Intel® Server System R2600SR

Technical Product Specification

An overview of product features, functions, and support specifications

Rev 1.0

February 2018

Intel® Server Products and Solutions

<Blank page>

Document Revision History

Date	Revision	Changes
February 2018	1.0	Initial release.

Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Learn more at Intel.com, or from the OEM or retailer.
You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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

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

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

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

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

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

Table of Contents

List of Figures

List of Tables

1. Introduction

The Intel® Server System R2600SR is a density-optimized, rack-mount, 2U, 4-node L9 product family designed to support high-density high-performance computing environments. Each system within the Intel® Server System R2600SR product family includes four independent and pre-configured compute nodes allowing for a power-on ready installation for any supported operating environment.

This Technical Product Specification (TPS) provides a high-level overview of the features, functions, and support specifications of the Intel® Server System R2600SR. For additional information, refer to the documents specified in the following table:

Table 1. Intel® Server System R2600SR reference documents

Document Title	Document Classification
Intel® Server System R2600SR Setup and Service Guide	Public
Intel® Server System R2600SR Configuration Guide	Public
Intel® Server System R2600SR Message and Code Reference Guide	Public
Intel® Server System R2600SR System Management Module Users Guide	Public
Intel® Server System R2600SR Technical Product Specification	Public

1.1 Document Outline

This document covers this information:

- **Introduction** Chapter 1
- **System Features Overview** Chapter 2
- **Compute Nodes** Chapter 3
- **Processor Support** Chapter 4
- **System Memory** Chapter 5
- **Rear System I/O** Chapter 6
- **Server Management** Chapter 7
- **System Power** Chapter 8
- **Cooling** Chapter 9
- **System Security** Chapter 10
- **Server Board Switches** Chapter 11
- **Serviceability Features** Chapter 12
- **Glossary of Terms** Appendix A

2. System Feature Overview

Each system within the Intel® Server System R2600SR product family includes four fully configured compute nodes and other common features. Available system models differ only by processor model, DIMM size, and SATA SSD size preconfigured within each compute node. Refer to the Intel® Server System R2600SR Configuration Guide for a complete list of available system models.

The following illustration represents the system view of the Intel® Server System R2600SR.

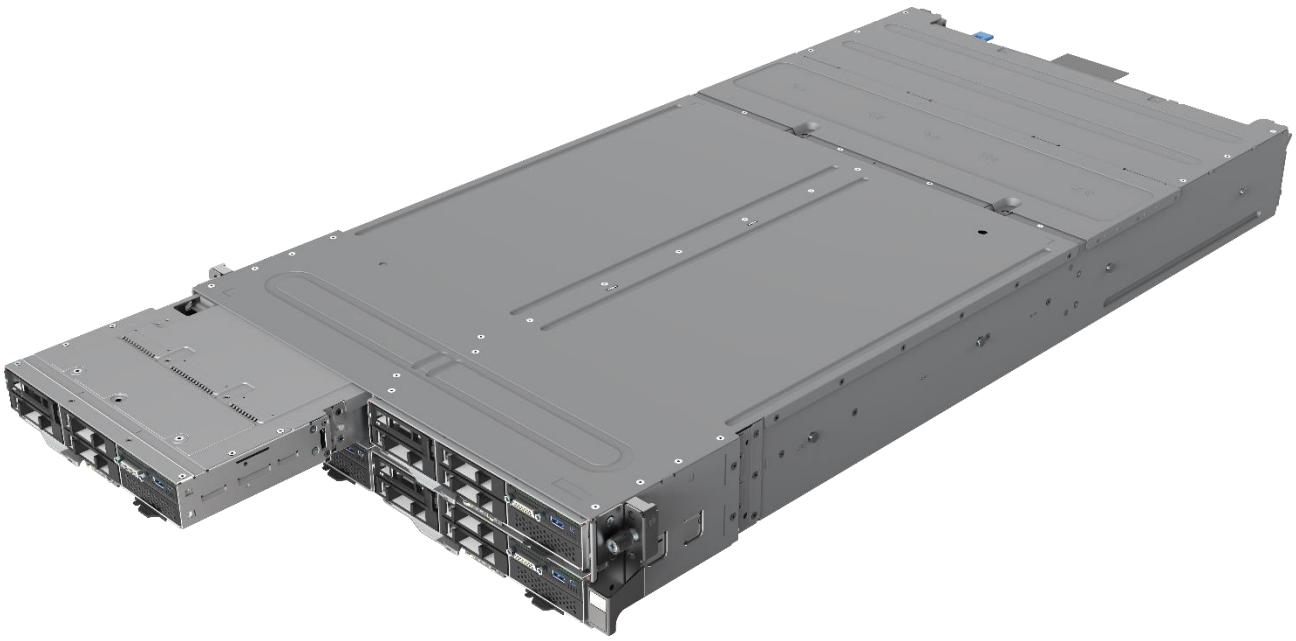

Figure 1. Intel® Server System R2600SR

The following table provides an overview of the feature set and other system specifications which the Intel® Server System R2600SR product family systems are designed to support.

Table 2. Intel® R2600SR server system features/specifications

Feature/Specification	Description
Chassis Type	2U, rack-mount, multi-node
Number of Compute Nodes	4
Chassis Dimensions	Height: 87.0 mm (3.5 inches) Depth: 891.5 mm (35.1 inches) Width: 488.0 mm (19.3 inches)
System Weight (configured as shipped)	Weight: 46.14 kg (101.5 lbs) with four fully configured compute nodes
Packaging Dimensions	Height: 285.0 mm (11.25 inches) Depth: 1087.0 mm (42.8 inches) Width: 587.0 mm (23.1 inches)
Packaged Weight (as shipped)	Weight: 62.3 kg (137.35 lbs)
Processor Support	2 – High TDP Intel® Xeon® processor Scalable family (per node) • Intel® Xeon® Gold 6154, 200 Watts, 24.75M Cache, 3.00 GHz • Intel® Xeon® Platinum 8168, 205 Watts, 33M Cache, 2.70 GHz • Intel® Xeon® Platinum 8180, 205 Watts, 38.5M Cache, 2.5 GHz • Supported processor configuration is system SKU-dependent
Chipset	Intel® C624 Chipset
Memory Support (per node)	12 - DDR4 DIMMs – 6 DIMMs per CPU • 2666 MHz RDIMM • 1.2V • 768 GB Max • 16GB or 32GB, depending on system configuration
Storage (per node)	1 – SATA SSD • Intel® SSD DC S4600 Series drive-sized 480GB, 960GB, or 1.9TB, depending on system configuration
PCI Expansion (per node)	1 – PCIe X16 Add-in Card Slot • Add-in card slot for each node is pre-populated with an Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16 Low Profile
Network Ports (per node)	2 – 10GbE Base-T, RJ45
Cooling	5 – Dual-rotor Hot-Swap System Fans with support for fan redundancy • Three 60x60x56mm fans • Two 80x80x80mm2 fans 1 – Fan per installed Power Supply Module (PSM)
Local KVM support (per node)	Supported via pre-installed KVM Module located on the front of each node 1 – USB 3.0 Port 1 – USB 3.1 Port 1 – KVM Breakout Cable Connector (1 cable per system) supporting the following: • Serial Port • Video Port • USB 2.0 Ports

Power	2 – 2000 watt AC power supplies
	• Power redundancy support
	• 80 PLUS Platinum
	• Input voltage 2000W
	  ○ Minimum: 200 VAC
	  ○ Maximum: 240 VAC
	• Input kilovolt-amperes (kVA), approximately:
	  ○ Minimum: 0.153 kVA
	  ○ Maximum: 2.61 kVA

Feature/Specification	Description
System Management Module (SMM) (Included)	• Offers RJ45 port for management of compute nodes and system over 1G Ethernet • Equipped with ASPEED* controller • Integrated systems management • BMC management • Power and cooling control • Remote browser and CLI-based remote user interfaces via dedicated GB Ethernet port • Hot-swappable • Four control LEDs: system error, identification, status, and system power

Rail Kit (Included)	• Tool-less installation
	• Allows for full system extraction from rack
	• Travel distance:
	  ○ 750mm for fan replacement access
	  ○ 920mm for full system removal
	• Maximum weight = 121 lbs / 55Kg

Serviceability	Modular chassis features for simplified serviceability • Fully independent compute node modules • Modular PCIe* Add-in card cassette for easy add-in card replacement • Modular Rear Shuttle for tool-less access to internal system features • Hot-Swap Power Supplies • Hot-Swap System Fans • Hot-Swap Solid State Disk (SSD) Storage
Operating Ambient Temperature Support	10°C – 35° Cat power-on
Acoustical Noise Emissions	With the maximum configuration of four compute nodes with two processors, full memory, full drives, and two (2) 2000W power supplies installed: • Operation: up to 6.8 bels • Idle: up to 6.2 bels
Heat Output (approximate)	Four maximal configuration nodes: 7564.4 BTU per hour (2610 watts)
Security (per node)	Onboard TPM 2.0 support – Not supported in China

2.1 System Feature Identification

The following illustrations provide quick visual reference for identifying key features found on systems in the Intel® Server System R2600SR product family. Corresponding sections within this document give additional information for these features.

