Intel® Server System S9200WK Product Family

S9200WK

Technical Product Specification

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

Rev 1.8

February 2022

Intel® Server Products and Solutions

<Blank page>

Document Revision History

Date	Revision	Changes
July 2019	1.0	Initial release.
November 2019	1.1	• Updated Table 3 • Added product weight information to section 2.6 • Updated Section 5.1 • Updated illustrations in Chapter 11
December 2019	1.2	• Add information on Intel® Compute Module S9232WK1HAC and mixed compute module configurations • Updated Appendix D. System Configuration Table for Thermal Compatibility table update • Updated Table 4. Product weight information
January 2020	1.3	• Updated Appedix E – Product Regulatory Information, included content to support EU Lot 9 product collaterals support links.
March 2020	1.4	• Added EU Lot 9 Support Summary content to Appendix E • Updated RAID support for Systems Features Table • Updated RAID support for U.2 and M.2 storage sections • Updated Section 11.1 Thermal Operation and Configuration Requirements • Updated Appendix D.System Configuration Table for Thermal Compatibility • Added section 10.3.5 Cooling type monitoring
May 2020	1.5	• Updated “Product weight Information” table • Updated chapter 14.3 “Power Supply Module options”
December 2020	1.6	• Added DDR4 Support Disclaimer
December 2021	1.7	• Updated “Compute module status LED state definitions” table
February 2022	1.8	• Removed the “Note” For additional information on Intel® Intelligent Power Node Manager usage and support.

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

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

Table of Contents

List of Figures

List of Tables

<Blank page>

1. Introduction

The Intel® Server System S9200WK is a density-optimized, rack-mount, 2U, multi node system product family designed to support high-density high-performance computing environments. Each system within the Intel® Server System S9200WK product family includes compute modules that are preconfigured and independent, allowing for a power-on ready installation. The Intel® Server System 9200WK product family offers options to support liquid cooled and air cooled configurations.

This Technical Product Specification (TPS) provides a high-level overview of the features, functions, and support specifications of the Intel® Server System S9200WK product family. For additional information, refer to the documents specified in the following table:

Table 1. Intel® Server System S9200WK reference documents

Document Title	Document Classification
Intel® Server System S9200WK Product Family Setup and Service Guide	Public
Intel® Server System S9200WK Product Family Configuration Guide	Public
Intel® Server System S9200WK Product Family Technical Product Specification	Public

2. System Overview

Each server system within the Intel® Server System S9200WK product family includes compute modules, power supplies, cooling components, and rails for rack or cabinet mounting. Configurable options for the compute modules include memory, storage, and network components. Refer to the Intel® Server System S9200WK product family Configuration Guide for a complete list of available options.

Figure 1. Intel® Server System S9200WK

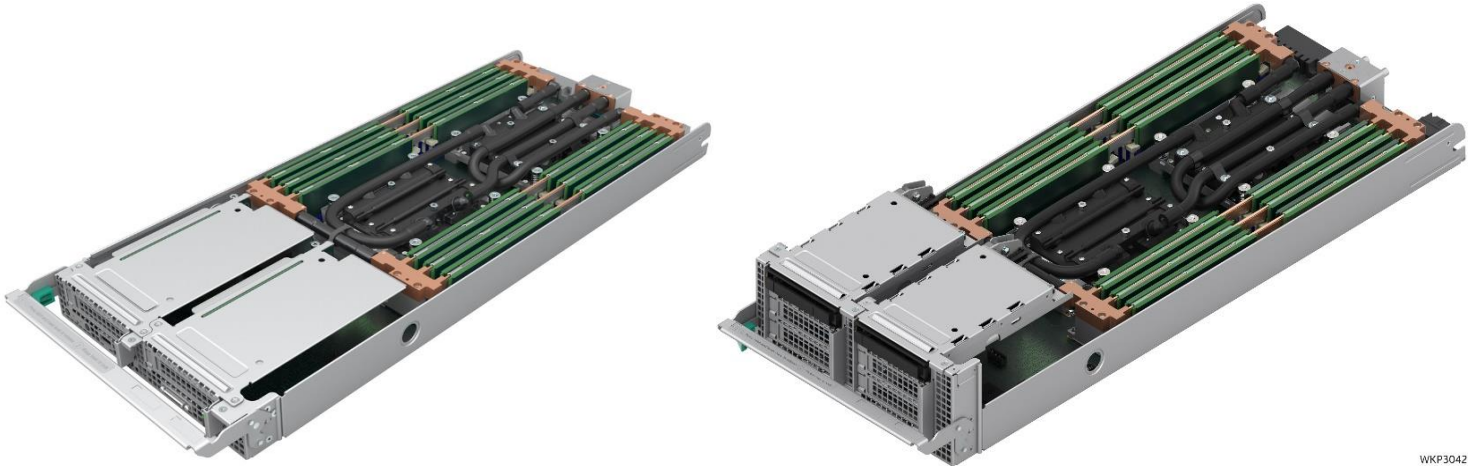

1U Compute module

2U Compute module

Figure 2. Compute module options supported with the Intel® Server System S9200WK product family

2.1 System Features Overview

The following table provides a high level overview of the feature set and system specifications that the Intel® Server System S9200WK product family systems are designed to support:

Table 2. System features/specifications table

Feature/Specification	Description
Chassis Type	2U, rack-mount, multi-node
Number of Compute Modules	Up to four 1U modules or up to two 2U modules
Chassis Dimensions	865 mm x 441.8 mm x 86.8 mm
Packaging Dimensions	1192 mm x 758 mm x 317 mm
Processor Support	Two Intel® Xeon® Platinum 9200 processors (per compute module)
Chipset	Intel® C621 Chipset
Memory Support (per compute module)	• Up to 24 DDR4 DIMMs – 12 DIMMs per processor • 2933 MT/s RDIMMs and LRDIMMs • 1.2V DIMMs • 3 TB Max memory supported • DIMM capacities of 8GB, 16GB, 32GB, or 64GB
Storage (per compute module)	**1U Compute Module** • Up to two SATA/PCIe* NVMe* M.2 Drives • VROC RAID unsupported for M.2 **2U Compute Module** • Up to two U.2 PCIe* NVMe* Drives • U.2 VROC RAID support; RAID boot on volume unsupported • Up to two SATA/PCIe* NVMe* M.2 Drives • VROC RAID unsupported for M.2
PCI Expansion (per compute module)	1U Compute Module • Two – PCIe* 3.0 x16 mechanical and electrical Add-in Card Slot 2U Compute Module • Four – PCIe* 3.0 x16 mechanical and electrical Add-in Card Slot
Network Ports (per compute module)	One 1GbE Base-T, RJ45 One dedicated management port, RJ45

Cooling	**Air-cooled configurations:**
	• Five Dual-rotor Hot-Swap System Fans with support for fan redundancy
	  • Three 60x60x56mm fans
	  • Two 80x80x80mm fans
	• One Fan per installed Power Supply Unit (PSU)
	
	Liquid-cooled configurations:
	• Three 60x60x56mm Dual-rotor Hot-Swap System Fans with support for fan redundancy
	• Liquid cooling loop (per compute module)
	• Liquid cooling plumbing connections on the back of the chassis
	• 1 Fan per installed Power Supply Unit (PSU)

Front panel support (per compute module)	One USB 3.0 Port One I/O breakout cable connector (1 cable per system) supporting the following: • Serial Port • Video Port • Two USB 2.0 Ports
Power	Three 1600 or 2100-watt AC power supplies (dependent on system configuration) • Power redundancy support (dependent on system configuration)
Rack Mount Kit	Tool-less installation

Feature/Specification	Description
(Included)	Fixed position
Serviceability	Modular chassis features for simplified serviceability • Fully independent compute modules • Hot-Swap Power Supplies • Hot-Swap System Fans • Hot-Swap U.2 Solid State Disk (SSD) Storage (2U compute modules only)
Operating Ambient Temperature Support	10°C – 35° C ambient temperature
Security (per compute module)	Onboard TPM 2.0 included (Rest of the world) **Note**: China only TPM not supported

2.2 System Feature Identification

All systems within the Intel® Server System S9200WK product family feature front loading compute modules. The following system configuration options are supported:

- Four 1U compute modules
- Two 2U compute modules
- One 2U compute module and two 1U compute modules

Note: Mixing compute modules is supported only for modules with the same type of cooling.

The following sections provide system views, identify key system features, and provide chassis dimensional data for all supported system configurations.

Figure 3. Compute module identification – four-node systems using 1U compute modules

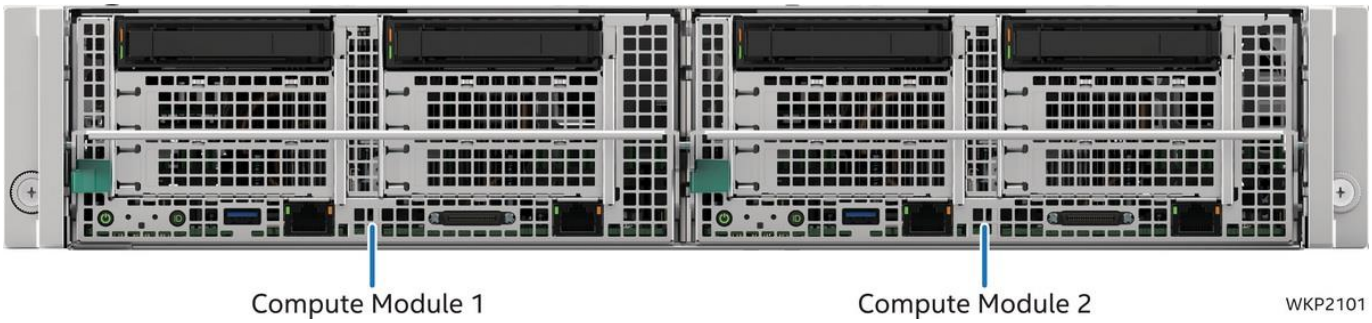

Figure 4. Compute module identification – two-node systems using 2U compute modules

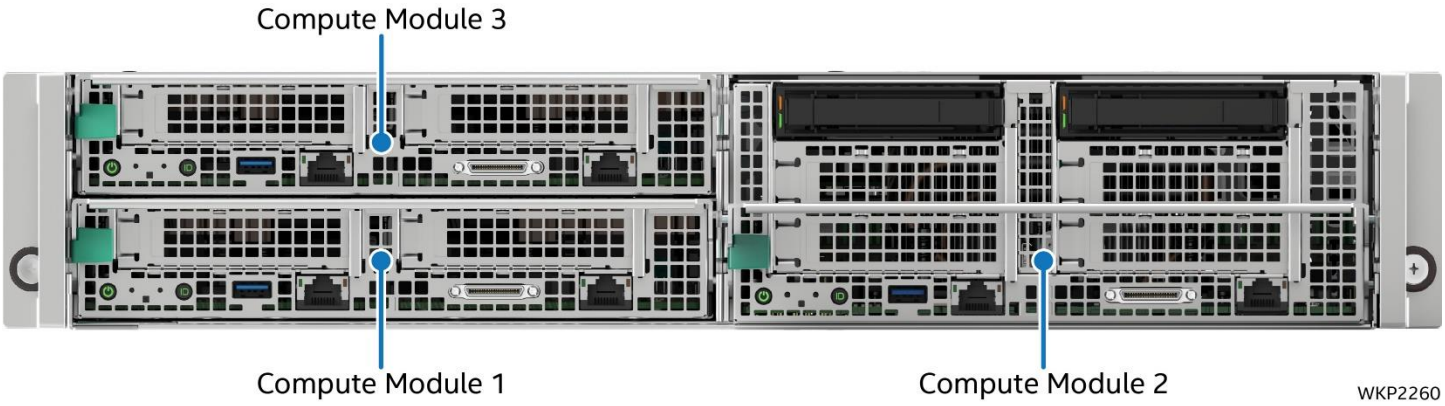

Figure 5. Compute module identification – mixed-node systems using two 1U and one 2U compute modules

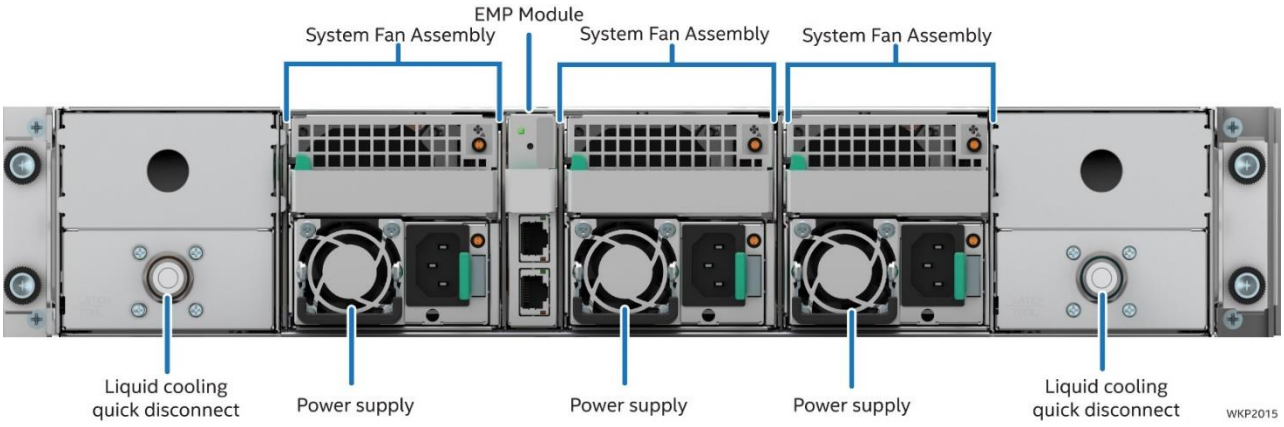

Figure 6. Liquid-cooled system back view

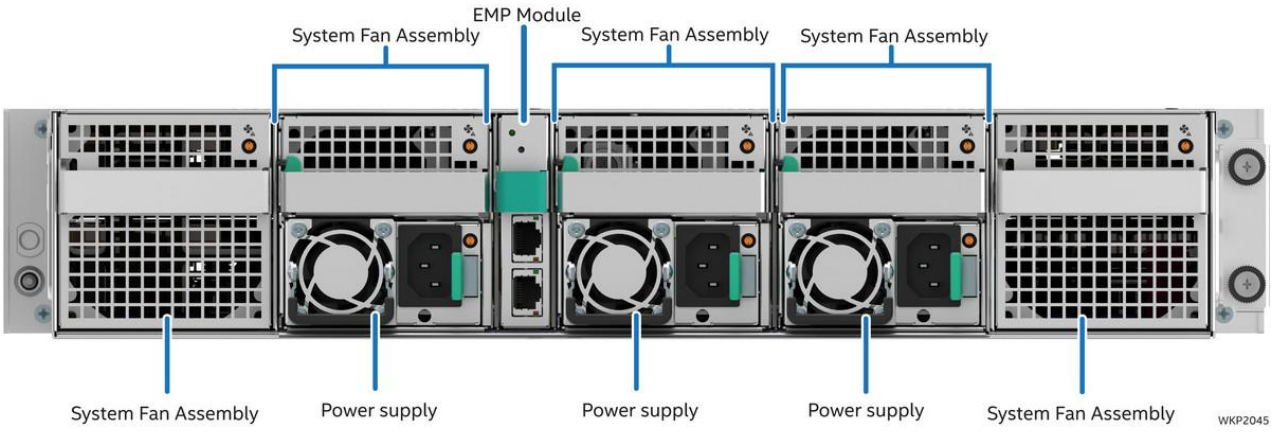

Figure 7. Air-cooled system back view

2.3 System dimensional data

2.3.1 Chassis Dimensions

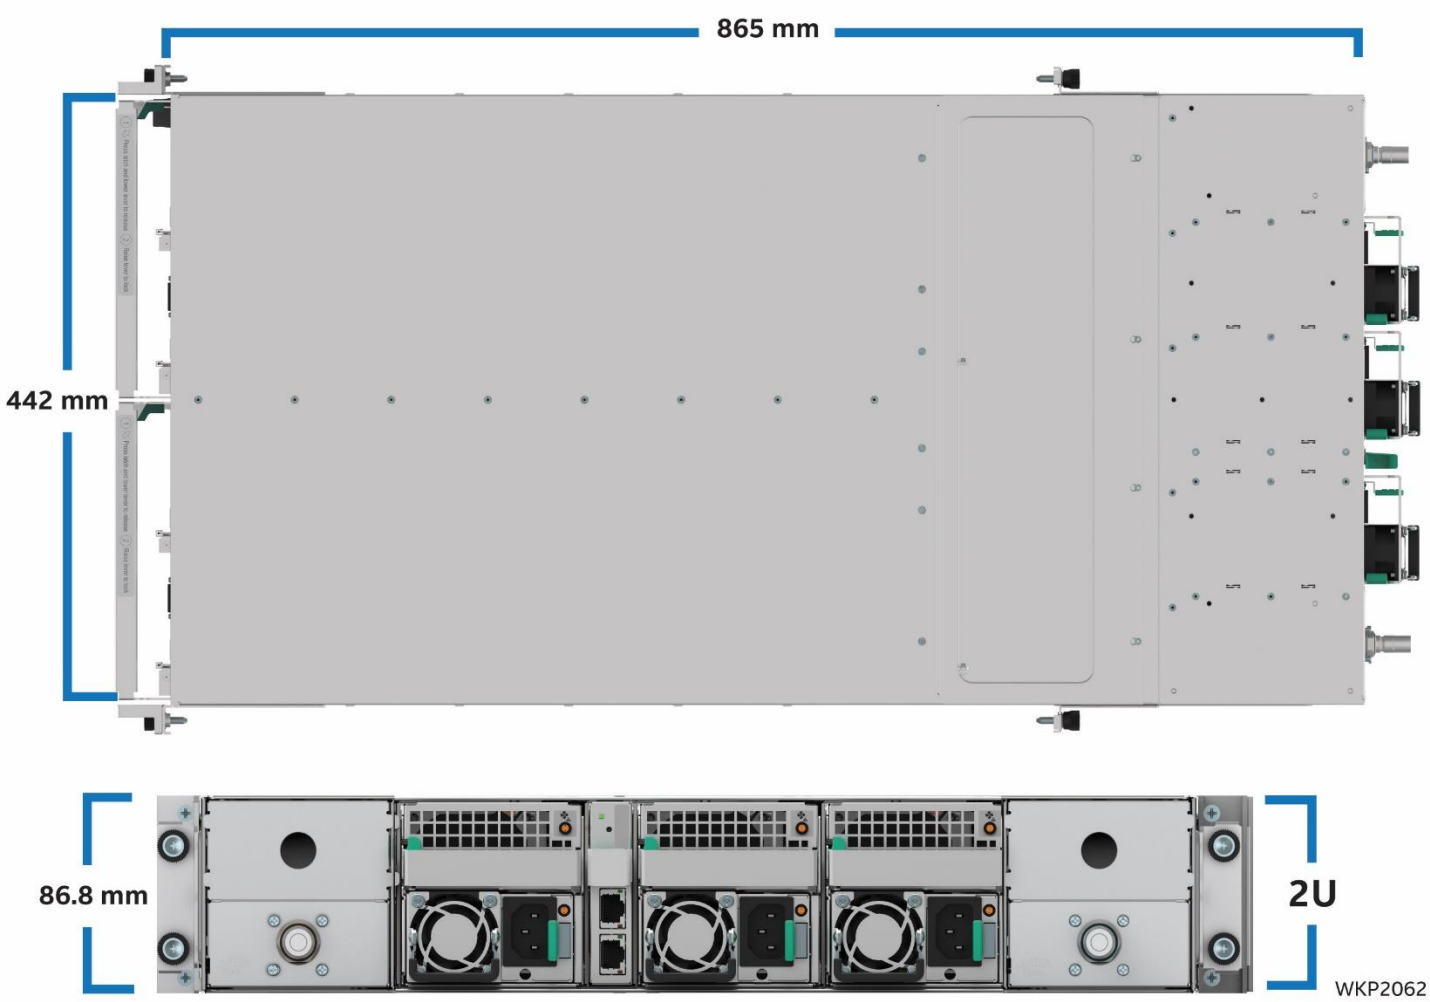

Figure 8. System chassis dimensions

2.3.2 Pull-Out Tab

The chassis includes a pull-out tab that can be used for asset tracking and identification. Figure 9 identifies the location of the pull-out tab, and Figure 10 identifies its features.

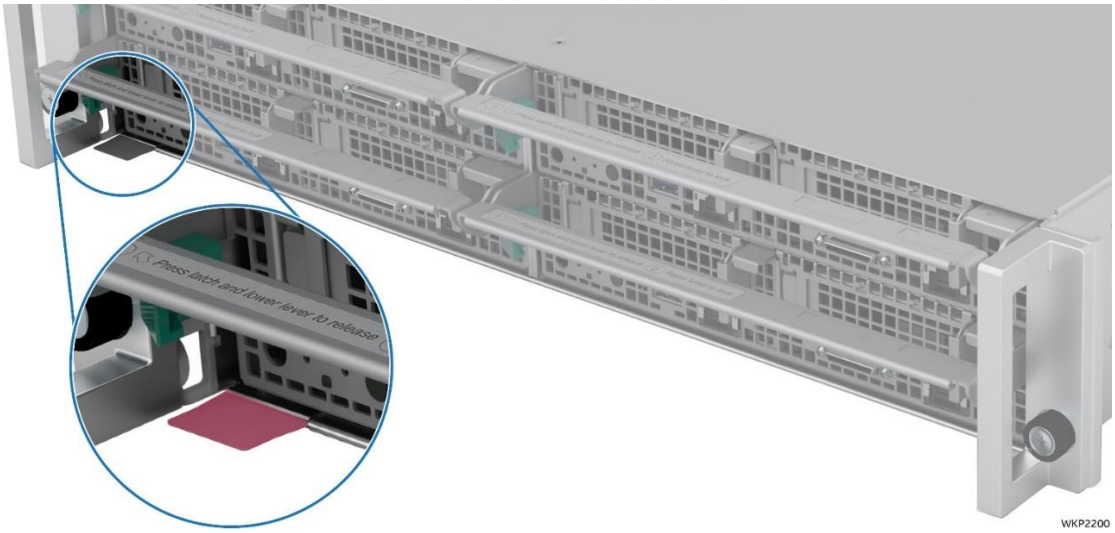

Figure 9. Pull-out tab location

Figure 10. Pull out tab dimensions

2.4 Rack and Cabinet Mounting Kit

The Intel® Server System S9200WK includes a fixed rail kit for installation into a 4-post rack or cabinet. Features and specifications for the rail kit are listed below:

- **Intel Product Order Code FCXXRAILKIT**– Fixed Rail Kit
  - Tool-less installation
  - Max supported weight: 330 lbs. (150 kg)

SAFETY NOTE: Due to the weight of a fully configured system, Intel recommends the following: Use a mechanical lift to aid with the installation of the system into the rack, and/or use at least two people to install the system into the rack, or remove all installed compute modules from the system before attempting to install the system into a rack or cabinet.

2.5 System Level Environmental Limits

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

Table 3. System environmental limits summary

Parameter		Limits
Temperature	**Operating**	ASHRAE Class A2 – Continuous Operation. 10 °C to 35 ° C (50 °F to 95 °F) with the maximum rate of change not to exceed 10 °C per hour. See Chapter 11 for more information.
	Shipping	-40 °C to 70 °C (-40 °F to 158 °F)
Altitude	**Operating**	Support operation up to 3050 m with ASHRAE class de-ratings.
Humidity	**Shipping**	50% to 90%, non-condensing with a maximum wet bulb of 28 °C (at temperatures from 25 °C to 35 °C)
Shock	**Operating**	Half sine, 2 G 11 msec

Parameter		Limits		
	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	**System Configuration**	4x S9256WK1HLC	2x S9248WK2HLC	2x S9248WK2HAC
	Server Sound Power Level	Up to 74.1 LWA (dBA)	Up to 72.9 LWA (dBA)	Up to 75.5 LWA (dBA)

The Intel® Server System S9200WK may have system configuration requirements that allow the system to operate within system thermal limits. For a complete list of supported system configurations, refer to Appendix D or the Intel® Server System S9200WK product family Power Budget and Thermal Configuration Tool.

2.6 System Packaging

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

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 the original packaging material must be used, including both boxes and each inner packaging component. In addition, all inner packaging components must be reinstalled in the proper location to ensure adequate protection of the system for subsequent shipment.

Note: The design of the inner packaging components does not prevent improper placement within the packaging assembly. There is only one correct packaging assembly that allows the package to meet the ISTA Test Procedure 3A (2008) limits. See the *Intel® Server System S9200WK Product Family Setup 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: 1192 mm
  - Breadth: 758 mm
  - Height: 317 mm
- **Inner box internal dimension**
  - Length: 1174 mm
  - Breadth: 742 mm
  - Height: 287 mm

Table 4. Product weight information

iPC	Product Type	Packaged Gross Weight (lbs)	Packaged Gross Weight (kg)	Un-packaged Net Weight (lbs)	Un-packaged Net Weight (kg)
S9248WK2HAC	Compute module	15.1	6.85	12.28	5.57
S9232WK2HAC	Compute module	15.1	6.85	12.28	5.57
S9232WK1HAC	Compute module	12.98	5.89	10.33	4.69
S9248WK2HLC	Compute module	17.81	8.07	14.9	6.7
S9232WK2HLC	Compute module	17.81	8.07	14.9	6.7
S9232WK1HLC	Compute module	16.57	7.51	13.88	6.3
S9248WK1HLC	Compute module	16.57	7.51	13.88	6.3
S9256WK1HLC	Compute module	16.57	7.51	13.88	6.3
FC2HLC21W3	Chassis	74.29	33.7	65.56	29.74
FC2HAC21W3	Chassis	59.37	26.93	50.57	22.94
FC2HAC16W3	Chassis	59.37	26.93	50.57	22.94

3. Compute Module Support

This chapter provides an overview of the compute modules within the system and their supported features. For additional information, see the *Intel® Server System S9200WK product family Configuration Guide*. The Intel® Server System S9200WK product family offers three compute module options:

- 1U liquid-cooled
- 1U air-cooled
- 2U liquid-cooled
- 2U air-cooled

Supported system configurations include four 1U liquid-cooled compute modules, two 2U liquid-cooled compute modules, two 2U air-cooled compute modules, and mixed configuration of two 1U and one 2U compute modules of the same cooling type. For a mixed configuration, Intel recommends locating the two 1U compute modules on the left-hand side and the one 2U compute module on the right-hand side as shown in Figure 5.

Each compute module within a system configuration is independently operated from the others. The following illustrations identify the features available with each compute module option.

