOS Setup menu options:

ADVANCED -> PCI CONFIGURATION -> VOLUME MANAGEMENT DEVICE

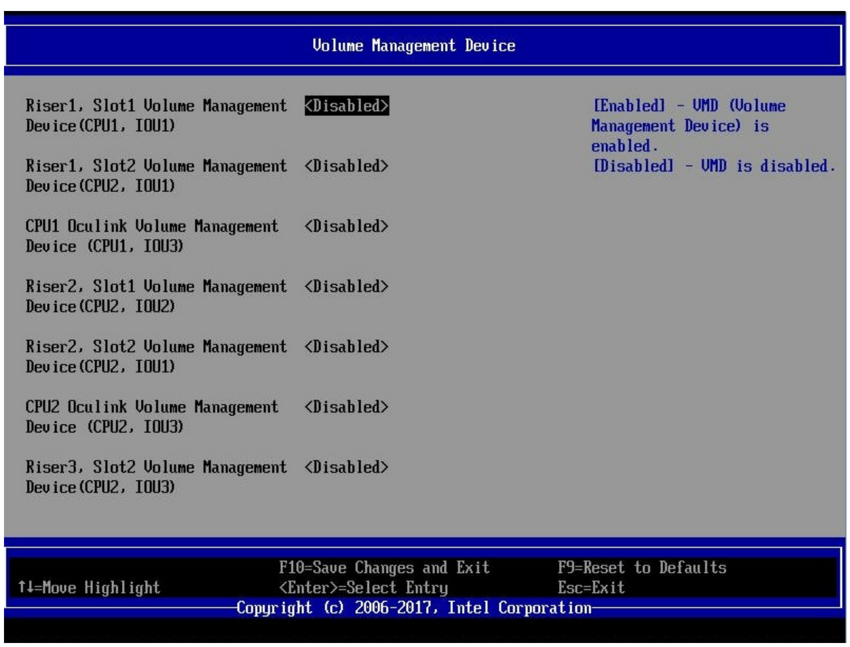

Figure 42. VMD Support Disabled in <F2> BIOS Setup

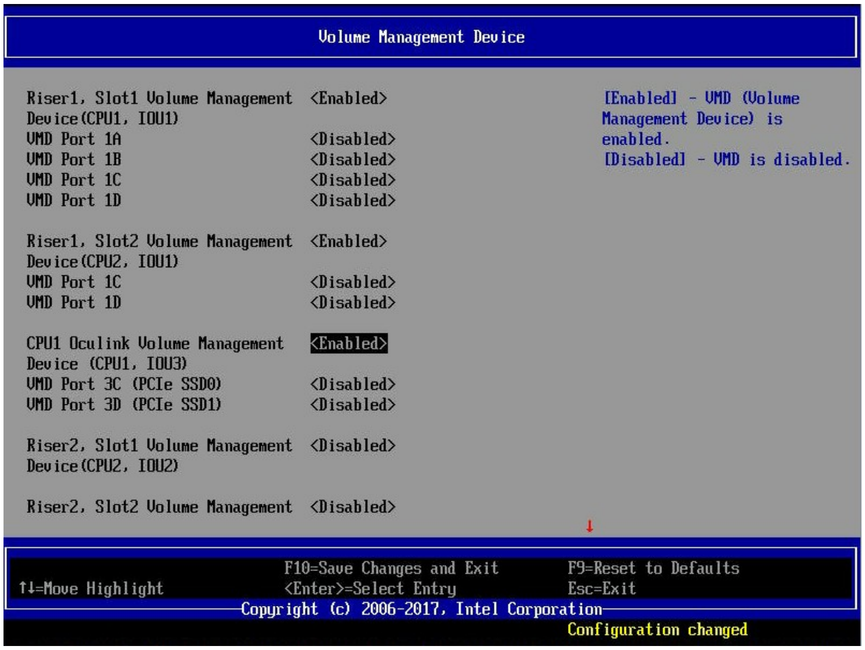

Figure 43. VMD Support Enabled in <F2> BIOS Setup

6.4.3 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)

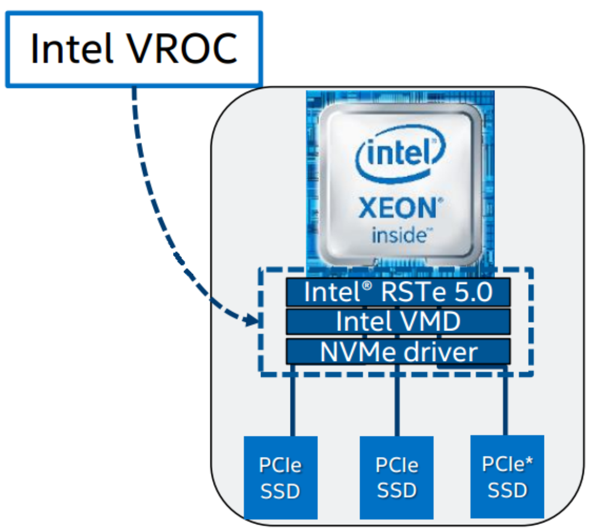

Figure 44. Intel® VROC Basic Architecture Overview

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 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 the following illustration.

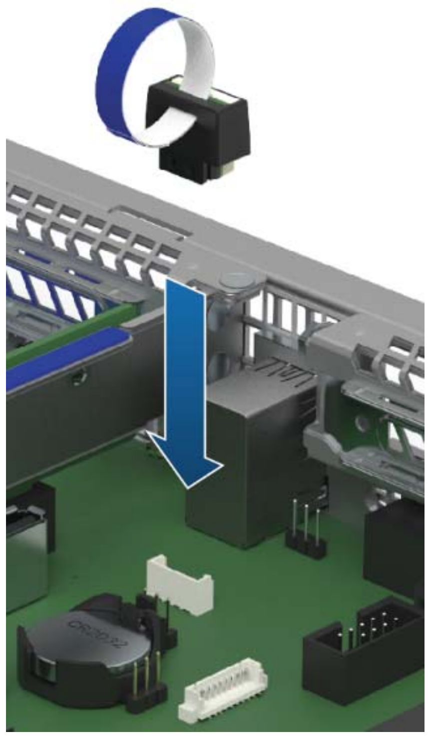

Figure 45. Intel® VROC Upgrade Key

The following table identifies available Intel® VROC upgrade key options.

Table 36. 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 35 are used for PCIe* NVMe SSDs only. For SATA RAID support, see Section 6.2.

6.4.4 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 Server Board OCuLink connectors and/or optional add-in PCIe Switch. 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.

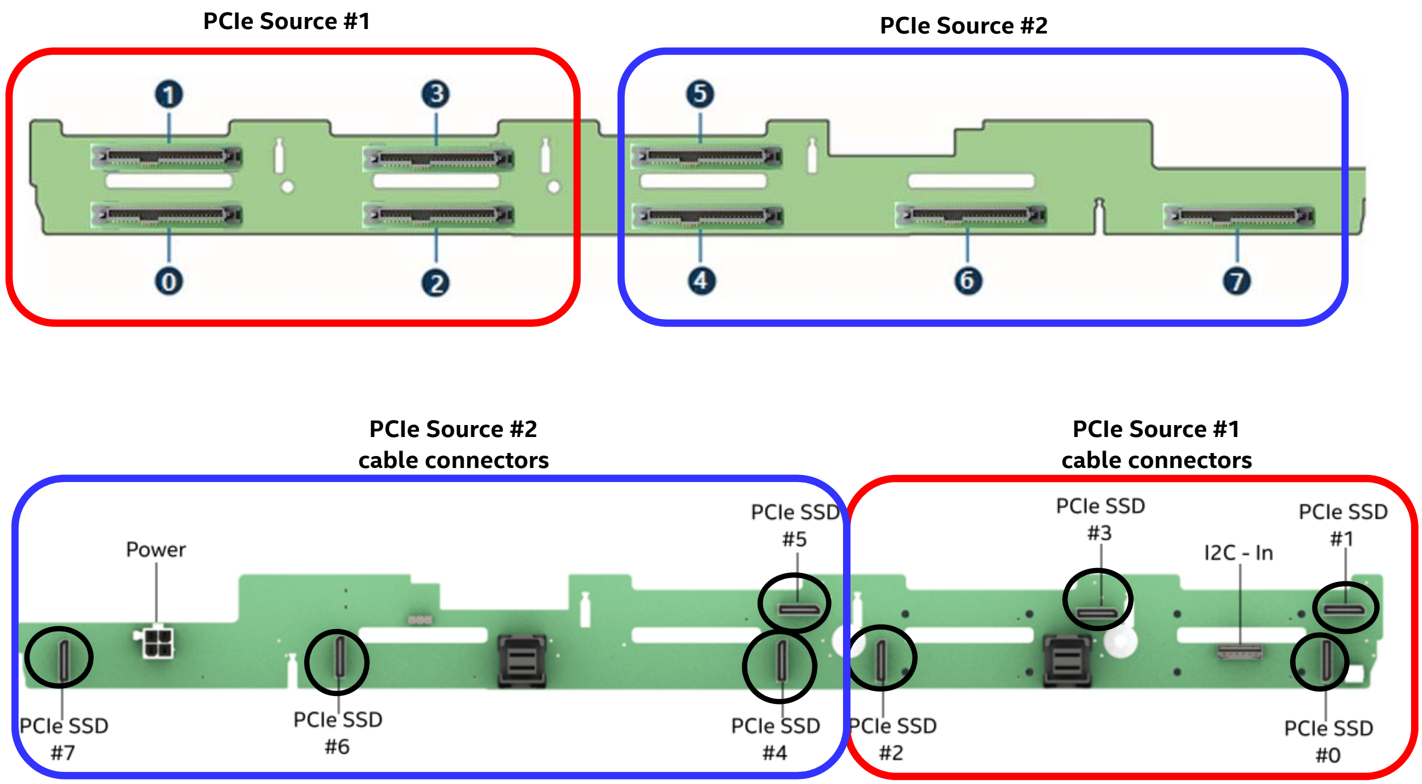

Figure 46. Backplane Cabling from Two PCIe Sources

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

6.4.4.1 Onboard PCIe OCuLink Connectors 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.

- 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.

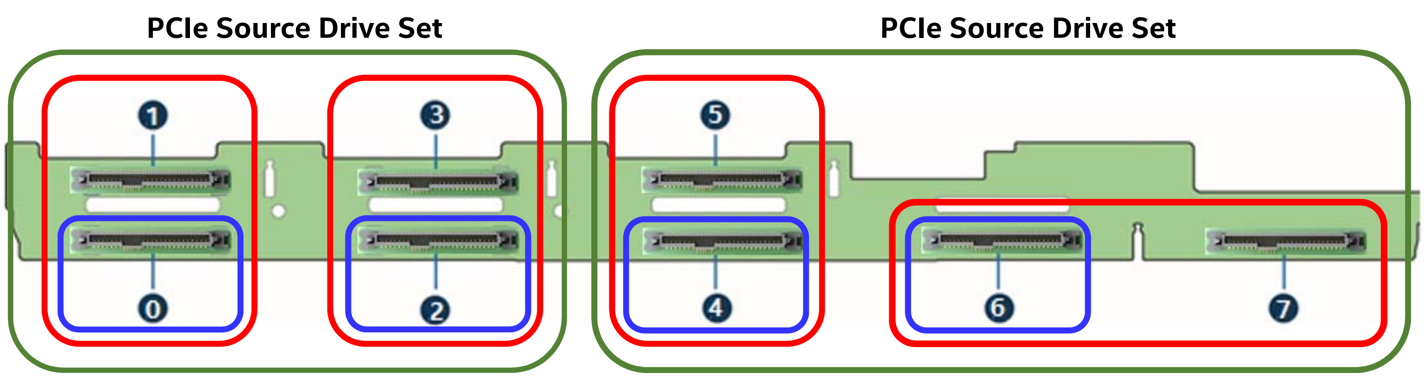

The following illustrations identify supported and unsupported drive populations associated with any defined drive pair of the 8x2.5" combo backplane when Intel VROC is used for NVMe drive management and the PCIe source to the backplane is from any of the onboard OCuLink connectors.