Figure 2. Compute node identification

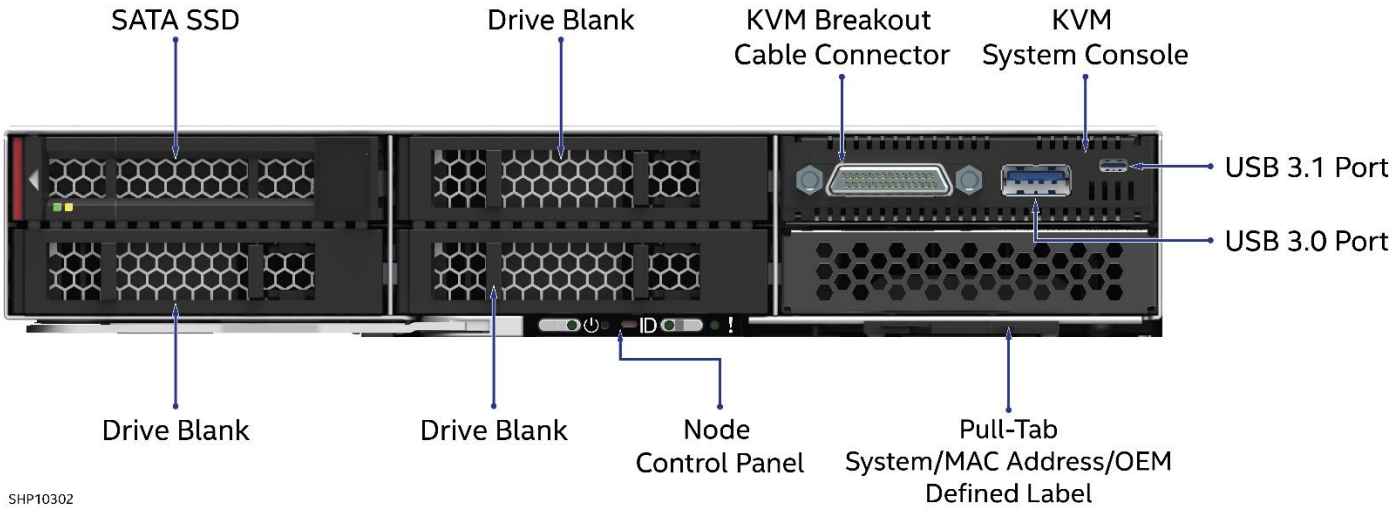

Figure 3. Compute node front panel feature identification

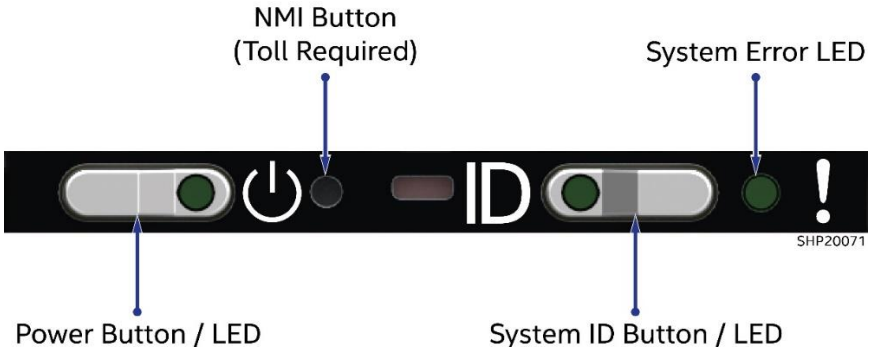

Figure 4. Compute node control panel feature identification

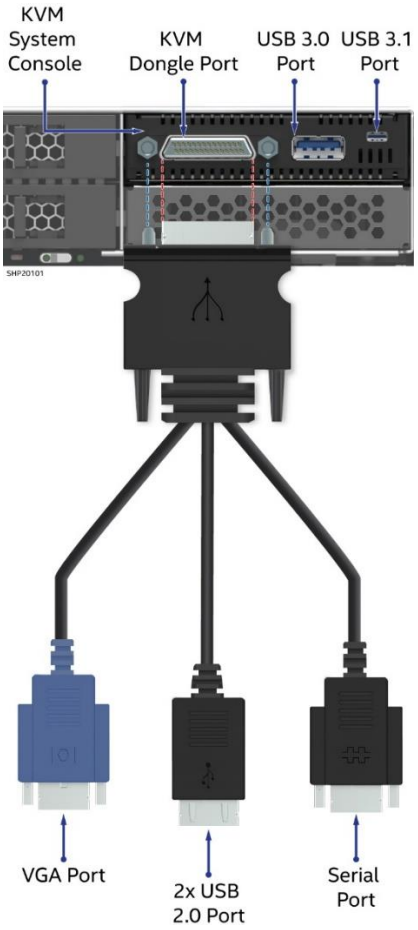

Figure 5. KVM module breakout cable identification

Figure 6. Network access tag identification

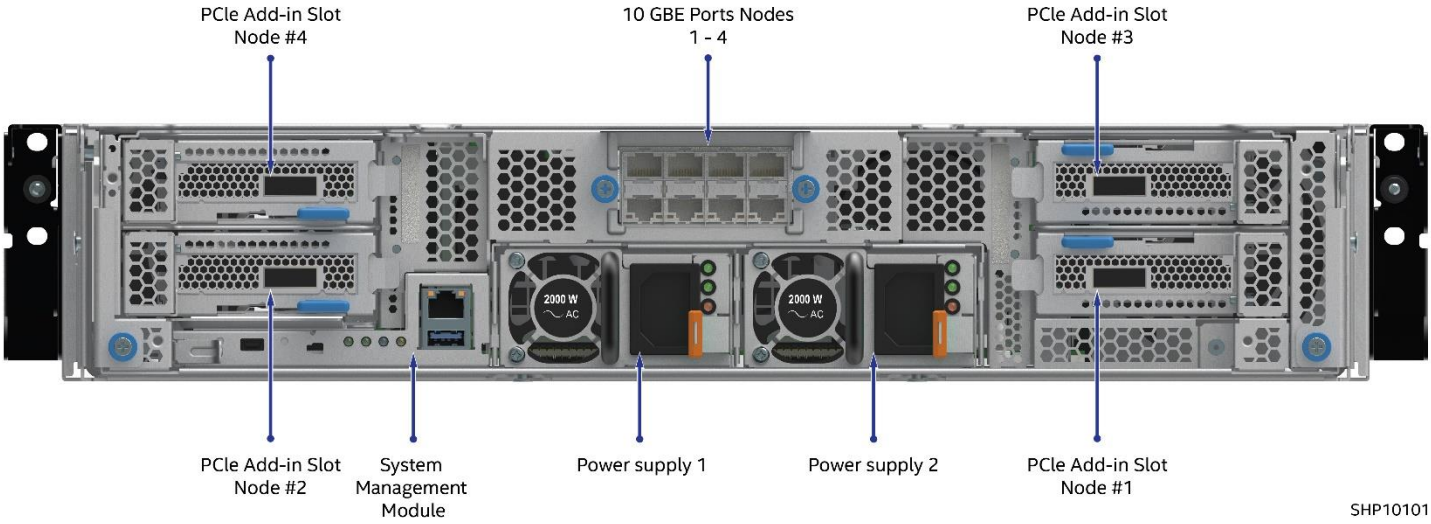

Figure 7. System back panel feature identification

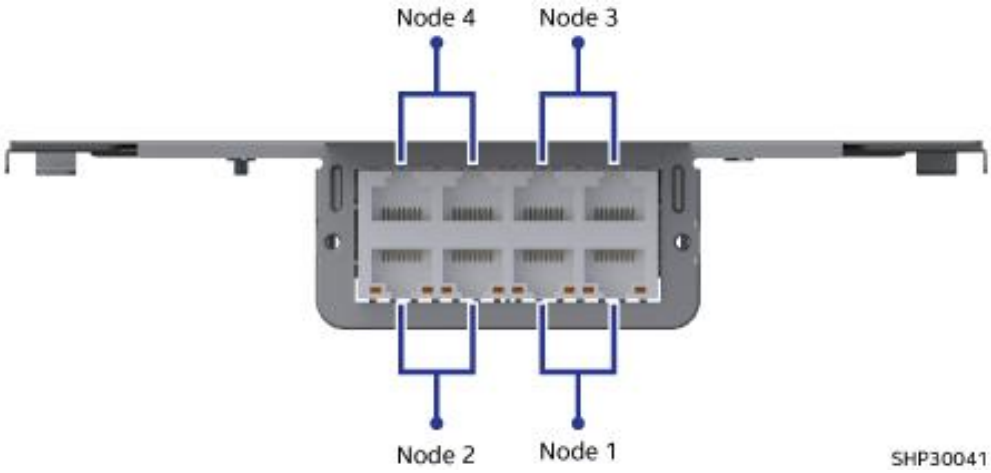

Figure 8. Network port identification

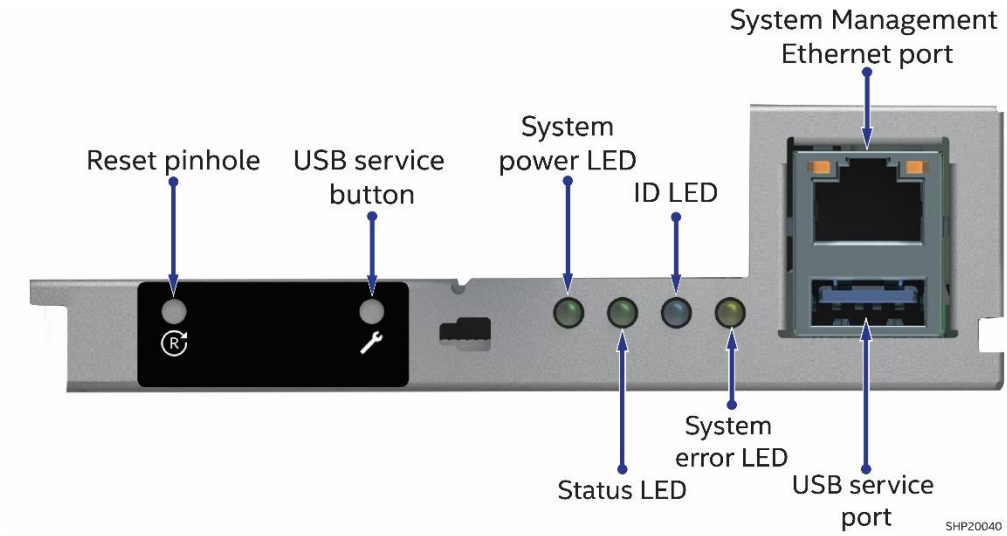

Figure 9. System Management Module (SMM) feature identification

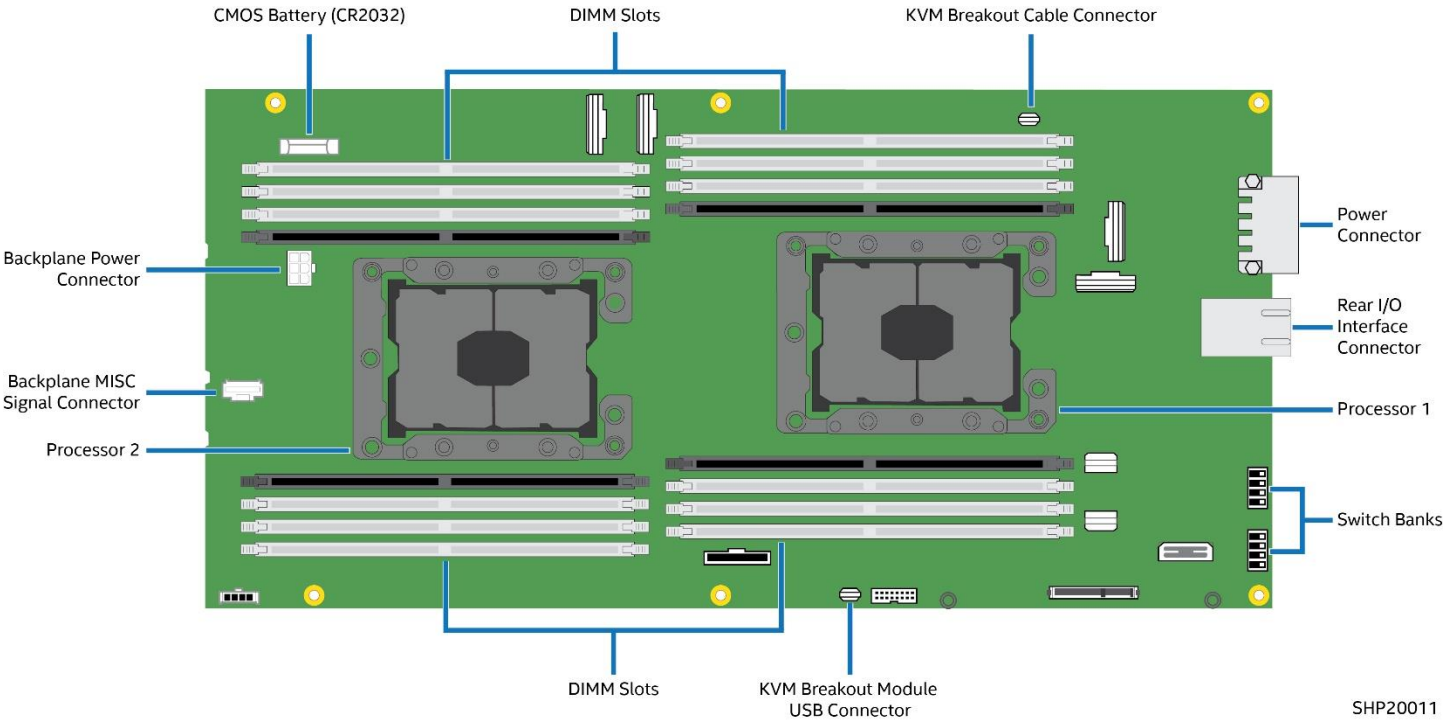

Figure 10. Server board feature identification

3. Compute Node

The compute node is a 1U density-optimized module designed to operate as an independent single-system node within a four-module, multi-node chassis. It is offered in multiple preconfigured system configurations. Reference the Intel® Server System R2600SR Configuration Guide for configuration specifics.

CPUs, Processor Heat Sink Modules (PHMs), and Dual Inline Memory Modules (DIMMs) are installed on the server board integrated in the compute node. Also pre-installed in the compute node are a Solid State Drive (SSD) and a Keyboard/Video/Mouse (KVM) Module.

Figure 11. Compute nodes

3.1 Solid State Drive (SSD) Support

Each compute node includes one 6 GB/sec SATA Solid State Drive (SSD) mounted to a hot-swap drive tray. The system model determines SSD capacity. See the Intel® Server System R2600SR Product Family Configuration Guide for a list of available system models.

