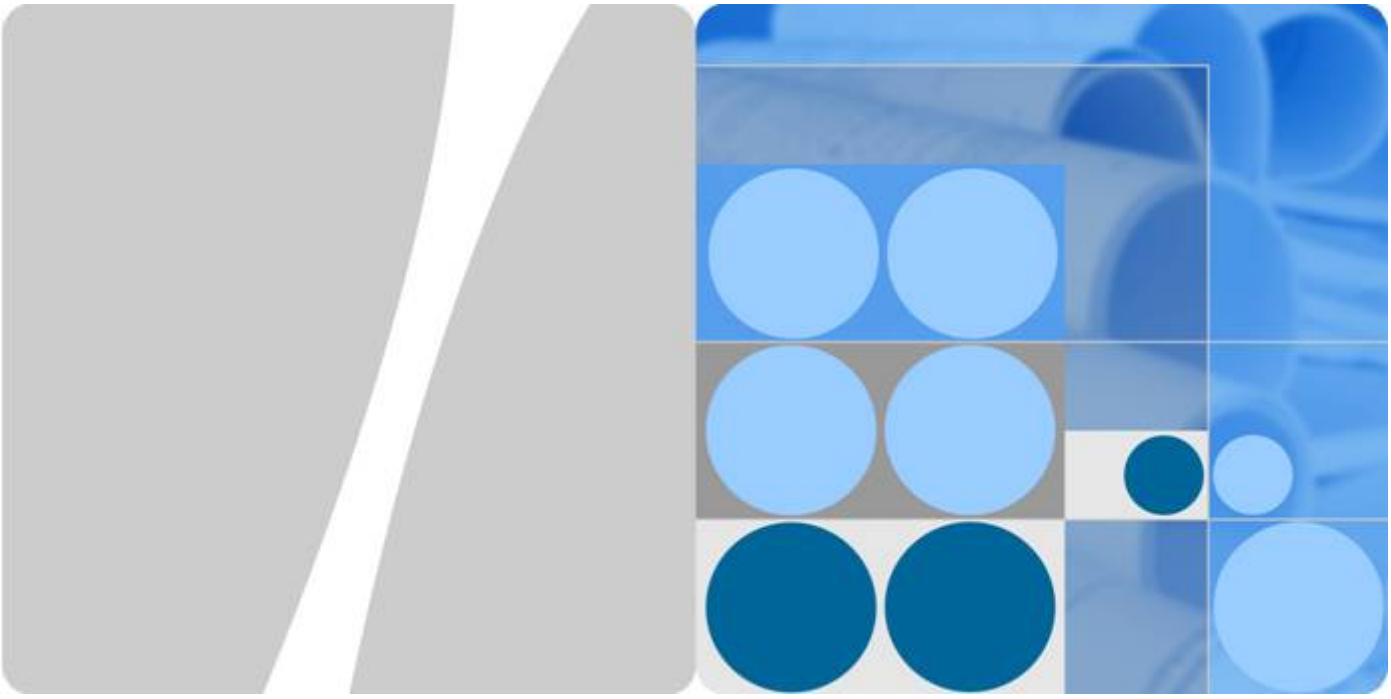

Huawei CH121 V5 Compute Node

V100R001

White Paper

Issue 02

Date 2017-10-12

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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://www.huawei.com

Email: support@huawei.com

About This Document

Purpose

This document describes the appearance, features, technical specifications, and configuration of the new-generation CH121 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

Change History

Changes between document issues are cumulative. The latest document issue contains all changes made in previous issues.

Issue 02 (2017-10-12)

This issue is the second official release.

Issue 01 (2017-07-18)

This issue is the first official release.

Contents

1 Introduction

1.1 Functions

The CH121 V5 is a half-width compute node that adopts new-generation Intel® Xeon® Scalable CPUs (that is, Purley Skylake CPUs) and large-capacity memory. It provides powerful computing capabilities and flexible scalability. The CH121 V5 compute nodes are installed in an E9000 chassis and are managed by the management module MM910 in a centralized manner.