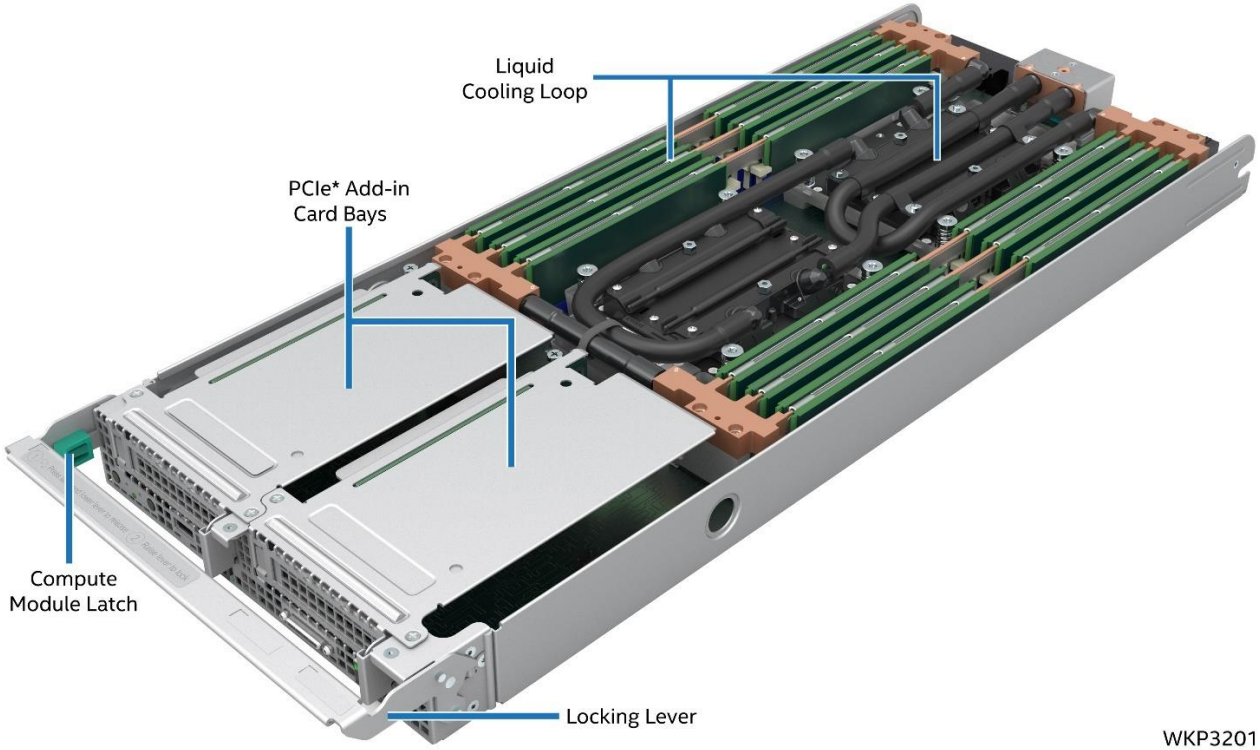

Figure 11. 1U liquid-cooled compute module feature overview

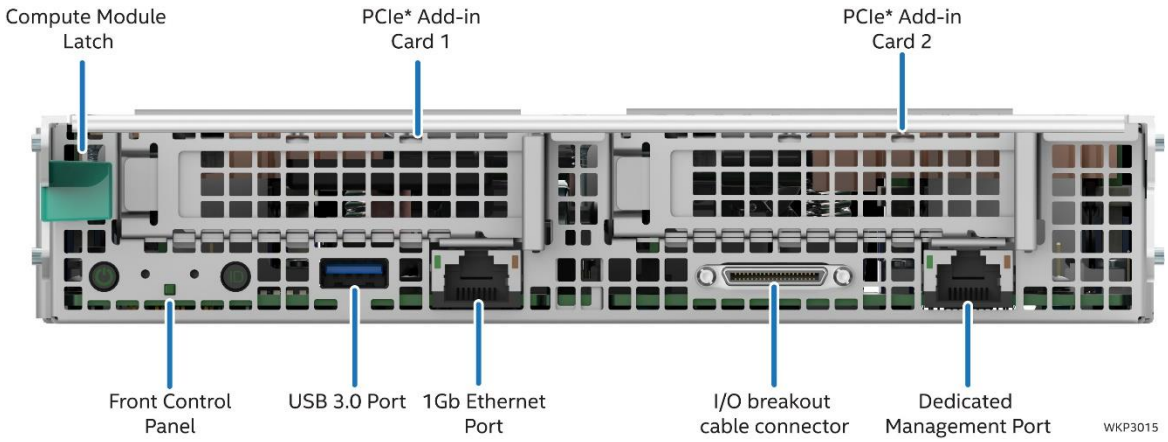

Figure 12. 1U compute module front features

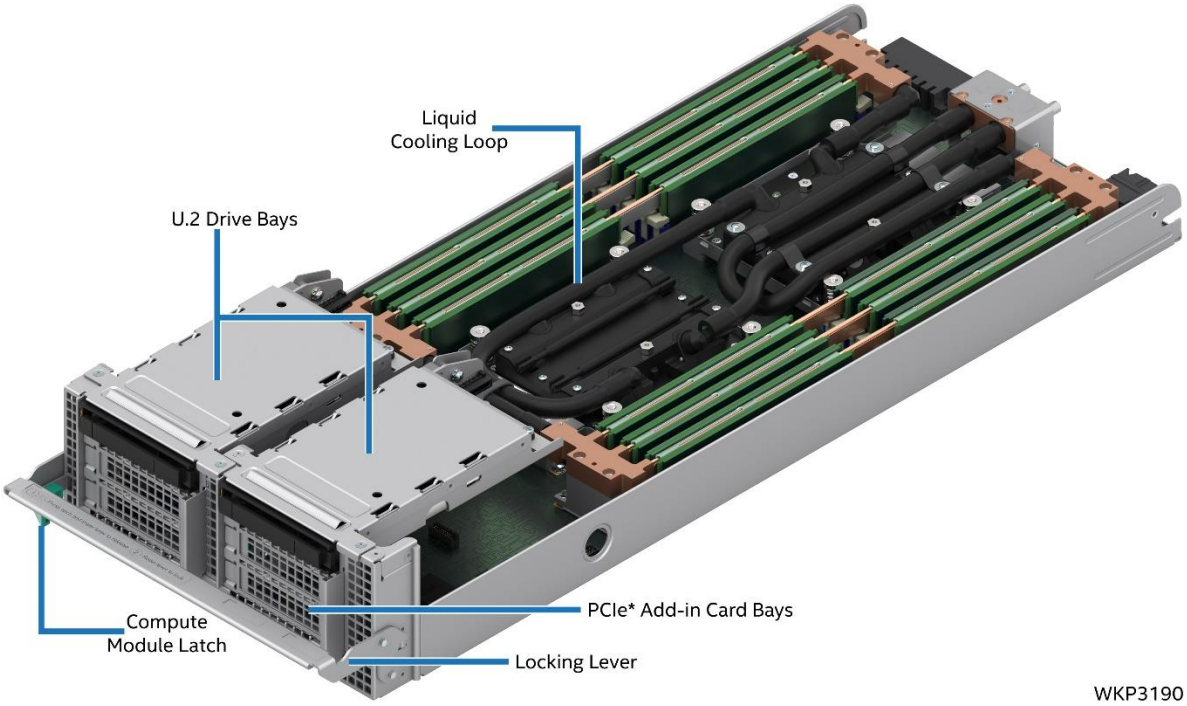

Figure 13. 2U liquid-cooled compute module feature overview

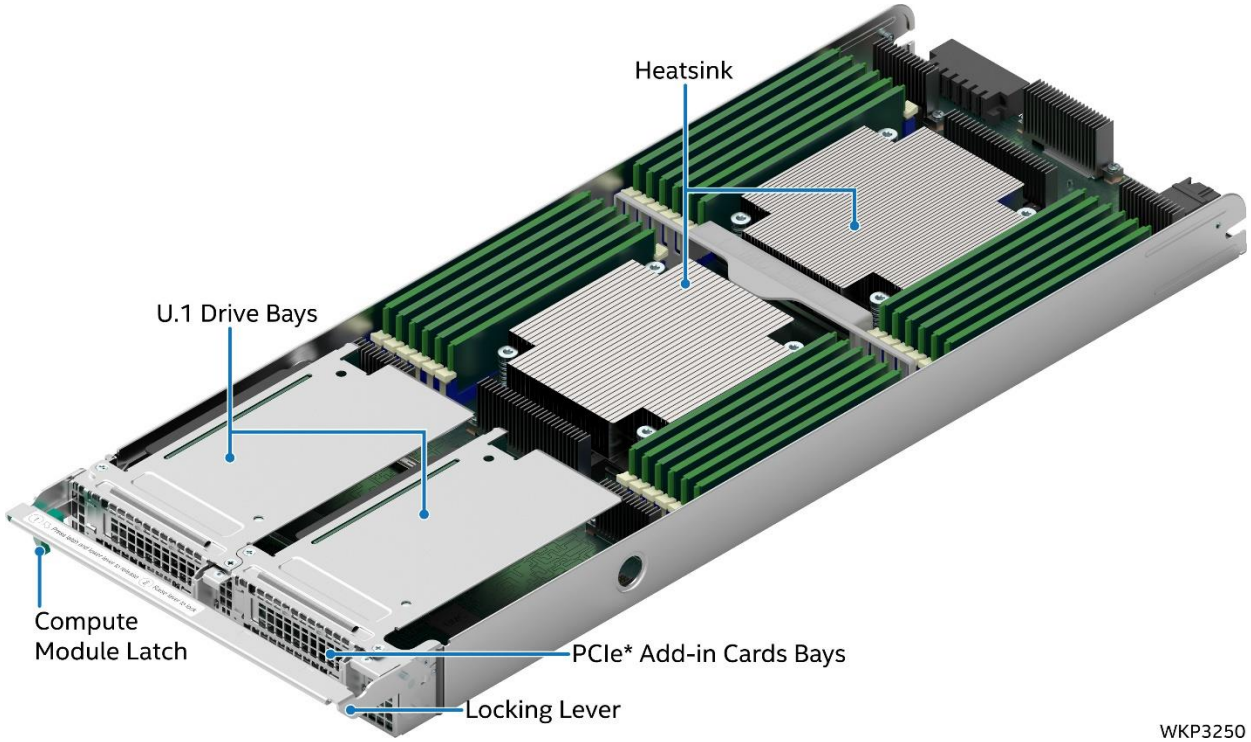

Figure 14. 1U air-cooled compute module feature overview

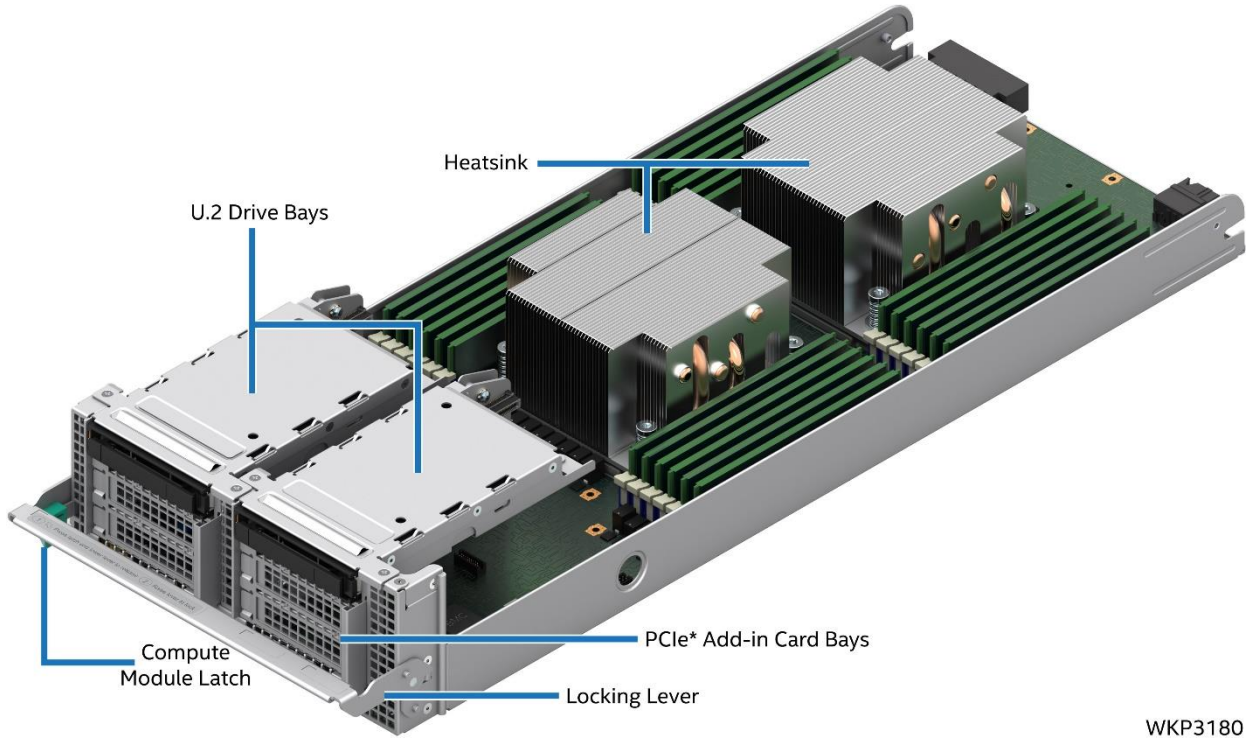

Figure 15. 2U air-cooled compute module feature overview

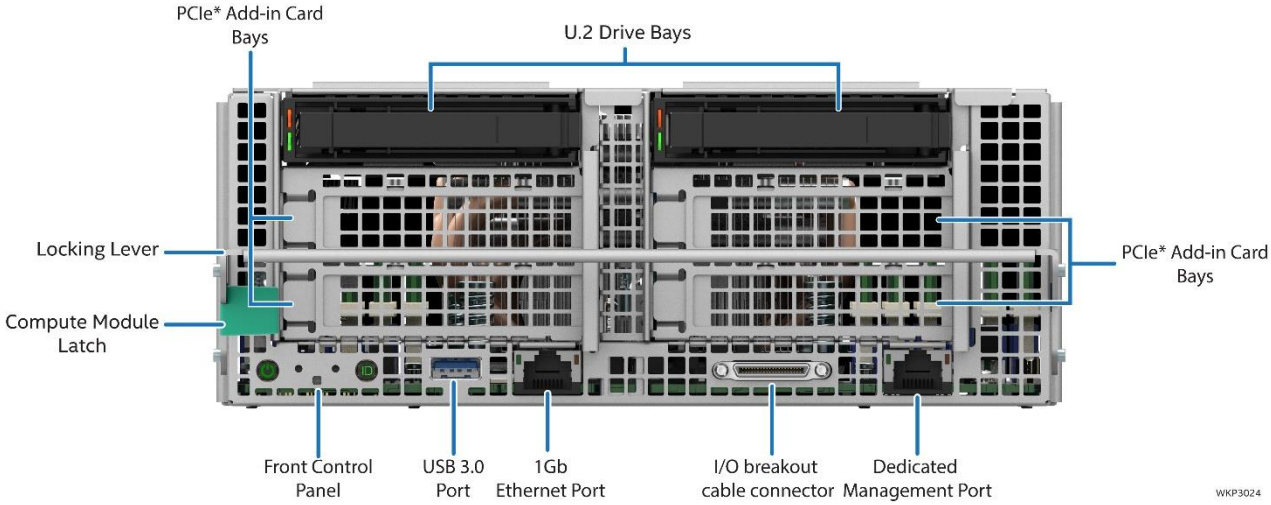

Figure 16. 2U compute module front features

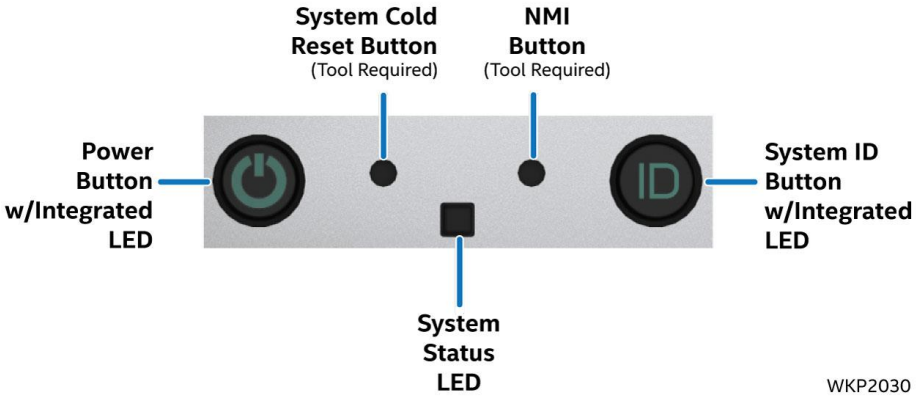

Figure 17. Front control panel features

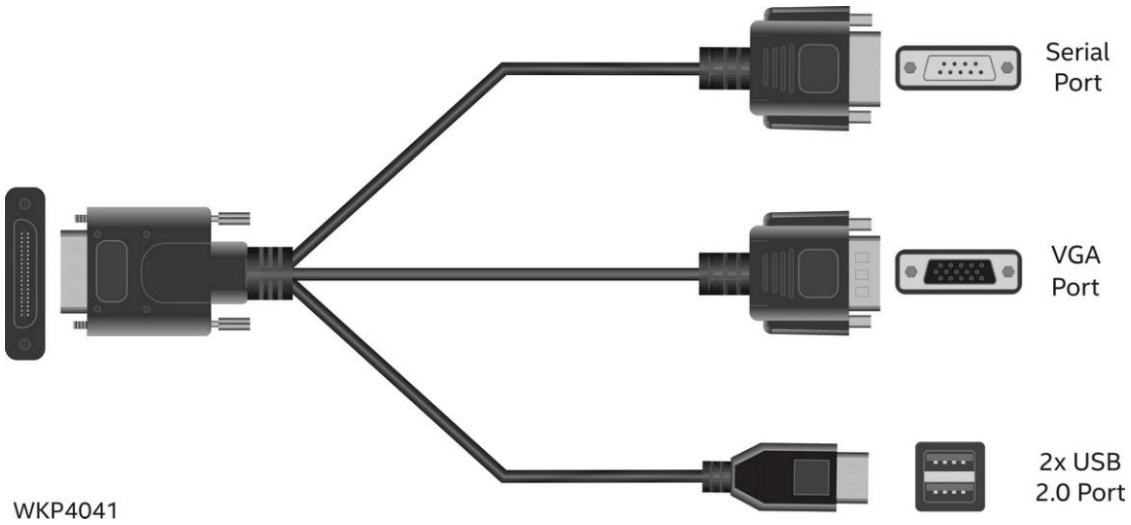

Figure 18. I/O breakout cable connector identification

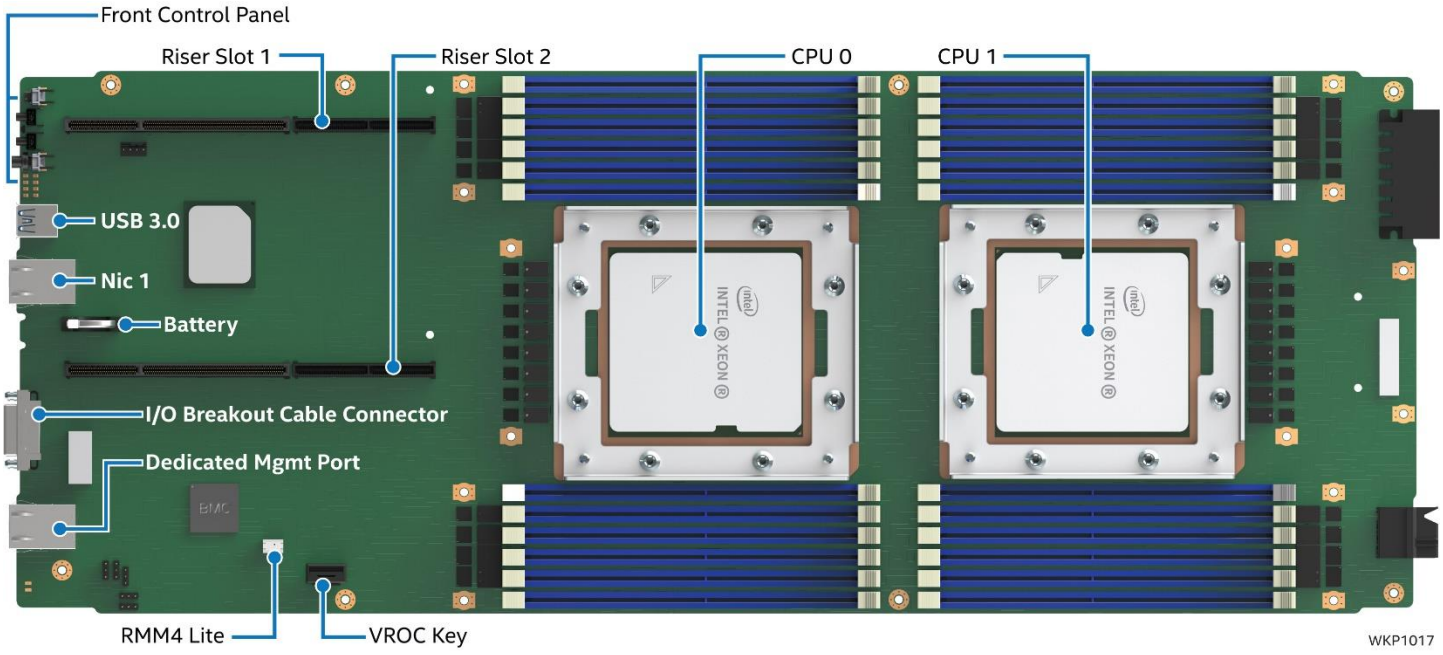

Figure 19. Compute module board feature overview

3.1 Compute Module Architecture Overview

The architecture of Intel® Server System S9200WK product family compute modules is developed around the integrated features and functions of the following components:

- Intel® Xeon® Platinum 9200 Processor
- Intel® C620 series chipset (PCH)
- Intel® Ethernet Controller I210
- ASPEED* AST2500 baseboard management controller (BMC).

Figure 20 provides an overview of the compute module server board architecture, showing the features and interconnects of each of the major sub-system components.

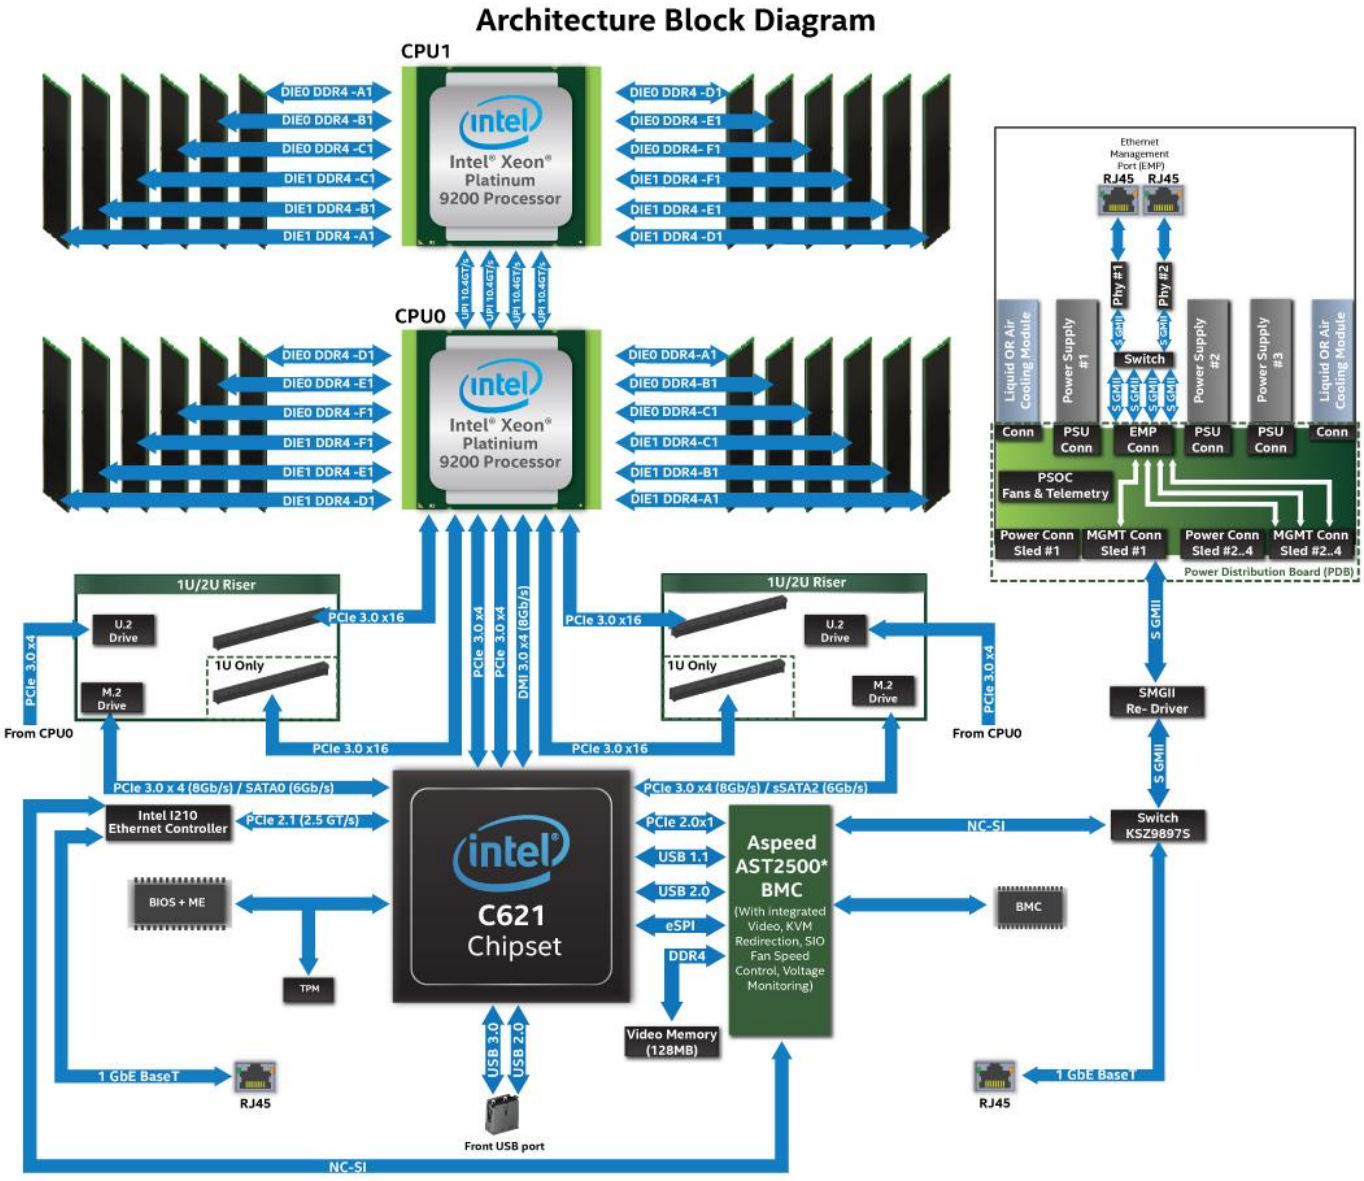

Figure 20. Compute module architecture overview

3.2 System Software Stack

The compute module includes a system software stack that consists of the following components:

- System BIOS
- BMC firmware
- Intel® Management Engine (Intel® ME) firmware
- Field replaceable unit (FRU) and sensor data record (SDR) data

Together, they configure and manage features and functions of the server system.

Many features and functions of the server system are managed jointly by the system BIOS and the BMC firmware, including:

- IPMI watchdog timer
- Messaging support, including command bridging and user/session support
- BIOS boot flags support

- Event receiver device: The BMC receives and processes events from the BIOS.
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS.
- Fault resilient booting (FRB): Fault resistant boot level 2 (FRB-2) is supported by the watchdog timer functionality.
- Front panel management: The BMC controls the system status LED and system ID LED. It supports secure lockout of certain front panel functionality and monitors button presses. The system--- ID LED is turned on using a front panel button or a command.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Integrated KVM
- Integrated remote media redirection
- Intel® Intelligent Power Node Manager support
- Sensor and system event log (SEL) logging additions/enhancements (e.g., additional thermal monitoring capability)
- Embedded platform debug feature, which allows capture of detailed data for later analysis by Intel

A factory installed system software stack is pre-programmed on each compute module. However, later revisions may be available. To ensure optimal system operation, Intel recommends the following:

- Power up the compute module and access the onboard <F2> BIOS Setup Utility to verify the version numbers of the installed system software stack: BIOS, BMC firmware, ME firmware, FRU and SDR.
- Check the following Intel website for possible updates: http://downloadcenter.intel.com
- Download and update the software stack if later revisions are available

System updates can be performed in a number of operating environments, including the UEFI shell using the UEFI-only system update package (SUP), or under different operating systems using the Intel® One Boot Flash Update (Intel® OFU) utility.

Refer to the following Intel documents for more in depth information concerning the system software stack and their functions:

- *Intel® Server Board S9200WK BIOS External Product Specification*: Intel NDA Required
- *Intel® Server System Integrated Baseboard Management Controller (BMC) Firmware External Product Specification for Intel® Servers Systems supporting the Intel® Xeon® processor Scalable family*: Intel NDA Required

3.2.1 Hot Keys Supported During POST

Certain hot keys are recognized during power-on self-test (POST). A hot key is a key or key combination that is recognized as an unprompted command input, where the operator is not prompted to press the hot key. In most cases, hot keys are recognized even while other processes are in progress.

BIOS supported hot keys are only recognized by the system BIOS during the system boot time POST process. Once the POST process has completed and hands off the boot process to the operating system, BIOS supported hot keys are no longer recognized.

Table 5 provides a list of available POST hot keys along with a description for each.

Table 5. POST hot keys

Hot Key	Function
<F2>	Enter the BIOS setup utility
<F6>	Pop-up BIOS boot menu
<F12>	Network boot

<Esc>	Switch from logo screen to diagnostic screen
<Pause>	Stop POST temporarily

3.2.1.1 POST Logo/Diagnostic Screen

If quiet boot is enabled in the BIOS setup utility, a splash screen is displayed with the standard Intel logo screen or a customized original equipment manufacturer (OEM) logo screen if one is present in the designated flash memory location. Both quiet boot and the logo screen displayed are enabled by default within the BIOS setup utility. However, pressing **<Esc>** hides the logo screen and displays the diagnostic screen instead.

If a logo is not present in the BIOS flash memory space, or if quiet boot is disabled within the system configuration, the POST diagnostic screen is displayed with a summary of system configuration information. The POST diagnostic screen is purely a text mode screen, as opposed to the graphics mode logo screen.

If console redirection is enabled in the BIOS setup utility, the quiet boot setting is disregarded and the text mode diagnostic screen is displayed unconditionally. This is due to the limitations of console redirection, which transfers data in a mode that is not graphics-compatible.

3.2.1.2 BIOS Boot Pop-Up Menu

The BIOS boot specification (BBS) provides a boot pop-up menu that can be invoked by pressing the **<F6>** key during POST. The BBS pop-up menu displays all available boot devices. The boot order in the pop-up menu is not the same as the boot order in the BIOS setup utility. The pop-up menu simply lists all of the available devices from which the system can be booted, and allows a manual selection of the desired boot device.

When an administrator password is installed in the BIOS setup utility, the administrator password is required to access the boot pop-up menu. If a user password is entered, the user is taken directly to the boot manager in the BIOS setup utility only allowing booting in the order previously defined by the administrator.

3.2.1.3 Entering BIOS Setup

To enter the BIOS setup utility using a keyboard (or emulated keyboard), press the **<F2>** function key during boot time when the OEM or Intel logo screen or the POST diagnostic screen is displayed.

The following instructional message is displayed on the diagnostic screen or under the quiet boot logo screen:

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

Note: With a USB keyboard, it is important to wait until the BIOS discovers the keyboard and beeps; until the USB controller has been initialized and the keyboard activated, key presses are not read by the system.

When the BIOS setup utility is entered, the main screen is displayed initially. However, if a serious error occurs during POST, the system enters the BIOS setup utility and displays the error manager screen instead of the main screen.

For additional BIOS setup utility information, refer to *Intel® Server Board S9200WK BIOS Setup User Guide*.

3.2.1.4 BIOS Update Capability

To bring BIOS fixes or new features into the system, it is necessary to replace the current installed BIOS image with an updated one. The BIOS image can be updated using a standalone IFLASH32 utility in the UEFI shell or using the OFU utility program under a supported operating system. Full BIOS update instructions are provided with update packages downloaded from the Intel website.

3.2.1.5 BIOS Recovery

If a system is unable to boot successfully to an OS, hangs during POST, or even hangs and fails to start executing POST, it may be necessary to perform a BIOS recovery procedure to replace a defective copy of the primary BIOS

The BIOS provides three mechanisms to start the BIOS recovery process, which is called recovery mode:

- The recovery mode jumper causes the BIOS to boot in recovery mode.
- At power on, if the BIOS boot block detects a partial BIOS update was performed, the BIOS automatically boots in recovery mode.
- The BMC asserts the recovery mode general purpose input/output (GPIO) in case of partial BIOS update and FRB-2 timeout.

The BIOS recovery takes place without any external media or mass storage device as it uses a backup BIOS image inside the BIOS flash in recovery mode.

Note: The recovery procedure is included here for general reference. However, if in conflict, the instructions in the BIOS release notes are the definitive version.

When the BIOS recovery jumper is set, the BIOS begins by logging a recovery start event to the system event log (SEL). It then loads and boots with a backup BIOS image residing in the BIOS flash device. This process takes place before any video or console is available. The system boots to the embedded UEFI shell, and a recovery complete event is logged to the SEL. From the UEFI shell, the BIOS can then be updated using a standard BIOS update procedure defined in update instructions provided with the system update package downloaded from the Intel website. Once the update has completed, switch the recovery jumper back to its default position and power cycle the system.

If the BIOS detects a partial BIOS update or the BMC asserts recovery mode GPIO, the BIOS boots in recovery mode. The difference is that the BIOS boots up to the error manager page in the BIOS setup utility. In the BIOS Setup utility, a boot device, shell or Linux*, for example, could be selected to perform the BIOS update procedure under shell or operating system environment.

Note: Before attempting a recovery boot, it is highly advisable to reference the BIOS Release Notes to verify the proper recovery procedure.

3.2.2 Field Replaceable Unit (FRU) and Sensor Data Record (SDR) Data

As part of the manufacturing process, FRU and SDR data is loaded into the compute modules within the system. This ensures that the embedded platform management system is able to monitor the appropriate sensor data and operate the system with best cooling and performance. This also ensures that auto-configuration occurs without the need to perform additional SDR updates or provide other user input to the system when any of the following components are added or removed:

- Compute module
- Memory
- Power supply
- Fan
- Power distribution board

Intel recommends updating the SDR to the latest available version whenever a system software update is performed.

Note: The system may not operate with best performance or best/appropriate cooling if the proper FRU and SDR data is not installed.

3.2.2.1 Loading FRU and SDR Data

The FRU and SDR data can be updated using a standalone FRUSDR utility in the UEFI shell, or can be done using the OFU utility program under a supported operating system. Full FRU and SDR update instructions are provided with the appropriate system update package (SUP) or OFU utility which can be downloaded from http://downloadcenter.intel.com.

4. Processor Support

Each compute module installed within the Intel® Server System S9200WK product family is configured with two non-removable Intel® Xeon® Platinum 9200 processors soldered to the server board. All compute modules within the system are configured with identical processors that meet high thermal design power (TDP) limits, ranging from 250W up to 400W, depending on the system configuration.

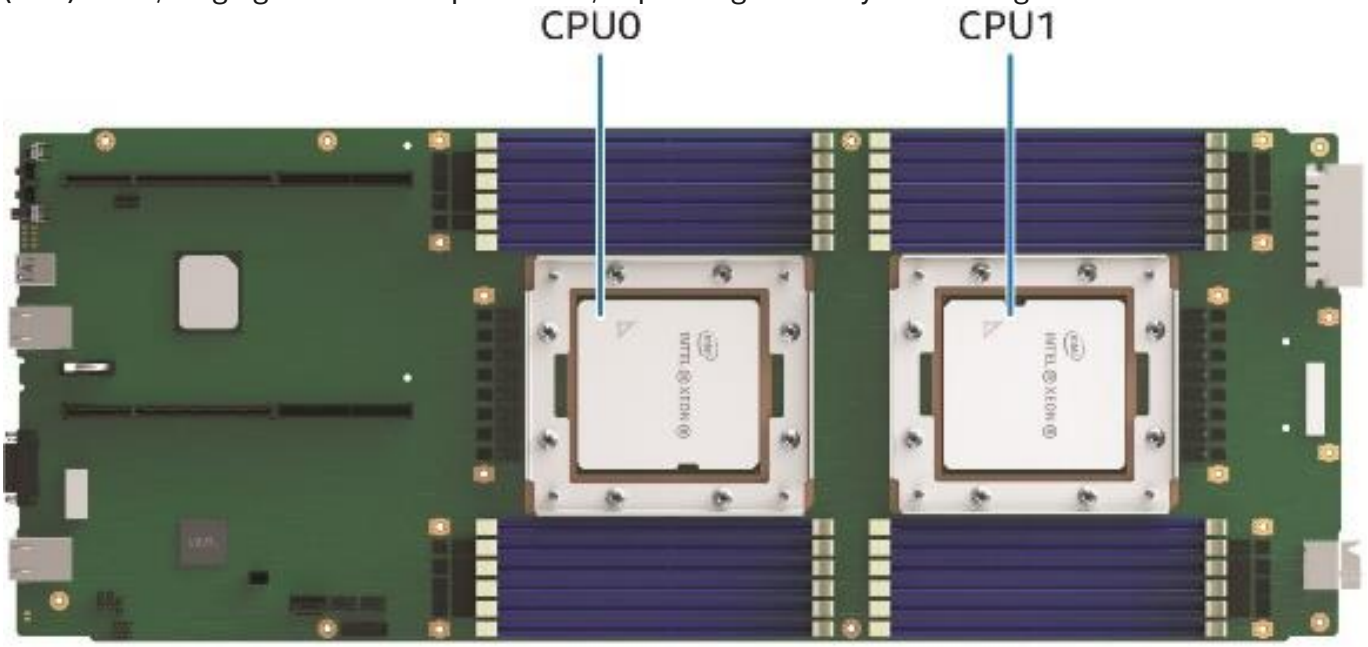

Figure 21. Processor identification

4.1 Intel® Xeon® Platinum 9200 Processor Family Overview

The Intel® Server System S9200WK product family supports the following processors:

- Intel® Xeon® Platinum 9221 Processor
- Intel® Xeon® Platinum 9222 Processor
- Intel® Xeon® Platinum 9242 Processor
- Intel® Xeon® Platinum 9282 Processor

The Intel® Xeon® Platinum 9200 processor family combines several key system components into a single processor package, including two dies with their cores, Integrated Memory Controllers (IMC), and Integrated IO Modules (IIO).

The processor includes many core, uncore features, and technologies including:

Core Features:

- Intel® Ultra Path Interconnect (UPI) – up to 10.4 GT/s
- Intel® Speed Shift Technology
- Intel® 64 Architecture
- Enhanced Intel® SpeedStep Technology
- Intel® Turbo Boost Technology 2.0
- Intel® Hyper-Threading Technology
- Intel® Virtualization Technology (Intel® VT-x)
- Intel® Virtualization Technology for Directed I/O (Intel® VT-d)
- Execute Disable Bit
- Intel® Trusted Execution Technology (Intel® TXT)

- Intel® Advanced Vector Extensions (Intel® AVX-512)
- Advanced Encryption Standard New Instructions (AES-NI)
- Intel® Deep Learning Boost through VNNI
- Intel® Speed Select Technology (select SKUs)
- Intel® Resource Director Technology

Uncore Features:

- Up to 80 PCIe* 3.0 lanes available routed to CPU0 – 79GB/s bi-directional pipeline
- 12 DDR4 memory channels supported per CPU
- DMI3/PCIe* 3.0 interface with a peak transfer rate of 8.0 GT/s.
- Non-Transparent Bridge (NTB) Enhancements – 3 full duplex NTBs and 32 MSI-X vectors
- Intel® Volume Management Device (Intel® VMD) – Manages CPU Attached NVMe SSDs
- Intel® Quick Data Technology
- Support for Intel® Node Manager 4.0

4.2 Processor Features Overview

4.2.1 Intel® 64 Instruction Set Architecture (Intel® 64)

64-bit memory extensions to the IA-32 architecture. Further details on Intel® 64 architecture and programming model can be found at http://developer.intel.com/technology/intel64/

4.2.2 Intel® Hyper-Threading Technology

The processor supports Intel® Hyper-Threading Technology (Intel® HT Technology), which allows an execution core to function as two logical processors. While some execution resources such as caches, execution units, and buses are shared, each logical processor has its own architectural state with its own set of general-purpose registers and control registers. This feature must be enabled via the BIOS and requires operating system support.

4.2.3 Enhanced Intel® SpeedStep® Technology

Processors in the Intel® Xeon® Platinum 9200 Processor family support Enhanced Intel® SpeedStep® Technology. The processors support multiple performance states, allowing the system to dynamically adjust processor voltage and core frequency as needed to enable lower power consumption and lower heat production. All controls for transitioning between states are centralized within the processor, allowing for an increased frequency of transitions for better operation.

The Enhanced Intel® Speedstep® Technology feature may be enabled/disabled by an option in the Processor Configuration Setup screen, but is enabled by default. If Enhanced Intel® Speedstep Technology is disabled, then the processor speed is set to the processor's Max TDP Core Frequency (nominal rated frequency).

4.2.4 Intel® Turbo Boost Technology 2.0

Intel® Turbo Boost Technology is featured on all processors in the Intel® Xeon® Platinum 9200 Processor family. Intel® Turbo Boost Technology opportunistically and automatically allows the processor to run faster than the marked frequency if the processor is operating below power, temperature, and current limits, resulting in increased performance for both multi-threaded and single-threaded workloads.

4.2.5 Intel® Virtualization Technology (Intel® VT-x)

Intel® Virtualization Technology (Intel® VT-x) provides improved performance and robustness for virtualization via hardware support in the core. Intel® VT-x specifications and functional descriptions are included in the Intel® 64 and IA-32 Architectures Software Developer's Manual.

4.2.6 Intel® Virtualization Technology for Directed I/O (Intel® VT-d)

Intel® Virtualization Technology for Directed I/O (Intel® VT-d) provides hardware support in the core and uncore implementations to support and improve I/O virtualization performance and robustness.

4.2.7 Execute Disable Bit

Intel's Execute Disable Bit functionality helps to prevent certain classes of malicious buffer overflow attacks when combined with a supported operating system. This allows the processor to classify areas in memory by where application code can execute and where it cannot. When malicious code attempts to insert code in the buffer, the processor disables code execution, preventing damage and further propagation.

4.2.8 Intel® Trusted Execution Technology for servers (Intel® TXT)

Intel® TXT defines platform-level enhancements that provide the building blocks for creating trusted platforms. The Intel® TXT platform helps to provide the authenticity of the controlling environment such that those wishing to rely on the platform can make an appropriate trust decision. The Intel® TXT platform determines the identity of the controlling environment by accurately measuring and verifying the controlling software.

4.2.9 Intel® Advanced Vector Extension (Intel® AVX-512)

The base of the 512-bit SIMD instruction extensions are referred to as Intel® AVX-512 foundation instructions. They include extensions of the AVX family of SIMD instructions but are encoded using a new encoding scheme with support for 512-bit vector registers, up to 32 vector registers in 64-bit mode, and conditional processing using opmask registers.

4.2.10 Advanced Encryption Standard New Instructions (AES-NI)

Intel® Advanced Encryption Standard New Instructions (AES-NI) is a set of instructions implemented in all processors in the Intel® Xeon® Platinum 9200 Processor family. This feature adds AES instructions to accelerate encryption and decryption operations used in the Advanced Encryption Standard. The Intel® AES-NI feature includes 6 additional Single Instruction Multiple Data (SIMD) instructions in the Intel® Streaming SIMD Extensions (SSE) instruction set.

The BIOS is responsible in POST to detect whether the processor has the AES-NI instructions available. Some processors may be manufactured without AES-NI instructions.

The AES-NI instructions may be enabled or disabled by the BIOS. AES-NI instructions are enabled unless the BIOS has explicitly disabled them.

4.2.11 Intel® Node Manager 4.0

The Intel® C620 series chipset Intel® Management Engine (Intel® ME) supports Intel® Intelligent Power Node Manager (Intel® NM) technology. The Intel® ME/Intel® NM combination is a power and thermal control capability on the platform, that exposes external interfaces that allow IT (through external management software) to query the ME about platform power capability and consumption, thermal characteristics, and specify policy directives like setting a platform power budget. The Intel® ME enforces these policy directives by controlling the power consumption of underlying subsystems using available control mechanisms (such as processor P/T states). The determination of the policy directive is done outside of the Intel® ME either by intelligent management software or by the IT operator.

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

- **Platform Power Monitoring and Limiting:** The Intel® ME/Intel® NM monitors platform power consumption and holds average power over duration. It can be queried to return actual power at any given instance. The power limiting capability is to allow external management software to address key IT issues by setting a power budget for each server.

- **Inlet Air Temperature Monitoring:** The Intel® ME/Intel® NM monitors server inlet air temperatures periodically. If there is an alert threshold in effect, then Intel® ME/Intel® NM issues an alert when the inlet (room) temperature exceeds the specified value. The threshold value can be set by policy.
- **Memory Subsystem Power Limiting:** The Intel® ME/Intel® NM monitors memory power consumption. Memory power consumption is estimated using average bandwidth utilization information.
- **Processor Power monitoring and limiting:** The Intel® ME/Intel® NM monitors processor or socket power consumption and holds average power over duration. It can be queried to return actual power at any given instant. The monitoring process of the Intel® ME will be used to limit the processor power consumption through processor P-states and dynamic core allocation.
- **Core allocation at boot time:** Restrict the number of cores for OS/VMM use by limiting how many cores are active at boot time. After the cores are turned off, the CPU will limit how many working cores are visible to the BIOS and OS/VMM. The cores that are turned off cannot be turned on dynamically after the OS has started. It can be changed only at the next system reboot.
- **Core allocation at run-time:** This particular use case provides a higher level processor power control mechanism to a user at runtime, after booting. An external agent can dynamically use or not use cores in the processor subsystem by requesting Intel® ME/Intel® NM to control them, specifying the number of cores to use or not use.

4.2.12 Intel® Deep Learning Boost

Intel® Deep Learning Boost on the Intel® Xeon® Platinum 9200 Processor family is designed to deliver more efficient Deep Learning (Inference) acceleration by expanding the capabilities of Intel® AVX-512 through Intel® Vector Neural Network Instructions (VNNI) dedicated to Deep Learning tasks. Consult the *Intel® 64 and IA-32 Architectures Software Developer's Manual* for details.

4.2.13 Intel® Speed Select Technology

Intel® Speed Select Technology, available on select Intel® Xeon® Platinum 9200 Processor family SKUs, offers three distinct operating voltage-frequency points for guaranteed base frequency (P1). This frequency is based on the number of active cores within the SKU only when thermal requirements are met. Intel® Speed Select Technology allows either a higher active core count with lower base frequency or a lower active core count with higher base frequency, providing multiple CPU personalities based on workload/VM needs.