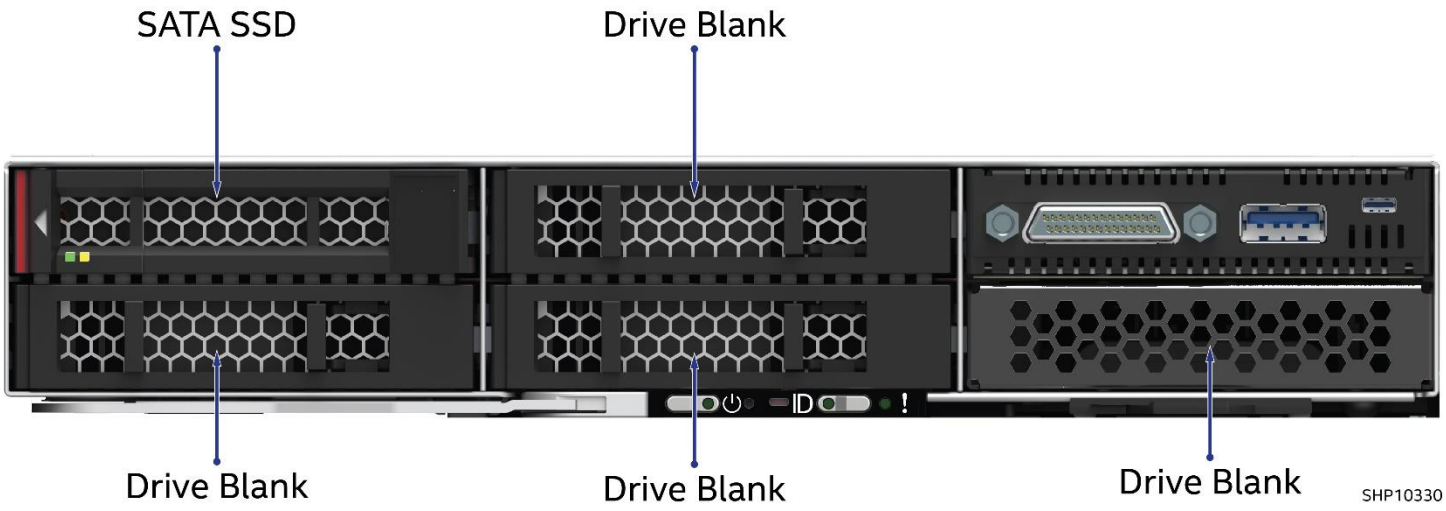

Figure 12. Compute node storage configuration

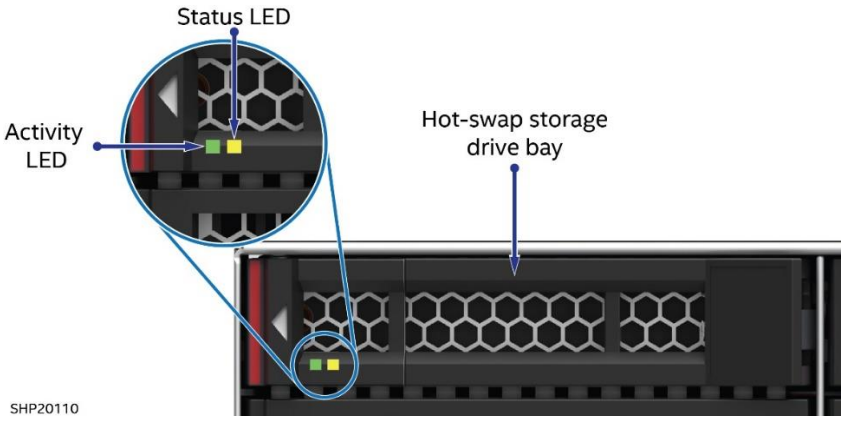

Figure 13. Drive LED definition

Three additional drive bays on each compute node are populated with drive blanks. Drive blanks do not support the installation of a hard disk drive (HDD) or a solid state drive (SSD).

Note: If an HDD or SSD is added to any compute node, it cannot be supported to assure that thermal support of high TDP processors is maintained to deliver best performance.

3.2 Keyboard/Video/Mouse (KVM) Module

Installed within the front panel of each compute node is a Keyboard/Video/Mouse (KVM) Module, which provides support for one USB 3.0 Port and one KVM Breakout Cable (one per system). When attached to the KVM Module, the KVM Breakout Cable supports one serial port, two USB 2.0 ports, and one VGA port. The KVM Module and the cable combined allows for the connection of a keyboard, video monitor, and mouse for local console support of an individual compute module.

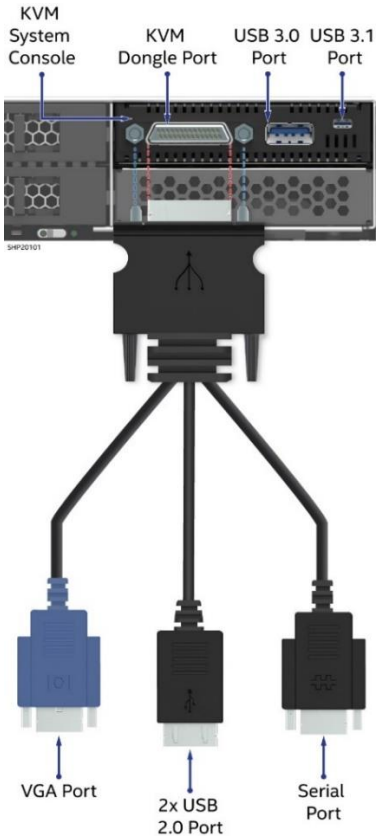

Figure 14. KVM breakout cable and port

3.3 Control Panel

Located along the bottom edge of each compute node is a control panel dedicated to that specific node. The control panel functions include a power button with a Status LED, a recessed Non-Maskable Interrupt (NMI) button, a Status ID button with LED, and a System Error LED.

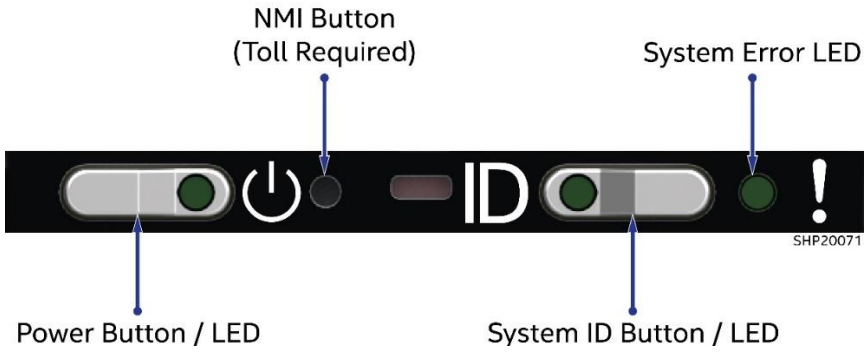

Figure 15. Compute node control panel feature identification

3.4 Network Access Tag

Each compute node installed in the system includes a network access tag that includes labels on the top with the following information:

- Network Port MAC Addresses
- IPv6: Link Local Address

The network access tag can also be used to affix OEM-defined labels.

Figure 16. Network access tag on the front of the compute node

4. Processor Support

Each of the four compute nodes installed within the Intel® Server System R2600SR is configured with two processors from the Intel® Xeon® processor Scalable family. All compute nodes within a common system are configured with identical processors that meet high thermal design power (TDP) limits, ranging from 200W up to 205W, depending on the system configuration.

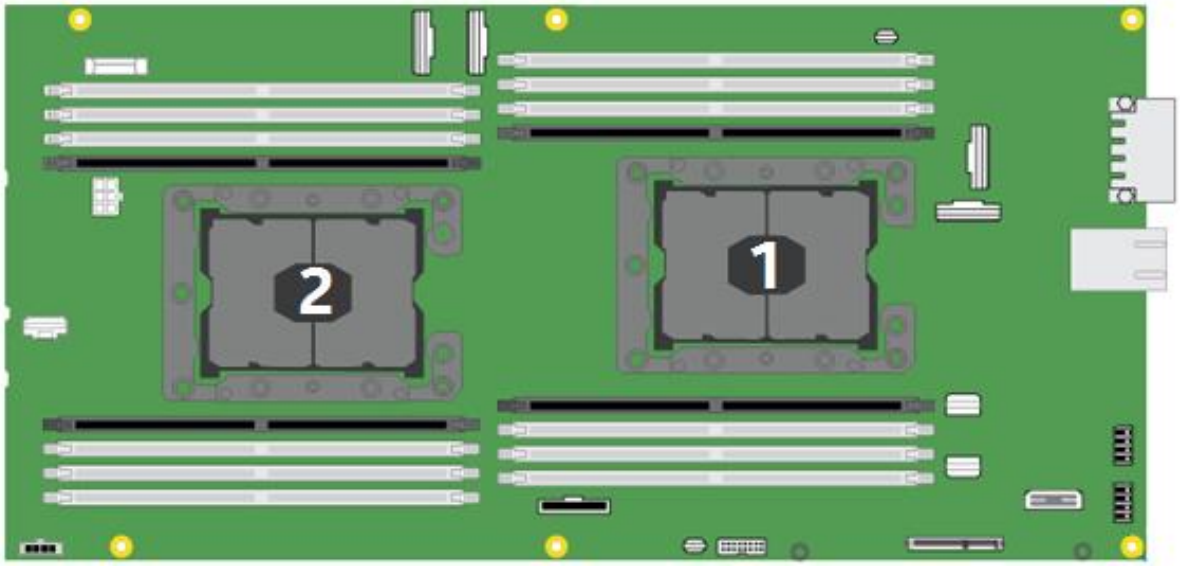

Figure 17. Processor identification

Note: Previous generation Intel® Xeon® processors and CPU heat sinks are not compatible with the compute nodes described in this document.

The server board in this system uses a Processor Heat Sink Module (PHM). The PHM is attached to the processor before the processor is installed onto the server board.

To enable cooling of the higher TDP processors installed in each compute node, two different sizes of heat sinks are used.

1. The heat sink assembled into the PHM towards the front of the compute node and identified as CPU #2 uses an aluminum, 85mm heat sink.
2. The heat sink assembled into the PHM at the rear of the compute node and identified as CPU #1 uses an aluminum, 108mm heat sink. The additional width of the PHM for CPU#1 ensures adequate cooling for that CPU at the back of the node; however, the additional width covers two of the DIMM slots for that CPU.

Attention: In order to ensure adequate cooling of both CPUs in the compute node, the wider (108mm) heat sink must be installed at the rear of the node in the CPU#1 position and the narrower (85mm) heat sink must be installed in the CPU#2 position.

85mm Heat Sink 108mm Heat Sink

Figure 18. Processor Heat Sinks (PHMs)

The following illustration identifies components associated with the processor assembly, including the PHM. Figure 19 is of the 85mm heat sink and Figures 20 and 21 are of the 108mm heat sink. As seen in the illustrations, both sizes of heat sinks are assembled and installed in the same manner. Labels on the top of each heat sink are identical.

Note: These illustrations do NOT represent the processor installation process.

Before installing a processor into the system, the processor must be assembled into the PHM.

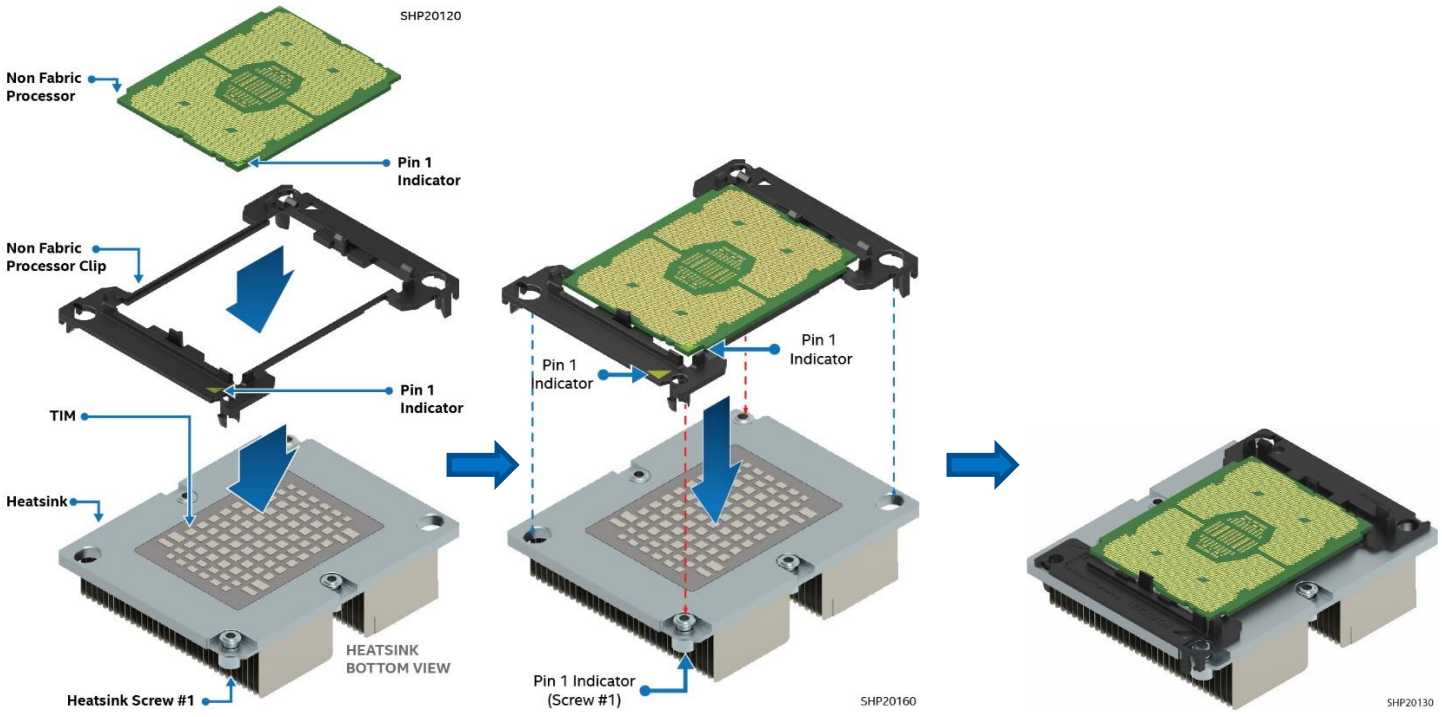

Figure 19. Processor Heat Sink Module (PHM) assembly

Two bolster plate guide pins of dissimilar sizes allow the PHM to be installed only one way onto the processor socket assembly.