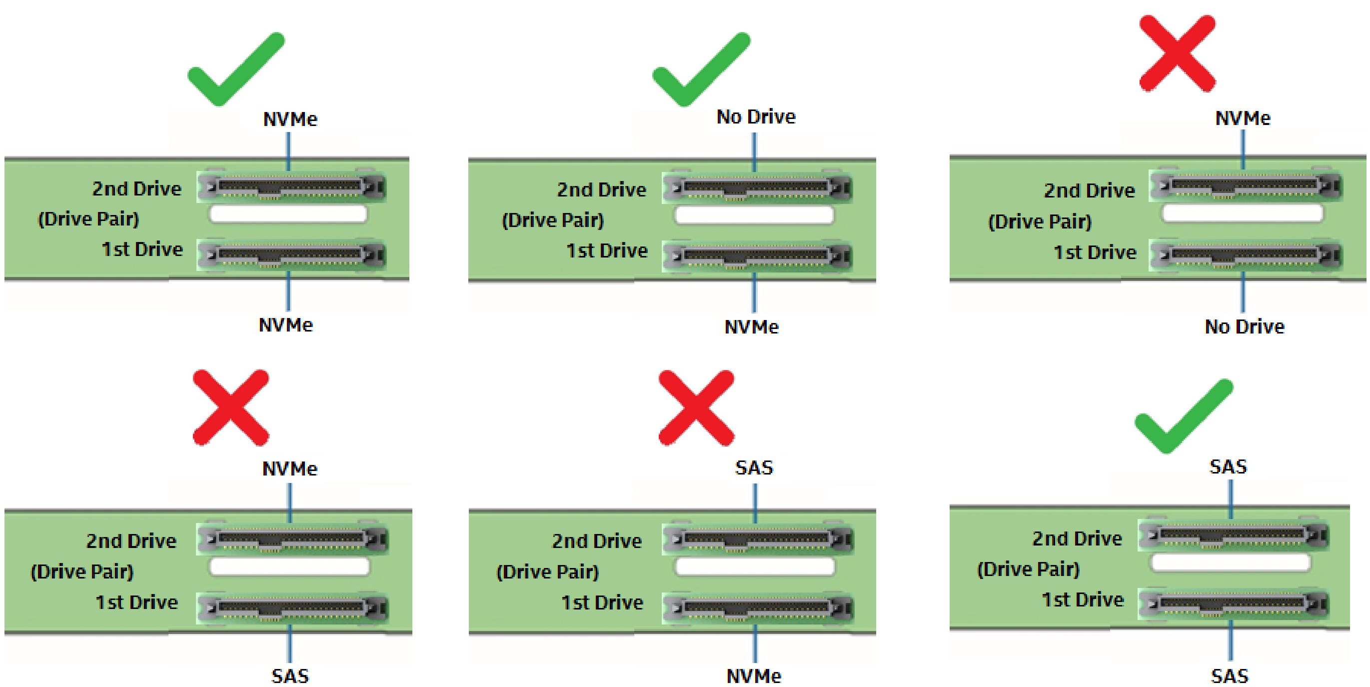

Where 1st Drive = drive connectors 0, 2, 4, or 6 and 2nd Drive = drive connectors 1, 3, 5, or 7

6.4.4.2 4 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 a 4 port PCIe Switch add-in card option.

- 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.

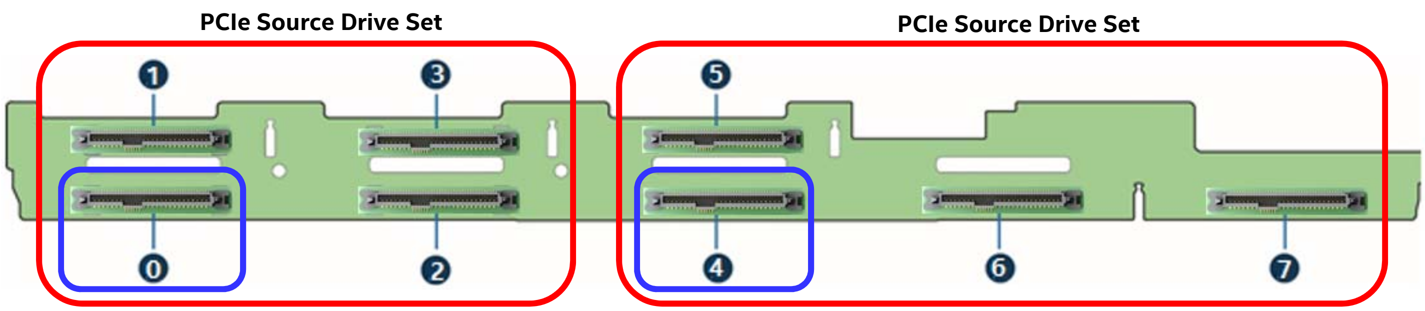

The following illustrations identify supported and unsupported drive populations associated with any defined drive set of the 8x2.5" combo backplane when an Intel® VROC key is installed to the server board and the PCIe source to the backplane is from an add-in PCIe Switch.

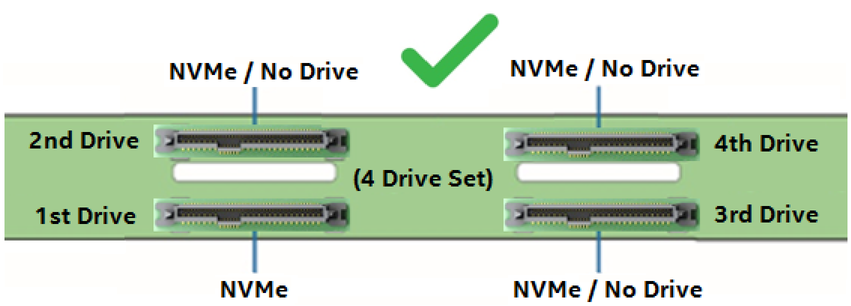

1st Drive = NVMe + NVMe in any sequential drive slot (No gaps)

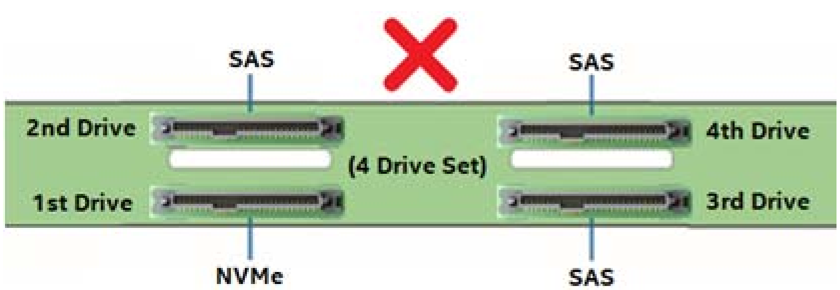

Mixing NVMe and SAS within a common drive set is not supported

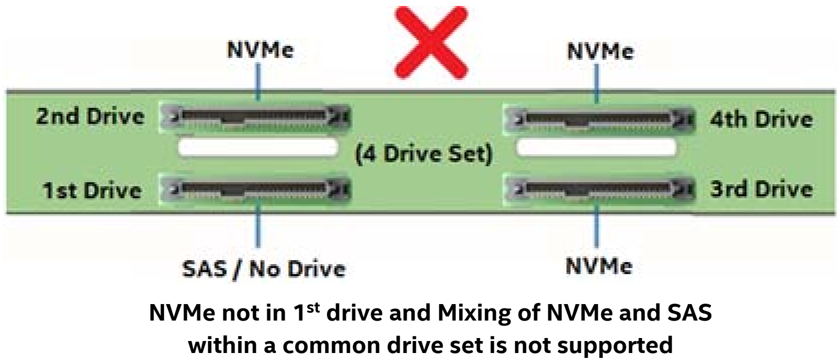

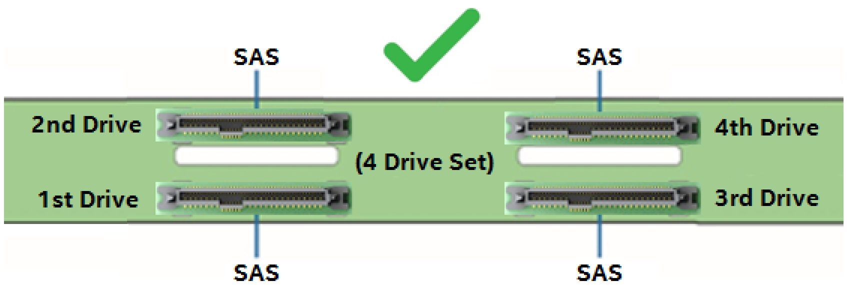

Where 1st Drive = drive connectors 0 or 4 on the backplane

6.5 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.