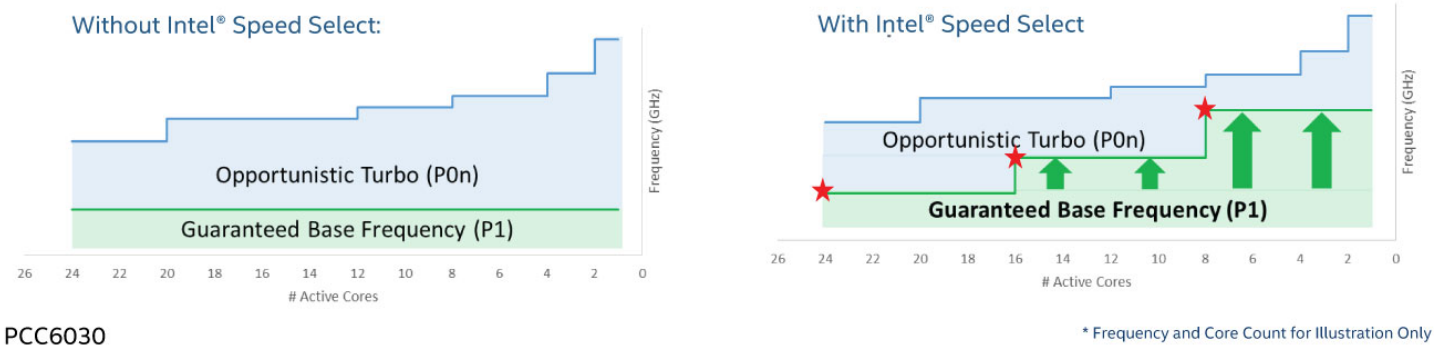

Figure 22. Intel® Speed Select Technology comparison

4.2.14 Intel® Resource Director Technology

Intel® Resource Directory Technology, available on the Intel® Xeon® Platinum 9200 Processor family, mitigates execution contention when several applications, containers, or virtual machines are sharing platform resources. Software threads are able to have memory bandwidth according to their priority, not just CPU time, and is achieved with the following features:

- **Cache Monitoring Technology (CMT):** monitors LLC (L3 cache) usage by each software thread, through Resource Monitoring ID (RMID).
- **Code Data Prioritization (CDP):** provides the capability to separate code from data in LLC using masks.
- **Memory Bandwidth Monitoring (MBM):** gives the OS/VMM the abilities of assigning RMID to software threads and read the memory bandwidth utilization for a given RMID.
- **Memory Bandwidth Allocation (MBA):** MBA is a new feature introduced in the Intel® Xeon® Platinum 9200 Processor family that enables software to control the amount of memory bandwidth a thread or core can consume based on credits.

4.3 Processor Initialization Error Summary

Table 6 describes processor conditions and recommended actions for Intel server systems designed around the Intel® Xeon® Platinum 9200 processor and Intel® C620 series chipset architecture. The errors fall into one of the following categories:

- **Fatal**: If the system cannot boot, POST halts and display the following message:
  `Unrecoverable fatal error found. System will not boot until the error is resolved`
  `Press <F2> to enter setup`
  When the **<F2>** key on the keyboard is pressed, the error message is displayed on the error manager screen and an error is logged to the system event log (SEL) with the POST error code.
  The “POST Error Pause” option setting in the BIOS setup does not have any effect on this error.
  If the system is not able to boot, the system generates a beep code consisting of three long beeps and one short beep. The system cannot boot unless the error is resolved. The faulty component must be replaced.
  The system status LED is 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.
- **Major**: An error message is displayed to the error manager screen and an error is logged to the SEL. If the BIOS setup option “Post Error Pause” is enabled, operator intervention is required to continue booting the system. If the BIOS setup option “POST Error Pause” is disabled, the system continues to boot.
- **Minor**: An error message may be displayed to the screen or to the BIOS setup error manager and the POST error code is logged to the SEL. The system continues booting in a degraded state. The user may want to replace the erroneous unit. The “POST Error Pause” option setting in the BIOS setup does not have any effect on this error.

Table 6. Processor error summary

Error	Severity	System Action when BIOS Detects the Error Condition
Processor microcode update failed	Major	Logs the POST error code into the SEL. Displays 816x: Processor 0x unable to apply microcode update message in the error manager or on the screen. Takes major error action. The system may continue to boot in a degraded state, depending on the "POST Error Pause" setting in setup, or may halt with the POST error code in the error manager waiting for operator intervention.
Processor microcode update missing	Minor	Logs the POST error code into the SEL. Displays 818x: Processor 0x microcode update not found message in the error manager or on the screen. The system continues to boot in a degraded state, regardless of the "POST Error Pause" setting in setup.

5. Memory Support

Each compute module within the Intel® Server System S9200WK product family supports up to 24 dual inline memory modules (DIMMs): 12 DIMMs per processor and one DIMM per memory channel.

Figure 23. Memory subsystem architecture

5.1 Supported Memory

The Intel® Server System S9200WK supports the following:

- DDR4 DIMMs only
- Error Correction Code (ECC) enabled Registered DIMMs (RDIMMs) or Load-Reduced DIMMs (LRDIMMs) only
- RDIMMs and LRDIMMs with thermal sensor on-DIMM (TSOD) only
- RDIMM sizes of 8 GB, 16 GB, and 32 GB.
- LRDIMM sizes of 64 GB and 128 GB
- DIMMs organized as Single Rank (SR), Dual Rank (DR), Quad Rank (QR), or Oct Rank (OR).

Refer to the *Intel® Server System S9200WK Product Family Configuration Guide* for compatible DIMMs.

5.2 Memory Slot Identification and Population Rules

The following memory population rules apply when installing DIMMs:

- Only multiples of 8 DIMMs are supported to be installed on a compute module (8,16,24). For more information, see Section 5.3 DIMM Population Guidelines for Best Performance.
- On the compute module, each DIMM slot is labeled by CPU #, die #, memory channel, and slot # such as CPU0_0_DIMM_A1.
- Mixing DIMMs of different frequencies and latencies is not supported within or across processors. If a mixed configuration is encountered, the BIOS attempts to operate at the highest common frequency and the lowest latency possible.
- Mixing of DDR4 DIMM Types (RDIMM, LRDIMM) within the processor attached DIMM slots or across processors is not supported. This is a Fatal Error Halt in Memory Initialization.

Intel DDR4 DIMM Support Disclaimer:

Intel validates and will only provide support for system configurations where all installed DDR4 DIMMs have matching "Identical" or "Like" attributes. See Table 7. A system configured concurrently with DDR4 DIMMs from different vendors will be supported by Intel if all other DDR4 "Like" DIMM attributes match.

Intel does not perform system validation testing nor will it provide support for system configurations where all populated DDR4 DIMMs do not have matching "Like" DIMM attributes as listed in Table 7.

Intel will only provide support for Intel server systems configured with DDR4 DIMMs that have been validated by Intel and are listed on Intel's Tested Memory list for the given Intel server product family.

Intel configures and ships pre-integrated L9 server systems. All DDR4 DIMMs within a given L9 server system as shipped by Intel will be identical. All installed DIMMs will have matching attributes as those listed in the *"Identical" DDR4 DIMM4 Attributes* column in Table 7.

When purchasing more than one integrated L9 server system with the same configuration from Intel, Intel reserves the right to use "Like" DIMMs between server systems. At a minimum "Like" DIMMS will have matching DIMM attributes as listed in the table below. However, the DIMM model #, revision #, or vendor may be different.

For warranty replacement, Intel will make every effort to ship back an exact match to the one returned. However, Intel may ship back a validated "Like" DIMM. A "Like" DIMM may be from the same vendor but may not be the same revision # or model #, or it may be an Intel validated DIMM from a different vendor. At a minimum, all "Like" DIMMs shipped from Intel will match attributes of the original part according to the definition of "Like" DIMMs in the following table.

Table 7. DDR4 DIMM Attributes Table for "Identical" and "Like" DIMMs

• DDR4 DIMMs are considered "Identical" when ALL listed attributes between the DIMMs match • Two or more DDR4 DIMMs are considered "Like" DIMMs when all attributes minus the Vendor, and/or DIMM Part # and/or DIMM Revision#, are the same.			
Attribute	**"Identical" DDR4 DIMM Attributes**	**"Like" DDR4 DIMM Attributes**	**Possible DDR4 Attribute Values**
Vendor	**Match**	**Maybe Different**	Memory Vendor Name
DIMM Part #	**Match**	**Maybe Different**	Memory Vendor Part #
DIMM Revision #	**Match**	**Maybe Different**	Memory Vendor Part Revision #
SDRAM Type	**Match**	**Match**	DDR4
DIMM Type	**Match**	**Match**	RDIMM, LRDIMM
Speed (MHz)	**Match**	**Match**	2666, 2933, 3200
Voltage	**Match**	**Match**	1.2V
DIMM Size (GB)	**Match**	**Match**	8GB, 16GB, 32GB, 64GB, 128GB, 256GB
Organization	**Match**	**Match**	1Gx72; 2Gx72; 4Gx72; 8Gx72; 16Gx72; 32Gx72
DIMM Rank	**Match**	**Match**	1R, 2R, 4R, 8R
DRAM Width	**Match**	**Match**	x4, x8
DRAM Density	**Match**	**Match**	8Gb, 16Gb

5.3 DIMM Population Guidelines for Best Performance

Processors within the Intel® Xeon® Platinum 9200 Processor family include two integrated memory controllers (IMC) per die, each supporting three memory channels.

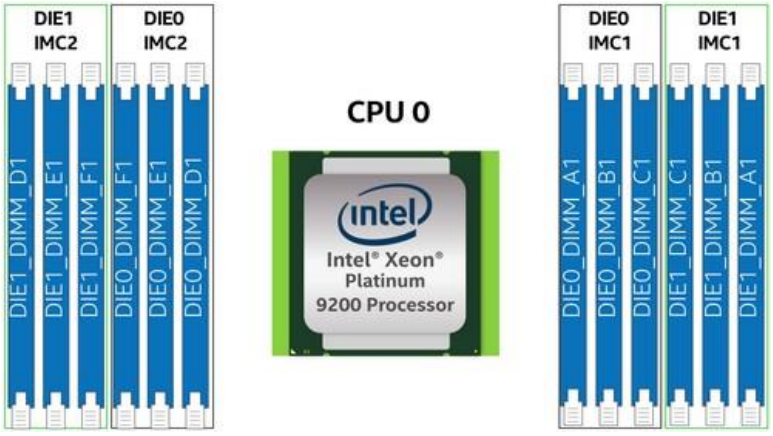

Memory population rules require that DIMMs within a compute module be populated in sets of 8 (8,16, and 24). For best performance each processor should have matching DIMM configurations. DIMMs should be populated using the following guidelines:

- DIMMs should be populated evenly across memory controllers within a processor.
- Memory channels should be populated in alphabetical order.
- For liquid cooled compute modules with DIMMs of 64GB capacity, DIMMs must be populated in pairs to each side of the memory heat spreaders within the liquid cooling loop, and a DIMM retention clip must be installed to secure them to the DIMM heat spreader.

The following illustrations show the recommended DIMM populations for best performance for liquid cooled and air cooled compute modules with 8 or 16 DIMM configurations. Refer to the *Intel® Server System S9200WK product family Setup and Service Guide* for DIMM installation instructions.

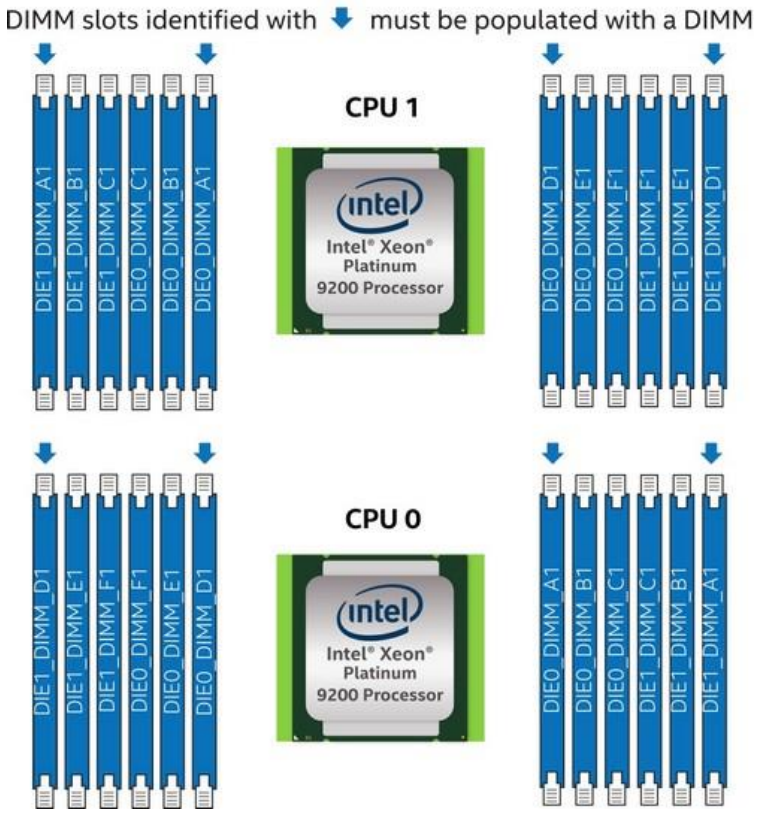

Figure 24. DIMM population for liquid-cooled compute modules with eight DIMMs of up to 32GB capacity

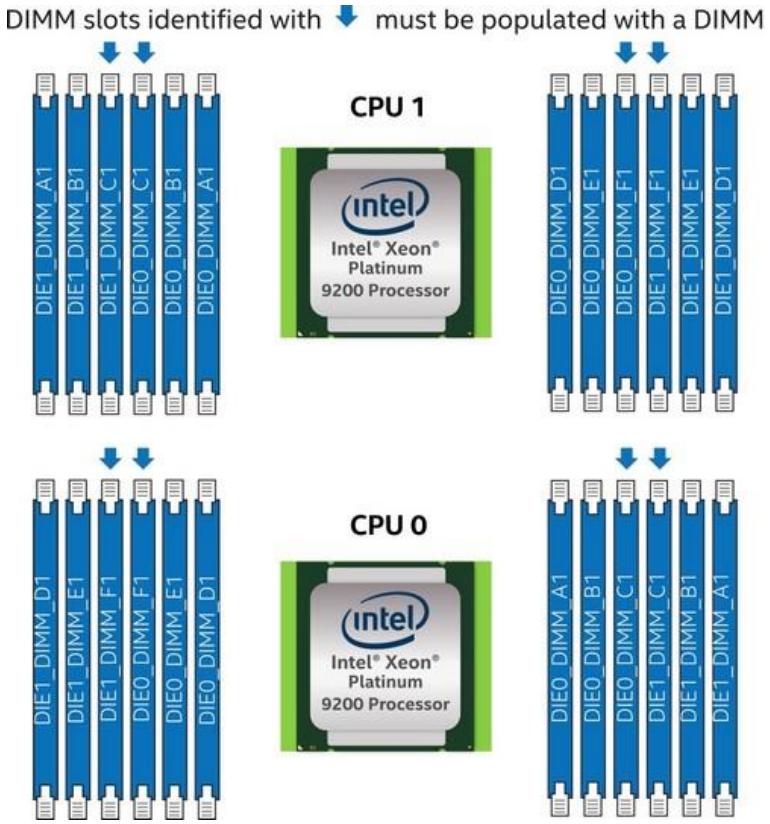

Figure 25. DIMM population for liquid-cooled compute modules with eight DIMMs of 64GB capacity

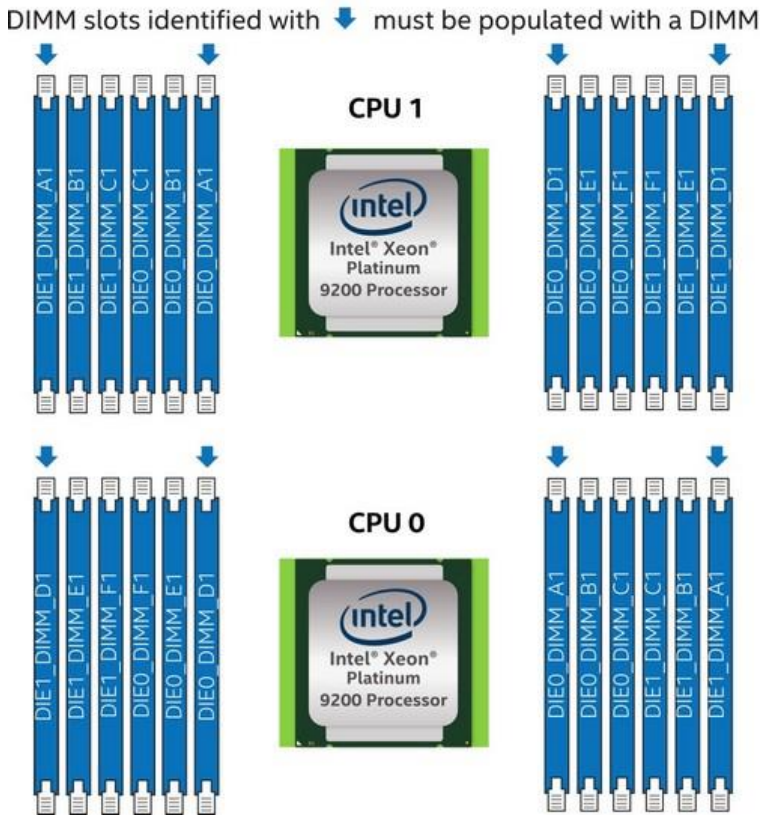

Figure 26. DIMM population for air-cooled compute modules with eight DIMMs

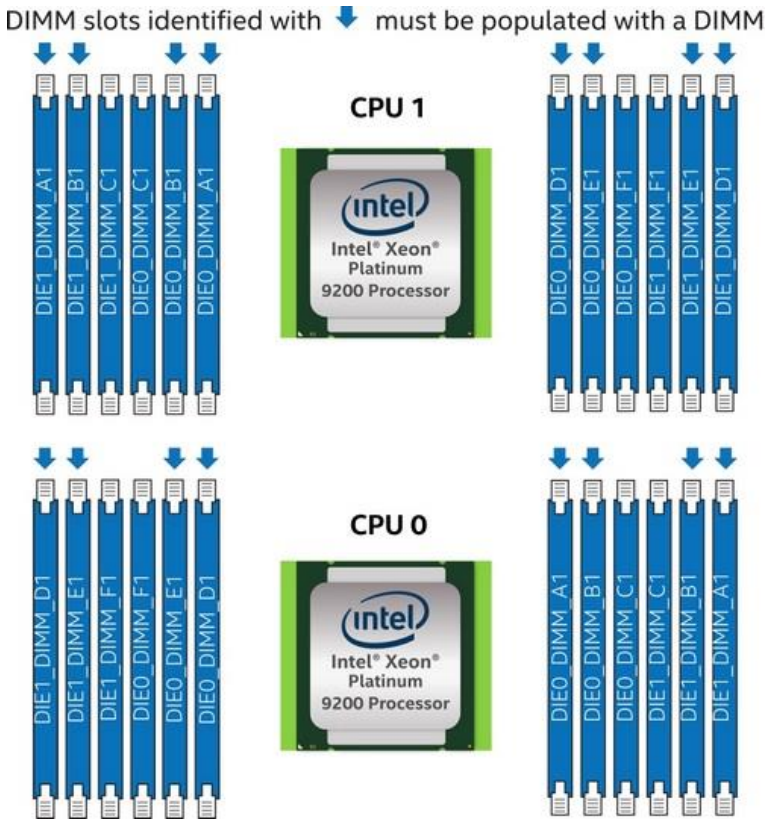

Figure 27. DIMM population for liquid and air-cooled compute modules with 16 DIMMs

5.4 Memory RAS Features

Supported memory RAS features are dependent upon the level of processor installed. Each processor within the Intel® Xeon® Platinum 9200 processor Family has support for advanced memory RAS features as defined in Table 8.

Table 8. Memory RAS features

RAS Feature	Description
Device Data Correction	x8 Single Device Data Correction (SDDC) via static virtual lockstep (Applicable to x8 DRAM DIMMs)
	Adaptive Data Correction (SR) (Applicable to x4 DRAM DIMMs)
	x8 Single Device Data Correction + 1 bit (SDDC+1) (Applicable to x8 DRAM DIMMs)
	SDDDC + 1, and ADDDC (MR) + 1 (Applicable to x4 DRAM DIMMs)
DDR4 Command/Address Parity Check and Retry	DDR4 Command/Address Parity Check and Retry: Is a DDR4 technology based CMD/ADDR parity check and retry with following attributes: CMD/ADDR Parity error “address” logging CMD/ADDR Retry
DDR4 Write Data CRC Protection	DDR4 Write Data CRC Protection detects DDR4 data bus faults during write operation.
Memory Demand and Patrol Scrubbing	Demand scrubbing is the ability to write corrected data back to the memory once a correctable error is detected on a read transaction. Patrol scrubbing proactively searches the system memory, repairing correctable errors. Prevents accumulation of single bit errors.
Memory Mirroring	Full Memory Mirroring: An intra IMC method of keeping a duplicate (secondary or mirrored) copy of the contents of memory as a redundant backup for use if the primary memory fails. The mirrored copy of the memory is stored in memory of the same processor socket's IMC. Dynamic (without reboot) failover to the mirrored DIMMs is transparent to the OS and applications.
	Address Range/Partial Memory Mirroring: Provides further intra socket granularity to mirroring of memory by allowing the firmware or OS to determine a range of memory addresses to be mirrored, leaving the rest of the memory in the socket in non-mirror mode.
Sparing **Rank Level Memory Sparing** **Multi-rank Level Memory Sparing**	Dynamic fail-over of failing Ranks to spare Ranks behind the same memory controller DDR ranks.
	With Multi Rank up to two ranks out of a maximum of eight ranks can be assigned as spare ranks.
iMC’s Corrupt Data Containment	Corrupt Data Containment is a process of signaling error along with the detected UC data. iMC's patrol scrubber and sparing engine have the ability to poison the UC data.
Failed DIMM Isolation	Ability to identify a specific failing DIMM thereby enabling the user to replace only the failed DIMM(s). In case of uncorrected error and lockstep mode, only DIMM-pair level isolation granularity is supported.
Memory Disable and Map Out for FRB	Allows memory initialization and booting to OS even when memory fault occurs.
Post Package Repair	Starting with DDR4 technology there is an additional capability available known as PPR (Post Package Repair). PPR offers additional spare capacity within the DDR4 DRAM that can be used to replace faulty cell areas detected during system boot time.

Note: RAS features may not be supported on all SKUs of a processor type.

5.4.1 DIMM Populations Rules and BIOS Setup for Memory RAS

- In order to support memory sparing and memory mirroring options, a compute module must be populated with at least 16 DIMMs.
- Memory sparing and memory mirroring options are enabled in the BIOS setup.
- Memory sparing and memory mirroring options are mutually exclusive. Only one operating mode may be selected in the BIOS setup.
- If a RAS mode is enabled, and the memory configuration is unable to support it during boot, the system will fall back to independent channel mode, logging and displaying errors.
- Rank sparing mode is only possible when all channels that are populated with memory that meet the requirement of having at least two single-rank or double-rank DIMMs installed, or at least one quad-rank DIMM installed, on each populated channel.
- Memory mirroring mode requires that for any channel pair that is populated with memory, the memory population on both channels of the pair must be identically sized.

6. PCIe* Support

The PCI Express* (PCIe*) interface of the compute modules within the Intel® Server System S9200WK product family is fully compliant with the PCIe* Base Specification, Revision 3.0 supporting the following PCIe* bit rates: 3.0 (8.0 GT/s), 2.0 (5.0 GT/s), and 1.0 (2.5 GT/s). All available PCIe* ports are routed to CPU0.

For specific features and functions supported by the PCIe* sub-system, see Chapter 7. Table 9 provides the PCIe* port routing information in the compute modules:

Table 9. CPU0 - PCIe* port routing

CPU 0					
DIE 0			**DIE 1**		
Port	**Width**	**Usage**	**Port**	**Width**	**Usage**
Port 1A	x4	Riser 2, 1U PCIe* Slot / 2U bottom PCIe* slot	Port 1A	x4	Riser 2, 2U top PCIe* Slot
Port 1B	x4	Riser 2, 1U PCIe* Slot / 2U bottom PCIe* slot	Port 1B	x4	Riser 2, 2U top PCIe* Slot
Port 1C	x4	Riser 2, 1U PCIe* Slot / 2U bottom PCIe* slot	Port 1C	x4	Riser 2, 2U top PCIe* Slot
Port 1D	x4	Riser 2, 1U PCIe* Slot / 2U bottom PCIe* slot	Port 1D	x4	Riser 2, 2U top PCIe* Slot
Port 2A	x4	Riser 2, 2U U.2 connector	Port 2A	x4	Riser 1, 1U PCIe* Slot / 2U bottom PCI slot
Port 3A	x4	Riser 1, 2U U.2 connector	Port 2B	x4	Riser 1, 1U PCIe* Slot / 2U bottom PCI slot
Port 4A	x4	Riser 1, 2U top PCIe* Slot	Port 2C	x4	Riser 1, 1U PCIe* Slot / 2U bottom PCI slot
Port 4B	x4	Riser 1, 2U top PCIe* Slot	Port 2D	x4	Riser 1, 1U PCIe* Slot / 2U bottom PCI slot
Port 4C	x4	Riser 1, 2U top PCIe* Slot			
Port 4D	x4	Riser 1, 2U top PCIe* Slot			

6.1.1 PCIe* Enumeration and Allocation

The BIOS assigns PCI bus numbers in a depth-first hierarchy, in accordance with the PCI Local Bus Specification, Revision 3.0. The bus number is incremented when the BIOS encounters a PCI-PCI bridge device.

Scanning continues on the secondary side of the bridge until all subordinate buses are assigned numbers. PCI bus number assignments may vary from boot to boot with varying presence of PCI devices with PCI-PCI bridges.

If a bridge device with a single bus behind it is inserted into a PCI bus, all subsequent PCI bus numbers below the current bus are increased by one. The bus assignments occur once, early in the BIOS boot process, and never change during the pre-boot phase.

The BIOS resource manager assigns the PIC-mode interrupt for the devices that are accessed by the legacy code. The BIOS ensures that the PCI BAR registers and the command registers for all devices are correctly set up to match the behavior of the legacy BIOS after booting to a legacy OS. Legacy code cannot make any assumption about the scan order of devices or the order in which resources are allocated to them. The BIOS automatically assigns IRQs to devices in the system for legacy compatibility. A method is not provided to manually configure the IRQs for devices.

6.1.2 Non-Transparent Bridge

The PCIe* Non-Transparent Bridge (NTB) acts as a gateway that enables high performance, low latency communication between two PCIe* Hierarchies, such as a local and remote system. The NTB allows a local

processor to independently configure and control the local system and provides isolation of the local host memory domain from the remote host memory domain, while enabling status and data exchange between the two domains. The NTB is discovered by the local processor as a Root Complex Integrated Endpoint (RCiEP).

Figure 28 shows two systems that are connected through an NTB. Each system is a completely independent PCIe* hierarchy. The width of the NT Link can be x16, x8, or x4 at the expense of other PCIe* root ports. Only Port A can be configured as an NT port.

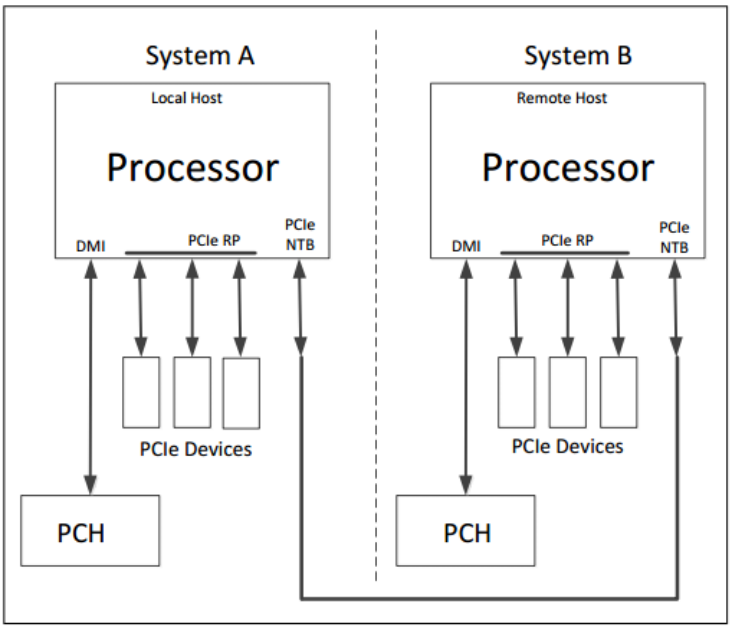

Figure 28. Two systems connected through an NTB

The specified processor family supports the following NTB features.

The NTB only supports one configuration/connection model:

- NT Port attached to another NT Port of the same component type and generation
- The NTB provides Direct Address Translation between the two PCIe* Hierarchies through two separate regions in Memory Space. Accesses targeting these Memory addresses are allowed to pass through the NTB to the remote system. This mechanism enables the following transactions flows through the NTB:
    - Both Posted Mem Writes and Non-Posted Mem Read transactions across the NTB
    - Peer-to-Peer Mem Read and Write transactions to and from the NTB

In addition, the NTB provides the ability to interrupt a processor in the remote system through a set of Doorbell registers. A write to a Doorbell register in the local side of the NTB will generate an interrupt to the remote processor. Since the NTB is designed to be symmetric, the converse is also true.

7. System I/O

7.1 Networking

The compute modules within the Intel® Server System S9200WK product family include two RJ45 connectors that provide support for the following features:

- 1Gb Network Interface Ethernet Port
- Dedicated 1Gb server management port

The following illustrations show the location of the different ports on the compute modules.

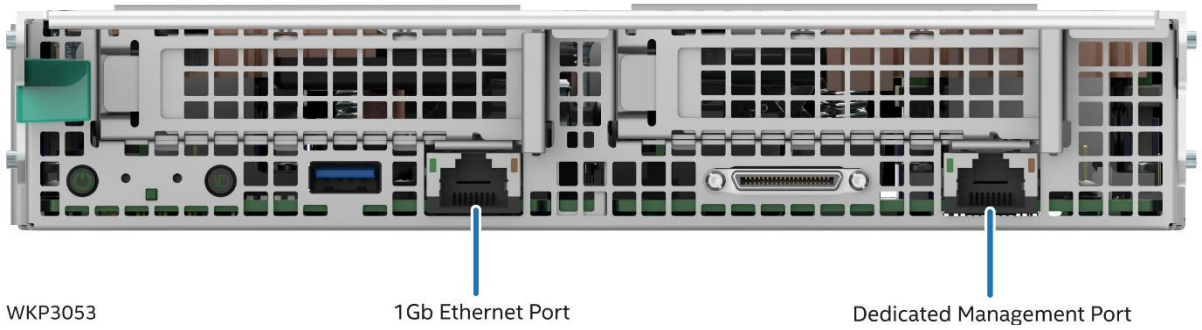

Figure 29. RJ45 connector identification for 1U compute modules

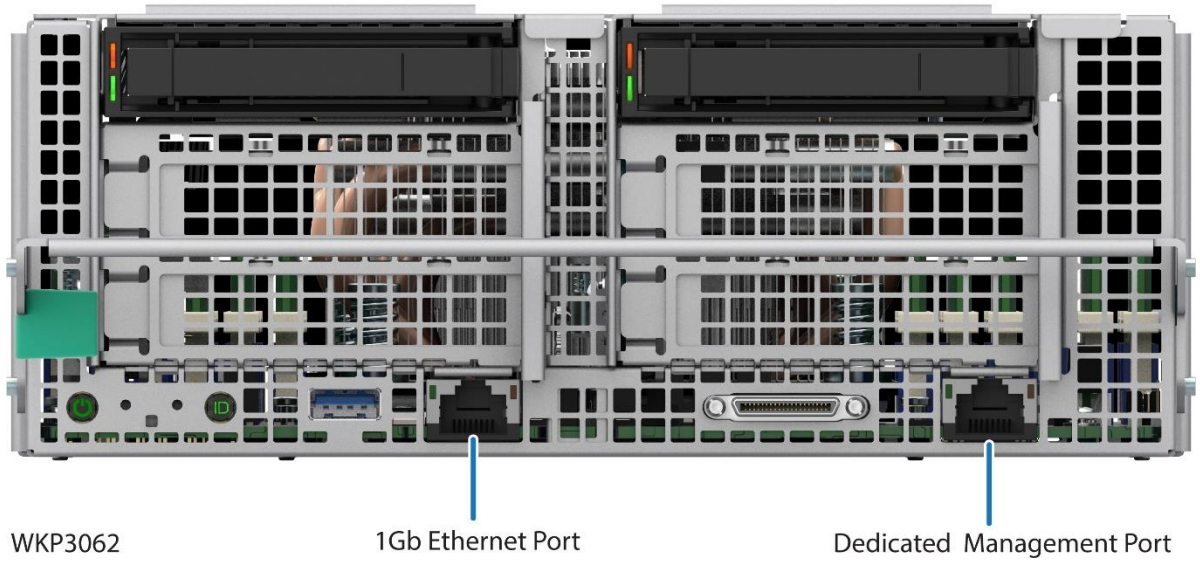

Figure 30. RJ45 connector identification for 2U compute modules

7.1.1 Network Interface

Network connectivity is provided by the onboard Intel® Ethernet Controller I210 providing 10/100/1000 Mb Ethernet support. The Intel® Ethernet Controller I210 is a single, compact, low-power component that offers a fully-integrated Gigabit Ethernet Media Access Control (MAC) and Physical Layer (PHY) port. The Intel® Ethernet Controller I210 uses the PCIe* architecture from the Intel® C620 series PCH and provides a single-port implementation.

Refer to the respective product data sheet for a complete list of supported features.

The RJ45 network interface connector includes two LEDs. The LED at the left of the connector is the link/activity LED and indicates a network connection when on and transmit/receive activity when blinking. The LED at the right of the connector indicates link speed as defined in Table 10.

Figure 31. RJ45 Network Interface Connector LEDs

Table 10. RJ45 Network Interface Connector port LED definition

LED	LED State	NIC State
Link/activity LED (left)	On	Active Connection
	Blinking	Transmit/Receive activity
Speed LED (right)	Off	10 Mbps
	Solid amber	100 Mbps
	Solid green	1000 Mbps

7.1.2 Dedicated Management port

Each compute module includes a dedicated 1 GbE RJ45 management port. The management port is active with or without the Intel® Remote Management Module 4 Lite (Intel® RMM4 Lite) key installed (optional). See Chapter 9 for additional information about server management support.

The dedicated management port includes two LEDs. The behavior of the LEDs is defined in Table 11**Error! Reference source not found.**.

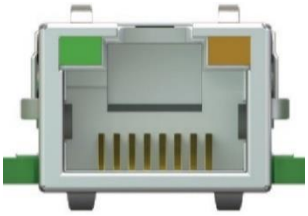

Figure 32. Dedicated management port LEDs

Table 11. Dedicated management port LED definition

Network Status	Left LED	Right LED
No link	Off	Off
10 Mbps link	On	On
10 Mbps activity	Blink	Blink
100 Mbps link	On	Off
100 Mbps activity	Blink	Off
1 Gpbs link	Off	On
1 Gbps activity	Off	Blink

7.1.3 MAC Address Definition

The compute modules within the Intel® Server System S9200WK product family have the following MAC addresses assigned at the factory:

- RJ45 Network Interface Connector (base MAC address)
- Dedicated management port (base MAC address + 1)

7.2 USB Support

The Intel® Server System S9200WK provides one USB 3.0 port on the front panel of each compute module and two USB 2.0 ports through a dedicated I/O breakout cable. One I/O breakout cable is included with each system. The following illustrations show the location of the ports located in the front panel.

