Intel® Server System M70KLP Family

Technical Product Specification

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

Rev. 1.0

April 2021

M70KLP

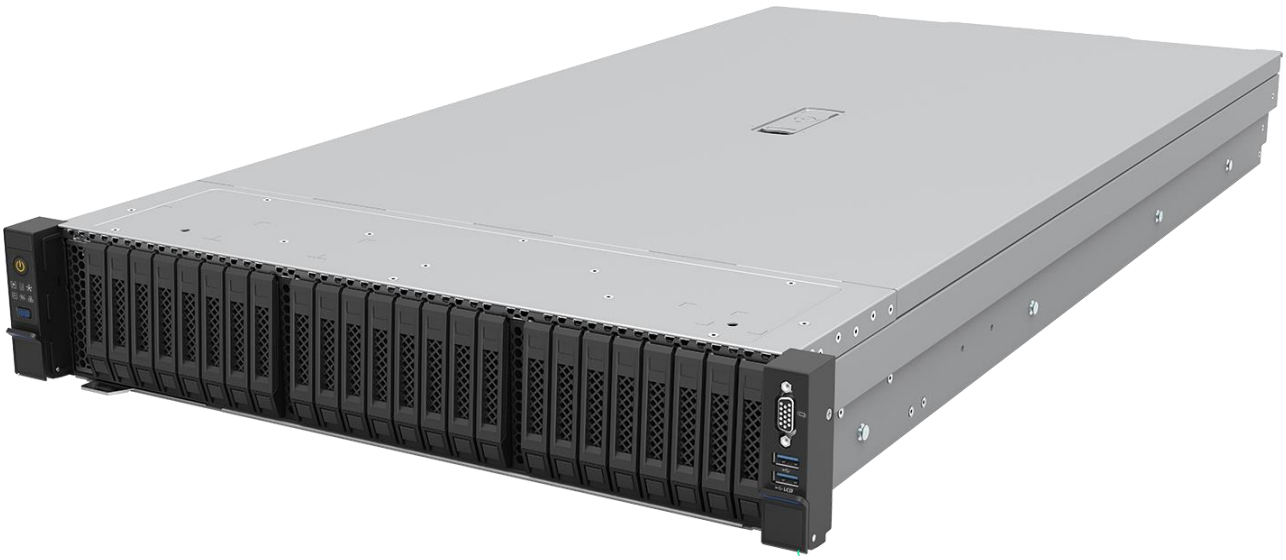

Delivering Breakthrough Data Center System Innovation – Experience What's Inside !

<Blank Page>

Document Revision History

Date	Revision	Changes
April 2021	1.0	1st Production External Release

Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

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.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications.

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, Xeon, Intel Optane, 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

1. Introduction

This technical product specification (TPS) provides a high-level overview of the features, functions, architecture, and support specifications of the Intel® Server System M70KLP product family.

The Intel® Server System M70KLP is a purpose-built system that delivers power and performance at a peak efficiency in a 2U rack mount server form factor. It features the 3rd Gen Intel® Xeon® Scalable processor family in a four-socket configuration, delivering high core count and new hardware-enhanced security features. Previous generation Intel® Xeon® processor and Intel® Xeon® Scalable processor families are not supported.

With support for up to 48 DDR4 DIMMs, the system provides high memory bandwidth for memory intensive workloads. Increase the amount of memory or add memory persistence by adding high capacity Intel® Optane™ persistent memory 200 series modules.

Flexible I/O capabilities include support for optional high-speed networking using Intel network adapters for OCP* 3.0 (small form factor); up to twenty-four (24) 2.5” hot swap capable front drive bays; and up to ten (10) PCIe* add-in cards with optional riser card options. Slim-PCIe connectors on the server board add additional flexibility by providing the option to increase the number of PCIe add-in cards to twelve (12) and/or provide NVMe* SSD support to the front drive bays.

Figure 1. Intel® Server System M70KLP (Standard System, 24-Drive Option Shown)

For additional information, refer to the product support collaterals specified in the following table:

Table 1. Intel® Server System M70KLP Reference Documents and Support Collaterals

Topic	Document Title or Support Collateral	Document Classification
For system integration instructions and service guidance	Intel® Server System M70KLP Family Service Guide	Public
For server configuration guidance and compatibility	*Intel® Server System M70KLP Family Configuration Guide*	Public
For in-depth technical information about this product family	*Intel® Server System M70KLP Family Technical Product Specification (TPS)*	Public
For information on the integrated BIOS Setup Utility	*Intel® Server System M70KLP Product Family BIOS Setup User Guide*	Public (Pending)
For information on the Integrated BMC Web Console	*Intel® Server System M70KLP Product Family Integrated BMC Web Console User Guide*	Public (Pending)
BMC and BIOS Security Best Practices	*Intel® Server Systems Baseboard Management Controller (BMC) and BIOS Security Best Practices White Paper* *https://www.intel.com/content/www/us/en/support/articles/000055785/server-products.html*	Public
Managing an Intel Server Overview	*Managing an Intel Server System 2020* *https://www.intel.com/content/www/us/en/support/articles/000057741/server-products.html*	Public
For technical information for Intel® Optane™ persistent memory 200	*Intel® Optane™ Persistent Memory 200 Series Operations Guide*	Intel Confidential
For setup information for Intel® Optane™ persistent memory 200	*Intel® Optane™ Persistent Memory Startup Guide*	Public
For latest system software updates: BIOS and Firmware	Intel® System Update Package (SUP) for Intel® Server System M70KLP** Intel® System Firmware Update Utility (SYSFWUPDT) - Various operating system support** Intel® System Firmware Update Utility User Guide	Public
To obtain full system information	Intel® SYSINFO Utility for Intel® Server System M70KLP Intel® System Information Utility User Guide	Public
Configure, Save and Restore various system options	Intel® SYSCFG Utility for Intel® Server System M70KLP – Various operating system support** Intel® System Configuration Utility User Guide	Public
To configure and manage Intel® RAID controllers	Intel® RAID Web Console 3 Utility – Various operating system support** Intel® RAID Web Console 3 Utility User Guide Intel® STORCLI Utility	Public
Product Warranty Information	Warranty Terms and Conditions (https://www.intel.com/content/www/us/en/support/services/000005886.html)	Public
Safety and Regulatory Compliance Information	*Intel® Server System M70KLP Product Family Technical Product Specification*	Public

2. System Overview

This chapter provides an overview of the server system features and environmental support specifications.

The Intel® Server System M70KLP family includes two base system configurations; a Standard system offering many configuration options, and a system with configurations and features specifically designed to support high-power graphics processor unit (GPU) add-in cards.

Figure 2. Intel® Server System M70KLP (Standard System, 24 Drive Option Shown)

Figure 3. Intel® Server System M70KLP (GPU Support Option)

The following table provides a high-level overview of the feature set and system specifications supported by each base system option.

Table 2. System Features / Specifications

System Options	Standard System - No GPU Support	System with GPU Support
Chassis Form Factor	2U, Rack Mount	
Chassis Dimensions	841 mm x 435 mm x 87 mm	
Processor Support	• Four (4) Socket P+ (4189 pin) processor sockets Supported Processors – 3rd Gen Intel® Xeon® Scalable processor family: • Gold 63xx (H) (HL) and Platinum 83xx (H) (HL) ○ Up to 28 Cores per processor / Up to 112 Cores per system ○ Six (6) UPI links per processor ○ UPI Speeds up to 10.4GT/s ○ Maximum supported processor TDP: ≤ 250W ** Supported 3rd Gen Intel® Xeon® Scalable processor SKUs must end in (H) or (HL). All other processor SKUs are not supported. ** Previous generation Intel® Xeon® processor and Intel® Xeon® Scalable processor families are not supported	
Chipset	Intel® C621 Chipset	
Memory Support	Up to 48 DIMMs – (12 DIMMs per processor socket) • 6 memory channels per processor • 2 DIMM slots per memory channel DDR4 – RDIMM, RDIMM-3DS, LRDIMM, LRDIMM-3DS Supported Memory Speeds in MT/s: DDR4 SDRAM DIMM Only Configurations • Platinum 83xx: 3200 (1 or 2 DPC); 2933 (1 or 2 DPC) (DPC – DIMMs per Channel) • Gold 63xx: 2933 (1 or 2 DPC) Intel® Optane™ persistent memory 200 Series (App Direct Mode Only) Memory Speeds in MT/s: 2666 (Intel® Optane™ PMem + DDR4 SDRAM configurations)	
PCIe* 3.0 Add-in Card Support Options	• Up to six (6) onboard (default) • Up to 10 with riser card options • Up to 12 with riser card options + Aux PCIe cabling option	• Up to four (4) onboard + • Up to four (4) with riser cards ○ 2 x GPU – (FH, FL, DW) ○ 2 x PCIe X8 – (FH, HL) with Aux PCIe cabling option
Network Support Options	OCP* 3.0 Small Form Factor Add-in card options: • See *Intel® Server System M70KLP Configuration Guide* for supported options.	
Front Access Drive Bay support	8, 16, or 24 Hot Swap Drive Bays • 2.5” SSDs or HDDs • SAS, SATA, NVMe*	8 Hot Swap Drive Bays • 2.5” SSDs or HDDs • SAS, SATA, NVMe
Internal M.2 SSD support	Up to two (2) internal mount M.2 SATA SSDs • 2280 and 22110 form factors supported	Up to two (2) internal mount M.2 SATA SSDs • 2280 and 22110 form factors supported
Rear Panel Features	• One (1) Rear Access OCP Add-in card bay • One (1) OCP slot power button for hot swap support – (NOT SUPPORTED) • Two (2) USB 3.0 Ports • One (1) VGA Connector • One (1) RJ45 Dedicated Management Port • One (1) 3.5mm Serial Port interface connector – (NOT SUPPORTED) • One (1) 3.5mm BMC Serial Port interface connector – (NOT SUPPORTED) • One (1) UID Button/LED • One (1) System Reset Button • One (1) Rear Access Dual Power Supply Module Bay	

System Options	Standard System - No GPU Support	System with GPU Support
Front Control and I/O Panel Features	Left Front Control Panel Features • System Power Button / LED • UID Button / LED • Various system feature/status LEDs Right Front I/O Panel Features • VGA connector • One (1) USB 3.0 connector • One (1) USB 2.0 connector	
Power Supply	Two (2) CRPS AC 2000W (80-Plus Platinum) power supply modules Supported operating modes: • 1+1 Redundant power (hot swappable) • 2+0 Combined Power (No power redundancy)	Two (2) CRPS AC 2000W (80-Plus Platinum) power supply modules Supported operating modes: • 1+1 Redundant Power (hot swappable) • 2+0 Combined Power (No power redundancy)
System Cooling Features	• Six (6) managed 60x60x56 mm system fans with support for fan redundancy • One Fan per installed power supply • Four (4) 2U CPU heat sinks • Standard air duct	• Six (6) managed 60x60x56 mm system fans with support for fan redundancy • One Fan per installed power supply • Four (4) 1U CPU heat sinks • Low Profile air duct
Management Support	• One (1) Dedicated RJ45 1Gb Management Port (Back Panel) • Standard System Management features o Virtual KVM over HTML5 o Integrated BMC Web Console o Redfish* o IPMI 2.0 o Out-of-band BIOS/BMC Update and Configuration o System Inventory o Autonomous Debug Log • Advanced System Management features (Optional) o Virtual Media Image Redirection (HTML5 and Java) o Virtual Media over network share and local folder o Active Directory support • Support for Intel® Data Center Management software (Intel® DCM)	
Serviceability Features	Tool-less (Removal / Installation) • Top cover • PCIe add-in cards and OCP 3.0 Add-in card • System fans – hot swappable • System fan housing • Power Distribution Board • Power Supply Module(s) – hot swappable in 1+1 redundant configuration • Front Drive Bay Backplane(s) • Front Mount Drives – Hot Swappable in fault tolerant RAID configurations. o Screws required to mount drive to drive carrier	
Operating Ambient Temperature Support	• 10 – 35° C ambient temperature	
Security	TPM 2.0 Option (Rest of the world) - **iPC KLPTPM** (**Note:** China only TPM not available) Intel® PFR – Intel® Platform Firmware Resilience Intel® CBnT – Converged Boot Guard and Intel® Trusted Execution (Intel® TXT)	
Rack Mount Kit (Included)	Rack mount rails Tool-less attachment to chassis and rack installation Full extension from rack 38 Kg maximum supported weight	

2.1 System Features Overview

The following illustration identifies the major features of the Intel® Server System M70KLP.

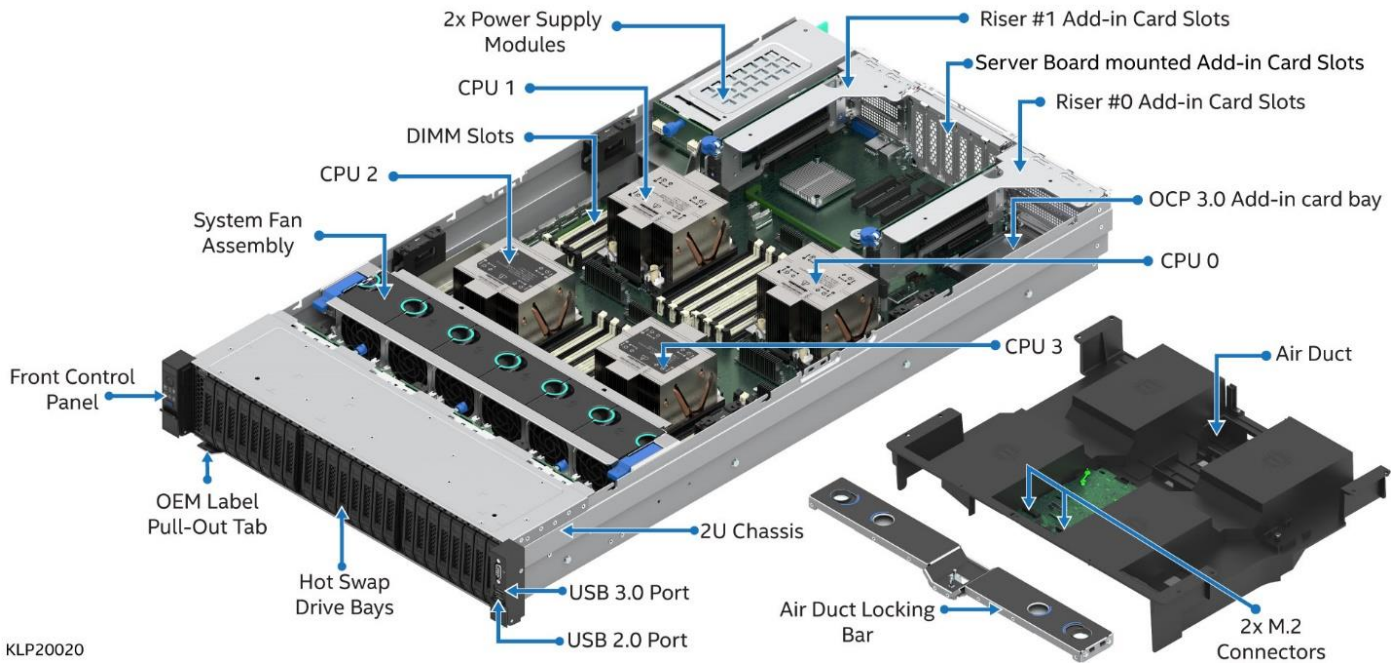

Figure 4. System Features Overview

The Standard and GPU Enabled systems will each include an air duct designed to meet the specific airflow and features requirements of the system.

The Standard air duct will accommodate 2U processor heat sinks and allow support for up to three (3) half-length PCIe add-in cards from Riser #0 and up to three (3) half-length PCIe add-in cards from Riser #1. Mounted atop the air duct are an M.2 interface board capable of support up to two (2) SATA M.2 SSDs, and mounting features to support up to three (3) Intel® RAID Maintenance Free Backup (RMFBU) modules.

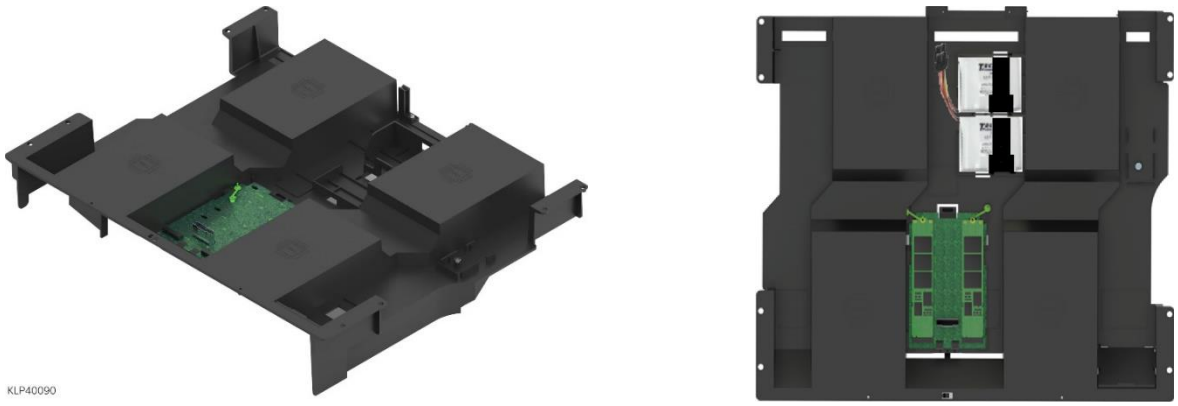

Figure 5. Standard System Air Duct

The low-profile air duct used in the GPU enabled system option will accommodate 1U processor heat sinks and allow support for one (1) GPU add-in card (Full-Height Full-Length Double-wide) + one (1) Full-Height Half-Length PCIe add-in card from each riser card assembly. Mounted atop the air duct is an M.2 interface board capable of supporting up to two (2) M.2 SATA SSDs, and a mounting feature to support one (1) Intel® RAID Maintenance Free Backup (RMFBU) module.

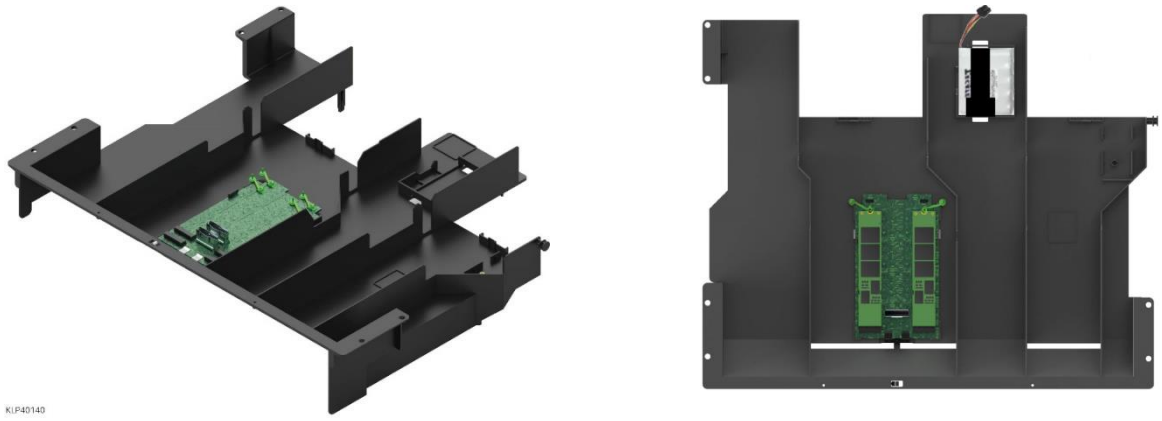

Figure 6. Low Profile Air Duct (GPU Systems Only)

The following illustrations identify the features found on the back panel for both the Standard and GPU enabled system options. All I/O connectors, buttons, LEDs, OCP support, and power supply bays will be common between the two system options.

The base Standard system option supports up to six (6) server board mounted PCIe* 3.0 add-in cards. Riser card options can add support for an additional six (6) PCIe 3.0 add-in cards[1].

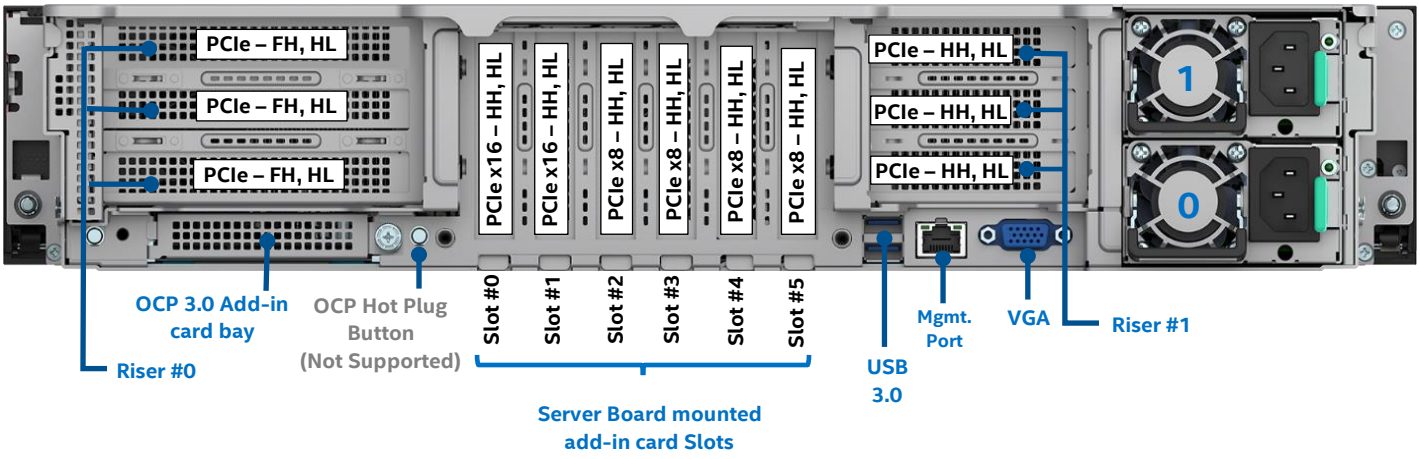

Figure 7. Standard System - Back Panel Features

[1] Some riser card add-in slots require that PCIe bus lanes be cabled to the riser cards from one or more of the PCIe slim-line connectors on the server board.

The GPU enabled system option supports up to four (4) server board mounted PCIe 3.0 add-in cards. Two PCIe riser cards can each support one (1) Full-Height Full-Length Double-Wide GPU add-in card + one (1) Full-Height Half Length PCIe 3.0 add-in card[1].

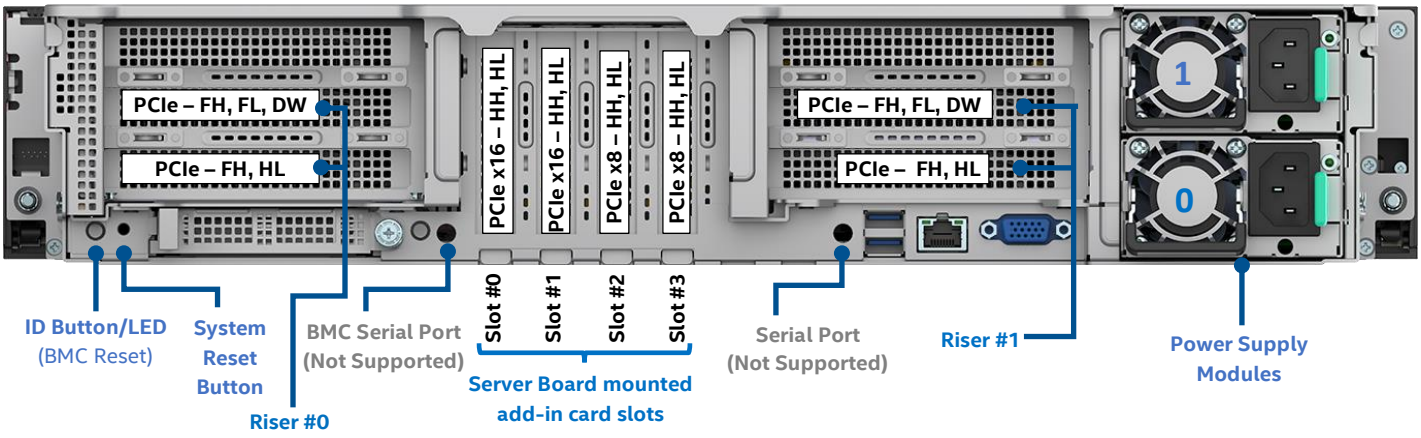

Figure 8. GPU Enabled System – Back Panel Features

The system has support for two AC 2000 Watt (80-Plus Platinum) hot-swap capable power supply modules. Supported power supply operating modes include: 1+1 Redundant and 2+0 combined power.

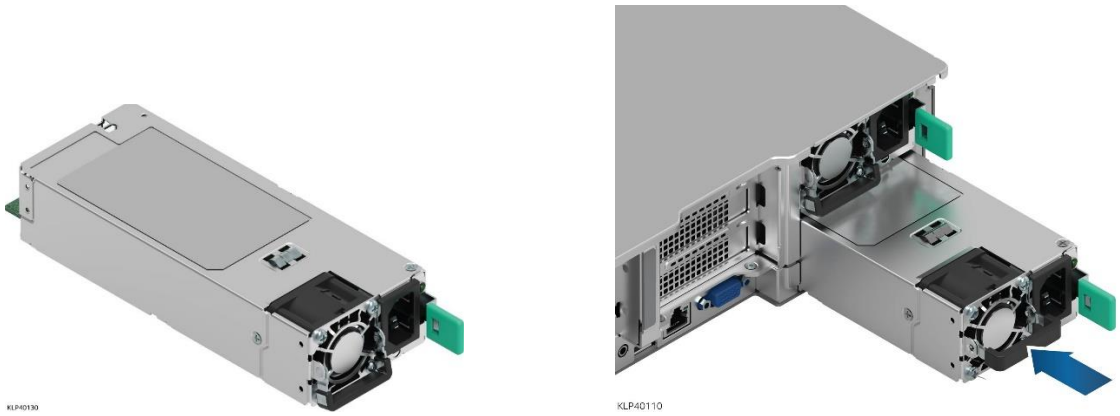

Figure 9. Power Supply Module

Important Note: All integrated L9 system configurations from Intel will draw greater than 1000W of power when running medium to heavy workloads. To support this power requirement, systems must be connected to a 220V AC input power source to operate. A system running a medium to heavy workload while connected to a 110V AC input power source will not operate correctly.

The system has support options for 8, 16, or 24 front drive bays. Each drive bay includes a drive carrier that must be populated with a 2.5" drive (SSD or HDD) or supplied drive blank. Each set of eight drive bays is supported by a common backplane. All drives attached to a common backplane must match media type (SSD or HDD). Mixing drive storage media types within a common backplane is not supported. In addition, all drives attached to a common backplane must share a common interface type (SATA/SAS or NVMe).

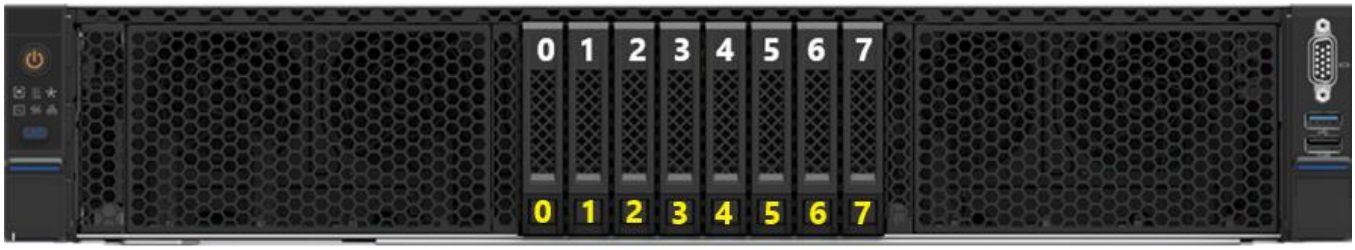

Figure 10. 8 x2.5" Drive Configuration

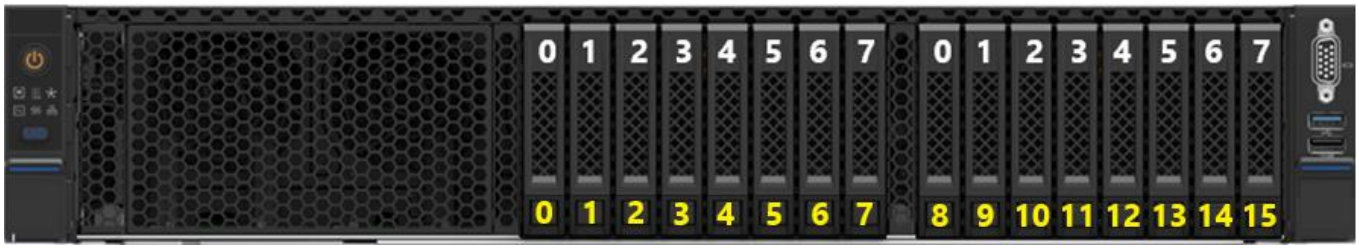

Figure 11. 16 x2.5" Drive Configuration

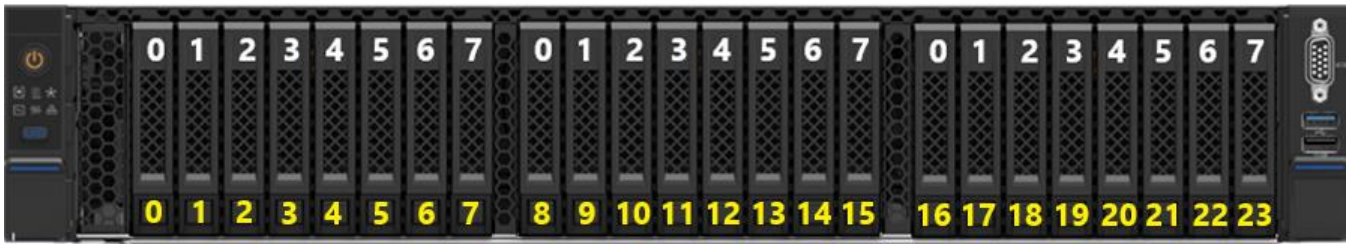

Figure 12. 24 x2.5" Drive Configuration

Embedded within the front system handles are a control panel (Left) and I/O panel (Right), with features identified below.