The CH121 V5 combines dense computing capabilities with an ultra-large memory capacity. It is optimized for virtualization, cloud computing, high-performance computing, and compute-intensive enterprise applications.

1.2 Appearance

Figure 1-1 Appearance

Installation Positions

The CH121 V5 is installed in a half-width slot at the front of the E9000 chassis. A chassis can house a maximum of 16 CH121 V5 compute nodes. **Figure 1-2** shows the installation positions and slot numbering in the chassis.

Figure 1-2 Installation positions and slot numbering

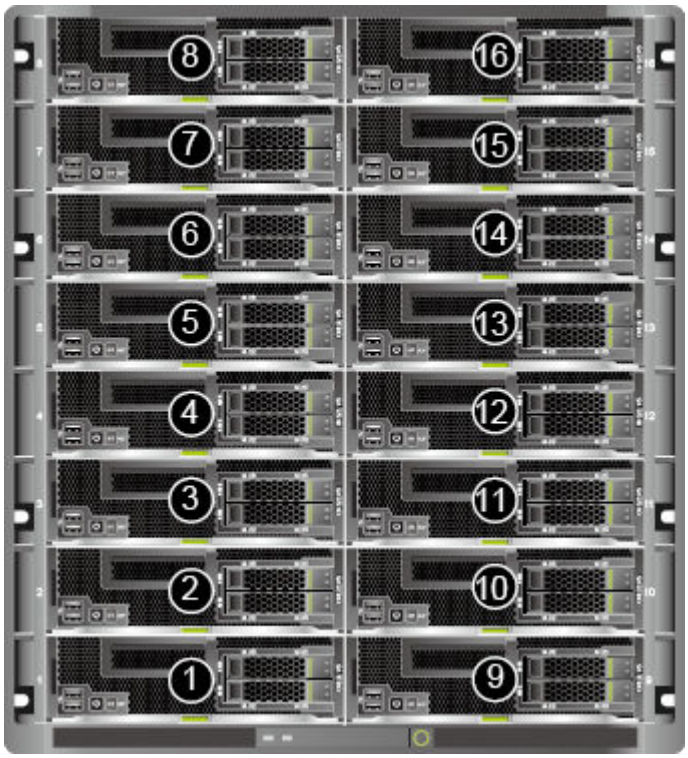

Front panel

Figure 1-3 shows the front panel of a CH121 V5.

Figure 1-3 Front panel

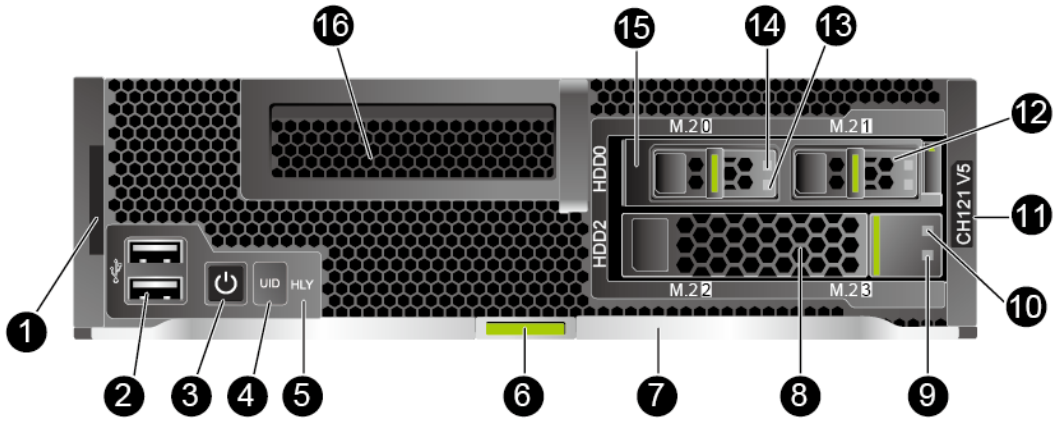

1	Slide-out label plate (with an ESN label)	2	USB 3.0 port
3	Power button/indicator	4	UID button/indicator
5	HLY indicator	6	Release button
7	Ejector lever	8	2.5-inch hard disk