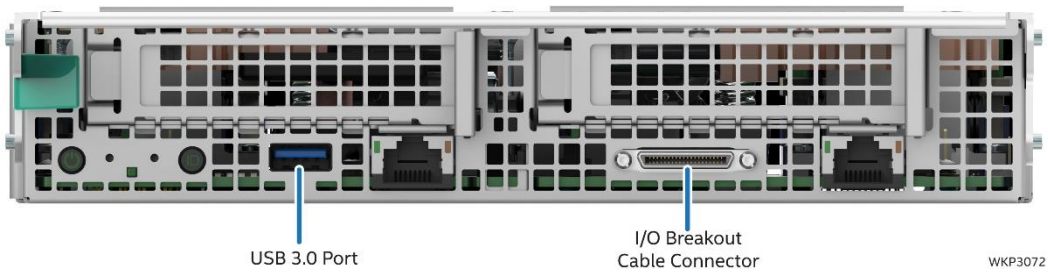

Figure 33. USB port location - 1U compute module

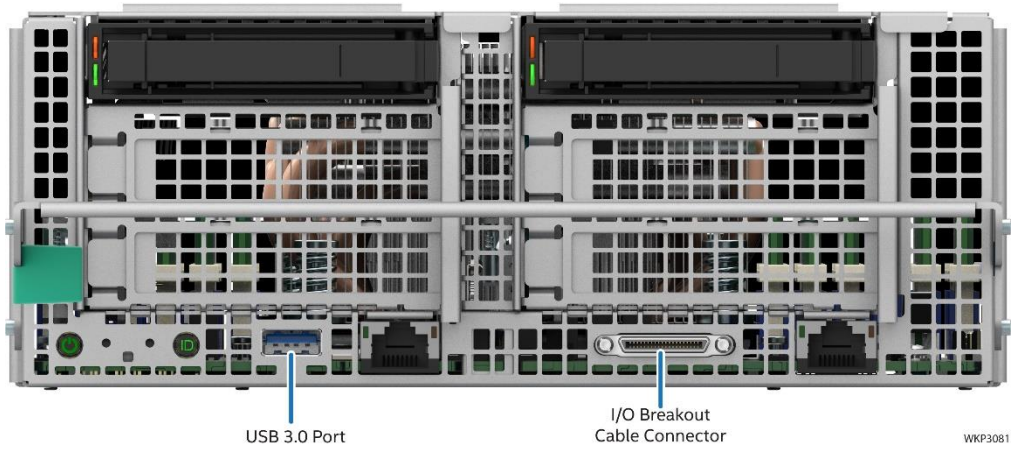

Figure 34. USB port location - 2U compute module

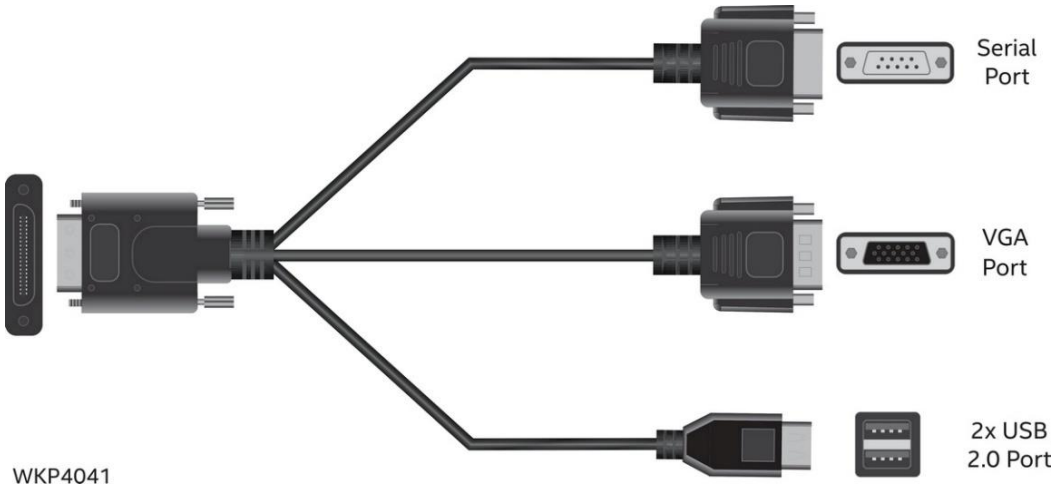

Figure 35. I/O breakout cable port identification

7.3 M.2 SSD Support

Both 1U and 2U compute modules within the Intel® Server System S9200WK family support up to two M.2 storage devices labeled “M.2x4_PCIE/M.2_SATA_6” on risers #1 and #2. Each M.2 connector supports PCIe* or SATA modules that conform to a 2280 (80 mm) or 22110 (110 mm) form factor.

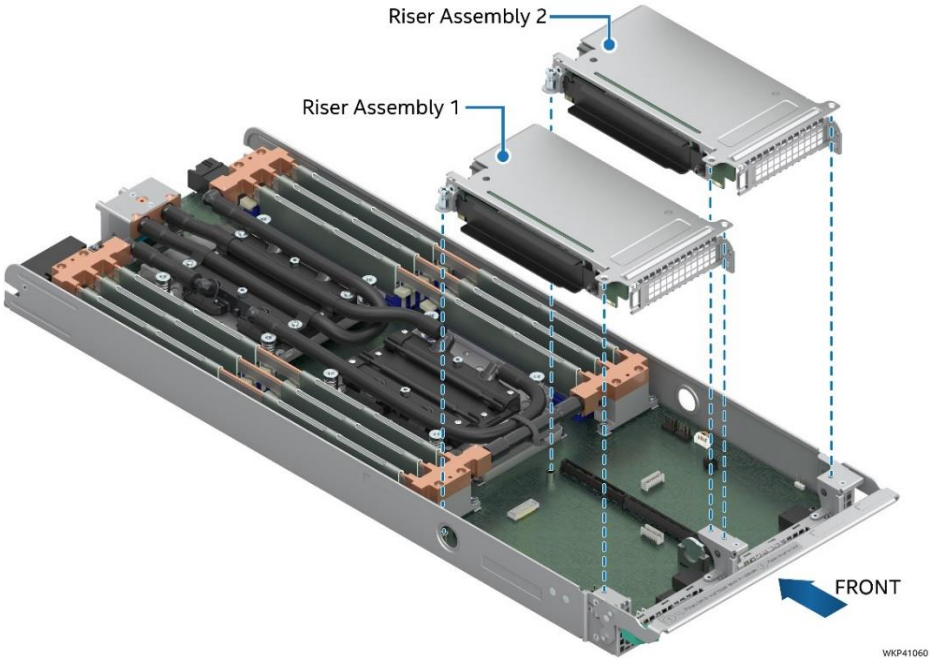

Figure 36. Riser location

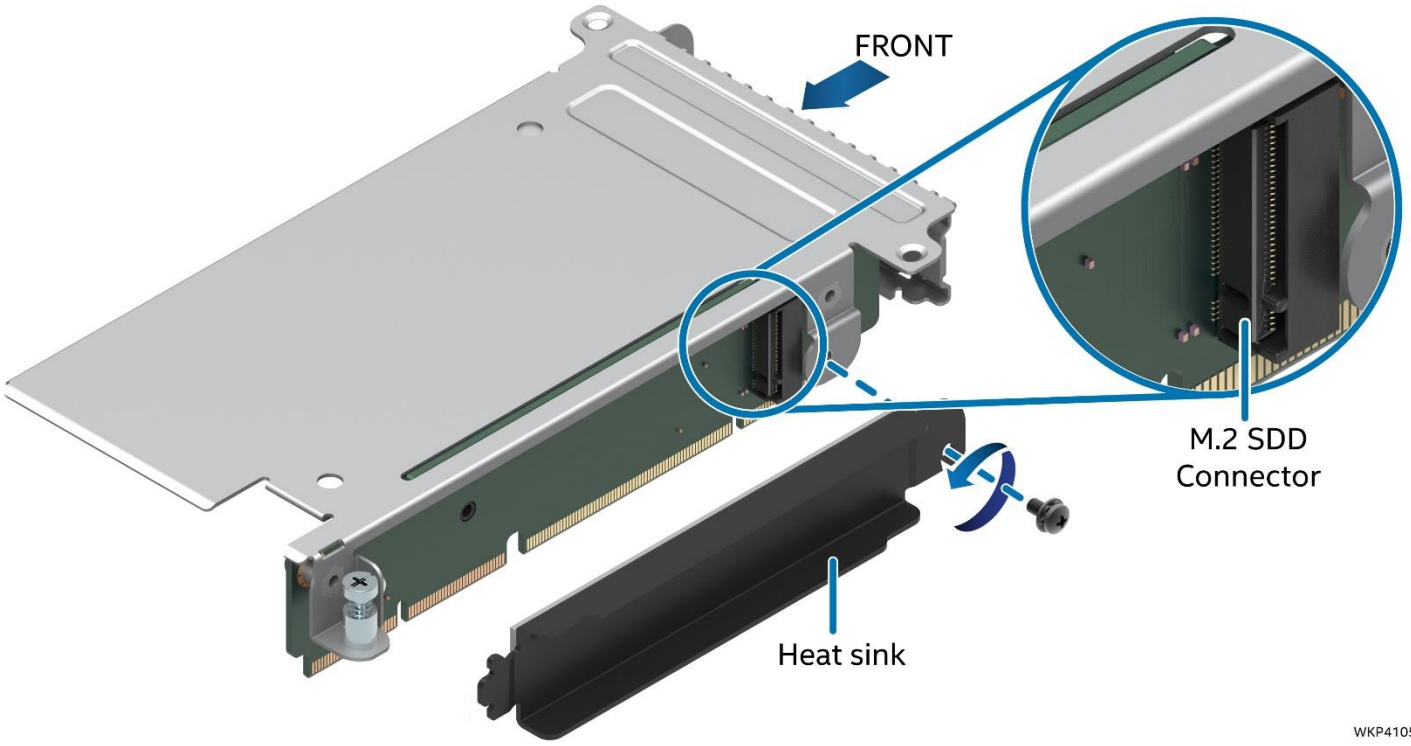

Figure 37. M.2 connector location on riser card

Note: Riser and connector location is the same for 1U and 2U compute modules.

The M.2 connector pinout definition is only made available by obtaining the board schematics directly from Intel (NDA required).

The M.2 connector on riser #1 is supported by PCIe* x4 bus lanes and SATA-0 from the chipset embedded SATA controller. The M.2 connector on riser #2 is supported by PCIe* x2 bus lanes and sSATA-2 from the chipset embedded sSATA controller. Intel® VROC (SATA RAID) is not supported for PCIe* or SATA M.2 SSDs.

Riser #1 and Riser #2 include a heat sink for M.2 SSDs. The included heat sink must be installed whether an M.2 SSD is present.

7.4 PCIe* NVMe* Drive Support

The 2U compute modules within the Intel® Server System S9200WK product family support up to two PCIe* NVMe U.2 storage drives through the hot-swap drive bays located in the upper part of the front of the compute module. Intel® VROC is supported for U.2 drives only. U.2 VROC RAID boot on volume is unsupported.

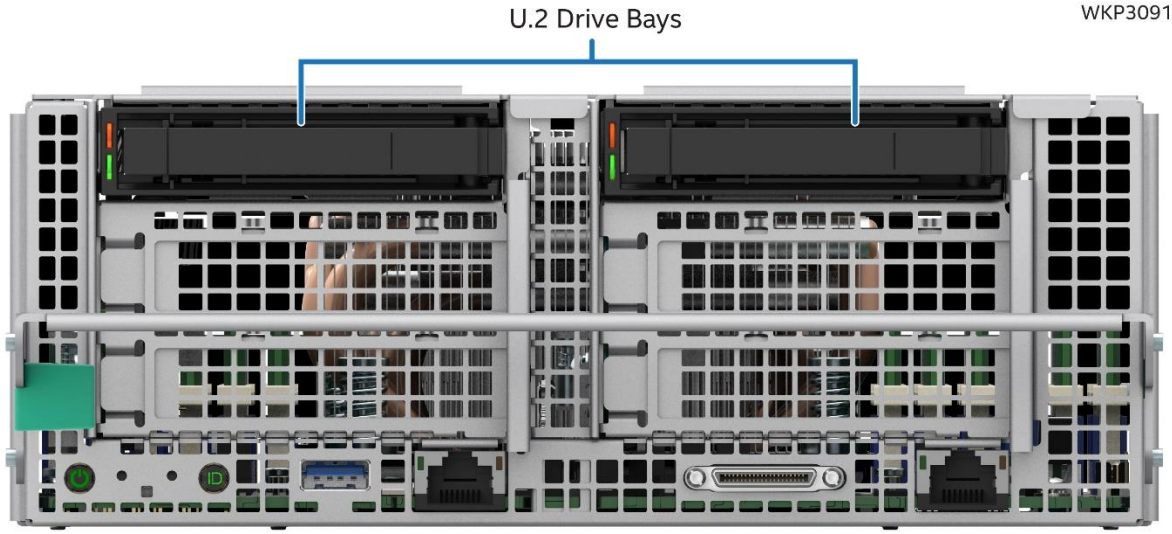

Figure 38. Locating the U.2 drive bays

Each installed NVMe drive is connected to the PCIe* bus through a dedicated U.2 connector attached to the riser assemblies in the compute module.

7.5 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.

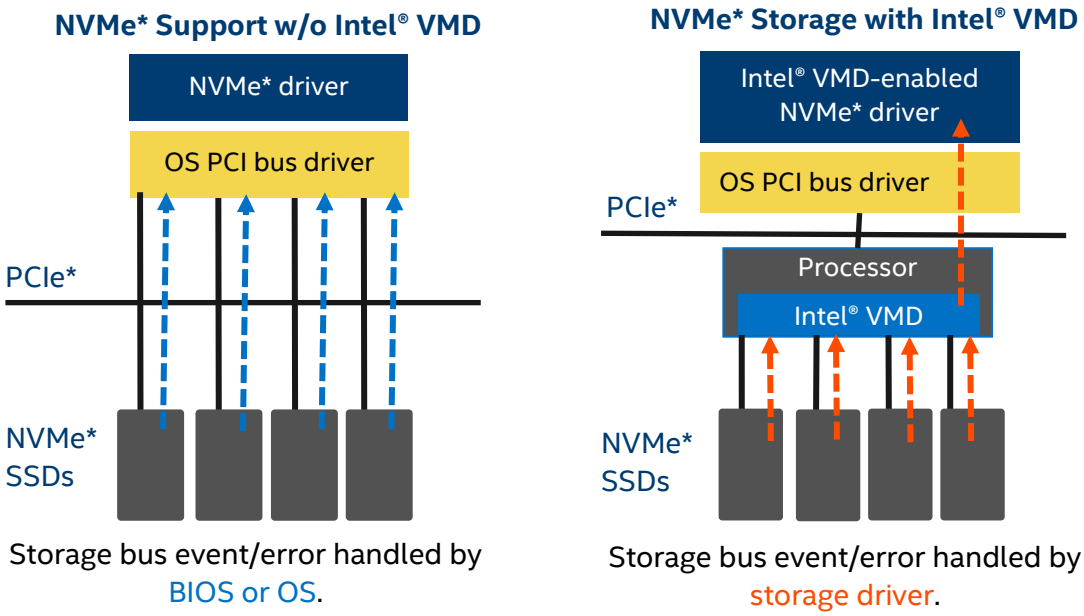

Figure 39. NVMe* storage bus event/error handling

Intel® VMD handles the physical management of NVMe storage devices as a standalone function but can be enhanced when Intel® VROC support options are enabled to implement RAID based storage systems. See Section 7.6 for more information.

The compute modules within the Intel® Server System S9200WK product family support the following VROC features and capabilities:

- 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 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.

7.5.1 Enabling Intel® VMD support

For installed NVMe devices to utilize the Intel® VMD features supported by the compute modules in the system, Intel® VMD must be enabled on the appropriate CPU PCIe* root ports in BIOS setup. By default, Intel® VMD support is disabled on all CPU PCIe* root ports in BIOS setup.

Table 9 provides the PCIe* root port mapping for all on-board PCIe* devices and riser card slots.

In <F2> BIOS setup, the Intel® VMD support menu can be found on the following menu tab: **Advanced** > **PCI Configuration** > **Volume Management Device.**

7.6 Intel® Virtual RAID on Chip (Intel® VROC) For NVMe*

Intel® VROC (VMD NVMe RAID) enables NVMe RAID and volume management..

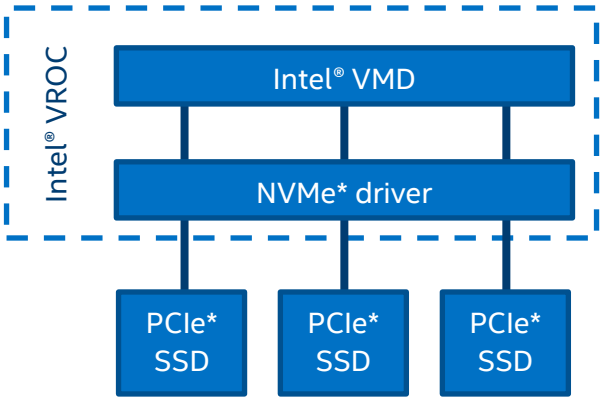

Figure 40. Intel® VROC basic architecture overview

The compute modules within the Intel® Server System S9200WK product family support the following Intel® VROC features:

- I/O processor with controller (ROC) and DRAM.
- 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.
- NVMe SSD hot plug and surprise removal on CPU PCIe* lanes.
- LED management for CPU PCIe* attached storage.
- RAID / storage management using representational state transfer (RESTful) application programming interfaces (APIs).
- Graphical user interface (GUI) for Linux.
- 4K native NVMe SSD support.
- U.2 RAID volume support
- U.2 Boot on VROC RAID volume unsupported
- M.2 VROC RAID unsupported

Enabling Intel® VROC support requires the installation of an optional upgrade key (iPC – VROCSTANMOD) in the compute module as shown in Figure 41. Table 12 identifies the supported RAID features by the optional upgrade key.

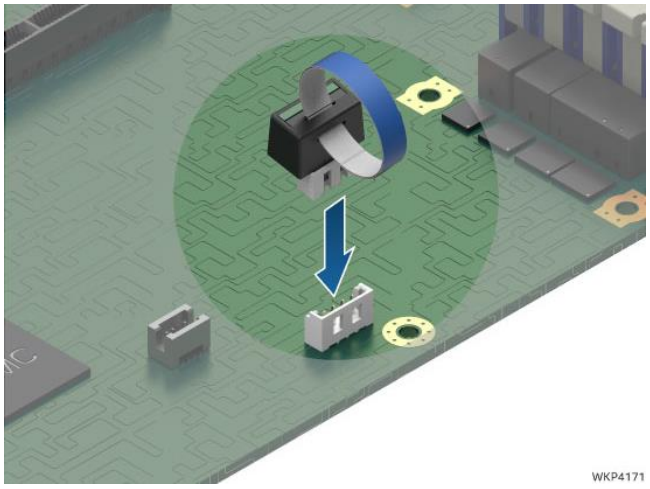

Figure 41. Intel® VROC upgrade key

Table 12. Optional VROC upgrade key - Supported NVMe* RAID features

NVMe* RAID Major Features	Support in the Intel® Server System S9200WK
CPU attached NVMe SSD – high perf.	√
Boot on RAID volume	Unsupported
Third party vendor SSD support	√
RAID 0/1	√
RAID 10	Unsupported
RAID 5	Unsupported
RAID write hole closed (RMFBU replacement)	Unsupported
Hot plug/ surprise removal (2.5” SSD form factor only)	√
Enclosure LED management	√

Note: Intel® VROC upgrade keys referenced in Table 12 are used for PCIe* NVMe SSDs only.

7.7 Onboard SATA Support

The compute modules utilize two chipset embedded AHCI SATA controllers, identified as “SATA” and “sSATA”, providing for up to two 6 Gb/sec SATA ports.

The AHCI SATA controller provides support for one SATA port accessed through an M.2 connector located on riser #1. The AHCI sSATA controller provides support for one SATA port accessed through an M.2 connector located on riser #2. Table 13 lists the supported features by the embedded SATA and sSATA controllers.

Table 13. SATA and sSATA controller feature support

Feature	Description	AHCI Mode
Native Command Queuing (NCQ)	Allows the device to reorder commands for more efficient data transfers	Supported
Auto Activate for DMA	Collapses a DMA Setup then DMA Activate sequence into a DMA Setup only	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

Feature	Description	AHCI Mode
Asynchronous Signal Recovery	Provides a recovery from a loss of signal or establishing communication after hot plug	Supported
6 Gb/s Transfer Rate	Capable of data transfers up to 6 Gb/s	Supported
ATAPI Asynchronous Notification	A mechanism for a device to send a notification to the host that the device requires attention	Supported
Host & Link Initiated Power Management	Capability for the host controller or device to request Partial and Slumber interface power states	Supported
Staggered Spin-Up	Enables the host the ability to spin up hard drives sequentially to prevent power load problems on boot	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

The SATA controller and the sSATA controller can be independently enabled and disabled and configured through the BIOS setup utility under the Mass Storage Controller Configuration menu screen. Table 14 identifies the supported setup options.

Table 14. SATA and sSATA controller BIOS setup utility options

SATA Controller	sSATA Controller	Supported
AHCI	AHCI	Yes
AHCI	Disabled	Yes
Disabled	AHCI	Yes
Disabled	Disabled	Yes

7.8 PCIe* Add-in Card Support

The compute modules within the Intel® Server System S9200WK product family support low profile PCIe* add-in cards. The 1U and 2U compute modules support a different number of add-in cards:

- 1U Compute Module
  - Supports up to Two PCIe* 3.0 x16 mechanical and electrical Add-in Cards
- 2U Compute Module
  - Supports up to Four PCIe* 3.0 x16 mechanical and electrical Add-in Cards

PCIe* add-in cards are supported through two riser cards installed at the front of the compute modules. Each riser card is part of a riser assembly that can be independently removed or replaced. The following illustrations show the location of the riser assemblies and add-in card bays.

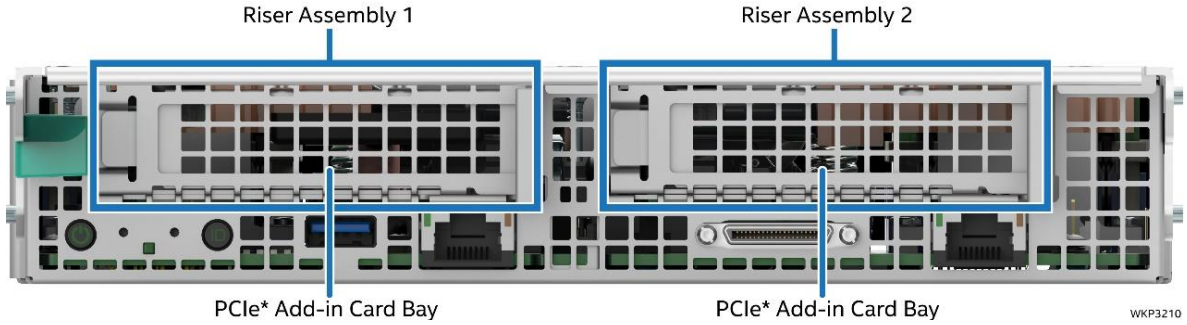

Figure 42. Riser assembly and add-in card bay identification – 1U compute modules

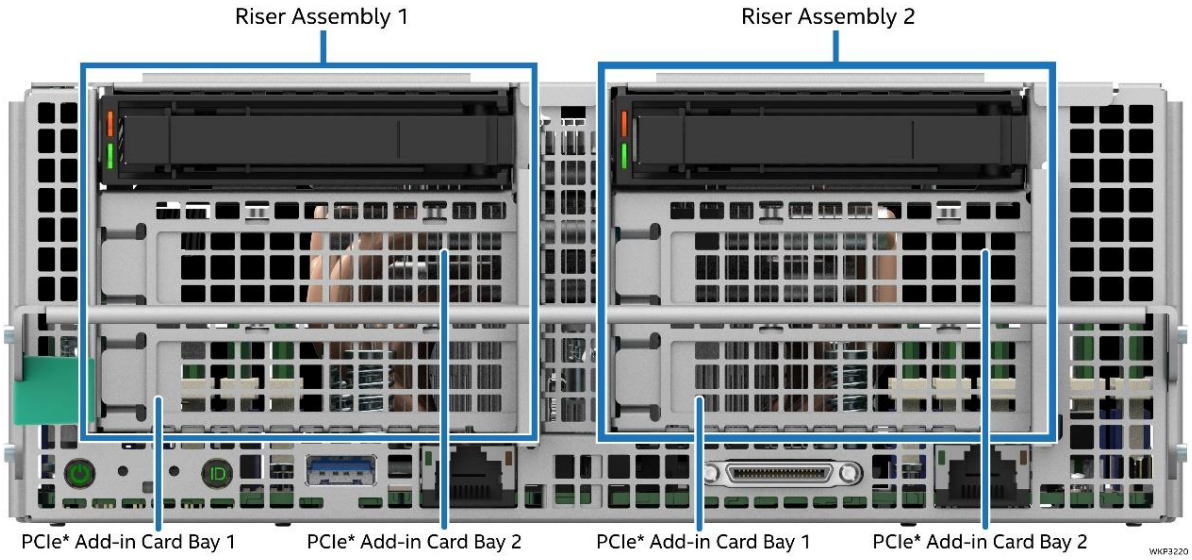

Figure 43. Riser assembly and add-in card bay identification – 2U compute modules

Note: The hot-swap drive bays are part of the 2U riser assembly.

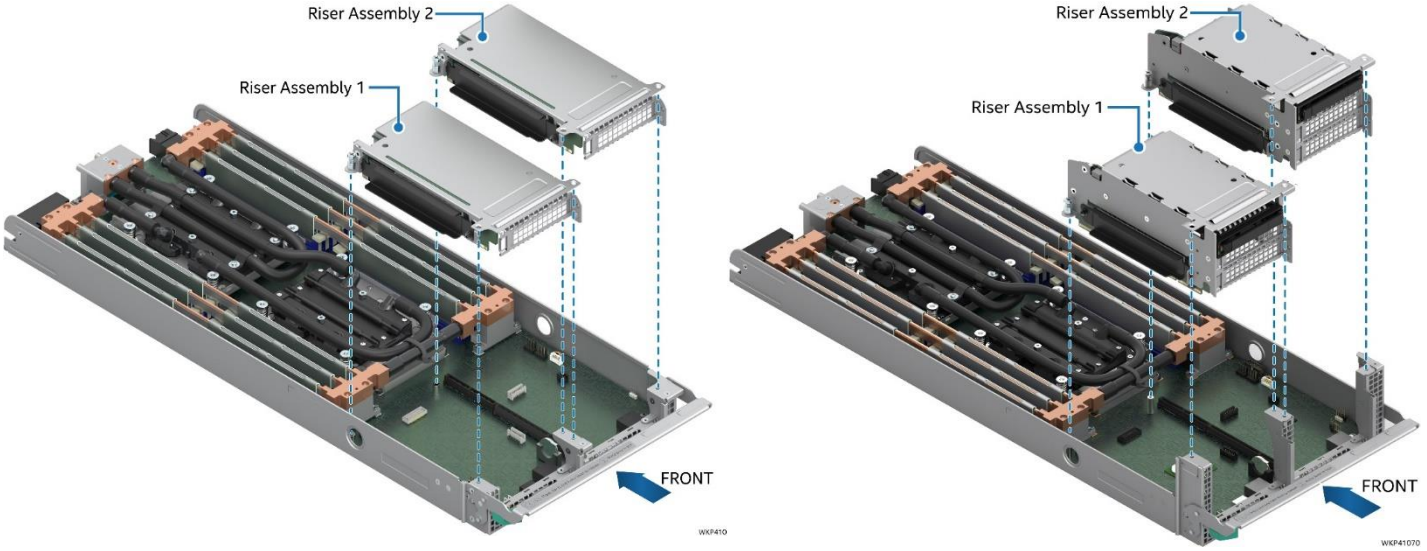

1U Compute module **2U Compute module**

Figure 44. Riser assembly location – 1U and 2U compute modules

The following illustration identifies the location of the slots in which the riser cards are installed in the compute modules, as well as their link width.

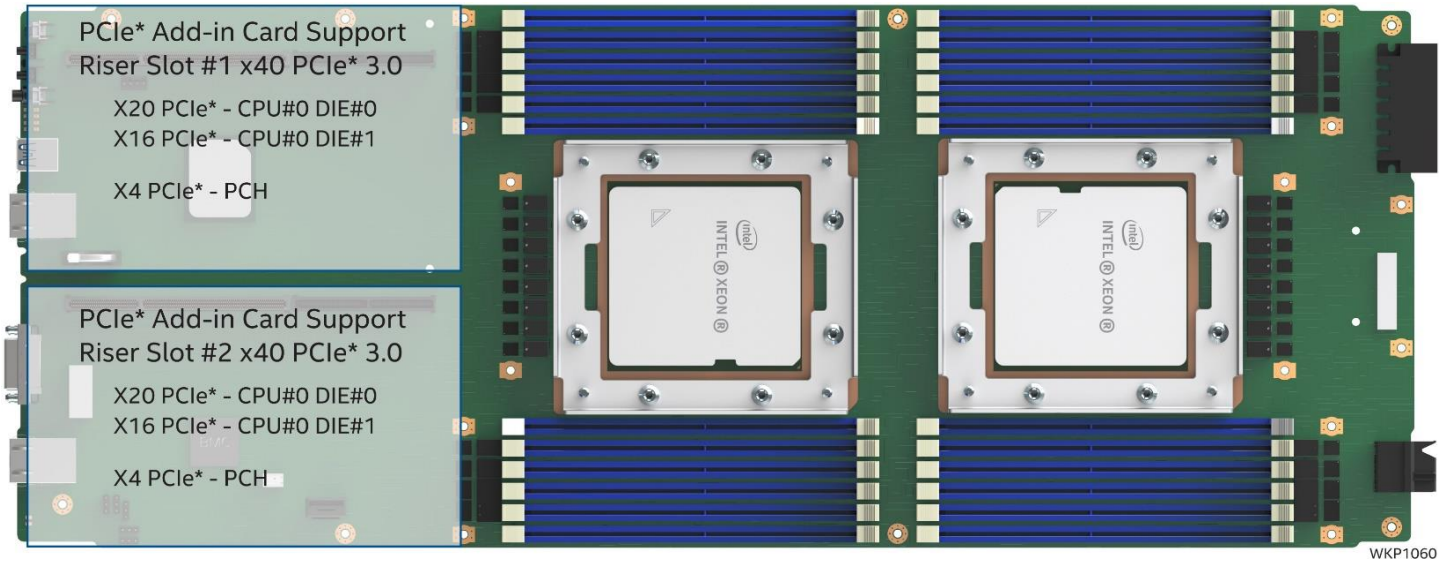

Figure 45. Riser card slot location – 1U and 2U compute modules

Table 15 and Table 16 provide the PCIe* bus routing for the supported riser cards.

Table 15. Riser slot #1 root port mapping

Usage	1U riser card	2U riser card
Bottom PCIe* slot	CPU 0, DIE 1: Ports 2A, 2B, 2C and 2D (x16 elec, x16 mech)	CPU 0, DIE 1: Ports 2A, 2B, 2C and 2D (x16 elec, x16 mech)
Top PCIe* slot	N/A	CPU 0, DIE 0: Ports 4A, 4B, 4C and 4D (x16 elec, x16 mech)
U.2 connector	N/A	CPU 0, DIE 0: Port 3A (x4 elec)

Table 16. Riser slot #2 root port mapping

Usage	1U riser card	2U riser card
Bottom PCIe* slot	CPU 0, DIE 0: Ports 1A, 1B, 1C and 1D (x16 elec, x16 mech)	CPU 0, DIE 0: Ports 1A, 1B, 1C and 1D (x16 elec, x16 mech)
Top PCIe* slot	N/A	CPU 0, DIE 1: Ports 1A, 1B, 1C and 1D (x16 elec, x16 mech)
U.2 connector	N/A	CPU 0, DIE 0: Port 2A (x4 elec)

Note: The PCIe* slots in 1U and 2U riser cards provide up to 25W of power.

7.8.1 Riser Card and Add-in Card Support for 1U Compute Modules

The 1U compute modules include a riser card within each riser assembly. Each riser card supports one add-in card. The following figures show the riser assembly and the riser card features.

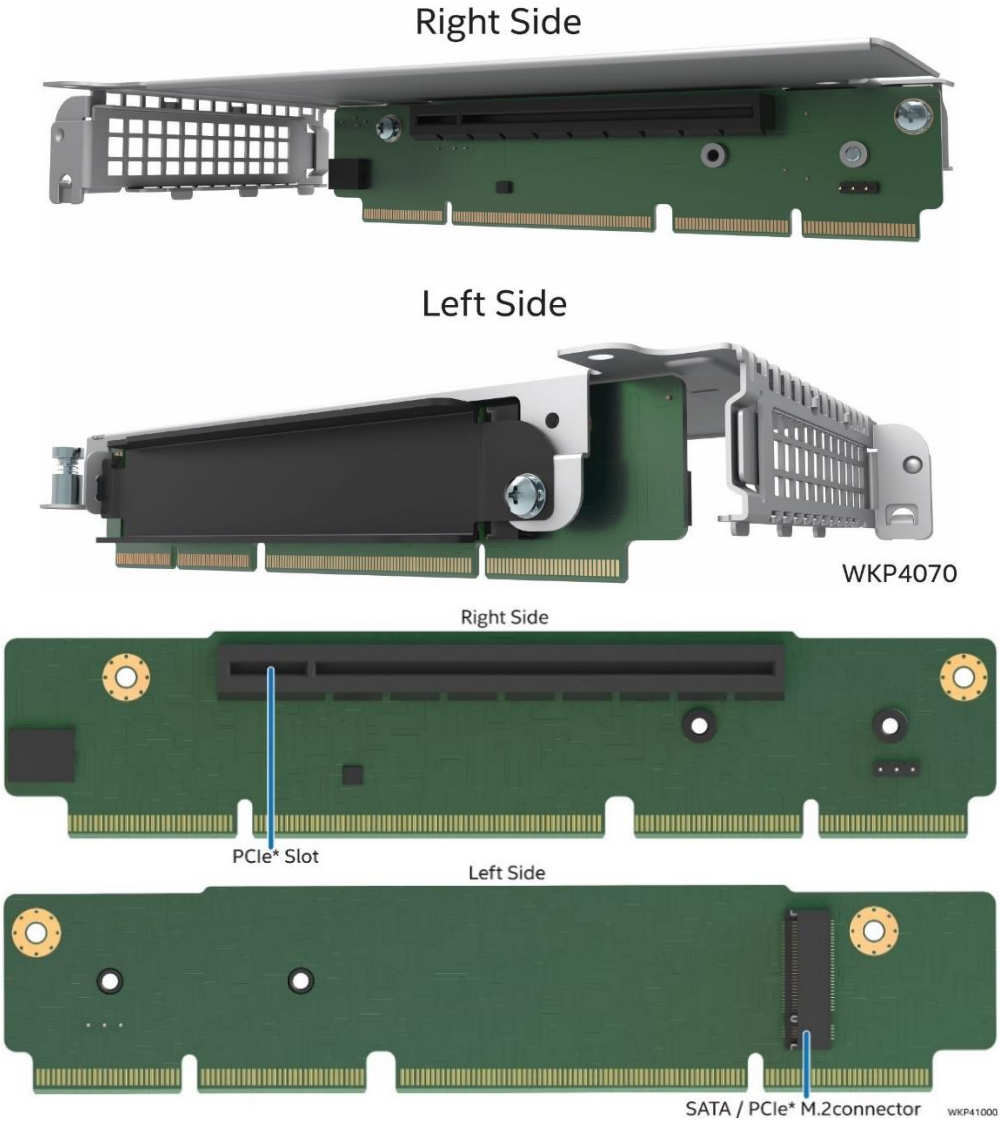

Figure 46. 1U riser assembly and riser card features

The features of the 1U riser card are listed below:

1- PCIe* slot compatible with low profile add-in cards (x16 elec, x16 mech)

1- M.2 connector compatible with 80mm and 100mm SATA and PCIe* and storage devices

7.8.2 Riser Card and Add-in Card Support for 2U Compute Modules

The 2U compute modules include a riser card within each riser assembly. Each riser card supports two add-in cards. The following figures show the riser assembly and the riser card features.