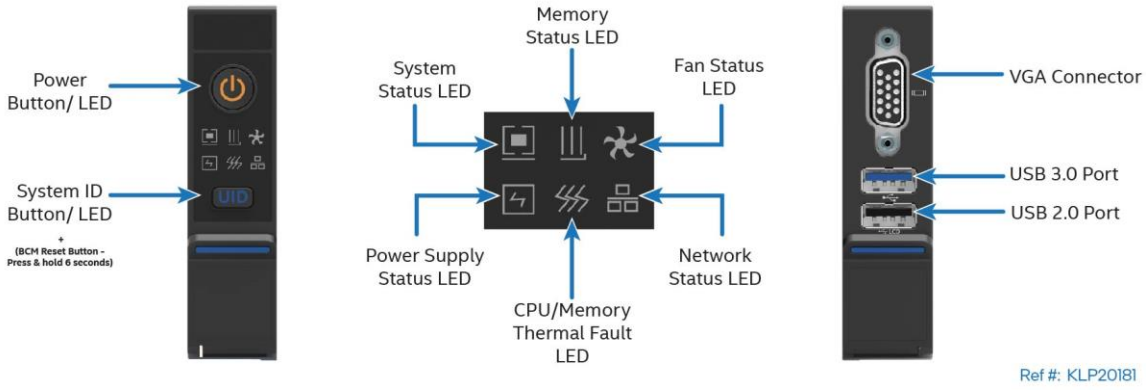

Figure 13. Front Control Panel Features

Integrated within the server system is the Intel® M70KLP server board. The following illustration identifies its supported features and available system option interconnects.

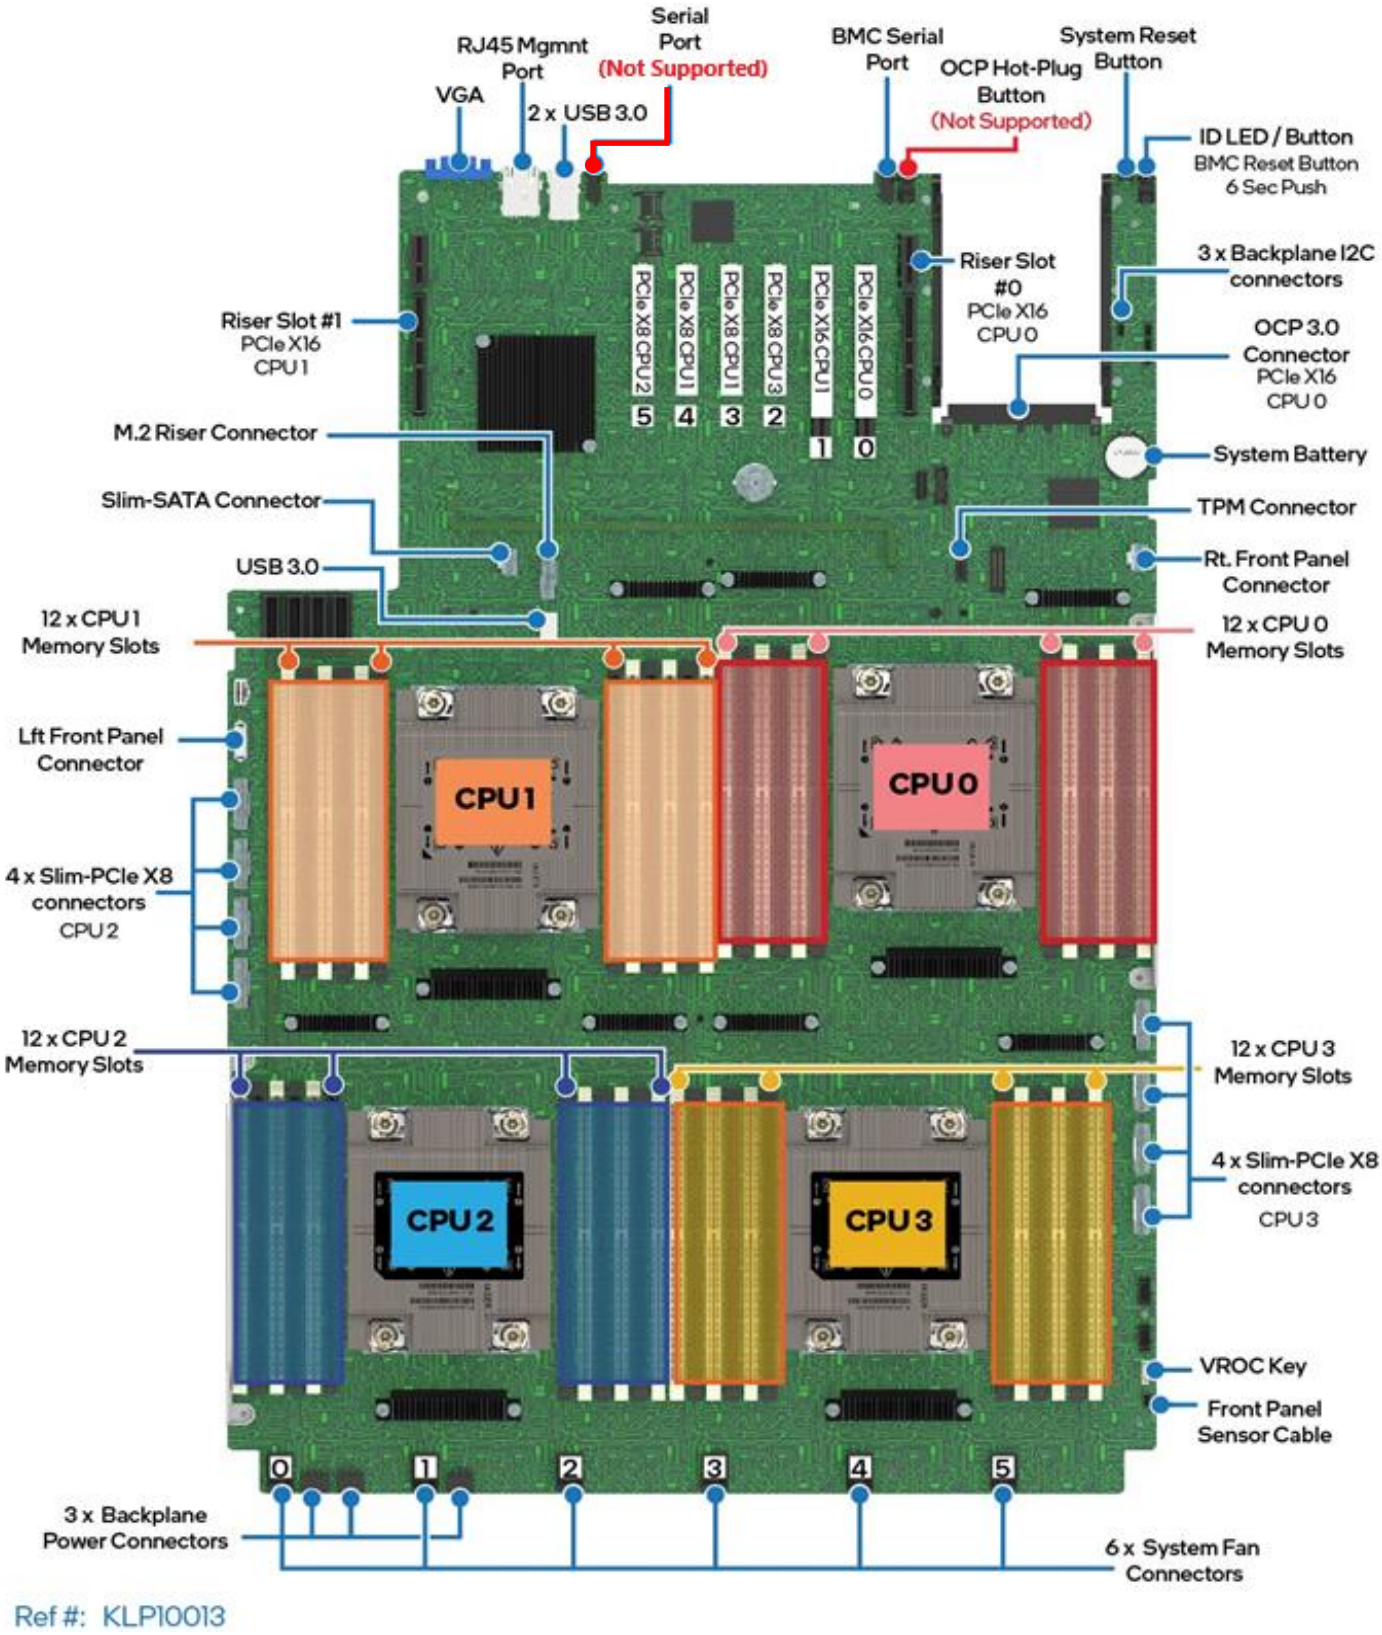

Figure 14. Server Board Features

2.2 Chassis Dimensions

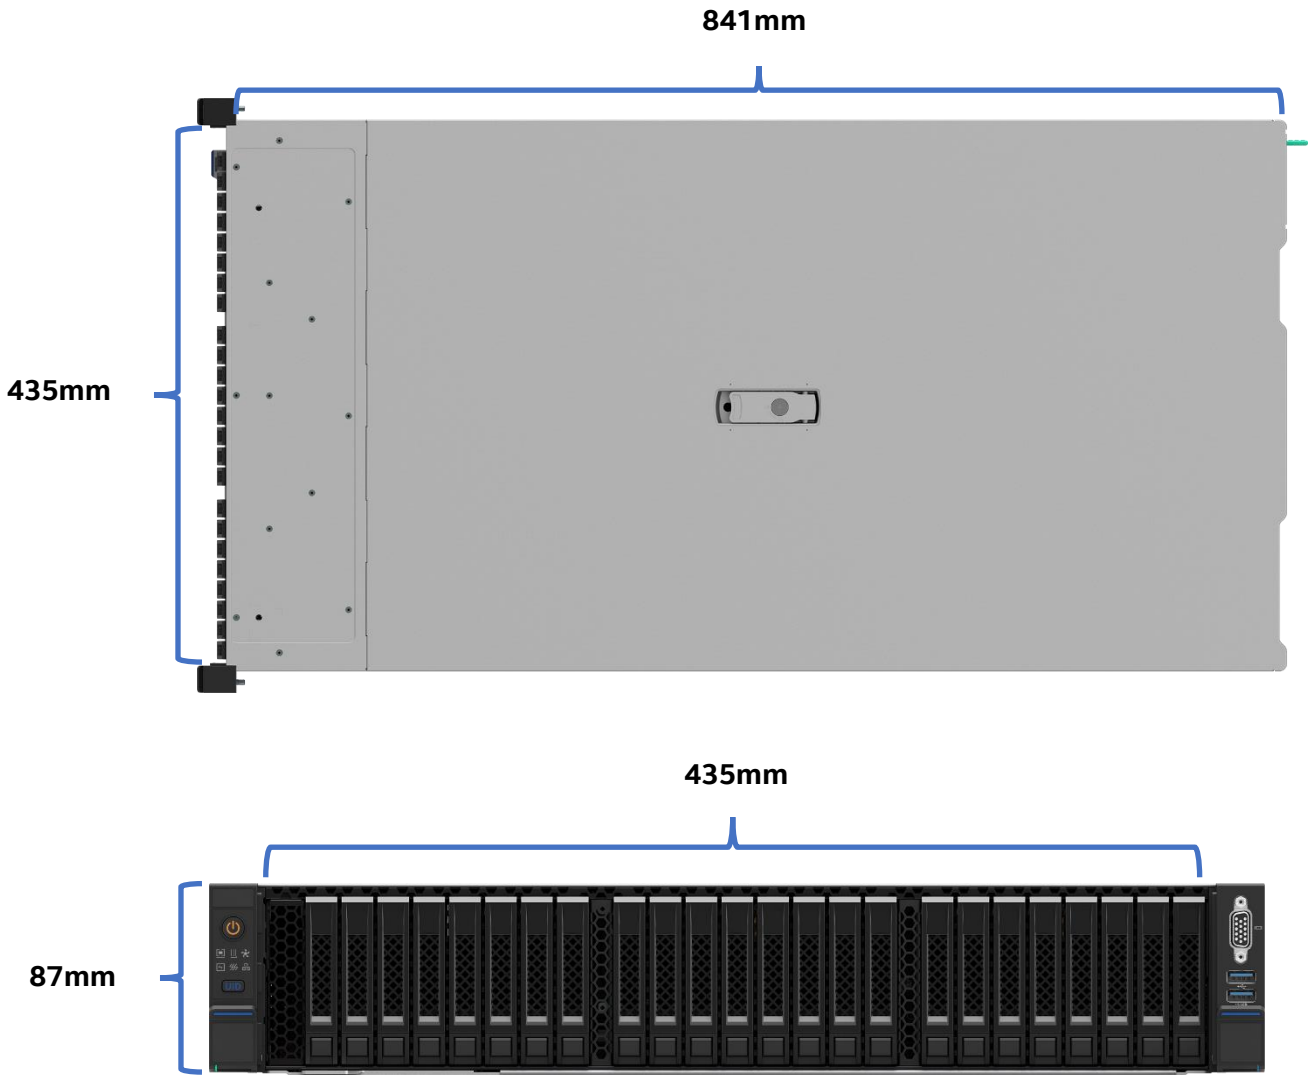

Figure 15. Chassis Dimensions

2.3 Rack and Cabinet Mounting Kit

The Intel® Server System M70KLP includes a rail kit for system installation into a 4-post rack or cabinet. The following installation guidelines should be observed.

- For proper system ventilation, leave a minimum of 15 cm clearance in the front and rear of the system.
- Servers are high-power electrical appliances. They should be installed into dedicated cabinets with vents or professional water-cooled cabinets to prevent system failures caused by overheating.
- If installing more than one server or component into a given rack or cabinet, begin installing them from the bottom and load the heaviest items first.
- Note the cabinet's load-bearing capacity, power source capacity, and heat dissipation capacity. Be sure not to install devices that go beyond the cabinet's capacity thresholds.
- For the convenience of using the front and rear ports of the system and to allow for cabling, leave a minimum clearance of 70 mm between the front of the server and the inside of the cabinet's front door. Leave 150 mm between the back of the system and the inner side of the cabinet's back door.
- Due to the weight of a system, Intel recommends carrying the system with two or more people supporting the system from the sides, using a mechanical lift, or a cart when moving the system from one location to another.
- If your system has rack handles installed, do not lift, or carry the system solely by the rack handles. These handles are intended for the sole purpose of pulling a system from or pushing it into a rack.
- When lifting or moving a system, it is best to grasp and lift it by all four corners using two or more people. Do not grasp and lift the system by two opposing diagonal corners. Doing so will flex the chassis that may damage the internal system components.
- With no other option available but to lift the system using only two points of contact, grasp and lift the system at the mid-point of each side of the system.

Download the *Intel® Server System M70KLP Family Service Guide* for complete installation instructions.

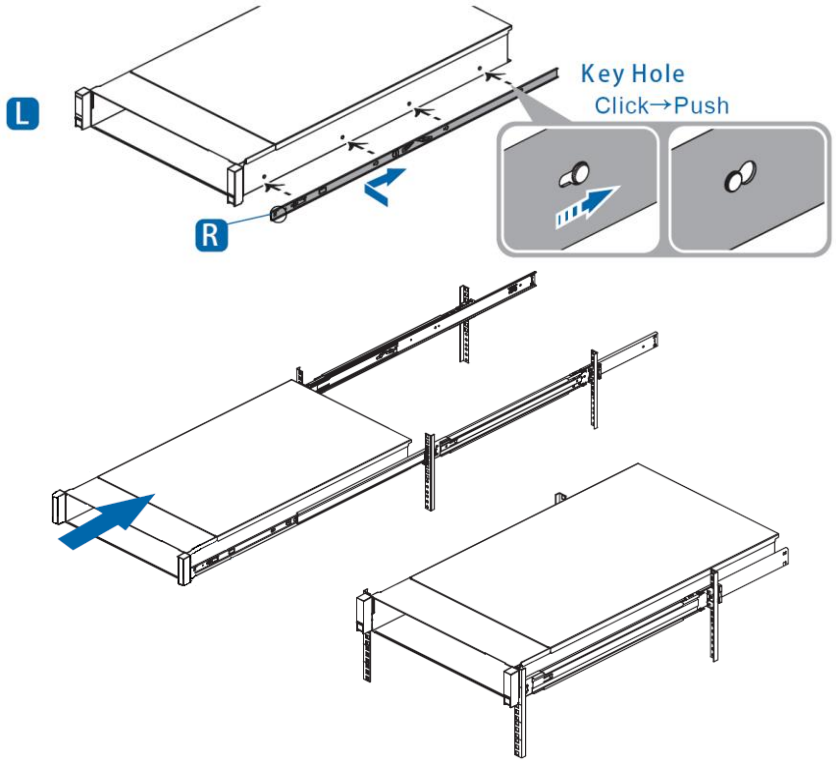

Figure 16. 4-post or Cabinet Rail Kit

Features and specifications for the rail kit are listed below:

- **Intel Product Order Code (iPC)** – **KLPRAILK** - Spare/Accessory Rail Kit
  - Tool-less attachment to system and installation into rack/cabinet
  - Rack installation front and rear post distance adjustment ranges from 610mm to 942mm
  - 890mm travel distance
  - Full extension from rack
  - Max supported weight: 38 Kg (86 lbs.)
  - Support for cable management arm accessory kit – iPC **KLP2UCMA**

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

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

2.4 System Environmental Limits Summary

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. 10–35 ° C (50–95 °F) with the maximum rate of change not to exceed 10 °C per hour.		
	Shipping	-40–70 °C (-40–158 °F)		
Altitude	Operating	Support operation up to 900 m with ASHRAE class de-ratings.		
Humidity	Shipping	50–90%, non-condensing with a maximum wet bulb of 28 °C (at temperatures 25–35 °C)		
Shock	Operating	Half sine, 2 G 11 msec		
	Unpackaged	Trapezoidal, 25 G, velocity change is based on packaged weight		
	Packaged	ISTA (International Safe Transit Association) Test Procedure 3A 2008 - TBD		
Vibration	Unpackaged	5–500 Hz, 2.20 G RMS random		
	Packaged	ISTA (International Safe Transit Association) Test Procedure 3A 2008 - TBD		
AC-DC	Voltage	90 V to 132 V and 180 V to 264 V		
	Frequency	47–63 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	8.0 kV		
	Contact Discharge	4.0 kV		
Acoustics Sound Power Measured	System Configuration	**Standard System Config**	**GPU Enabled System**	
	Server Sound Power Level (Idle)	Up to 7.5 LWA (dBA)	Up to 7.5 LWA (dBA)	
	Server Sound Power Level (70% Loading)	Up to 8.8 LWA (dBA)	Up to 8.8 LWA (dBA)	

2.5 Architectural Overview

The architecture of Intel® Server System M70KLP product family is developed around the integrated features and functions of the following components:

- 3rd Gen Intel® Xeon® Scalable Processor Family – Platinum 83xx and Gold 63xx (H) (HL) SKUs only
- Intel® C621A chipset (PCH)
- ASPEED* AST2500 baseboard management controller (BMC).

Figure 17 provides an overview of the server board architecture, showing the features and interconnects of each of the major subsystem components.

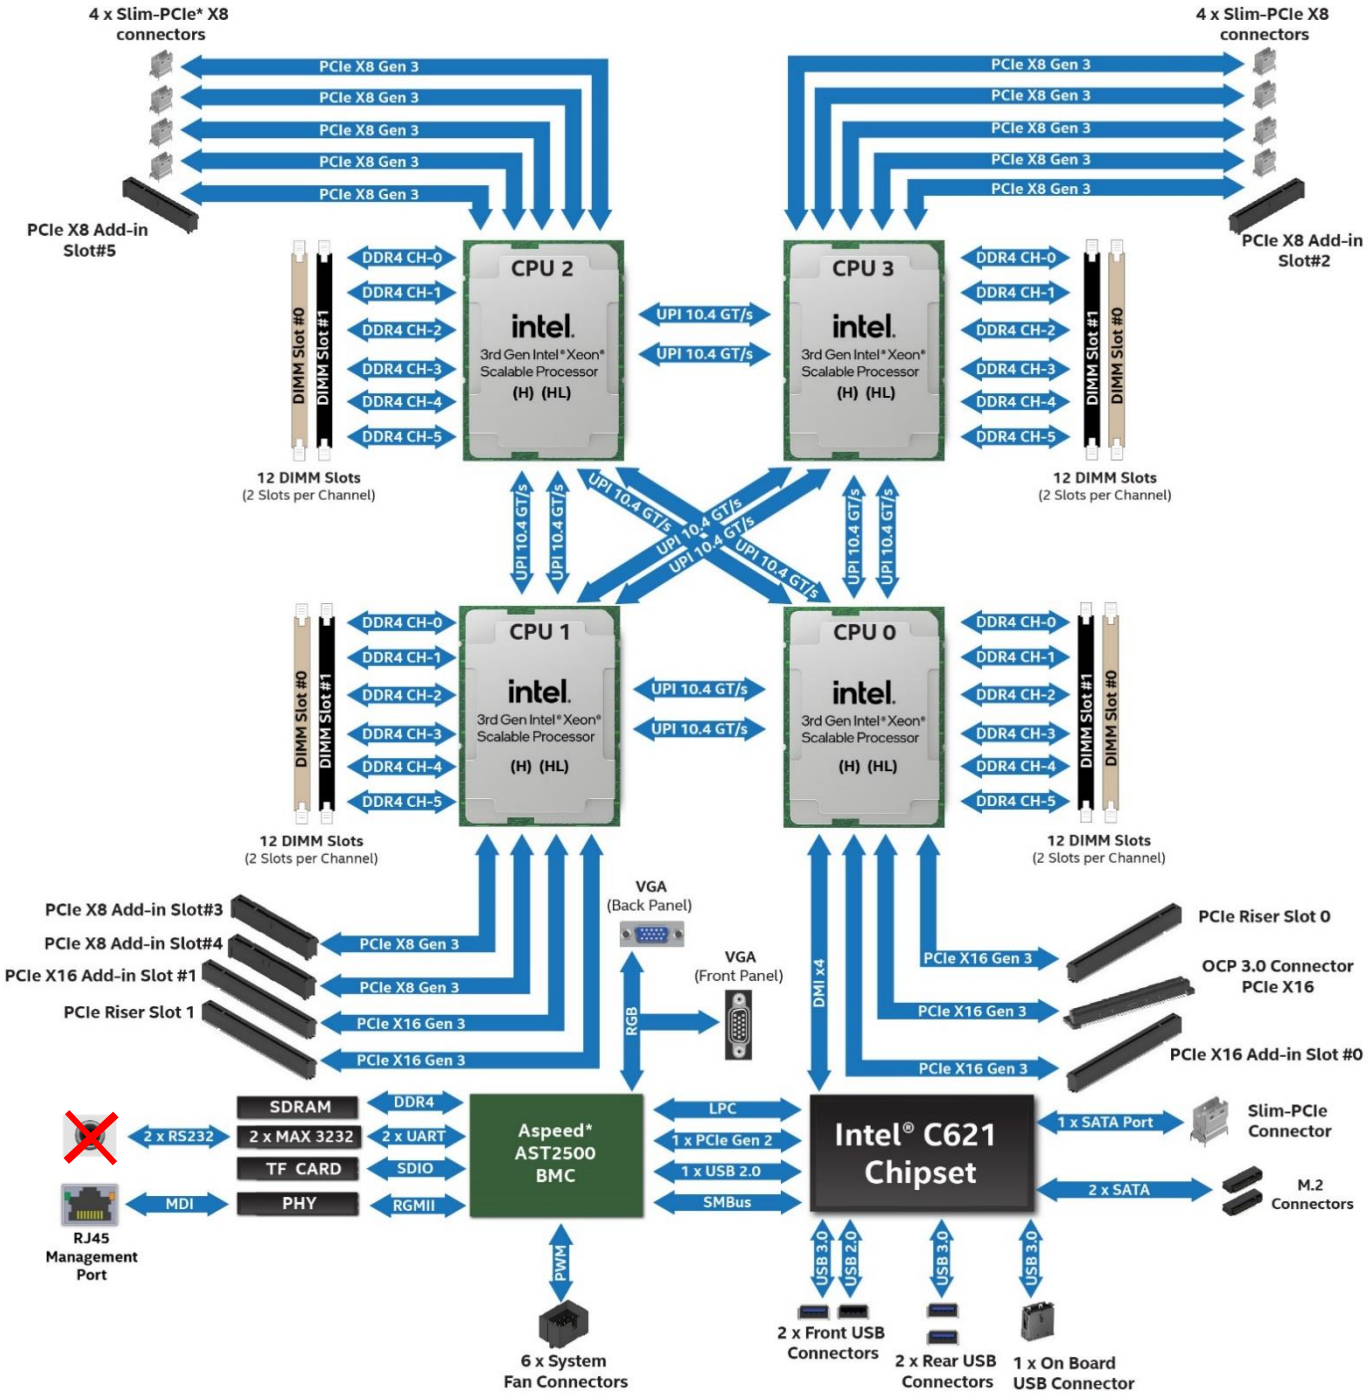

Figure 17. Architectural Block Diagram

3. System Software and Utilities

Several embedded and stand-alone software utilities are available to configure, customize, support, and update the Intel® Server System M70KLP. These utilities include:

- BIOS Setup utility
- AMI* BIOS Configuration Program (AMIBCP)
- AMI* ChangeLogo* utility
- Intel® System Configuration utility (SysCfg)
- System Update Package (SUP)
- Intel® System Firmware Update utility (SYSFWUPDT)
- Intel® System Information Retrieval utility (SYSINFO).

The following sections provide a brief description of each.

3.1 <F2> BIOS Setup Utility

The BIOS Setup utility is an embedded text-based utility used to view and configure system settings that determine how the system operates, set security features, set remote server management access parameters, set boot management options, and access the error management display screen.

To enter the BIOS Setup 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

When the BIOS Setup utility is first entered, the utility's Main Menu page is displayed, providing system information and access to various sub-menus. However, should a serious POST error occur, accessing the BIOS Setup utility will automatically redirect the initial display to the Error Manager screen within the BIOS Setup utility, where additional error information can be found.

It is also possible to boot directly to the BIOS Setup utility using the IPMI 2.0 command *Get/Set System Boot Options*. For details, see the IPMI 2.0 specification.

For further BIOS Setup Utility information, download the *Intel® Server System M70KLP BIOS Setup Utility User Guide*.

3.2 AMI* BIOS Configuration Program (AMIBCP) – Creating a BIOS Image for Customized Settings

The AMI* BIOS Configuration Program (AMIBCP) is a powerful customization utility that enables OEMs/ODMs to customize the BIOS ROM image without intervening on the source code and rebuilding the BIOS. With AMIBCP, it is possible to obtain multiple ROM image flavors ready for production.

The AMIBCP utility for this system is custom and is only available under NDA and upon request from Intel.

3.3 AMI* ChangeLogo Utility – Customizing the BIOS Splash Screen

The AMI* ChangeLogo utility allows developers to easily change splash screen logos displayed by BIOS at boot using the GUI or CLI. The full screen "splash" logo and small logos appearing on the main screen during POST can be replaced with custom logos.

The ChangeLOGO utility for this system is custom and is only available under NDA and upon request from Intel.

3.4 Intel® System Configuration Utility (SYSCFG)

The Intel® System Configuration Utility (SYSCFG) is a command-line utility that supports the following features:

- Save selective BIOS and/or firmware settings to a file
- Write BIOS and Firmware settings from a file to a server
- Configure selected firmware settings
- Configure selected BIOS settings
- Configure selected system settings
- Display selected firmware settings
- Display selected BIOS settings

Generally, this Intel utility can be used on any Intel server platform. However, the Intel® Server System M70KLP has a version of this utility that is only compatible for use on this platform. The SYSCFG utility used on other Intel server systems is NOT compatible with this server system.

Two ways to identify the M70KLP version of the utility are:

- Check the utility zip pkg name string
  The M70KLP Utility zip pkg string includes a '**klp**' identifier.
  Example: *Syscfg_V14_2**klp**_Build*x*_AllOS.zip*

- Check the utility header after launching the tool
  The M70KLP utility screen header contains the '**klp**' identifier as shown in Figure 18.

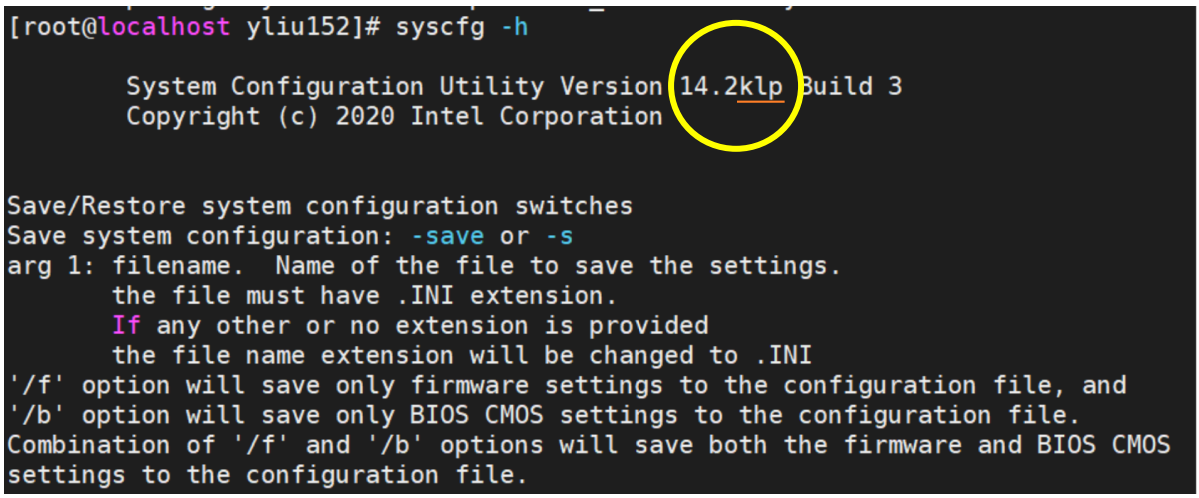

Figure 18. SYSCFG Utility for Intel® Server System M70KLP – Utility Header Screenshot

For further SYSCFG utility information, download the *Intel® System Configuration Utility User Guide.*

3.5 System Update Package (SUP) for Intel® Server System M70KLP

The SUP is a set of uEFI-based utilities and files bundled together and used to update the system BIOS and other embedded system firmware. Included within the compressed file package is a README file providing complete system update instructions and a STARTUP.NSH script file that automates the entire system update process with little or no user intervention.

System Update Advisory: This system supports Intel® Platform Firmware Resilience (Intel® PFR – see Section 11.3). In support of this security feature, the update procedure and behavior of the system during the system update process have changed from previous generation Intel server systems.

Updating the full system software stack may take 20–30+ minutes depending on which components are being updated using the recommended update procedure (automated using the STARTUP.NSH script file).

Using the included STARTUP.NSH script, the system will reboot and appear to power off for approximately 10+ minutes for each of the following updates: BIOS, BMC FW, and CPLD. Throughout each reboot period, the display screen may be blank and all front panel power and reset buttons will be disabled. This is normal and expected behavior. **DO NOT disconnect the AC power cords or remove the USB storage device holding the update files from the system or attempt to reboot the system yourself**. The system will power back on automatically after each reboot period and will resume the system update process until all the updates have successfully completed.