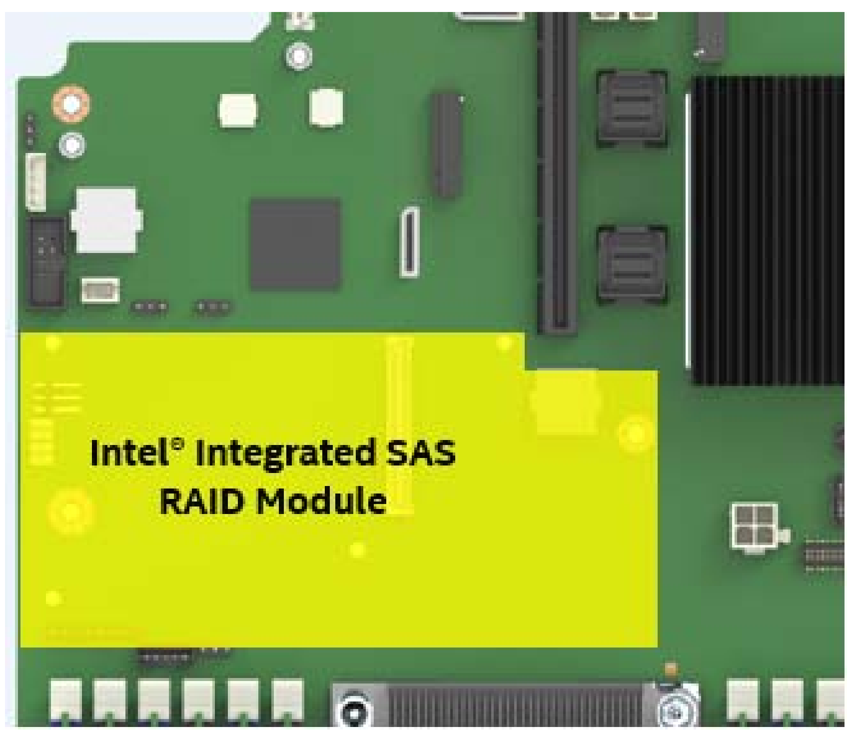

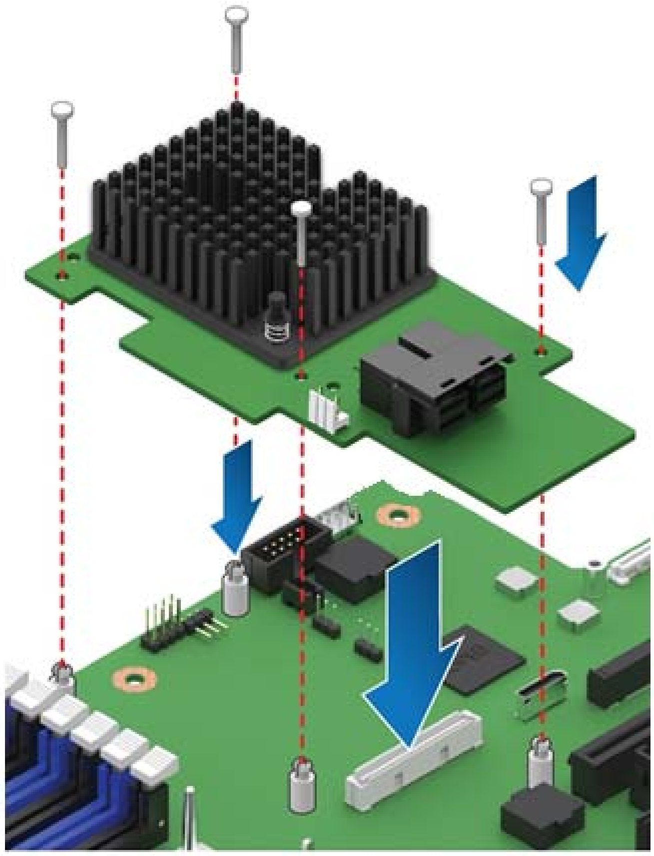

Figure 47. Intel® Integrated RAID Module

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.

6.5.1 Intel® RAID Maintenance Free Backup Unit (RMFBU) Support

The 1U system has support for one or two Intel® RAID Maintenance Free Backup Units (RMFBU).

Figure 48. Support for single Intel® RAID Maintenance Free Backup Unit (Standard Option)

Support for two RMFBUs will require the use of an optional bracket capable of supporting stacked RMFBUs.
Intel Accessory Kit order code – AWTAUXBBUBKT

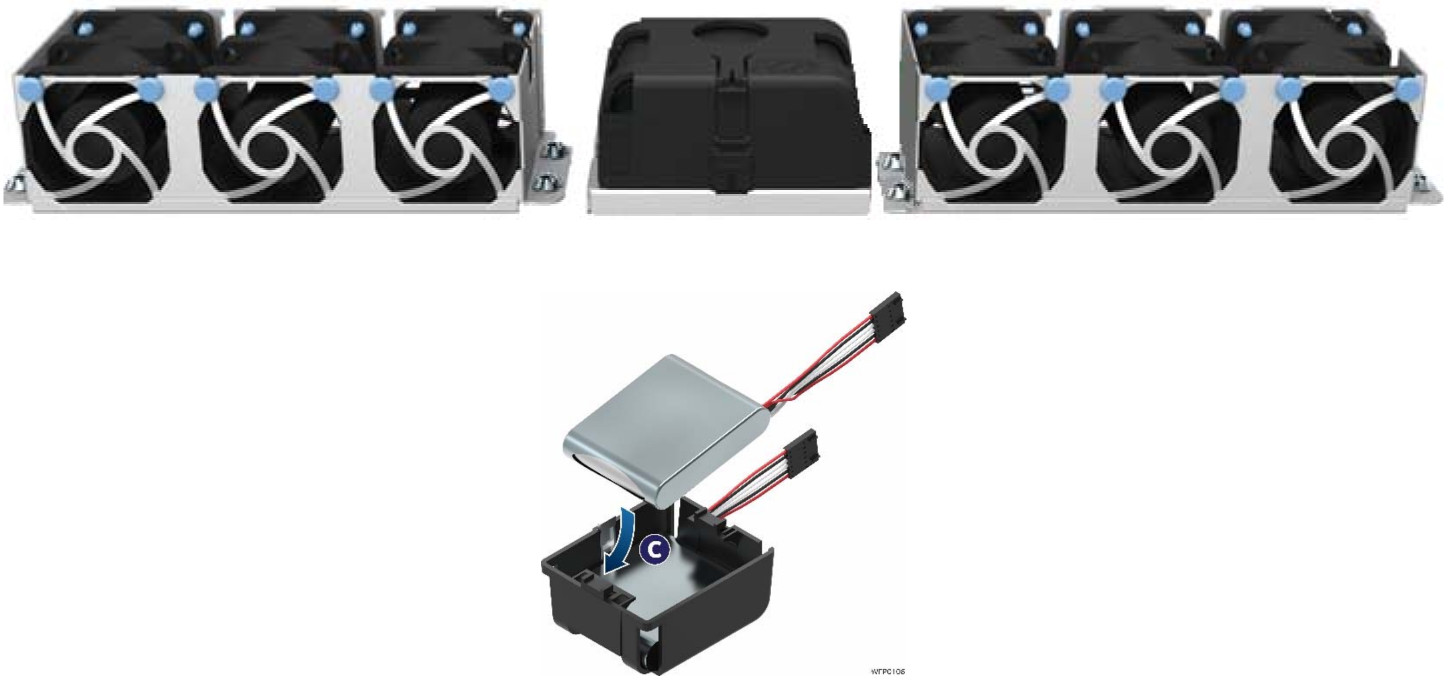

Figure 49. Support for dual Intel® RAID Maintenance Free Backup Units (Optional Accessory)

7. Front I/O Panel and Control Panel Overview

All system configurations include an I/O Panel and Control Panel and on the front of the system. This section describes the features and functions of both.

7.1 I/O Panel Features

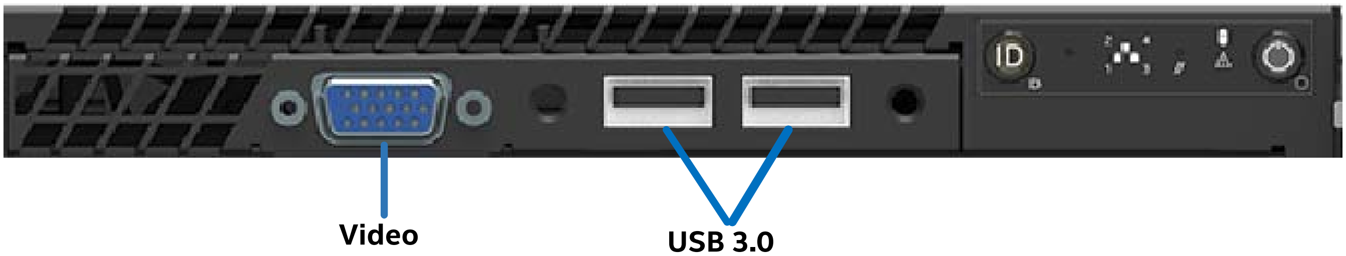

Figure 50. Front I/O Panel Features

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 onboard 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.

7.2 Control Panel Features

The system includes a control panel that provides push button system controls and LED indicators for several system features. This section will provide a description for each front control panel feature.

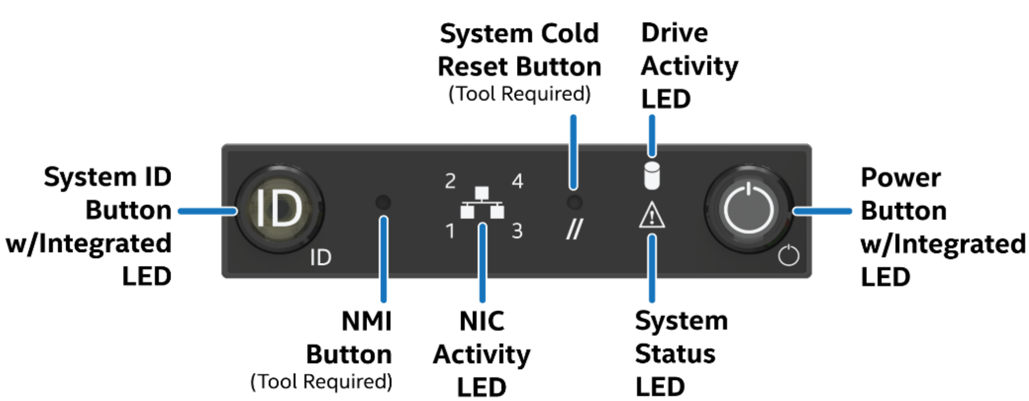

Figure 51. Front Control Panel Options

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 will cause 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 an 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 onboard Network Interface Controller (NIC). When a network link is detected, the LED will turn on solid. The LED will blink when 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 will reboot and re-initialize the system. To prevent an inadvertent system reset, the actual Reset 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.

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, the other is located on the back edge of the server board, viewable from the back of the system. Both LEDs are tied together and will show the same state. The System Status LED states are driven by the onboard platform management subsystem. The following table provides a description of each supported LED state.

Table 37. 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 will send a signal to the integrated BMC, which will either power on or power 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 38. 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	Steady On	System and operating system are up and running.

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

8. PCIe* Riser Card Support

The server board provides three riser card slots identified as: Riser Slot #1, Riser Slot #2, and Riser Slot #3. In a 1U system, PCIe* Add-in cards can only be installed into single slot riser cards installed to Riser Slots #1 and #2. Riser Slot #3 is not supported in a 1U system. Per the PCIe specification, each riser card slot can support a maximum 75W of power. The PCIe* bus interface for each riser card slot is supported by each of the two installed processors.

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.