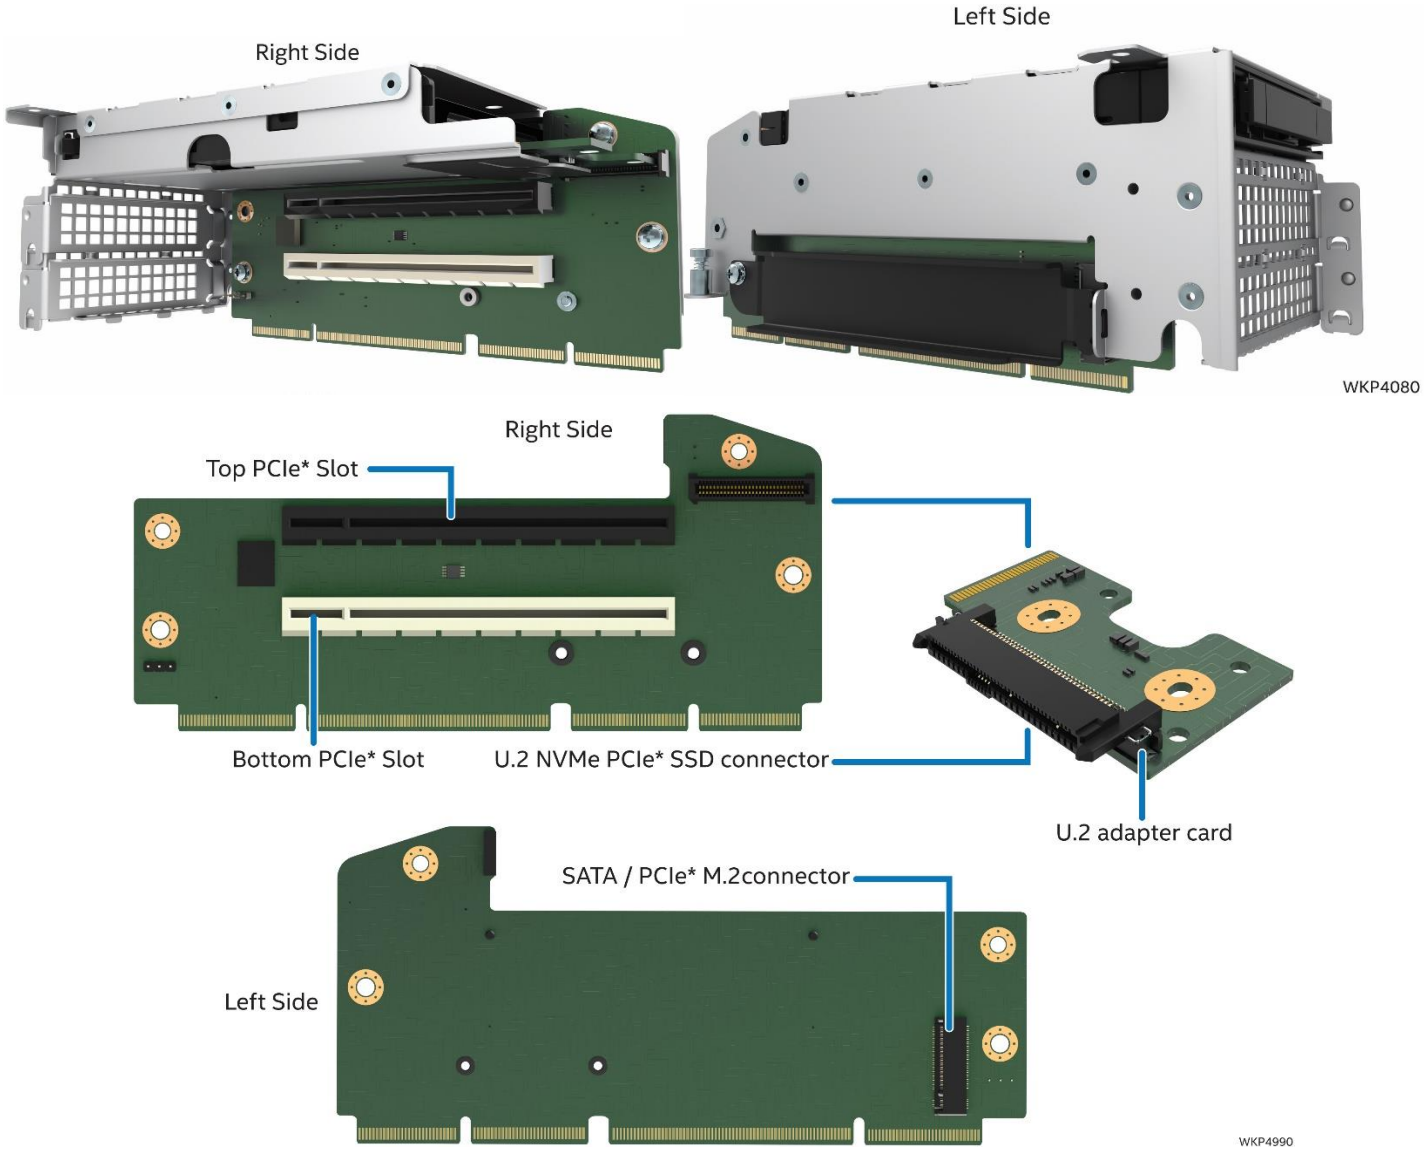

Figure 47. 2U riser assembly and riser card features

The features of the 2U riser card are listed below:

2- PCIe* slot compatible with low profile add-in cards (x16 elec, x16 mech)

1- M.2 connector compatible with 80mm and 100mm SATA and PCIe* and storage devices

1- U.2 connector for PCIe* NVMe SSD storage devices through a dedicated adapter card.

8. Control Panel Features

The compute modules include a control panel that provides push button controls and LED indicators for several features. This section provides a description for each front control panel feature.

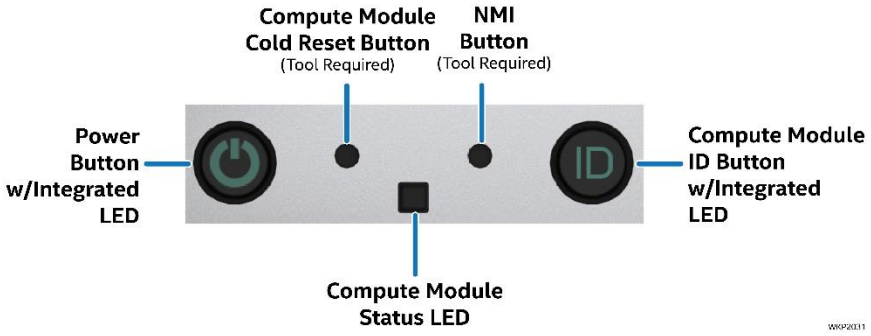

Figure 48. Front control panel features

- **Compute module ID button w/ integrated LED** – Toggles the integrated blue ID LED on and off. The Compute module ID LED is used to identify an individual compute module within a chassis for maintenance when installed in a rack of similar server systems. The Compute Module ID LED can also be toggled on and off remotely using the IPMI "Chassis Identify" command which causes the LED to blink for 15 seconds.
- **NMI Button** – When the NMI button is pressed, it puts the compute module in a halt state and issues a non-maskable interrupt (NMI). This can be useful when performing diagnostics for a given issue where a memory download is necessary to help determine the cause of the problem. To prevent an inadvertent compute module 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.
- **Compute module cold reset button** – When pressed, this button reboots and re-initializes the compute module.
- **Compute module status LED** – The compute module status LED is a bi-color (green/amber) indicator that shows the current health of the compute module. The compute module status LED states are driven by the integrated platform management subsystem. Table 17 provides a description of each supported LED state.

Table 17. Compute module status LED state definitions

LED State	Compute module State	BIOS Status Description
Off	Compute module is not operating.	• Compute module AC power is off.
Solid green	Compute module is operating normally.	• Compute module is running (in S0 State) and its status is healthy. The compute module 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 compute module ID LED 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. • Compute module is in S5 soft-off state. • Compute module is in EuP Lot6 off mode.
Blinking green	Compute module is operating in a degraded state although still functioning, or compute module 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 compute module no longer has spared DIMMs (a redundancy lost condition). • In mirrored configuration, when memory mirroring takes place and the compute module loses memory redundancy. • Battery failure. • BMC executing in uBoot. (Indicated by compute module ID LED blinking at 3Hz). Compute module in degraded state (no manageability). BMC uBoot is running but has not transferred control to BMC Linux*. The compute module 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	Compute module is operating in a degraded state with an impending failure warning, although still functioning. Compute module 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. • Storage 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.

LED State	Compute module State	BIOS Status Description

Solid amber	Critical/non-recoverable – compute module is halted. Fatal alarm – compute module has failed or shut down.	• CPU CATERR signal asserted.
		• MSID mismatch detected (CATERR also asserts for this case).
		• CPU 0 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 (compute module ID LED shows blue/solid-on for this condition).
		• Both uBoot BMC firmware images are bad (compute module ID LED shows blue/solid-on for this condition).
		• 240 VA fault.
		• Fatal Error in processor initialization:
		  ○ Processor family not identical
		  ○ Processor model not identical
		  ○ Processor core/thread counts not identical
		  ○ Processor cache size not identical
		  ○ Unable to synchronize processor frequency
		  ○ Unable to synchronize UPI link frequency
		• Uncorrectable memory error in a non-redundant mode.

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

Table 18. Power/sleep LED functional states

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

9. Basic And Advanced Management Features

The integrated Baseboard Management Controller (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 19. Intel® Remote Management Module 4 (Intel® RMM4) options

Intel Product Code (iPC)	Description	Kit Contents	Benefits
AXXRMM4LITE2	Intel® Remote Management Module 4 Lite	Intel® RMM4 Lite Activation Key	Enables remote keyboard, video, and mouse (KVM) and media redirection.

On the server board within the compute module, the Intel® RMM4 Lite key is installed at the location shown in Figure 49.

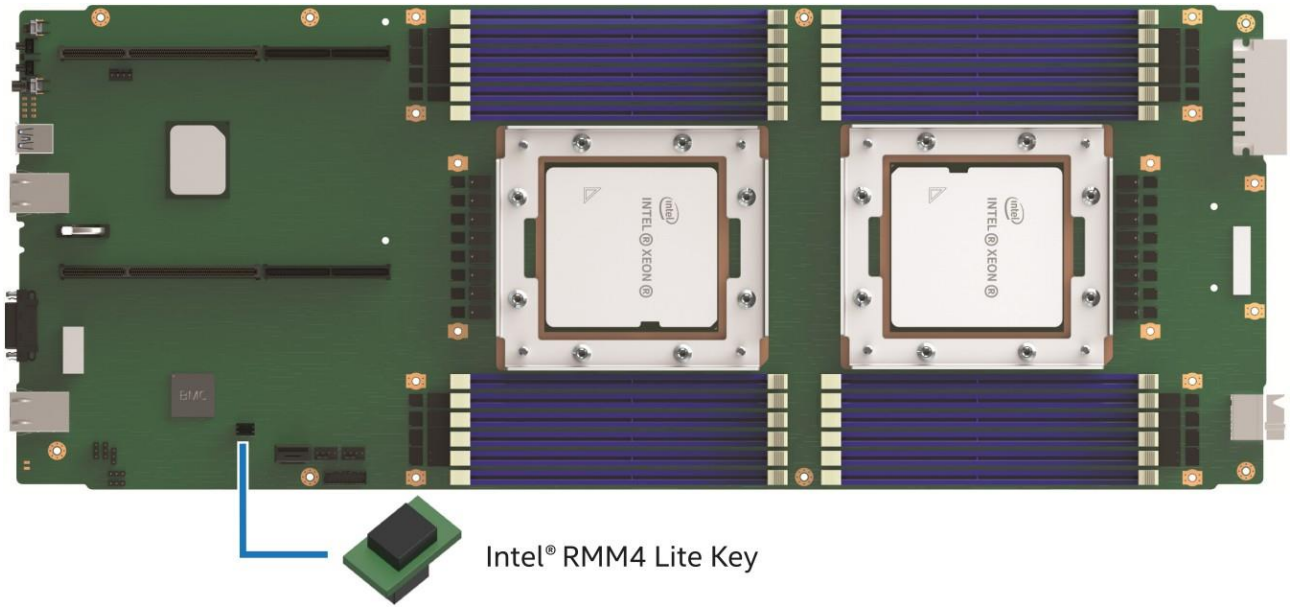

Figure 49. Intel® RMM4 Lite activation key connector location

Basic and Advanced management features are available after AC power is applied to the system. With AC power present, stand-by power within the compute module is enabled, applying power to the integrated BMC which then enables all available management features. Basic and Advanced management features are enabled whether the compute module is powered on or not.

Table 20 identifies both basic and advanced server management features.

Table 20. Basic and advanced server management features overview

Feature	Basic	Advanced w/ Intel® 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

Feature	Basic	Advanced w/ Intel® RMM4 Lite Key
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

9.1 Dedicated Management Port

The compute modules include a dedicated 1GbE RJ45 management port. The management port is active with or without the Intel® RMM4 Lite key installed.

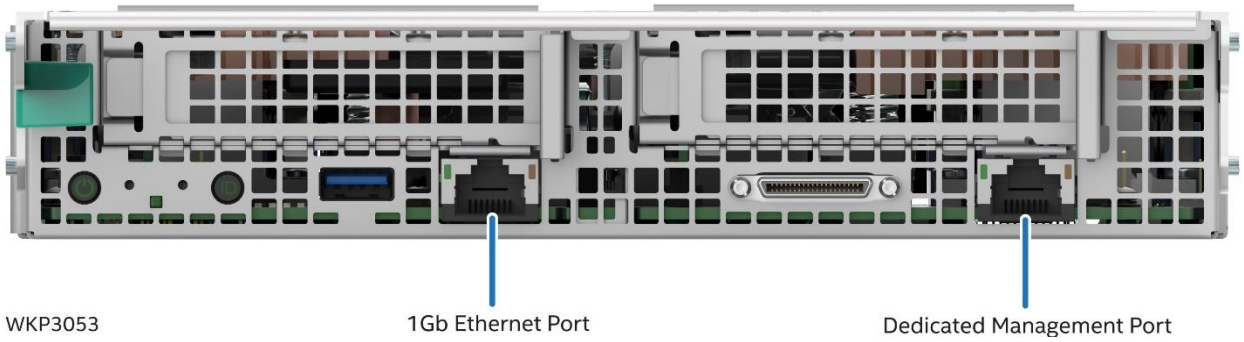

Figure 50. Dedicated management port location – 1U compute modules

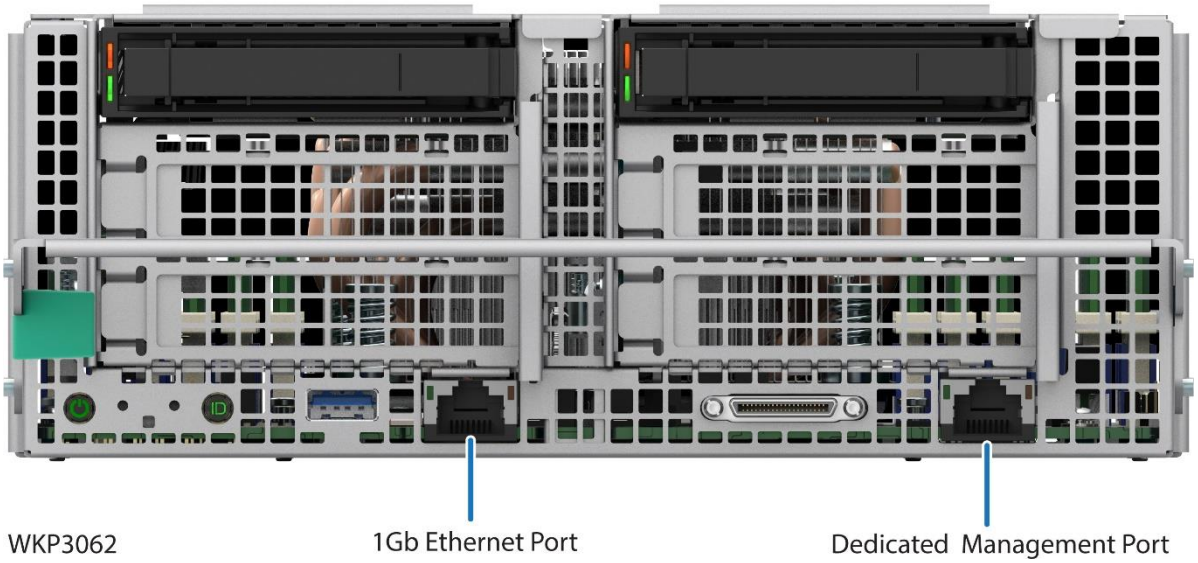

Figure 51. Dedicated management port location – 2U compute modules

9.2 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. 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. Encryption using 128-bit SSL is supported. User authentication is based on user ID and password.

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

Additionally, the web interface can:

- Present all the basic features to the users.
- Power on, power off, and reset the compute module 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 sensor data with a configurable refresh rate.
- Provide online help
- Display/clear SEL (display is in easily understandable human readable format).
- Support major industry-standard browsers.
- Automatically time out GUI session after a user-configurable inactivity period. By default, this inactivity period is 30 minutes.
- Provide embedded platform debug feature, allowing the user to initiate a "debug dump" to a file that can be sent to Intel for debug purposes.
- Provide a virtual front panel with 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 Intel® ME sensor data. Only sensors that have associated SDRs loaded are displayed.
- Save the SEL to a file.
- Force HTTPS connectivity for greater security. This is provided through a configuration option in the user interface.
- Display processor and memory information that is available over IPMI over LAN.

- Get and set Intel® Node Manager (Intel® NM) power policies
- Display the power consumed by the compute module.
- View and configure VLAN settings.
- Warn user that the reconfiguration of IP address causes 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.
- Force into BIOS setup on a reset (server power control).
- Provide the compute module 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.
- Provide 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.

9.3 Advanced Management Feature Support

The integrated baseboard management controller has support for advanced management features which are enabled when an optional Intel® RMM4 Lite is installed. The Intel® RMM4 Lite add-on 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 Lite hardware.

- Key features of the Intel® RMM4 Lite add-on include:
- **KVM redirection** from either the dedicated management NIC or the compute module standard NIC used for management traffic and up to two KVM sessions. KVM 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** 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 compute module. Once mounted, the remote device appears to the server just like a local device, allowing system administrators or users to install software (including operating systems), copy files, update BIOS, or boot the compute module from this device.

9.3.1 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, or an HTML 5 compatible web browser 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 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 for user definable keyboard macros.

KVM redirection feature supports the following resolutions and refresh rates:

- 640x480 at 60 Hz, 72 Hz, 75 Hz, 85 Hz
- 800x600 at 60 Hz, 72 Hz, 75 Hz, 85 Hz
- 1024x768 at 60 Hz, 72 Hz, 75 Hz, 85 Hz
- 1152x864 at 75 Hz
- 1280x800 at 60 Hz
- 1280x1024 at 60 Hz
- 1440x900 at 60 Hz
- 1600x1200 at 60 Hz

9.3.1.1 Availability

The remote KVM session is available even when the compute module is powered off (in stand-by mode). No restart of the remote KVM session is required during a reset or power on/off. A BMC reset – for example, due to a BMC watchdog initiated reset or BMC reset after BMC firmware update – does require the session to be re-established.

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

9.3.1.2 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.

9.3.1.3 Usage

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

9.3.1.4 Force-enter BIOS Setup

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

9.3.2 Media Redirection

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

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

The following list describes additional media redirection capabilities and features.

- The operation of remotely mounted devices is independent of the local devices on the compute module. 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 compute module.
- 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 compute module is powered off (in standby mode). No restart of the remote media session is required during a reset or power on/off. A BMC reset (for example, due to an BMC reset after BMC FW update) requires 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 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 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 BIOS.

9.3.2.1 Availability

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

9.3.3 Remote Console

The remote console is the redirected screen, keyboard, and mouse of the remote host compute module. To use the remote console window of the managed host compute module, the browser must include a Java* Runtime Environment (JRE) 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 CD-ROM media redirection, and #5123 for floppy and USB media redirection. When encryption is enabled, the protocol uses ports #7582 for KVM, #5124 for CD-ROM media redirection, and #5127 for floppy and 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 Network Address Translation (NAT) settings have to be configured accordingly.

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

9.3.4 Performance

The remote display accurately represents the local display. The feature adapts to changes in 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 does degrade performance. Enabling video compression provides the fastest response while disabling compression provides better video quality. For the best possible KVM performance, a 2 Mbps 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.

9.4 I/O Breakout Cable

Present within the front panel of each compute module is a 40-pin connector providing support for an I/O breakout cable as shown in Figure 49. Each system configuration includes one I/O breakout cable providing a single compute module with additional support for one serial port, two USB 2.0 ports, and one VGA port as needed.

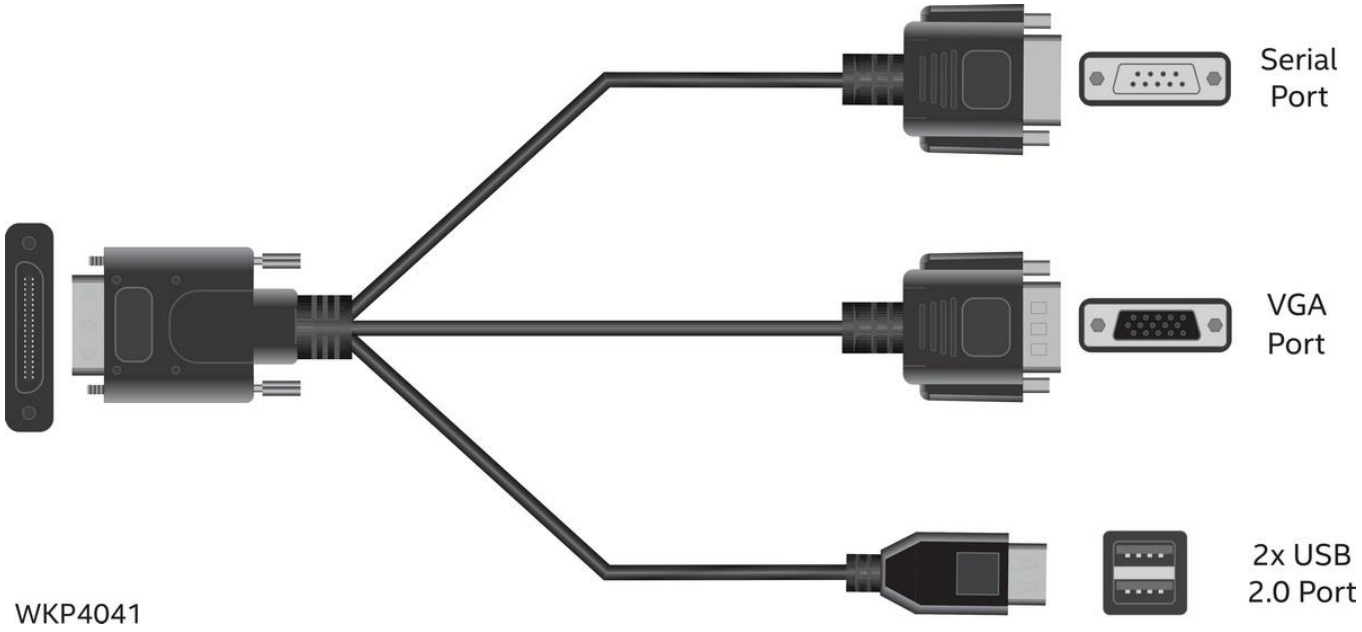

Figure 52. I/O breakout cable port identification

10. Platform Management

Platform management is supported by several hardware and software components integrated on the server system that work together to:

- Control system functions – power system, ACPI, system reset control, system initialization, front panel interface, system event log.
- Monitor various compute module and system sensors and regulate platform thermals and performance to maintain (when possible) server functionality in the event of component failure and/or environmentally stressed conditions.
- Monitor and report system health.
- Provide an interface for Intel® Server Management software applications.

This chapter provides a high level overview of the platform management features and functionality implemented on the compute modules in the system.

The *Intel® Server System BMC Firmware External Product Specification* (EPS) and the *Intel® Server System BIOS External Product Specification* (EPS) for Intel® Server Products based on the Intel® Xeon® processor Scalable product families should be referenced for more in-depth and design level platform management information.

10.1 Management Feature Set Overview

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

10.1.1 IPMI 2.0 Features Overview

The baseboard management controller (BMC) supports the following IPMI 2.0 features:

- IPMI watchdog timer.
- Messaging support, including command bridging and user/session support.
- Chassis device functionality, including power/reset control and BIOS boot flags support.
- Event receiver device to receive and process events from other platform subsystems.
- Access to system Field Replaceable Unit (FRU) devices using IPMI FRU commands.
- System Event Log (SEL) device functionality including SEL Severity Tracking and Extended SEL.
- Storage of and access to system Sensor Data Records (SDRs).
- Sensor device management and polling to monitor and report system health.
- IPMI interfaces
    - Host interfaces including system management software (SMS) with receive message queue support and server management mode (SMM)
    - Intelligent platform management bus (IPMB) interface
    - LAN interface that supports the IPMI-over-LAN protocol (RMCP, RMCP+)
- Serial-over-LAN (SOL)
- ACPI state synchronization to state changes provided by the BIOS.
- Initialization and runtime self-tests including making results available to external entities.

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

10.1.2 Non-IPMI Features Overview

The BMC supports the following non-IPMI features.

- In-circuit BMC firmware update.

- Fault resilient booting (FRB) including FRB-2 supported by the watchdog timer functionality.
- Chassis intrusion detection (dependent on platform support).
- Fan speed control with SDR, fan redundancy monitoring, and support.
- Enhancements to fan speed control.
- Power supply redundancy monitoring and support.
- Hot-swap fan support.
- Acoustic management and support for multiple fan profiles.
- Test commands for setting and getting platform signal states.
- Diagnostic beep codes for fault conditions (dependent on platform support).
- System globally unique identifier (GUID) storage and retrieval.
- Front panel management including system status LED and system ID LED (turned on using a front panel button or command), secure lockout of certain front panel functionality, and button press monitoring.
- Power state retention.
- Power fault analysis.
- Intel® Light-Guided Diagnostics.
- Power unit management including support for power unit sensor and handling of power-good dropout conditions.
- DIMM temperature monitoring facilitating new sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Sending and responding to Address Resolution Protocols (ARPs) (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP) (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support.
- Email alerting.
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server.
    - Human-readable SEL.
    - Additional system configurability.
    - Additional system monitoring capability.
    - Enhanced online help.
- Integrated keyboard, video, and mouse (KVM).
- Enhancements to KVM redirection.
    - Support for higher resolution.
- Integrated remote media redirection.
- Lightweight Directory Access Protocol (LDAP) support.
- Intel® Intelligent Power Node Manager support.
- Embedded platform debug feature which allows capture of detailed data for later analysis.
    - Creation of password protected files accessible by Intel only.
- Provisioning and inventory enhancements.
    - Inventory data/system information export (partial SMBIOS table).
- DCMI 1.5 compliance.
- Management support for Power Management Bus (PMBus*) 1.2 compliant power supplies.
- BMC data repository (managed data region feature).
- Support for an Intel® Local Control Panel display.
- System airflow monitoring.
- Exit air temperature monitoring.
- Ethernet controller thermal monitoring.
- Global aggregate temperature margin sensor.
- Memory thermal management.
- Power supply fan sensors.

- ENERGY STAR* server support.
- Smart ride through (SmaRT) / closed-loop system throttling (CLST).
- Power supply cold redundancy.
- Power supply firmware update.
- Power supply compatibility check.
- BMC firmware reliability enhancements:
    - Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC.
    - BMC system management health monitoring.
- Monitor compute node cooling type

10.2 Platform Management Features and Functions

10.2.1 Power Subsystem

The compute modules in the system support several power control sources that can initiate power-up or power-down activity, as detailed in Table 21**Error! Reference source not found.**

Table 21. Power control sources

Source	External Signal Name or Internal Subsystem	Capability
Power button	Front panel power button	Turns power on or off
BMC watchdog timer	Internal BMC timer	Turns power off, or power cycle
BMC chassis control commands	Routed through command processor	Turns power on or off, or power cycle
Power state retention	Implemented by means of BMC internal logic	Turns power on when AC power returns
Chipset	Sleep S4/S5 signal (same as POWER_ON)	Turns power on or off
CPU thermal	Processor Thermtrip	Turns power off
PCH thermal	PCH Thermtrip	Turns power off
WOL (Wake On LAN)	LAN	Turns power on

10.2.2 Advanced Configuration and Power Interface (ACPI)

The compute modules in the server system support the Advanced Configuration and Power Interface (ACPI) states described in Table 22.

Table 22. ACPI power states

State	Supported	Description
S0	Yes	Working. • Front panel power LED is on (not controlled by the BMC). • Fans spin at the normal speed, as determined by sensor inputs. • Front panel buttons work normally.
S1	No	Not supported.
S2	No	Not supported.
S3	No	Supported only on workstation platforms. See appropriate platform specific Information for more information.
S4	No	Not supported.
S5	Yes	Soft off. • Front panel buttons are not locked. • Fans are stopped. • Power-up process goes through the normal boot process. • Power, reset, front panel non-maskable interrupt (NMI), and ID buttons are unlocked.

10.2.3 System Initialization

During system initialization, both the BIOS and the BMC initialize the items described in the following sections.

10.2.3.1 Processor Tcontrol Setting

Processors used with this chipset implement a feature called Tcontrol, which provides a processor-specific value that can be used to adjust the fan-control behavior to achieve optimum cooling and acoustics. The BMC reads these from the CPU through PECI Proxy mechanism provided by Intel® ME. The BMC uses these values as part of the fan-speed-control algorithm.

10.2.3.2 Fault Resilient Booting (FRB)

Fault resilient booting (FRB) is a set of BIOS and BMC algorithms and hardware support that allow a multiprocessor system to boot even if the bootstrap processor (BSP) fails. Only FRB-2 is supported using watchdog timer commands.

FRB-2 refers to the FRB algorithm that detects system failures during POST. The BIOS uses the BMC watchdog timer to back up its operation during POST. The BIOS configures the watchdog timer to indicate that the BIOS is using the timer for the FRB2 phase of the boot operation.

After the BIOS has identified and saved the BSP information, it sets the FRB-2 timer use bit and loads the watchdog timer with the new timeout interval.

If the watchdog timer expires while the watchdog use bit is set to FRB-2, the BMC (if so configured) logs a watchdog expiration event showing the FRB-2 timeout in the event data bytes. The BMC then hard resets the system, assuming the BIOS-selected reset as the watchdog timeout action.

The BIOS is responsible for disabling the FRB-2 timeout before initiating the option ROM scan and before displaying a request for a boot password. If the processor fails and causes an FRB-2 timeout, the BMC resets the system.

The BIOS gets the watchdog expiration status from the BMC. If the status shows an expired FRB-2 timer, the BIOS enters the failure in the system event log (SEL). In the OEM bytes entry in the SEL, the last POST code generated during the previous boot attempt is written. FRB-2 failure is not reflected in the processor status sensor value.

The FRB-2 failure does not affect the front panel LEDs.

10.2.3.3 Post Code Display

The BMC, upon receiving standby power, initializes internal hardware to monitor port 80h (POST code) writes. Data written to port 80h is output to the system POST LEDs. The BMC deactivates POST LEDs after POST completes. Refer to Appendix B for a complete list of supported POST code diagnostic LEDs.

10.2.4 Watchdog Timer

The BMC implements a fully IPMI 2.0 compatible watchdog timer. For details, see the *Intelligent Platform Management Interface Specification Second Generation v2.0*. The NMI/diagnostic interrupt for an IPMI 2.0 watchdog timer is associated with an NMI. A watchdog pre-timeout SMI or equivalent signal assertion is not supported.

10.2.5 System Event Log (SEL)

The BMC implements the system event log as specified in the *Intelligent Platform Management Interface Specification v2.0*. The SEL is accessible regardless of the system power state through the BMC's in-band and out-of-band interfaces.

The BMC allocates 95,231 bytes (approx. 93 KB) of non-volatile storage space to store system events. The SEL timestamps may not be in order. Up to 3,639 SEL records can be stored at a time. Because the SEL is

circular, any command that results in an overflow of the SEL beyond the allocated space will overwrite the oldest entries in the SEL, while setting the overflow flag.

10.3 Sensor Monitoring

The BMC monitors system hardware and reports system health. The information gathered from physical sensors is translated into IPMI sensors as part of the IPMI sensor model. The BMC also reports various system state changes by maintaining virtual sensors that are not specifically tied to physical hardware. This section describes general aspects of BMC sensor management as well as describing how specific sensor types are modeled. Unless otherwise specified, the term “sensor” refers to the IPMI sensor-model definition of a sensor.

- Sensor scanning
- BIOS event-only sensors
- Margin sensors
- IPMI watchdog sensor
- BMC watchdog sensor
- BMC system management health monitoring
- VR watchdog timer
- System airflow monitoring sensors - valid for Intel® server chassis only
- Fan monitoring sensors
- Thermal monitoring sensors
- Voltage monitoring sensors
- CATERR sensor
- LAN leash event monitoring
- CMOS battery monitoring
- NMI (diagnostic interrupt) sensor

10.3.1 Sensor Re-arm Behavior

10.3.1.1 Manual versus Automatic Re-arm Sensors

Sensors can be re-armed either manually or automatically. An automatic re-arm sensor re-arms (clears) the assertion event state for a threshold or offset if that threshold or offset is de-asserted after having been asserted. This allows a subsequent assertion of the threshold or an offset to generate a new event and associated side-effect. An example side-effect is boosting fans due to an upper critical threshold crossing of a temperature sensor. The event state and the input state (value) of the sensor track each other. Most sensors are of the auto-rearm type.

A manual re-arm sensor does not clear the assertion state even when the threshold or offset becomes de-asserted. In this case, the event state and the input state (value) of the sensor do not track each other. The event assertion state is "sticky". The following methods can be used to re-arm a sensor:

- Automatic re-arm – Only applies to sensors that are designated as auto re-arm.
- IPMI command – Re-arm sensor event.
- BMC internal method – The BMC may re-arm certain sensors due to a trigger condition. For example, some sensors may be re-armed due to a system reset. A BMC reset re-arms all sensors.
- System reset or DC power cycle – Re-arms all system fan sensors.

10.3.2 Thermal Monitoring

The BMC provides monitoring of component and board temperature sensing devices. This monitoring capability is instantiated in the form of IPMI analog/threshold or discrete sensors, depending on the nature of the measurement.

For analog/threshold sensors, with the exception of processor temperature sensors, critical and non-critical thresholds (upper and lower) are set through SDRs and event generation enabled for both assertion and de-assertion events.

For discrete sensors, both assertion and de-assertion event generation are enabled.

Mandatory monitoring of platform thermal sensors includes:

- Inlet temperature (physical sensor is typically on system front panel or HDD back plane),
- Board ambient thermal sensors,
- Processor temperature,
- Memory (DIMM) temperature,
- CPU voltage regulator down (VRD) hot monitoring, and
- Power supply unit (PSU) inlet temperature (only supported for PMBus*-compliant PSUs).

Additionally, the BMC firmware may create virtual sensors that are based on a combination or aggregation of multiple physical thermal sensors and the application of a mathematical formula to thermal or power sensor readings.

10.3.3 Memory Thermal Management

The system memory is the most complex subsystem to manage thermally, as it requires substantial interactions between the BMC, BIOS, and the embedded memory controller hardware. This section provides an overview of this management capability from a BMC perspective.

10.3.3.1 Memory Thermal Throttling

The system supports thermal management through closed loop throttling (CLTT) only. Throttling levels are changed dynamically to cap throttling based on memory and system thermal conditions as determined by the system and DIMM power and thermal parameters. The BMC fan speed control functionality is related to the memory throttling mechanism used.

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

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

Intel® Server Systems supporting the Intel® Xeon® processor Scalable family support a type of CLTT, called a hybrid CLTT, for which the integrated memory controller estimates the DRAM temperature in between actual reads of the TSODs. Hybrid CLTT is used on all Intel Server Systems supporting the Intel® Xeon processor Scalable family that have DIMMs with thermal sensors. Therefore, the terms Dynamic-CLTT and Static-CLTT are really referring to this “hybrid” mode. Note that if the IMC’s polling of the TSODs is interrupted, the temperature readings that the BMC gets from the IMC are these estimated values.

10.3.3.2 Dynamic (Hybrid) CLTT

The system supports dynamic (memory) CLTT for which the BMC firmware dynamically modifies thermal offset registers in the IMC during runtime based on changes in system cooling (fan speed). For static CLTT, a fixed offset value is applied to the TSOD reading to get the die temperature; however this is does not provide results as accurate when the offset takes into account the current airflow over the DIMM, as is done with dynamic CLTT.

To support this feature, the BMC firmware derives the air velocity for each fan domain based on the PWM value being driven for the domain. Since this relationship is dependent on the chassis configuration, a method must be used which supports this dependency (for example, through OEM SDR) that establishes a lookup table providing this relationship.

BIOS has an embedded lookup table that provides thermal offset values for each DIMM type and air velocity range (three ranges of air velocity are supported). During system boot, BIOS provides three offset values (corresponding to the three air velocity ranges) to the BMC for each enabled DIMM. Using this data the BMC firmware constructs a table that maps the offset value corresponding to a given air velocity range for each DIMM. During runtime the BMC applies an averaging algorithm to determine the target offset value corresponding to the current air velocity and then the BMC writes this new offset value into the IMC thermal offset register for the DIMM.

10.3.4 Power Management Bus (PMBus*)

The Power Management Bus (PMBus*) is an open standard protocol that is built upon the SMBus* 2.0 transport. It defines a means of communicating with power conversion and other devices using SMBus-based commands. A system must have PMBus-compliant power supplies installed in order for the BMC or Intel® ME to monitor them for status and/or power metering purposes.

For more information on PMBus*, see the System Management Interface Forum website, http://www.powersig.org/.

1.1.2 System cooling type determination

The BMC determines the cooling type configuration of each compute module and chassis by reading the FRU data during system boot. The BMC will alert and halt the system when an incompatible configuration is detected. This feature is designed to prevent system damage and malfunction by ensuring that the correct Air-Cooled or Liquid Cooled system configuration is selected during the boot sequence.

11. Thermal Management

The Intel® Server System S9200WK product family supports liquid and air cooling for the compute modules, and only air cooling for the chassis. This chapter provides an overview of the thermal management features and capabilities of the system.

