

# Huawei CH221 V5 Compute Node V100R001

## White Paper

**Issue** 03  
**Date** 2018-12-04



**Copyright © Huawei Technologies Co., Ltd. 2018. All rights reserved.**

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

## **Huawei Technologies Co., Ltd.**

Address: Huawei Industrial Base  
Bantian, Longgang  
Shenzhen 518129  
People's Republic of China

Website: <http://e.huawei.com>

# About This Document

## Purpose

This document describes the appearance, features, technical specifications, and configuration of the new-generation CH221 V5 compute node of the Huawei E9000 server.

## Intended Audience

This document is intended for:

- Huawei presales engineers
- Channel partner presales engineers
- Enterprise presales engineers

## Symbol Conventions

The symbols that may be found in this document are defined as follows.

Symbol	Description
	Indicates an imminently hazardous situation which, if not avoided, will result in death or serious injury.
	Indicates a potentially hazardous situation which, if not avoided, could result in death or serious injury.
	Indicates a potentially hazardous situation which, if not avoided, may result in minor or moderate injury.
	Indicates a potentially hazardous situation which, if not avoided, could result in equipment damage, data loss, performance deterioration, or unanticipated results. NOTICE is used to address practices not related to personal injury.

Symbol	Description
 NOTE	<p>Calls attention to important information, best practices and tips.</p> <p>NOTE is used to address information not related to personal injury, equipment damage, and environment deterioration.</p>

## Change History

Issue	Date	Description
03	2018-12-04	This issue is the third official release. Modified <b>1.5 Hardware Structure</b> .
02	2018-11-28	This issue is the second official release. Modified <b>1.5 Hardware Structure</b> .
01	2018-08-22	This issue is the first official release.

---

# Contents

---

<b>About This Document.....</b>	<b>ii</b>
<b>1 Introduction.....</b>	<b>1</b>
1.1 Functions.....	1
1.2 Appearance.....	2
1.3 Ports.....	5
1.4 Indicators.....	5
1.5 Hardware Structure.....	6
1.6 Logical Structure.....	13
1.7 Technical Specifications.....	14
<b>2 Features.....</b>	<b>16</b>
<b>3 Components.....</b>	<b>19</b>
<b>4 Management.....</b>	<b>26</b>
<b>5 Warranty.....</b>	<b>28</b>
<b>6 Certifications.....</b>	<b>29</b>

# 1 Introduction

---

## 1.1 Functions

This section describes the basic functions and components of the CH221 V5 compute node (CH221 V5).

The CH221 V5 is a full-width compute node that uses the new-generation Intel® Xeon® scalable processors (Skylake processors for the Purley platform). It provides PCIe resources for the compute node and supports GPU cards of a variety of specifications. The CH221 V5 is installed in an E9000 chassis and centrally managed by the MM910 management module.

The CH221 V5 can meet users' requirements for quick graphic processing, artificial intelligence (AI), and PCIe card acceleration applications. It is ideal for oil exploration, animation rendering, scientific computing, and seismic processing.

## 1.2 Appearance

### Appearance

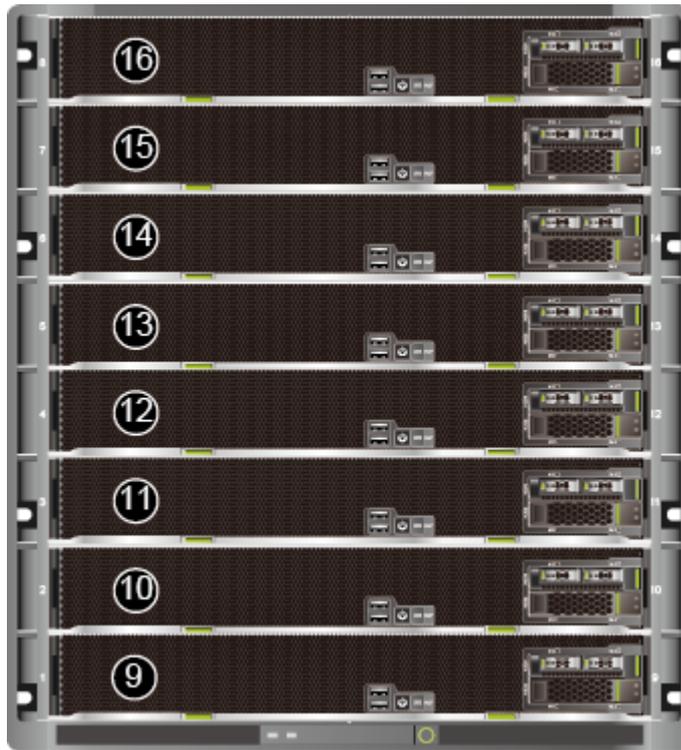
Figure 1-1 Appearance



### Installation Positions

An E9000 chassis can hold a maximum of eight CH221 V5 compute nodes in the full-width slots at the front of the chassis.

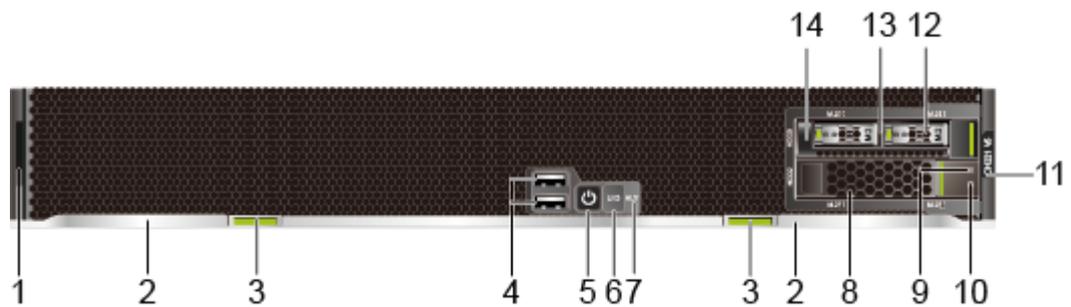
**Figure 1-2** Installation positions and slot numbering



## Front Panel

Each hard drive slot can be installed with one hard disk drive (HDD), SAS/SATA/NVMe solid-state drive (SSD), or M.2 module. The CH221 V5 supports mixed use of HDDs, SSDs, and M.2 modules.