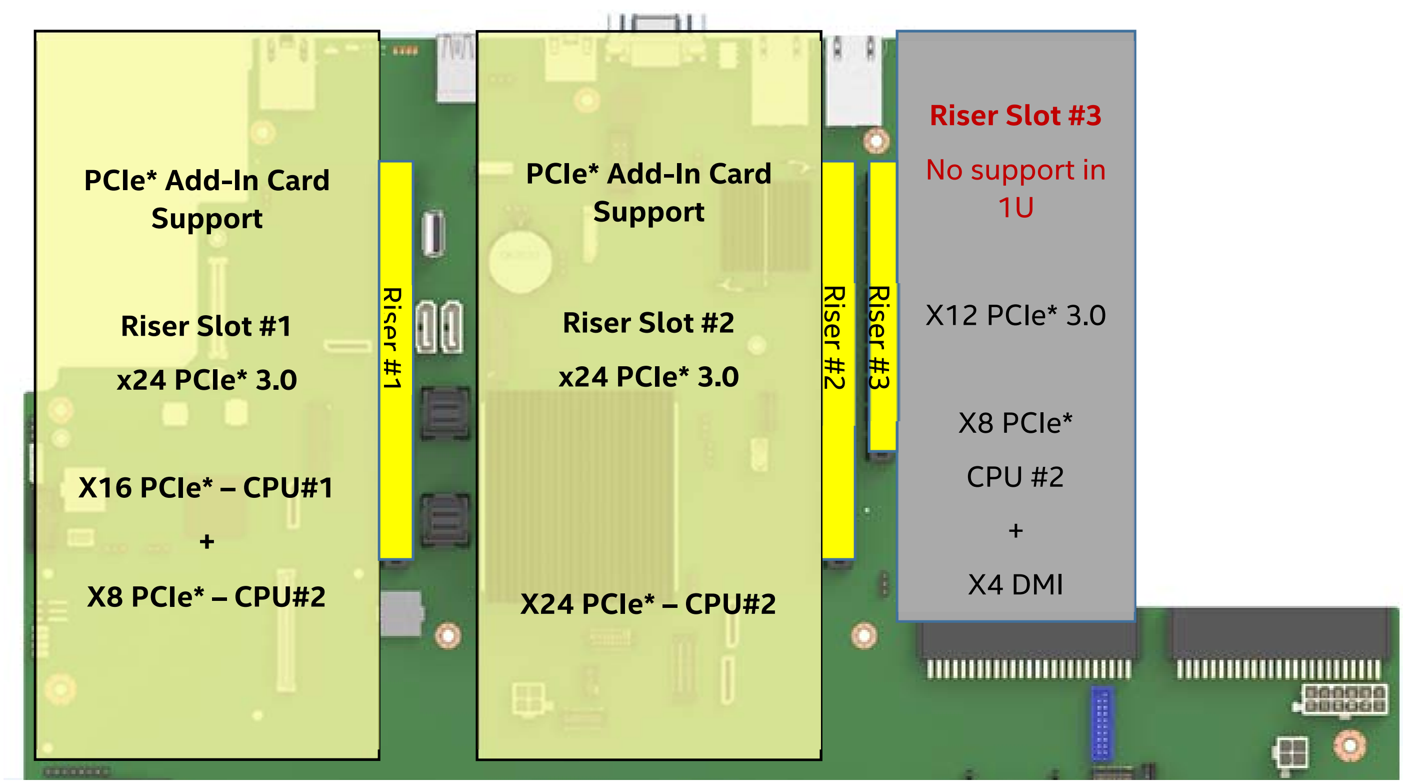

Figure 52. PCIe* Add-in Card Support

The following tables provide the PCIe* bus routing for all supported risers cards.

Table 39. Riser Slot #1 & #2 – PCIe* Root Port Mapping

Riser Slot #1 & #2 – Riser Card Options	
1U - 1-Slot Riser Card, iPC – F1UL16RISER3APP	
PCIe* Riser #1 CPU #1 – Ports 1A thru 1D (x16 elec, x16 mech)	PCIe* Riser #2 CPU #2 – Ports 2A thru 2D (x16 elec, x16 mech)

Available 1U Riser Card:

One PCIe* add-in card slot – PCIe* x16, x16 mechanical

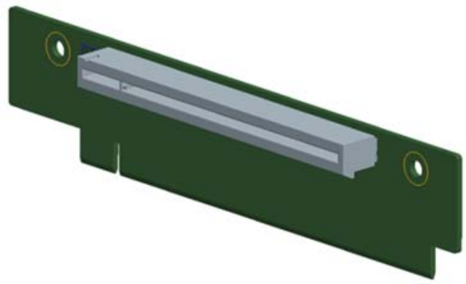

Figure 53. 1U One Slot PCIe* Riser Card (iPC – F1UL16RISER3APP)

Each riser card is mounted to a bracket assembly which is inserted into a riser card slot on the server board.

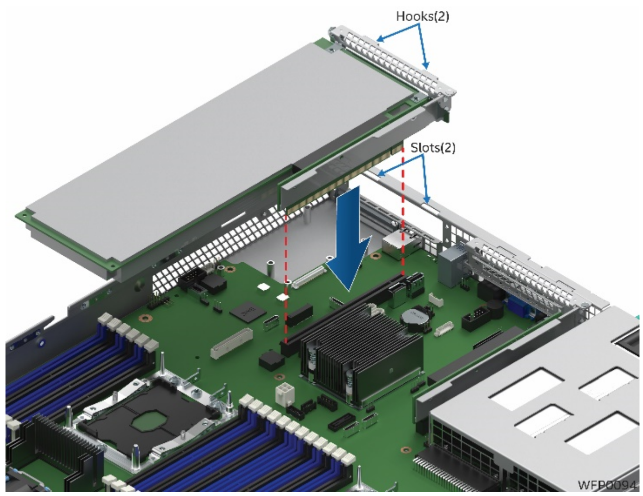

Figure 54. Add-in Card Support

Each riser card assembly has support for a single full height, ½ length PCIe* add-in card.

Note: Add-in cards that exceed the PCI specification for ½ length PCI add-in cards (167.65mm or 6.6in) may interfere with other installed devices on the server board.

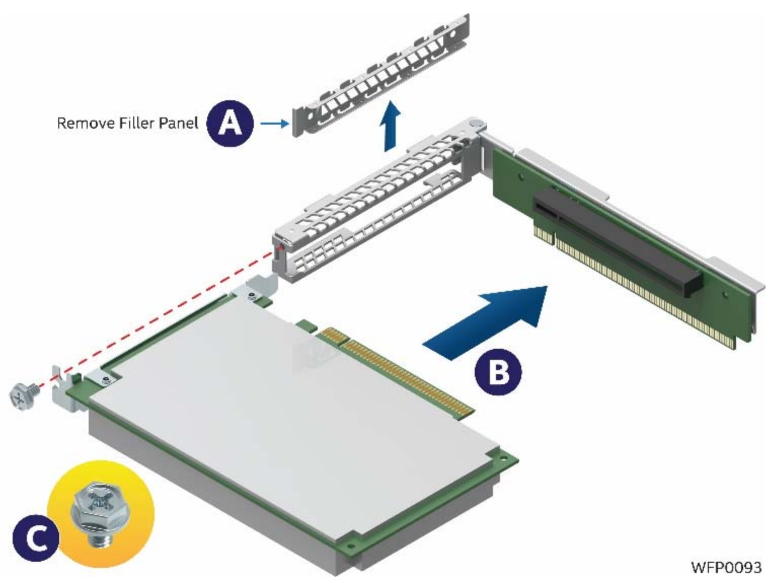

Figure 55. Riser Card Assembly

9. OCP* Compatible Intel Ethernet Network Adapter Support

The Intel® Server Board S2600WF Product Family offers a line of LAN KR OCP mezzanine modules that follows the OCP 2.0 form factor.

The optional OCP mezzanine module can be installed onto the connector (labeled "OCP_IO_Module") on the server board, as shown in the following illustration.

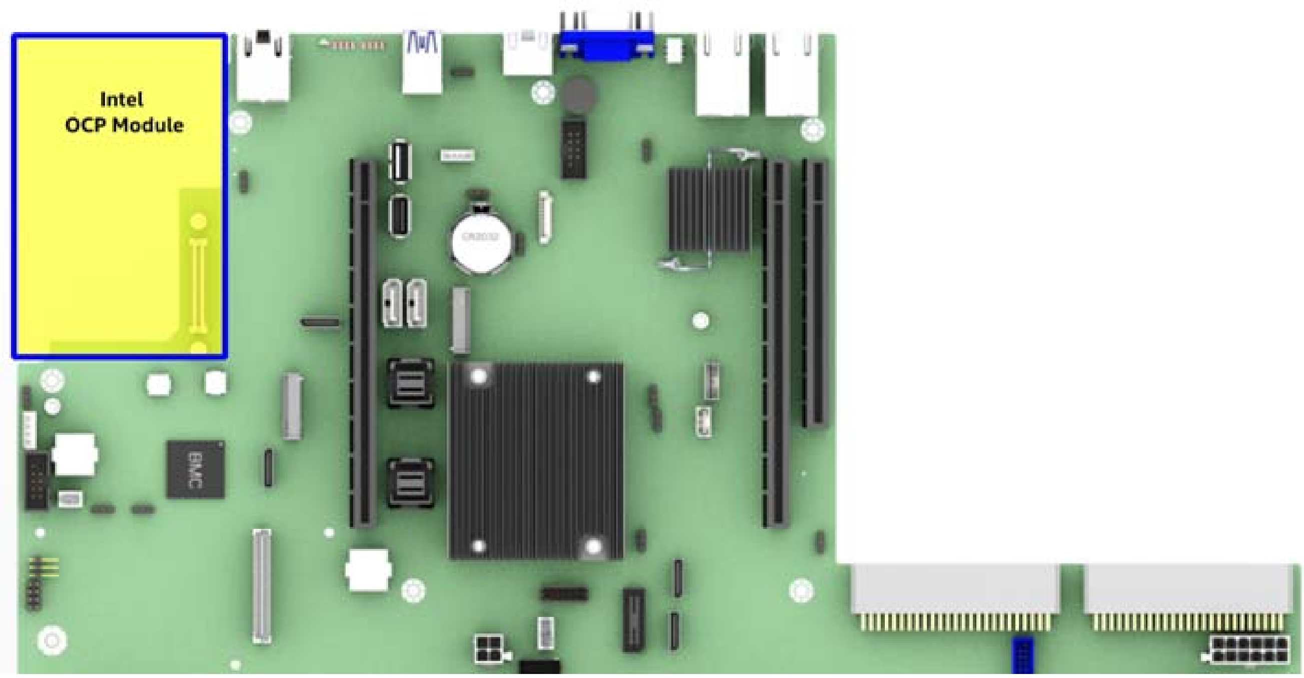

Figure 56. OCP Module Placement

Table 40. Supported Intel® OCP Modules

Description	Intel Product Code
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.

10. 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 41. Intel® Remote Management Module 4 (RMM4) Options

Intel Product Code	Description	Kit Contents	Benefits
AXXRMM4LITE2	Intel® Remote Management Module 4 Lite	RMM4 Lite 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 42. 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.