9	Hard disk activity indicator	10	Hard disk fault indicator
11	Product model	12	M.2 FRU
13	Activity indicator of an M.2 FRU	14	Fault indicator of an M.2 FRU
15	M.2 module (consisting of one M.2 adapter and two M.2 FRUs)	16	PCIe panel

NOTE

Each hard disk slot supports an HDD, NVMe SSD, or M.2 module. The CH121 V5 supports mixed configuration of HDDs, NVMe SSDs, and M.2 modules.

Figure 1-3 is for reference only.

ESNs

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

Figure 1-4 shows the ESN format.

Figure 1-4 ESN example

No.	Description
1	Indicates the ESN ID (two digits).
2	Indicates the item identification code (eight characters).
3	Indicates the vendor code (two characters).
4	Indicates the year and month of production (two characters). The first character indicates the year. Digits 1 to 9 indicate 2001 to 2009, and letters A to Z indicate 2010 to 2035. The second character indicates the month. Digits 1 to 9 indicate January to September, and letters A to C indicate October to December.
5	Indicates the sequence number (six digits).
6	Indicates RoHS compliance (one character).
7	Indicates the internal model number of the compute node.

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

The front panel indicators on the CH121 V5 display its working status.

Table 1-2 Indicators on the front panel

Indicator	Meaning	Color	State Description
PWR 	Power button/ indicator	Yellow and green	● Off: The compute node is not powered on. ● Blinking yellow: The power button is locked. When iBMC is being started during the compute node power-on process, the power button is locked. ● Steady yellow: The compute node is to be powered on. ● Steady green: The compute node is properly powered on. **NOTE** ● When the compute node is powered on, pressing this button will shut down the OS properly or cause the compute node to become unresponsive. ● When the compute node is powered on, holding down this button for 6 seconds will power off the compute node. ● When the compute node is ready to be powered on, you can press this button to power on the compute node.

Indicator	Meaning	Color	State Description
UID	UID button/ indicator	Blue	The Unit Identification (UID) indicator helps locate a compute node in a chassis. The MM910 can be used to remotely activate the UID indicator to pinpoint the location of a compute node in a chassis. ● On: indicates that the compute node has been located. ● Blinking: distinguishes the compute node from multiple compute nodes that have also been located. ● Off: the compute node has not been powered on or is not being located. **NOTE** ● You can press this button to turn on or off the UID indicator. ● You can hold down the UID button for 4 to 6 seconds to reset iBMC.
HLY	Health status indicator	Red and green	● Off: The compute node is not powered on. ● Steady green: The compute node hardware is operating properly. ● Blinking red (at 1 Hz): A major alarm is generated for the compute node. ● Blinking red (at 5 Hz): A critical alarm is generated for the compute node, or the compute node is not securely installed.
	Activity indicator of a hard disk or M.2 FRU	Green	● Off: The hard disk or M.2 FRU is not detected or is faulty. ● Blinking green: Data is being read from or written to hard disk or M.2 FRU. ● Steady green: The hard disk or M.2 FRU is inactive.
	Fault indicator of a hard disk or M.2 FRU	Yellow	● Off: The hard disk or M.2 FRU is operating properly. ● Blinking yellow: The hard disk or M.2 FRU is being located, or RAID is being rebuilt. ● Steady yellow: The hard disk or M. 2 FRU is not detected or is faulty.

Indicator	Meaning	Color	State Description
	Activity indicator of an NVMe SSD	Green	● Off: The hard disk is faulty or not detected. ● Blinking green: Data is being read from, written to the hard disk, or synchronized between hard disks. ● Steady green: The hard disk is idle.
	Fault indicator of an NVMe SSD	Yellow	● Off: The hard disk is operating properly. ● Steady yellow: The hard disk is faulty or not detected.