**Figure 1-3** Front panel



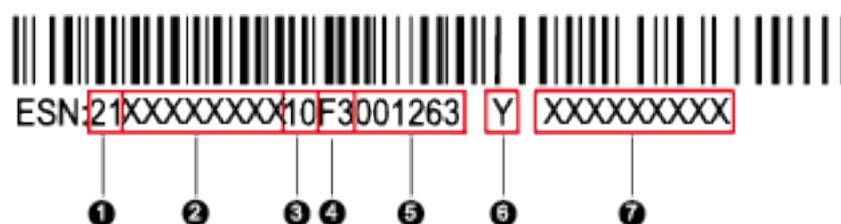
1	Slide-out label plate (with an ESN label)	2	Ejector levers
3	Ejector release buttons	4	USB 3.0 port
5	Power button/indicator	6	Unit Identification (UID) button/indicator
7	HLY indicator	8	2.5-inch drive

9	Hard drive fault indicator	10	Hard drive activity indicator
11	Product model	12	M.2 field replaceable unit (FRU)
13	M.2 adapter	14	M.2 module (consisting of one M.2 adapter and two M.2 FRUs)

## ESN

An Equipment Serial Number (ESN) uniquely identifies a compute node. An ESN is required when you apply for technical support from Huawei.

**Figure 1-4** ESN example



Callout No.	Description
1	ESN ID (two characters), which is 21 by default.
2	Material identification code (eight characters), that is, processing code.
3	Vendor code (two characters). The code <b>10</b> indicates Huawei, and other values indicate outsourcing vendors.
4	Year and month (two characters). <ul style="list-style-type: none"> <li>The first character indicates the year, where: <ul style="list-style-type: none"> <li>Digits 1 to 9 indicate 2001 to 2009, respectively.</li> <li>Letters A to H indicate 2010 to 2017, respectively.</li> <li>Letters J to N indicate 2018 to 2022, respectively.</li> <li>Letters P to Y indicate 2023 to 2032, respectively.</li> </ul> </li> </ul> <p><b>NOTE</b> The years from 2010 are represented by upper-case letters excluding I, O, and Z because the three letters are similar to digits 1, 0, and 2.</p> <ul style="list-style-type: none"> <li>The second character indicates the month, where: <ul style="list-style-type: none"> <li>Digits 1 to 9 indicate January to September, respectively.</li> <li>Letters A to C indicate October to December, respectively.</li> </ul> </li> </ul>
5	Serial number (six characters).

6	RoHS compliance (one character). <b>Y</b> indicates environmental protection processing.
7	Internal model, that is, product name.

## 1.3 Ports

**Table 1-1** Panel port description

Port	Quantity	Description
USB port	2	The panel provides two USB 3.0 ports, which are compatible with USB 2.0.

## 1.4 Indicators

**Table 1-2** Indicators on the front panel

Indicator	Meaning	Color	State Description
PWR 	Power button/indicator	Yellow and green	<ul style="list-style-type: none"> <li>● Off: The compute node is not powered on.</li> <li>● Blinking yellow: The power button is locked. When iBMC is being started during the compute node power-on process, the power button is locked.</li> <li>● Steady yellow: The compute node is to be powered on.</li> <li>● Steady green: The compute node is properly powered on.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>● When the compute node is powered on, pressing the power button will shut down the OS properly or cause the compute node to become unresponsive.</li> <li>● When the compute node is powered on, holding down the power button for 6 seconds will power off the compute node.</li> <li>● When the compute node is ready to be powered on, you can press the power button to power on the compute node.</li> </ul>

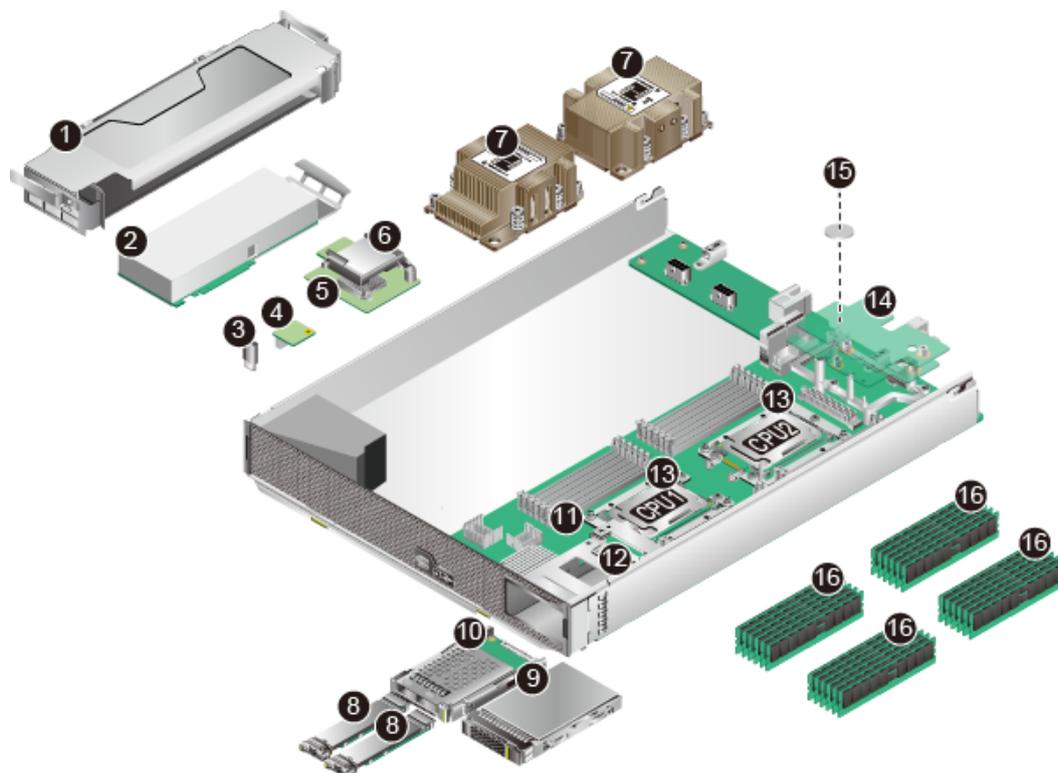
Indicator	Meaning	Color	State Description
UID	UID button/ indicator	Blue	<p>The Unit Identification (UID) indicator helps locate a compute node in a chassis.</p> <ul style="list-style-type: none"> <li>● On: indicates that the compute node has been located.</li> <li>● Blinking: distinguishes the compute node from multiple compute nodes that have also been located.</li> <li>● Off: the compute node has not been powered on or is not being located.</li> </ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"> <li>● You can remotely control the UID indicator status (off, on, or blinking) by using the MM910.</li> <li>● Press the UID button to turn on or off the UID indicator.</li> <li>● Hold down the UID button for 4 to 6 seconds to reset iBMC.</li> </ul>
HLY	Health status indicator	Red and green	<ul style="list-style-type: none"> <li>● Off: The compute node is not powered on.</li> <li>● Steady green: The compute node hardware is operating properly.</li> <li>● Blinking red (at 1 Hz): A major alarm is generated for the compute node.</li> <li>● Blinking red (at 5 Hz): A critical alarm is generated for the compute node, or the compute node is not securely installed.</li> </ul>
	Activity indicator of a hard drive	Green	<ul style="list-style-type: none"> <li>● Off: The drive is not detected or is faulty.</li> <li>● Blinking green: Data is being read from or written to the drive.</li> <li>● Steady green: The drive is inactive.</li> </ul>
	Fault indicator of a hard drive	Yellow	<ul style="list-style-type: none"> <li>● Off: The drive is operating properly.</li> <li>● Blinking yellow: The drive is being located, or RAID is being rebuilt.</li> <li>● Steady yellow: The drive is not detected or is faulty.</li> </ul>

## 1.5 Hardware Structure

This section describes the components, PCIe devices, and mainboard layout of the CH221 V5.