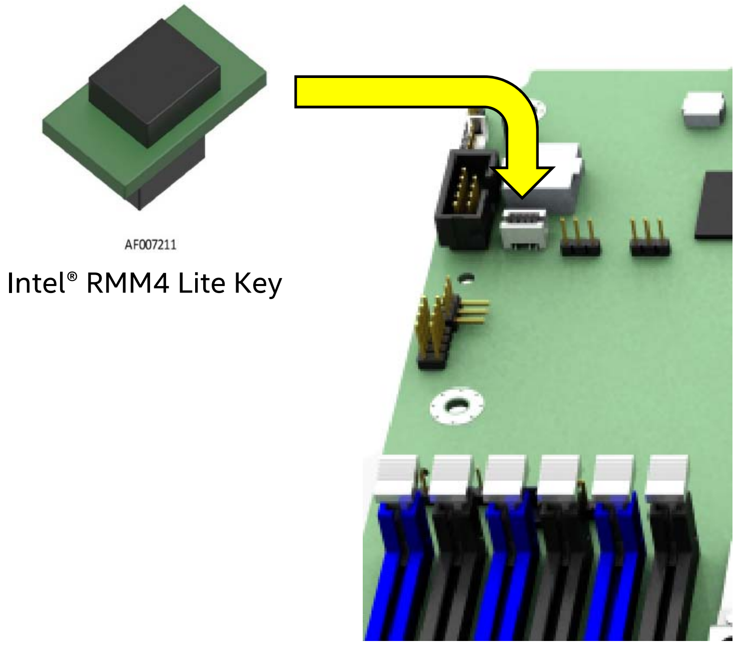

Figure 57. Intel® RMM4 Lite Activation Key Placement

10.1 Dedicated Management Port

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

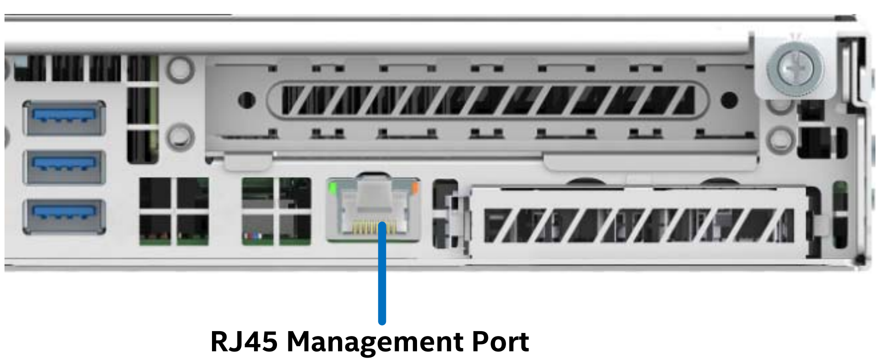

Figure 58. Dedicated Management Port

10.2 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 onboard 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*.

10.3 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.

10.3.1 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 60Hx, 72Hz, 75Hz, 85Hz
- 1280x960 at 60Hz
- 1280x1024 at 60Hz
- 1600x1200 at 60Hz
- 1920x1080 (1080p) at 60Hz
- 1920x1200 (WUXGA+) at 60Hz
- 1650x1080 (WSXGA+) at 60Hz

10.3.2 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.

10.3.3 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.

10.3.4 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.

10.3.5 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.

10.3.6 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.

10.3.7 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.

10.3.8 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.

10.3.8.1 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.

10.3.8.2 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 R1000WF Product Family and should be kept in mind while configuring your server system.

- Intel highly recommends that system integrators and system service personnel reference the ***Intel® Server System R1000WF Product Family System Integration and Service Guide*** for complete system assembly and component installation and removal instructions.
- The process to install and remove processors has changed from all previous generation Intel Server Systems. Before attempting to install or remove a processor for the first time, Intel highly recommends that system integrators and system service personnel reference the ***Intel® Server System R1000WF Product Family System Integration and Service Guide*** for complete processor installation and removal instructions
- 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 board supports the Intel® Xeon® processor Scalable family with a Thermal Design Power (TDP) of up to and including 165 Watts in the 1U Intel Server System product family. Previous generations of the Intel® Xeon® processors and their heat sinks, are not supported in this product family.
- Processors must be installed in order. CPU 1 must be populated for the server board to operate.
- To support proper air flow and system thermals when the system is operational, both processor heat sinks must be installed, regardless of whether the system is configured with two processors or not.
- To support proper air flow and system thermals when the system is operational, all 1U system configurations must have specific DIMM slots populated with a DIMM or supplied DIMM blank. See Section 4.1 for complete System Thermal Configuration requirements.
- The riser card slots on the server board 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 is a bank of eight diagnostic LEDs (4 Green and 4 Amber) that display POST Codes when the system is first powered-on or reset. If the server board hangs during POST, the LEDs display the last POST event run before the hang. See Appendix B to determine how to decode the Diagnostic LEDs when the system is locked during the POST process.
- The System Status LED will be set to a steady Amber color for all Fatal Errors that are detected during processor initialization. A steady Amber System Status LED indicates that an unrecoverable system failure condition has occurred
- 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 will 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, ME Firmware and FRU/SDR data. The latest system software can be downloaded from http://downloadcenter.intel.com.

Appendix B. POST Code Diagnostic LED Decoder

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

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.

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

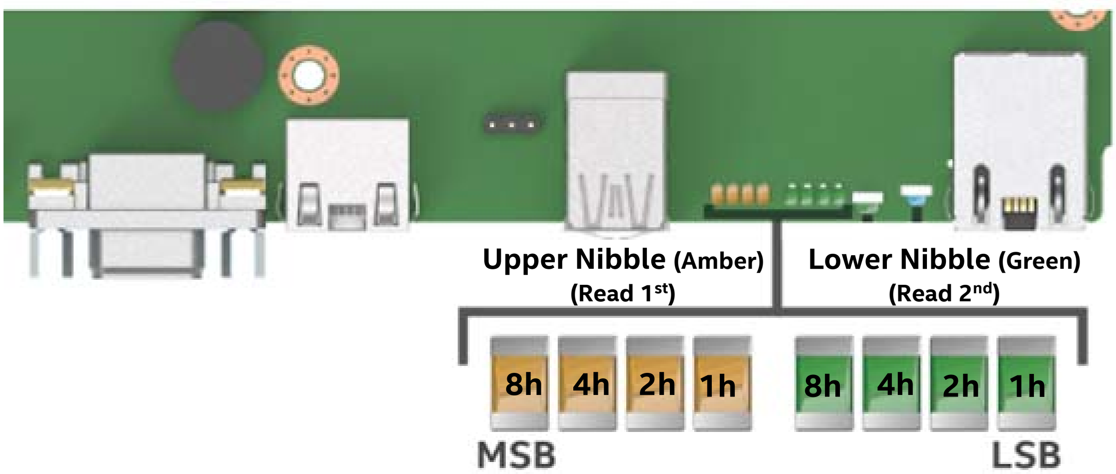

Figure 59. POST Diagnostic LED Location

In the following example, the BIOS sends a value of AC to the diagnostic LED decoder. The LEDs are decoded as shown in the table below, 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 43. 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 44. MRC Progress Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, **0** = LED Off								
	Upper Nibble (Amber)				**Lower Nibble (Green)**				
	MSB							LSB	
	8h	**4h**	**2h**	**1h**	**8h**	**4h**	**2h**	**1h**	
MRC Progress Codes									
B0h	1	0	1	1	0	0	0	0	Detect DIMM population
B1h	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2h	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3h	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4h	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5h	1	0	1	1	0	1	0	1	Program registers on the channel level
B6h	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7h	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8h	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9h	1	0	1	1	1	0	0	1	Hardware memory test and init
BAh	1	0	1	1	1	0	1	0	Execute software memory init
BBh	1	0	1	1	1	0	1	1	Program memory map and interleaving
BCh	1	0	1	1	1	1	0	0	Program RAS configuration
BFh	1	0	1	1	1	1	1	1	MRC is done

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 will 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 3 long beeps as identified in Table 48.

Table 45. MRC Fatal Error Codes

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, **0** = LED Off								
	Upper Nibble (Amber - Read 1st)				**Lower Nibble** (Green - Read 2nd)				
	MSB							LSB	
	8h	**4h**	**2h**	**1h**	**8h**	**4h**	**2h**	**1h**	
MRC Fatal Error Codes									
E8h	1	1	1	0	1	0	0	0	No usable memory error 01h = No memory was detected from SPD read, or invalid config that causes no operable memory. 02h = Memory DIMMs on all channels of all sockets are disabled due to hardware memtest error. 03h = No memory installed. All channels are disabled.
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel Trusted Execution Technology and is inaccessible
EAh	1	1	1	0	1	0	1	0	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
EBh	1	1	1	0	1	0	1	1	Memory test failure 01h = Software memtest failure. 02h = Hardware memtest failed.
EDh	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.
EFh	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 46. POST Progress Codes

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off								Description
	Upper Nibble (Amber - Read 1st)				Lower Nibble (Green - Read 2nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
SEC Phase									
01h	0	0	0	0	0	0	0	1	First POST code after CPU reset
02h	0	0	0	0	0	0	1	0	Microcode load begin
03h	0	0	0	0	0	0	1	1	CRAM initialization begin
04h	0	0	0	0	0	1	0	0	EI Cache When Disabled
05h	0	0	0	0	0	1	0	1	SEC Core at Power on Begin
06h	0	0	0	0	0	1	1	0	Early CPU initialization during Sec Phase.
UPI RC (Fully leverage without platform change)									
A1h	1	0	1	0	0	0	0	1	Collect info such as SBSP, Boot Mode, Reset type etc
A3h	1	0	1	0	0	0	1	1	Setup minimum path between SBSP & other sockets
A7h	1	0	1	0	0	1	1	1	Topology discovery and route calculation
A8h	1	0	1	0	1	0	0	0	Program final route
A9h	1	0	1	0	1	0	0	1	Program final IO SAD setting
AAh	1	0	1	0	1	0	1	0	Protocol layer and other uncore settings
ABh	1	0	1	0	1	0	1	1	Transition links to full speed operation
ACh	1	0	1	0	1	1	0	0	Phy layer setting
ADh	1	0	1	0	1	1	0	1	Link layer settings
AEh	1	0	1	0	1	1	1	0	Coherency settings
AFh	1	0	1	0	1	1	1	1	UPI initialization done
07h	0	0	0	0	0	1	1	1	Early SB initialization during Sec Phase.
08h	0	0	0	0	1	0	0	0	Early NB initialization during Sec Phase.
09h	0	0	0	0	1	0	0	1	End Of Sec Phase.
0Eh	0	0	0	0	1	1	1	0	Microcode Not Found.
0Fh	0	0	0	0	1	1	1	1	Microcode Not Loaded.
PEI Phase									
10h	0	0	0	1	0	0	0	0	PEI Core
11h	0	0	0	1	0	0	0	1	CPU PEIM
15h	0	0	0	1	0	1	0	1	NB PEIM
19h	0	0	0	1	1	0	0	1	SB PEIM
MRC Progress Codes									

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off Upper Nibble (Amber - Read 1st) MSB				Lower Nibble (Green - Read 2nd)			LSB	Description
	8h	4h	2h	1h	8h	4h	2h	1h	
31h	0	0	1	1	0	0	0	1	Memory Installed
32h	0	0	1	1	0	0	1	0	CPU PEIM (CPU Init)
33h	0	0	1	1	0	0	1	1	CPU PEIM (Cache Init)
4Fh	0	1	0	0	1	1	1	1	Dxe IPL started
DXE Phase									
60h	0	1	1	0	0	0	0	0	DXE Core started
61h	0	1	1	0	0	0	0	1	DXE NVRAM Init
62h	0	1	1	0	0	0	1	0	DXE Setup Init
63h	0	1	1	0	0	0	1	1	DXE CPU Init
65h	0	1	1	0	0	1	0	1	DXE CPU BSP Select
66h	0	1	1	0	0	1	1	0	DXE CPU AP Init
68h	0	1	1	0	1	0	0	0	DXE PCI Host Bridge Init
69h	0	1	1	0	1	0	0	1	DXE NB Init
6Ah	0	1	1	0	1	0	1	0	DXE NB SMM Init
70h	0	1	1	1	0	0	0	0	DXE SB Init
71h	0	1	1	1	0	0	0	1	DXE SB SMM Init
72h	0	1	1	1	0	0	1	0	DXE SB devices Init
78h	0	1	1	1	1	0	0	0	DXE ACPI Init
79h	0	1	1	1	1	0	0	1	DXE CSM Init
80h	1	0	0	0	0	0	0	0	DXE BDS Started
81h	1	0	0	0	0	0	0	1	DXE BDS connect drivers
82h	1	0	0	0	0	0	1	0	DXE PCI Bus begin
83h	1	0	0	0	0	0	1	1	DXE PCI Bus HPC Init
84h	1	0	0	0	0	1	0	0	DXE PCI Bus enumeration
85h	1	0	0	0	0	1	0	1	DXE PCI Bus resource requested
86h	1	0	0	0	0	1	1	0	DXE PCI Bus assign resource
87h	1	0	0	0	0	1	1	1	DXE CON_OUT connect
88h	1	0	0	0	1	0	0	0	DXE CON_IN connect
89h	1	0	0	0	1	0	0	1	DXE SIO Init
8Ah	1	0	0	0	1	0	1	0	DXE USB start
8Bh	1	0	0	0	1	0	1	1	DXE USB reset
8Ch	1	0	0	0	1	1	0	0	DXE USB detect
8Dh	1	0	0	0	1	1	0	1	DXE USB enable
91h	1	0	0	1	0	0	0	1	DXE IDE begin