3.6 Intel® System Firmware Update Utility (SYSFWUPDT)

The System Firmware Update utility provides the ability to update the system BIOS and Firmware while the server is running its host operating system. This utility is a command-line tool and it requires users to have administrator (Windows) or root (Linux) privileges.

The SYSFWUPDT utility supports following features.

- BIOS Update – Update PFR BIOS in the system, tool transfers the bin file to BMC and the real update will start on next reboot by default.
- BMC Update - Update Server Management (SM) firmware (FW) of the Baseboard Management Controller (BMC), and on next BMC reset the new BMC firmware will be loaded.
- CPLD Update
- Intel® Optane™ PMem Firmware Update (If system is configured with devices)
- Recovery Update
- Modify specific FRU field
- Display BIOS/Intel ME/BMC/Base Board/System/FRU/SDR/SMBIOS information
- Restore BIOS Default setting

Note: When the system is rebooted to perform the update for the BIOS, BMC FW, or CPLD, the system may appear to power off. The display screen will be blank (dark) and the front panel buttons will be disabled. This is expected behavior and can last for up to 10 minutes for each update. Do not unplug the system from AC Power and do not disrupt the update process in any way. Wait for the system to complete each update and automatically reboot the system.

Generally, this Intel utility can be used on any Intel server platform. However, the Intel® Server System M70KLP has a version of this utility that is only compatible for use on this platform. The SYSFWUPDT utility used on other Intel server systems is NOT compatible with this server system.

There are two ways to identify the M70KLP version of the utility:

- Check the utility zip pkg name string
  The M70KLP Utility zip pkg string contains the '**klp**' identifier.
  Example: Sysfwupdt_V14_2**klp**_Buildx_AllOS.zip

- Check the utility header after launching the tool
  The M70KLP utility screen header contains the '**klp**' identifier as shown in Figure 19.

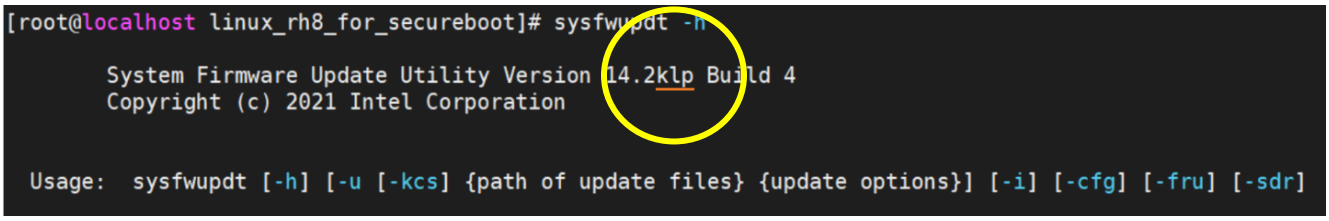

Figure 19. SYSFWUPDT Utility for Intel® Server System M70KLP – Utility Header Screenshot

For further SYSFWUPDT Utility information, download the *Intel® System Firmware Update Utility User Guide.*

3.7 Intel® System Information Retrieval Utility (SYSINFO)

The System Information Retrieval Utility is a command-line utility that provides the ability to collect system information, as fully described in the IPMI and BMC specifications. Running the utility requires that the user have Windows* administrator or Linux* root permissions.

The SYSINFO Utility collects the following system information and writes the data to a log file:

- Platform Firmware Inventory
- Sensors
- Sensor Data Records (SDR)
- Baseboard FRU
- System Boot Order
- BMC User Settings
- BMC LAN Channel Settings
- BMC SOL Channel Settings
- BMC Power Restore Policy Settings
- BMC channel settings
- SMBIOS Type 1, Type 2, Type 3
- Memory
- Processor
- Storage Devices - Hard Disk Drives (HDD) and Solid State Drives (SSD)
- Operating System Information
- Device Manager Information (such as drivers)
- List of Software Installed
- Operating System Event Log
- PCI Bus Device Information
- RAID settings and RAID log
- BIOS Settings (per the BIOS setup)
- Power Telemetry (if available)

Generally, this Intel utility can be used on any Intel server system. However, the Intel® Server System M70KLP has a version of this utility that is only compatible for use on this server product family. The SYSINFO utility used on other Intel server systems is NOT compatible with this server product family.

There are two ways to identify the M70KLP version of the utility:

- Check the utility zip file name string
  The M70KLP Utility zip file name string contains the '**klp**' identifier.
  Ex: Sysinfo_V14_2**klp**_Buildx_AllOS.zip

- Check the utility screen header after launching the tool
  The M70KLP utility screen header contains the '**klp**' keyword as shown in Figure 79.

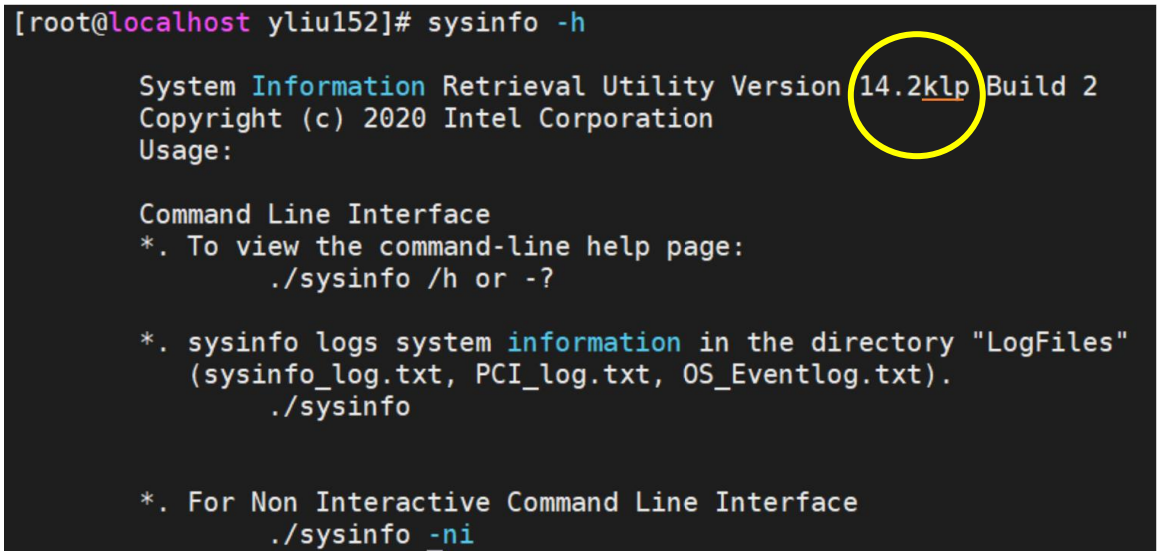

Figure 20. SYSINFO Utility for Intel® Server System M70KLP – Utility Header Screenshot

For further SYSINFO Utility information, download the Intel® System Information Retrieval Utility *User Guide.*

4. Processor Support

The server board integrated within the Intel® Server System M70KLP product family includes four Socket P+ 4189 pin processor sockets identified on the server board as CPU 0 through CPU 3, as shown below.

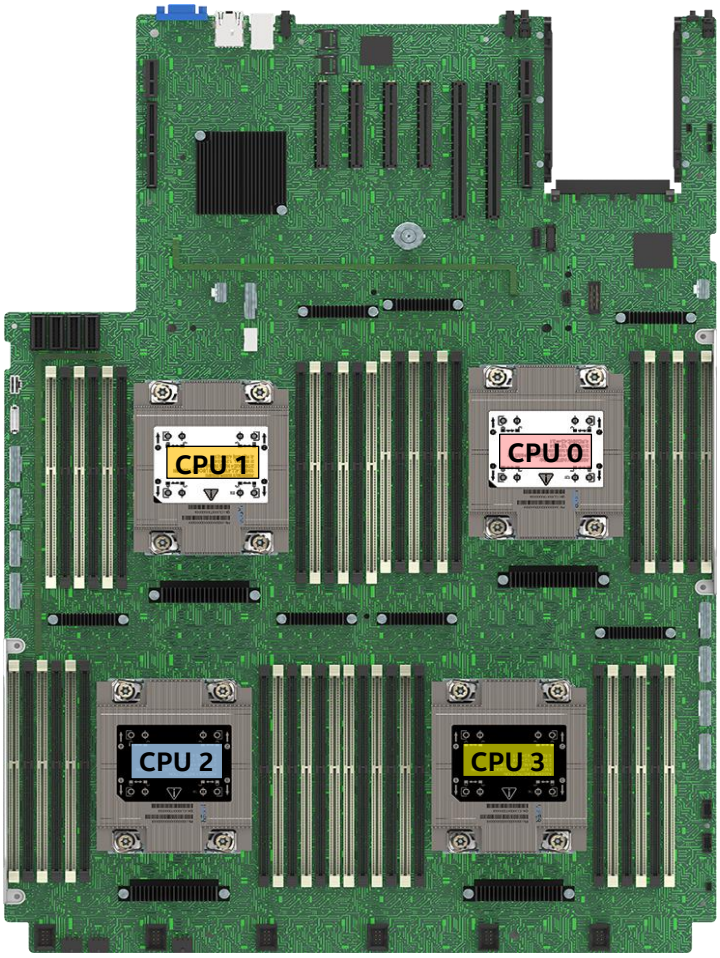

Figure 21. Processor Socket Identification

The system includes four (4) processor heat sinks. The heat sinks for CPU2 and CPU3 (CPUs closest to front of the system) support high airflow and are identified by their Black labels. The heat sinks for CPU 0 and CPU1 (CPUs closest to back of the system) have White labels. White and Black heat sinks must be installed to their designated CPU sockets to support required system airflow requirements.

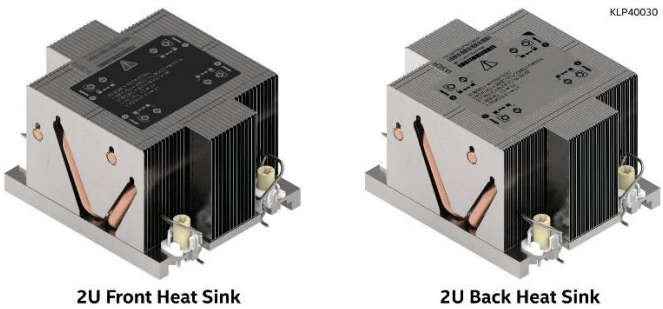

Figure 22. Processor Heat Sink Identification

In the Standard system option, all four processor heat sink modules (PHM) use 2U heat sinks. In the GPU enabled system option, all four PHMs use 1U heat sinks.

4.1 Processor Heat Sink Module (PHM)

Processors in this system are part of an assembly referred to as a PHM (Processor Heat sink Module). A PHM consists of a processor, a processor carrier clip, and the processor heat sink, which is preassembled into a single module before placement onto the processor socket assembly on the server board. The PHM concept reduces the risk of damaging pins within the processor socket during the installation process. Reference the *Intel® Server System M70KLP Service Guide* for complete PHM assembly and installation instructions.

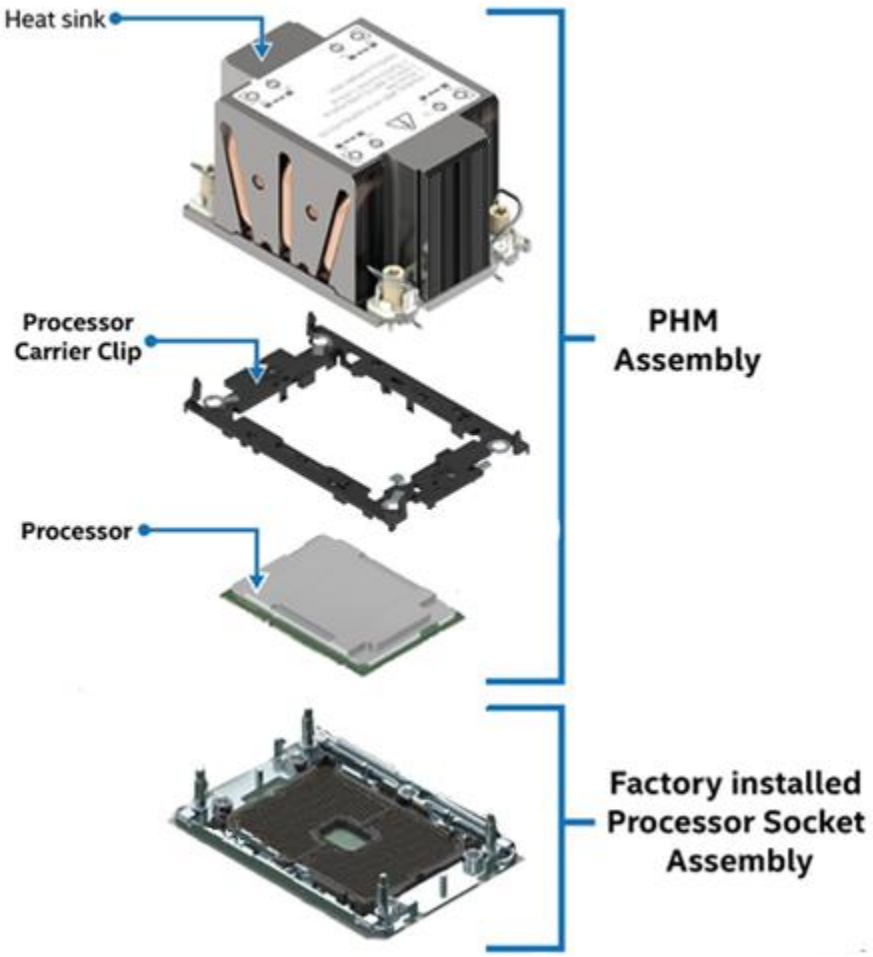

Figure 23. Processor Heat Sink Module (PHM) Assembly

4.2 3rd Gen Intel® Xeon® Scalable Processor Family – Platinum, Gold – H, HL

The Intel® Server System M70KLP product family supports the 3rd Gen Intel® Xeon® Scalable processor family. Within this processor family, only Platinum and Gold processor SKUs that end in H or HL with a TDP of ≤ 250W will be supported. All other processor SKUs within the 3rd Gen Intel® Xeon® Scalable processor family and all previous generation Intel® Xeon® processors are not supported.

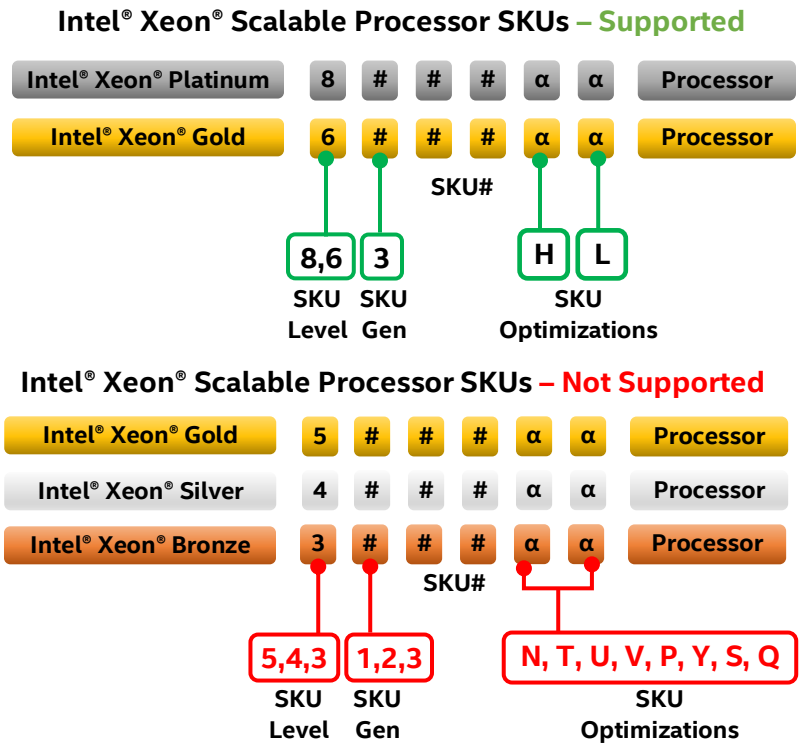

Figure 24. Intel® Xeon® Scalable Processor SKU Compatibility Chart

Table 4. Processor Features

Processor Feature	Detail
Maximum TDP[1]	250W
L1D Cache	32 KB/core
L2 Cache (MLC)	1 MB/core
L3 Cache (LLC)	1.375 MB/core
Node controller support	No
New instructions	VNNI, Bfloat16
Intel® SGX and Intel® TME security features	No
# DDR4 channels	6
DDR4 max speeds[1]	3200
# UPI	6
UPI speeds	Up to 10.4 GT/s
Intel® Optane™ Persistent Memory Support	Intel® Optane™ Persistent Memory 200 Series. ** Only App Direct Mode is supported
PCIE I/O	PCIe 3.0, 48 lanes (x16, x8, x4)
Physical/Virtual Address Bits	46/48
SKU shelves	SKUs in Platinum, Gold

Note: (1) Feature may vary between processor SKUs.

Reference 3rd Gen Intel® Xeon® Scalable Processor specification sheets and product briefs for additional information

4.3 Processor Population Rules

When more than one processor is installed, the following population rules apply:

- All processors must have identical extended family, extended model number and processor type
- All processors must have the same number of cores
- All processors must have the same cache sizes for all levels of processor cache memory
- All processors must support identical DDR4 memory frequencies

Note: Processors with different steppings can be mixed in a system if the rules mentioned above are met.

Population rules are applicable to any combination of processors within the 3rd Gen Intel® Xeon® processor Scalable family.

Note: The server board may support processor configurations consisting of different processors that meet the defined criteria. However, Intel does not perform validation testing of this configuration. In addition, Intel does not ensure that a server system configured with unmatched processors will operate reliably. The system BIOS attempts to operate with processors that are not matched but are generally compatible. For optimal system performance in multi-processor configurations, Intel recommends that identical processors be installed.

5. Memory Support

This chapter describes the architecture that drives the memory subsystem, supported memory types, memory population rules, and supported memory RAS features.

The Intel® Server System M70KLP has support for up to 48 DDR4 DIMMs; 12 DIMM slots per processor.

Each processor supports six memory channels identified as 0 through 5. Each memory channel has support for two DIMM slots identified as Slot 0 (White slots) and Slot 1 (Black slots).

On the server board, DIMM slots are identified by CPU #, Memory Channel #, and Slot #.

Examples: (CPU 0, Ch 0, Slot 0), (CPU 0, Ch 0, Slot 1), (CPU 0, Ch 1, Slot 0), etc.

DIMM slots adjacent to their respective processor have the following layout on the server board.

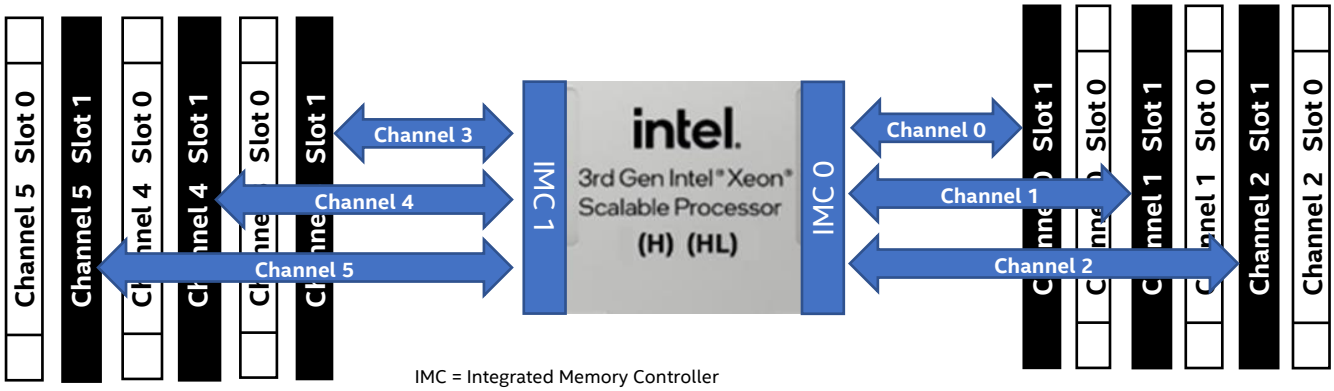

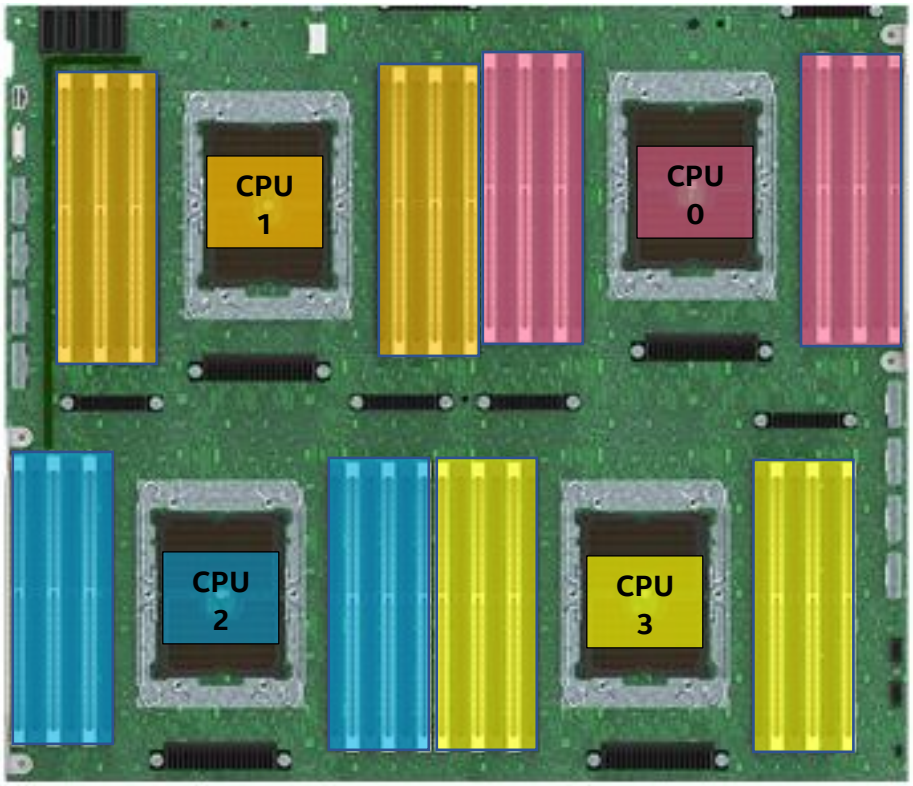

Figure 25. DIMM Slot Layout

5.1 Supported Memory

The Intel® Server System M70KLP supports standard DDR4 RDIMMs and LDRIMMs, and Intel® Optane™ persistent memory 200 series modules.

Note: Previous generation Intel® Optane™ persistent memory modules are not supported.

5.1.1 Standard DDR4 DIMM Support

The following figure shows a standard DDR4 DIMM module.

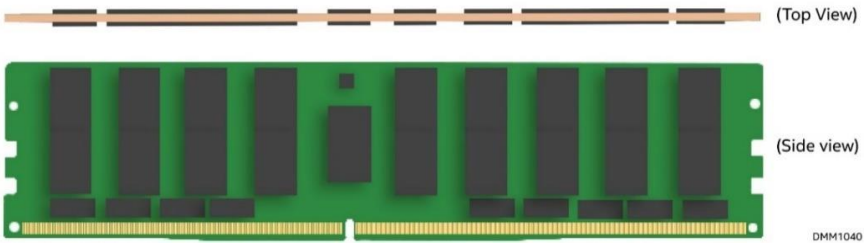

Figure 26. Standard SDRAM DDR4 DIMM Module

The server supports DDR4 DIMMs with the following specifications:

- All DDR4 DIMMs must support ECC
- Registered DDR4 (RDIMM), 3DS-RDIMM, Load Reduced DDR4 (LRDIMM), 3DS-LRDIMM
  Note: 3DS = 3-Dimensional Stacking
- RDIMMs and LRDIMMs with thermal sensor On DIMM (TSOD)
- DIMM speeds of up to 3200 MT/s
- DIMM capacities of 8 GB, 16 GB, 32 GB, 64 GB, and 128 GB
- RDIMMs organized as Single Rank (SR), Dual Rank (DR)
- 3DS-RDIMM organized as Quad Rank (QR), or Oct Rank (OR)
- LRDIMMs organized as Quad Rank (QR)
- 3DS-LRDIMM organized as Quad Rank (QR), or Oct Rank (OR)

Table 5. Supported DDR4 DIMM Memory

Type	Ranks per DIMM and Data Width	DIMM Capacity (GB)		Maximum Speed (MT/s) at 1.2 V	
		8 Gb DDR4 Density	16 Gb DDR4 Density	1 DPC	2 DPC
RDIMM	SR x8	8	16	3200	3200
	SR x4	16	32	3200	3200
	DR x8	16	32	3200	3200
	DR x4	32	64	3200	3200
3DS-RDIMM	QR/OR x4	64 (2H) 128 (4H)	64 (2H) 128 (4H)	3200	3200
LRDIMM	QR x4	64	128	3200	3200
3DS-LRDIMM	QR/OR x4	128 (4H)	128 (2H)	3200	3200

Note: SR = Single Rank, DR = Dual Rank, QR = Quad Rank, OR = Oct Rank, DPC= DIMMs per Channel

Table 6. Maximum Supported Standard SDRAM DIMM Speeds by Processor Shelf

3rd Gen Intel® Xeon® Scalable processor family	Maximum DIMM Speed (MT/s) by Processor Shelf			
	Platinum 83xx Processors (H) (HL)	Gold 63xx Processors (H) (HL)	Gold 53xx Processors	Silver 43xx Processors
Supported Maximum DIMM Speed	3200	2933	Not supported	Not supported

5.1.2 Intel® Optane™ Persistent Memory 200 Series Support

The server system supports Intel® Optane™ persistent memory (PMem) 200 series.

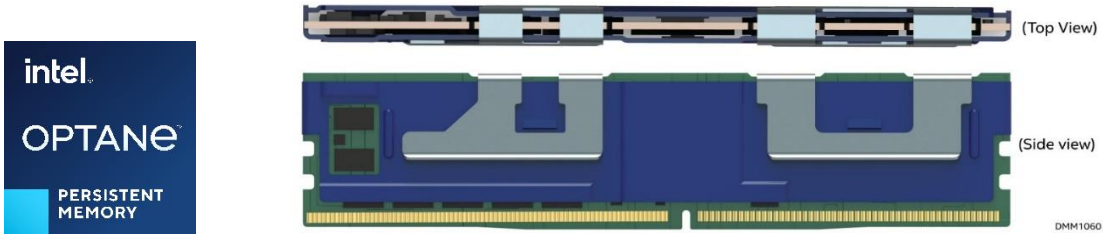

Figure 27. Intel® Optane™ Persistent Memory 200 Series

Intel® Optane™ PMem is an innovative technology that delivers a unique combination of affordable large memory capacity and data persistence (non-volatility). It represents a new class of memory and storage technology architected specifically for data center usage. Intel® Optane™ PMem 200 series enables higher density (capacity per DIMM) DDR4-compatible memory modules with near-DRAM performance and advanced features not found in standard SDRAM.

Intel® Optane™ PMem 200 Series support the following features:

- DDR4 Pin Compatible
- Available PMem Capacities – 128, 256, 512 GB
- Up to 3 TB per processor socket
- Up to 2666 MT/sec
- TDP = 15W
- AES256 Bit Encryption
- Secure Erase
- Data persistence in power failure event – ADR, eADR (optional)

Supported operating modes:

- App Direct (AD) mode

5.1.2.1 Intel® Optane™ Persistent Memory 200 Series – Memory Mode (MM)

Not Supported

5.1.2.2 Intel® Optane™ Persistent Memory 200 Series – App Direct (AD) Mode

In App Direct mode, the operating system sees Intel® Optane™ persistent memory and DDR4 SDRAM DIMMs as two separate pools of memory. It can direct which type of data read or write is suitable for SDRAM or Intel® Optane™ PMem. Operations that require the lowest latency and do not need permanent data storage can be executed on SDRAM DIMMs, such as database “scratch pads”. Data that needs to be made persistent or structures that are very large can be routed to Intel® Optane™ persistent memory. The App Direct mode must be used to make data persistent in memory. This mode requires an operating system or virtualization environment enabled with a persistent memory-aware file system.

App Direct mode requires both driver and explicit software support. To ensure operating system compatibility, visit https://www.intel.com/content/www/us/en/architecture-and-technology/optane-memory.html.

5.1.2.3 Intel® Optane™ PMem configuration using the <F2> BIOS Setup Utility

Following the installation of Intel® Optane™ PMem devices into the system, they need to be configured using the <F2> BIOS Setup utility. The BIOS Setup utility includes several Intel® Optane™ PMem configuration options across multiple BIOS Setup screens. The following illustration provides the BIOS Setup screen navigation directing the user to the main Intel® Optane™ PMem configuration screen.

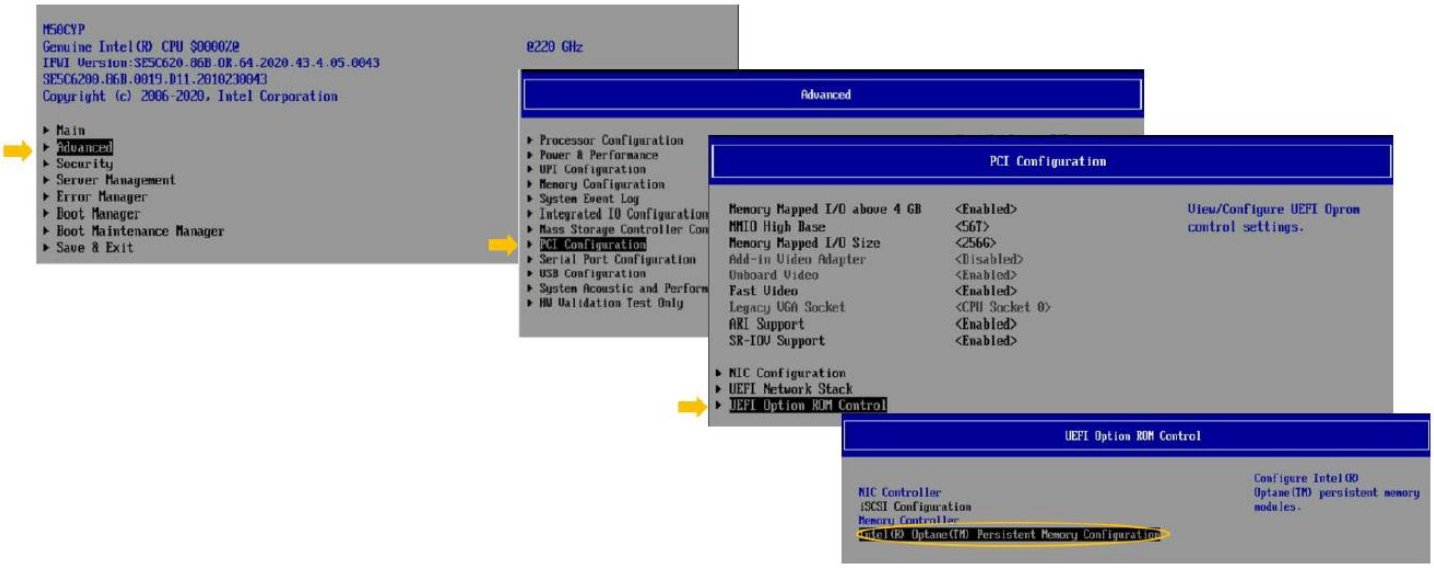

Figure 28. <F2> BIOS Setup Screen Navigation for Intel® Optane™ PMem Setup Options

The main Intel® Optane™ PMem Configuration screen provides links to the various device information and setup screens.