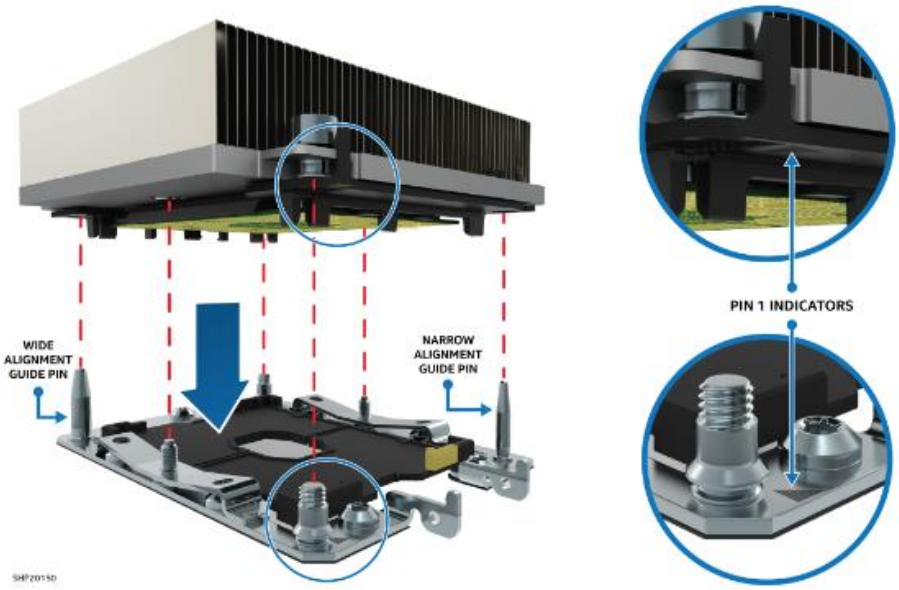

Figure 20. Mounting the PHM to the server board processor socket

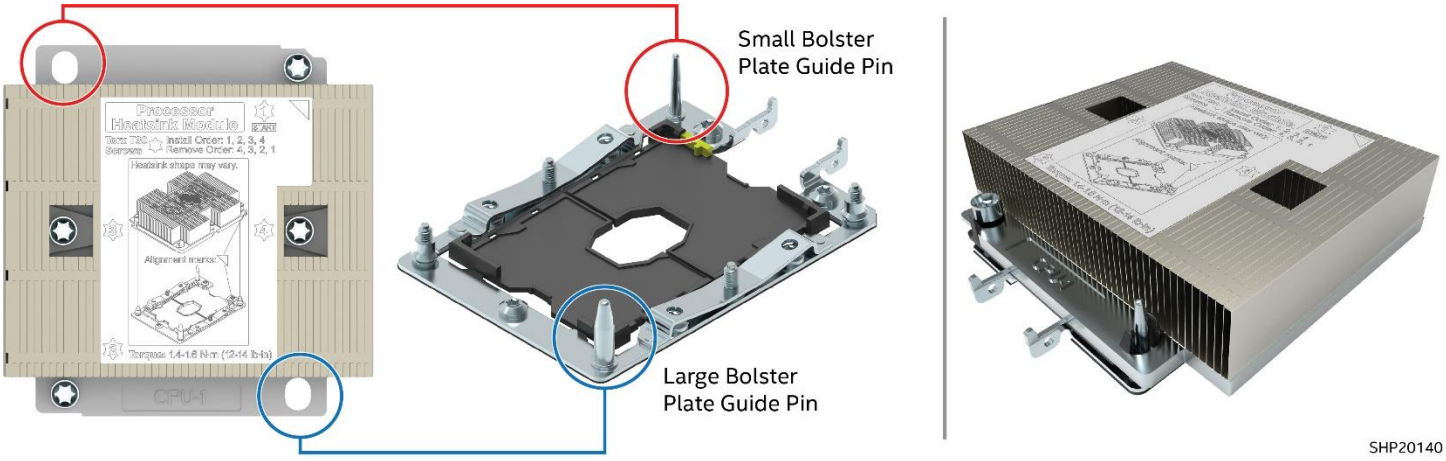

Figure 21. Processor assembly and installation

The PHM is properly installed when it is securely seated over the two bolster plate guide pins and sits evenly over the processor socket. Once the PHM is properly seated over the processor socket assembly, the four heat sink TORX screws must be tightened in the order specified on the label affixed to the top side of the processor heat sink. Heat sink screws should be tightened to 12 In-Lbs torque.

Caution: Failure to tighten/untighten the heat sink screws in the specified order may cause damage to the processor socket assembly.

Refer to the Intel® Server System R2600SR Product Family Setup and Service Guide for detailed processor assembly and installation instructions.

4.1 Intel® Xeon® Processor Scalable Family Overview

The Intel® Server System R2600SR product family supports these processors from the Intel® Xeon® processor Scalable family:

- Intel® Xeon® Platinum 81XX processor – where **XX** defines the specific processor model
- Intel® Xeon® Gold 61xx processor – where **XX** defines the specific processor model

Note: The table below identifies features each processor is capable of supporting. The architecture and design targets of the server board may use a subset of these features.

Table 3. Intel® Xeon® processor scalable family features

Feature	81xx Platinum	61xx Gold	Notes
# of Intel® UPI Links	3	3	Server board design supports 2 UPIs.
UPI Speed	10.4 GT/s	10.4 GT/s	
Supported Topologies	2S-2UPI 2S-3UPI 4S-2UPI 4S-3UPI 8S- 3UPI	2S-2UPI 2S-3UPI 4S-2UPI 4S-3UPI	
Node Controller Support	Yes	Yes	
# of Memory Channels	6	6	
Max DDR4 Speed	2666	2666	
Memory Capacity	768GB 1.5TB (select SKUs)	768GB 1.5TB (select SKUs)	Server board supports 768 GB max.
RAS Capability	Advanced	Advanced	
Intel® Turbo Boost	Yes	Yes	
Intel® Hyper-Threading	Yes	Yes	
Intel® AVX-512 ISA Support	Yes	Yes	
Intel® AVX-512 - # of 512b FMA Units	2	2	
# of PCIe* Lanes	48	48	

The Intel® Xeon® processor Scalable family combines key system components into one processor package, including the CPU cores, Integrated Memory Controller (IMC), and Integrated IO Module (IIO), as listed below.

Core Features:

- Intel® Ultra Path Interconnect (UPI) – up to 10.4 GT/s
  - Server board design provides support for two UPIs between CPU 1 and CPU 2.
- 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)

Uncore Features:

- Up to 48 PCIe* lanes 3.0 lanes per CPU – 79GB/s bi-directional pipeline
- Six channels DDR4 memory support per CPU
- On-package integration of next generation Intel® Omni-Path Fabric Controller – select SKUs
- DMI3/PCI Express* 3.0 interface with a peak transfer rate of 8.0 GT/s.
- Non-Transparent Bridge (NTB) Enhancements – three full duplex NTBs and 32 MSI-X vectors
- Intel® Volume Management Device (Intel® VMD)
- Intel® Quick Data Technology
- Support for Intel® Node Manager 4.0

5. System Memory

Each compute node in an Intel® Server System R2600SR product family supports up to twelve (12) Dual Inline Memory Modules (DIMMs): six DIMMs per installed processor and one DIMM per memory channel. The server board can support 16 DIMMs (8 DIMMs per processor). This product family supports only twelve (12) DIMMs (6 per processor) for proper cooling of the higher TDP processors installed. Black-colored DIMM slots are not used because of Processor Heat Sink Module (PHM) interference and air flow requirements.

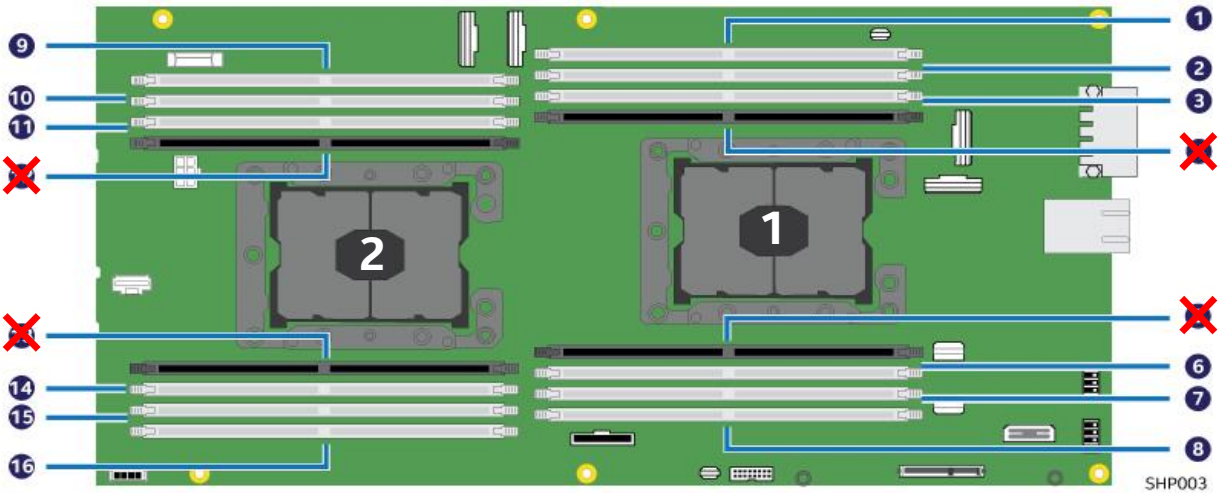

Figure 22. Dual Inline Memory Module (DIMM) slot identification

5.1 Supported Memory

Each compute node is configured with this memory type:

- Industry-standard double-data rate 4 (DDR4), 2666 MT/s, PC4-21300 (single-rank or dual-rank), unbuffered or synchronous dynamic random-access memory (SDRAM) DIMMs with error-correcting code (ECC) in one of the two capacities listed below.

Depending on the configuration selected, each node is prepopulated with:

- 12 - 16GB DDR4 2667 RDIMMs (2R x 8, 1.2V), or
- 12 - 32GB DDR4 2667 RDIMMs (2R x 4, 1.2V)

5.2 Memory Population

Each compute node has support for these memory configurations:

- Independent Memory Mode (Default Normal Mode)
- Memory Mirroring Mode
- Memory Rank Sparing Mode

The compute node supports Independent Memory Mode by default when DIMMs are installed in this order: 6, 14, 3, 11, 7, 15, 2, 10, 8, 16, 1, 9

The compute node supports Memory Mirroring and Memory Rank Sparing modes only when it has the proper memory configuration installed and one of the memory modes is enabled in BIOS Setup.

- To support a Memory Mirroring Mode, identical DIMMs must be installed in these population sets: (6, 7, 8), (14, 15, 16), (1, 2, 3), (9, 10, 11)
- To support a Memory Rank Sparing Mode, DIMMs must be installed in this order: 6, 14, 3, 11, 7, 15, 2, 10, 8, 16, 1, 9

6. Rear System I/O

6.1 PCIe* Add-in Card Support

The system provides each compute node with a PCIe* Riser Card supporting one X16 PCIe* Add-in Card Slot accessible from the back of the system.