## Components

Figure 1-5 CH221 V5 components



1	PCIe riser module (supporting one dual-slot GPU)	2	(Optional) Standard PCIe card
3	(Optional) USB flash drive	4	(Optional) TPM card
5	RAID controller card	6	(Optional) Supercapacitor
7	Heat sinks	8	(Optional) M.2 FRUs
9	Drive	10	(Optional) M.2 adapter
11	Mainboard	12	Drive cage
13	CPUs	14	Mezzanine cards
15	BIOS battery	16	DIMMs

Table 1-3 Component description

No.	Name	Description
1	PCIe riser module	Provides two PCIe 3.0 x16 slots. The slot width is x16 for slot 1 and x8 for slot 2.

No.	Name	Description
2	(Optional) Standard PCIe card	Standard PCIe card. A dual-slot GPU card, such as V100, or P40, can be configured. The interface is PCIe 3.0 x16.
3	(Optional) USB flash drive	The mainboard provides one built-in USB port for connecting to a USB 3.0 device smaller than the following dimensions (H x D x W): 33.9 mm x 14.5 mm x 7.12 mm (1.33 in. x 0.57 in. x 0.28 in.)
4	(Optional) TPM card	TPM 2.0 is supported. The Trusted Platform Module (TPM) complies with Trusted Computing Group (TCG) standards and is used for device identification, authentication, and encryption.
5	RAID controller card	Provides two SAS or SATA ports for drive expandability. The CH221 V5 supports the following RAID controller cards: <ul style="list-style-type: none"> <li>● LSI SAS3008: supports RAID 0 and 1.</li> <li>● LSI SAS3108: supports RAID 0, 1, 5, 6, 10, 50, and 60. When used with the CH221 V5, the LSI SAS3108 supports only RAID 0 and 1.</li> <li>● Avago SAS3408: supports RAID 0 and 1.</li> <li>● Avago SAS3508: supports RAID 0, 1, 5, 6, 10, 50, and 60. When used with the CH221 V5, the Avago SAS3508 supports only RAID 0 and 1.</li> </ul> For details, use the <a href="#">Huawei Server Compatibility Checker</a> .
6	(Optional) Supercapacitor	Protects cache data from power failures if the Avago SAS3508 RAID controller card is used.
7	Heat sink	Dissipates heat from CPUs. Each CPU is configured with one heat sink.  If the CPU power consumption is less than or equal to 165 W, the heat sink of CPU 2 is higher than that of CPU 1.  If the CPU power consumption is greater than 165 W, the two heat sinks are of the same height. However, the heat sink of CPU 2 has a higher fin density than that of CPU 1.

No.	Name	Description
8	(Optional) M.2 FRU	<ul style="list-style-type: none"> <li>● An M.2 SSD with a SATA 3.0 interface is a high-speed, compact storage device with high stability.</li> <li>● The CH221 V5 supports the M.2 FRUs of the following specifications:                             <ul style="list-style-type: none"> <li>- <b>2280</b>: 22 mm x 80 mm (0.86 in. x 3.15 in.)</li> <li>- <b>2242</b>: 22 mm x 42 mm (0.86 in. x 1.65 in.)</li> </ul> </li> <li>● Each 2.5-inch drive slot can hold two M.2 FRUs and one M.2 adapter. A drive slot does not support mixed configuration of a drive and an M.2 FRU. A maximum of four M.2 FRUs can be configured.</li> <li>● Each M.2 FRU is hot-swappable and can be independently installed and removed.</li> </ul>
9	Drive	<p>A CH221 V5 supports a maximum of two 2.5-inch HDDs or SAS/SATA/NVMe SSDs, or M.2 modules. Mixed configuration of HDDs, SSDs, and M.2 module is supported. Each drive is hot-swappable and can be installed and removed separately.</p> <p><b>NOTE</b> The BIOS in Legacy mode does not support 4K native (4Kn) drives.</p>
10	(Optional) M.2 adapter	Provides two SATA 3.0 ports for connecting to the two M.2 FRUs in a 2.5-inch drive slot.
11	Mainboard	<p>Integrates components, including the CPUs, DIMMs, drive interface modules, power control module, iBMC, logic module, chipset, LAN on motherboard (LOM), and display adapter.</p> <p>The compute node chipset uses the Intel® C622 chip.</p> <p>The LOM (X722) is integrated in the PCH and provides two 10 Gbit/s Ethernet ports to connect to the Base network ports on the switch modules in slots 2X and 3X. The LOM supports WOL and PXE functions.</p> <p>The display adapter is integrated on the Hi1710 chip of the iBMC and provides 32 MB display memory. The maximum resolution is 1600 x 1200 at 60 Hz with 16 M colors.</p>
12	Drive cage	Houses hard drives.

No.	Name	Description
13	CPU	<p>The CH221 V5 supports one or two CPUs.</p> <ul style="list-style-type: none"> <li>Supported models: Intel® Xeon® scalable processors (Bronze 3100, Silver 4100, Gold 5100/6100, and Platinum 8100) with a maximum of 28 cores</li> <li>Built-in memory controllers: supports six DDR4 memory channels (two DIMMs per channel) and memory speeds of 2400 MT/s and 2666 MT/s</li> <li>Built-in PCIe controllers: supports PCIe 3.0 and 48 lanes per CPU</li> <li>Two Ultra Path Interconnect (UPI) links, each of which provides 10.4 GT/s transmission speed between the CPUs</li> <li>Maximum frequency: 3.6 GHz</li> </ul>
14	Mezzanine card	<p>Connects to the CH221 V5 to switch or pass-through modules through the midplane.</p> <p>A CH221 V5 has two mezzanine cards:</p> <ul style="list-style-type: none"> <li>Upper mezzanine card (Mezz 1): CPU 1 outputs PCIe 3.0 x16 to Mezz 1, which connects to slots 2X and 3X at the rear of the E9000 chassis.</li> <li>Lower mezzanine card (Mezz 2): CPU 2 outputs PCIe 3.0 x16 to Mezz 2, which connects to slots 1E and 4E at the rear of the E9000 chassis.</li> </ul> <p>For details about the mapping between mezzanine cards and CPUs, see <a href="#">Figure 1-7</a>.</p>
15	BIOS battery	Supplies power to the real time clock (RTC) when the CH221 V5 is powered off.
16	DIMM	<ul style="list-style-type: none"> <li>Maximum number of DIMM slots: 24 (12 for each CPU).</li> <li>DIMM types: registered dual in-line memory module (RDIMM) and load-reduced DIMM (LRDIMM)</li> <li>Maximum memory speed: 2666 MT/s</li> <li>Memory protection technology for reliability, availability, serviceability (RAS): advanced error-correcting code (ECC), memory mirroring, and memory sparing</li> </ul>

## PCIe Devices

[Table 1-4](#) provides the mapping between PCIe slots and CPUs, and the PCIe standards supported by the CH221 V5.