Checkpoint	Diagnostic LED Decoder								Description
	1 = LED On, 0 = LED Off								
	Upper Nibble (Amber - Read 1st)				Lower Nibble (Green - Read 2nd)				
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
92h	1	0	0	1	0	0	1	0	DXE IDE reset
93h	1	0	0	1	0	0	1	1	DXE IDE detect
94h	1	0	0	1	0	1	0	0	DXE IDE enable
95h	1	0	0	1	0	1	0	1	DXE SCSI begin
96h	1	0	0	1	0	1	1	0	DXE SCSI reset
97h	1	0	0	1	0	1	1	1	DXE SCSI detect
98h	1	0	0	1	1	0	0	0	DXE SCSI enable
99h	1	0	0	1	1	0	0	1	DXE verifying SETUP password
9Bh	1	0	0	1	1	0	1	1	DXE SETUP start
9Ch	1	0	0	1	1	1	0	0	DXE SETUP input wait
9Dh	1	0	0	1	1	1	0	1	DXE Ready to Boot
9Eh	1	0	0	1	1	1	1	0	DXE Legacy Boot
9Fh	1	0	0	1	1	1	1	1	DXE Exit Boot Services
C0h	1	1	0	0	0	0	0	0	RT Set Virtual Address Map Begin
C2h	1	1	0	0	0	0	1	0	DXE Legacy Option ROM init
C3h	1	1	0	0	0	0	1	1	DXE Reset system
C4h	1	1	0	0	0	1	0	0	DXE USB Hot plug
C5h	1	1	0	0	0	1	0	1	DXE PCI BUS Hot plug
C6h	1	1	0	0	0	1	1	0	DXE NVRAM cleanup
C7h	1	1	0	0	0	1	1	1	DXE ACPI Enable
0h	0	0	0	0	0	0	0	0	Clear POST Code
S3 Resume									
40h	0	1	0	0	0	0	0	0	S3 Resume PEIM (S3 started)
41h	0	1	0	0	0	0	0	1	S3 Resume PEIM (S3 boot script)
42h	0	1	0	0	0	0	1	0	S3 Resume PEIM (S3 Video Repost)
43h	0	1	0	0	0	0	1	1	S3 Resume PEIM (S3 OS wake)
BIOS Recovery									
46h	0	1	0	0	0	1	1	0	PEIM which detected forced Recovery condition
47h	0	1	0	0	0	1	1	1	PEIM which detected User Recovery condition
48h	0	1	0	0	1	0	0	0	Recovery PEIM (Recovery started)
49h	0	1	0	0	1	0	0	1	Recovery PEIM (Capsule found)
4Ah	0	1	0	0	1	0	1	0	Recovery PEIM (Capsule loaded)
E8h	1	1	1	0	1	0	0	0	No Usable Memory Error:

Checkpoint	Diagnostic LED Decoder 1 = LED On, 0 = LED Off Upper Nibble (Amber - Read 1st)				Lower Nibble (Green - Read 2nd)				Description
	MSB							LSB	
	8h	4h	2h	1h	8h	4h	2h	1h	
E9h	1	1	1	0	1	0	0	1	Memory is locked by Intel® Trusted Execution Technology and is inaccessible.
EAh	1	1	1	0	1	0	1	0	DDR4 Channel Training Error:
EBh	1	1	1	0	1	0	1	1	Memory Test Failure
EDh	1	1	1	0	1	1	0	1	DIMM Configuration/Population Error
EFh	1	1	1	0	1	1	1	1	Indicates a CLTT table structure error
B0h	1	0	1	1	0	0	0	0	Detect DIMM population
B1h	1	0	1	1	0	0	0	1	Set DDR4 frequency
B2h	1	0	1	1	0	0	1	0	Gather remaining SPD data
B3h	1	0	1	1	0	0	1	1	Program registers on the memory controller level
B4h	1	0	1	1	0	1	0	0	Evaluate RAS modes and save rank information
B5h	1	0	1	1	0	1	0	1	Program registers on the channel level
B6h	1	0	1	1	0	1	1	0	Perform the JEDEC defined initialization sequence
B7h	1	0	1	1	0	1	1	1	Train DDR4 ranks
B8h	1	0	1	1	1	0	0	0	Initialize CLTT/OLTT
B9h	1	0	1	1	1	0	0	1	Hardware memory test and init
BAh	1	0	1	1	1	0	1	0	Execute software memory init
BBh	1	0	1	1	1	0	1	1	Program memory map and interleaving
BCh	1	0	1	1	1	1	0	0	Program RAS configuration
BFh	1	0	1	1	1	1	1	1	MRC is done

Appendix C. POST Code Errors

Most error conditions encountered during POST are reported using POST Error Codes. These codes represent specific failures, warnings, or are informational. POST Error Codes may be displayed in the Error Manager 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 are handed by a Diagnostic LED display with a system halt.

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

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

  Note that for 0048 “Password check failed”, the system halts, and then after the next reset/reboot will display the error code on the Error Manager screen.
- **Fatal:** The system halts during post at a blank screen with the text “Unrecoverable fatal error found. System will not boot until the error is resolved” and “Press <F2> to enter setup” The POST Error Pause option setting in the BIOS setup does not have any effect with this class of error.

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

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

Table 47. POST Error Messages and Handling

Error Code	Error Message	Action Message	Response
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_E3failed 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 users on error conditions. The beep code is followed by a user-visible code on the POST Progress LEDs.

Table 48. 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 & 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 49. 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.

Appendix D. System Configuration Table for Thermal Compatibility

This section provides system configuration compatibility data based on various supported system operating thermal limits. Two tables are provided. The first table identifies supported system configurations while the system is in "normal" operating mode; all systems fans are present, on-line, and operational. The second table identifies supported system configurations while the system is in a "fan fail" mode; one system fan or system fan rotor is no longer on-line or operational and fan redundancy is lost.

The following notes communicate support criteria associated with specific configurations identified in the following tables. Each relevant note to a configuration is identified by reference number in the table. Listed notes that are not specified in the table will reflect support criteria for a similar 2U based system within the Intel® Server Board S2600WF product family, details of which can be found in the Intel® Server System R2000WF Technical Product Specification.

Thermal configuration table notes:

1. The 27°C configuration alone is limited to elevations of 900m or less. Altitudes higher than 900m need to be de-rated to ASHRAE Class 2 levels.
2. To support system fan redundancy, the system must be configured with two power supplies to maintain sufficient cooling. Concurrent system and power supply fan failures is not supported.
3. Processor and memory throttling may occur which may impact system performance. CPU reliability is not impacted.
4. In fan fail mode, Intel® Ethernet Connection X557-T2 and Intel® Ethernet Connection X527-DA2, both dual-port OCP mezzanine cards are only supported in the specified base system model configured with 120W processors and DRx4 memory.
5. Use of the designated PCIe* 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 requirements.
6. For ASHRAE Class 3 and Class 4 support, the following power supply margining is required to meet thermal specifications:
   a) For dual power supply configurations, the power budget must fit within a single power supply rated load and be installed in a dual configuration, or
   b) For single power supply configurations, the power budget must be sized with 30% margin to single power supply rated load.
7. PCIe* SSD AIC FF devices can only be supported in the bottom add-in card slot on Riser Slot #1 and Riser Slot #2.
8. The Intel® RAID Maintenance Free Backup Unit (AXXRMFBUx) can support a case temperature of up to 45°C with the system operating in normal mode and up to 55°C with the system operating in a fan fail mode. The case temperature of Intel® Smart RAID Battery (AXXRSBBUx) can support up to 45°C in both normal and fan fail mode. Excursions over these specs may result in a reliability impact.
9. LBG cooling is not sufficient at higher ambient temperatures, performance could be limited to reduce preheating.
10. System cooling capability testing was carried out in environmental lab controlled conditions according to ASHRAE standard.
11. Performance depends on data center environmental temperature and relative humidity levels controls provided by end user.
12. It is the responsibility of the system integrator to both consider the thermal configuration matrix and power budget tool documents in order to arrange end use configuration.
13. M.2 drives may see performance impact under heavy work load.
14. Light workload is assuming 70% write, 30% read, 100% Random, 100% access, 8kb transfer rate, and IO "delay" of 8.

15. See TPS for thermal limitations.
16. AEP electrically/system not supported on the program.
17. PSU Unit 2 supported up to max CPU TDP of 150 W.
18. PSU Unit 2 supported up to max CPU TDP of 140 W.
19. PSU Unit 2 supported up to max CPU TDP of 125 W.
20. PSU Unit 2 supported up to max CPU TDP of 105 W.
21. PSU Unit 2 supported up to max CPU TDP of 85 W.
22. Intel NVME drives were used for thermal testing.
23. During fan fail mode CPU throttling might occur due power limiting feature for PSU 2.
24. PSU 2 supported up to a max CPU TDP of 135 W (Applies for Fabric sku only).
25. PSU 2 supported up to max CPU TDP of 113 W (Applies for Fabric sku only).