Figure 23. PCIe* add-in card identification

For ease of serviceability, each riser card is mounted to a modular cassette to allow for tool-less extraction from the system.

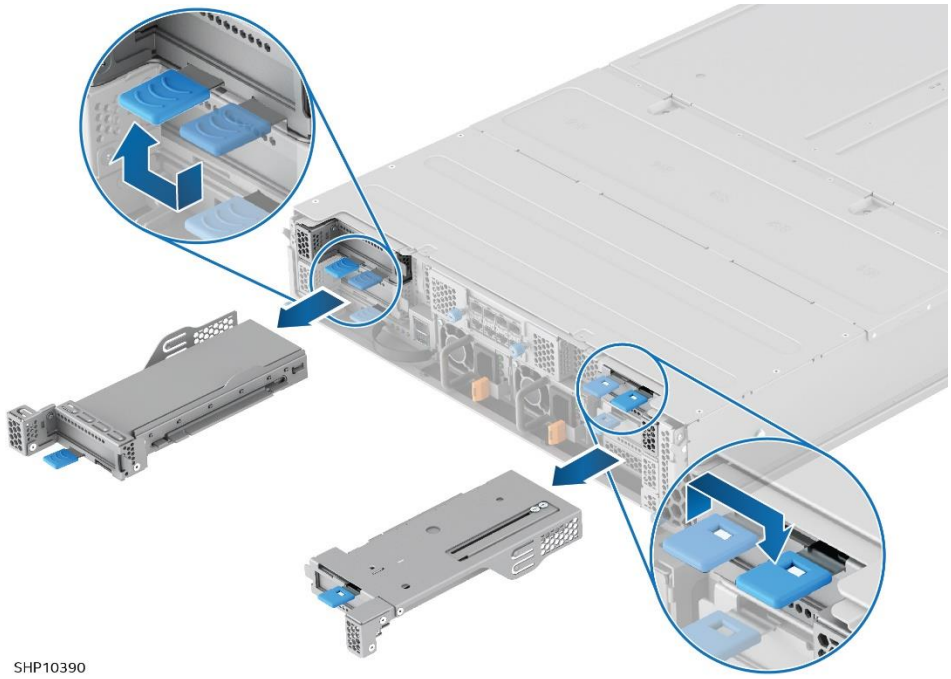

Figure 24. PCIe* riser card cassette modules

Note: The PCIe Riser Card cassette modules are **not** hot-swappable. Power off the corresponding compute node before removing the PCIe Riser Card cassette module from the system.

The PCIe Riser Card cassette can support a low-profile X16 PCIe add-in card.

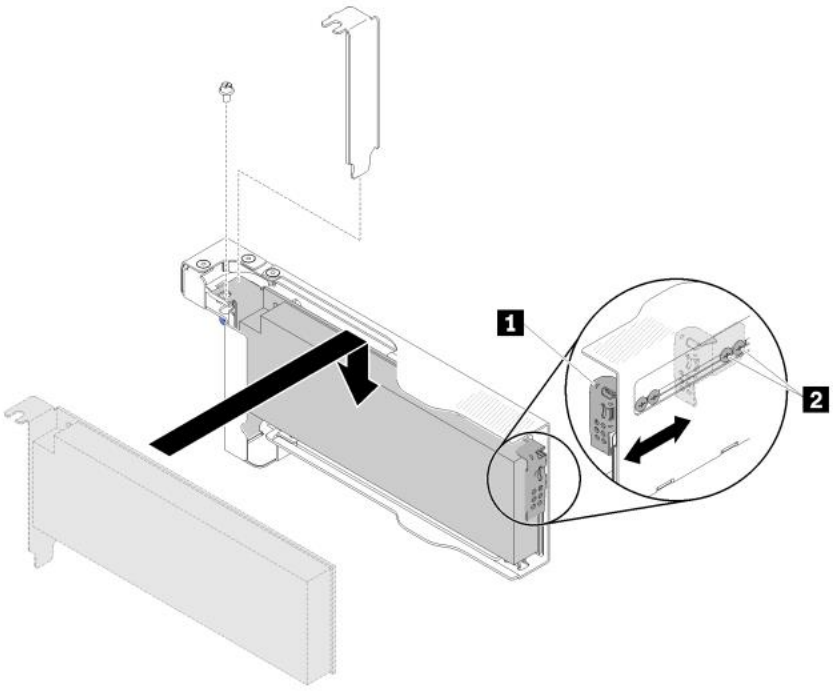

Figure 25. PCIe* add-in card installation to the riser card cassette

6.1.1 Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16

The PCIe Riser Card supported by each compute node may be prepopulated with a single port 100 Series Intel® Omni-Path Host Fabric Interface adapter, shown below.

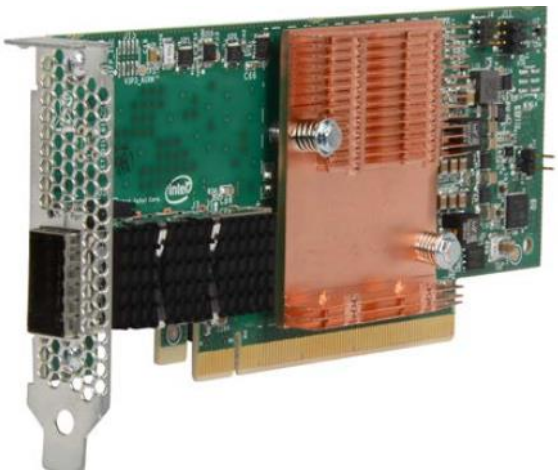

Figure 26. Intel® Omni-Path Host Fabric Interface Adapter

The following table provides detail about the Intel® Omni-Path Host Fabric Interface Adapter add-in card.

Table 4. Intel® Omni-Path Host Fabric Interface Adapter 100 Series specifications

Specification	Description
Board Model	100HFA016LS
Bus Interface	PCI Express* 3.0 x16
Interfaces	QSFP28 quad small form factor pluggable passive copper, optical transceivers, and active optical cables
Max Data Rate	100 GBPS

6.2 Network Ports

Each compute node has two dedicated 10 GB ports routed to the 8-port Ethernet I/O Module (EIOM) at the rear of the system, as shown below. The ports are connected to the integrated Intel® Ethernet Connection X722 Controller.

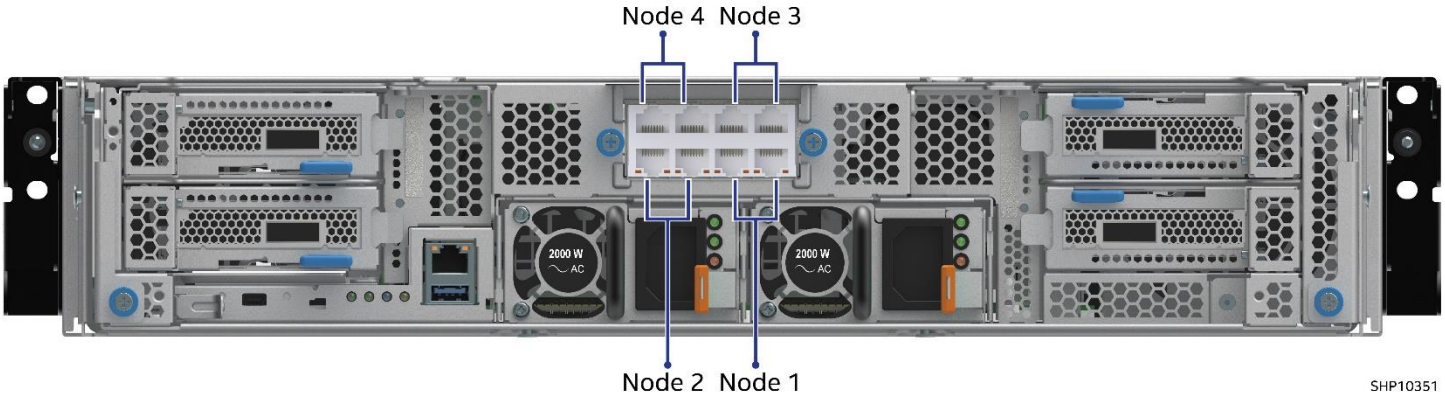

Figure 27. Network port identification

7. Server Management

Server Management is accomplished through the System Management Module (SMM), which performs the following tasks:

- Node Status Report
- Enclosure Power and Fan Status Report
- Enclosure Power and Fan Configuration Management
- Enclosure Vital Product Data (VPD) Information Report
- Enclosure Event Log Display, Save, and Clear
- SMM Management and Settings Backup/Restore

The SMM firmware is displayed as a series of web pages. It supports Transport Layer Security (TLS) 1.2 for data encryption over the network and certificate management.

The SMM web interface is accessed through an Ethernet connection (10/100/1000 Mbit) by establishing a session with the SMM. Through the SMM web interface, the following six system functions are managed:

1. **Enclosure Summary**
   The Enclosure Summary displays overall system status and information in two sections: Front and Rear.

   Front Under the Enclosure Front section, the following node status is detailed:
   - **Node:** Indicates node numbering.
   - **Height:** Indicates node height.
   - **Status:** Indicates the status of the node.
   - **Reset/Reseat**: Is used to perform virtual reset/virtual reseat.

   Rear Under the Enclosure Rear section, three major sections show the Enclosure Rear status:
   - **Management Module:** Indicates the status of the System Management Module (SMM).
   - **Current PSU:** Indicates the status of power supplies.
   - **Fans:** Indicates the status of system fans.

2. **Power**
   The Power section provides status for five power-related conditions/functions:

   - **Power Overview:** Displays the enclosure level power consumption, the node level power consumption, and power consumption of subsystems, which includes power subsystem (PSUs), and thermal subsystem (System fans).
   - **PSU Configuration:** Allows users to set the redundancy mode for power supply units.
   - **Power Cap:** Allows users to set power capping/saving.
   - **Voltage Overview:** Monitors the voltage rail on SMM.
   - **Power Restore Policy:** Allows user to enable the power restore policy.

3. **Cooling**
   The Cooling section contains three sections for the monitoring and management of system cooling:

   - **Cooling Overview:** Monitors system fan speed.
   - **PSU Fan Speed:** Monitors power supply fan speed.
   - **Acoustic Mode:** Allows acoustic attenuation to be adjusted to one of five different acoustic modes.

4. **System Information**
   The System Information section lists six categories of fixed VPD data:

   - **Enclosure VPD:** Displays system/chassis-level information.
   - **PDM VPD:** Provides information related to the Power Distribution Module installed in the system shuttle.
   - **SMM VPD:** Displays System Management Module-related information.
   - **PSU VPD:** Displays Power Supply Module-related information.
   - **EIOM VPD:** Displays Ethernet I/O Module-related information.
   - **PIOR Right/Left VPD:** Displays PCIe I/O Riser-related information for left or right side riser card.

5. **Event Log**
   The Event Log section displays System Event Log (SEL) information. The SEL records enclosure/chassis level information, warning, and critical events that can aid in system issue resolution. A maximum of 4090 event entries can be logged.

6. **Configuration**
   The Configuration section list twelve categories which manage the system and the SMM module are:

   - **Firmware Update:** Allows for importing and loading of firmware images onto the system.
   - **SMTP:** Defines the SMTP traps to be monitored for selected events.
   - **SNMP:** Defines the SNMP traps to be monitored for selected events.
   - **PEF:** Set SMTP/SNMP trap event types in the PEF page.
   - **Network Configuration:** Allows network parameters to be changed.
   - **Time Settings:** Allows the configuration of the system time and date.
   - **User Account:** Defines and manages user roles.
   - **Account Security:** Defines and sets advanced account security rules.
   - **Services:** Allows configuration of HTTPS ports to connect/enable/disable an IPMI service state.
   - **Web Certificate:** Displays current certificate information.
   - **NTP:** Allows configuration of Network Time Protocol (NTP) and time zone settings.
   - **Backup and Restore:** Backs up and restores configuration data.

Refer to the Intel® Server System R2600SR Product Family System Management Module User Guide for a detailed explanation of system management capabilities and procedures.

7.1 System Management Module (SMM) Overview

The following figures provide a visual representation of the System Management Module (SMM), its location within the rear of the system, and feature identification.