### NOTE

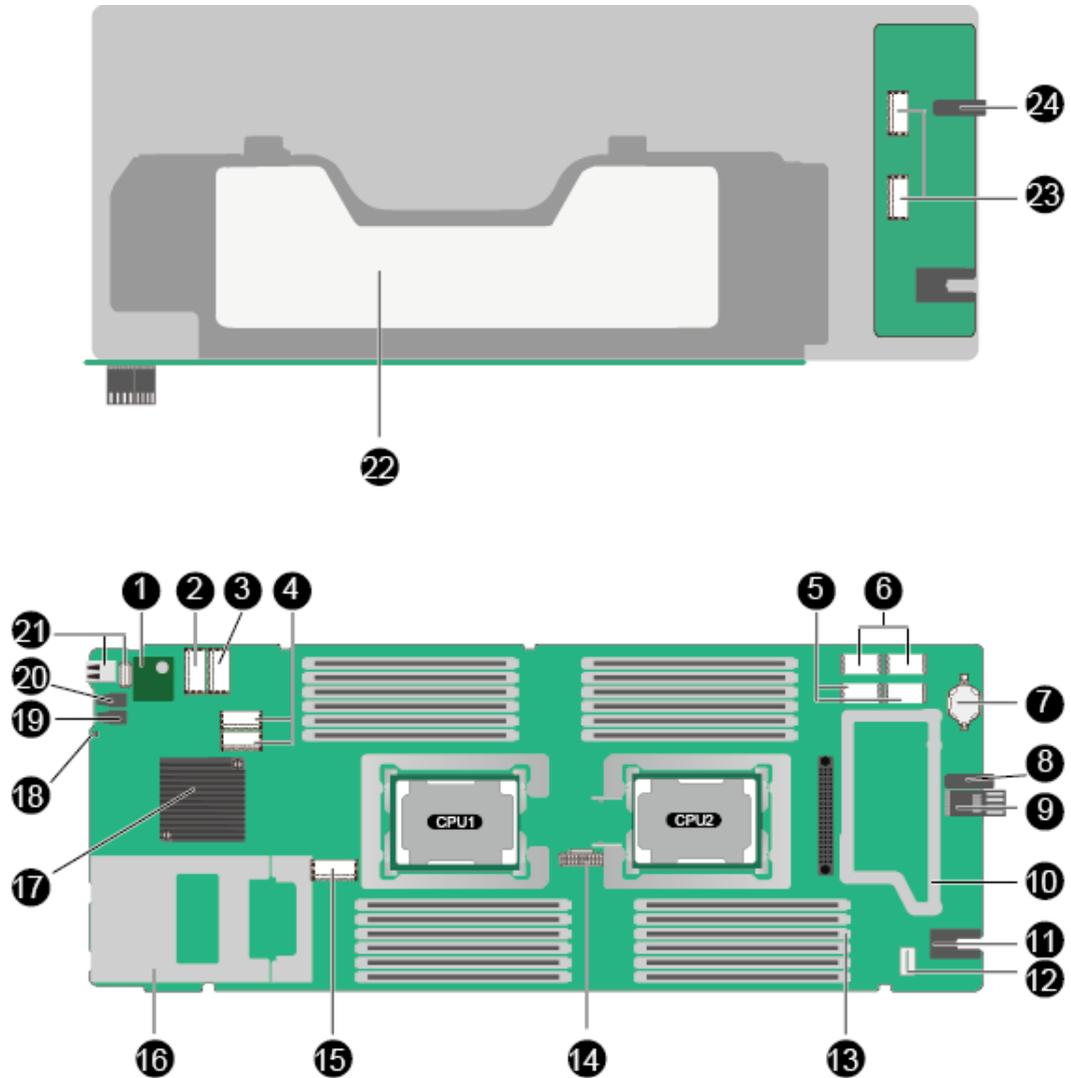
If CPU 2 is not installed, the related PCIe devices are unavailable.

**Table 1-4** PCIe device description

PCIe Device	Displayed on the BMC User Interface	CPU Socket	PCIe Standard	Connector Width	Bus Width	Port No.	Slot Size
RAID controller card	\	CPU 1	PCIe 3.0	x8	x8	Port 1A	Non-standard device
Mezz 1	\	CPU 1	PCIe 3.0	x16 or (x8 + x8)	x16 or (x8 + x8)	Port 2A or (Port 2A+Port 2C)	Non-standard device
Mezz 2	\	CPU 2	PCIe 3.0	x16 or (x8 + x8)	x16 or (x8 + x8)	Port 2A or (Port 2A+Port 2C)	Non-standard device
PCIe slot 1	PCI Slot1	CPU 1	PCIe 3.0	x16	x16	Port 3A	Full-height half-length (FHHL) or dual-slot GPU card (If a dual-slot GPU card is configured, slot2 cannot be used.)
PCIe slot 2	PCI Slot2	CPU 2	PCIe 3.0	x16	x8	Port 3C	FHHL
NVMe SSD 0	disk 0	CPU 2	PCIe 3.0	x4	x4	Port 1C	2.5-inch drive
NVMe SSD 2	disk 2	CPU 2	PCIe 3.0	x4	x4	Port 1D	2.5-inch drive