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

			R1304 HDD WF				**R1208 HDD WF**				**R1208 NVME WF**			
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 Note 6)		1100W AC	●	●	●	●	●	●	●	●	●	●	●	●
		750W DC	17, 24	19, 24	20, 25	21	17, 24	19, 24	20	21				
CPU TDP	**173 W**	Intel® Xeon® Platinum 8176F _28C	●				●				3			
	165 W	Intel® Xeon® Platinum 8176 _28C	●	3			3	3			3			
		Intel® Xeon® Platinum 8170_26C	●	3			3	3			3			
		Intel® Xeon® Gold 6150_18C	●	●			●	3			3			
		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® Xeon® Gold 6136_12C	●	●	●		●	●			●			
		Intel® Xeon® Platinum 8158_12C	●	●	●		●	●			●			
		Intel® Xeon® Gold 6142_16C	●	3	3		●	3			●			
	140 W	Intel® Xeon® Gold 6132_14C	●	●	●		●	●	●		●	●		

		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	•	•	•	•	•	•	3	
		Intel® Xeon® Gold 6130_16C	•	•	•	3	•	•	•	•	•	•	3	
		Intel® Xeon® Platinum 8153_16C	•	•	•	3	•	•	•	•	•	•	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	•	•	•	•	•	•	•	•	•	•	•	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 16)		RDIMM-2Rx8,1Rx4, 1Rx8	•	•	•	•	•	•	•	•	•	•	•	•
		RDIMM-DRx4	•	•	•	•	•	•	•	•	•	•	•	•
		LRDIMM-QRx4 DDP	•	•	3	3	•	•	3	3	•	•	3	3

Add-in Cards (See Note 5)		Riser #1 - Bottom Slot (1U riser and 2U riser)	●	●	●	●	●	●	●	●	●	●	●	●
		Riser #2 - Bottom Slot (1U riser and 2U riser)	●	●	●	●	●	●	●	●	●	●	●	●
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	●	●	●	●	●	●	●	●	●	●	●	●
		16 port tri mode mezzanine	●	●	12	12	●	●	12	12	●	●	12	12
		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)	●	●	●	●	●	●	●	●	●	●	●	●
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 22)		1600GB/2TB									●	●	●	●
		800GB									●	●	●	●
		600GB									●	●	●	●
		400GB									●	●	●	●
		200GB									●	●	●	●
SATA HDD		Up to 2 TB	●	●	●	●	●	●	●	●	●	●	●	●
PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		1600GB/2TB	●	●	●	●	●	●	●	●	●	●	●	●
		800GB	●	●	●	●	●	●	●	●	●	●	●	●
		600GB	●	●	●	●	●	●	●	●	●	●	●	●
		400GB	●	●	●	●	●	●	●	●	●	●	●	●
		200GB	●	●	●	●	●	●	●	●	●	●	●	●
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	●	●	13	13	●	●	13	13	●	●	13	13
		Non-Intel M.2	●	●	15	15	●	●	15	15	●	●	15	15
M.2 Light Workload (See Note 14)		Intel 120G and 80G SATA	●	●	●	13	●	●	●	13	●	●	●	13
		Non-Intel M.2	●	●	●	15	●	●	●	15	●	●	●	15

Thermal Configuration Matrix
Fan Fail Mode

"●" = Full Support without limitation
"4,5" (Cell with number) = Conditional support with limitation (See notes above)
" " (Blank) = No support

			R1304 HDD WF				R1208 HDD WF				R1208 NVME WF			
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 Note 6)		1100W AC	2	2	2	2	2	2	2	2	2	2	2	2
		750W DC	2, 20, 24	2, 21, 24	2	2	2, 20, 24	2, 21	2	2				
CPU TDP	**173 W**	Intel® Xeon® Platinum 8176F _28C	23	23			23				23			
	165 W	Intel® Xeon® Platinum 8176 _28C	23	23			23				23			
		Intel® Xeon® Platinum 8170_26C	23	23			23				23			
		Intel® Xeon® Gold 6150_18C	23	23			23				23			
		Intel® Xeon® Platinum 8170F_26C	23	23			23	23			23			
	160 W	Intel® Xeon® Platinum 8160F_24C	23	23			23	23			23			
		Intel® Xeon® Gold 6148F_20C	23	23			23	23			23			
		Intel® Xeon® Gold 6142F_16C	23	23			23	23			23			
	150 W	Intel® Xeon® Platinum 8164_26C	3	23	23		23	23			23			
		Intel® Xeon® Platinum 8160_24C	3	23	23		23	23			23			
		Intel® Xeon® Gold 6148_20C	3	23	23		23	23			23			
		Intel® Xeon® Gold 6136_12C	3	23	23		23	23			23			
		Intel® Xeon® Platinum 8158_12C	3	23	23		23	23			23			

		Intel® Xeon® Gold 6142_16C	3	23	23		23	23			23			
	140 W	Intel® Xeon® Gold 6132_14C	3	3	23		3	23	23		3	23		
		Intel® Xeon® Gold 6152_22C	3	3	23		3	23	23		3	23		
		Intel® Xeon® Gold 6140_18C	3	3	23		3	23	23		3	23		
	135 W	Intel® Xeon® Gold 6138F_20C	3	3	23		3	23	23		•	23		
		Intel® Xeon® Gold 6130F_16C	3	3	23		3	23	23		•	23		
		Intel® Xeon® Gold 6126F_12C	3	3	23		3	23	23		•	23		
	130 W	Intel® Xeon® Gold 6134_8C	•	3	23		•	23	23		•	23		
	125 W	Intel® Xeon® Gold 6138_20C	•	3	23	23	•	3	23	23	•	3	23	
		Intel® Xeon® Gold 6130_16C	•	3	23	23	•	3	23	23	•	3	23	
		Intel® Xeon® Platinum 8153_16C	•	3	23	23	•	3	23	23	•	3	23	
		Intel® Xeon® Gold 6126_12C	•	3	23	23	•	3	23	23	•	3	23	
	115 W	Intel® Xeon® Gold 6128_6C	•	•	•	23	•	•	23	23	•	•	23	
	113 W	Intel® Xeon® Gold 5117F_14C	•	•	23	23	•	•	23	23	•	•	23	
	105 W	Intel® Xeon® Gold 5122_4C	•	•	3	23	•	•	3	23	•	•	3	23
		Intel® Xeon® Platinum 8156_4C	•	•	3	23	•	•	3	23	•	•	3	23
		Intel® Xeon® Gold 5120_14C	•	•	3	23	•	•	3	23	•	•	3	23
		Intel® Xeon® Gold 5118_12C	•	•	3	23	•	•	3	23	•	•	3	23
	85 W	Intel® Xeon® Gold 5115_10C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Silver 4116_12C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Silver 4114_10C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Silver 4110_8C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Silver 4108 _8C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Bronze 3106 _8C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Bronze 3104_6C	•	•	•	23	•	•	•	23	•	•	•	23
		Intel® Xeon® Silver 4112_4C	•	•	•	23	•	•	•	23	•	•	•	23
Memory Type		RDIMM-2Rx8,1Rx4, 1Rx8	•	•	•	•	•	•	•	•	•	•	•	•

(See Note 16)		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 #2 - Bottom Slot (1U riser and 2U riser)	●	●	●	●	●	●	●	●	●	●	●	●
SAS and OCP Modules (See Note 4)		Intel® Integrated RAID Modules (Mezzanine cards)	●	●	●	●	●	●	●	●	●	●	●	●
		16 port tri mode mezzanine	●	12	12	12	●	12	12	12	●	12	12	12
		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)	●	●	●	●	●	●	●	●	●	●	●	●
2.5" SFF NVMe SSD (DC P3700/P3500) (See Note 22)		1600GB/2TB									●	●	●	●
		800GB									●	●	●	●
		600GB									●	●	●	●
		400GB									●	●	●	●
		200GB									●	●	●	●
SATA HDD		Up to 2 TB	●	●	●	●	●	●	●	●	●	●	●	●
PCIe SSD AIC FF (DC P3700/P3500) (See Note 8)		1600GB/2TB	●	●	●	●	●	●	●	●	●	●	●	●
		800GB	●	●	●	●	●	●	●	●	●	●	●	●
		600GB	●	●	●	●	●	●	●	●	●	●	●	●
		400GB	●	●	●	●	●	●	●	●	●	●	●	●
		200GB	●	●	●	●	●	●	●	●	●	●	●	●
M.2 Sustained Max Power (TDP)		Intel 120G and 80G SATA	●	13	13	13	●	13	13	13	●	13	13	13
		Non-Intel M.2	●	15	15	15	●	15	15	15	●	15	15	15
M.2 Light Workload (See Note 14)		Intel 120G and 80G SATA	●	13	13	13	●	13	13	13	●	13	13	13
		Non-Intel M.2	●	15	15	15	●	15	15	15	●	15	15	15