1.5 Hardware Structure

This topic describes the components, PCIe devices, and mainboard layout of the CH121 V5.

Components

Figure 1-5 Components of the CH121 V5

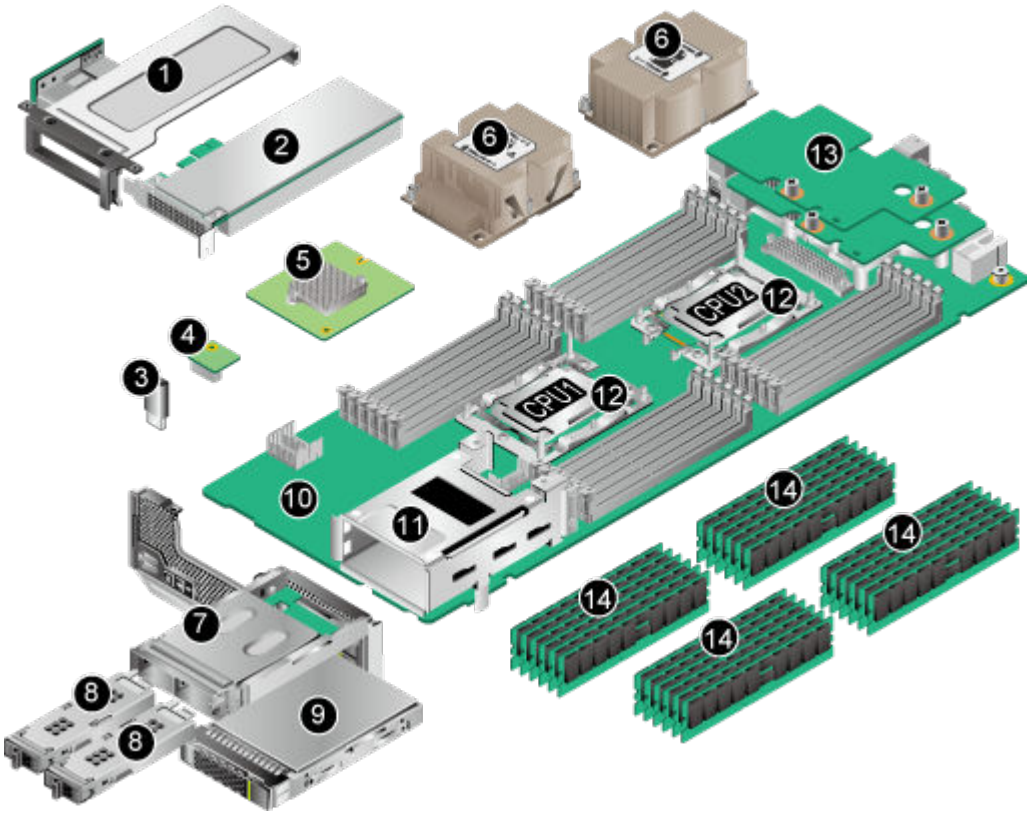

1	(Optional) Half-height half-length PCIe riser card	2	(Optional) Half-height, half-length standard PCIe card
3	(Optional) USB flash drive	4	(Optional) TPM card
5	RAID controller card	6	Heat sink
7	(Optional) M.2 adapter	8	(Optional) M.2 FRU
9	Hard disk	10	Mainboard
11	Hard disk tray	12	CPU
13	Mezzanine card	14	DIMM