## Mainboard Layout

Figure 1-6 Positions of the connectors and components

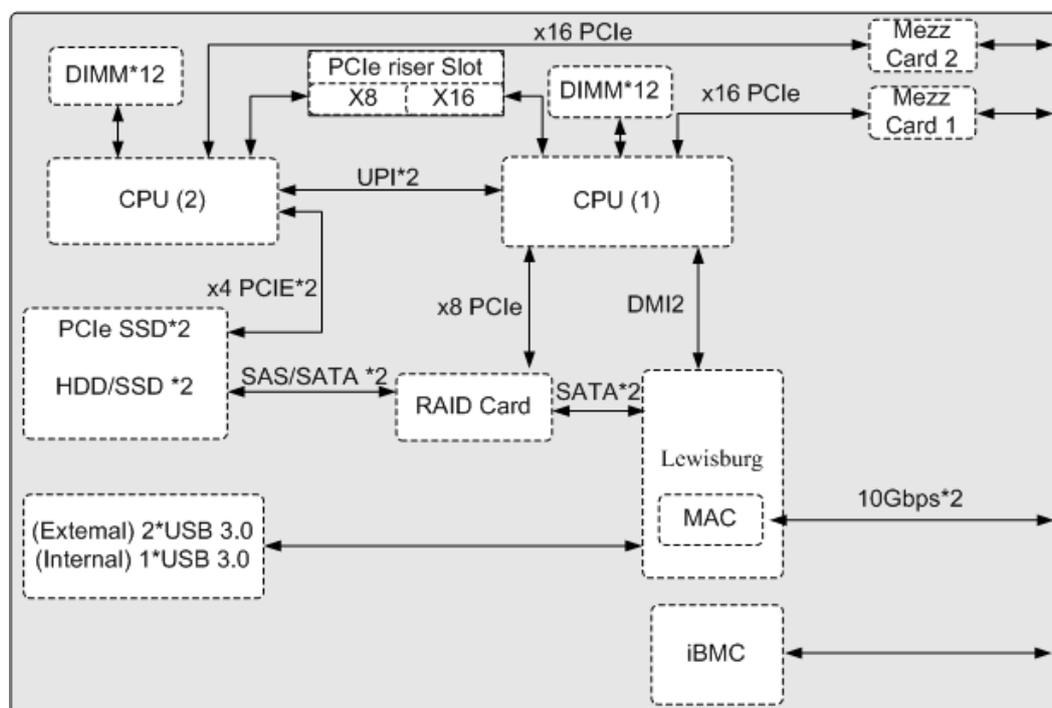


1	TPM card connector	2	PCIe riser module connector (PCIe 3.0 x16)
3	PCIe riser module connector (x8)	4	RAID controller card connectors
5	Mezzanine card 2 connectors	6	Mezzanine card 1 connectors
7	BIOS battery	8	Positioning sleeve
9	Midplane signal connector	10	Mezzanine card tray
11	Power connector	12	SoftRAID key connector
13	DIMM slots	14	CPU2 OPA sideband signal port

15	Drive backplane connector	16	Drive cage
17	Platform Controller Hub (PCH)	18	HLY indicator
19	UID button/indicator	20	Power button/indicator
21	USB 3.0 ports	22	PCIe card cage
23	GPU auxiliary power connector	24	Positioning pin

## 1.6 Logical Structure

Figure 1-7 Logical structure of the CH221 V5



The CH221 V5 uses Intel® Xeon® Scalable new-generation CPUs. Each CPU supports 12 DIMMs. The CPUs are interconnected through UPI links carrying up to 10.4 GT/s. CPU 1 connects to the Lewisburg chip through the DMI2 link at a speed of 8 GT/s. Through PCIe links, both CPUs connect to mezzanine cards that provide service ports.

The Lewisburg PCH is a next-generation Intel® southbridge chip used on server platforms. It supports external I/O interfaces and bus expansion, and is integrated with two MAC chips to provide two 10 Gbit/s interfaces.

The hard drive interface module consists of a RAID card and a hard drive backplane (HDD/SSD in the preceding figure). The hard drive interface module connects to the CPUs through the PCIe interface.

The iBMC provides device management functions, such as compute node power control, slot ID query, power supply monitoring, and KVM over IP.

## 1.7 Technical Specifications

**Table 1-5** Technical specifications

Category	Item	Specifications
Mechanical specifications	Dimensions (H x W x D)	60.46 mm x 423 mm x 537.2 mm (2.38 in. x 16.65 in. x 21.15 in.)
	Color	<ul style="list-style-type: none"> <li>● Front panel: black</li> <li>● Casing: silver</li> </ul>
	Weight	<ul style="list-style-type: none"> <li>● Net weight: 11 kg (24.26 lb)</li> <li>● Packing materials: 3.3 kg (7.28 lb)</li> </ul>
Environmental specifications	Temperature	<ul style="list-style-type: none"> <li>● Operating temperature: 5°C to 40°C (41°F to 104°F) (ASHRAE Class A3 compliant)</li> <li>● Storage temperature: - 40°C to +65°C ( - 40°F to +149°F)</li> <li>● Maximum change rate: 20°C (36°F)/hour</li> </ul> <p><b>NOTE</b> When Intel® Xeon® scalable processors (Platinum 8180/8180M/8168 or Gold 6154/6144/6146) are configured or the GPU V100, or P40 is configured, the maximum operating temperature is 35°C (95°F).</p>
	Relative humidity (RH, non-condensing)	<ul style="list-style-type: none"> <li>● Operating humidity: 5% to 85%</li> <li>● Storage humidity: 5% to 95%</li> <li>● Maximum change rate: 20%/hour</li> </ul>
	Maximum altitude	40°C@900 m (104°F@2952 ft.) When the altitude is higher than 900 m (2952.8 ft), the highest operating temperature decreases by 1°C (1.8°F) as the altitude increases by 300 m (984.24 ft). HDDs are not supported at an altitude of over 3000 m (9842.5 ft).
	Corrosive air pollutant	Maximum growth rate of the corrosion product thickness: <ul style="list-style-type: none"> <li>● Copper test piece: 300 Å/month (in compliance with the ANSI/ISA-71.04-2013 gaseous corrosion level G1)</li> <li>● Silver test piece: 200 Å/month</li> </ul>

Category	Item	Specifications
	Particle pollutant	<ul style="list-style-type: none"> <li>● The ISO14664-1 Class 8 requirements are met. A professional organization is required to monitor particle contaminants in the equipment room.</li> <li>● There is no explosive, conductive, magnetic, or corrosive dust in the equipment room.</li> </ul>
Input power specifications	Rated input voltage	12 V DC
Power consumption	Maximum power consumption	<p>955 W</p> <p>Configuration for testing the maximum power consumption:</p> <ul style="list-style-type: none"> <li>● CPU: Platinum 8168 x 2</li> <li>● Memory: 24 x 64 GB PC4 2666 MT/s</li> <li>● Drive: SATA SSD x 2</li> <li>● RAID controller card: LSI SAS3008</li> <li>● NIC: <ul style="list-style-type: none"> <li>- Mezz 1: MZ710</li> <li>- Mezz 2: MZ620</li> </ul> </li> <li>● PCIe card: V100</li> </ul>