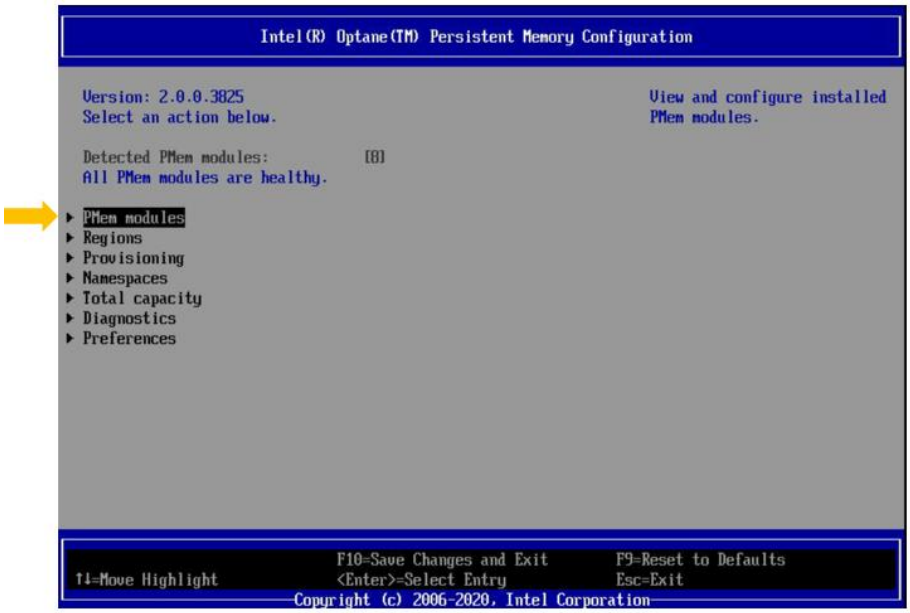

Figure 29. Intel® Optane™ PMem Configuration Menu in <F2> BIOS Setup

Refer to the appropriate Intel® Optane™ Persistent Memory setup documentation and Intel® Server System M70KLP BIOS Setup User Guide for further information regarding Intel® Optane™ PMem configuration and security settings.

5.2 Memory Population Overview

The Intel® Server System M70KLP includes 48 memory slots (12 slots per processor) and can support memory configurations that consist of only DDR4 SDRAM DIMMs or can support configurations consisting of both DDR4 SDRAM DIMMs and Intel® Optane™ persistent memory.

The following sections provide memory population rules that will result in best performance and compatibility for both configuration types.

5.2.1 DDR4 DRAM DIMM Population Rules

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 #
DRAM 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

The following DDR4 SDRAM population rules apply for most reliable operation.

- Mixed DDR4 DIMM rules:

Note: Intel only supports mixed DDR4 DRAM DIMM configurations as defined in the ***Intel DDR4 Support Disclaimer*** above.

  - Mixing DDR4 DIMMs of different speeds and latencies is not supported within or across processors. If a mixed configuration is encountered, the BIOS attempts to operate at the highest common speed and the lowest latency possible.
  - x4 and x8 DDR4 DIMMs may be mixed in the same channel.
  - Mixing of DDR4 DIMM types (RDIMM, LRDIMM, 3DS-RDIMM, 3DS-LRDIMM) within or across processors is not supported. This situation is a Fatal Error Halt in Memory Initialization.
- For a single DDR4 DIMM in a dual-slot channel, populate slot 0 (white slot).

Chan 5		Chan4		Chan 3					Chan 0		Chan 1		Chan 2	
Slot 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	IMC 1	CPU	IMC 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	Slot 0
DDR4		DDR4		DDR4						DDR4		DDR4		DDR4

- For multiple DDR4 DIMMs per channel:
  - For RDIMM, LRDIMM, 3DS-RDIMM, and 3DS-LRDIMM, always populate DIMMs with higher electrical loading in slot 0 (white slot).
  - When populating a quad-rank DDR4 DIMM with a single- or dual-rank DDR4 DIMM in the same channel, the quad-rank DDR4 DIMM must be populated farthest from the processor. Incorrect DIMM placement results in an MRC error code. A maximum of 8 logical ranks can be used on any one channel, as well as a maximum of 10 physical ranks loaded on a channel.
- Memory slots associated with a given processor socket cannot be used if no processor is installed within the socket.
- Processor sockets are self-contained and autonomous. However, all memory subsystem support (such as memory RAS and error management) in the BIOS Setup are applied commonly for each installed processor.
- For best system performance, memory must be installed in all six channels for each installed processor.
- For best performance, when NOT populating all memory channels, DDR4 DIMMs must be populated symmetrically between IMCs into memory slots on both sides of the CPU for each installed CPU
  - Example: when populating 4 DDR4 DIMMs to a given CPU, the DIMMs must be installed to memory channels 0, 1, 3, and 4 as shown below.

Chan 5		Chan4		Chan 3					Chan 0		Chan 1		Chan 2	
Slot 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	IMC 1	CPU	IMC 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	Slot 0
		DDR4		DDR4						DDR4		DDR4		

- For best system performance in a multi-processor configuration, the installed DDR4 DIMM population must be the same for each installed processor.

5.2.2 Intel® Optane™ Persistent Memory 200 Series Population Rules

All operating modes:

- Platform requires a maximum of 1 PMem module per memory channel and at least one PMem module per populated CPU.
- All populated CPUs within a platform must have an identical PMem population.
- Mixing PMem modules of different capacities anywhere on the platform, across or within same CPU socket, is NOT supported.
- When populating a memory channel with both a DDR4 DRAM DIMM and a PMem module, populate the PMem module into the Slot 1 (Black) memory slot.

Chan 5		Chan4		Chan 3		IMC 1	CPU	IMC 0	Chan 0		Chan 1		Chan 2	
Slot 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1				Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	Slot 0
DDR4	PMem	DDR4	PMem	DDR4	PMem				PMem	DDR4	PMem	DDR4	PMem	DDR4

- When populating a memory channel with only a PMem module, populate the Slot 0 (White) memory slot.
- When populating a single PMem module for each installed CPU, the PMem module must be installed into memory channel 0
- Mixing of different DDR4 DRAM DIMM types on the platform is not supported nor validated. Ensure all DDR4 and PMem capacity size and types across all CPUs on the platform are the same.

Memory mode:

- Not Supported

App direct mode:

- A minimum of one DDR4 DRAM DIMM per memory channel for each installed CPU.
- A minimum of one PMem module per system.
- When populating DDR4 DRAM + PMem, Intel recommends devices have a DDR4 DIMM capacity to PMem module capacity ratio of 1:1, 1:2, or 1:4

  Examples)

  - 1:1 = 128 GB DRAM DIMM: 128 GB PMem Module
  - 1:2 = 64 GB DRAM DIMM: 128 GB PMem Module
  - 1:4 = 32 GB DRAM DIMM: 128 GB PMem Module

- PMem modules must be populated symmetrically into memory slots on both sides of the CPU
  - Example: when populating 2 PMem modules to a given CPU, the PMem modules must be installed to memory channels 0 and 3 as shown below.

Chan 5		Chan4		Chan 3		IMC 1	CPU	IMC 0	Chan 0		Chan 1		Chan 2	
Slot 0	Slot 1	Slot 0	Slot 1	Slot 0	Slot 1				Slot 1	Slot 0	Slot 1	Slot 0	Slot 1	Slot 0
				DDR4	PMem				PMem	DDR4				

5.3 Memory RAS Support

The Intel® Server M70KLP family supports Reliability, Availability, Serviceability (RAS) features for both DDR4 memory and Intel® Optane™ persistent memory 200 series modules. The following sections provide an overview of supported RAS features associated with each memory type.

5.3.1 DDR4 Memory RAS Features Overview

The embedded memory controllers of the 3rd Gen Intel® Xeon® Scalable processor support memory RAS features. Depending on the processor SKU, Memory RAS feature support varies between Standard and Advanced options, as identified in the following table. All processor SKUs supported by the Intel® Server System M70KLP70 product family have support for the Advanced memory RAS features.

Table 8. DDR4 Memory RAS Features

RAS Feature	Description	Standard RAS SKU	Advanced RAS SKU
Memory Single Device Data Correction (SDDC)	SDDC feature allows managing DRAM persistent (hard) failures where the whole DRAM device is failed. UEFI-FW invokes HW capability to copy the failed device content to the ECC device (also known as device sparing on to the ECC device). Subsequent soft error detection continues, but hardware cannot correct further errors (this is equivalent of legacy device-tagging capability). At this stage, immediate service call is required. In case of 'advanced RAS' SKUs of the processor, additional single bit error detection and correction is feasible and immediate service call is not required. Supported with x4 DIMMs only.	Yes	Yes
Adaptive Data Correction - Single Region (ADC-SR)	ADC-SR feature allows managing DRAM persistent (hard) failures using a technique called Adaptive Virtual Lockstep (AVLS) provided the hard failure is confined to a bank region. If first hard failure within a DRAM device occurs at bank region granularity then it is mapped out using 'bank level AVLS' method. UEFI-FW invokes this method using 'buddy bank'. Further soft error correction continues normally. If second hard failure occurs within a DRAM device in the same channel, then UEFI-FW can invoke SDDC flow as described in R_MEM01. Platform specific firmware is required to identify hard failure at bank region granularity. ADC-SR is supported only with x4 DRAM based DIMMs and can be activated on a per channel basis. This feature cannot be enabled along with "DDR4 Memory Multi-rank Sparing" or DDR4 Write Data CRC Check and Retry.	Yes	Yes
Adaptive Double Device Data Correction - Multiple Region (ADDDC-MR, +1)	The "ADDDC-MR,+1" feature allows managing DRAM persistent (hard) failures more efficiently as compare to the ADC-SR feature. Up to two DRAM hard failures can be corrected within different bank, rank, or DIMM and occur in a time-staggered manner. Both DRAM hard failures map out the affected bank/rank region by invoking adaptive virtual lockstep (AVLS) algorithm thus creating spare bank/rank region for replacement. If third hard failure occurs within a DRAM device in the same channel, then UEFI-FW can invoke SDDC flow. Subsequently, single bit errors can still be detected and corrected. Platform specific firmware is required to identify hard failure at bank/rank region granularity. ADDDC-MR is supported only with x4 DRAM based DIMMs and can be activated on a per channel basis.	No	Yes
DDR4 Command/ Address Parity Check and Retry	DDR4 CMD/ADDR parity check and retry with following attributes: CMD/ADDR Parity error "address" logging CMD/ADDR Retry In Mirror mode, the controller will fail over to secondary channel.	Yes	Yes

RAS Feature	Description	Standard RAS SKU	Advanced RAS SKU
DDR4 Write Data CRC Check and Retry	DDR4 write data CRC check within the DRAM device and signaling an event back to CPU/IMC for retry. DIMM will signal the CRC miss-match using PAR_ALERT signal. When enabled, two additional bursts are added (a total of 10 bursts) to transfer the write CRC bits. Improves the FRU isolation coverage due to bus transient/ persistent errors. This feature cannot be enabled along with “ADC-SR” or “ADDDC- MR,+1”.	Yes	Yes
Memory Data Scrambling with Command and Address	Scrambles the data with address and command in "write cycle" and unscrambles the data in "read cycle". Addresses reliability by improving signal integrity at the physical layer. Additionally, assists with detection of an address bit error.	Yes	Yes
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 pro-actively searches the system memory, repairing correctable errors. Prevents accumulation of single-bit errors that may result in uncorrected error.	Yes	Yes
DDR4 Memory Multi Rank Sparing	Dynamic fail-over of failing Ranks to spare Ranks behind the same memory controller. UEFI-FW enforces the restriction where up to two ranks out of a maximum of eight ranks can be assigned as spare ranks. This feature cannot be enabled along with “ADC-SR”, “ADDDC-MR,+1”, Memory Mirroring - Intra IMC, and Address Range/ Partial Memory Mirroring.	Yes	Yes
Memory Mirroring - Intra IMC	Memory Mirroring is a 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 behind the same integrated memory controller (IMC). Dynamic (without reboot) fail-over to the mirrored DIMMs is transparent to the SW/OS. This feature cannot be enabled along with “DDR4 Memory Multi-rank Sparing”.	Yes	Yes
Address Range/ Partial Memory Mirroring	In this mode, a subset of memory is mirrored, leaving the rest of the memory in non-mirror mode. Such partial memory mirroring mode allows mirroring of critical memory address range (called as ‘more reliable memory’) without incurring the cost of full memory mirroring. Address ranges need to be within the same IMC. Provides operating system level interface for requesting ‘more reliable memory’ as a percent of full visible memory address space. Optionally, Platform firmware can configure the ‘more reliable memory’ range during boot time using EFI utilities. This feature cannot be enabled along with “DDR4 Memory Multi-rank Sparing”.	No	Yes
DDR4 Post Package Repair (PPR)	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.	Yes	Yes
Memory SMBus Hang Recovery	It allows system recovery in case SMBus fails to respond during runtime thus preventing system crash. UEFI firmware is notified via an SMI and it may be able to recover the SMBus by resetting and re-activating the link. Memory SMBus is actively used by the processor for thermal monitoring and implementing CLTT.	Yes	Yes
Memory Disable/ map-out for FRB	Allows memory initialization and completion of booting flow even when memory fault occurs accomplishing Fault Resilient Booting (FRB).	Yes	Yes
Memory Thermal Throttling	It prevents DIMMs from overheating while balancing power and performance. Management Controller monitors the memory DIMM temperature and can temporarily slow down the memory access rates to reduce the DIMM temperature if needed.	Yes	Yes
MEMHOT Pin Support for Error Reporting	It is primarily used to implement Memory Thermal Throttling solution. Refer to Power Controller Unit (PCU) Functional Description Chapter for further details.	Yes	Yes

Memory population rules and BIOS Setup requirements are defined below:

- Memory sparing and memory mirroring options are enabled in BIOS Setup.
- Memory sparing and memory mirroring options are mutually exclusive in this product. Only one operating mode at a time may be selected in BIOS Setup.
- If a RAS mode has been enabled and the memory configuration is not able to support it during boot, the system will fall back to independent channel mode and log and display errors.
- Rank sparing mode is only possible when all channels that are populated with memory have 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.

5.3.2 Intel® Optane™ Persistent Memory 200 Series RAS Features Overview

Intel® Optane™ PMem 200 series devices support their own Reliability, Availability, and Serviceability (RAS) features. The PMem controller and its embedded firmware work in conjunction with the CPU iMC to provide the following RAS features.

Table 9. Intel® Optane™ PMem 200 Series RAS Features

RAS Features
PMem module Error Detection and Correction
PMem module Device Failure Recovery (SDDC – Single Device Data Correct)
PMem module Package Sparing (DDDC – Double Device Data Correct)
PMem module Patrol Scrubbing
PMem module Address Error Detection
PMem module Data Poisoning (Corrupt Data Containment)
PMem module Viral Mode
PMem module Address Range Scrub (ARS)
PMem module Faulty PMem module Isolation
PMem module Error Reporting

Availability – Ensures consistent uptime / availability with all RAS features

Serviceability (Fault Management)

- Error Reporting
  - Improved error handling to reduce PMem fatal conditions
  - Thermal throttle performance loss – tracks performance loss due to thermal throttling condition
  - eADR – confirms final status from last dirty shutdown
- Faulty DIMM Isolation

See Intel® Optane™ persistent memory 200 series documentation for full RAS feature details.

6. PCI Express (PCIe*) SubSystem Overview

3rd Gen Intel® Xeon® Scalable processors can provide up to 48 PCIe* 3.0 bus lanes to the host system. PCIe bus lanes from each processor are used to support various system features as shown in the following diagram. The system must be configured with four processors to support all possible PCIe functionality of the system.

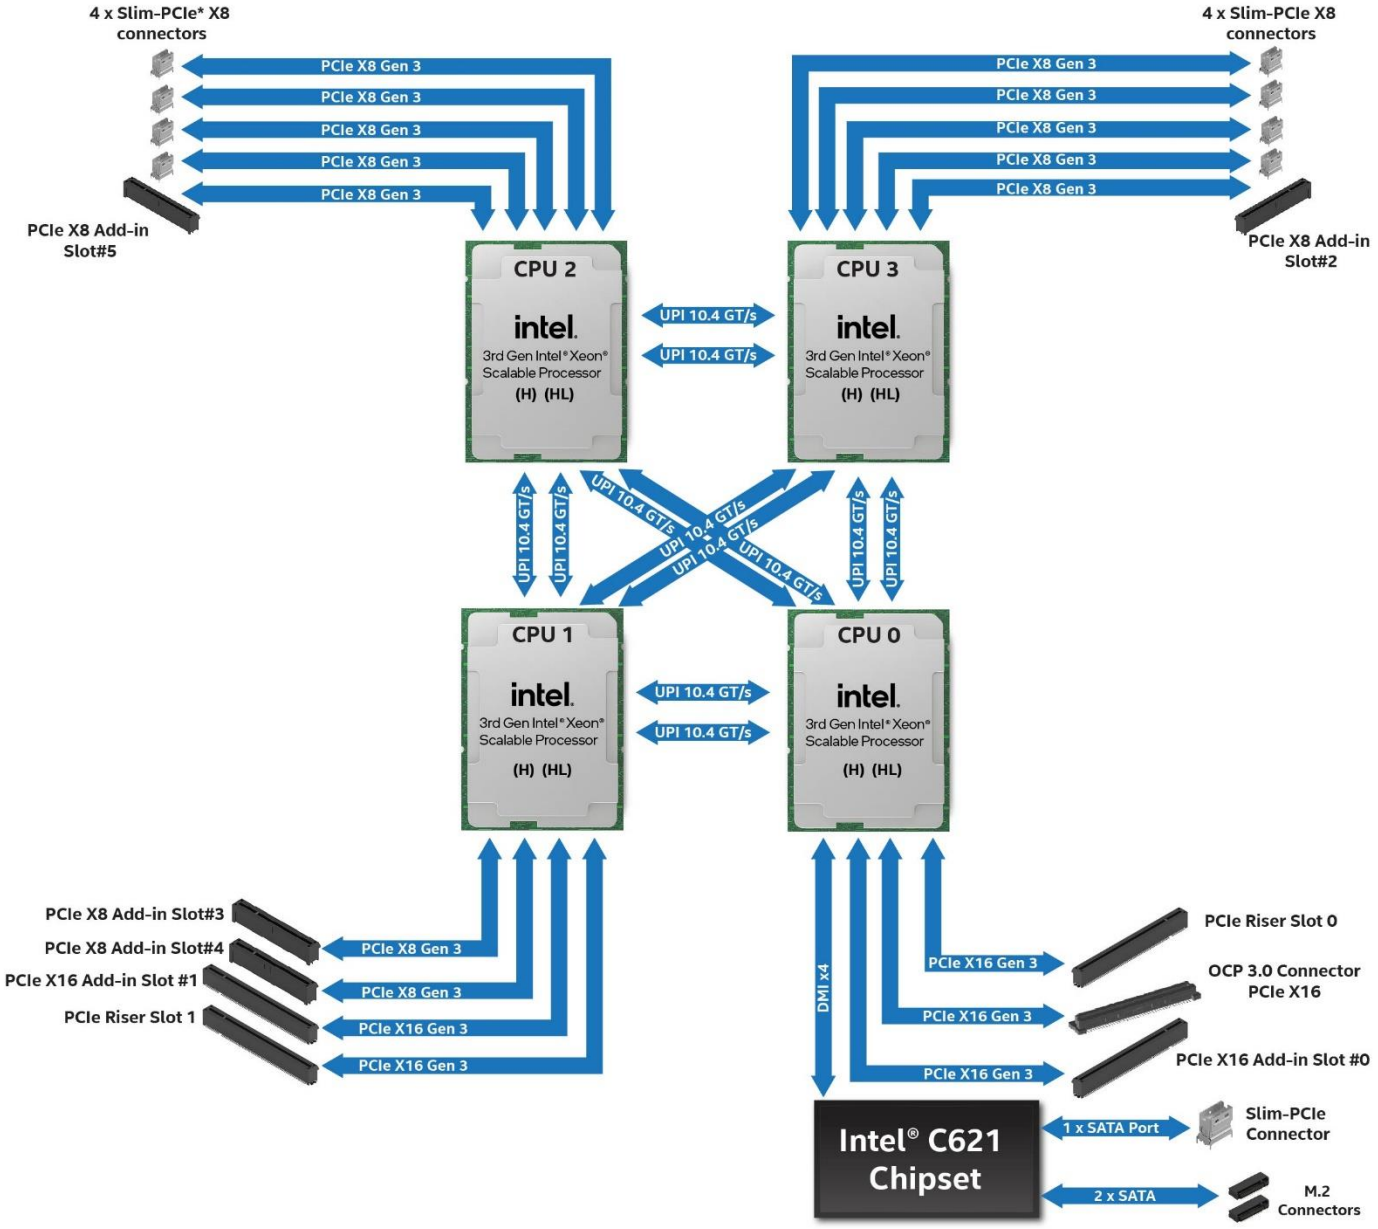

Figure 30. PCIe* Subsystem Architectural Block Diagram

Table 10. PCIe* Port Map

CPU #	Port #	Port Range	# Bus lanes	Onboard PCIe* Connector
CPU0	Port 1	CPU0 PE1 <15:0>	PCIE X16	Riser Slot #0
	Port 2	CPU0 PE2 <15:0>	PCIE X16	OCP* Connector
	Port 3	CPU0 PE3 <15:0>	PCIE X16	PCIe Add-in Slot #0
CPU1	Port 1	CPU1 PE1 <15:8>	PCIE X8	PCIe Add-in Slot #4
	Port 1	CPU1 PE1 <7:0>	PCIE X8	PCIe Add-in Slot #3
	Port 2	CPU1 PE2 <15:0>	PCIE X16	PCIe Add-in Slot #1
	Port 3	CPU1 PE3 <15:0>	PCIE X16	Riser Slot #1
CPU2	Port 1	CPU2 PE1 <15:8>	PCIE X8	Slim-PCIe #1
	Port 1	CPU2 PE1 <7:0>	PCIE X8	Slim-PCIe #0
	Port 2	CPU2 PE2 <15:8>	PCIE X8	PCIe Add-in Slot #5
	Port 3	CPU2 PE3 <15:8>	PCIE X8	Slim-PCIe #3
	Port 3	CPU2 PE3 <7:0>	PCIE X8	Slim-PCIe #2
CPU3	Port 1	CPU3 PE1 <15:8>	PCIE X8	Slim-PCIe #5
	Port 1	CPU3 PE1 <7:0>	PCIE X8	Slim-PCIe #4
	Port 2	CPU3 PE2 <15:8>	PCIE X8	Slim-PCIe #6
	Port 2	CPU3 PE2 <7:0>	PCIE X8	Slim-PCIe #7
	Port 3	CPU3 PE3 <15:8>	PCIE X8	PCIe Add-in Slot #2

6.1 OCP* 3.0 Add-in Card Support

The Intel® Server System M70KLP supports one (1) OCP 3.0 (Small form factor) add-in card. Supported OCP 3.0 cards have a pull-tab to remove the card from the system and a thumbscrew to secure it to the system. Other OCP 3.0 small form factor cards that use a latch or internal lock to secure and remove a card from the system are not supported. See the *Intel® Server Board M70KLP Configuration Guide* for a complete list of supported cards.

Figure 31. OCP* 3.0 Small Form Factor Add-in Cards

Located on the system back panel is the OCP add-in card bay. After removing the filler panel from the chassis, the card is installed from outside of the chassis into the card bay where it is blind mated with a card edge connector on the server board. When fully seated, a captive thumb screw is used to fasten the card to the chassis to ensure the card remains seated when the system is transported.

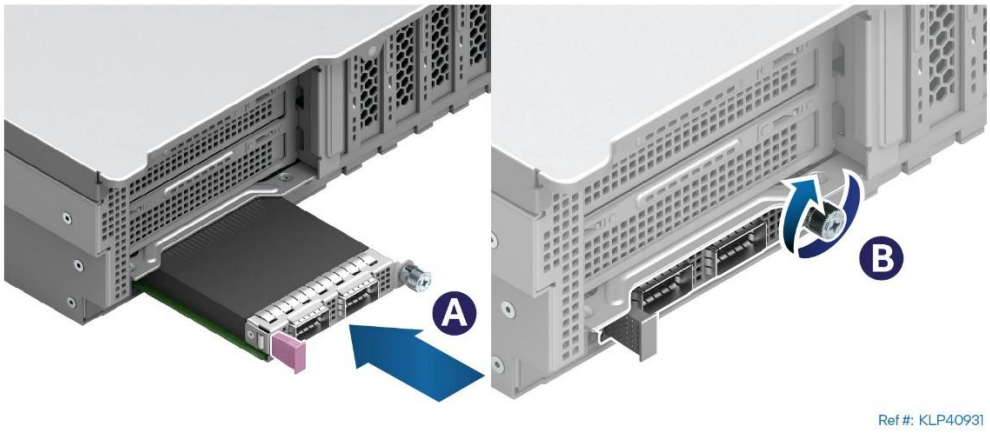

Figure 32. OCP* 3.0 Add-in Card Bay

The OCP card is NOT hot swappable. The system must be powered off and disconnected from AC power to install or remove the card.

Note: The system includes an OCP Hot Plug button located next to the OCP Card bay on the system back-panel. The functionality of this OCP Hot Plug feature has been disabled and is **NOT SUPPORTED**.

6.2 Server Board – PCIe* Add-in Card Support

The server board within the Intel® Server System M70KLP includes six (6) onboard PCIe 3.0 add-in card slots: four (4) PCIe X8 + two (2) PCIe X16. A silkscreen label next to each add-in slot identifies the slot by number (PCIe 0 through PCIe 5) and by the CPU where its PCIe lanes are routed from (CPU0 through CPU3). Each onboard add-in card slot supports a half-height, half-length PCIe add-in card.

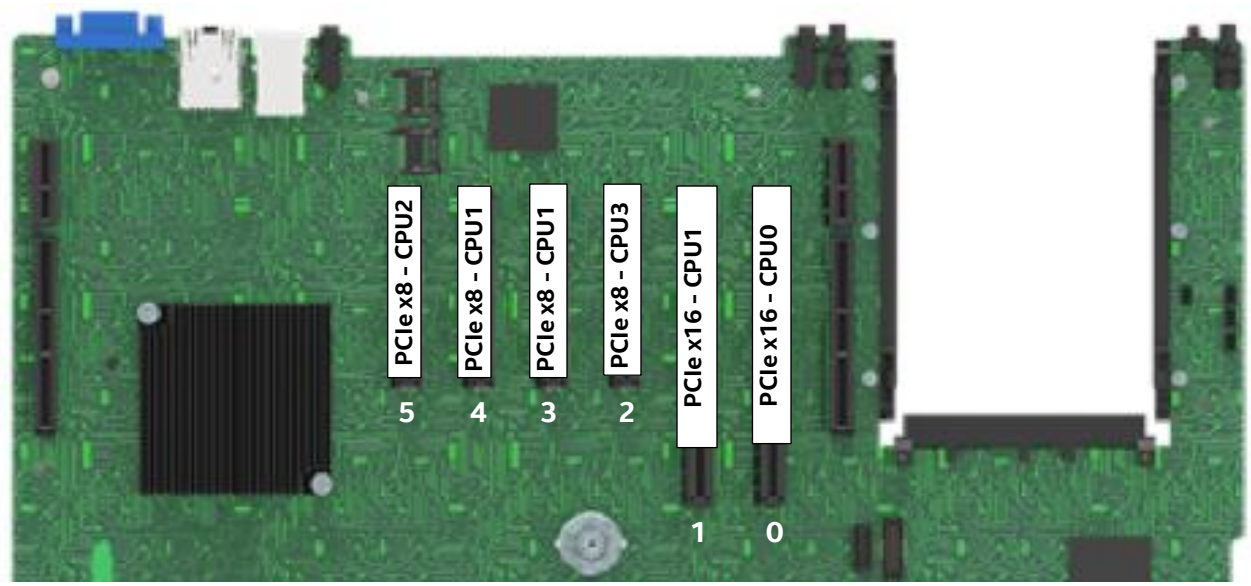

Figure 33. Onboard PCIe* Add-in Slot Identification

6.3 PCIe* Riser Slots

The server board within the Intel® Server System M70KLP includes two PCIe riser card slots identified as Riser slot #0 and Riser slot #1.

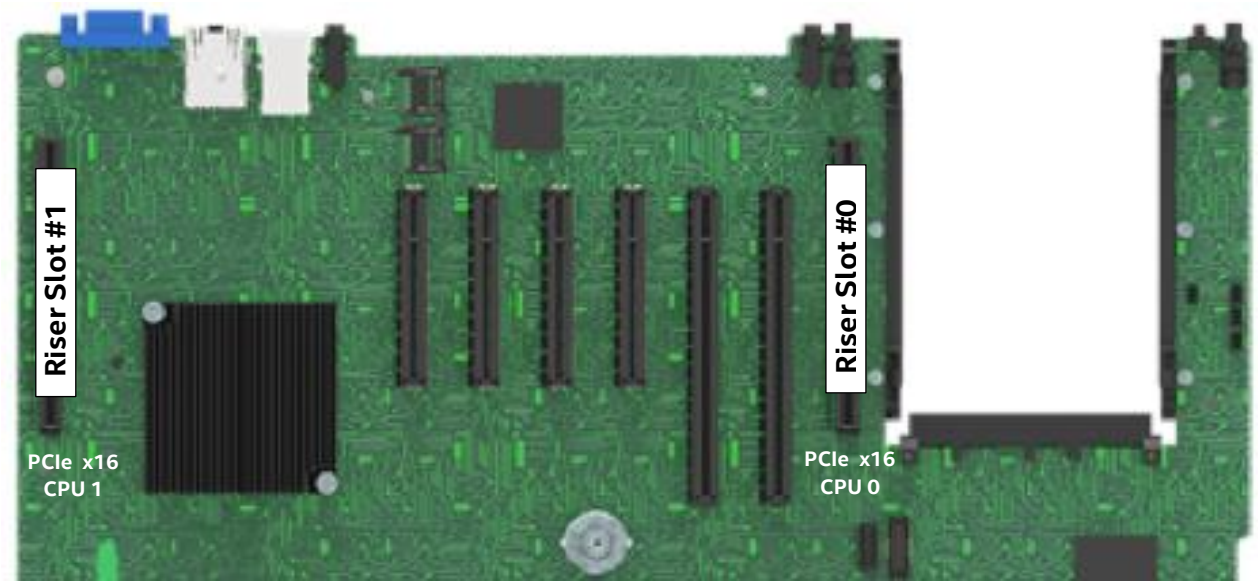

Figure 34. Riser Slot Identification

Each riser slot supports X16 PCIe bus lanes. PCIe bus lanes for Riser Slot #0 are supported by CPU 0, and PCIe bus lanes for Riser Slot #1 are supported by CPU 1.