Table 1-3 Component description

No.	Name	Description
1	(Optional) Half-height half-length PCIe riser card	The PCIe riser card provides one PCIe 3.0 x16 slot.
2	(Optional) Half-height, half-length standard PCIe card	The standard PCIe card (half-height and half-length) can be a PCIe SSD card of an 800 GB, 1.2 TB, 1.6 TB, 3.2 TB, or 6.4 TB capacity The interface is PCIe 3.0 x16.
3	(Optional) USB flash drive	The mainboard provides a built-in USB port for connecting to a USB 3.0 device within the dimensions of 33.9 mm x 14.5 mm x 7.12 mm (H x D x W).
4	(Optional) TPM card	TPM 2.0 is supported. The TPM is a security card that complies with the Trusted Computing Group (TCG) standards. It enhances platform security by preventing viruses or unauthorized operations.
5	RAID controller card	The RAID controller card connects to external hard disks to expand the storage capacity of the compute node. The RAID controller card provides two SAS or SATA ports to connect to 2.5-inch HDDs/SSDs or four SATA ports to connect to M.2 FRUs. The CH121 V5 supports the LSI SAS3008 RAID controller card, which supports RAID 0 and 1.
6	Heat sink	A heat sink cools a CPU. Each CPU is configured with one heat sink. If CPU power consumption is less than or equal to 165 W, the heat sink for CPU 2 is higher than that for CPU 1. If CPU power consumption is greater than 165 W, the two heat sinks are of the same height, but the heat sink for CPU 2 has a higher fin density than that for CPU 1.

No.	Name	Description
7	(Optional) M.2 adapter	An M.2 adapter can be installed into a 2.5-inch hard disk slot to provide two SATA 3.0 ports, which are used for connecting to two M.2 FRUs.
8	(Optional) M.2 FRU	An M.2 SATA 3.0 SSD is a high-performance, compact, and high-stability storage component. The M.2 2280 (22 mm x 80 mm) standard form factor is supported. Each 2.5-inch hard disk slot can be configured with two M. 2 FRUs, which require one M.2 adapter. A hard disk slot does not support mixed configuration of a 2.5-inch HDD/SSD and an M.2 FRU. A maximum of four M.2 FRUs are supported. Each M.2 FRU is hot-swappable and can be installed and removed separately.
9	Hard disk	The compute node supports a maximum of two 2.5-inch HDDs or SSDs (SAS, SATA, or NVMe SSDs). Mixed configuration of HDDs, SSDs, and M.2 modules is supported. Each HDD or SSD is hot-swappable and can be installed and removed separately. **NOTE** If the BIOS is in Legacy mode, 4K native (4Kn) disks are not supported.
10	Mainboard	The mainboard holds the CPUs, DIMMs, hard disk interface module, power control module, iBMC (integrated baseboard management controller), logic module, chipset, LOM,and display adapter. The compute node chipset is the Platform Controller Hub (PCH) using the Intel® C622 chip. A video chip with 32 MB display memory is integrated into the Hi1710 chip of iBMC. The maximum resolution is 1600 x 1200 at 60 Hz with 16 M colors.
11	Hard disk tray	Houses hard disks.
12	CPU	The mainboard supports one or two CPUs. ● Intel® Xeon® Scalable full-series CPUs (Bronze 3100, Silver 4100, Gold 5100/6100, and Platinum 8100) with up to 28 cores ● Each CPU is integrated with memory controllers and supports six DDR4 memory channels. Each channel supports two DDR4 DIMMs at 2400 or 2666 MT/s. ● Each CPU integrates a PCIe controller with 48 lanes and supports PCIe 3.0. ● Every two CPUs are interconnected through two Ultra Path Interconnect (UPI) links, with each link delivering up to 10.4 GT/s. ● The maximum core frequency supported by the CPUs reaches 3.6 GHz.