Figure 28. System Management Module (SMM)

Figure 29. SMM installed

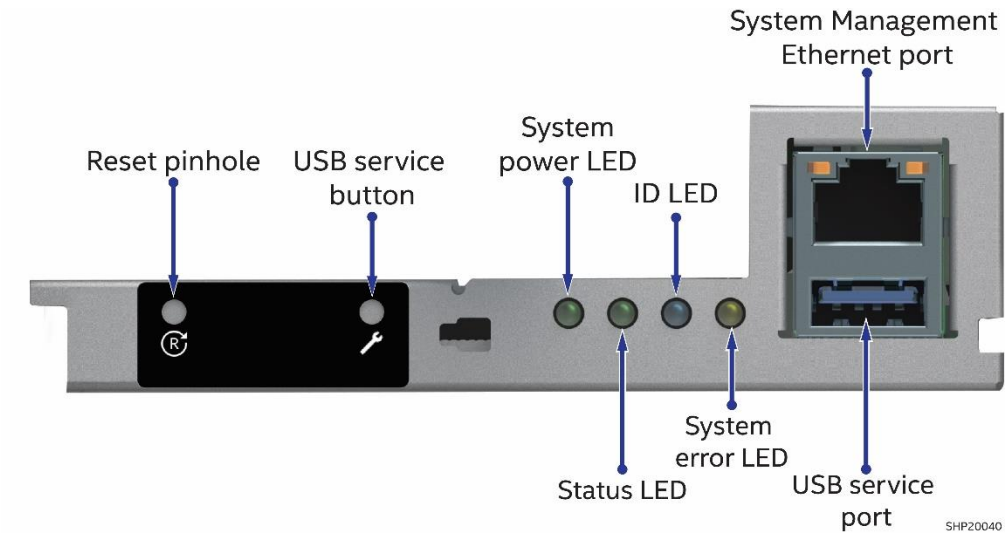

Figure 30. SMM feature identification with LEDs

7.1.1 SMM LEDs

Four LEDs on the SMM provide SMM operating status: System Error, System Identification, System Status, and System Power. Each LED is described below with its color, protocol, and any actions to be taken.

System Error LED (Yellow): When lit, indicates a system error has occurred. Check the event log for more information.

System Identification LED (Blue): When lit, indicates the SMM location. To control the LED for locating the enclosure, use the following commands:

- Turn on:

```
ipmitool.exe -I lanplus -H <SMM's IP> -U USERID -P PASSW0RD raw 0x32 0x97 0x01
0x01
```

- Turn off:

```
ipmitool.exe -I lanplus -H <SMM's IP> -U USERID -P PASSW0RD raw 0x32 0x97 0x01
0x00
```

System Status LED (Green): When lit, indicates SMM operating status. This LED flashes, with the flash rate indicating the status of the SMM:

- Flashes normally: The SMM is working normally.
- Flashes slowly: The speed of the flash indicates the following states.
    - A flash once every 0.25 seconds indicates the SMM firmware is in the preboot process.
    - A flash once every 1 second indicates the SMM firmware is ready.

System Power LED (Green): When lit, indicates that SMM has power.

Note: The default SMM IP address is 192.168.70.100

Refer to the Intel® Server System R2600SR Product Family System Management Module User Guide for a detailed explanation of system management capabilities and procedures.

8. System Power

The Intel® Server System R2600SR comes standard with two (2) hot-plug power supplies which act as a redundant pair to ensure the system remains powered even if one power supply fails or is disconnected. These power supplies are 2000W, 80 PLUS Platinum-certified for energy efficiency.

Note: The system-supported 2000W power supplies auto-range from 200 – 240V AC only. Lesser rated power supplies and lower voltages are not supported by the system.

Figure 31. Power Supply Module (PSM) identification

Each power supply has three LED indicator lights to the right of the power cord which light, blink, or light and blink to indicate power supply status. The LEDs are labeled as AC, DC, and Error (!). Refer to the Intel® Server System R2600SR Product Family Setup and Service Guide for information on these status indicators.

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 with a matching slot connector on the I/O shuttle.

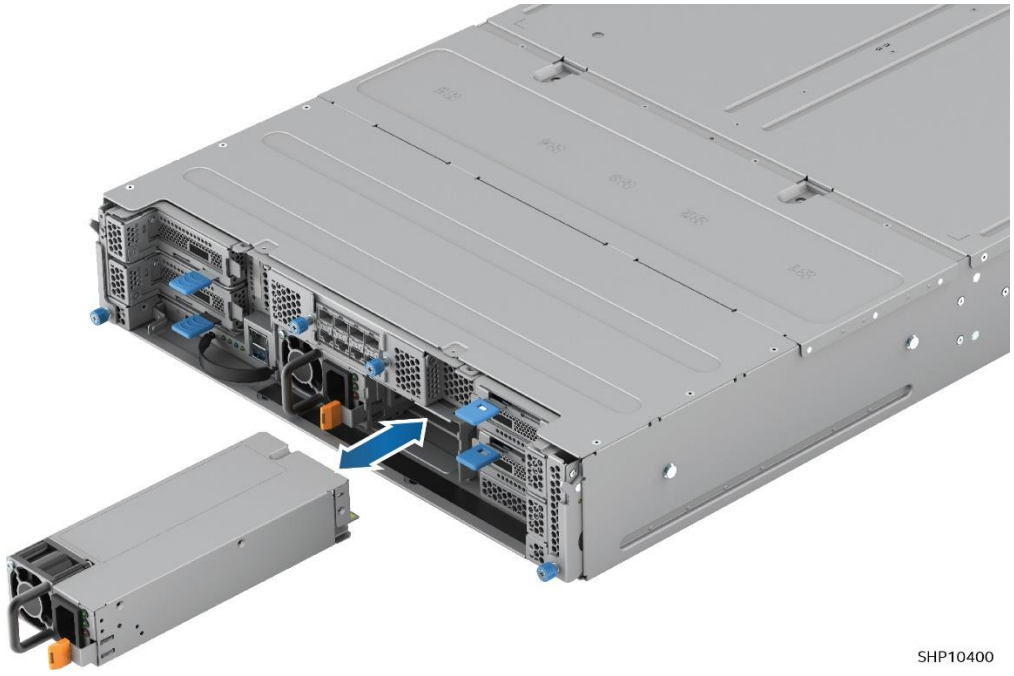

Figure 32. PSM tool-less removal and insertion

9. Cooling

The system includes five, dual rotor hot-swap fans which are used to cool all components. In addition, each Power Supply Module (PSM) has its own integrated fan.

The five hot-swap fans have the following specifications:

- Three (3) 60mm x 60mm x 56mm, dual rotor, hot-swap fans
- Two (2) 80mm x 80mm x 80mm, dual rotor, hot-swap fans

The fans are accessible by removing the panel on the top of the system, as shown below.

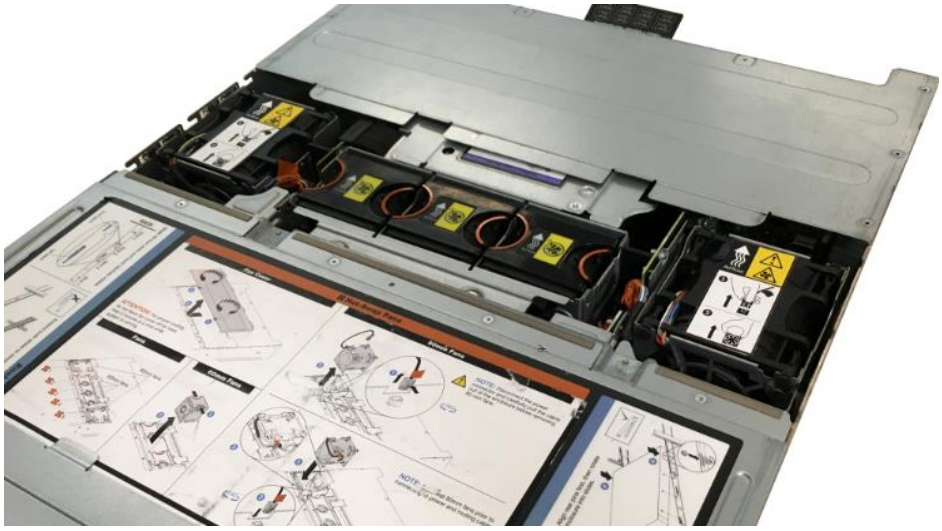

Figure 33. System fans

Cooling by the five system fans and the power supply fans is monitored through the System Management Module (SMM).

System fan health (by fan rotor) is monitored and data updated every 30 seconds. Updates include actual fan speed, lower non-critical and lower critical fan speed along with a visual indication noting if the fan is "Healthy" or in "Fault" mode. Power supply fan speed is also monitored by the PSM with actual speed, % of max speed, and status (Normal, Not Present, or Fault) noted.

In the event of a fan failure, either single or dual rotor, the system fan power control enables power management (capping) to ensure continuous operation of the system. CPU performance while under fan power control may or may not be affected depending on individual system workload and temperature.

Six Acoustic Mode options are selectable to reduce the noise level of the system during run-time:

- **None**: Fan speeds change as required for optimal cooling
- **Mode 1**: Highest acoustic attenuation (lowest cooling)
- **Mode 2**: Higher acoustic attenuation
- **Mode 3**: Intermediate acoustic attenuation
- **Mode 4**: Low acoustic attenuation (higher cooling)
- **Mode 5**: Aggressive cooling mode (no acoustic attenuation)

Additional information is available in the Intel® Server System R2600SR Product Family System Management Module User Guide.

10. System Security

The Intel® Server System R2600SR supports the following system security features:

- Administrator and power-on passwords
- Trusted Platform Module (TPM) supporting both TPM 1.2 and TPM 2.0

Note: The system shipping default system security level is TPM 1.2. During initial system setup, the recommended TPM level for system security is TPM 2.0.

10.1 System User Names and Administrator and Power-on Passwords

System user name and password requirements are largely consistent among the Unified Extensible Firmware Interface (UEFI), BMC Controller, and the System Management Module (SMM). The following rules apply to these password requirements:

User Name Rules (use in establishing user names for user accounts):

- May be from 1 to 32 characters in length
- The upper limit may be constrained by other factors, in particular, endpoints (For example, the maximum length on XCC is 16 because of IPMI restrictions)
- May contain only the following characters: A-Z, a-z, 0-9, . (period), - (hyphen), _ (underscore)

Password Complexity Rules (use in establishing passwords for user accounts):

- Note that the password is set during installation. The following characters are not supported: ~ ' & <> / [ ] { } ;
- White space characters also are not allowed
- May contain only ASCII characters from 0x21 to 0x7E, inclusive of the following:
  - A-Z, a-z, 0-9
  - All other allowable characters are referred to as special characters
- May be from 8 to 20 characters in length
- Must contain at least one letter
- Must contain at least one number
- Must contain at least two of the following:
  - An uppercase letter
  - A lowercase letter
  - A special character
- May contain at most two consecutive occurrences of the same character
- Must not repeat or be a reverse of the corresponding user name

User-Controllable Settings

The User Controllable Settings Table shows additional information concerning passwords.

Table 5. User-controllable settings

Rule	Description	Value			Comments
		Min	Default	Max	
Minimum password length	The minimum number of characters that can be used to specify a valid password.	8	10	20	UEFI allows PAP and POP passwords to be empty. If non-empty, however, these must conform to the minimum password length requirement.
Force user to change password on first access	Determine if a user is required to change the password when logging in to the management server for the first time.	–	Yes	–	Doesn't apply to UEFI as it doesn't have user IDs per se (it has just a PAP and a POP password). Note that LXCA and LXEM require the user to set a password during installation.
Password expiration period	The amount of time in days that a user may use a password before it must be changed. Smaller values reduce the amount of time for attackers to guess passwords. If set to 0, passwords never expire.	0	90	365	
Password expiration warning period	The amount of time in days before password expiration when users begin to receive warnings about the impending user password expiration. If set to 0, users receive no warnings.	0	5	Password expiration period	
Minimum password change interval	The minimum amount of time in hours that must elapse before a user may change a password again after changing it once. The value specified for this setting cannot exceed the value specified for the password expiration period. A small value allows users to more quickly use old passwords. If set to 0, passwords may be changed immediately.	0	24	240	
Minimum password reuse cycle	The minimum number of times that a user must enter a unique password when changing the password before reusing passwords. A higher number enhances security. If set to 0, passwords may be reused immediately.	0	5	10	
Maximum number of login failures	The maximum number of times a user can attempt to log in with an incorrect password before being locked out of the user account. The number specified for the lockout period after maximum login failures determines how long the user account is locked out. Accounts that are locked cannot gain access to the system even if a valid password is provided. The failed login counter is reset to zero after a successful login. If set to 0, accounts are never locked.	0 (LXCA) 0 (other)	20 (LXCA) 5 (other)	100 (LXCA) 10 (other)	

Rule	Description	Value			Comments
		Min	Default	Max	
Lockout period after maximum login failures	The minimum amount of time in minutes that must pass before a user that was locked out can attempt to log back in again. If set to 0, the account remains locked until an administrator explicitly unlocks it. A setting of 0 can expose the system to a greater extent to serious denial of service attacks, where deliberate failed login attempts can leave accounts permanently locked.	0	60	2,880	
Web inactivity session timeout	The amount of time in minutes that a user session established with the management server can be inactive before being logged out. If set to 0, the web session never expires.	0 (LXCA)	1440 (LXCA)	1440 (LXCA)	
		0 (other)	20 (other)	1440 (other)	

10.2 Trusted Platform Module (TPM) TPM 2.0 Support

The Trusted Platform Module (TPM) provides platform security functions such as hash, encryption, and secure storage. It works in conjunction with the processor's Intel® Trusted Execution Technology (TXT) functionality. TPM2.0 is the next generation of TPM; it provides multiple benefits over the former TPM1.2 specification. The Intel® Server System R2600SR uses the SPI version of TPM2.0 on the server board.