The fully integrated system is designed to operate at external ambient temperatures of between 10°C and 35° C as stated in Table 2. Working with integrated platform management, several features within the system are designed to capture heat from direct contact, or move air from front to back through the system and over critical components to prevent them from overheating, allowing the system to operate with optimal performance.

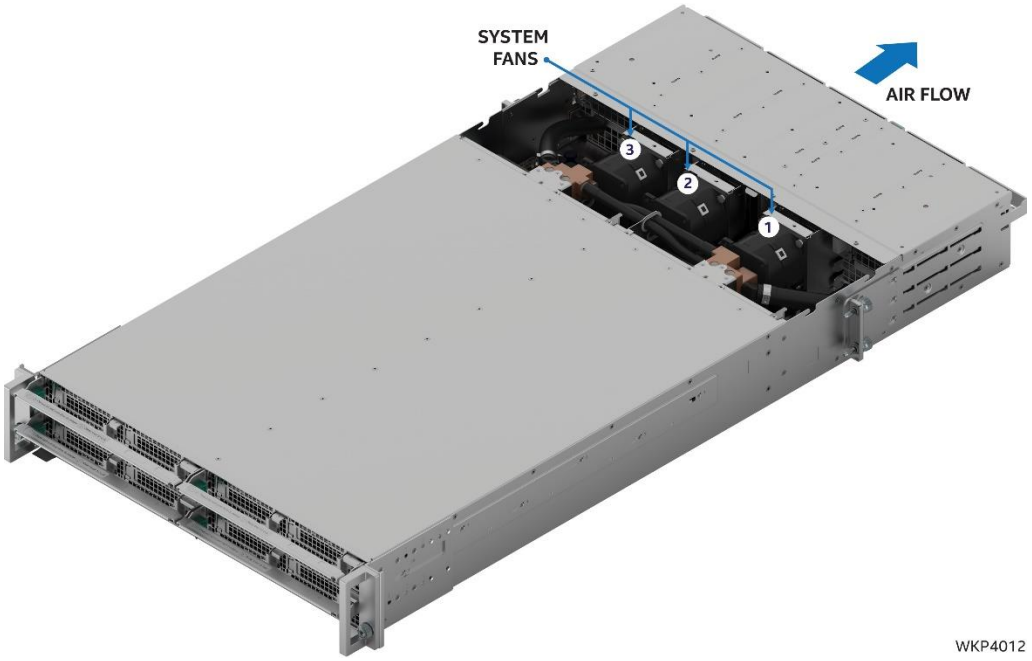

Figure 53. System airflow and fan identification – Liquid cooled configuration

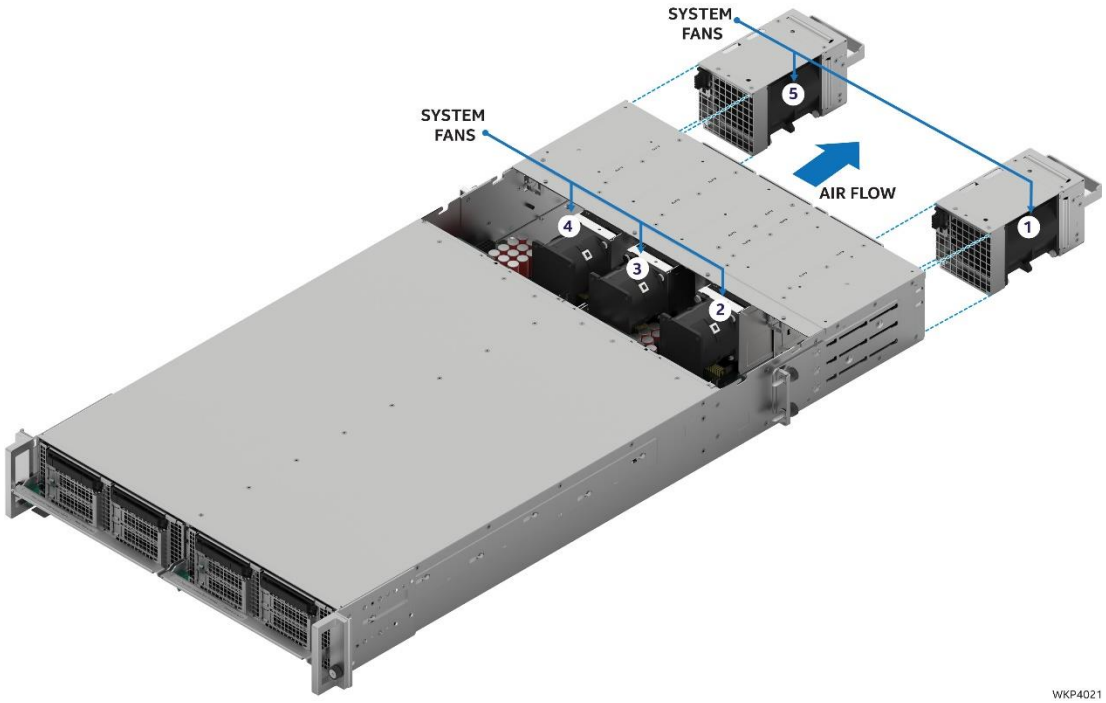

Figure 54. System airflow and fan identification – Air cooled configuration

The installation and functionality of several system components is used to maintain system thermals. Depending on system cooling configurations, the involved components are listed below:

- Liquid cooled configuration
  - Three managed 60 mm dual rotor hot-swap system fans
  - Fans integrated into each installed power supply module
  - Liquid cooling loop
  - Add-in card bays
  - Populated DIMM slots
  - Board component heat sinks.
  - Populated drive carriers (for 2U compute modules)
- Air cooled configuration
  - Three managed 60 mm dual rotor hot-swap system fans
  - Two managed 80 mm dual rotor hot-swap system fans
  - Fans integrated into each installed power supply module
  - Populated drive carriers (for 2U compute modules)
  - Add-in card bays
  - Populated DIMM slots
  - Board component heat sinks.
  - CPU heat sinks
  - Air duct

Drive carriers can be populated with a storage device or supplied drive blank.

11.1 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:

- For air cooled configurations:
  - The system is designed to sustain operations at an ambient temperature of up to 35 °C for 2U form factor and 27 °C for 1U form factor compute modules (refer details on Appendix D).

Note: 1U air cooled compute node cannot support 14W (64GB) DIMMs at 27 °C.

Table 23 Supported Ambient air conditions per DIMM on 1U AC compute node at 900m elevation

DIMM size	DIMM Power (W)	Max Ambient (C)
16GB	4	34
	6	32
32GB	8	30
	10	28
	12	27
64GB	14	25
	16	23

- For liquid cooled configurations:
  - The system is designed to sustain operations at an ambient temperature of up to 35 °C (ASHRAE Class A2)
  - The system is designed to support facility coolant supply temperature of 2 °C to 32 °C (ASHRAE Class W3)
- For air cooled mixed configurations: The system is designed to sustain operations at an ambient temperature of up to 32 °C
  - 1U/2U AC assumes 1U AC nodes with 250 W TDP and 2U AC nodes with 350 W TDP.
- For liquid cooled mixed configurations: The system is designed to sustain operations at an ambient temperature of up to 35 °C
  - 1U/2U LC assumes 1U LC nodes with 400 W TDP and 2U LC nodes with 400 W TDP.

11.2 Thermal Management Overview

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

For system configurations that support fan redundancy, should a single rotor failure occur in a system fan or power supply fan, integrated platform management changes the state of the system status LED to blinking green, reports an error to the system event log, and automatically adjusts fan speeds as needed to maintain system temperatures below maximum thermal limits. Fan redundancy is lost if more than one fan rotor is in a failed state.

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

In the event that system temperatures should continue to increase with the system fans operating at their maximum speed, platform management may begin to throttle bandwidth of either the memory subsystem, the processors, or both to keep components from overheating in order to keep the system operational. Throttling of these sub-systems continues until system temperatures are reduced below preprogrammed limits.

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

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

Note: For proper system thermal management, Sensor Data Records (SDRs) for any given system configuration must be updated when performing a system software stack update. SDRs are loaded using the FRUSDR utility which is part of the System Update Package (SUP) or One-boot Firmware Update (OFU) package which can be downloaded from http://downloadcenter.intel.com.

11.2.1 Fan Speed Control

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

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

A fan domain has three states: sleep, boost, and nominal. 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. The factors for the boost state are listed below in order of precedence, high to low. If any of these conditions apply, the fans are set to a fixed boost state speed.

- An associated fan is in a critical state or missing. The SDR describes which fan domains are boosted in response to a fan failure or removal in each domain. If a fan is removed when the system is in fans-off mode, it is not detected and there is not any fan boost until the system comes out of fans-off mode.
- 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.

For more information on nominal fan speed, see Section 11.2.1.5.

11.2.1.1 Programmable Fan PWM Offset

The compute modules provide a BIOS setup option to boost the system fan speed by a programmable positive offset setting. Setting the **Fan PWM Offset** option causes the BMC to add the offset to the fan speeds to which it would otherwise be driving the fans. This 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 and non-Intel chassis and better support of ambient temperatures higher than 35 °C.

11.2.1.2 Hot-Swappable Fans

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

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

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

11.2.1.3 Fan redundancy Detection

The BMC supports redundant fan monitoring and implements a fan redundancy sensor. A fan redundancy sensor generates events when its associated set of fans transitions between redundant and non-redundant states, as determined by the number and health of the fans. The definition of fan redundancy is

configuration dependent. The BMC allows redundancy to be configured on a per fan redundancy sensor basis through OEM SDR records.

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

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

11.2.1.4 Fan Domains

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

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

The same device may drive multiple PWM signals.

11.2.1.5 Nominal Fan Speed

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

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

The PWM duty cycle value for a domain is computed as a percentage using one or more instances of a stepwise linear algorithm and a clamp algorithm. The transition from one computed nominal fan speed (PWM value) to another is ramped over time to minimize audible transitions. The ramp rate is configurable by means of the OEM SDR.

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

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

11.2.1.6 Thermal and Acoustic Management

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

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

11.2.1.7 Thermal Sensor Input to Fan Speed Control

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

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

- Front panel temperature sensor [1]
- CPU margin sensors [2, 4, 5]
- DIMM thermal margin sensors [2, 4]
- Exit air temperature sensor [1, 7, 9]
- PCH temperature sensor [3, 5]
- Onboard Ethernet controller temperature sensors [3, 5]
- 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]
- Riser card temperature sensors

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

Figure 55 shows a high-level representation of the fan speed control structure that determines fan speed.

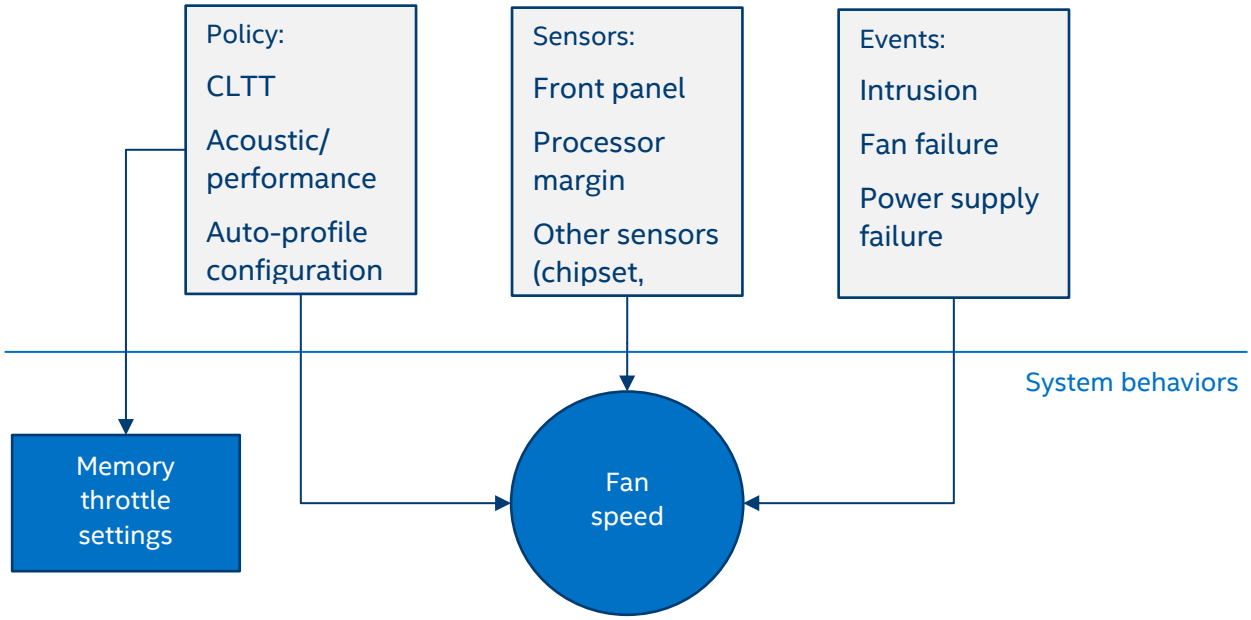

Figure 55. High-level fan speed control model

11.3 System Fans

The system includes different fans depending on cooling configurations, listed below:

- Liquid cooled configuration
  - Three managed 60 mm dual rotor hot-swap system fans
  - Fans integrated into each installed power supply module
- Air cooled configuration
  - Three managed 60 mm dual rotor hot-swap system fans
  - Two managed 80 mm dual rotor hot-swap system fans
  - Fans integrated into each installed power supply module

The system fans and power supply module fans provide the primary airflow for the system.

The system is designed for fan redundancy when configured with three power supply modules, all system fan rotors are operational, and ambient air remains at or below 35 °C. Should a single 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 fan rotor is in a failed state.

The system is capable of supporting up to a total of five system fans in air cooled configurations and up to three system fans in liquid cooled configurations. All system fans are mounted in an individual fan assembly module Each system fan includes support for the following:

- Each individual fan is hot-swappable
- Each fan is connected to a matching 10-pin connector located on the power distribution board
- Each fan is designed for tool-less insertion and extraction from the system chassis
- Each fan has a tachometer signal that allows the integrated BMC to monitor its status
- Fan speed for each fan is controlled by integrated platform management. As system thermals fluctuate high and low, the integrated BMC firmware increases and decreases the speeds to specific fans within the fan assembly to regulate system thermals.
- The system fans are mounted inside a fan cage assembly that can be removed from the back of the system chassis.
- The fan cage assembly includes a fan fail LED visible from the back of the system chassis.

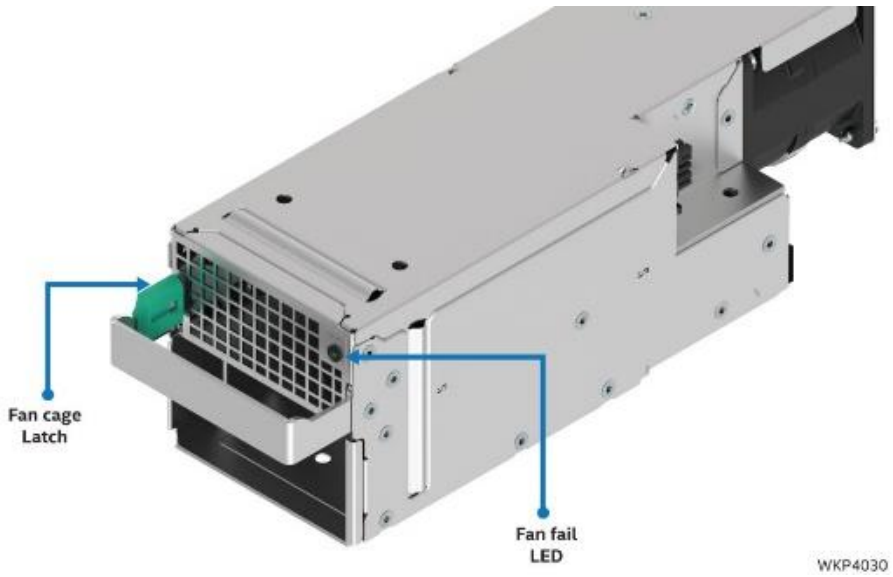

Figure 56. System fan cage assembly

Each system fan connects to the power distribution board through a 2x5 pin connector. The following table lists the connector pinout.

Table 24. System fan connector pinout

Pin	Signal Name	Pin	Signal Name
1	P12V_FAN_ROTOR1	6	FAN_PRSNT_N
2	GND	7	FAN_CATH_LED
3	FAN_TACH1_FLT	8	FAN_ANODE_LED
4	FAN_TACH2_FLT	9	P12V_FAN_ROTOR2
5	FAN_PWM_FLT	10	GND

11.4 Power Supply Module Fans

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

11.5 Liquid Cooling Support

Select system configurations within the Intel® Server System S9200WK product family support liquid cooling on the installed compute modules through an integrated liquid cooling loop. The liquid cooling loop is designed as a passive component in terms of coolant flow that captures heat from the processors and memory inside the compute module.

The coolant flow for the liquid cooling loop is supported through liquid cooling plumbing connections installed in the back of the server chassis. The liquid cooling loops from all installed compute modules in the system are connected in parallel to the chassis plumbing connections. The chassis plumbing connections include two Staubli* SCG 06 quick disconnect couplings to connect to coolant supply and return.

Figure 57. Chassis plumbing couplings identification

11.5.1 Liquid Cooling Loop Components

The dedicated liquid cooling loop in the compute modules consists of the following components:

- Two cold plate assemblies with integrated retention mechanism
- 12 memory heat spreaders
- Blind-mate quick disconnect couplings
- Liquid cooling tubes (filled with liquid coolant)

11.5.2 Liquid Cooling Loop Operation

The dedicated liquid cooling loop kit has three basic elements: retention and cold plate units, memory cooling heat spreaders, and coolant tubes to transport the liquid. The cold plates are a direct replacement for the standard CPU air cooled heat sink. The cold plates and cooling pads transfer the heat from the CPU and memory respectively to the cooling liquid. A non-Intel coolant distribution unit (not included) connected

to the chassis plumbing continuously pushes the warm liquid to the exterior and draws cool liquid back to the cold plates. The flow rate in conjunction with the enhanced heat transfer cold plates and heat spreaders capture heat from the surface of the processor and the installed memory. The system fans pull air across the system exhausting heat from the other components on the compute modules and chassis to the outside of the system. Figure 58 shows the liquid cooling loop kit flow path as the liquid coolant passes through the compute module.

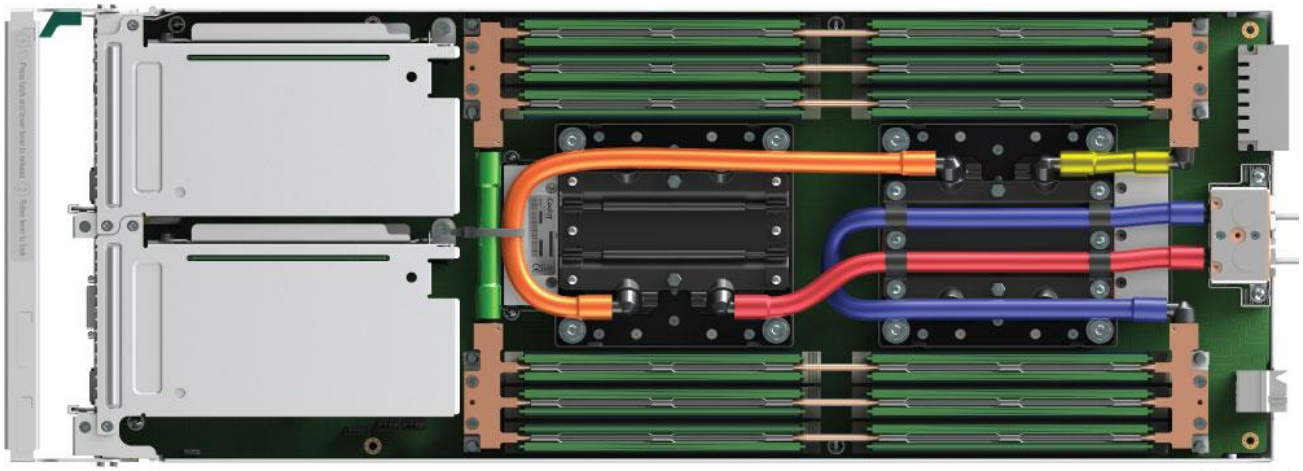

Flow order	Component
First	1st group of memory heat spreaders
Second	2nd group of memory heat spreaders
Third	CPU2 cold plate
Fourth	CPU1 cold plate
Fifth	To coolant return

Figure 58. Liquid cooling loop flow map

Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for specific instructions regarding the removal and installation of the liquid cooling loop.

11.5.3 Liquid Cooling Loop Specification

The cold plate assembly in the liquid cooling loop mounts directly on top of the processors. The retention mechanism for installation on top of the processors is integrated into the cold plate. The liquid coolant contained within the tubes, is a mixture of demineralized water and propylene glycol with the following benefits: Anti-Freeze, Anti-Corrosion and Anti-Bacterial. Table 25 and

Table 26 list the features and specifications of the liquid cooling loop kit.

Table 25. Liquid cooling loop specifications

Specification	Value
Cold plate material	Copper
Thermal Interface Material	Dowsil TC-5622
Memory heat spreader Thermal Interface Material	Laird 300H
Dimensions	
Height	34.05 mm
Weight	6.94 lbs (3.15 kg)

Table 26. Liquid cooling loop operation temperature specifications

Specification	Value
Operating Liquid Temperature	Minimum: 2°C Maximum: 65°C
Operating Air temperature	32°C
Coolant flow rate	0.4 – 1.16 l/min (per compute module)
Operating humidity	5-95%
Storage temperature	-40 to 70 °C
Storage humidity	5-95%

Important: The liquid cooling loop and the chassis plumbing are pre charged with liquid coolant.

11.6 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 compute modules with platform specific environmental limits, configuration data, and the appropriate sensor data records (SDRs) for use by these management features.

As part of the system manufacturing process, a default software stack is loaded that contains FRU and SDR data, but may not be the latest available version. Intel recommends updating the SDR to the latest available as part of a planned system software update.

The FRUSDR utility for the given server platform can be downloaded as part of the system update package (SUP) or one-boot firmware update (OFU) package from http://downloadcenter.intel.com.

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

See Appendix D for flow rate requirements for liquid cooling configurations.

12. System Security

The compute modules within the Intel® Server System S9200WK product family support a variety of security options designed to prevent unauthorized access or tampering of settings. Security options supported include:

- Password protection
- Front panel lockout
- Trusted Platform Module (TPM) support
- Intel® Trusted Execution Technology (Intel® TXT)

12.1 Password Protection

The BIOS setup utility, accessed during POST, includes a Security tab where options to configure passwords, front panel lockout, and TPM settings, can be found.

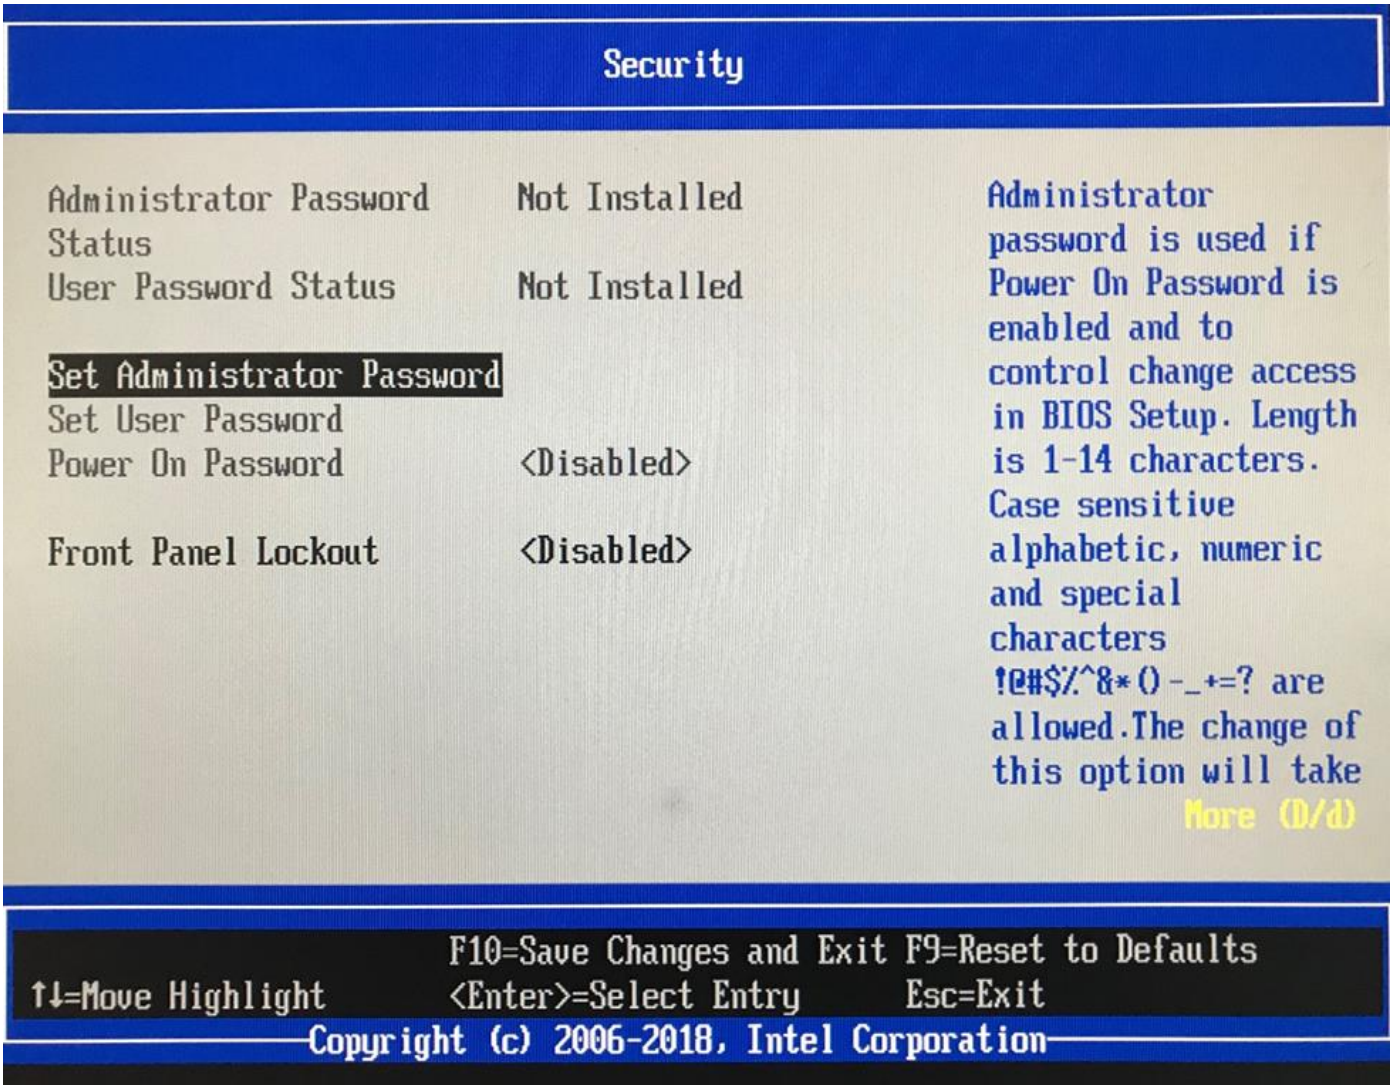

Figure 59. BIOS setup Security tab

12.1.1 Password Setup

The BIOS uses passwords to prevent unauthorized access to the compute module. Passwords can restrict entry to the BIOS setup utility, restrict use of the Boot Device popup menu during POST, suppress automatic

USB device re-ordering, and prevent unauthorized power on. It is strongly recommended that an administrator password be set. A system with no administrator password set allows anyone who has access to the compute modules to change BIOS settings.

An administrator password must be set in order to set the user password.

The maximum length of a password is 14 characters and can be made up of a combination of alphanumeric (a-z, A-Z, 0-9) characters and any of the following special characters:

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

Passwords are case sensitive.

The administrator and user passwords must be different from each other. An error message is displayed and a different password must be entered if there is an attempt to enter the same password for both. The use of strong passwords is encouraged, but not required. In order to meet the criteria for a strong password, the password entered must be at least eight characters in length, and must include at least one each of alphabetic, numeric, and special characters. If a weak password is entered, a warning message is displayed, and the weak password is accepted. Once set, a password can be cleared by changing it to a null string. This requires the administrator password, and must be done through BIOS setup. Clearing the administrator password also clears the user password. Passwords can also be cleared by using the password clear jumper on the compute module. For more information on the password clear jumper, see Section 13.2.

Resetting the BIOS configuration settings to default values (by any method) has no effect on the administrator and user passwords.

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

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

12.1.2 System Administrator Password Rights

When the correct administrator password is entered when prompted, the user has the ability to perform the following actions:

- Access the BIOS setup utility.
- Configure all BIOS setup options in the BIOS setup utility.
- Clear both the administrator and user passwords.
- Access the Boot Menu during POST.

If the Power On Password function is enabled in BIOS setup, the BIOS halts early in POST to request a password (administrator or user) before continuing POST.

12.1.3 Authorized System User Password Rights and Restrictions

When the correct user password is entered, the user has the ability to perform the following actions:

- Access the BIOS setup utility.
- View, but not change, any BIOS setup options in the BIOS setup utility.
- Modify system time and date in the BIOS setup utility.

If the Power On Password function is enabled in BIOS setup, the BIOS halts early in POST to request a password (administrator or user) before continuing POST.

Configuring an administrator password imposes restrictions on booting the system, and configures most setup fields to read-only if the administrator password is not provided. The boot popup menu requires the

administrator password to function, and the USB reordering is suppressed as long as the administrator password is enabled. Users are restricted from booting in anything other than the boot order defined in setup by an administrator.

12.2 Front Panel Lockout

If enabled in BIOS setup from the Security screen, this option disables the following front panel features:

- The off function of the power button.
- System reset button.

If front panel lockout is enabled, power off and reset must be controlled via a system management interface.

12.3 Trusted Platform Module (TPM) Support

The Trusted Platform Module (TPM) is a hardware-based security device that addresses the growing concern about boot process integrity and offers better data protection. TPM protects the system startup process by ensuring it is tamper-free before releasing system control to the operating system. A TPM device provides secured storage to store data, such as security keys and passwords. In addition, a TPM device has encryption and hash functions. The compute modules implement TPM as per *TPM PC Client Specifications revision 2.0,* published by the Trusted Computing Group (TCG).

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

After the BIOS completes the measurement of its boot process, it hands off control to the operating system loader and, in turn, to the operating system. If the operating system is TPM-enabled, it compares the BIOS TPM measurements to those of previous boots to make sure the system was not tampered with before continuing the operating system boot process. Once the operating system is in operation, it optionally uses TPM to provide additional system and data security (for example, Microsoft Windows 10* supports Bitlocker* drive encryption).

12.3.1 TPM Security BIOS

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

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

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

12.3.2 Physical Presence

Administrative operations to the TPM require TPM ownership or physical presence indication by the operator to confirm the execution of administrative operations. The BIOS implements the operator presence indication by verifying the setup administrator password.

A TPM administrative sequence invoked from the operating system proceeds as follows:

1. A user makes a TPM administrative request through the operating system's security software.
2. The operating system requests the BIOS to execute the TPM administrative command through TPM ACPI methods and then resets the system.
3. The BIOS verifies the physical presence and confirms the command with the operator.
4. The BIOS executes TPM administrative command, inhibits BIOS setup entry, and boots directly to the operating system which requested the TPM command.

12.3.3 TPM Security Setup Options

The BIOS TPM setup allows the operator to view the current TPM state and to carry out rudimentary TPM administrative operations. Performing TPM administrative options through the BIOS setup requires TPM physical presence verification.

Using the BIOS TPM setup, the operator can turn TPM functionality on or off and clear the TPM ownership contents. After the requested TPM BIOS setup operation is carried out, the option reverts to No Operation.

The BIOS TPM setup also displays the current state of the TPM, whether TPM is enabled or disabled and activated or deactivated. Note that while using TPM, a TPM-enabled operating system or application may change the TPM state independently of the BIOS setup. When an operating system modifies the TPM state, the BIOS Setup displays the updated TPM state.

The BIOS setup **TPM Clear** option allows the operator to clear the TPM ownership key and allows the operator to take control of the system with TPM. You use this option to clear security settings for a newly initialized system or to clear a system for which the TPM ownership security key was lost.

12.4 Intel® Trusted Execution Technology

The Intel® Xeon® Platinum 9200 Processor family supports Intel® Trusted Execution Technology (Intel® TXT), which is a robust security environment. Designed to help protect against software-based attacks, Intel® TXT integrates new security features and capabilities into the processor, chipset, and other platform components. When used in conjunction with Intel® Virtualization Technology (Intel® VT), Intel® TXT provides hardware-rooted trust for virtual applications.

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

Intel® TXT requires a computer system with Intel Virtualization Technology enabled (both VT-x and VT-d), an Intel® TXT -enabled processor, chipset, and BIOS, Authenticated Code Modules, and an Intel® TXT compatible measured launched environment (MLE). The MLE could consist of a virtual machine monitor, an OS, or an application. In addition, Intel® TXT requires the system to include a TPM v2.0, as defined by the *Trusted Computing Group TPM PC Client Specifications, Revision 2.0*.

When available, Intel® TXT can be enabled or disabled in the processor by a BIOS setup option. For general information about Intel® TXT, visit http://www.intel.com/technology/security/.

13. Reset and Recovery Jumpers

The compute modules within the Intel® Server System S9200WK product family include several jumper blocks which can be used to configure, protect, or recover specific features of the compute modules. Figure 60 identifies the location of each jumper block on the compute module server board. Pin 1 of each jumper block can be identified by the arrowhead (▼) silkscreened on the compute module server board next to the pin.

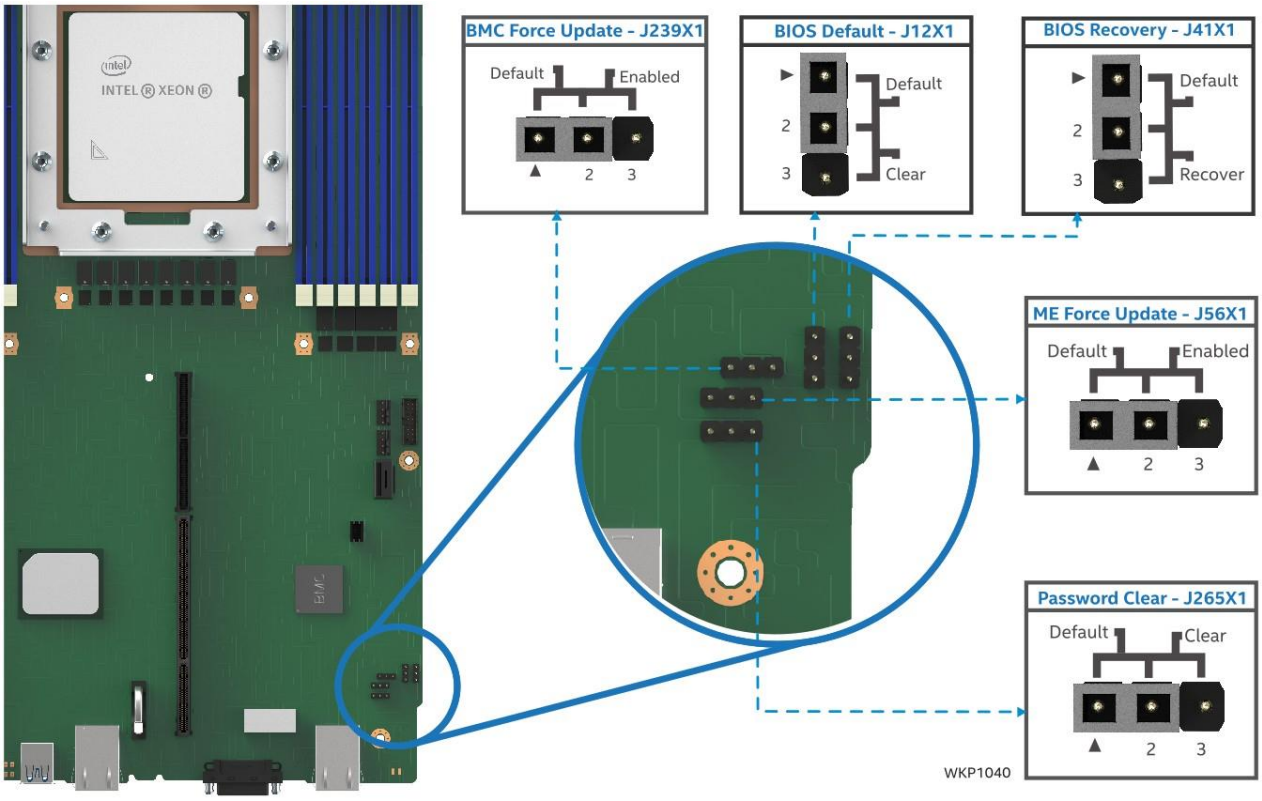

Figure 60. Reset and recovery jumper block location

The following sections describe how each jumper block is used.