The riser slots have different use options depending on the desired system configuration.

- To increase the number of PCIe cards that can be used or to use full height / full length PCIe add-in cards, the riser slots can be configured with a riser card option (See Section 6.4).
- To increase the number of PCIe NVMe drives in the front drive bay, each riser slot can be configured with a PCIe bus breakout cable, which routes the X16 PCIe lanes of the riser slot to a backplane mounted to the front drive bay (See Section 7.6).

6.4 Riser Card Card Support

The Intel® Server System M70KLP has the option to support up to two riser card assemblies consisting of a riser card bracket, a riser card, and Auxiliary PCIe signal cables (as necessary).

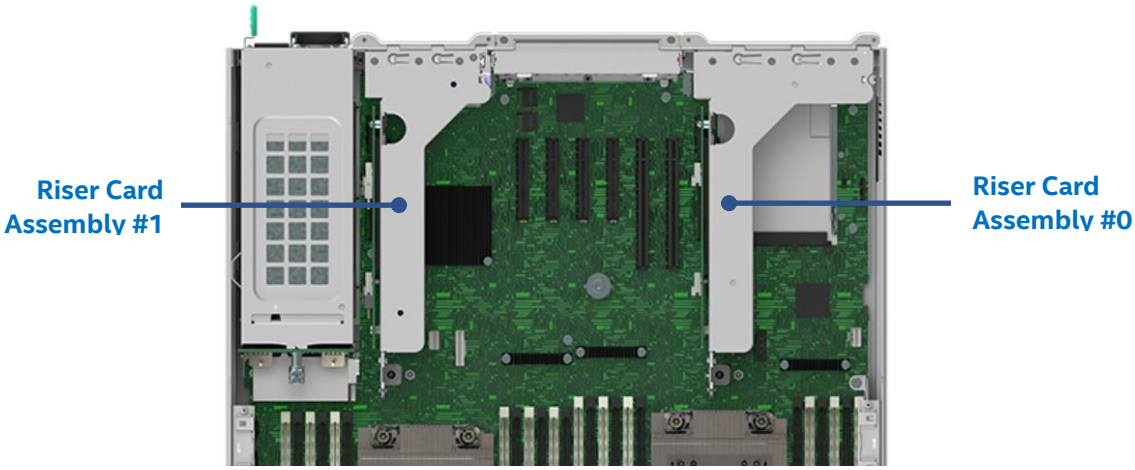

Figure 35. Riser Card Assembly Identification

6.4.1 Riser Slot #0 Assembly Bracket Overview

With the Standard system options (No GPU support), all available add-in cards slots of riser card #0 can accommodate full-height, half-length PCIe add-in cards.

In the GPU enabled system, with its low-profile air duct, riser card #0 supports one (1) GPU add-in card (full-height, full length, double wide) + one (1) full-height, half-length PCIe add-in card. The riser bracket includes an attachment brace to accommodate the extra weight typical of a GPU card. The brace prevents the card from any vertical movement, which can cause possible slot damage while the system is in transit. Different brace options are available to accommodate the mounting requirements of different GPU cards.

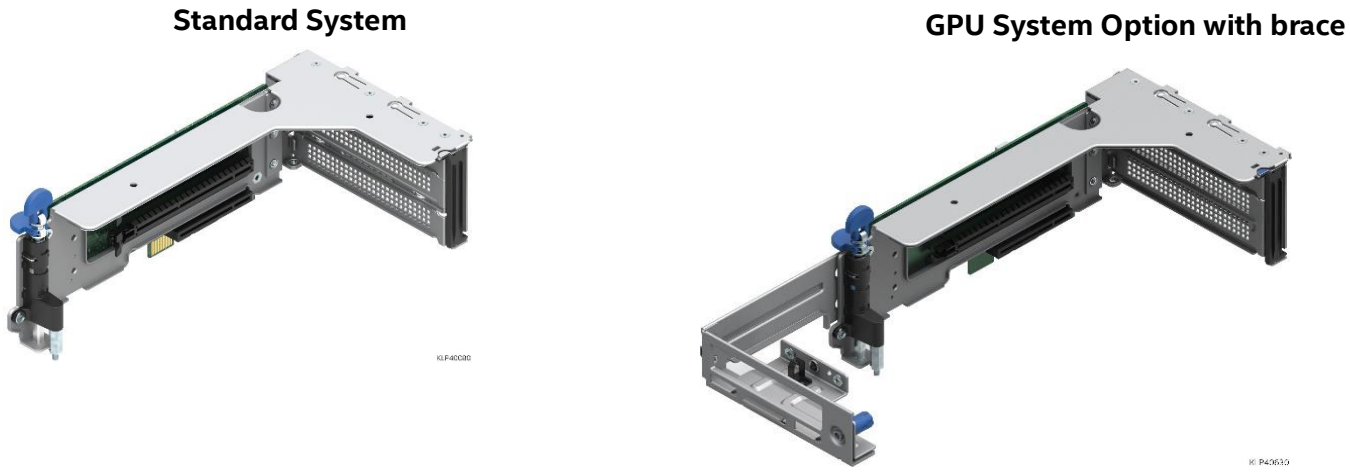

Figure 36. Riser Card #0 Brackets – Standard and GPU Support with Brace

6.4.2 Riser Slot #1 Assembly Bracket Overview

The riser card brackets for Riser Slot #1 and the chassis back panels differ between the Standard system options and the GPU enabled system option.

With the Standard system options (no GPU support), all available add-in cards slots of riser card #1 accommodate half-height, half-length PCIe add-in cards. Using this assembly bracket, all six (6) vertical PCIe add-in slots on the server board are accessible.

Figure 37. Standard System - Riser Slot #1 Low Profile Bracket

In the GPU enabled system, with its low-profile air duct, riser card #1 supports one (1) GPU add-in card (full-height, full length, double wide) + one (1) full-height, half-length PCIe add-in card. The riser bracket includes an attachment brace to accommodate the extra weight typical of a GPU card. The brace prevents the card from any vertical movement, which can cause possible slot damage while the system is in transit. Different brace options are available to accommodate the mounting requirements of different GPU cards.

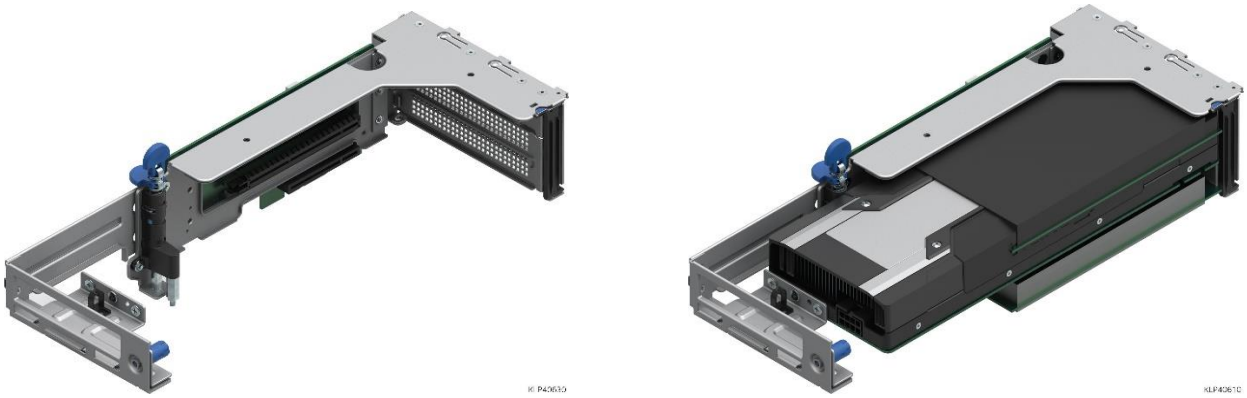

Figure 38. GPU Enabled System - Riser Slot #1 Full Height Bracket (with and without GPU Card)

Using this assembly bracket, only four (4) of six (6) vertical PCIe add-in slots on the server board are accessible.

6.5 Riser Card Options

For system configurations that require additional PCIe add-in card support or need support for full-height add-in cards, two 3-slot riser card options are available.

6.5.1 3-Slot (3 x PCIe X8) Riser Card

The 3-Slot (3 x PCIe X8) riser card is offered as a configurable option. It can be installed into either PCIe riser slot on the server board and their supporting brackets. On the front side of the riser card are three (3) PCIe X8 (mechanical, electrical) add-in card slots. Add-in card slots are identified as 0 through 2, with Slot 0 being closest to the server board.

When the riser card is installed into a riser slot, the X16 PCIe lanes from the riser slot is bifurcated to two X8 bus lanes to support add-in slots 0 and 1. PCIe add-in slot 2 (top slot) is NOT functional unless an Auxiliary PCIe cable is routed to the riser card from one of the X8 Slim-PCIe connectors on the server board.

On the back side of the riser card is a Slim-PCIe connector used as a PCIe interconnect between the riser card and any of the X8 Slim-PCIe cable connectors on the server board. See Section 6.6.

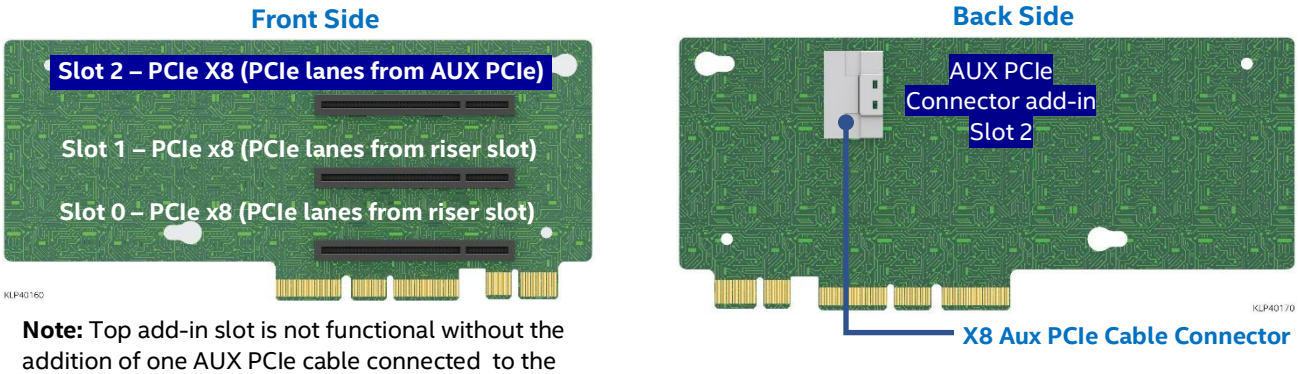

Note: Top add-in slot is not functional without the addition of one AUX PCIe cable connected to the backside of the riser card.

Figure 39. 3-Slot (3 x PCIe X8) Riser Card

6.5.2 3-Slot (2 x PCIe X16 + 1 x PCIe X8) Riser Card

The 3-Slot (2 x PCIe X16 + 1 x PCIe X8) riser card is offered as a configurable option. It can be installed into either PCIe riser slot on the server board and their supporting brackets. On the front side of the riser card are two (2) PCIe X16 (mechanical, electrical) add-in card slots + one (1) PCIe X8 (mechanical, electrical) add-in card slot. Add-in card slots are identified as 0 through 2, with Slot 0 being closest to the server board.

When the riser card is installed into a riser slot, all X16 PCIe lanes from the riser slot are routed to support the middle X16 PCIe add-in Slot 1. PCIe add-in slots 0 (bottom) and 2 (Top) are NOT functional unless Auxiliary PCIe cables are routed to the riser from any of the X8 Slim-PCIe connectors on the server board.

On the back side of the riser card are three (3) Slim-PCIe cable connectors used as PCIe interconnects between the riser card and any of the X8 Slim-PCIe cable connectors on the server board. See Section 6.6.

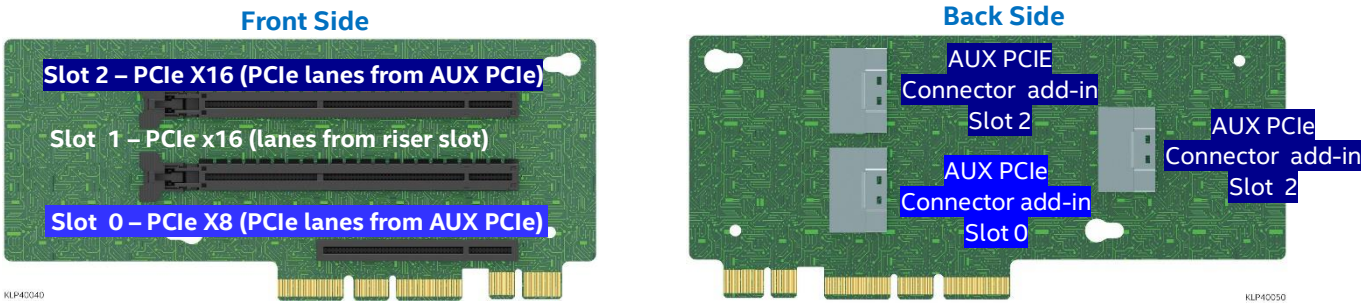

Note: Top and bottom add-in slots are not functional without the addition of three AUX PCIe cables connected to the backside of the riser card.

Figure 40. PCIe x16 Riser Card

To enable PCIe X8 add-in Slot 0 (bottom slot), an Auxiliary PCIe cable must be connected to Slim-PCIe connector 0 on the back side of the riser card and routed to any available X8 Slim-PCIe connector on the server board.

To enable PCIe X16 add-in Slot 2 (top slot), two Auxiliary PCIe cables must be connected to Slim-PCIe connectors 1 and 2 on the back side of the riser card and routed to any two available X8 Slim-PCIe connectors on the server board.

6.6 Server Board X8 Slim-PCIe Connectors

The Intel® M70KLP server board includes eight (8) X8 Slim-PCIe cable connectors. A bank of four connectors (labeled Slim0 – Slim3) on the left edge of the server board is supported by CPU 2. A bank of four connectors (labeled Slim4 – Slim7) on the right edge of the server is supported by CPU 3. See Figure 41.

Each Slim-PCIe connector has X8 PCIe 3.0 bus lanes routed to it from their designated processor.

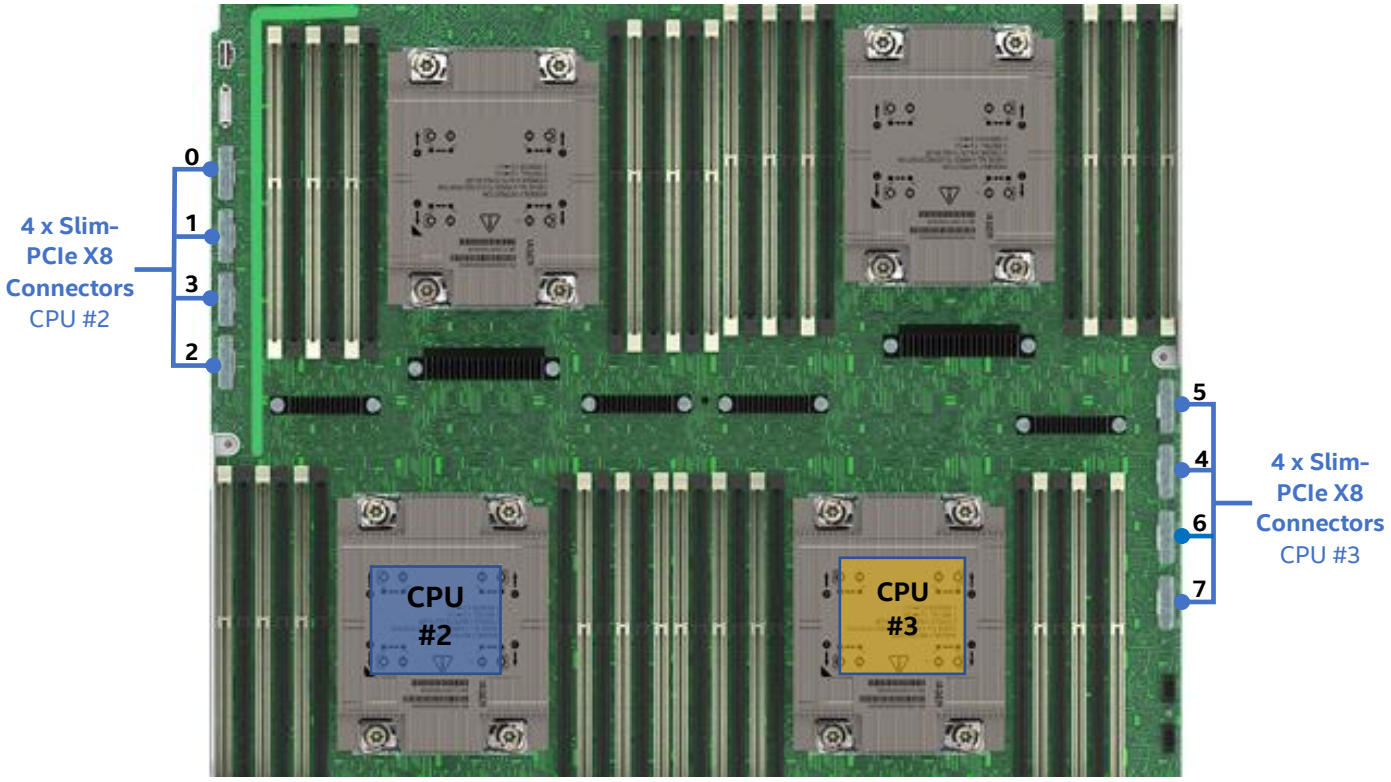

Figure 41. On-board Slim-line PCIe Connectors

When cabled, each X8 Slim-PCIe connector provides the flexibility of either increasing the number of PCIe add-in cards being supported by a riser card or adds support for PCIe NVMe SSDs to the front drive bay.

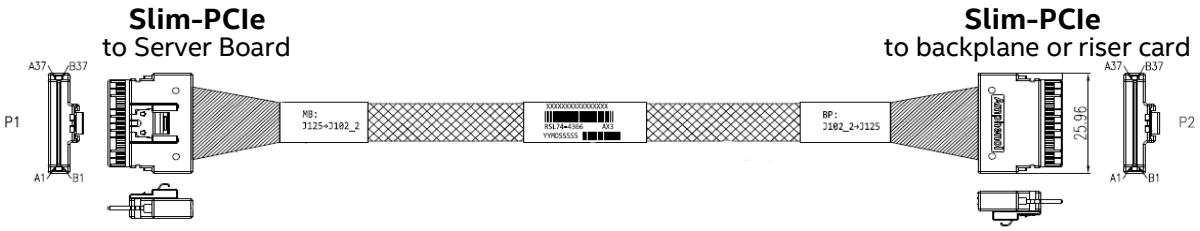

Figure 42. Auxiliary PCIe* Cable

7. Data Storage Options

The Intel® Server System M70KLP includes support for different data storage options. This chapter provides an overview of each.

7.1 Internal M.2 SATA SSD support

Mounted atop of the system air duct is an M.2 SSD interface board that is cabled to a SATA interface connector on the server board. The interface board provides the option to support up to two (2) 80 mm or 110 mm SATA M.2 SSDs.

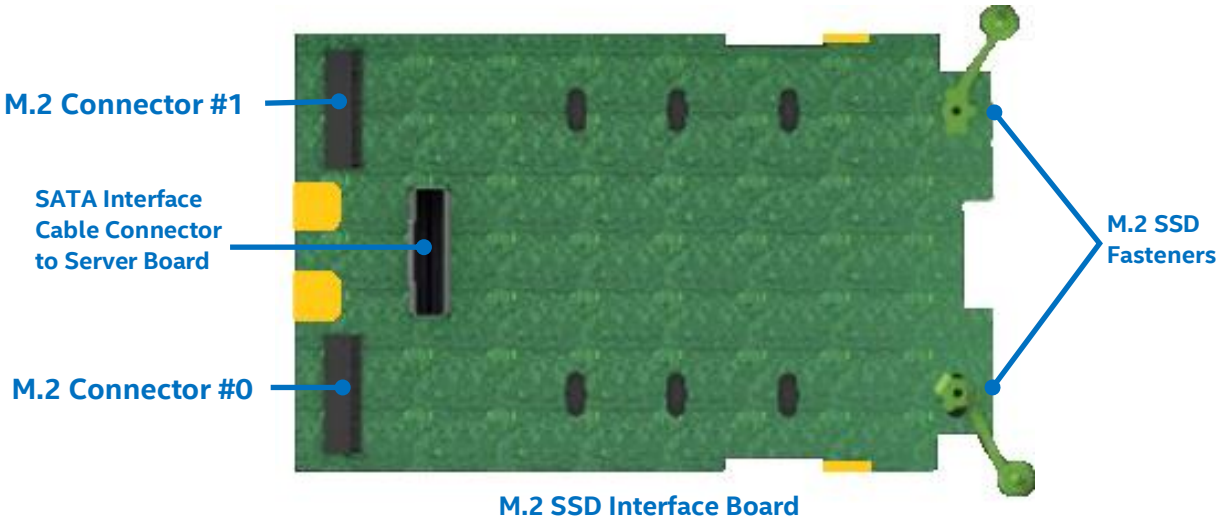

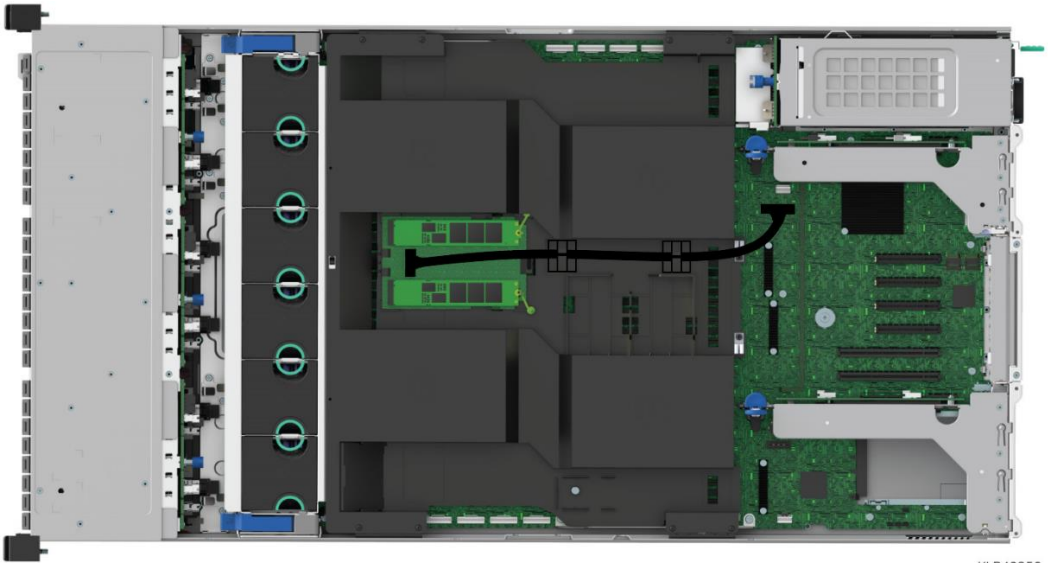

Figure 43. M.2 SATA SSD Support

SATA signals for the M.2 SSDs are supported by the Intel® C621 Chipset.

7.2 Internal USB 3.0

Located on the server board, directly behind CPU 1, is a Type A USB 3.0 connector that can be used to support a USB flash drive.

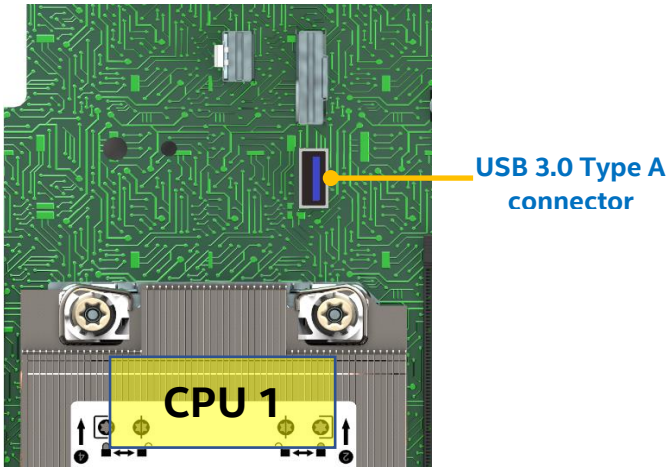

Figure 44. Internal USB 3.0 Type A Connector

7.3 Front Drive Bays

The system has support options for 8, 16, or 24 front drive bays. Each drive bay includes a drive carrier that must be populated with a 2.5” drive (SSD or HDD) or supplied drive blank. Each set of eight drive bays is supported by a common backplane. All drives attached to a common backplane must match media type (SSD or HDD). Mixing drive storage media types within a common backplane is not supported. In addition, all drives attached to a common backplane must share a common interface type (SATA/SAS or NVMe).

Figure 45. 8, 16, and 24 Drive Bay Options

Each installed backplane is mounted internally to the back side of the drive bay using two key-hole studs and one captive thumbscrew allowing for tool-less removal and installation.

On the front side of the backplane are eight drive interface connectors. Each connector supports SATA/SAS or PCIe NVMe drives. Data signals to each drive connector are routed from designated cable connectors on the back side of the backplane.

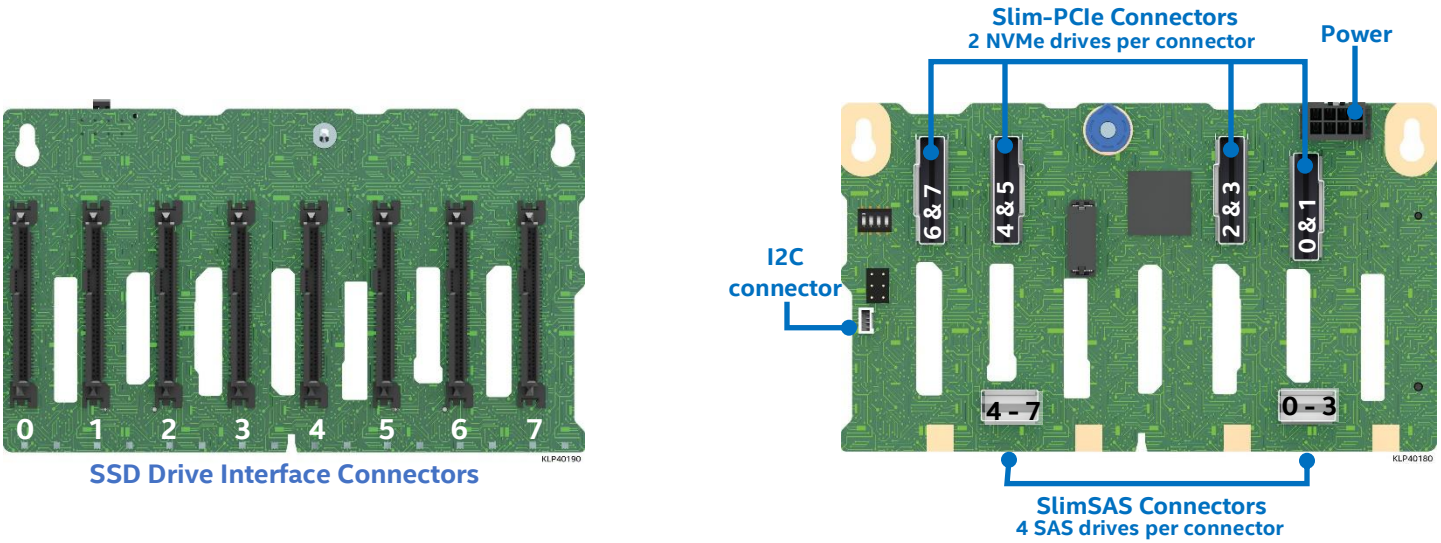

Figure 46. Hot Swap Backplane – Front and Back Views

On the back side of the backplane are the storage interface connectors:

- Two 4-port SlimSAS* cable connectors to support a SATA/SAS interface
- Four 2-port Slim-PCIe connectors to support a PCIe interface.

For each backplane installed, only one interface type (SATA/SAS or PCIe) can be in use. Concurrent use of the SAS/SATA and NVMe interface connectors on a common backplane is not supported.

Note: System configurations with two or more backplanes installed can support different storage interfaces between backplanes.

Each Slim-PCIe connector provides X8 PCIe bus lanes to support two X4 PCIe NVMe drive connectors on the front side of the backplane. Each SlimSAS connector supports four SATA/SAS drive connectors on the front side of the backplane

All 2.5” drives are mounted to a drive carrier using four (4) screws. Drive removal from the chassis and installation into the chassis is tool-less. However, installing a drive into the drive carrier requires a Phillips* head screwdriver.

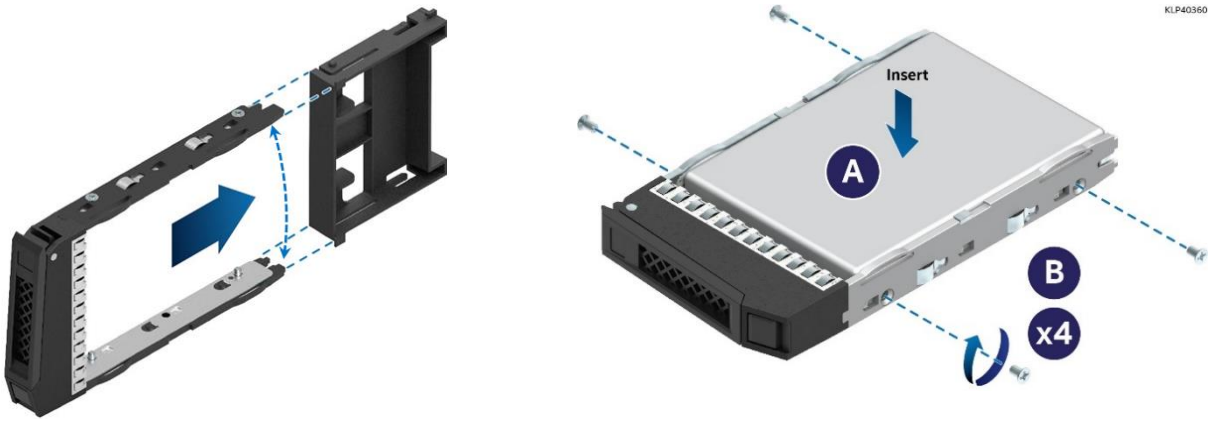

Figure 47. Drive Carrier

Note: All drive bays must be populated and all drive carriers installed within the front drive bay must be assembled with a drive or supplied drive blank while a system is operational.