Appendix E. System Cable Routing Diagrams

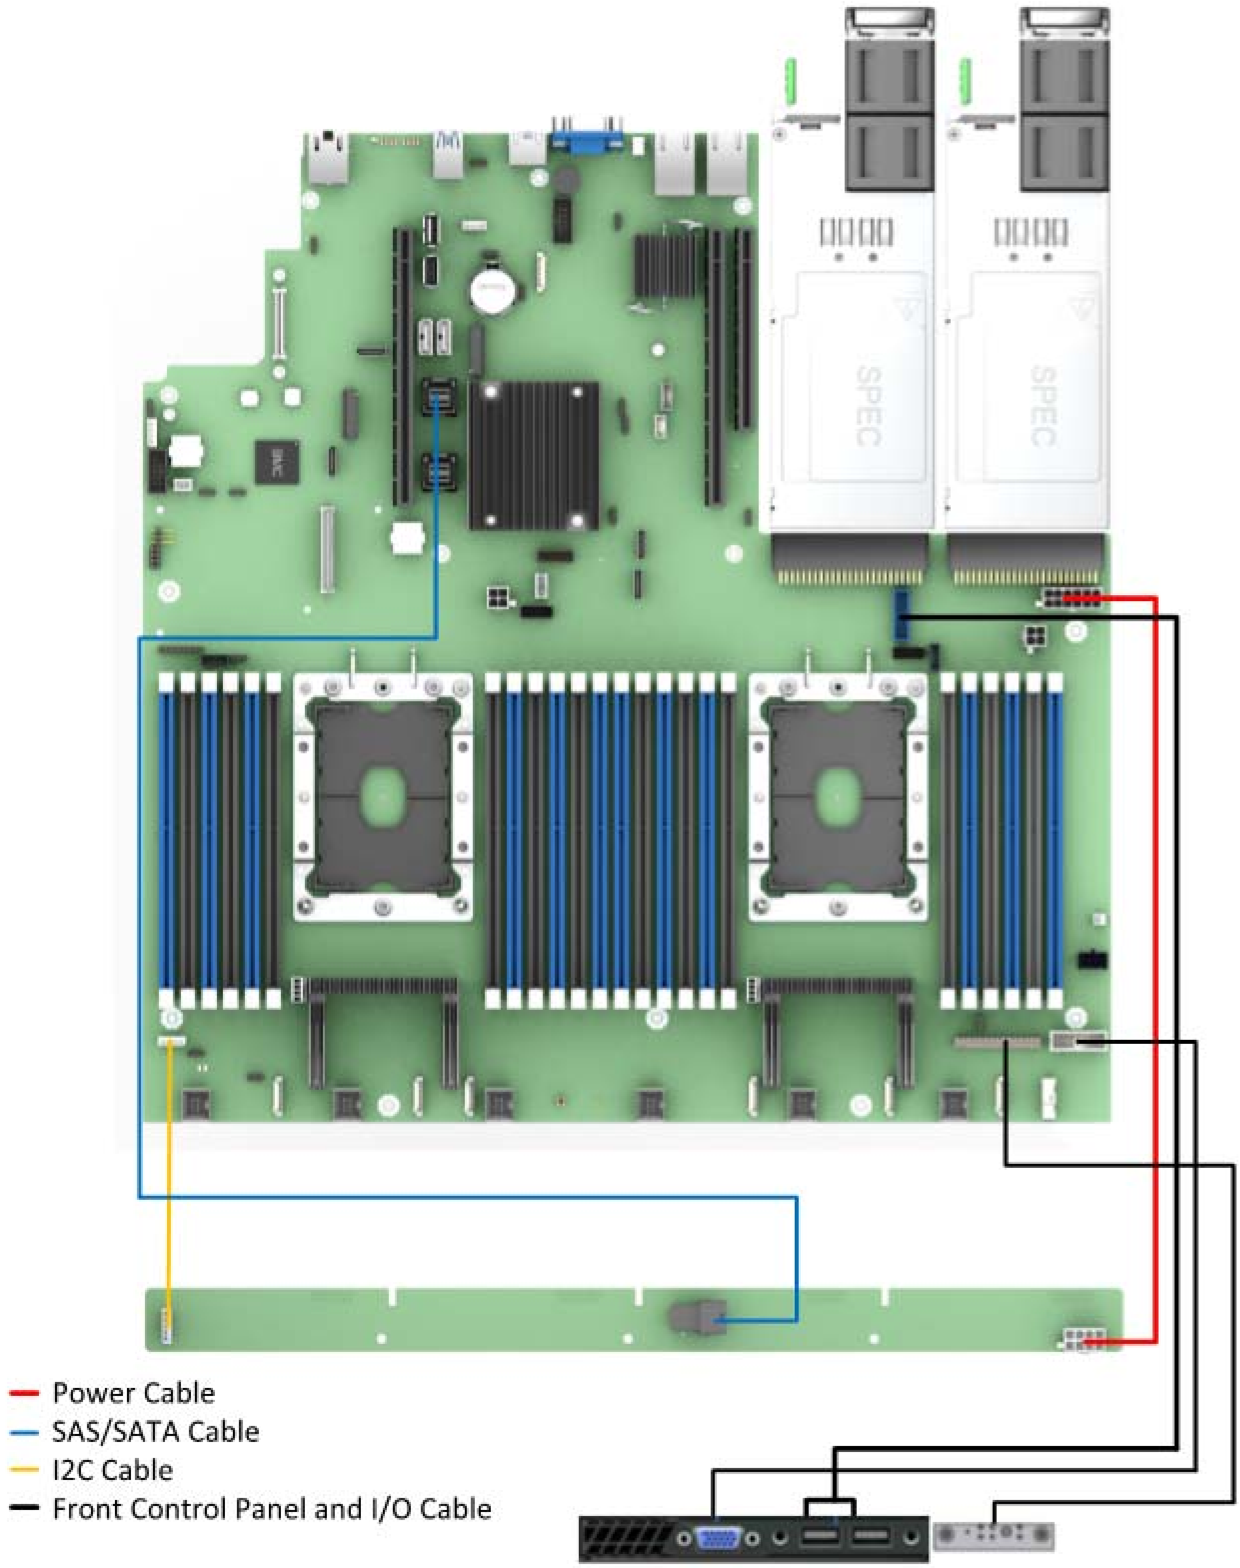

Figure 60. R1304WFxxx

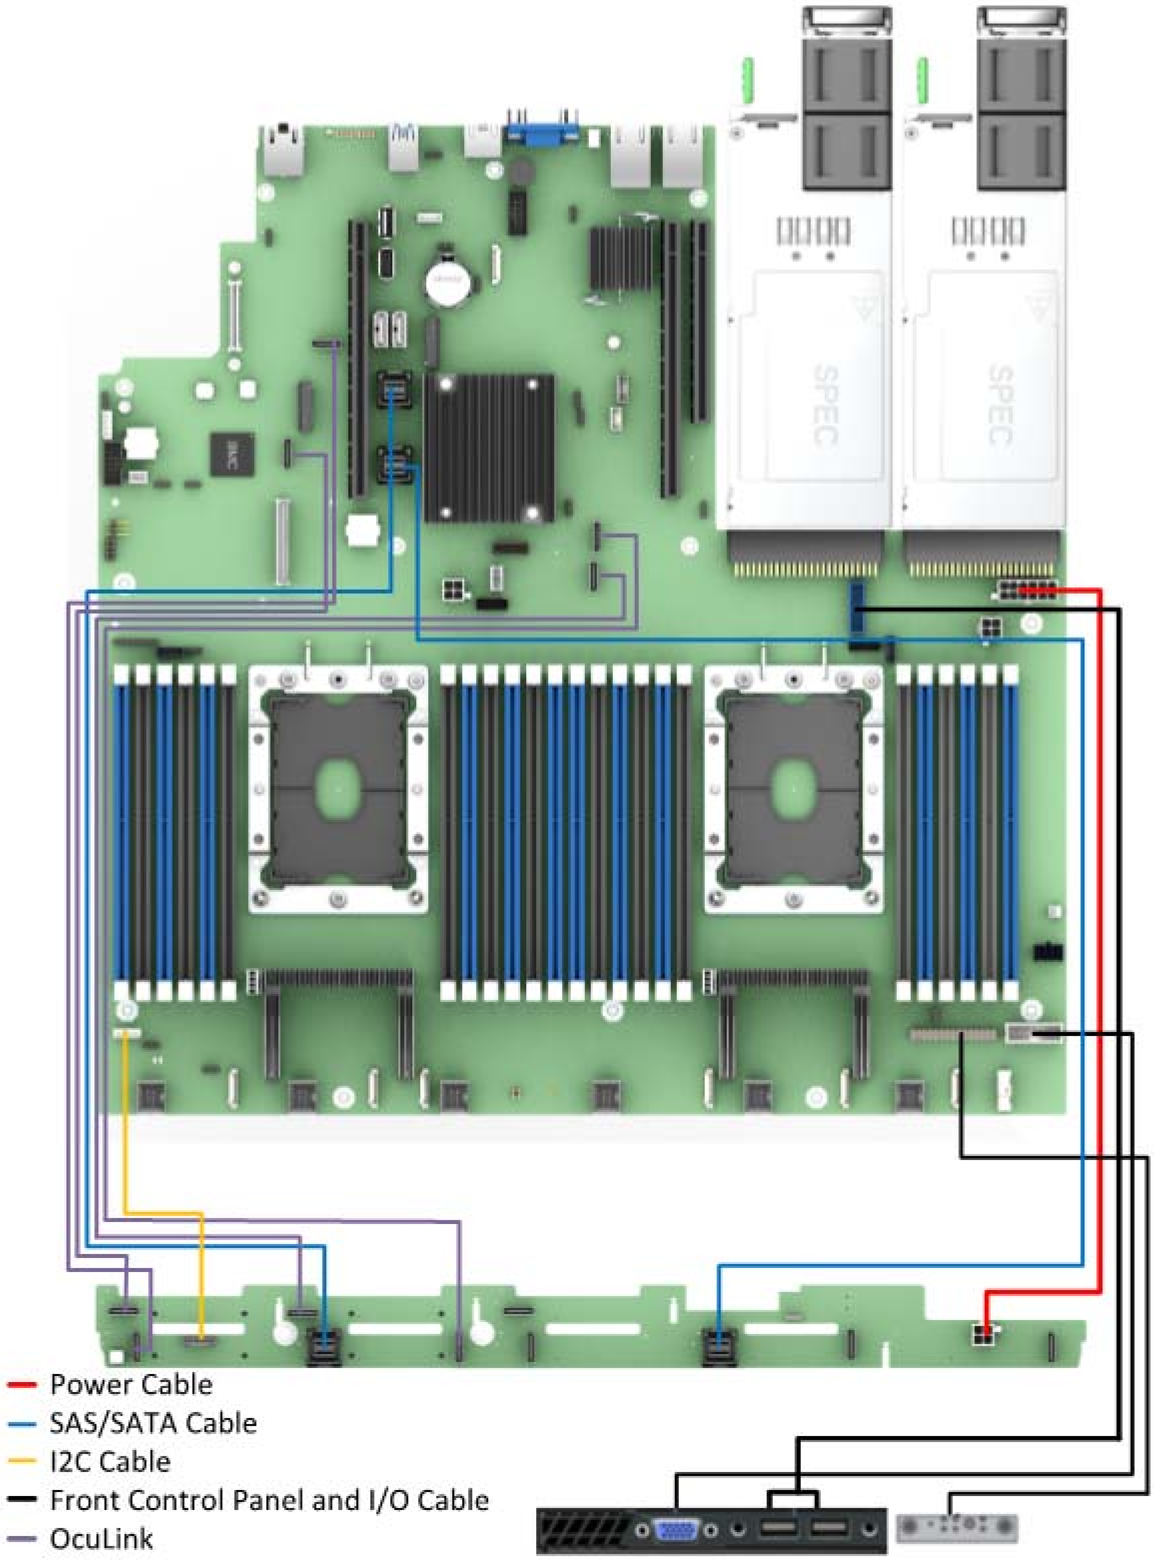

Figure 61. R1208WFxxx

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 R1000WF product family.

The tables provide the following data for each identified component.

Component Type

Three types of memory components are used on the server board assembly. These include:

- 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.
- Volatile: Volatile memory is cleared automatically when power is removed from the system.
- 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

The size of each component includes sizes in bits, Kbits, bytes, kilobytes (KB) or megabytes (MB).

Board Location

The physical location of each component is specified in the Board Location column. The board location information corresponds 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:

- 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.

- 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 will maintain user passwords to control this access. These passwords are stored in the BMC flash.

Table 50. Intel® Server Board S2600WFT (iPN - H48104-XXX)
Intel® Server Board S2600WFO (iPN - H87179-XXX)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	512MB	U3D1	No(BIOS)	BIOS Flash
Non-Volatile	512MB	U1D1	No(FW)	BMC Flash
Non-Volatile	4MB	U5M1	No	10 GB NIC EEPROM
Non-Volatile	N/A	U1E3	No	CPLD
Volatile	4GB	U1D2	No	BMC SDRAM
Non-Volatile	8GB	U8N1	No	BMC eMMC

Table 51. 1U 1 Slot PCIe* Riser Card (iPN – H39531-xxx)

Component Type	Size	Board Location	User Data	Name
N/A	N/A	None	No	N/A

Table 52. Front Panel Board (iPN – H29366-xxx)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256x8	None	Yes	PSOC / Microcontroller

Table 53. 1U 4 x 3.5" Hot Swap Back Plane option (iPN – G97162-XXX)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	None	Yes	PSOC / Microcontroller / FRU
Non-Volatile	1024x8	None	No	SAS Re-Driver Settings

Table 54. 1U 8 x 2.5" SAS Hot Swap Back Plane option (iPN – H88382-XXX)

Component Type	Size	Board Location	User Data	Name
Non-Volatile	16384x8	None	Yes	PSOC / Microcontroller / FRU
Non-Volatile	1024x8	None	No	SAS Re-Driver Settings

Table 55. Intel® Remote Management Module Lite Accessory Option (iPC – AXXRMM4LITE)

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

Glossary

Word/Acronym	Definition
BMC	Baseboard Management Controller
BIOS	Basic Input/Output System
CLST	Closed Loop System Throttling
CMOS	Complementary Metal-oxide-semiconductor
CPU	Central Processing Unit
DDR4	Double Data Rate 4th edition
DIMM	Dual In-line Memory Module
DOM	Disk-on-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
HDD	Hard Disk Drive
I2C	Inter-integrated Circuit bus
LCD	Liquid Crystal Display
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
OCP	Over-current Protection
OTP	Over-temperature Protection
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

Word/Acronym	Definition
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
SFF	Small Form Factor
SSD	Solid State Drive
TDP	Thermal Design Power
TPM	Trusted Platform Module
TPS	Technical Product Specification
USB	Universal Serial Bus
VLSI	Very Large Scale Integration
VSB	Voltage Standby

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.