13.1 BIOS Default Jumper Block

This jumper resets BIOS options, configured using the <F2> BIOS Setup Utility, back to their original default factory settings.

Note: This jumper does not reset administrator or user passwords. To reset passwords, the password clear jumper must be used.

To use the BIOS default jumper, perform the following steps:

1. Power down the compute module.
2. Remove the compute module from the chassis and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
3. Move the “BIOS DFLT” jumper from pins 1-2 (normal operation) to pins 2-3 (set BIOS defaults).
4. Wait five seconds then move the jumper back to pins 1-2.
5. Re-install the riser assemblies.
6. Re-Install compute module in the chassis.

Note: The system automatically powers on after AC is applied to the system.

7. Power on the compute module and press **<F2>** during POST to access the BIOS setup utility to configure and save desired BIOS options.

After resetting BIOS options using the BIOS default jumper, the Error Manager Screen in the BIOS setup utility displays two errors:

- 0012 System RTC date/time not set
- 5220 BIOS Settings reset to default settings

Also, the system time and date may need to be reset.

13.2 Password Clear Jumper Block

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

To use the password clear jumper, perform the following steps:

1. Power down the compute module.
2. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
3. Move the password clear jumper from pins 1-2 (default) to pins 2-3 (password clear position).
4. Re-install the riser assemblies and re-install the compute module in the chassis.
5. Power on the compute module and press **<F2>** during POST to access the BIOS setup utility.
6. Verify the password clear operation was successful by viewing the Error Manager screen. Two errors should be logged:
    - 5221 Passwords cleared by jumper
    - 5224 Password clear jumper is set
7. Exit the BIOS setup utility and power down the compute module.
8. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
9. Move the password clear jumper back to pins 1-2 (default).
10. Re-install the riser assemblies and re-install the compute module in the chassis.
11. Power up the compute module.
12. It is strongly recommended to boot into BIOS setup immediately, navigate to the Security tab, and set the administrator and user passwords if intending to use BIOS password protection.

13.3 Intel® Management Engine (Intel® ME) Firmware Force Update Jumper Block

When the Intel® ME firmware force update jumper is moved from its default position, the Intel® ME is forced to operate in a reduced minimal operating capacity. This jumper should only be used if the Intel® ME firmware has gotten corrupted and requires re-installation.

Note: System update files are included in the system update packages (SUP) posted to Intel's download center website at http://downloadcenter.intel.com.

To use the Intel® ME firmware force update jumper, perform the following steps:

1. Power down the compute module.
2. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
3. Move the "ME FRC UPD" jumper from pins 1-2 (default) to pins 2-3 (force update position).
4. Re-install the riser assemblies and re-install the compute module in the chassis.
5. Power on the compute module.
6. Boot to the EFI shell.
7. Change directories to the folder containing the update files.
8. Update the Intel® ME firmware using the following command:

```
iflash32 /u /ni  <version#>_ME.cap
```

9. When the update has successfully completed, power off the compute module.
10. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
11. Move the "ME FRC UPD" jumper back to pins 1-2 (default).
12. Re-install the riser assemblies and re-install the compute module in the chassis.
13. Power on the compute module.

13.4 BMC Force Update Jumper Block

The BMC force update jumper is used to put the BMC in boot recovery mode for a low-level update. It causes the BMC to abort its normal boot process and stay in the boot loader without executing any Linux* code. This jumper should only be used if the BMC firmware has become corrupted and requires re-installation.

Note: System update files are included in the SUP posted to Intel's download center website at http://downloadcenter.intel.com.

To use the BMC force update jumper, perform the following steps:

1. Power down the compute module.
2. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
3. Move the "BMC FRC UPD" jumper from pins 1 - 2 (default) to pins 2 - 3 (force update position).
4. Re-install the riser assemblies and re-install the compute module in the chassis.
5. Power on the compute module.
6. Boot to the EFI shell.
7. Change directories to the folder containing the update files.
8. Update the BMC firmware using the following command:

```
FWPIAUPD -u -bin -ni -b -o -pia -if=usb <file name.BIN>
```

9. When the update has successfully completed, power off the compute module.
10. Power down the compute module.
11. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
12. Move the "BMC FRC UPD" jumper back to pins 1-2 (default).

13. Re-install the riser assemblies and re-install the compute module in the chassis.
14. Power on the compute module.
15. Boot to the EFI shell.
16. Change directories to the folder containing the update files.
17. Re-install the board/system SDR data by running the FRUSDR utility.
18. After the SDRs have been loaded, reboot the compute module.

13.5 BIOS Recovery Jumper

When the BIOS recovery jumper block is moved from its default pin position (pins 1–2), the compute module boots using a backup BIOS image to the UEFI shell, where a standard BIOS update can be performed. See the BIOS update instructions that are included with the SUP downloaded from Intel's download center website. This jumper is used when the BIOS has become corrupted and is non-functional, requiring a new BIOS image to be loaded on to the compute module.

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

Note: System update files are included in the SUP posted to Intel's download center website at http://downloadcenter.intel.com.

To use the BIOS recovery jumper, perform the following steps:

1. Power down the compute module.
2. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
3. Move the "BIOS Recovery" jumper from pins 1 – 2 (default) to pins 2 – 3 (BIOS recovery position).
4. Re-install the riser assemblies and re-install the compute module in the chassis.
5. Power on the compute module. The system automatically boots to the EFI shell.
6. Update the BIOS using the standard BIOS update instructions provided with the system update package.
7. After the BIOS update has successfully completed, power down the compute module.
8. Remove the compute module from the chassis, and remove the riser assemblies from the compute module. Refer to the *Intel® Server System S9200WK Product Family Setup and Service Guide* for instructions.
9. Move the BIOS recovery jumper back to pins 1 – 2 (default).
10. Re-install the riser assemblies and re-install the compute module in the chassis.
11. Power on the compute module. During POST, press **<F2>** to access the BIOS setup utility to configure and save desired BIOS options.

14. Power

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

14.1 Power distribution board (PDB)

The Intel® Server System S9200WK product family includes a power distribution board in the back of the chassis that provides power and management connections for several components in the system. Figure 61 shows the location of the PDB in the system and .

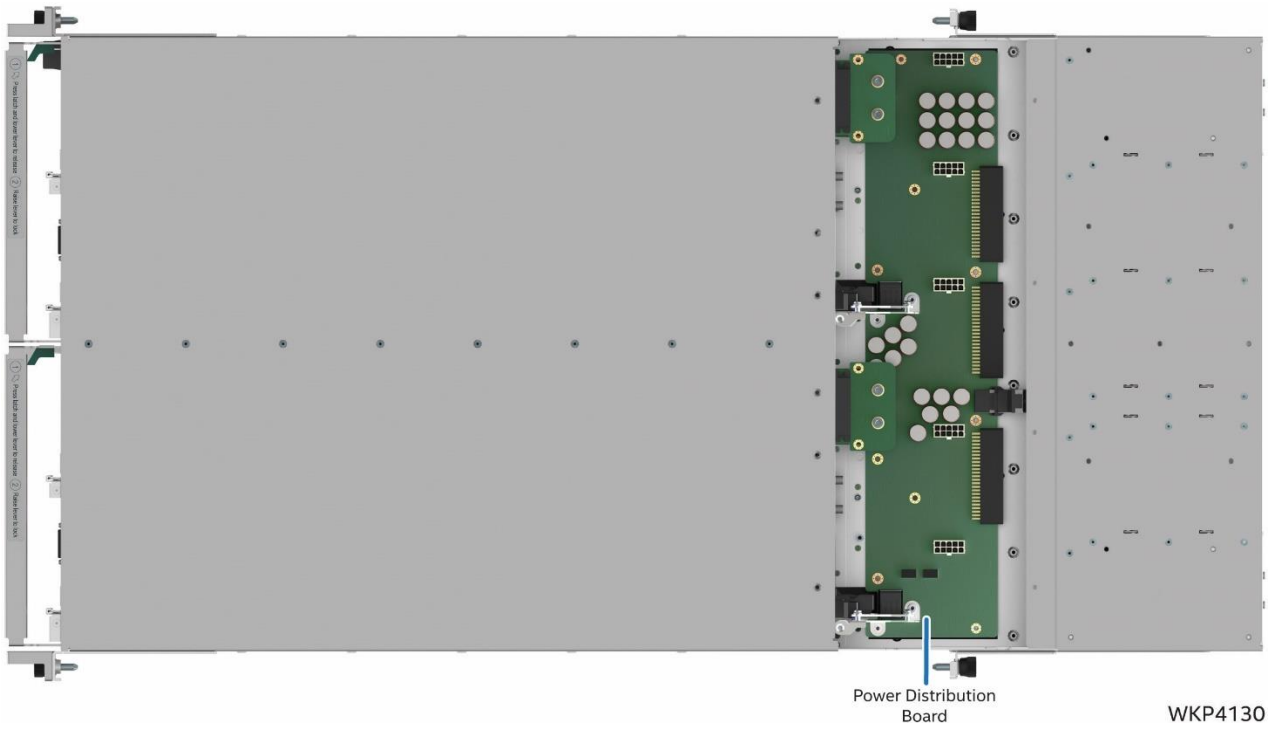

Figure 61. Power distribution board identification

The power distribution board includes the following connectors:

- Three power supply interface connectors in the rear
- One EMP Module docking port in the rear
- Five system fan connectors on top
- Two management risers with high speed management ports, for communication between the compute modules, the EMP module and system fans
- Four power connectors in the front, for providing power to the compute modules

Figure 62 identifies the different connectors in the power distribution board.

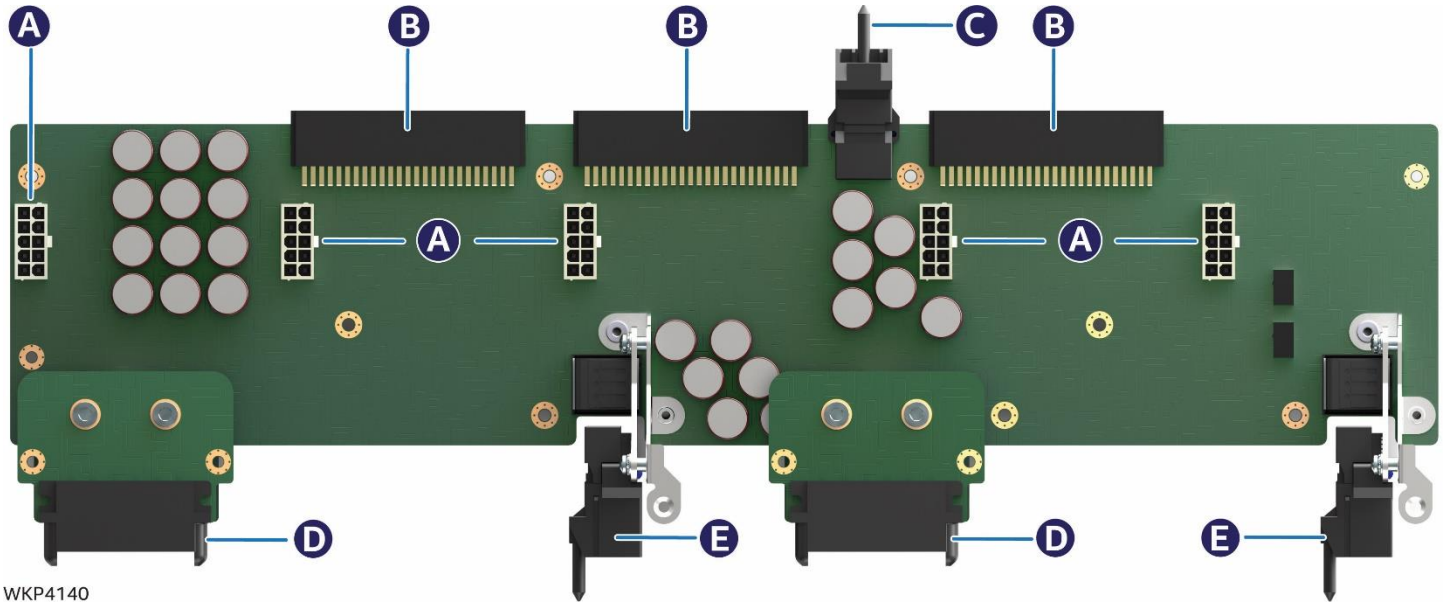

Connector Identifier	Description
A	System fan connectors
B	Power supply connectors
C	EMP module docking port
D	Compute module power connectors
E	Management risers

Figure 62. Power distribution board connector identification

14.2 Power Supply Configurations

The server system can have up to three power supply modules installed and can support the following power supply configurations:

- 2+1 redundant power
- 3+0 combined power (non-redundant)

2+1 redundant power and 3+0 combined power configurations are automatically configured depending on the amount of compute modules in the system. If the server system has two compute modules, the three power supplies will work in a 2+1 power configuration. If the server system has four compute modules, then the three power supplies will work in a 3+0 power configuration. Should system thermal levels exceed programmed limits, platform management attempts to keep the system operational. For details, see Section 14.6 and Chapter 11.

Note: 3+0 combined power configurations do not support redundancy.

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.

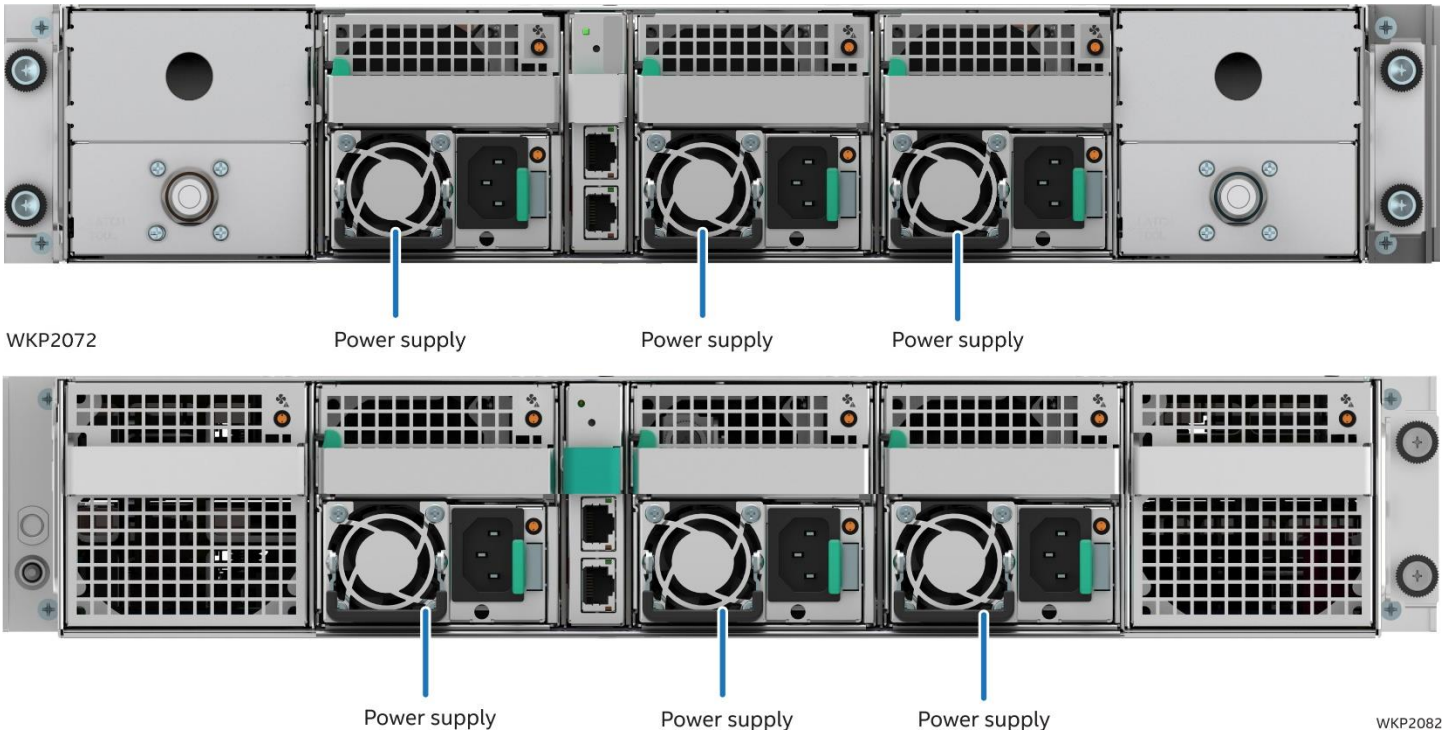

Figure 63. Power supply module identification

In the event of a power supply failure, the redundant 2+1 power supply configuration has support for hot-swap extraction and replacement of the failed power supply. The AC input is auto-ranging and power factor corrected

14.3 Power Supply Module Options

There are two power supply options available for the Intel® Server System S9200WK product family:

- 1600W AC 80 PLUS* Titanium – iPC AXX1600TCRPS
- 2100W AC 80 PLUS Platinum – iPC FCXX2100CRPS

Dimmensions: CRPS185: 40mm x 74mm x 185mm (H x W x D)

Output connector: Card edge compatible with FCi 2x25

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

Important: For systems with more than two compute modules installed, the system must be connected to 220 V input voltage in order to operate properly.

14.3.1 Power Supply Module Efficiency

The following tables provide the required minimum efficiency level at various loading conditions. These are provided at four 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 27. 1600W and 2100W AC power supplies efficiency (80 PLUS* Titanium)

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

14.3.2 Power Cord Specification Requirements

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

Table 28. AC power cord specifications

Cable Type	SJT
Wire Size	14 AWG
Temperature Rating	105ºC
Amperage Rating	15 A
Voltage Rating	250 V

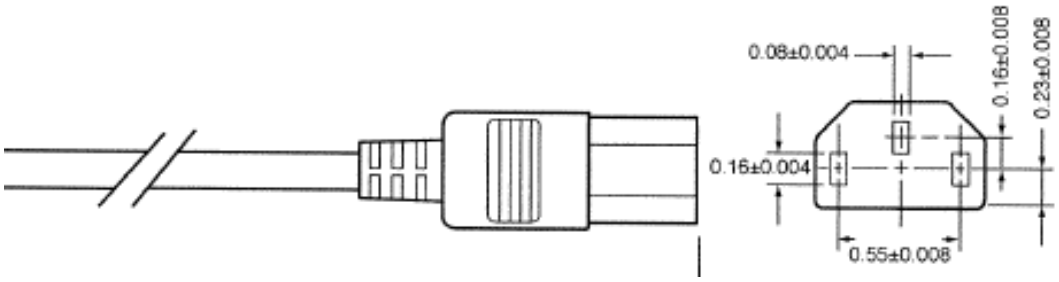

Figure 64. AC power cord specification

14.4 AC Power Supply Input Specifications

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

14.4.1 Power Factor

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

Table 29. AC power factor

Output Power	10% load	20% load	50% load	100% load
Power Factor	> 0.90	> 0.96	> 0.98	> 0.99

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

14.4.2 Power Supply Status LED

There is a single bi-color LED to indicate power supply status. The LED operation is defined in Table 30.

Table 30. LED indicators

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

14.5 Cold Redundancy Support

14.5.1 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 they need to power on. Depending upon which position (1, 2, or 3) the system defines that power supply to be in, the cold standby configuration slightly changes the load share threshold that the power supply will power on at.

Table 31. Example load share threshold for activating supplies

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

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

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

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

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

14.5.3 BMC Requirements

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

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

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

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

14.5.4 Power Supply Turn on Function

Powering on and off of the cold standby power supplies is only controlled by each PSU sensing the Vshare bus. Once a power supply turns on after crossing the enable threshold, it lowers itself to the disable

threshold. The system defines the position of each power supply in the cold redundant operation. It does this each time the system is powered on, a power supply fails, or a power supply is added to the system.

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

14.6 Closed Loop System Throttling (CLST)

The Intel® Server System S9200WK product family has support for Closed Loop System Throttling (CLST) which prevents the system from crashing if a power supply module is overloaded or over heats. Should system power reach a pre-programmed power limit, CLST throttles system memory and/or processors to reduce power. System performance is impacted should this occur.

14.7 Smart Ride Through (SmaRT)

The Intel® Server System S9200WK product family supports Smart Ride Through Throttling (SmaRT) which increases the reliability for a system operating in a heavy power load condition and to remain operational during an AC line dropout event.

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.

Appendix A. Getting Help

To obtain support for an issue with the server system, follow these steps:

1. Visit the following Intel support web page: http://www.intel.com/support/

   This web page provides 24x7 support when you need it to get the latest and most complete technical support information on all Intel® Enterprise Server and Storage Platforms. Information available at the support site includes:

   — Latest BIOS, firmware, drivers and utilities

   — Product documentation, setup and service guides

   — Full product specifications, technical advisories and errata

   —Compatibility documentation for memory, hardware add-in cards, and operating systems

   — Server and chassis accessory parts list for ordering upgrades or spare parts

   — A searchable knowledgebase to search for product information throughout the support site

2. If a solution cannot be found at Intel's support site, send an email to Intel's technical support center using the online form available at:

   http://www.intel.com/p/en_US/support/contactsupport

3. Lastly, contact an Intel support representative using one of the support phone numbers available at http://www.intel.com/support/feedback.htm?group=server (charges may apply).

Intel also offers Channel Program members around-the-clock 24x7 technical phone support on Intel® server boards, server chassis, server RAID controller cards, and Intel® Server Management at: http://www.intel.com/reseller/.

Note: Access to the 24x7 number requires a login to the reseller site.

Warranty Information

To obtain warranty information, visit http://www.intel.com/p/en_US/support/warranty.

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 compute modules include a blank of eight POST Code Diagnostic LEDs on the front edge of the compute module server board as shown in the Figure below.

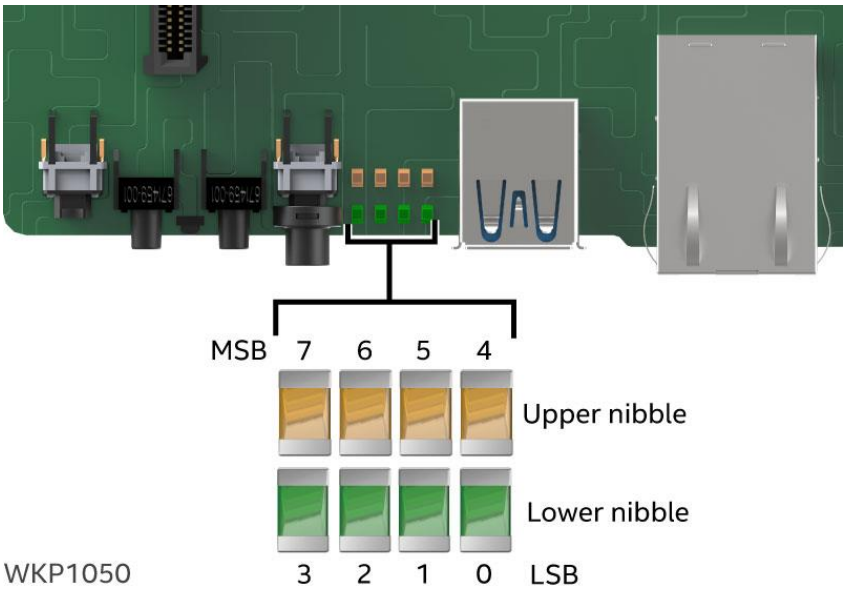

During the system boot process, Memory Reference Code (MRC) and System BIOS execute a number of memory initialization and platform configuration processes, each of which is assigned a specific hex POST code number.

As each routine is started, the given POST code number is displayed to the POST Code Diagnostic LEDs on the front edge of the compute module 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 Diagnostics 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: Diagnostic LEDs are best read and decoded when viewing the LEDs from the front of the system.

In the following example, the BIOS sends a value of ACh to the diagnostic LED decoder. The LEDs are decoded as shown in Table 32.

Table 32. POST progress code LED example

	Upper Nibble AMBER LEDs				Lower Nibble GREEN LEDs			
	MSB							LSB
Binary Value	**1**	**0**	**1**	**0**	**1**	**1**	**0**	**0**
LED State	**ON**	OFF	**ON**	OFF	**ON**	**ON**	OFF	OFF
Hex Value	**8h**	**4h**	**2h**	**1h**	**8h**	**4h**	**2h**	**1h**
Hex Result	**Ah**				**Ch**			

Upper nibble bits = 1010b = Ah; Lower nibble bits = 1100b = Ch; the two are concatenated as **AC**h

Early POST Memory Initialization MRC Diagnostic Codes

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

The MRC Progress Codes are displayed to the Diagnostic LEDs that show the execution point in the MRC operational path at each step.

Table 33. MRC progress codes

Checkpoint	Diagnostic LED Decoder **1** = LED On, **0** = LED Off **Upper Nibble (Amber)**				**Lower Nibble (Green)**				Description
	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, the MRC will display a fatal error code on the diagnostic LEDs, and a system halt command is executed. Fatal MRC error halts do NOT charge the state of the System Status LED, and they do NOT get logged as SEL events. The following table lists all MRC fatal errors that are displayed to the Diagnostic LEDs.

Table 34. MRC fatal error codes

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**	
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

BIOS POST Progress Codes

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

Table 35. POST 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	
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									
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

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	
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
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

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	
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:
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 Display screen, and are always logged to the System Event Log (SEL). Logged events are available to System Management applications, including Remote and Out of Band (OOB) management.

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

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

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

Table 36. POST Error Codes and Messages

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

Error Code	Error Message	Action message	Response
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 Die 0 disabled		Major
8131	Processor 01 Die 1 disabled		Major
8132	Processor 02 Die 0 disabled		Major
8133	Processor 02 Die 1 disabled		Major
8160	Processor 01 Die 0 unable to apply microcode update		Major
8161	Processor 01 Die 1 unable to apply microcode update		Major
8162	Processor 02 Die 0 unable to apply microcode update		Major
8163	Processor 02 Die 1 unable to apply microcode update		Major
8170	Processor 01 Die 0 failed Self Test (BIST)		Major
8171	Processor 01 Die 1 failed Self Test (BIST)		Major
8172	Processor 02 Die 0 failed Self Test (BIST)		Major
8173	Processor 02 Die 1 failed Self Test (BIST)		Major
8180	Processor 01 Die 0 microcode update not found		Minor
8181	Processor 01 Die 1 microcode update not found		Minor
8182	Processor 02 Die 0 microcode update not found		Minor
8183	Processor 02 Die 1 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	Baseboard Management Controller Sensor Data Record empty.	Please update right SDR.	Major
84FF	System Event Log full	Please clear SEL through EWS or SELVIEW utility.	Minor
85FC	Memory component could not be configured in the selected RAS mode		Major
8501	Memory Population Error	Please plug DIMM at right population.	Major
8520	Memory failed test/initialization CPU0_0_DIMM_A1	please remove the disabled DIMM.	Major
8521	Memory failed test/initialization CPU0_0_DIMM_A2	please remove the disabled DIMM.	Major
8522	Memory failed test/initialization CPU0_0_DIMM_A3	please remove the disabled DIMM.	Major
8523	Memory failed test/initialization CPU0_0_DIMM_B1	please remove the disabled DIMM.	Major
8524	Memory failed test/initialization CPU0_0_DIMM_B2	please remove the disabled DIMM.	Major
8525	Memory failed test/initialization CPU0_0_DIMM_B3	please remove the disabled DIMM.	Major
8526	Memory failed test/initialization CPU0_0_DIMM_C1	please remove the disabled DIMM.	Major
8527	Memory failed test/initialization CPU0_0_DIMM_C2	please remove the disabled DIMM.	Major
8528	Memory failed test/initialization CPU0_0_DIMM_C3	please remove the disabled DIMM.	Major
8529	Memory failed test/initialization CPU0_0_DIMM_D1	please remove the disabled DIMM.	Major
852A	Memory failed test/initialization CPU0_0_DIMM_D2	please remove the disabled DIMM.	Major
852B	Memory failed test/initialization CPU0_0_DIMM_D3	please remove the disabled DIMM.	Major
852C	Memory failed test/initialization CPU0_0_DIMM_E1	please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
852D	Memory failed test/initialization CPU0_0_DIMM_E2	please remove the disabled DIMM.	Major
852E	Memory failed test/initialization CPU0_0_DIMM_E3	please remove the disabled DIMM.	Major
852F	Memory failed test/initialization CPU0_0_DIMM_F1	please remove the disabled DIMM.	Major
8530	Memory failed test/initialization CPU0_0_DIMM_F2	please remove the disabled DIMM.	Major
8531	Memory failed test/initialization CPU0_0_DIMM_F3	please remove the disabled DIMM.	Major
8532	Memory failed test/initialization CPU0_0_DIMM_G1	please remove the disabled DIMM.	Major
8533	Memory failed test/initialization CPU0_0_DIMM_G2	please remove the disabled DIMM.	Major
8534	Memory failed test/initialization CPU0_0_DIMM_G3	please remove the disabled DIMM.	Major
8535	Memory failed test/initialization CPU0_0_DIMM_H1	please remove the disabled DIMM.	Major
8536	Memory failed test/initialization CPU0_0_DIMM_H2	please remove the disabled DIMM.	Major
8537	Memory failed test/initialization CPU0_0_DIMM_H3	please remove the disabled DIMM.	Major
8538	Memory failed test/initialization CPU0_1_DIMM_A1	please remove the disabled DIMM.	Major
8539	Memory failed test/initialization CPU0_1_DIMM_A2	please remove the disabled DIMM.	Major
853A	Memory failed test/initialization CPU0_1_DIMM_A3	please remove the disabled DIMM.	Major
853B	Memory failed test/initialization CPU0_1_DIMM_B1	please remove the disabled DIMM.	Major
853C	Memory failed test/initialization CPU0_1_DIMM_B2	please remove the disabled DIMM.	Major
853D	Memory failed test/initialization CPU0_1_DIMM_B3	please remove the disabled DIMM.	Major
853E	Memory failed test/initialization CPU0_1_DIMM_C1	please remove the disabled DIMM.	Major
853F (Go to 85C0)	Memory failed test/initialization CPU0_1_DIMM_C2	please remove the disabled DIMM.	Major
8540	Memory disabled.CPU0_0_DIMM_A1	please remove the disabled DIMM.	Major
8541	Memory disabled.CPU0_0_DIMM_A2	please remove the disabled DIMM.	Major
8542	Memory disabled.CPU0_0_DIMM_A3	please remove the disabled DIMM.	Major
8543	Memory disabled.CPU0_0_DIMM_B1	please remove the disabled DIMM.	Major
8544	Memory disabled.CPU0_0_DIMM_B2	please remove the disabled DIMM.	Major
8545	Memory disabled.CPU0_0_DIMM_B3	please remove the disabled DIMM.	Major
8546	Memory disabled.CPU0_0_DIMM_C1	please remove the disabled DIMM.	Major
8547	Memory disabled.CPU0_0_DIMM_C2	please remove the disabled DIMM.	Major
8548	Memory disabled.CPU0_0_DIMM_C3	please remove the disabled DIMM.	Major
8549	Memory disabled.CPU0_0_DIMM_D1	please remove the disabled DIMM.	Major
854A	Memory disabled.CPU0_0_DIMM_D2	please remove the disabled DIMM.	Major
854B	Memory disabled.CPU0_0_DIMM_D3	please remove the disabled DIMM.	Major
854C	Memory disabled.CPU0_0_DIMM_E1	please remove the disabled DIMM.	Major
854D	Memory disabled.CPU0_0_DIMM_E2	please remove the disabled DIMM.	Major
854E	Memory disabled.CPU0_0_DIMM_E3	please remove the disabled DIMM.	Major
854F	Memory disabled.CPU0_0_DIMM_F1	please remove the disabled DIMM.	Major
8550	Memory disabled.CPU0_0_DIMM_F2	please remove the disabled DIMM.	Major
8551	Memory disabled.CPU0_0_DIMM_F3	please remove the disabled DIMM.	Major
8552	Memory disabled.CPU0_0_DIMM_G1	please remove the disabled DIMM.	Major
8553	Memory disabled.CPU0_0_DIMM_G2	please remove the disabled DIMM.	Major
8554	Memory disabled.CPU0_0_DIMM_G3	please remove the disabled DIMM.	Major
8555	Memory disabled.CPU0_0_DIMM_H1	please remove the disabled DIMM.	Major
8556	Memory disabled.CPU0_0_DIMM_H2	please remove the disabled DIMM.	Major
8557	Memory disabled.CPU0_0_DIMM_H3	please remove the disabled DIMM.	Major
8558	Memory disabled.CPU0_1_DIMM_A1	please remove the disabled DIMM.	Major
8559	Memory disabled.CPU0_1_DIMM_A2	please remove the disabled DIMM.	Major
855A	Memory disabled.CPU0_1_DIMM_A3	please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
855B	Memory disabled.CPU0_1_DIMM_B1	please remove the disabled DIMM.	Major
855C	Memory disabled.CPU0_1_DIMM_B2	please remove the disabled DIMM.	Major
855D	Memory disabled.CPU0_1_DIMM_B3	please remove the disabled DIMM.	Major
855E	Memory disabled.CPU0_1_DIMM_C1	please remove the disabled DIMM.	Major
855F (Go to 85D0)	Memory disabled.CPU0_1_DIMM_C2	please remove the disabled DIMM.	Major
8560	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_A1		Major
8561	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_A2		Major
8562	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_A3		Major
8563	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_B1		Major
8564	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_B2		Major
8565	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_B3		Major
8566	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_C1		Major
8567	Memory encountered a Serial Presence Detection (SPD) failure.CPU0_0_DIMM_C2		Major
8568	Memory encountered a Serial Presence Detection (SPD) failure.CPU0_0_DIMM_C3		Major
8569	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_D1		Major
856A	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_D2		Major
856B	Memoryencountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_D3		Major
856C	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_E1		Major
856D	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_E2		Major
856E	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_E3		Major
856F	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_F1		Major
8570	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_F2		Major
8571	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_F3		Major
8572	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_G1		Major
8573	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_G2		Major
8574	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_G3		Major
8575	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_H1		Major
8576	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_H2		Major