Pressing the release button on the lower face of the drive carrier releases a handle that is used to pull the drive from the drive bay.

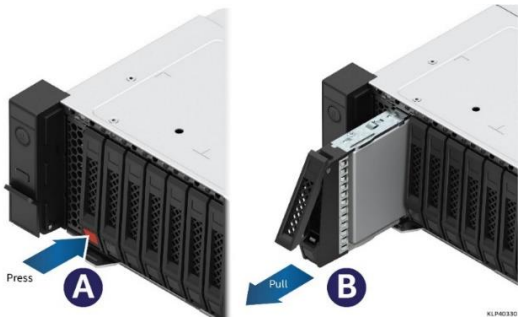

Figure 48. Drive Carrier Removal

The handle must be in the open position when installing the drive into the drive bay. When the drive is fully inserted into the drive bay, the handle is then pushed closed. This procedure guides the drive to its fully installed position and securely latches it to the drive bay.

Figure 49. Drive Carrier Installation

On the backplane next to each of the eight drive interface connectors are Drive Fault and Drive Activity LEDs. Light pipes integrated within each drive carrier are used to project the light from these LEDs making them visible on the front face plate of the drive carrier as shown in the following illustration.

Label	Description	Description
1	Activity Status LED	Steady green: Normal Flashing green: Read and write activity
2	Fault Alarm LED	Steady red: A failure has occurred Steady blue: Drive positioning Steady pink: RAID rebuilding

Figure 50. Drive LEDs

7.4 Cable Routing

Proper cable routing and management within the system is necessary to correctly install the system air duct. In addition, managing how cables are routed within the system is critically important for proper airflow and maintaining system thermals. Internal cables routed improperly can block airflow to critical components and areas within the chassis. This situation can cause the components to heat up beyond their thermal limits, which can impact system performance and longevity.

When routing cables to or from the area behind the front drive bay, they must be routed using cable channels along the chassis sidewalls as shown with yellow arrows below. No cables should be routed through the center of the system.

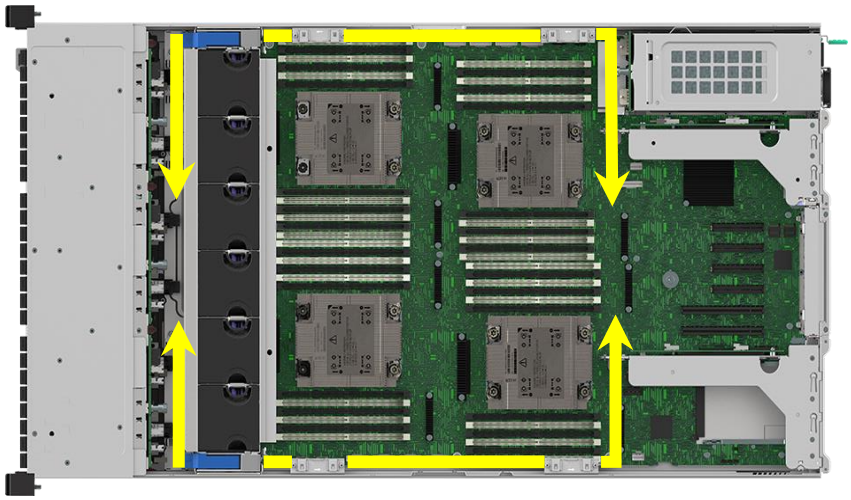

Figure 51. Cable Routing

Several features designed within the chassis help to manage cables and create an aesthetically clean looking system.

The system fan housing is modular and tool-less, allowing for easy removal and installation. The fan housing must be removed to route cables to or from the area behind the front drive bay.

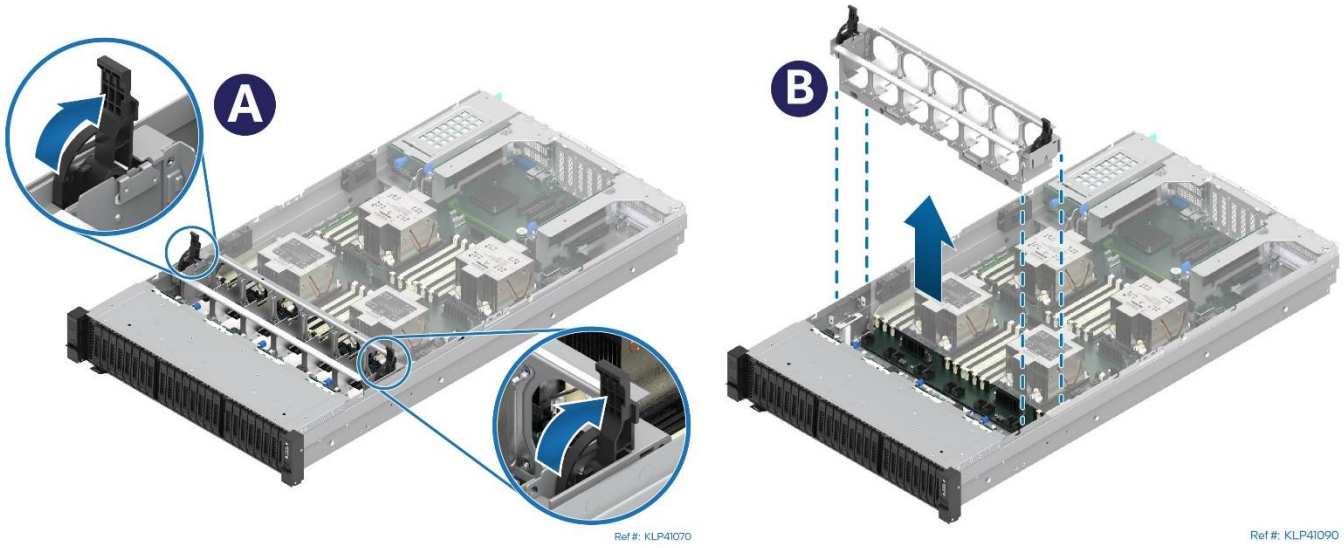

Figure 52. Removable System Fan Bracket

Mounted to each chassis sidewall are two cable management brackets as shown below.

Open bracket to install cables

Close bracket to secure cables

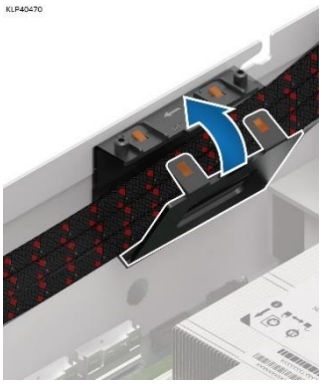

Cables routed and secured

Figure 53. Internal Cable Management

7.5 SAS / SAS RAID Support

All backplanes and drive bays configured within an Intel® Server System M70KLP have support for 2.5" SAS or SATA drives. However, the system has no embedded SATA, SAS, or SAS RAID controller for use on the front drive bay. To support SATA or SAS drives, one or more PCIe add-in SAS or SAS RAID controller(s) must be installed to any open PCIe add-in slot on the server board or riser card.

Up to two 4-port mini-SAS HD connectors from an add-in Host Bus Adapter (HBA) or RAID card can be cabled to a single backplane using cables with connectors that match those of the add-in card (mini SAS HD) and backplane (Slim-SAS).

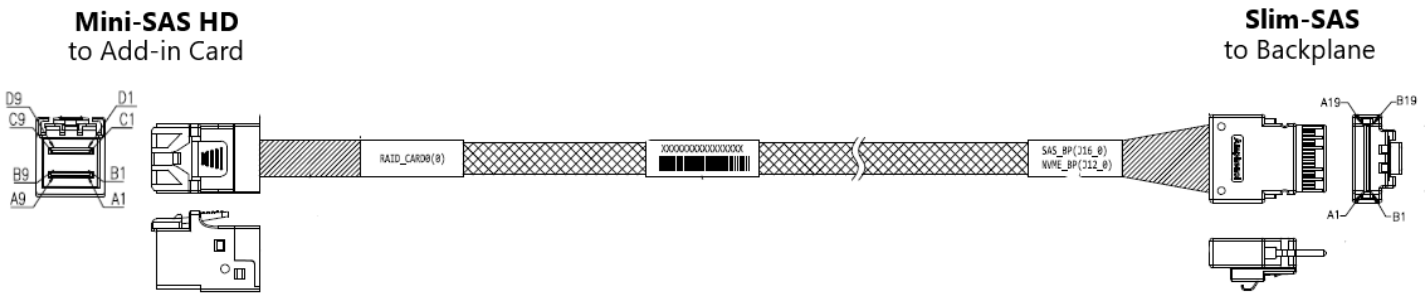

Figure 54. Mini-SAS HD to Slim-SAS Cable

7.5.1 Intel® SAS RAID RMFBU Support

Many Intel SAS RAID cards have support for an optional RAID Maintenance Free Backup Unit (RMFBU). During a power loss event, this kit prevents possible data loss by providing the RAID card with enough backup power to save any unwritten data.

The standard system can support up to three RFMBUs mounted to the top side of the system air duct. The mounting location closest to the back edge of the air duct is designed to support two stacked RMFBU modules. In the GPU enabled system, one RMFBU can be installed to the air duct.

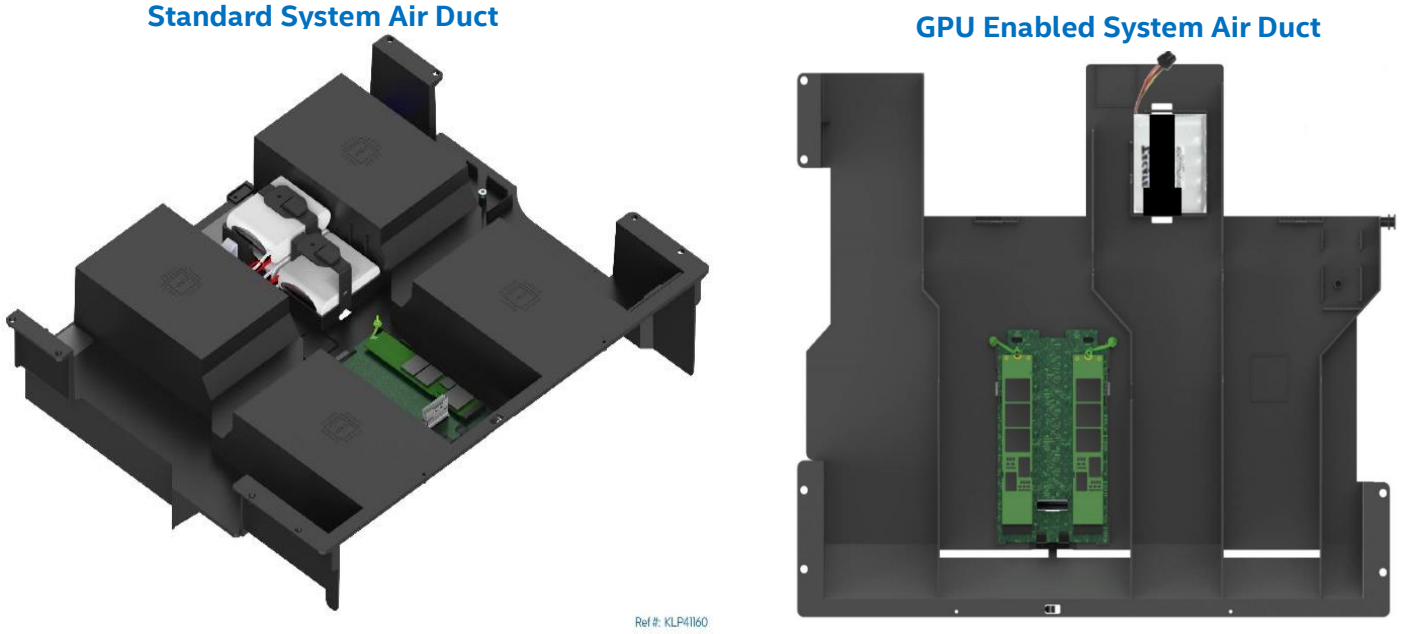

Figure 55. Intel® RAID RMFBU Accessory Support

All RMFBU modules are secured to the air duct using a hook and latch strap.

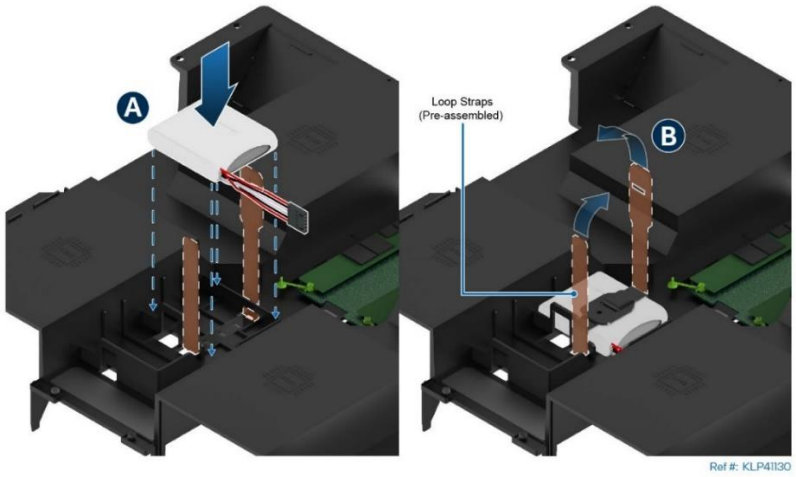

Figure 56. Intel® RMFBU Installation

Note: The mounting bracket included with the Intel RMFBU accessory kit is not used in this server system.

7.6 PCIe* NVMe* Support

All backplanes and drive bays configured within an Intel® Server System M70KLP have support for 2.5" U.2 PCIe NVMe SSDs. Using an Auxiliary PCIe cable, X8 PCIe bus lanes from any of eight Slim-PCIe connectors on the server board are cabled to matching Slim-PCIe connectors on the backplane. Each Slim-PCIe connector on the backplane supports two (2) PCIe NVMe drives. See Figure 46.

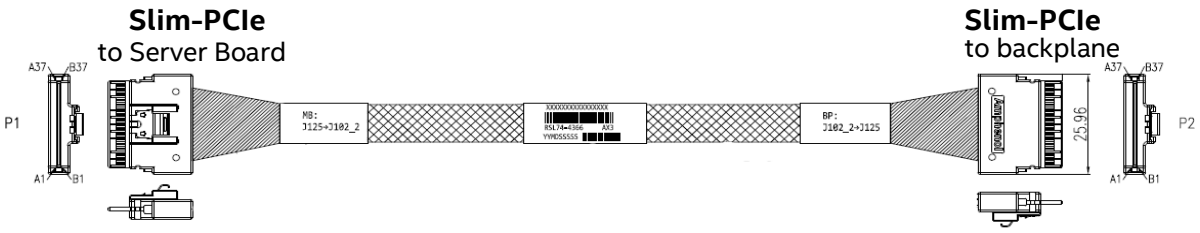

Figure 57. Auxiliary PCIe* Cable

Up to sixteen NVMe drives are supported when using all eight Slim-PCI connectors on the server board.

To support greater than sixteen NVMe drives, PCIe bus lanes from one or both PCIe riser slots can be cabled to a given backplane using a PCIe bus break-out cable that plugs directly into the riser slot.

Note: Riser slots on the server board can be populated with a Riser Card or a PCIe Breakout cable. PCIe bus lanes from the riser slot cannot be routed to a backplane when a riser card is installed into it.

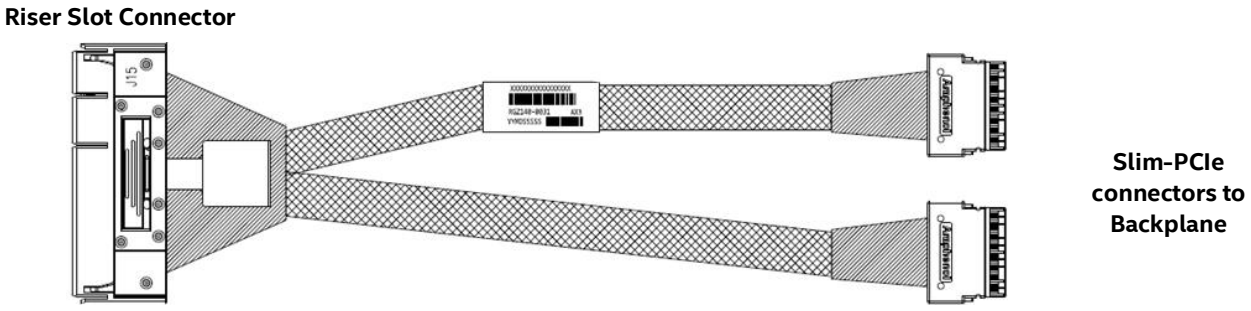

Figure 58. PCIe* Break-out Cable - Riser Slot to Backplane

A PCIe bus break-out cable routes all X16 PCIe bus lanes of the riser slot to two Slim-PCIe connectors on a backplane, supporting up to four NVMe drives.

Using a PCIe break-out cable in both riser slots adds support for up to an additional eight NVMe drives.

7.6.1 Intel® Volume Management Device (Intel® VMD)

Intel® Xeon® Scalable processors support an Intel® Volume Management Device (Intel® VMD), which is an integrated controller inside the CPU PCIe root complex. With Intel® VMD, NVMe SSDs are directly connected to the CPU, allowing the full performance potential of fast storage devices to be realized.

Intel® VMD 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.

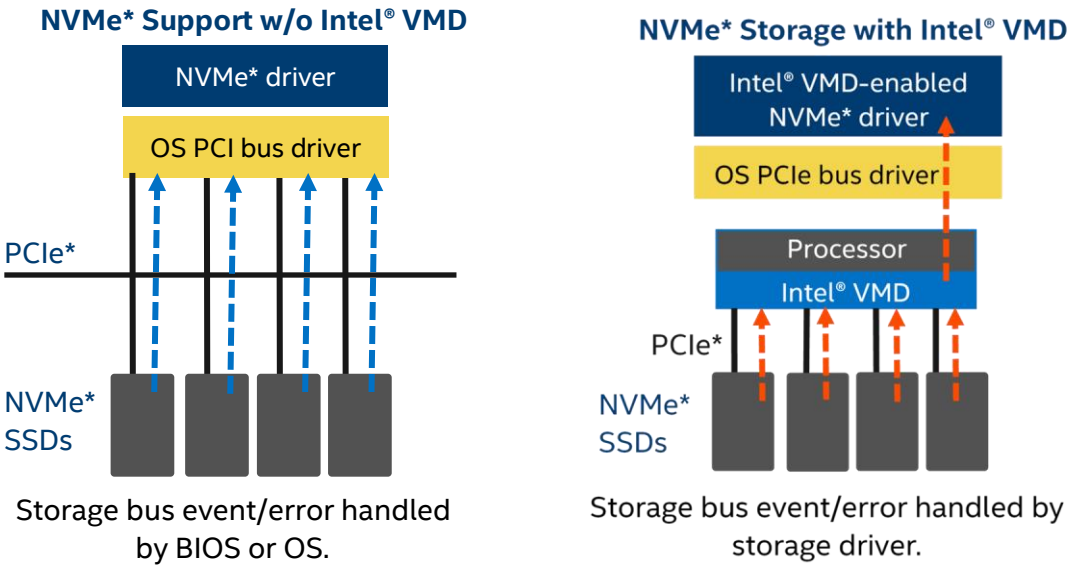

Figure 59. NVMe* Storage Bus Event / Error Handling

Intel® VMD 1.0 supports the following 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* NVMe* SSDs.
- Support for PCIe* NVMe* SSDs only (no network interface controllers (NICs), graphics cards, and so on)
- Maximum of 128 PCIe* bus numbers per domain.
- Support for Management Component Transport Protocol (MCTP) over SMBus* only.
- Support for MMIO only (no port mapped I/O).
- Does not support NTB, Quick Data Tech, Intel® Omni-Path Architecture (Intel® OPA), or SR-IOV.
- Correctable errors do not bring down the system.

Each installed NVMe SSD communicates with the VMD controller of a given processor through X4 PCIe bus lanes routed to the onboard PCIe interface to which the drive is connected. Each set of X4 PCIe bus lanes are associated with a PCIe root port within the processor's PCIe controller.

By default, Intel® VMD support is disabled on all processor PCIe* root ports. To enable VMD controller support for NVMe devices, Intel® VMD must be enabled on the appropriate processor PCIe* root ports in BIOS Setup.

7.6.1.1 PCIe* Root Port Setup for NVMe* Support

The following table provides the PCIe* root port mapping information for each processor. Line items in bold black identify all the connectors on the server board that can be configured for NVMe support.

Table 11. PCIe* Port Map for Intel® VMD Setup

CPU #	Port #	Port Range	# Bus lanes	Onboard PCIe Connector
CPU0	**Port 1**	**CPU0 PE1 <15:0>**	**PCIE X16**	**Riser Slot #0**
	Port 2	CPU0 PE2 <15:0>	PCIE X16	OCP Connector
	Port 3	CPU0 PE3 <15:0>	PCIE X16	PCIe Add-in Slot #0
CPU1	Port 1	CPU1 PE1 <15:8>	PCIE X8	PCIe Add-in Slot #4
	Port 1	CPU1 PE1 <7:0>	PCIE X8	PCIe Add-in Slot #3
	Port 2	CPU1 PE2 <15:0>	PCIE X16	PCIe Add-in Slot #1
	Port 3	**CPU1 PE3 <15:0>**	**PCIE X16**	**Riser Slot #1**
CPU2	**Port 1**	**CPU2 PE1 <15:8>**	**PCIE X8**	**Slim-PCIe #1**
	Port 1	**CPU2 PE1 <7:0>**	**PCIE X8**	**Slim-PCIe #0**
	Port 2	CPU2 PE2 <15:8>	PCIE X8	PCIe Add-in Slot #5
	Port 3	**CPU2 PE3 <15:8>**	**PCIE X8**	**Slim-PCIe #3**
	Port 3	**CPU2 PE3 <7:0>**	**PCIE X8**	**Slim-PCIe #2**
CPU3	**Port 1**	**CPU3 PE1 <15:8>**	**PCIE X8**	**Slim-PCIe #5**
	Port 1	**CPU3 PE1 <7:0>**	**PCIE X8**	**Slim-PCIe #4**
	Port 2	**CPU3 PE2 <15:8>**	**PCIE X8**	**Slim-PCIe #6**
	Port 2	**CPU3 PE2 <7:0>**	**PCIE X8**	**Slim-PCIe #7**
	Port 3	CPU3 PE3 <15:8>	PCIE X8	PCIe Add-in Slot #2

To determine which PCIe root ports to enable for NVMe support requires knowledge of the PCIe cable routing between the backplane used for NVMe SSDs and the onboard PCIe connectors they are cabled to.

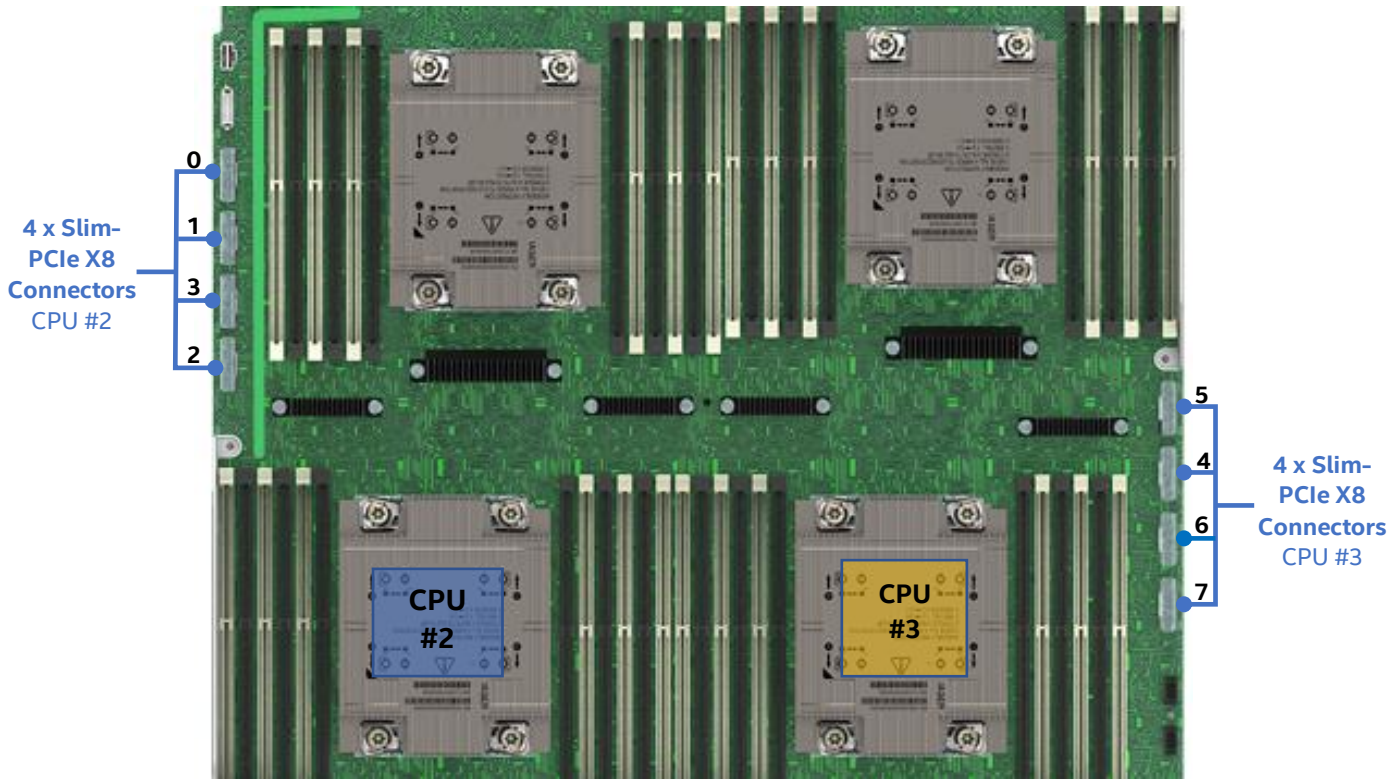

Figure 60. Onboard Slim-PCIe Connectors

With NVMe cable routing determined, then the <F2> BIOS Setup utility can be used to enable the appropriate PCIe root ports.

In BIOS Setup, the Intel® VMD support menu can be found using the following menu map:

Advanced > **PCI Configuration** > **Volume Management Device.**

7.6.1.2 Enabling NVMe* RAID Support

Intel® VMD handles the physical management of NVMe* storage devices as a stand-alone function but can be enhanced when optional Intel® Virtual RAID on CPU (Intel® VROC) support features are enabled to implement RAID-based storage systems.

Intel® VROC supports the following:

- 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 operating system for error handling.
- Protected R5 data from operating system crash.
- RAID/storage management using Representational State Transfer (RESTful) application programming interfaces (APIs).
- Graphical user interface (GUI) for Linux*.
- 4K native NVMe* SSD support.

Enabling the RAID features of Intel® VROC requires the installation of an optional upgrade key on to the server board as shown in the following figure.

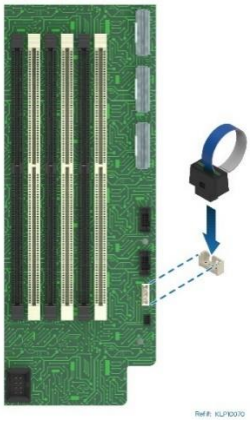

Figure 61. Intel® VROC 7.5 Key Insertion

Available Intel® VROC upgrade key options include the following:

Table 12. Optional VROC Upgrade Key - Supported NVMe* RAID Features

NVMe* RAID Major Features	Standard Intel® VROC 7.5 Key (iPC – VROCSTANMOD)	Premium Intel® VROC 7.5 Key (iPC – VROCPREMMOD)
Processor-attached NVMe* SSD – high performance	Yes	Yes
Boot on RAID volume	Yes	Yes
Third party vendor SSD support	Yes	Yes
RAID 0/1/10	Yes	Yes
RAID 0/1/5/10	No	Yes
RAID write hole closed (RMFBU replacement)	No	Yes
Hot plug/ surprise removal (2.5” SSD form factor only)	Yes	Yes
Enclosure LED management	Yes	Yes

8. Front Control and I/O Panel Features Overview

Integrated within the front system handles are a control panel (Left) and I/O panel (Right). The left control panel includes Power and System ID/BMC Reset buttons and System Status LEDs. The right I/O Panel includes interface connectors for USB and Video.

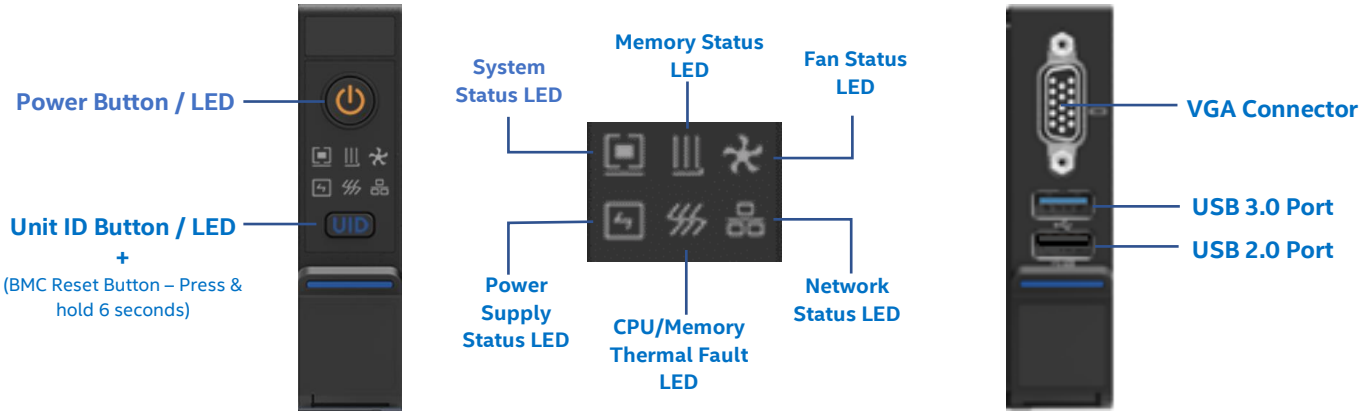

Figure 62. Front Control and I/O Panels

Table 13. Front Panel Button and LED Operation

Button / LED	Description
Power button with LED	**The power LED is green when power is on.**
	The power LED is orange when the system is in standby.
	Press the button for 4 seconds to force a system shutdown.
Unit ID Button with LED / BMC RST button	**The button is used to turn on/off the Blue UID LED.**
	Press and hold the button for 6 seconds to force the BMC to restart.
System fault indicator	LED is OFF when system is in a normal state
	LED is steady RED if the system is in a fault condition.
	LED will Blink RED when an alarm is reported.
Memory fault indicator	LED is OFF when system is in a normal state
	LED is steady RED if memory is in a fault condition.
	LED will Blink RED when an alarm is reported.
Fan fault indicator	LED is OFF when system is in a normal state
	LED is steady RED with a fan RPM read failure.
	LED Blinks RED in the case of RPM read exception.
Power supply fault indicator	LED is OFF when system is in a normal state
	LED Is steady RED when a power supply fails.
	LED Blinks RED when the power supply is abnormal.
CPU / Memory thermal fault indicator	LED is OFF when system is in a normal state
	LED is steady RED when the CPU or memory are overheated.
Network status LED	**LED is steady green or blinks green when the network connection is normal.**
	LED Is off in the case of no network connection.
	*Note: It only indicates the working status of the OCP card (4S).