The key advantages of TPM2.0 are:

- No special provisioning process is required.
- Authorization is unified for more flexibility and relatively easy-to-revoke keys.
- The SHA-2 and AES encryption algorithms are included.
- A "one size fits all" approach to international security is allowed.

TXT is a hardware solution that validates the behavior of key components within a server or PC at startup. Known as the "root of trust," TXT allows the system to check for consistency in behaviors and launch time configurations against a "known good" sequence.

The boot guard provides a hardware-based Static Root of Trust for Measurement (RTM) and Root of Trust for Verification (RTV) using Intel architectural components. This is accomplished by cryptographically verifying the first portion of OEM BIOS code executed out of reset.

Notes:

Once TPM is enabled, its deactivation takes effect only after the AC power cycle.

The BMC toggles the TPM_EN_VAR variable only based on a command from the management network connection through a Secure Shell (SSH) connection.

The BMC can read the status of the TPM disable pin from the Keyboard Controller Style (KCS) interface but cannot modify it.

For additional details, refer to the Intel® Server System BIOS External Product Specification at the Remote Desktop Connection (RDC).

11. Server Board Switches

The server board within each compute node includes two banks of switches to enable/disable a number of server board functions.

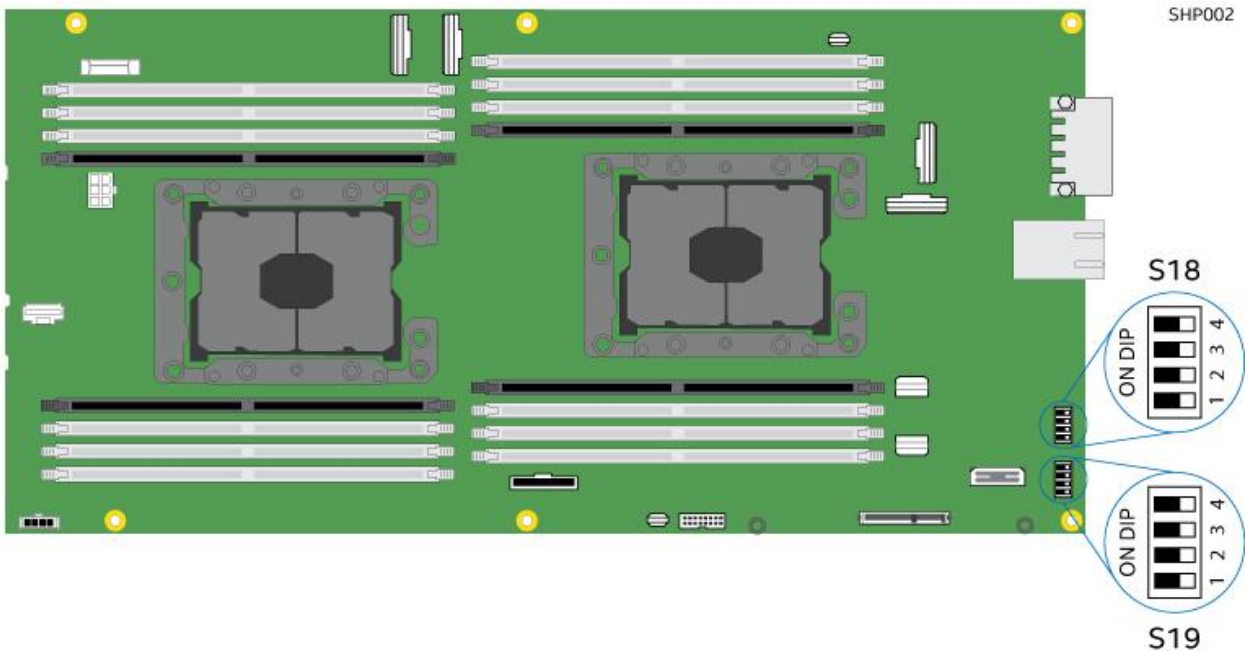

Figure 34. Server board switches

The following table details supported functionality affected by server board switches.

Table 6. Server board functionality and switch positions

Switch Block	Switch	Switch Name	Usage Description	
			On	Off
S18	2	Integrated BMC Boot Backup	Normal (Default)	The compute node boots by using a backup of the integrated BMC firmware.
	3	Integrated BMC Force Update	Normal (Default)	Enables the integrated BMC forced update.
	4	TPM Physical Presence	Normal (Default)	Indicates a physical presence to the system TPM.
S19	1	System UEFI (BIOS) Backup	Normal (Default)	Enables system BIOS backup.
	2	Password Override	Normal (Default)	Overrides the power-on password.
	3	CMOS Clear	Normal (Default)	Clears the real time clock (RTC) registry.

Note: Any switch not defined in the table above is reserved or unused.

To change a switch setting, the compute node must be powered off and removed from the system.

12. Serviceability Features

Each system within the Intel® Server System R2600SR product family includes four fully configured compute nodes and other common components and features. Available system models differ only by the processor model, DIMM size, and SATA SSD size preconfigured within each compute node.

Each system includes several user-serviceable components. The following figures show these components. Refer to the Intel® Server System R2600SR Product Family Setup and Service Guide for service information.

12.1 Compute Node

The compute node is accessible from the front of the system. It can be removed by releasing the locking lever and pulling it straight out of the system chassis. Installation is in the reverse order.

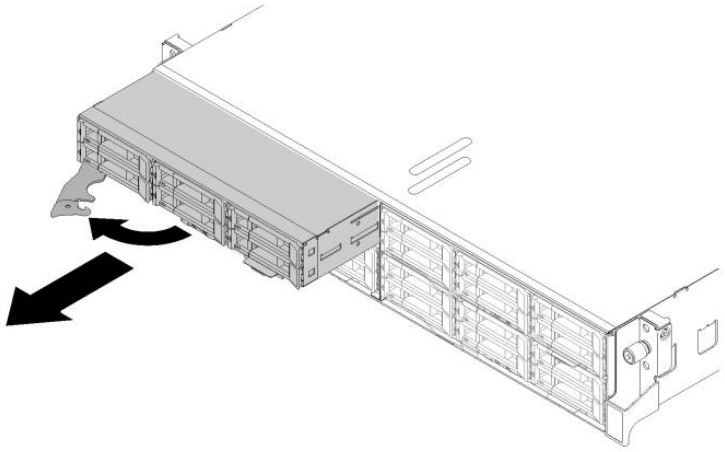

Figure 35. Removing the compute node

12.2 Keyboard/Video/Mouse (KVM) Module

Each compute node offers local manageability through the location of the Keyboard/Video/Mouse (KVM) module on the front of the node. To remove the KVM module, remove the compute module cover and the screw securing it. Disconnect the cables from the server board and push out the module from the rear. Disconnect the cables from the KVM module. Installation is in reverse order.

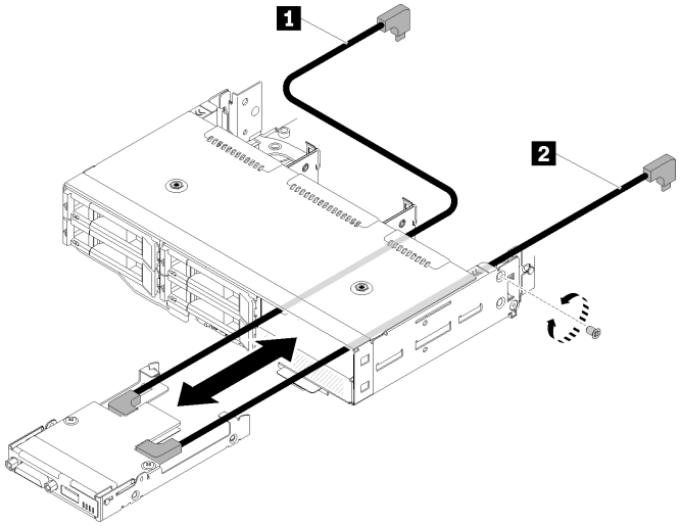

Figure 36. Removing the KVM module

12.3 System Fans

The system includes five hot-swap fans used to cool all components. The five system fans have the following specifications:

- Three (3) 60mm x 60mm x 56mm hot-swap fans
- Two (2) 80mm x 80mm x 80mm hot-swap fans

Access the fans by simply removing the panel on the top of the system, as shown below. To remove the three 60mm fans, grasp them on the top by the orange locking tabs, squeeze and pull straight up. To remove the 80mm fans, disconnect the fan cable and pull straight up, using the semi-circular hold on either side of the fan. Installation is in reverse order.

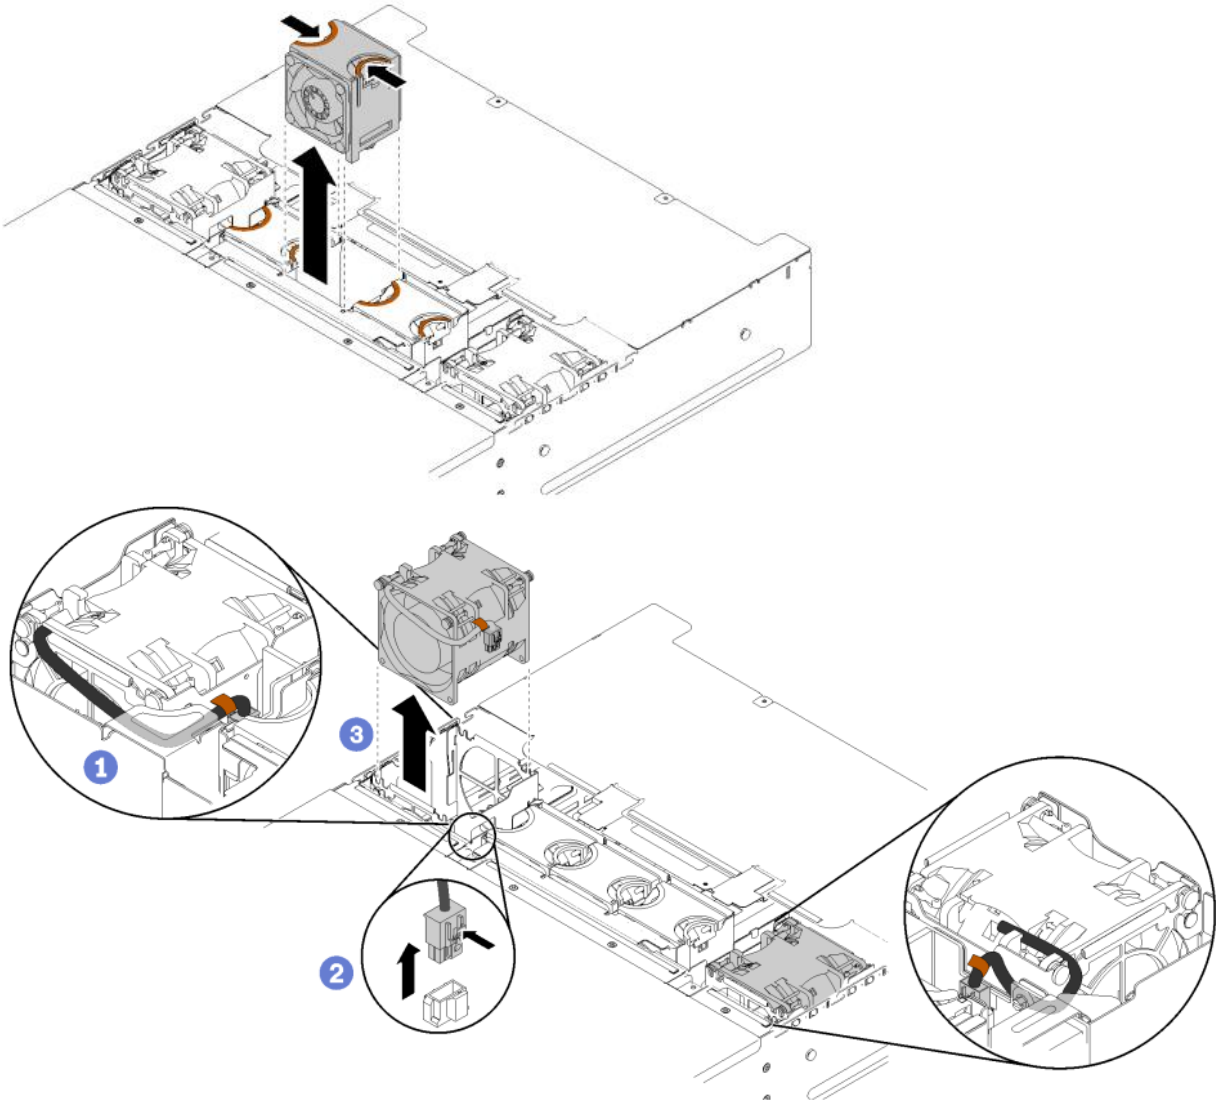

Figure 37. Removing system fans

12.4 Hot-Swap Solid State Drive (SSD)

Each compute node is equipped with a single Solid State Drive (SSD). The drive carrier and SSD can be removed by depressing the release lever to the left, grasping the carrier lever, and pulling the carrier straight out. Installation is in reverse order.