# 2 Features

## Performance and Scalability

- Each Intel® Xeon® scalable processor (Bronze 3100, Silver 4100, Gold 5100/6100, or Platinum 8100) ensures high system performance by providing up to 28 cores, 3.6 GHz frequency, 38.5 MB L3 cache, and two 10.4 GT/s UPI links.
- Each compute node supports two CPUs, 56 cores, and 112 threads to maximize the concurrent execution of multithreaded applications.
- Intel® Turbo Boost Technology 2.0 enables CPU cores to run faster than the frequency specified in the thermal design power (TDP) configuration if they are operating below power, current, and temperature specification limits.
- Intel® Hyper-Threading Technology boosts performance for multithreaded applications by allowing each core to concurrently process up to two threads.
- Intel® virtualization technology integrates hardware-level virtualization functions, allowing OS vendors to better use hardware to address virtualization workloads.
- Intel® advanced vector extensions 512 (AVX-512) improves floating-point computing performance for compute-intensive applications.
- A CH221 V5 fully configured with 24 LRDIMMs provides supreme speed, high availability, and a maximum memory capacity of 3 TB.
- Two Intel® Xeon® scalable processors provide a maximum theoretical memory bandwidth of 256 GB/s (64 bits/8 x 2666 MHz x 6 channels x 2), 66.7% higher than the processors of the previous generation.
- The use of all SSDs provides higher I/O performance than the use of all HDDs or a combination of HDDs and SSDs. An SSD supports up to 100 times more I/O operations per second (IOPS) than a typical HDD.
- The CH221 V5 supports 96-lane PCIe 3.0 (8 GT/s per lane) to provide 20% higher I/O bandwidth than the previous 80-lane PCIe.
- Intel® integrated I/O technology enables the PCIe 3.0 controller to be integrated into Intel® Xeon® scalable processors. This shortens I/O latency and improves overall system performance.
- The CH221 V5 supports mezzanine cards with a variety of network ports to meet requirements for flexible networking.
- The CH221 V5 supports standard PCIe cards.

- A LOM with two 10GE ports is supported.

## Availability and Serviceability

- The CH221 V5 provides memory mirroring and memory module backup to prevent system shutdown caused by uncorrectable memory errors.
- Hot-swappable drives are provided to protect data and prolong system service life using RAID properties.
- The UID and HLY indicators on the panel and the intelligent baseboard management system (iBMC) WebUI help technical support personnel promptly obtain the status of key components and locate failed (or failing) components. This simplifies O&M, accelerates troubleshooting, and improves system availability.
- The use of SSDs offers better reliability than HDDs, ensuring continued system performance.
- The Huawei proprietary iBMC monitors system parameters in real time, triggers alarms, and performs recovery actions in case of failures, minimizing system downtime.

## Manageability and Security

- The built-in iBMC system monitors server operating status and provides remote management.
- The integrated industry-standard Unified Extensible Firmware Interface (UEFI) increases setup, configuration, and update efficiency and simplifies fault handling.
- The optional TPM 2.0 provides advanced encryption functions, such as digital signatures and remote authentication.
- The Advanced Encryption Standard–New Instruction (AES NI) algorithm allows faster and stronger encryption.
- Intel® Execute Disable Bit (EDB) prevents certain types of malicious buffer overflow attacks when working with a supported OS.
- Intel® Trusted Execution technology provides enhanced security by using hardware-based defense against malicious software attacks, allowing each application to run in isolated space.

## Energy Efficiency

- The Intel® Xeon® Scalable Platinum 8100 processors provide higher performance over the previous-generation processors, while fitting into 60 W higher TDP limit.
- Intel® intelligent power capability allows a CPU to be powered on or off based on service requirements.
- Low-voltage Intel® Xeon® processors consume less energy and apply to data centers and telecommunication environments that have power and thermal limitations.
- The 1.2 V DDR4 RDIMMs consume 20% to 30% less power than 1.35 V DDR3 RDIMMs.
- SSDs consume 80% less power than HDDs.
- The hexagonal ventilation holes on the ejector levers of the CH221 V5 provide better ventilation than round holes, significantly increasing system cooling efficiency.
- The efficient voltage regulator down (VRD) power supplies minimize the energy loss in DC/DC power conversion.

- The CH221 V5 provides power capping and power control functions.

# 3 Components

---

This section describes the software and hardware supported by the CH221 V5.

Use the [Huawei Server Compatibility Checker](#) to obtain information about the software and hardware supported by the CH221 V5.

## CPU

The CH221 V5 supports one or two CPUs.

- CPU models supported: Intel® Xeon® scalable processors, including Bronze 3100, Silver 4100, Gold 5100/6100, and Platinum 8100, with up to 28 cores.
- Built-in memory controllers: supports six DDR4 memory channels (two DIMMs per channel) and memory speeds of 2400 MT/s and 2666 MT/s.
- Built-in PCIe controllers: supports PCIe 3.0 and 48 lanes per CPU.
- Links between CPUs: two UPI links, each of which provides 10.4 GT/s transmission speed.
- Maximum frequency: 3.6 GHz

## Memory

The mainboard provides up to 24 DIMM slots (12 for each CPU). At least one DIMM must be configured.

### DIMM Configuration Rules

Observe the following rules when configuring DIMMs:

- The CH221 V5 supports DIMMs of 8 GB, 16 GB, 32 GB, 64 GB, or 128 GB. When fully configured with DIMMs, a CH221 V5 provides a maximum memory capacity of 3 TB.
- The maximum number of DIMMs to be installed varies depending on the CPU type, DIMM type, and rank quantity. For details, see **Maximum number of DIMMs** in [Table 3-1](#).

 **NOTE**

- The DIMM slots of CPU 1 must be configured with DIMMs. If DIMMs are installed only in the DIMM slots of CPU 2, the compute node cannot power on. [Figure 1-6](#) shows the positions of CPUs 1 and 2.
- The number of DIMMs supported by each channel is calculated as follows:  
Number of DIMMs supported by each channel ≤ Number of ranks supported by each memory channel/Number of ranks supported by each DIMM
- If different types of RDIMMs or LRDIMMs are installed in the same chassis for capacity expansion purposes, the memory RAS features will be affected.

Observe the following rules when you need to use different types of DIMMs in a compute node:

- Do not use RDIMMs and LRIMMs together.
- Do not use DIMMs of 128 GB together with DIMMs of other capacities.
- Install the DIMM more ranks in slots farther away from the CPU when DIMMs with different rank quantities need to be installed in the same channel. For example, install the dual-rank DIMM in slot 1A1 and the single-rank DIMM in slot 1A2.
- DIMMs of different speeds can be mixed in any way. However, the memory speed is the lowest speed of the DIMMs configured.
- RAS features, such as memory mirroring, SDDC, SDDC+1, and DDDC, are not supported if x4 and x8 DIMMs are used together.
- All DIMMs operate at the same speed, which is the smaller value of:
  - Memory speed supported by a specified CPU.
  - The smallest values among the maximum operating speed of specific DIMMs. For details, see the maximum operating speed in [Table 3-1](#).