9. Thermal Management Overview

Managing airflow within the system is critical to having it operate with best performance and for maintaining long term reliability. The system pulls cool air (10–35°C) in from the front, channels it over and through several high heat generating components and areas within the chassis, and then pushes the hot air out the back.

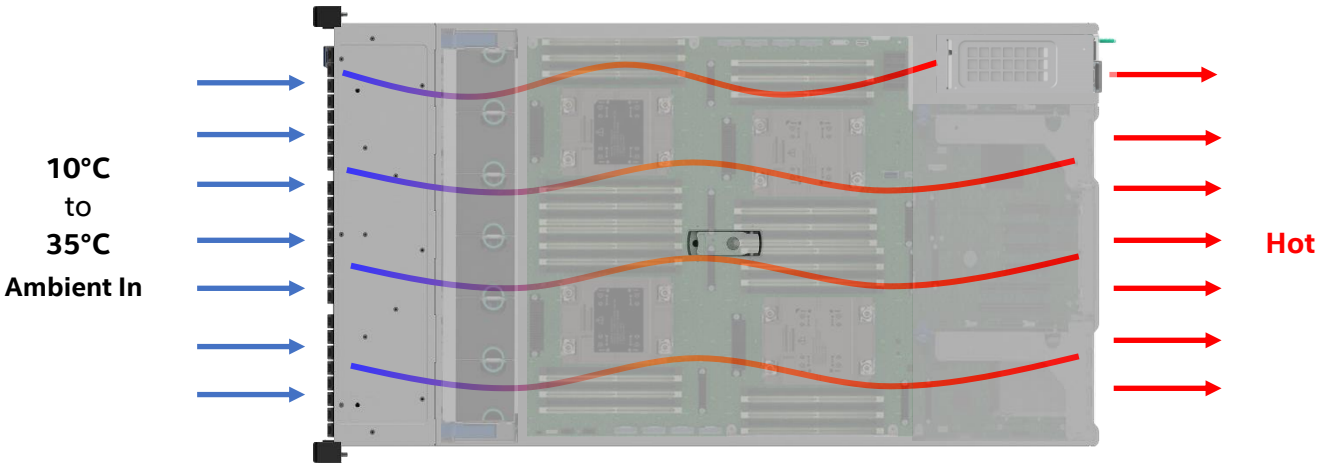

Figure 63. System Airflow

System fans, power supply fans, an air duct, population rules for various system components, thermal sensors, and embedded platform management are used to monitor temperatures and manage airflow.

9.1 System Fans

Six managed system fans are used to drive air through the system.

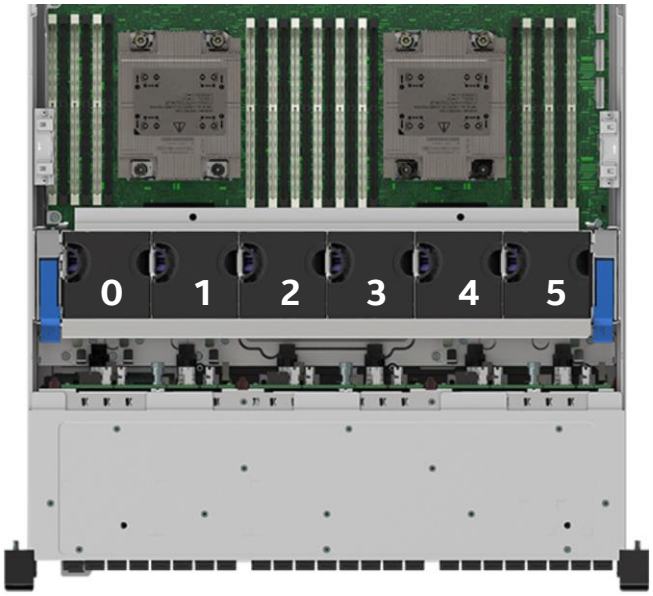

Figure 64. System Fan Identification

System fans have no cables and replacement requires no tools.

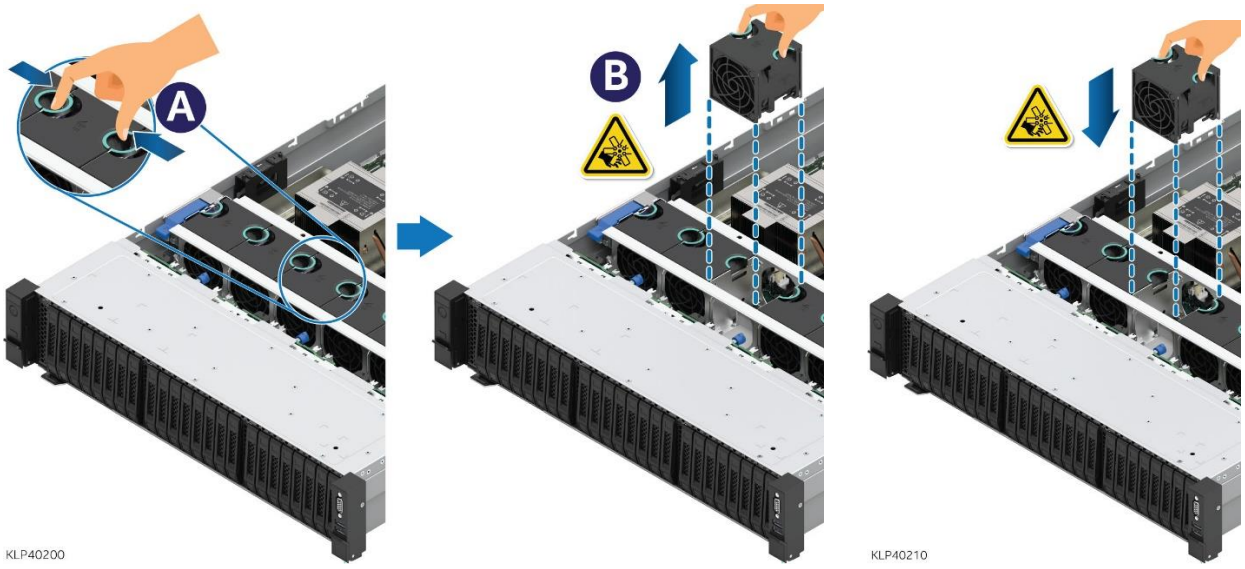

Figure 65. System Fan Replacement

Each system fan is managed by the baseboard management controller (BMC), which is responsible for monitoring temperature sensors of various components and areas within the system. As ambient and internal temperatures increase and decrease, the BMC will vary the speed of each system fan using pulse width modulation to increase or decrease the airflow as necessary. If the BMC detects that a non-critical thermal threshold is exceeded with system fans already operating at 100%, platform management will impose operational limits (also known as throttling) to either the memory or processors or both to reduce the operating temperature within the system. System performance will be impacted should throttling occur. Throttling will continue until the non-critical thermal threshold is no longer exceeded. Should the system exceed a critical thermal threshold, platform management will initiate a system shutdown.

Note: Non-critical and critical thresholds for various monitored system features are pre-determined and pre-programmed within the BMC firmware and are not configurable.

The system supports fan redundancy that allows the system to maintain necessary airflow for best system performance when a system fan fails. Fan redundancy is lost with a single fan failure. A failed system fan should be replaced as soon as possible.

Each system fan is hot swappable, meaning that a failed fan can be replaced without first powering down the system. Should a system fan fail, the BMC will increase the speed of the remaining fans to maintain the airflow necessary to keep the system operating at best performance. The remaining system fans will continue to operate at the higher speeds until the failed fan is replaced.

Note: The system top cover must be removed from the system when replacing a system fan. A failed system fan can be replaced with the system still operational. However, when hot swapping a system fan with the top cover removed, the procedure must be performed as quickly as possible to ensure the system remains operating below its critical thermal limit. See the *Intel® Server System M70KLP Service Guide* for complete fan replacement instructions.

9.2 Power Supply Fan

Each installed power supply module includes a fan that pulls air from within the system, through the power supply, and out the back. Thermal monitoring will vary the fan speed. Should a power supply fan fail, the resulting increase in temperature within the power supply will initiate its thermal protection circuit to shut itself down. Power supply fans cannot be replaced. The power supply module must be replaced if its fan stops working.

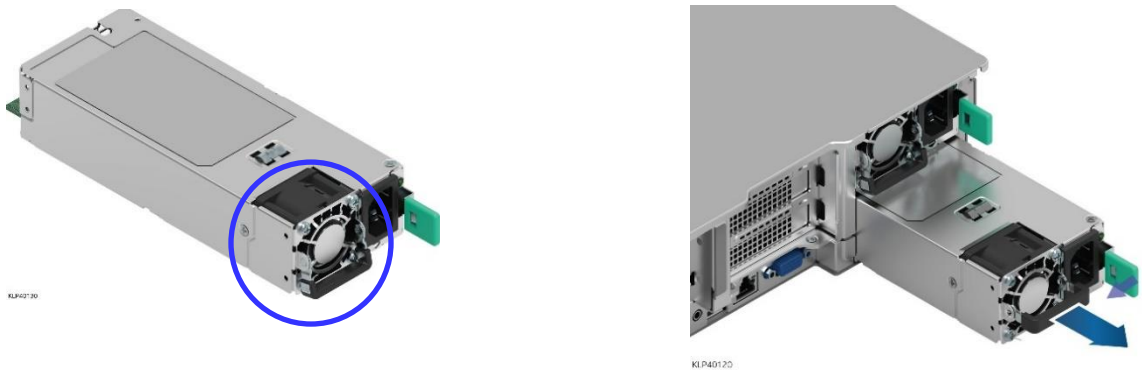

Figure 66. Power Supply Fan

Note: A power supply module can be hot swapped in a dual power supply configuration where one of the power supplies is shut down and the system is still operational.

9.3 System Requirements for Thermal Compliance

For best system performance and to maintain the airflow necessary to keep the system operating below the non-critical thermal limit, the following requirements must be met when the system is operational:

1. The system top cover must be installed. The only exception to this requirement is when hot swapping a system fan. When performing this procedure, it must be completed as quickly as possible.
2. The system air duct must be installed.
3. All 48 memory slots must be populated with a memory device (DIMM or Intel® Optane™ PMem) or DIMM blank. Pre-installed DIMM blanks should only be removed when replacing it with a memory device.

Figure 67. DIMM Blank

4. All front drive bays must be populated with a drive carrier.
5. All drive carriers installed within the front drive bay must be assembled with a drive (HDD or SSD) or a supplied drive blank.
6. All system fans and power supply fans must be operational.
7. Ambient air flowing into the system must be between 10–35 °C (50–95 °F).
8. For proper system ventilation, leave a minimum of 15 cm clearance in the front and rear of the system.

10. System Power Overview

The Intel® Server System M70KLP supports up to two hot-swap capable power supply modules. Power supplies are installed from outside of the chassis into a single bay on the system back panel.

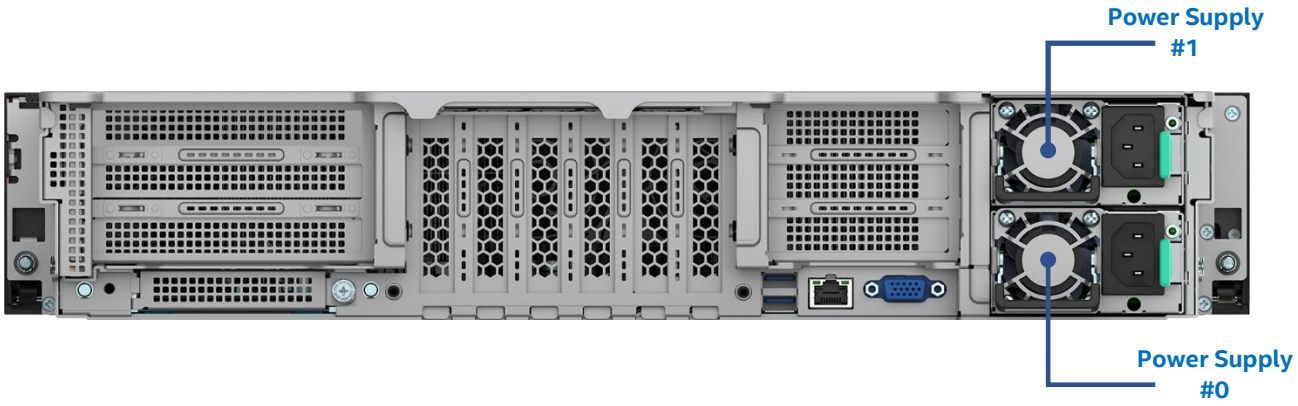

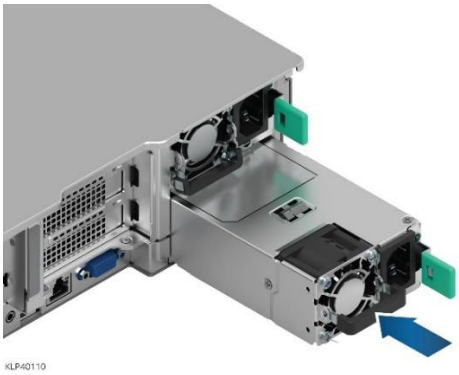

Figure 68. Power Supply Module Bay

When installing a power supply module, its card edge interface connector is blind mated with a matching slot connector on a power distribution board (PDB) mounted within the power supply bay.

10.1 Power Supply Configurations

The system supports the following power configurations:

- **1+0** (One power supply – no redundancy, only supported if the power draw of the system stays below the maximum power limit of the power supply)
- **1+1** (Two power supplies – redundant power, hot swap supported, only supported if the power draw of the system stays below the maximum power limit of one power supply)
- **2+0** (Two power supplies – combined power, no redundancy, enabled when the power draw of the system is greater than the maximum power limit of a single power supply)

Embedded platform management automatically determines and configures the power supply configuration based on the number of power supplies installed and the total power draw of the system.

10.1.1 Single Power Supply (1+0) Power Configuration

In a single power supply configuration, total available power to the system is limited to the maximum power capacity of the power supply. Anytime the system power draw exceeds the power limit of the power supply, platform management will limit I/O operations, also referred to as “throttling”, to system memory, processors, or both, to try to reduce the total system power draw. System performance will be degraded

should throttling occur. Throttling will remain enabled until the sustained system power draw falls below the maximum limit of the power supply. Should the system power draw go beyond the maximum limit of the power supply with throttling enabled, then platform management will shut down the system.

10.1.2 Dual Power Supply 1+1 Power Configuration

In a dual power supply configuration, if the total power draw from the system is less than or equal to the maximum power capacity of a single power supply, platform management will automatically configure the system to support 1 + 1 redundant power. In a redundant power configuration, if one of the power supplies shuts down, the backup or secondary power supply will automatically engage and provide the necessary power to maintain optimal system operation.

Note: When platform management detects a power supply that has shut down, several system errors and system status change events are logged to the system event log. The Power Supply Fault LED indicator on the front panel changes to Blinking Red, denoting that fan redundancy is lost, but the system is still in an operational state. In addition, the system power configuration changes to non-redundant 1 + 0 (see Section10.1.1) until the failed power supply is replaced.

The power supplies are hot swappable, allowing a failed power supply to be replaced without having to first power down the system. After replacing a failed power supply, platform management will automatically change the power configuration to either 1 + 1 or 2 + 0 depending on the total system power draw at the time the new power supply was detected.

10.1.3 Dual Power Supply 2+0 Power Configuration

In a dual power supply configuration, if the total power draw from the system is or becomes greater than the maximum power capacity of a single power supply, platform management will automatically configure the system to support a combined power (2 + 0) configuration. In this configuration, power from both power supplies will be used to supply the system with power to support optimal system operation.

In combined 2+0 power mode, total power available at peak power levels may be less than twice the rated power of the installed power supply modules. With the power supplies operating at peak levels, the amount of heat generated will prevent them from supplying maximum rated power.

If a power supply should fail and is shut down, platform management will attempt to reduce the total system power draw to below the maximum power limit of a single power supply by limiting system I/O operations, also referred to as “throttling”, to system memory, processors, or both. With throttling enabled, the system will still be operational but performance will be degraded until the failed power supply is replaced.

The power supplies are hot swappable, allowing a failed power supply to be replaced without having to first power down the system. After replacing a failed power supply, platform management automatically changes the power configuration to either 1 + 1 or 2 + 0, depending on the total system power draw at the time the new power supply was detected.

If a power supply failure occurs and shuts down and platform management is not able to reduce the power draw below the maximum power limit of a single power supply, then it will shut down the system.

Note: When platform management detects a power supply that has shut down, several system errors and system status change events are logged to the system event log. The Power Supply Status LED changes to Blinking Red if power redundancy is lost or it will change to Red if the system was shut down due to a failed power supply.

10.2 Power Supply Cold Redundancy

In dual power supply 1 + 1 redundant power configurations, by default, the BMC enables support for Cold Redundancy mode. Cold redundancy mode can put the redundant power supply into a low power (almost off) standby state. This is done to save energy at system idle while still being able to turn back on fast enough (within 100 µsec) in case of a power supply failure to keep the system operating normally.

In cold redundancy mode, the BMC assigns and identifies each power supply as either “Active” or “Cold Stand-by”. The Active power supply provides the system with power. The Cold Stand-by power supply is placed in a low power standby state and is a backup to the Active power supply in case of failure.

To support highest long-term reliability of each power supply, the BMC schedules a rolling re-configuration where installed power supplies alternate between being the “Active” and the “Cold Stand-by”. A rolling configuration allows for equal loading over the lifetime of each power supply.

10.3 Power Supply Options

The Intel® Server System M70KLP family supports a 2000-Watt AC (80-Plus Platinum) CRPS power supply module.

Figure 69. 2000W AC Power Supply Module

AC power supplies are auto-ranging and power factor corrected.

The following sections provide an overview of select power supply features and functions.

10.3.1 AC Power Cord Specifications

Power cords attached to a power supply must meet the following specifications.

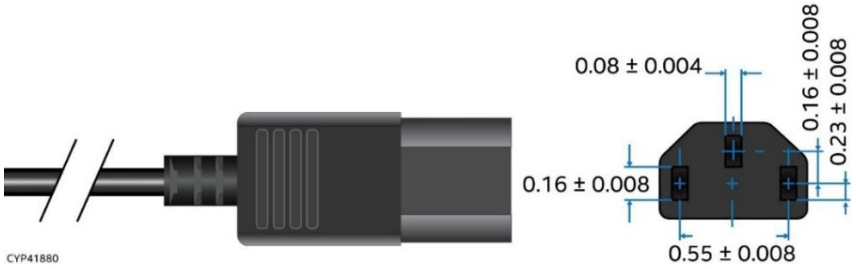

Figure 70. AC Power Cord Specification

Table 14. AC Power Cord Specification

Item	Description
Cable Type	SJT
Wire Size	14 AWG
Temperature Rating	105 ºC
Amperage Rating	10 A at 240 V
Voltage Rating	240 VAC

10.3.2 Power Supply Status LED

On the back panel of each installed power supply module is a Status LED that identifies the state the power supply is in. See Table 15 for a description of each supported LED state.

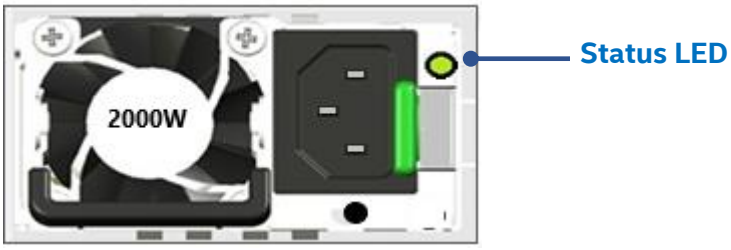

Figure 71. Power Supply Status LED

Table 15. Power Supply Status LED States

Condition	LED
No AC power to power supply (AC Cord unplugged or AC lost)	OFF
12V output mode normal	Green
AC input is normal, PSU operates in standby mode.	1 Hz blink Green
Sleep - PSU in cold redundant power state	0.33Hz blink Green, 1s off, 2s Green
Power supply critical event causing a shutdown; failure, over current, short circuit, over voltage, fan failure, over temperature.	Amber
Power supply warning events where the power supply continues to operate. high temp, high power, high current, slow FAN	1 Hz blink Amber
Power supply firmware updating mode	2 Hz blink Green

10.3.3 Power Supply Specification

Parameter		Description
		PSU (2000W)
Maximum Output Power (W)	High Line (180V~264Vac)	2000
	DC Line (180V~320Vdc)	1800
	Low Line (90V~132Vac)	1000
Maximum 12V Output Current (A)	High Line (180V~264Vac)	166
	DC Line (180V~320Vdc)	150
	Low Line (90V~132Vac)	83

Important Note: All fully integrated L9 system configurations from Intel will have a >1000W power draw with the system operating at moderate to heavy workloads. To support this power requirement, systems must be connected to a 220V AC input power source to operate. A system with a >1000W power draw that is connected to a 110V AC input power source will not operate correctly.

10.3.3.1 Efficiency

Efficiency shall be tested at AC input voltages of 230VAC/50 Hz as defined by 80 PLUS® standards.

Table 16. 2000W Power Supply Efficiency Data (80-Plus Platinum)

Input Voltage	DC Output Load			
	10% load	20% load	50% load	100% load
230VAC	88%	94%	94%	91%
270Vdc	87%	93%	93%	90%
Note: Fan loading is not included. Efficiency to be measured at 20~25 degree C ambient after power supply unit has run for 20 minutes.				

10.3.3.2 Protection Circuits

Designed into the power supply are several protection circuits that will cause the main DC power output of the power supply to shut down should a given fault condition occur. Stand-by power and the I^2C communications bus will continue to operate under a +12VDC fault protection condition. When a protection circuit shuts down the +12VDC output, the green status LED on the back panel of the power supply shall change to an Amber status.

Embedded protection circuits include:

- Current Limit
- Over Voltage / Under Voltage Protection
- Over Temperature Protection

10.4 Auxiliary 12-Volt Power

Each of the two riser card slots on the server board can supply a maximum of 100W power. When a riser card is installed, the 100W is split to support all add-in slots of the riser, with a single add-in slot maximum of 75W. For GPU, or other high-power add-in cards, that require >75W of power to operate, additional 12-Volt Auxiliary power must be cabled to the card. In the Intel® Server System M70KLP, additional AUX 12-Volt power for add-in cards can be drawn from two 8-pin (Black 2x4) AUX power connectors on a power distribution board found mounted to the power supply bay.

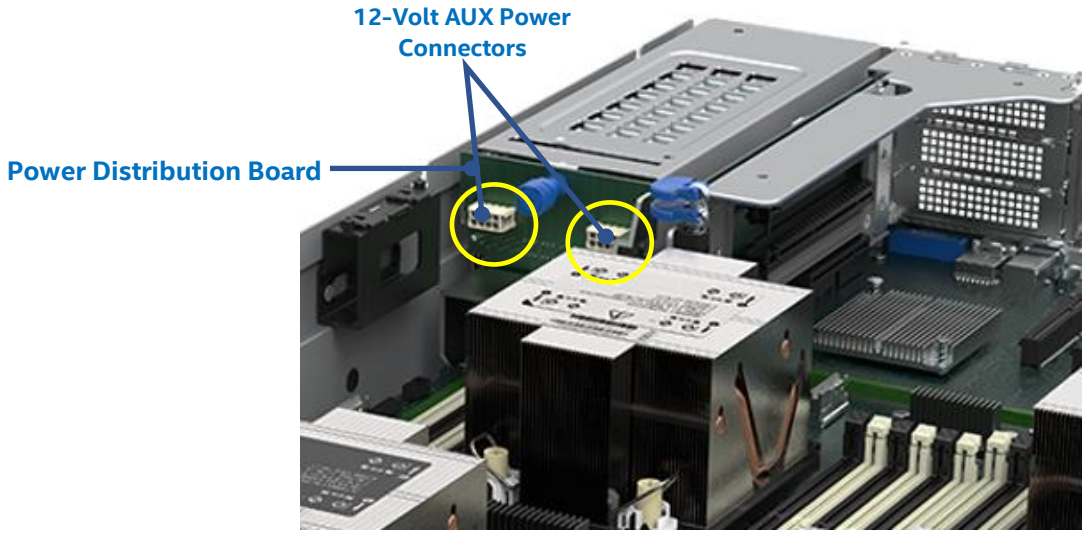

Figure 72. AUX 12-Volt Power Connectors

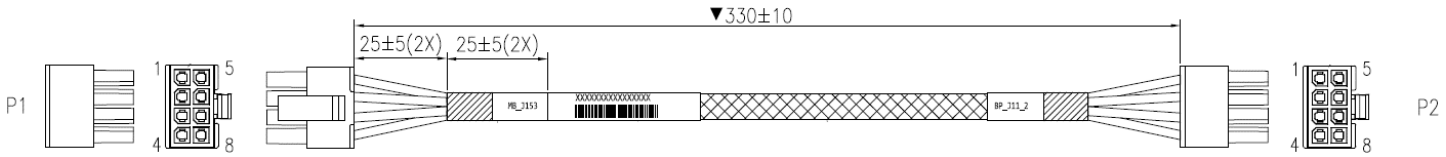

Figure 73. AUX 12-Volt Power Cable

Table 17. AUX 12-Volt Power Cable Pinout

Signal Name	Connector P1 Pin #	Wire Color	Connector P2 Pin #	Signal Name
GND	1	Black	1	GND
GND	2	Black	2	GND
GND	3	Black	3	GND
GND	4	Black	4	GND
P12V	5	Yellow	5	P12V
P12V	6	Yellow	6	P12V
P5V	7	Red	7	P5V
P5V	8	Red	8	P5V

11. System Security Features Overview

The Intel® Server System M70KLP family supports a variety of security options designed to prevent unauthorized access or tampering of settings. Security options supported include:

- Password protection
- Front panel lockout
- Intel® Platform Firmware Resilience (Intel® PFR)
- Trusted Platform Module (TPM) support
- Intel® CBnT – Converged Boot Guard and Trusted Execution Technology (Intel® TXT)
- Unified Extensible Firmware Interface (UEFI) Secure Boot Technology

11.1 Password Protection

The <F2> BIOS Setup utility includes a Security tab where options to configure passwords, front panel lockout, and TPM settings are found.

Figure 74. BIOS Setup Security tab

11.1.1 Password Setup

The BIOS uses passwords to prevent unauthorized access or modification of system settings and options. 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 system 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 server module to change BIOS settings.

An administrator password must be set and enabled before setting a user password.

The maximum length of a password is 14 characters. The minimum length is one character. The password 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. 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 server board. For more information on the password clear jumper, see Section TBD.

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 system 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.

11.1.2 System Administrator Password Rights

When the correct administrator password is entered, the user may 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.

11.1.3 Authorized System User Password Rights and Restrictions

When the correct user password is entered, the user can 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 if the administrator password is enabled. Users are restricted from booting in anything other than the boot order defined in setup by an administrator.

11.2 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.

11.3 Intel® Platform Firmware Resilience (Intel® PFR)

As the intensity, sophistication, and disruptive impact of security attacks continue to escalate, data centers are driving a holistic approach to protect their critical infrastructure. This approach includes protecting server systems at the firmware level, the lowest layers of the platform where threats are most difficult to detect. To address this, Intel has developed Intel® Platform Firmware Resilience (Intel® PFR) technology where platforms can provide security starting with power-on, system boot, and OS load activities.

The Intel® Server System M70KLP family supports Intel® PFR technology, a hardware-enhanced platform security that uses an Intel® FPGA to protect, detect, and correct platform firmware.

- **Protect:** Monitors and filters malicious traffic on system buses. All platform firmware is attested safe before code execution.
- **Detect:** Verifies integrity of platform firmware images before executing. Performs boot and runtime monitoring to assure server is running a known good firmware.
- **Recover:** Automatically restores corrupted firmware from a protected gold recovery image within minutes.

Critical firmware elements protected in an Intel® Server System M70KLP family include: BIOS, SPI Descriptor, BMC, Intel® Management Engine (Intel® ME), and Power Supply firmware. This capability to mitigate firmware corruption is an important industry innovation and provides an optimal solution for security-sensitive organizations.

Intel® PFR fully supports the National Institute of Standards and Technology (NIST*) proposed firmware resiliency guidelines (SP800-193) that have wide industry support.

11.4 Trusted Platform Module (TPM) Support

The Trusted Platform Module (TPM) option 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 Intel® Server System M70KLP family implements TPM as per *TPM PC Client Specifications revision 2.0*, published by the Trusted Computing Group (TCG).

A TPM device is optionally installed on the server board. It 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 the measurements 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 the TPM to provide additional system and data security (for example, Microsoft Windows* 10 supports BitLocker* drive encryption).

11.4.1 Trusted Platform Module (TPM) Security BIOS

The BIOS TPM support conforms to the TCG (Trusted Computing Group) PC Client Specific Implementation Specification for Conventional BIOS the TCG PC Client Specific 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) 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, see the *TCG PC Client Specific Implementation Specification*, the *TCG PC Client Platform Physical Presence Interface Specification*, and the *Microsoft Windows* BitLocker* Requirements* documents.

11.4.2 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 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 that requested the TPM command.

11.4.3 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 BIOS TPM 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. 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. This option is used to clear security settings for a newly initialized system or to clear a system for which the TPM ownership security key was lost.

11.5 Intel® CBnT – Converged Intel® Boot Guard and Trusted Execution Technology (Intel® TXT)

Previous generation Intel servers supported Intel® Boot Guard and Intel® Trusted Execution Technology (Intel® TXT).

Boot Guard

- Provides mechanism to authenticate the initial BIOS Code, before BIOS starts
- Hardware-based Static Root of Trust for Measurement (SRTM)
- Defends against attackers replacing/modifying the platform firmware

Intel® TXT

- Provides the ability to attest the authenticity of a platform configuration and OS environment; Establish trust
- Hardware-based Dynamic Root of Trust for Measurement (DRTM)
- Defends against software-based attacks aimed at stealing sensitive information

The two security features combined included some redundancies and inefficiencies between them. With this product generation, Intel rearchitected and fused together the two technologies into Intel® CBnT (Converged Intel® Boot Guard and Trusted Execution Technology). Combining the two technologies into one made them more efficient, eliminated redundancies between them, simplified their implementation, and provided stronger protections.