Error Code	Error Message	Action message	Response
8577	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_0_DIMM_H3		Major
8578	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_A1		Major
8579	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_A2		Major
857A	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_A3		Major
857B	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_B1		Major
857C	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_B2		Major
857D	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_B3		Major
857E	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_C1		Major
857F (Go to 85E0)	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_C2		Major
85C0	Memory failed test/initialization CPU0_1_DIMM_C3	please remove the disabled DIMM.	Major
85C1	Memory failed test/initialization CPU0_1_DIMM_D1	please remove the disabled DIMM.	Major
85C2	Memory failed test/initialization CPU0_1_DIMM_D2	please remove the disabled DIMM.	Major
85C3	Memory failed test/initialization CPU0_1_DIMM_D3	please remove the disabled DIMM.	Major
85C4	Memory failed test/initialization CPU0_1_DIMM_E1	please remove the disabled DIMM.	Major
85C5	Memory failed test/initialization CPU0_1_DIMM_E2	please remove the disabled DIMM.	Major
85C6	Memory failed test/initialization CPU0_1_DIMM_E3	please remove the disabled DIMM.	Major
85C7	Memory failed test/initialization CPU0_1_DIMM_F1	please remove the disabled DIMM.	Major
85C8	Memory failed test/initialization CPU0_1_DIMM_F2	please remove the disabled DIMM.	Major
85C9	Memory failed test/initialization CPU0_1_DIMM_F3	please remove the disabled DIMM.	Major
85CA	Memory failed test/initialization CPU0_1_DIMM_G1	please remove the disabled DIMM.	Major
85CB	Memory failed test/initialization CPU0_1_DIMM_G2	please remove the disabled DIMM.	Major
85CC	Memory failed test/initialization CPU0_1_DIMM_G3	please remove the disabled DIMM.	Major
85CD	Memory failed test/initialization CPU0_1_DIMM_H1	please remove the disabled DIMM.	Major
85CE	Memory failed test/initialization CPU0_1_DIMM_H2	please remove the disabled DIMM.	Major
85CF	Memory failed test/initialization CPU0_1_DIMM_H3	please remove the disabled DIMM.	Major
85D0	Memory disabled.CPU0_1_DIMM_C3	please remove the disabled DIMM.	Major
85D1	Memory disabled.CPU0_1_DIMM_D1	please remove the disabled DIMM.	Major
85D2	Memory disabled.CPU0_1_DIMM_D2	please remove the disabled DIMM.	Major
85D3	Memory disabled.CPU0_1_DIMM_D3	please remove the disabled DIMM.	Major
85D4	Memory disabled.CPU0_1_DIMM_E1	please remove the disabled DIMM.	Major
85D5	Memory disabled.CPU0_1_DIMM_E2	please remove the disabled DIMM.	Major
85D6	Memory disabled.CPU0_1_DIMM_E3	please remove the disabled DIMM.	Major
85D7	Memory disabled.CPU0_1_DIMM_F1	please remove the disabled DIMM.	Major
85D8	Memory disabled.CPU0_1_DIMM_F2	please remove the disabled DIMM.	Major
85D9	Memory disabled.CPU0_1_DIMM_F3	please remove the disabled DIMM.	Major
85DA	Memory disabled.CPU0_1_DIMM_G1	please remove the disabled DIMM.	Major
85DB	Memory disabled.CPU0_1_DIMM_G2	please remove the disabled DIMM.	Major
85DC	Memory disabled.CPU0_1_DIMM_G3	please remove the disabled DIMM.	Major
85DD	Memory disabled.CPU0_1_DIMM_H1	please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
85DE	Memory disabled.CPU0_1_DIMM_H2	please remove the disabled DIMM.	Major
85DF	Memory disabled.CPU0_1_DIMM_H3	please remove the disabled DIMM.	Major
85E0	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_C3		Major
85E1	Memory encountered a Serial Presence Detection (SPD) failure. CPU0_1_DIMM_D1		Major
85E2	Memory encountered a Serial Presence Detection (SPD) failure.CPU0_1_DIMM_D2		Major
85E3	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_D3		Major
85E4	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_E1		Major
85E5	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_E2		Major
85E6	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_E3		Major
85E7	Memory encountered a Serial Presence Detection (SPD) failure.CPU0_1_DIMM_F1		Major
85E8	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_F2		Major
85E9	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_F3		Major
85EA	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_G1		Major
85EB	Memory encountered a Serial Presence Detection(SPD) failure. CPU0_1_DIMM_G2		Major
85EC	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_G3		Major
85ED	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_H1		Major
85EE	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_H2		Major
85EF	Memory encountered a Serial Presence Detection(SPD) failure.CPU0_1_DIMM_H3		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
9500	Memory disabled.CPU1_0_DIMM_A1	please remove the disabled DIMM.	Major
9501	Memory disabled.CPU1_0_DIMM_A2	please remove the disabled DIMM.	Major
9502	Memory disabled.CPU1_0_DIMM_A3	please remove the disabled DIMM.	Major
9503	Memory disabled.CPU1_0_DIMM_B1	please remove the disabled DIMM.	Major
9504	Memory disabled.CPU1_0_DIMM_B2	please remove the disabled DIMM.	Major
9505	Memory disabled.CPU1_0_DIMM_B3	please remove the disabled DIMM.	Major
9506	Memory disabled.CPU1_0_DIMM_C1	please remove the disabled DIMM.	Major
9507	Memory disabled.CPU1_0_DIMM_C2	please remove the disabled DIMM.	Major
9508	Memory disabled.CPU1_0_DIMM_C3	please remove the disabled DIMM.	Major
9509	Memory disabled.CPU1_0_DIMM_D1	please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
950A	Memory disabled.CPU1_0_DIMM_D2	please remove the disabled DIMM.	Major
950B	Memory disabled.CPU1_0_DIMM_D3	please remove the disabled DIMM.	Major
950C	Memory disabled.CPU1_0_DIMM_E1	please remove the disabled DIMM.	Major
950D	Memory disabled.CPU1_0_DIMM_E2	please remove the disabled DIMM.	Major
950E	Memory disabled.CPU1_0_DIMM_E3	please remove the disabled DIMM.	Major
950F	Memory disabled.CPU1_0_DIMM_F1	please remove the disabled DIMM.	Major
9510	Memory disabled.CPU1_0_DIMM_F2	please remove the disabled DIMM.	Major
9511	Memory disabled.CPU1_0_DIMM_F3	please remove the disabled DIMM.	Major
9512	Memory disabled.CPU1_0_DIMM_G1	please remove the disabled DIMM.	Major
9513	Memory disabled.CPU1_0_DIMM_G2	please remove the disabled DIMM.	Major
9514	Memory disabled.CPU1_0_DIMM_G3	please remove the disabled DIMM.	Major
9515	Memory disabled.CPU1_0_DIMM_H1	please remove the disabled DIMM.	Major
9516	Memory disabled.CPU1_0_DIMM_H2	please remove the disabled DIMM.	Major
9517	Memory disabled.CPU1_0_DIMM_H3	please remove the disabled DIMM.	Major
9518	Memory disabled.CPU1_1_DIMM_A1	please remove the disabled DIMM.	Major
9519	Memory disabled.CPU1_1_DIMM_A2	please remove the disabled DIMM.	Major
951A	Memory disabled.CPU1_1_DIMM_A3	please remove the disabled DIMM.	Major
951B	Memory disabled.CPU1_1_DIMM_B1	please remove the disabled DIMM.	Major
951C	Memory disabled.CPU1_1_DIMM_B2	please remove the disabled DIMM.	Major
951D	Memory disabled.CPU1_1_DIMM_B3	please remove the disabled DIMM.	Major
951E	Memory disabled.CPU1_1_DIMM_C1	please remove the disabled DIMM.	Major
951F	Memory disabled.CPU1_1_DIMM_C2	please remove the disabled DIMM.	Major
9520	Memory disabled.CPU1_1_DIMM_C3	please remove the disabled DIMM.	Major
9521	Memory disabled.CPU1_1_DIMM_D1	please remove the disabled DIMM.	Major
9522	Memory disabled.CPU1_1_DIMM_D2	please remove the disabled DIMM.	Major
9523	Memory disabled.CPU1_1_DIMM_D3	please remove the disabled DIMM.	Major
9524	Memory disabled.CPU1_1_DIMM_E1	please remove the disabled DIMM.	Major
9525	Memory disabled.CPU1_1_DIMM_E2	please remove the disabled DIMM.	Major
9526	Memory disabled.CPU1_1_DIMM_E3	please remove the disabled DIMM.	Major
9527	Memory disabled.CPU1_1_DIMM_F1	please remove the disabled DIMM.	Major
9528	Memory disabled.CPU1_1_DIMM_F2	please remove the disabled DIMM.	Major
9529	Memory disabled.CPU1_1_DIMM_F3	please remove the disabled DIMM.	Major
952A	Memory disabled.CPU1_1_DIMM_G1	please remove the disabled DIMM.	Major
952B	Memory disabled.CPU1_1_DIMM_G2	please remove the disabled DIMM.	Major
952C	Memory disabled.CPU1_1_DIMM_G3	please remove the disabled DIMM.	Major
952D	Memory disabled.CPU1_1_DIMM_H1	please remove the disabled DIMM.	Major
952E	Memory disabled.CPU1_1_DIMM_H2	please remove the disabled DIMM.	Major
952F	Memory disabled.CPU1_1_DIMM_H3	please remove the disabled DIMM.	Major
9530	Memory failed test/initialization CPU1_0_DIMM_A1	please remove the disabled DIMM.	Major
9531	Memory failed test/initialization CPU1_0_DIMM_A2	please remove the disabled DIMM.	Major
9532	Memory failed test/initialization CPU1_0_DIMM_A3	please remove the disabled DIMM.	Major
9533	Memory failed test/initialization CPU1_0_DIMM_B1	please remove the disabled DIMM.	Major
9534	Memory failed test/initialization CPU1_0_DIMM_B2	please remove the disabled DIMM.	Major
9535	Memory failed test/initialization CPU1_0_DIMM_B3	please remove the disabled DIMM.	Major
9536	Memory failed test/initialization CPU1_0_DIMM_C1	please remove the disabled DIMM.	Major
9537	Memory failed test/initialization CPU1_0_DIMM_C2	please remove the disabled DIMM.	Major

Error Code	Error Message	Action message	Response
9538	Memory failed test/initialization CPU1_0_DIMM_C3	please remove the disabled DIMM.	Major
9539	Memory failed test/initialization CPU1_0_DIMM_D1	please remove the disabled DIMM.	Major
953A	Memory failed test/initialization CPU1_0_DIMM_D2	please remove the disabled DIMM.	Major
953B	Memory failed test/initialization CPU1_0_DIMM_D3	please remove the disabled DIMM.	Major
953C	Memory failed test/initialization CPU1_0_DIMM_E1	please remove the disabled DIMM.	Major
953D	Memory failed test/initialization CPU1_0_DIMM_E2	please remove the disabled DIMM.	Major
953E	Memory failed test/initialization CPU1_0_DIMM_E3	please remove the disabled DIMM.	Major
953F	Memory failed test/initialization CPU1_0_DIMM_F1	please remove the disabled DIMM.	Major
9540	Memory failed test/initialization CPU1_0_DIMM_F2	please remove the disabled DIMM.	Major
9541	Memory failed test/initialization CPU1_0_DIMM_F3	please remove the disabled DIMM.	Major
9542	Memory failed test/initialization CPU1_0_DIMM_G1	please remove the disabled DIMM.	Major
9543	Memory failed test/initialization CPU1_0_DIMM_G2	please remove the disabled DIMM.	Major
9544	Memory failed test/initialization CPU1_0_DIMM_G3	please remove the disabled DIMM.	Major
9545	Memory failed test/initialization CPU1_0_DIMM_H1	please remove the disabled DIMM.	Major
9546	Memory failed test/initialization CPU1_0_DIMM_H2	please remove the disabled DIMM.	Major
9547	Memory failed test/initialization CPU1_0_DIMM_H3	please remove the disabled DIMM.	Major
9548	Memory failed test/initialization CPU1_1_DIMM_A1	please remove the disabled DIMM.	Major
9549	Memory failed test/initialization CPU1_1_DIMM_A2	please remove the disabled DIMM.	Major
954A	Memory failed test/initialization CPU1_1_DIMM_A3	please remove the disabled DIMM.	Major
954B	Memory failed test/initialization CPU1_1_DIMM_B1	please remove the disabled DIMM.	Major
954C	Memory failed test/initialization CPU1_1_DIMM_B2	please remove the disabled DIMM.	Major
954D	Memory failed test/initialization CPU1_1_DIMM_B3	please remove the disabled DIMM.	Major
954E	Memory failed test/initialization CPU1_1_DIMM_C1	please remove the disabled DIMM.	Major
954F	Memory failed test/initialization CPU1_1_DIMM_C2	please remove the disabled DIMM.	Major
9550	Memory failed test/initialization CPU1_1_DIMM_C3	please remove the disabled DIMM.	Major
9551	Memory failed test/initialization CPU1_1_DIMM_D1	please remove the disabled DIMM.	Major
9552	Memory failed test/initialization CPU1_1_DIMM_D2	please remove the disabled DIMM.	Major
9553	Memory failed test/initialization CPU1_1_DIMM_D3	please remove the disabled DIMM.	Major
9554	Memory failed test/initialization CPU1_1_DIMM_E1	please remove the disabled DIMM.	Major
9555	Memory failed test/initialization CPU1_1_DIMM_E2	please remove the disabled DIMM.	Major
9556	Memory failed test/initialization CPU1_1_DIMM_E3	please remove the disabled DIMM.	Major
9557	Memory failed test/initialization CPU1_1_DIMM_F1	please remove the disabled DIMM.	Major
9558	Memory failed test/initialization CPU1_1_DIMM_F2	please remove the disabled DIMM.	Major
9559	Memory failed test/initialization CPU1_1_DIMM_F3	please remove the disabled DIMM.	Major
955A	Memory failed test/initialization CPU1_1_DIMM_G1	please remove the disabled DIMM.	Major
955B	Memory failed test/initialization CPU1_1_DIMM_G2	please remove the disabled DIMM.	Major
955C	Memory failed test/initialization CPU1_1_DIMM_G3	please remove the disabled DIMM.	Major
955D	Memory failed test/initialization CPU1_1_DIMM_H1	please remove the disabled DIMM.	Major
955E	Memory failed test/initialization CPU1_1_DIMM_H2	please remove the disabled DIMM.	Major
955F	Memory failed test/initialization CPU1_1_DIMM_H3	please remove the disabled DIMM.	Major
9560	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_A1		Major
9561	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_A2		Major
9562	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_A3		Major

Error Code	Error Message	Action message	Response
9563	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_B1		Major
9564	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_B2		Major
9565	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_B3		Major
9566	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_C1		Major
9567	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_C2		Major
9568	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_C3		Major
9569	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_D1		Major
956A	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_D2		Major
956B	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_D3		Major
956C	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_E1		Major
956D	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_E2		Major
956E	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_E3		Major
956F	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_F1		Major
9570	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_F2		Major
9571	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_F3		Major
9572	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_G1		Major
9573	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_G2		Major
9574	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_G3		Major
9575	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_H1		Major
9576	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_H2		Major
9577	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_0_DIMM_H3		Major
9578	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_A1		Major
9579	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_A2		Major
957A	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_A3		Major
957B	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_B1		Major
957C	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_B2		Major

Error Code	Error Message	Action message	Response
957D	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_B3		Major
957E	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_C1		Major
957F	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_C2		Major
9580	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_C3		Major
9581	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_D1		Major
9582	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_D2		Major
9583	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_D3		Major
9584	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_E1		Major
9585	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_E2		Major
9586	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_E3		Major
9587	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_F1		Major
9588	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_F2		Major
9589	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_F3		Major
958A	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_G1		Major
958B	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_G2		Major
958C	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_G3		Major
958D	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_H1		Major
958E	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_H2		Major
958F	Memory encountered a Serial Presence Detection(SPD) failure. CPU1_1_DIMM_H3		Major
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

Appendix D. System Configuration Table for Thermal Compatibility

This section provides a table listing system configuration compatibility data based on various supported system operating thermal limits. The table identifies supported system configurations while the system is in “normal” operating mode, meaning all systems fans are present, on-line, and operational.

The following notes apply to the information presented in the table:

1. 1.. System cooling capability testing was carried out in environmental lab controlled conditions according to ASHRAE standard.

2. The system performance depends on data center environmental temperature and relative humidity levels controls provided by end user.

3. Liquid minimum flowrates depend on supply temperature referred in the chart.

4. CPU and memory loss of performance – throttling is expected. Throttling mechanism caused by the CPU and memories itself, when these components reach the max temperature (from spec) they will start doing low level throttling - watt per watt until the temperature become stable below the threshold (5, 10, 15% of throttling etc.).

5. For combined configurations:

- 1U/2U AC assumes 1U AC nodes with 250 W TDP and 2U AC nodes with 350 W TDP.
- 1U/2U LC assumes 1U LC nodes with 400 W TDP and 2U LC nodes with 400 W TDP.

6. PSU reaching 61 degrees Centigrade, self-protect mechanism activated causing loss of performance (throttling) in the system. There is a self-protect mechanism in the PSU itself called CLST (Closed Loop System Throttling), when the sensor inside the PSU reach this temperature it will cause an aggressive throttling to the memories and CPUs in order to reduce the quantity of power and preheating through the unit. Throttling expected is more than 50% of the stress power.

7. U.2 supported just in the 2U LC or AC node SKUs.

Thermal Config Matrix ***Normal operating mode*** **"●"= Full Support without limitation;** **" "(Blank)=Not support;**		**1U LC - 4 sleds**		**2U LC - 2 sleds**		**2U AC - 2 sleds**		**1U AC - 4 sleds**		**1U/2U AC-Combined (See note 5)**		**1U/2U LC - Combined (See note 5)**	
ASHRAE (See note1, 2)	Max Ambient	**27°C**	**35°C**	**27°C**	**35°C**	**27°C**	**35°C**	**27°C**	**32°C**	**27°C**	**32°C**	**27°C**	**35°C**
	Altitude	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**
PSU	1600 W AC	●	●	●	●	●	●					●	●
	2100 W AC	●	●	●	●	●	●	●	6	●	6	●	●
CPU	400 W	●	●	●	●							●	●
	350 W	●	●	●	●	●	●					●	●
	250 W	●	●	●	●	●	●	●	●	●	●	●	●
Memory	8GB RDIMM Crucial 2933 1Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	8GB RDIMM Micron 3200 1Gb x 8	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Crucial 2933 2Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Crucial 2933 1Rx4	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Micron 2933 2Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Micron 2933 1Rx4	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Samsung 2933 1Rx4	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Samsung 2933 2Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM SK Hynix 2933 2Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	16GB RDIMM Micron 3200 2Gb x 4	●	●	●	●	●	●	●	6	●	6	●	●
	16GB RDIMM Micron 3200 1Gb x 8	●	●	●	●	●	●	●	●	●	●	●	●
	32GB RDIMM Crucial 2933 1Rx8	●	●	●	●	●	●	●	●	●	●	●	●
	32GB RDIMM Micron 2933 2Rx4	●	●	●	●	●	●	●	6	●	6	●	●

Thermal Config Matrix ***Normal operating mode*** **"●"= Full Support without limitation;** **" "(Blank)=Not support;**		**1U LC - 4 sleds**		**2U LC - 2 sleds**		**2U AC - 2 sleds**		**1U AC - 4 sleds**		**1U/2U AC-Combined (See note 5)**		**1U/2U LC - Combined (See note 5)**	
ASHRAE (See note1, 2)	Max Ambient	**27°C**	**35°C**	**27°C**	**35°C**	**27°C**	**35°C**	**27°C**	**32°C**	**27°C**	**32°C**	**27°C**	**35°C**
	Altitude	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**	**900**
	32GB RDIMM Samsung 2933 2Rx4	●	●	●	●	●	●	●	6	●	6	●	●
	32GB RDIMM Micron 3200 2Gb x 4	●	●	●	●	●	●	●	6	●	6	●	●
	64GB LRDIMM Micron 2933 4Rx4	●	●	●	●	●	●	6	6	●	6	●	●
	64GB LRDIMM Samsung 2933 4Rx4	●	●	●	●	●	●	●	6	●	6	●	●
	64GB LRDIMM Crucial 2933 4Rx4	●	●	●	●	●	●	6	6	●	6	●	●
	128GB LRDIMM Samsung 2933 4Rx4	●	●	●	●	●	●	●	6	●	6	●	●
Add-in Cards - Riser #1 and #2	Up to 20 W	●	●	●	●	●	●	●		●		●	●
U.2 - Riser #1 and #2. **Only supported in 2U form factor**	Up to 25 W	7	7	●	●	●	●	7	7	●	●	●	●
M.2- Riser #1 and #2	Up to 6 W	●	●	●	●	●	●	●	●	●	●	●	●

System Coolant Flow Rate Requirements

For liquid cooled system configurations, a minimum coolant flow rate is required depending on the coolant temperature. Figure 65 shows the required coolant flow rates per coolant temperature according to different CPU and memory configurations. The flow rate is measured in liters per minute with a maximum supported flow rate of 4.6 liters/minute and the temperature is measured in celsius degrees with a maximum supported temperature of 45 degrees. The different flow rates in Figure 65 represent the minimum required conditions to keep the system operating within supported maximum thermal limits. Flow rates below the ones listed are not supported and may cause system throttling.

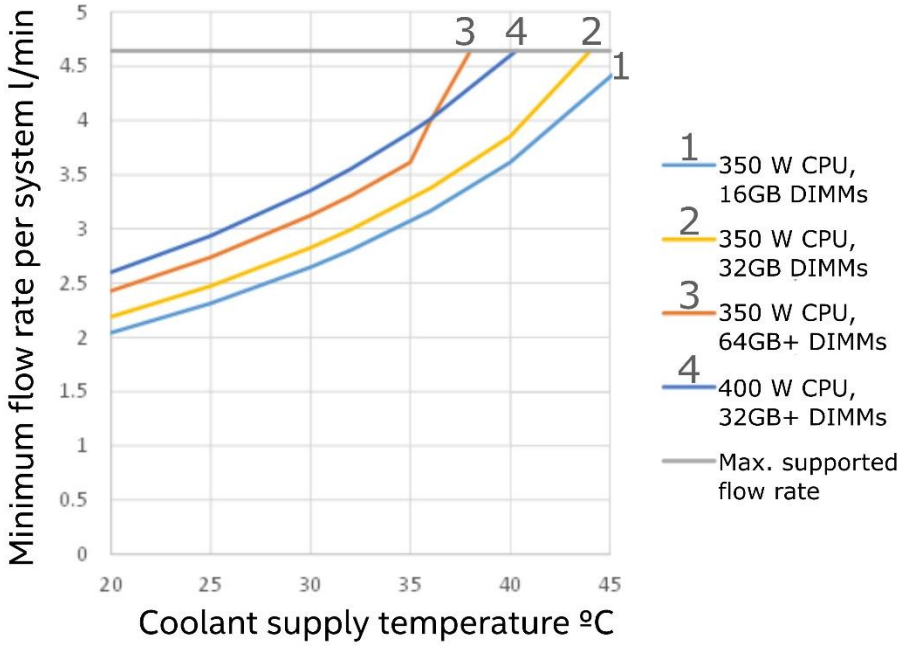

Figure 65. System coolant flow rate requirements

Appendix E. Product Regulatory Information

This product has been evaluated and certified as Information Technology Equipment (ITE), which may be installed in offices, schools, computer rooms, and similar commercial type locations. The suitability of this product for other product certification categories and/or environments (such as: medical, industrial, telecommunications, NEBS, residential, alarm systems, test equipment, etc.), other than an ITE application, will require further evaluation and may require additional regulatory approvals.

Intel has verified that all L3, L6, and L9 server products[1] **as configured and sold by Intel** to its customers comply with the requirements for all regulatory certifications defined in the following table. It is the Intel customer's responsibility to ensure their final server system configurations are tested and certified to meet the regulatory requirements for the countries to which they plan to ship and or deploy server systems into.

	Intel® Server S9200WK Family	**NOTES**
	"Walker Pass"	Intel Project Code Name
	L9 System	Product Integration Level
	FC2H	Product family identified on certification
Regulatory Certification		
RCM DoC Australia & New Zealand	✓	
CB Certification & Report (International - report to include all CB country national deviations)	✓	
China CCC Certification	○	
CU Certification (Russia / Belarus / Kazakhstan)	✓	
Europe CE Declaration of Conformity	✓	
FCC Part 15 Emissions Verification (USA & Canada)	✓	
Germany GS Certification	✓	
India BIS Certification	●	Only applicable to select OEM defined SKUs
International Compliance – CISPR32 & CISPR24	✓	
Japan VCCI Certification	✓	
Korea KC Certification	✓	
Mexico Certification	✓	
NRTL Certification (USA & Canada)	✓	
South Africa Certification	✓	
Taiwan BSMI Certification	✓	
Ukraine Certification	✓	

Table Key

Not Tested / Not Certified	○
Tested / Certified – Limited OEM SKUs only	●
Testing / Certification (Planned)	(Date)
Tested / Certified	✓

[1] An L9 system configuration is a power-on ready server system with NO operating system installed.
An L6 system configuration requires additional components to be installed in order to make it power-on ready. L3 are component building block options that require integration into a chassis to create a functional server system.

EU Directive 2019/424 (Lot 9)

Beginning on March 1, 2020, an additional component of the European Union (EU) regulatory CE marking scheme, identified as EU Directive 2019/424 (Lot 9), will go into effect. After this date, all new server systems shipped into or deployed within the EU must meet the full CE marking requirements including those defined by the additional EU Lot 9 regulations.

Intel has verified that all L3, L6, and L9 server products[2] **as configured and sold by Intel** to its customers comply with the full CE regulatory requirements for the given product type, including those defined by EU Lot 9. It is the Intel customer's responsibility to ensure their final server system configurations are SPEC® SERT™ tested and meet the new CE regulatory requirements.

Visit the following website for additional EU Directive 2019/424 (Lot9) information:
https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019R0424

In compliance with the EU Directive 2019/424 (Lot 9) materials efficiency requirements, Intel makes available all necessary product collaterals as identified below:

- **System Disassembly Instructions**
  - *Intel® Server System S9200WK Product Family Setup and Service Guide*
  - **https://www.intel.com/content/www/us/en/support/articles/000055308/server-products/server-systems.html?productId=192334&localeCode=us_en**

- **Product Specifications**
  - *Intel® Server System S9200WK Product Family Technical Product Specification (This document)*
  - **https://www.intel.com/content/www/us/en/support/articles/000055047/server-products/server-systems.html?productId=192334&localeCode=us_en**

- **System BIOS/Firmware and Security Updates**
  - System Update Package (SUP) – uEFI only
  - Intel® One Boot Flash Update (OFU) – Various OS Support
  - https://www.intel.com/content/www/us/en/support/products/192334/server-products/server-systems/intel-server-system-s9200wk-family.html

- **Intel® Solid State Drive (SSD) Secure Data Deletion and Firmware Updates**
  - Note: for system configurations that may be configured with an Intel SSD
  - Intel® Solid State Drive Toolbox
  - https://downloadcenter.intel.com/download/29205?v=t

- **Intel® RAID Controller Firmware Updates and other support collaterals**
  - Note: for system configurations that may be configured with an Intel® RAID Controller
  - https://www.intel.com/content/www/us/en/support/products/43732/server-products/raid-products.html

[2] An L9 system configuration is a power-on ready server system with NO operating system installed.
An L6 system configuration requires additional components to be installed in order to make it power-on ready. L3 are component building block options that require integration into a chassis to create a functional server system

EU Lot9 Support Summary – FC2H – Intel® Server System S9200WK Family ("Walker Pass")

DISCLAIMER *– The information contained within the following tables is for reference purposes only and is intended to provide Intel customers with a template to report product information necessary for (EU) 2019/424 (Lot 9) server conformity assessment. The information provided herein does not represent any final shipping server system test results, and customer's actual test results for shipping server configurations may differ from the information provided. Use of this information is at the sole risk of the user, and Intel assumes no responsibility for customer's server system level regulation compliance to EU 2019/424 (Lot 9).*

<table>
<tr><th colspan="2">Product Info.</th></tr>
<tr><td>Product Type</td><td>Server</td></tr>
<tr><td>Manufacturer Name</td><td>Intel Corporation</td></tr>
<tr><td>Registered trade name and address</td><td>Intel
2200 Mission College Blvd
Santa Clara, CA 95054-1594, USA</td></tr>
<tr><td>Product model number and model numbers for low end performance and high-end performance configure if applicable</td><td>FC2H</td></tr>
<tr><td>Product Launch Date</td><td>2019</td></tr>
<tr><td>Test System Manufacture Date</td><td>2019</td></tr>
<tr><td>Power Supply options efficiency at 10%, 20%, 50% and 100% of rated output power</td><td>AXX1600TCRPS (PSSF162205A) – 1600W AC (80+ Titanium)
FCXX2100CRPS (PSSF212201A) – 2100W AC (80+ Platinum)

<table><tr><td>Model</td><td>10%</td><td>20%</td><td>50%</td><td>100%</td></tr><tr><td>PSSF162205A</td><td>90.95%</td><td>94.57%</td><td>96.25%</td><td>95.15%</td></tr><tr><td>PSSF212201A</td><td>90.95%</td><td>93.90%</td><td>94.89%</td><td>93.31%</td></tr></table></td></tr>
<tr><td>PSU factor at 50% of rated load level</td><td>AXX1600TCRPS 1600W AC: 0.99
FCXX2100CRPS 2100W AC: 1.00</td></tr>
<tr><td>PSU Rated Power Output
(Server Only)</td><td>AXX1600TCRPS: 1600W
FCXX2100CRPS: 2100W</td></tr>
<tr><td>Idle state power (Watts) - Server only</td><td>Refer to the following table</td></tr>
<tr><td>List of all components for additional idle power allowances (Server only)</td><td>Refer to the following table</td></tr>
<tr><td>Maximum power (Server only)</td><td>Refer to the following table</td></tr>
<tr><td>Declared operating condition class</td><td>Class A2 – Continuous Operation. 10℃ to 35℃ with the maximum rate of change not to exceed 10℃ per hour.</td></tr>
<tr><td>Idle State Power (watts) at the higher boundary temp (Server Only)</td><td>Refer to the following table</td></tr>
<tr><td>the active state efficiency and the performance in active state of the server (server only)</td><td>Refer to the following table</td></tr>
<tr><td>Information on the secure data deletion functionality</td><td>Refer to the following table</td></tr>
<tr><td>for blade server, a list of recommended combinations with compatible chassis (Server only)</td><td>Not Applicable</td></tr>
<tr><td>If Product Model Is Part Of A Server Product Family, a list of all model configurations that are represented by the model shall be supplied (Server only)</td><td>Not Applicable</td></tr>
</table>

Energy Efficiency Data of FC2H – 2 (Dual) CPUs Installed Configuration

Test Configurations				
			2 CPUs Low-end Config.	**2 CPUs High-end Config.**
Details	Chassis	Model	FC2H	FC2H
	Node / Motherboard (MB)	# of nodes installed in system	**2 Nodes**	**4 Nodes**
		Model	**2U 32C Air Cooling**	**1U 56C Liquid Cooling**
	Processor	# of Processors per node / MB	**2 per node**	**2 per node**
		Processor Model	**Intel® Xeon® Scalable Platinum 9221**	**Intel® Xeon® Scalable Platinum 9282**
	Memory	# of DIMMs installed per node	**24 DIMMs per Node** (1 DIMM / mem. channel)	**24 DIMMS per Node** (1 DIMM / mem. channel)
		Capacity per DIMM (GB)	**8 GB**	**32 GB**
		Total Memory (GB) per node	**192 GB per node**	**768 GB per node**
	SSD	Total # of SSDs installed	**4**	**8**
	Power Supply (PSU)	Total # of PSU installed	**3**	**3**
		Model	**AXX1600TCRPS**	**FCXX2100CRPS**
	System Software Revisions installed to each Node		BIOS: **22010091** BMC: **2.30** FRU / SDR: **0.23**	
Data Summary				
Measured and Calculated Server Allowance	P Base		**40**	**40**
	Additional CPU		**134.69**	**235.32**
	Additional Power Supply		**10**	**10**
	Storage Devices		**10**	**10**
	Additional Memory		**33.84**	**137.52**
	Additional I/O Device (10Gx 15W/2Port on MB)		**0**	**0**
	$Perf_{cpu}$		**19.24**	**33.62**
Limits/ Results	Idle power allowances (Watts)		**228.5**	**432.8**
	Idle power tested (Watts Per node)		**163.7**	**165.8**
	Minimum Eff_{ACTIVE}		**8**	**8**
	Eff_{ACTIVE} tested		**31**	**35**
Other test results	Idle Power at Higher Temp. (per Node) @ 35 degree C		**173.5**	**180.4**
	Max Power (Watts Per Node)		**371.95**	**320.30**

Other Information:

Chemical Declaration

- Neodymium Not Applicable. (No HDD offered by Intel)
- Cobalt Not Applicable. (No BBUs. Coin battery is out of scope)

Appendix F. *Product Family Statements of Volatility*

This Appendix describes the volatile and non-volatile components on the Intel® Server System S9200WK product family. It is not the intention of this document to include any components not directly mounted to the server board in the compute modules or riser cards used within the compute modules or supported Intel chassis, these may include: processors, memory, storage devices, or add-in cards.

The tables in this appendix provide the following data for each identified component.

Component Type

Three types of memory components are used on the compute module 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 silkscreen in the compute module server board.

User Data

The flash components on the server boards within the compute modules 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 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 compute modules 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® compute module. 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.

The server board within the Intel® compute modules includes several components that can be used to store data. A list of those components is included in Table 37.

Table 37. Server board components

Component Type	Size	Board Location	User Data	Name
Non-Volatile	256Mb	U111X1	No (BIOS)	BIOS Flash
Non-Volatile	256Mb	U62X1	No (FW)	BMC FW
Non-Volatile	16Mb	U25X1	No	LAN Flash
Non-Volatile	768 B	EU63X1	No (TPM)	TPM
Non-Volatile		U37X1	No	FPGA
Volatile	4Gb	E29X1	No	BMC SDRAM

System boards within Intel® compute modules may include components that can be used to store data. The following tables provide a list of components associated with specific system boards supported by this product family. For server board components, see the section above.

Table 38. System Boards Components

Component Type	Size	Board Location	User Data	Name
Non-Volatile	2k Bytes	U1	No	1U riser card FRU
Non-Volatile	2k Bytes	U1	No	2U riser card FRU

System boards within Intel® server chassis contain components that can be used to store data. A list of components for the system boards in the chassis is included in the table below. For server board components and compute module components, see the sections above.

Table 39. Server Chassis Components

Component Type	Size	Board Location	User Data	Name
Non-Volatile	2k Bytes	U25	No	Power distribution board FRU
Non-Volatile	128k Bytes	U30	No	Power distribution board PIC MCU

Appendix G. Glossary

Term	Definition
1U	One rack unit (1.75 in.)
2U	Two rack units (3.5 in.)
AC	Alternating current
ACM	Authenticating code mode
ACPI	Advanced configuration and power interface
AES-NI	Intel® Advanced Encryption Standard New Instructions
BBS	BIOS boot specification
BDS	Boot device selection
BIOS	Basic input output system; non-volatile firmware
BIST	Built-in self test
BMC	Baseboard management controller
BSP	A boot strap processor that is selected at boot time to be the primary processor in a multi-processor system.
CDP	Code data prioritization
Chassis	Casing containing the server compute modules and fans/liquid cooling plumbing
CLTT	Closed loop thermal monitoring
CMOS	Complimentary Metal-oxide-semiconductor
CMT	Cache monitoring technology
CPU	Central processing unit
CRAM	Configuration RAM. A programmable bit inside an FPGA that controls its behavior.
CSM	Compatibility support module
DC	Direct current, the flow of electric charge in only one direction.
DIMM	Dual inline memory module
DQ	Data quality
DQS	Bi-directional data strobe
DXE	Driver execution environment, a component of Intel® Platform Innovation Framework for EFI architecture.
ECC	Error correction code
EI	Enhanced Intel
EMP module	Ethernet Management Port module
ESD	Electrostatic discharge
FRB	Fault-resilient booting
FRB-2	Fault-resistant booting 2
FRU	Field replaceable unit
GT/s	Gigatransfers per second
HPC	High performance computing
IIO	Integrated IO modules
IMC	Integrated memory controller, a memory controller integrated into the processor chip
IPL	Initial program load
IPMB	Intelligent platform management bus
ISTA	International Safe Transit Association
JRE	Java runtime environment
JEDEC	Joint Electron Device Engineering Council, an industry organization for memory standards
KVM	Keyboard, video, and mouse, an attachment that mimics those devices and connects them to a remote I/O user.
LAN	Local area network
LED	Light emitting diode
LRDIMM	Load reduced dim memory modules have buffer registes for both the address and data between the SDRAM modules and the system's memory controller.
MAC	Media access control
MBA	Memory bandwith allocation
MBM	Memory bandwith monitoring

Intel® ME	Intel® Management Engine
MRC	Memory Reference Code
MSB	Most significant bit
MSID	CPU ICC mismatch
MT/s	Megatransfers per second
NIC	Network interface card
NDA	Non-disclosure agreement
Intel® NM	Intel® Intelligent Power Node Manager
NMI	Non-maskable interrupt
NTB	Non-transparent bridge
NVMe*	Non-volatile memory express
NVRAM	Non-volatile random access memory
OEM	Original equipment manufacturer
OFU	One-boot flash update
OLTT	Open loop thermal throttling
OOB	Out of band
OR	Oct-rank, memory DIMM organization in eight ranks
OS	Operating system
PCH	Platform controller hub
PCI	Periphera component interconnect, or CI local bus standard a.k.a. “Conventional PCI”
PCIe*	Peripheral component interconnect express, an updated form of PCI offering better throughput and error management
PEI	Pre EFI initialization, a component of the Intel® Platform Innovation Framework for EFI architecture
PEIM	PEI module
PERR	Parity error
PHY	Physical port
POST	Power on self test, the BIOS activity from Power On until the OS bot starts
PSU	Power supply unit
QR	Quad rank, the memory DIMM organization in four ranks
Rack	Casing containing one or multiple chassis
RAID	Redundant Array of Inexpensive Disks. These provide data security by spreading data over multiple disk drives. RAID 0, RAID 1, RAID 10, and RAID 5 are different patterns of data on varying numbers of disks to provide varying degrees of security and performance.
RAM	Random access memory
RAS	Reliability, availability, and serviceability
RDIMM	Registered DIMM (also referred to as buffered) memory modules have an address buffer register between the SDRAM modules and the system’s memory controller.
ROM	Read only memory
Intel® RMM4	Intel® Remote Management Module 4
RT	Runtime, a component of Intel® Platform Innovation Framework for EFI architecture
RTC	Real time clock
SAD	Source address decoder
SATA	Serial ATA, a high speed serial data version of the disk ATA interface
SB	Southbound
SBSP	System boot-strap processor
SCSI	Small computer system interface, a connetion normaly used for disks of various tyes
SDDC	Single device data correction
SDR	Sensor data record
SEC	Security, a component of the Intel® Patform Innovation Framework for EFI architecture.
SEL	System event log
SERR	System error
SIO	Super I/o
SMM	Server management mode

SOL	Serial over lan
SPD	Serial presence detect
SSD	Solid state drive
SUP	System update package
TDP	Thermal design power
TIM	Thermal Interface Material
TPM	Trusted platform module
TSOD	Thermal sensor on-DIMM
Intel® TXT	Intel® Trusted Execution Technology
Intel® UPI	Intel® Ultrapath® Connect
USB	Universal serial bus, a standard serial expansion bus meant for connecting peripherals.
VGA	Videographics array
Intel® VMD	Intel® Volume Management Device
Intel® VNNI	Intel® Vector Neural Network Instructions
VR	Voltage regulator
Intel® VROC	Intel® Virtual Raid on Chip
Intel® VT-d	Intel® Virtualization Technology for Direct I/O
Intel® VT-x	Intel® Virtualization Technology