**Table 3-1** Memory configuration rules for Intel® Xeon® scalable processors

Parameter	DIMM		
	Dual rank	Quad rank	Octal rank
Rank	Dual rank	Quad rank	Octal rank
Rated speed (MT/s)	2666	2666	2400
Rated voltage (V)	1.2	1.2	1.2
Operating voltage (V)	1.2	1.2	1.2
Maximum number of DIMMs <sup>a</sup>	24	24	24 <sup>b</sup>
Maximum capacity per RDIMM (GB)	32	64	128
Maximum memory capacity <sup>b</sup> (GB)	768	1536	3072 <sup>b</sup>
Maximum memory capacity at the maximum operating speed (GB)	768	1536	3072 <sup>b</sup>

Parameter		DIMM		
Maximum operating speed (MT/s)	One DIMM per channel	2666	2666	2400
	Two DIMMs per channel	2666	2666	2400

- a: The maximum number of DIMMs in this table is based on two-CPU configuration. If the compute node uses one CPU, the maximum number of DIMMs is half the values given in this table.
- b: The M series CPUs support the memory capacity of 1.5 TB/socket. The non-M series CPUs support the memory capacity of 768 GB/socket. When 128 GB of memory is configured, the CPU type must be considered.
- This table is for reference only. For details about the components that can be purchased, consult the local Huawei sales representatives.

### DIMM Slot Configuration Rules

Use the [Huawei Server product Memory Configuration Assistant](#) to configure DIMMs.

[Figure 3-1](#) and [Figure 3-2](#) show the DIMM installation. For details about the DIMM slot numbers, see [Figure 1-6](#). Unbalanced memory configuration is not recommended.

**Figure 3-1** DIMM installation sequence (1 CPU)

CPU	Channel	DIMM Slot	Number of DIMMs (✓: recommended ○: not recommended)												
			✓	✓	✓	✓	○	✓	○	✓	○	○	○	✓	
			1	2	3	4	5	6	7	8	9	10	11	12	
CPU1	A	DIMM000(1A1)	•	•	•	•	•	•	•	•	•	•	•	•	•
		DIMM001(1A2)													
	B	DIMM010(1B1)		•	•	•	•	•	•	•	•	•	•	•	•
		DIMM011(1B2)									•	•	•	•	•
	C	DIMM020(1C1)				•	•	•	•	•	•	•	•	•	•
		DIMM021(1C2)									•	•	•	•	•
	D	DIMM030(1D1)					•	•	•	•	•	•	•	•	•
		DIMM031(1D2)									•	•	•	•	•
	E	DIMM040(1E1)					•	•	•	•	•	•	•	•	•
		DIMM041(1E2)									•	•	•	•	•
	F	DIMM050(1F1)									•	•	•	•	•
		DIMM051(1F2)													•



CPU	Memory Channel	DIMM Slot
	1F	DIMM041(1E2)
		DIMM050(1F1)
		DIMM051(1F2)
CPU2	2A	DIMM100(2A1)
		DIMM101(2A2)
	2B	DIMM110(2B1)
		DIMM111(2B2)
	2C	DIMM120(2C1)
		DIMM121(2C2)
	2D	DIMM130(2D1)
		DIMM131(2D2)
	2E	DIMM140(2E1)
		DIMM141(2E2)
	2F	DIMM150(2F1)
		DIMM151(2F2)

## Storage

The CH221 V5 supports two 2.5-inch HDDs or SSDs and allows mixed configuration of an HDD and an SSD. Each drive can be independently installed, removed, and hot-swapped.

Each 2.5-inch drive slot can be configured with one M.2 module (consisting of one M.2 adapter and two M.2 FRUs). Each M.2 FRU can be installed, uninstalled, and hot swapped independently.

### NOTE

The drive installed with the OS cannot be moved to another compute node; otherwise, a virtual flash drive or CD/DVD-ROM cannot be mounted via the KVM.

The CH221 V5 supports the following RAID controller cards:

- LSI SAS3008: supports RAID 0 and 1.
- LSI SAS3108: supports RAID 0, 1, 5, 6, 10, 50, and 60. When used with the CH221 V5, the LSI SAS3108 supports only RAID 0 and 1.
- Avago SAS3408: supports RAID 0 and 1.
- Avago SAS3508: supports RAID 0, 1, 5, 6, 10, 50, and 60. When used with the CH221 V5, the Avago SAS3508 supports only RAID 0 and 1.

For details, use the [Huawei Server Compatibility Checker](#).

**Table 3-3** RAID level comparison

RAID Level	Reliability	Read Performance	Write Performance	Minimum Number of Hard Drives	Drive Utilization
RAID 0	Low	High	High	2	100%
RAID 1	High	Low	Low	2	50%

## I/O Expansion

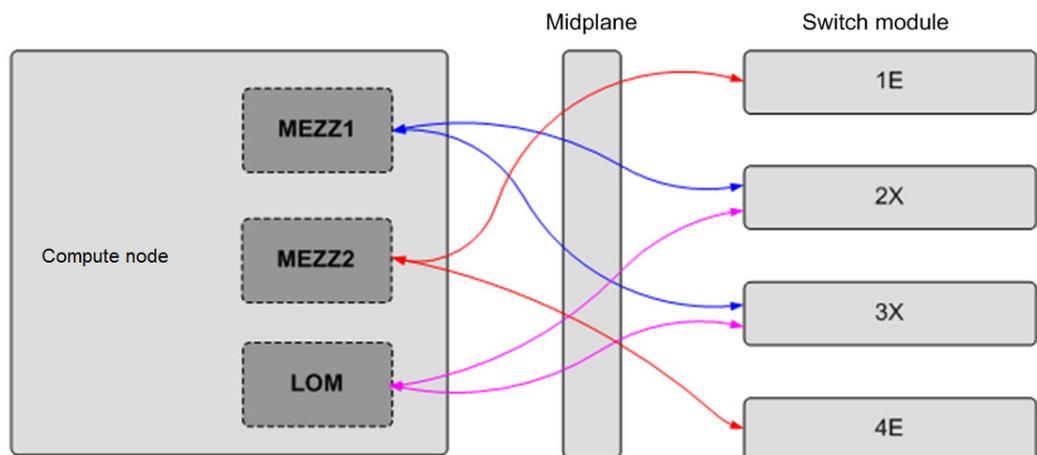
The CH221 V5 supports a variety of PCIe mezzanine cards, which are connected to switch modules through the midplane. You can select from the following PCIe mezzanine cards based on the I/O card type and transmission rate:

- GE expansion cards
- 10GE expansion cards
- 40GE expansion cards
- FC/FCoE expansion cards
- RoCE expansion cards
- IB expansion cards

## I/O Modules Matching the LOM

The LOM NIC can connect to I/O modules (switch modules).