For more information visit http://www.intel.com/technology/security/.

12. System Management

The Intel® Server System M70KLP utilizes the baseboard management controller (BMC) features of an ASpeed* AST2500 server management processor. The BMC supports multiple system management features including intra-system sensor monitoring, fan speed control, system power management, and system error handling and messaging. It also provides remote platform management capabilities including remote access, monitoring, logging, and alerting features.

In support of system management, the system includes a dedicated management port and support for two system management tiers and optional system management software.

- Standard management features (Included)
- Advanced management features ($$ Optional)
- Intel® Data Center Manager (DCM) support ($$ Optional)

The following subsections provide a brief description of each.

12.1 Management Port

The Intel® Server System M70KLP includes a dedicated 1 Gb/s RJ45 management port used to access embedded system management features remotely.

Figure 75. Management Port

Note: The management port is dedicated for system management access purposes only. It is not intended or designed to support standard LAN data traffic.

To access the server remotely using the management port requires network parameters to be configured using the <F2> BIOS Setup utility.

12.1.1 Configuring System Management Port Using <F2> BIOS Setup

1. During the system power-on POST process, press **<F2>** when prompted to go to the BIOS Setup utility main menu page.
2. Navigate to the ***Server Management*** tab and select ***BMC Network Configuration*** to enter the BMC LAN Configuration screen (Figure 76).

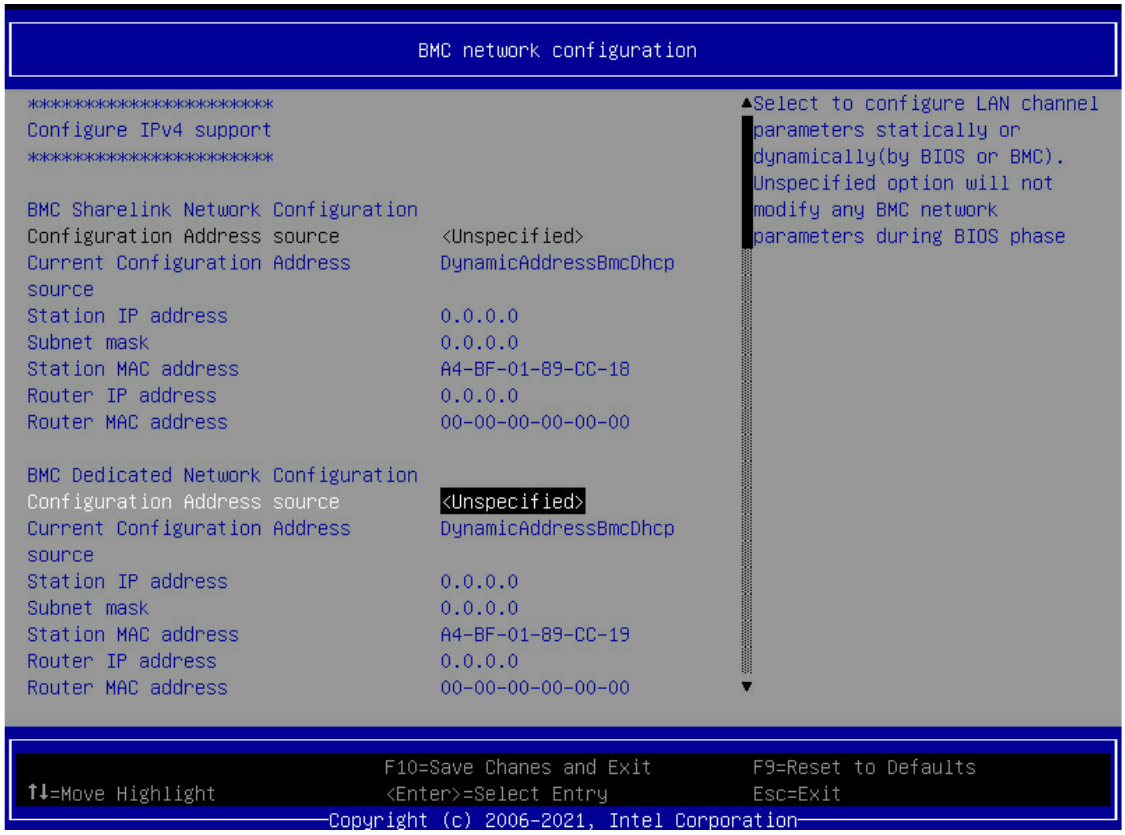

Figure 76. BIOS Setup BMC LAN Configuration Screen

3. The system is configured using the **BMC Dedicated Network Configuration**
4. For an IPv4 network:
   - If configuring the server management BMC LAN, scroll to **Baseboard LAN configuration** > **IP source** and then select either **Static** or **Dynamic**. If **Static** is selected, configure the **IP address**, **Subnet mask**, and **Gateway IP** as needed.
5. For an IPv6 network:
   - If configuring the server management BMC LAN, scroll to **Baseboard LAN IPv6 configuration** > **IP source** and then select **Enabled**. Then scroll to **IPV6 source** and select either **Static** or **Dynamic**. If **Static** is selected, configure the **IPV6 address**, **Gateway IPV6**, and **IPV6 Prefix Length** as needed.
6. The default **User Name** and **Password** is **admin/admin**.
7. If there is a need to change the User and Password, select **BMC User Settings** to enter the User Configuration screen (Figure 77).
8. Under **Add User**, set the following settings as desired:
   - **User Name** – Enter the desired name. Note that the anonymous user cannot be changed.
   - **User password** – Enter the desired password twice.
   - **Channel No** – Select the Channel
   - **Privilege** – Select the privilege to be used. (Administrator privilege is required to use KVM or media redirection)
9. Press **<F10>** to save the configured settings and exit BIOS Setup. The server reboots with the new LAN settings.

Figure 77. BIOS Setup User Configuration Screen

Once the management port is configured, the server can be accessed remotely to perform system management features defined in the following sections.

12.2 Standard System Management Features

The following system management features are supported on the Intel® Server System M70KLP by default.

- Virtual KVM over HTML5
- Integrated BMC Web Console
- Redfish
- IPMI 2.0
  - Node Manager
- Out-of-band BIOS/BMC Update and Configuration
- System Inventory
- Autonomous Debug Log

The following subsections provide a brief description for each feature.

12.2.1 Virtual KVM over HTML5

The BMC firmware supports keyboard, video, and mouse redirection (KVM) over LAN. This feature is available remotely from the embedded web server as a HTML5 application. USB1.1 or USB 2.0 based mouse and keyboard redirection are supported. It is also possible to use the KVM-redirection (KVM-r) session concurrently with media-redirection (media-r). This feature allows a user to interactively use the keyboard, video, and mouse (KVM) functions of the remote server as if the user were physically at the managed server.

KVM redirection consoles support the following keyboard layouts: English, Chinese (traditional), Japanese, German, French, Spanish, Korean, Italian, and United Kingdom. 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.

12.2.2 Integrated BMC Web Console

The BMC firmware has an embedded web server that can remotely serve web pages to any supported browser. This web console allows administrator to view system information including firmware versions, server health, diagnostic information, power statistics. The web console enables configuration of the BMC and BIOS. It provides the ability for users to perform power actions, launch KVM and set up virtual media redirection.

Enter the configured IP address of the BMC management port into the web browser to open the Integrated BMC Web Console module login page (See Figure 78). To use a secure connection, type:

```
https://<IPaddress_or_Hostname>/
```

Enter the username and password and select a language option. For example:

- Username: `root`
- Password: `superuser`
- Language: **English**

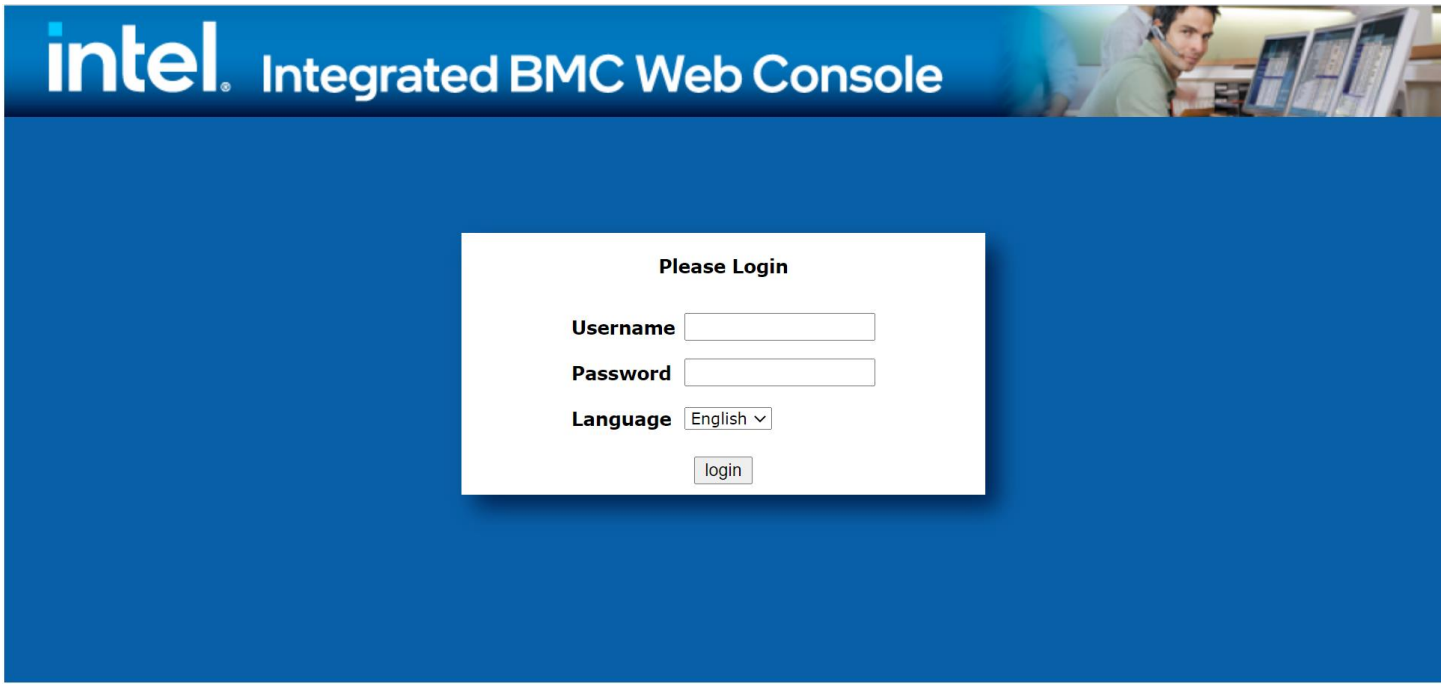

Figure 78. Integrated BMC Web Console Login Page

Click the **Login** button to view the home page.

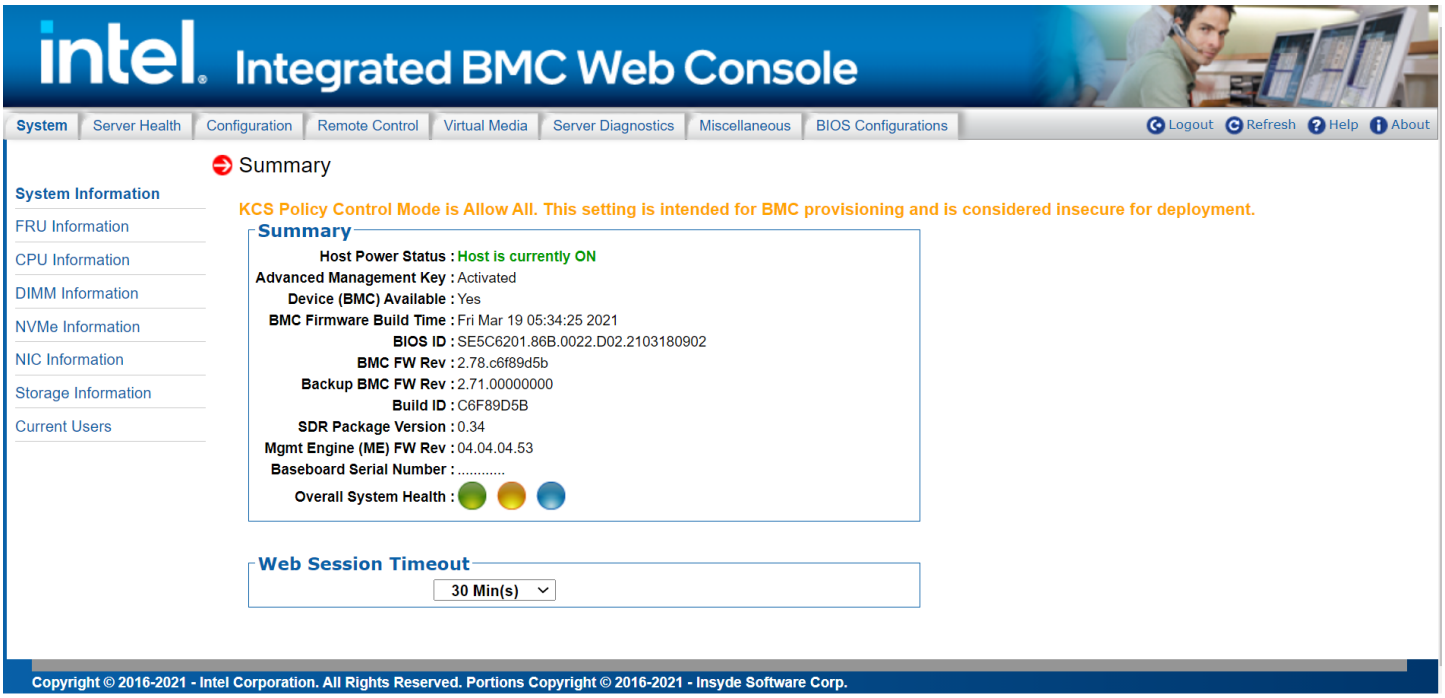

Figure 79. Integrated BMC Web Console – Main Console View

For setup and additional information about this utility, download *the Intel® Server System M70KLP Integrated BMC Web Console User Guide*.

12.2.3 Redfish* Support

The BMC currently supports version 1.7 and schema version 2019.1. DMTF's Redfish* is a standard designed to deliver simple and secure management for converged, hybrid IT and the Software Defined Data Center (SDDC). Both human readable and machine capable, Redfish leverages common Internet and web services standards to expose information directly to the modern tool chain.

12.2.4 IPMI 2.0 Support

The BMC is IPMI 2.0 compliant including support for Intel® Dynamic Power Node Manager. IPMI defines a set of interfaces used by system administrators for out-of-band management of computer systems and monitoring of their operation.

12.2.5 Out-of-Band BIOS / BMC Update and Configuration

The BMC supports Redfish schemas and embedded web console features that allow administrators to update the BMC, BIOS, Intel ME and SDR firmware. The BMC firmware also includes Power Supply and Back plan firmware. The BMC update will happen immediately and cause a BMC reset to occur at the end. The BIOS and Intel ME firmware is staged in the BMC and will be updated on the next reboot. The BMC also supports Redfish and embedded web console feature to view and modify BIOS settings. On each boot, BIOS provides all its settings and active value to the BMC to be displayed. BIOS also checks if any changes are requested and performs those changes.

12.2.6 System Inventory

The BMC supports Redfish schemas and embedded web console pages to display system inventory. This inventory includes FRU information, CPU, Memory, NVMe, networking, and storage. When applicable, the firmware version will also be provided.

12.2.7 Autonomous Debug Log

The BMC has a debug log that can be downloaded to facilitate support issues. This debug log can be downloaded from the embedded web console or via syscfg and SDPTool utilities. The debug log contains configuration data including SDR, SEL, BMC configuration, PCI configuration, power supply configuration and power supply black box data. The debug log also contains SMBIOS data and the POST codes from the last two system boots. Finally, when the system has a catastrophic error condition leading to a system shutdown, the BMC will hold the CPU in reset long enough to collect processor machine check registers, memory controller machine check registers, I/O global error registers, and other processor state info.

12.2.8 Security Features

The BMC contains several security features including OpenLDAP and Active Directory, security logs, ability to turn off any remote port, SSL certificate upload, VLAN support, and KCS control. The BMC also supports full user management with the ability to define password complexity rules. Each BMC release is given a security version number to prevent firmware downgrades from going to lower security versions. Intel provides a best practices security guide, available at
https://www.intel.com/content/www/us/en/support/articles/000055785/server-products.html

12.3 Advanced System Management Features

Purchasing an optional Advanced System Management product key (iPC ADVSYSMGMTKEY), unlocks the following advanced system management features:

- Virtual Media Image Redirection (HTML5 and Java)
- Virtual Media over network share and local folder
- Active Directory support
- Included single system license for Intel® Data Center Manager (Intel® DCM)
  - Intel® Data Center Manager (Intel® DCM) is a software solution that collects and analyzes the real-time health, power, and thermals of a variety of devices in data centers helping you improve the efficiency and uptime. For more information go to
    https://www.intel.com/content/www/us/en/software/intel-dcm-product-detail.html
- Future Feature Additions – Tentative availability Q4 2021
  - ❖ Full system firmware update including drives, memory, and RAID
  - ❖ Storage and network device monitoring
  - ❖ Out-of-band hardware RAID Management for latest Intel RAID cards

The Advanced System Management product key can be purchased and pre-loaded onto the system when ordering a fully integrated server system directly from Intel using its online Configure-to-Order (CTO) tool. Or, the Advanced System Management product key can be purchased separately and installed later. When purchasing the product key separately from the system, instructions will be provided on where to register the product key with Intel. A license file is then downloaded onto the system where the Integrated BMC Web Console or the SYSCFG utility are used to upload the key to the BMC firmware to unlock the advanced features.

12.3.1 Virtual Media Image Redirection (HTML5 and Java)

The BMC supports media redirection of local folders and .IMG and .ISO image files. This redirection is supported in both HTML5 and Java remote console clients. When the user selects “Virtual Media over HTML5”, a new web page will be displayed that provides the user interface to select which type of source media (image file or file folder*) to mount, and then allows the user to select the desired media to make available to the server system. After the type and specific media are selected, the interface provides a mount/unmount interface so the user can connect the media to or disconnect the media from the server system. Once connected, the selected image file or file folder is presented in the server system as standard

removable media and may be interacted within the normal fashion based on the operating system running on the server system. This feature allows system administrators the ability to install software (including operating system installation), copy files, perform firmware updates, and so on from media on their remote workstation.

Note: The file folder share is presented to the server system as a UDF file system; the server system operating system must be able to interact with UDF file systems for this feature to be used with the operating system.

12.3.2 Virtual Media over network share and local folder

In addition to supporting virtual media redirection from the remote workstation (see Section 12.3.1), the BMC also supports media redirection of file folders and .IMG and .ISO files hosted on a network file server accessible to the BMC network interface. The current version supports Samba shares (Microsoft* Windows* file shares). Future versions will add support for NFS shares. This virtual media redirection is more effective for mounting virtual media at scale, instead of processing all files from the workstation's drive through the HTML5 application and over the workstation's network. Each BMC makes a direct network file share connection to the file server and accesses files across that network share directly.

12.3.3 Active Directory support

The BMC supports Active Directory. Active Directory (AD) is a directory service developed by Microsoft* for Windows domain networks. This feature allows users to login to the web console or Redfish* using an active directory username instead of local authentication. The feature allows administrators to only change passwords on this single domain account instead on every remote system.

12.4 Intel® Datacenter Manager (DCM) Support

Intel® DCM is a solution for out-of-band monitoring and managing the health, power, and thermals of servers and a variety of other types of devices.

What can you do with Intel® DCM?

- Automate health monitoring
- Improve system manageability
- Simplify capacity planning
- Identify underutilized servers
- Measure energy use by device
- Pinpoint power/thermal issues
- Create power-aware job scheduling tasks
- Increase rack densities
- Set power policies and caps
- Improve data center thermal profile
- Optimize application power consumption
- Avoid expensive PDUs and smart power strips

For more information, go to

https://www.intel.com/content/www/us/en/software/intel-dcm-product-detail.html

Appendix A – System Regulatory Information

This product has been evaluated and certified as Information Technology Equipment (ITE) that 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, and so on), 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[2] **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 System M70KLP Family	Notes
	2U "Kelton Pass"	Intel Project Code Name
	L9 System	Product integration level
	TBD	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	April 2021	Testing In Progress
India BIS Certification	April 2021	Testing In Progress
International Compliance – CISPR32 & CISPR24	✓	
Japan VCCI Certification	✓	
Korea KC Certification	✓	
Mexico Certification	April 2021	Testing In Progress
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	✓

[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 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), went 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[3] **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 M70KLP Family System Service Guide*
    https://www.intel.com/content/www/us/en/support/products/208840/server-products/server-systems/intel-server-system-m70klp-family.html
- **Product Specifications**
  - *Intel® Server System M70KLP Family Technical Product Specification (TPS)*
    https://www.intel.com/content/www/us/en/support/products/208840/server-products/server-systems/intel-server-system-m70klp-family.html
- **System BIOS/Firmware and Security Updates – Intel® Server Board M70KLP family**
  - System Update Package (SUP) – uEFI only
    https://downloadcenter.intel.com/product/208840/Intel-Server-System-M70KLP-Family
- **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

[3] 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 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) – Support Summary

Intel® Server System M70KLP Family (Kelton Pass)

a template to report information needed 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 this list. Use of this information is at the sole risk of the user, and Intel assumes no responsibility for customers 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>M70KLP</td></tr>
<tr><td>Year of Launch</td><td>2021</td></tr>
<tr><td>PSU efficiency at 10%, 20%, 50% and 100% of rated output power</td><td>Great Wall GW-CRPS2000DW – 2000W AC (Platinum)
80Plus ID Number SO-1316
<table>
<tr><th>Model</th><th>100%</th><th>50%</th><th>20%</th><th>10%</th></tr>
<tr><td>GW-CRPS2000DW</td><td>92.87 %</td><td>94.74 %</td><td>92.59 %</td><td>91.44 %</td></tr>
</table></td></tr>
<tr><td>PSU factor at 50% of rated load level</td><td>Great Wall GW-CRPS2000DW: 1.00</td></tr>
<tr><td>PSU Rated Power Output
(Server Only)</td><td>Great Wall GW-CRPS2000DW: 2000W</td></tr>
<tr><td>Idle state power (Server only)
(Watts)</td><td>Refer to the following table - Energy Efficiency Data</td></tr>
<tr><td>List of all components for additional idle power allowances (server only)</td><td>Refer to the following table - Energy Efficiency Data</td></tr>
<tr><td>Maximum power (Server only)</td><td>Refer to the following table - Energy Efficiency Data</td></tr>
<tr><td>Declared operating condition class

** Refer to Appendix D - System Configuration Table for Thermal Compatibility for specific ASHRAE support by system configuration</td><td>• ASHRAE Class A2-Continuous Operation 10 °C to 35 °C with the maximum rate of change not to exceed 10 °C per hour</td></tr>
<tr><td>Idle State Power (watts) at the higher boundary temp (Server Only)</td><td>Refer to the following table - Energy Efficiency Data</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 - Energy Efficiency Data</td></tr>
<tr><td>Information on the secure data deletion functionality</td><td>Refer to the following table - Energy Efficiency Data</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 M70KLP4S2UHH – 4 CPU Installed Configuration

Configuration				
			4-CPU Low-end config.	4-CPU High-end Config.
Details	Chassis	Model	M70KLP4S2UHH	M70KLP4S2UHH
	Motherboard (MB)	# of MBs installed in system	1	1
		Model	M70KLP2SB	M70KLP2SB
	Processor	# of Processors per MB	4	4
		Processor Model	Intel® Xeon® Scalable Gold 6328HL	Intel® Xeon® Scalable Platinum 8356H
	Memory	# of DIMMs installed per MB	24 (1 DIMM / Mem. Channel)	48 (2 DIMM / Mem. Channel)
		Capacity per DIMM (GB)	64 GB	64 GB
		Total Memory (GB) per MB	1534 GB	3072 GB
	SSD	Total # of SSDs installed	2	2
	Power Supply (PSU) Chassis	Total # of PSU installed	2	2
		Model	Great Wall GW-CRPS2000DW 2000W	Great Wall GW-CRPS2000DW 2000W
	System Software Revisions installed to each Node or MB		BIOS: SE5C620.86B.01.01.0007.121220201907 BMC: 1.06.ad2174e7	
Data Summary				
Measured and Calculated Server Allowance	P Base		n/a	n/a
	Additional CPU		n/a	n/a
	Additional Power Supply		n/a	n/a
	Storage Devices		n/a	n/a
	Additional Memory		n/a	n/a
	Additional I/O Device (10Gx 15W/2Port on MB)		n/a	n/a
	Perfcpu		n/a	n/a
Limits/ Results	Idle power allowances (W)		n/a	n/a
	Idle power tested (W) Per node		182.8 W	188.2 W
	Minimum Eff_{ACTIVE}		n/a	n/a
	Eff_{ACTIVE} tested		33.4	21.8
Other test result	Idle Power at Higher Temp. (per Node) @ 35 degree C		181.6 W	198.4 W
	Max Power (Per Node)		1,079.3 W	1,426.9 W

Other Information:

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

Appendix B – Statement of Volatility

The tables in this section are used to identify the volatile and non-volatile memory components for system boards used within the Intel® Server System M70KLP product family.

The tables provide the following data for each identified component.

- **Component Type**: Three types of components are on the server board assembly:
  - **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**: Size of each component in bits, kilobits (Kbits), megabits (Mbits), bytes, kilobytes (KB), or megabytes (MB).
- **Board Location**: Board location is the physical location of each component corresponding to information on the server board silkscreen.
- **User Data**: The flash components on the server boards do not store user data from the operating system. No operating system level data is retained in any listed components after AC power is removed. The persistence of information written to each component is determined by its type as described in the table.

  Each component stores data specific to its function. Some components may contain passwords that provide access to that device's configuration or functionality. These passwords are specific to the device and are unique and unrelated to operating system passwords. The specific components that may contain password data are:
  - **BIOS**: The server board BIOS provides the capability to prevent unauthorized users from configuring BIOS settings when a BIOS password is set. This password is stored in BIOS flash and is only used to set BIOS configuration access restrictions.
  - **BMC**: The server boards support an Intelligent Platform Management Interface (IPMI) 2.0 conformant baseboard management controller (BMC). The BMC provides health monitoring, alerting and remote power control capabilities for the Intel server board. The BMC does not have access to operating system level data.

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

Table 18. Server Board Components

Component Type	Size	Board location	User Data	Name
Non-Volatile	128 MB	U208	No(BMC FW)	BMC Flash
Non-Volatile	64 MB	U205	No(BIOS FW)	BIOS Flash
Non-Volatile	32KB	U72	Yes	Board Information/FRU
Volatile	512 MB	U65	No	BMC Firmware SDRAM
Non-Volatile	2Kb	U63	No	USB-TF Configuration Flash（not used）

Table 19. Backplane Components

Component Type	Size	Board location	User Data	Name
Non-Volatile	32KB	U10	Yes	Board Information/FRU

Table 20. Power Distribution Board Components

Component Type	Size	Board location	User Data	Name
Non-Volatile	32KB	U72	Yes	Board Information/FRU

Table 21. M.2 Interface Board Components

Component Type	Size	Board location	User Data	Name
Non-Volatile				N/A - None
Volatile				N/A - None

Appendix C – Glossary of Terms

Term	Definition
ACPI	Advanced Configuration and Power Interface
AD	App Direct
AMIBCP	AMI* BIOS Configuration Program
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
AVLS	Adaptive Virtual Lockstep
BMC	Baseboard Management Controller
BIOS	Basic Input/Output System
Intel® CBnT	Converged Boot Guard and Intel® Trusted Execution (Intel® TXT)
CPU	Central Processing Unit
DCM	Data Center Manager
DDR4	Double Data Rate 4th edition
DHCP	Dynamic Host Configuration Protocol
DIMM	Dual In-line Memory Module
DPC	DIMMs per Channel
ECC	Error Correcting Code
EFI	Extensible Firmware Interface
FRB	Fault Resilient Boot
FRU	Field Replaceable Unit
GPGPU	General Purpose Graphic Processing Unit
GPU	graphics processor unit
GUI	Graphical User Interface
I^2C	Inter-integrated Circuit bus
IMC	Integrated Memory Controller
IIO	Integrated Input/Output
iPC	Intel Product Code
IPMI	Intelligent Platform Management Interface
ISTA	International Safe Transit Association
KVM	Keyboard, Video, Mouse
LED	Light Emitting Diode
LRDIMM	Load Reduced DIMM
MCTP	Management Component Transport Protocol
MKTME	Multi-key Total Memory Encryption
MRC	Memory Reference Code
NIC	Network Interface Controller
NVMe*	Non-Volatile Memory Express (NVMe*) is an optimized, high-performance scalable storage interface designed to address the needs of enterprise systems that use PCIe®-based solid-state storage. NVMe* provides efficient access to non-volatile memory storage devices. NVMe* allows host hardware and software to take advantage of the levels of parallelism possible in modern SSDs.
NTB	Non-Transparent Bridge
OEM	Original Equipment Manufacturer
OCP*	Open Compute Project*
OR	Oct (8) Rank
PCH	Peripheral Controller Hub
PCI	Peripheral Component Interconnect
PCIe*	Peripheral Component Interconnect Express*

Term	Definition
Intel® PFR	Intel® Platform Firmware Resilience
PHM	Processor Heat sink Module
PMBus	Power Management Bus
PMem	Persistent Memory Module
POST	Power-on Self-Test
PSU	Power Supply Unit
PWM	Pulse Width Modulation
QR	Quad Rank
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RAS	Reliability, Availability, and Serviceability
RDIMM	Registered DIMM
RMFBU	RAID Maintenance Free Backup
ROC	RAID On Chip
SAS	Serial Attached SCSI
SATA	Serial Advanced Technology Attachment
SEL	System Event Log
SCSI	Small Computer System Interface
SDPTool	Intel® Server Debug and Provisioning Tool
SDR	Sensor Data Record
SMBus	System Management Bus
SMI	Structure of Management Information
SOL	Serial-over-LAN
SR	Single Rank
SSD	Solid State Device
SYSCFG	Intel® System Configuration Utility
SYSFWUPDT	Intel® System Firmware Update Utility
SYSINFO	Intel® System Information Retrieval Utility
SUP	System Update Package
TCG	Trusted Computing Group
TDP	Thermal Design Power
TPM	Trusted Platform Module
TPS	Technical Product Specification
TSOD	thermal sensor On DIMM
Intel® TXT	Intel® Trusted Execution Technology
USB	Universal Serial Bus
Intel® VMD	Intel® Volume Management Device
UEFI	Unified Extensible Firmware Interface
VSB	Voltage Standby
Intel® VROC	Intel® Virtual RAID on CPU