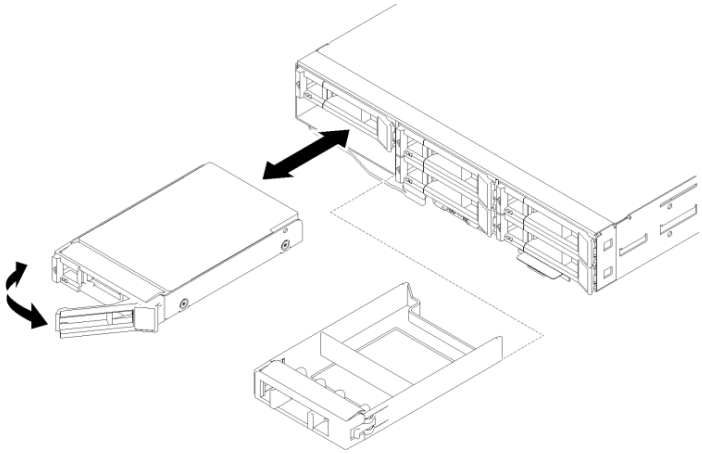

Figure 38. Removing the hot-swap drive carrier

Note: The system supports hot swap of SSD devices. However, each compute node is preconfigured to support only one SSD. Hot-swapping this SSD may cause service interruption.

12.5 Power Supplies

The system comes standard with two hot-plug power supplies which act as a redundant pair to ensure the enclosure remains powered even if one Power Supply Module (PSM) fails or is disconnected. These AC power supplies are 80 PLUS Platinum-certified for energy efficiency. A PSM can be removed by depressing the orange tab on the release lever to the left and pulling straight out by the black handle. Installation is in reverse order.

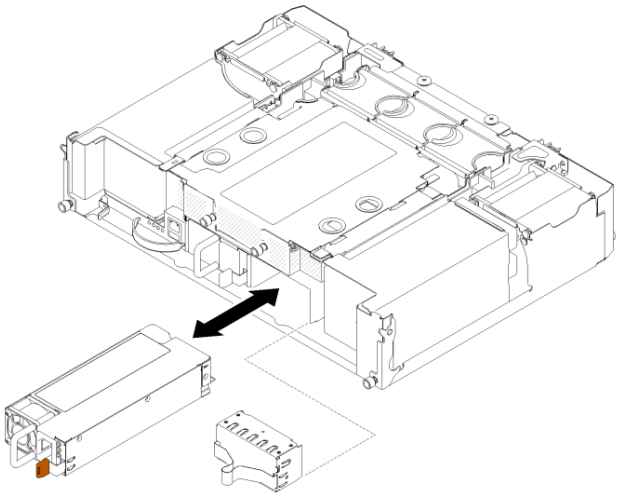

Figure 39. Removing a Power Supply Module (PSM)

12.6 PCIe* Add-in Cards

The system provides each compute node with a PCIe* Riser Card supporting one X16 PCIe* Add-in Card Slot accessible from the back of the system. Each riser card is mounted to a modular cassette, allowing for tool-less extraction from the system.

Remove the modular cassettes by releasing the blue locking lever on a cassette. Cassettes are mounted in the rear of the system opposite each other, but releasing them is the same motion. On the left side, depress the lever down and move it left. On the right, push the lever up to the right. Installation is in reverse order.

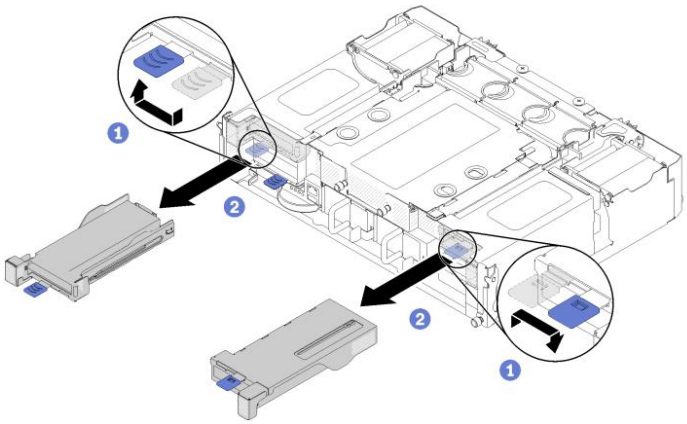

Figure 40. Removing a PCIe* modular cassette

Note: PCIe riser card cassette modules are NOT hot-swappable. Remove them from the system only after powering off the corresponding compute node.

12.7 System Management Module (SMM) Replacement

Before replacing a System Management Module (SMM), it is highly recommended to back up the system settings in order to migrate it to the new SMM.

Remove the SMM by pushing the silver release tab to the right and pulling the module straight out using the black plastic handle on the bottom right of the SMM. Installation is in reverse order.

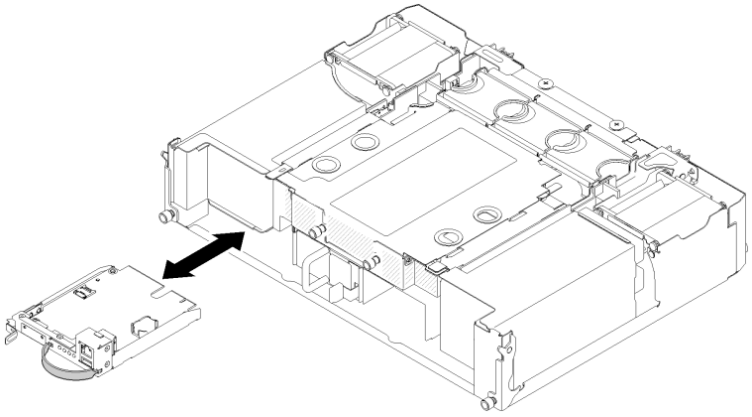

Figure 41. Removing the SMM

12.8 Rear Shuttle Assembly

The rear shuttle assembly houses the system fans, power supplies, SMM, PCIe cassettes, and Ethernet I/O Module (EIOM). Remove it by loosening the two blue screws on the lower outside corners of the system chassis and lifting up on the release levers. The shuttle is then pulled straight out with the release levers. Installation is in reverse order.

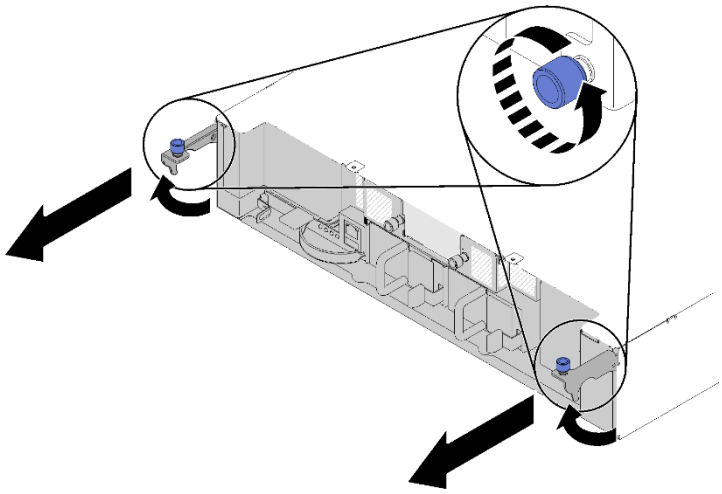

Figure 42. Removing the rear shuttle assembly

Appendix A. Glossary of Terms

Term	Definition
AES-NI	Advanced Encryption Standard New Instructions
ASU	Advanced Settings Utility
Intel® AVX-512	Intel® Advanced Vector Extensions
BELS	Acoustic power level measurement – logarithmic measure of the power of sound relative to a reference value
BMC	Baseboard Management Controller
BIOS	Basic Input/Output System
CLI	Command Line Interface
CMOS	Complementary Metal-Oxide Semiconductor
CPU	Central Processing Unit
DDR	Double-data Rate
DDR4	Double-data Rate 4
DIMM	Dual In-line Memory Module (a plug-in memory module with signal and power pins on both sides of the internal printed circuit board (front and back)).
DMI	Desktop Management Interface
DOM	Disk on Module
ECC	Error Correcting Code
EDS	External Design Specification
EIOM	Ethernet I/O Module
FMA	Floating-point Multiply Add
GB	Gigabyte
GBE	Gigabyte Ethernet
GBPS	Gigabytes Per Second
GT/s	GigaTransfers per second
GUI	Graphical User Interface
IIO	Integrated I/O Module
IMC	Internal Memory Controller
IP	Internet Protocol
IPMI	Intelligent Platform Management Interface
ISA	Instruction Set Architecture
KCS	Keyboard Controller Style/Knowledge-Centered Support
KVA	kilovolt-amperes (kVA)
KVM	Keyboard/Video/Mouse (An attachment that mimics those devices and connects them to a remote I/O user)
LAN	Local Area Network
LED	Light-Emitting Diode
MAC	Media Access Control
MB	Megabyte
MM	Millimeter
NIC	Network Interface Card
NMI	Non-Maskable Interrupt
NTB	Non-Transparent Bridge
OEM	Original Equipment Manufacturer
PCI	Peripheral Component Interconnect (or PCI Local Bus Standard – also called “Conventional PCI”)

Term	Definition
PCIe*	Peripheral Component Interconnect Express* (an updated form of PCI offering better throughput and error management)
PHM	Power Heat Sink Module
PSM	Power Supply Module
QSFP	Quad Small Form Factor-Pluggable
RAM	Random Access Memory
RAS	Reliability, Availability, and Serviceability
RDC	Remote Desktop Connection/Resource Design Center
RDIMM	Registered DIMM (also called buffered). (Memory modules have an address buffer register between the SDRAM modules and the system's memory controller.)
RTC	Real-Time Clock
RTM	Static Root of Trust for Measurement
RTV	Root of Trust for Verification
SATA	Serial ATA (High-speed serial data version of the disk ATA interface)
SDRAM	Synchronous Dynamic Random Access Memory
SKU	Stock-Keeping Unit
SMM	System Management Module
SPI	
SSD	Solid State Device
SSH	Secure Shell (Unix program)
TCM	Trusted Cryptographic Module
TDP	Thermal Design Power
TPM	Trusted Platform Module
TPS	Technical Product Specification
Intel® TXT	Intel® Trusted Execution Technology
UEFI	Unified Extensible Firmware Interface
Intel® UPI	Intel® UltraPath Interconnect
USB	Universal Serial Bus (standard serial expansion bus meant for connecting peripherals)
Intel® VT-d	Intel® Virtualization Technology for Directed I/O
Intel® VT-x	Intel® Virtualization Technology
Intel® VMD	Intel® Volume Management Device
VGA	Video Graphics Array