**Figure 3-3** Connection between the compute node NIC slots and I/O modules



The compute node connects to the Fabric ports of switch modules through the network ports provided by the mezzanine cards on the compute node. Mezz 1 connects to Fabric ports of switch module in slots 2X and 3X. Mezz 2 connects to Fabric ports of switch module in slots 1E and 4E. The mezzanine card port types vary depending on the mezzanine card model.

The LOM (X722) is integrated in the PCH and provides two 10 Gbit/s Ethernet ports to connect to the Base network ports on the switch modules in slots 2X and 3X. The LOM supports WOL and PXE functions.

 **NOTE**

- If a compute node is forcibly powered off, the Wake On LAN (WOL) function of the LOM will become invalid.
- Flow control takes effect only when it is enabled on both the LOM and the connected switch modules.

**Table 3-4** I/O modules matching the LOM

I/O Module	I/O Module Slot	LOM	Remarks
CX916	2X/3X	√	-
	1E/4E	×	The LOM cannot communicate with the I/O module in slot 1E or 4E.
CX920	2X/3X	√	-
	1E/4E	×	The LOM cannot communicate with the I/O module in slot 1E or 4E.

## PSU

The CH221 V5 is powered by the power supply units (PSUs) in the E9000 chassis, without any independent power supply.

## Peripherals

The CH221 V5 supports the USB CD/DVD drive.

## OSs and Software

Use the [Huawei Server Compatibility Checker](#) to obtain information about the OSs and virtualization software supported by the CH221 V5.

# 4 Management

Huawei iBMC, a remote management system for servers, is integrated on E9000 compute nodes. It reliably monitors and manages hardware, seamlessly communicates with chassis management modules, and complies with IPMI V2.0 standards. The management modules can also be used to manage compute nodes in a chassis.

iBMC supports the followings:

- KVM and text console redirection
- Remote virtual media
- IPMI V2.0
- Common information model (CIM)
- Web-based browser login

**Table 4-1** iBMC specifications

Feature	Description
Management interface	Supports various management interfaces for system integration. It can be integrated with any standard management system through the following interfaces: <ul style="list-style-type: none"><li>● IPMI V2.0</li><li>● CLI</li><li>● HTTPS</li><li>● Redfish</li><li>● SNMP</li></ul>
Fault detection	Helps detect and accurately locates hardware faults.
System watchdog	Supports BIOS POST, OS watchdog, and automatic system reset after fault timeout. You can enable or disable these features individually.
Boot device configuration	Supports out-of-band configuration for boot devices.

Feature	Description
Alarm management	To ensure that the compute node runs properly without interruption, iBMC supports alarm management, reports alarms in various formats (such as SMTP), and supports syslog service.
Integrated KVM	Allows remote troubleshooting and maintenance using KVM or KVM over IP. The maximum resolution is 1600 x 1200.
Integrated virtual media	Virtualizes local media devices or images for remote compute nodes to facilitate OS installation. The virtual DVD drive supports a transmission rate of up to 8 MB/s.
WebUI	The browser-based WebUI allows quick configuration and information queries. The following web browsers are supported: <ul style="list-style-type: none"><li>● Internet Explorer 8.0</li><li>● Firefox 9.0</li><li>● Chrome 13.0</li><li>● Safari</li></ul>
Fault reproduction	Reproduces faults to help diagnose them quickly.
Screenshots and videos	iBMC allows you to view screenshots and videos without login, facilitating preventive maintenance inspection (PMI).
Black Box	Allows you to enable or disable the black box function and download black box data.
DNS/LDAP	Supports domain management and directory services, significantly simplifying network and configuration management.
Dual-image backup	If iBMC software fails, it starts again from a backup image.
Asset management	Provides intelligent asset management.
Intelligent power management	Supports power capping to increase deployment density and uses dynamic energy saving technology to lower OPEX.

# 5 Warranty

---

According to the *Huawei Warranty Policy for Servers & Storage Products (Warranty Policy for short)*, Huawei provides a three-year warranty for the server, a one-year warranty for DVD-RW drives and batteries, and a three-month warranty for software media. The *Warranty Policy* stipulates warranty terms and conditions, including the available services, response time, terms of service, and disclaimer.

The warranty terms and conditions may vary by country, and some services and/or parts may not be available in all countries. For more information about warranty services in your country, contact Huawei technical support or a local Huawei representative office.

# 6 Certifications

This section describes the certifications that the E9000 has passed.

**Table 6-1** Certifications

Country /Region	Certification	Standard
Europe	WEEE	2002/96/EC, 2012/19/EU
Europe	RoHS	2002/95/EC, 2011/65/EU, EN 50581: 2012
Europe	REACH	EC NO. 1907/2006
Europe	CE	Safety: EN 60950-1: 2006+A11: 2009+A1: 2010+A12: 2011 EMC: <ul style="list-style-type: none"><li>● EN 55022: 2010</li><li>● CISPR 22: 2008</li><li>● EN 55024: 2010</li><li>● CISPR 24: 2010</li><li>● ETSI EN 300 386 V1.6.1: 2012</li><li>● ETSI ES 201 468 V1.3.1: 2005</li></ul>
China	RoHS	SJ/T-11363-20006 SJ/T-11364-20006 GB/T 26572-2011
China	China Environmental Labeling	GB/T24024: 2001 idt ISO14024: 1999 HJ 2507-2011
Australia	C-tick	AS/NZS CISPR22: 2009
America	UL	UL 60950-1
America	FCC	FCC Part 15 (Class A)

Country /Region	Certification	Standard
America	NTRL-UL	UL 60950-1, 2nd Edition, 2011-12-19 (Information Technology Equipment - Safety - Part 1: General Requirements) CSA C22.2 No.60950-1-07, 2nd Edition, 2011-12 (Information Technology Equipment-Safety-Part 1: General Requirements)
Canada	IC	ICES-003 Class A
Nigeria	SONCAP	IEC 60950-1: 2005 (2nd Edition) + A1: 2009 EN 60950-1: 2006+A11: 2009+A1: 2010 + A12: 2011
Kingdom of Saudi Arabia (KSA)	SASO	IEC 60950-1: 2005 (2nd Edition) + A1: 2009 EN 60950-1: 2006+A11: 2009+A1: 2010 + A12: 2011
Global	CB	IEC 60950-1
Japan	VCCI	VCCI V-4: 2012