No.	Name	Description
13	Mezzanine card	The mainboard provides two mezzanine card connectors to connect to the switch or pass-through modules through the midplane. ● The upper mezzanine card is Mezz 1, and the lower one is Mezz 2. ● Socket CPU 1 provides PCIe 3.0 x16 bandwidth for connecting to Mezz 1, and socket CPU 2 provides PCIe 3.0 x16 bandwidth for connecting to Mezz 2. **Figure 1 Logical structure of the CH121 V5** shows the connections between mezzanine cards and CPUs. ● Mezz 1 connects to slots 2X and 3X at the rear of the E9000 chassis. ● Mezz 2 connects to slots 1E and 4E at the rear of the E9000 chassis.
14	DIMM	Up to 24 DIMM slots for installing DIMMs (12 DIMMs for each CPU). ● Maximum memory speed: 2666 MT/s. ● Memory protection technology for reliability, availability, serviceability (RAS): advanced error checking and correcting (ECC), memory mirroring, SDDC, and memory sparing. ● DIMM type: registered DIMM (RDIMM) and load-reduced DIMM (LRDIMM).

PCIe Devices

Table 1-4 describes the mapping between PCIe slots and CPUs, and the PCIe specifications of the CH121 V5.

NOTE

If socket CPU 2 is vacant, its corresponding PCIe devices do not work.

Table 1-4 PCIe devices

PCIe Device	CPU	PCIe Standard	Connector Bandwidth	Bus Width	Port	Device 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

PCIe Device	CPU	PCIe Standard	Connector Bandwidth	Bus Width	Port	Device Size
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 Riser 1	CPU 1	PCIe 3.0	x16	x16	Port 3A	Half-height half-length
NVMe SSD 0	CPU 2	PCIe 3.0	x4	x4	Port 1C	2.5-inch disk
NVMe SSD 2	CPU 2	PCIe 3.0	x4	x4	Port 1D	2.5-inch disk

Mainboard Layout

Figure 1-6 Positions of the connectors and other components

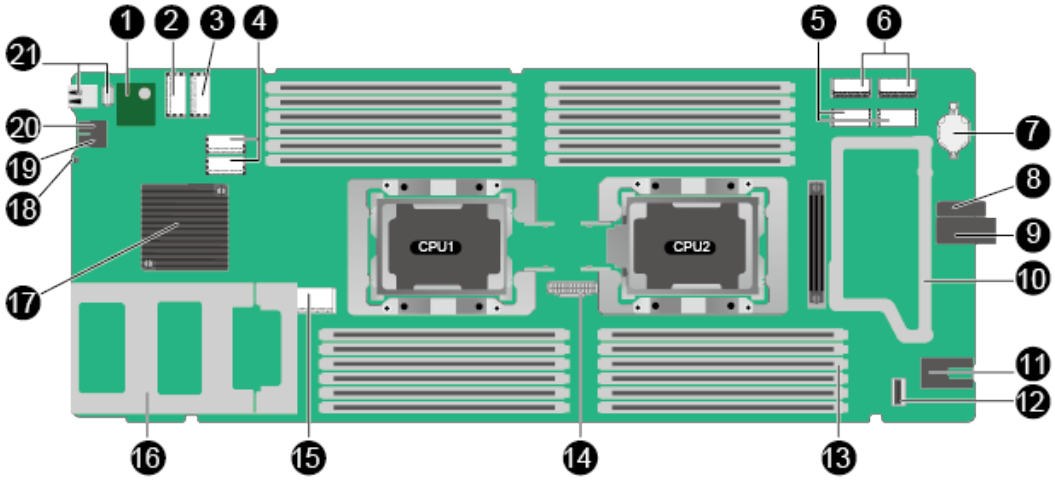

1	TPM card connector	2	PCIe riser card connector (PCIe 3.0 x16)
3	PCIe riser card connector (x8)	4	RAID controller card connector
5	Mezzanine card connector 2	6	Mezzanine card connector 1
7	BIOS battery	8	Positioning pin
9	Midplane signal connector	10	Mezzanine card tray
11	Power connector	12	SoftRAID key connector
13	DIMM slot	14	CPU2 OPA sideband signal interface
15	HDD backplane connector	16	HDD tray
17	PCH	18	HLY indicator

19	UID button/indicator	20	Power button/indicator
21	USB 3.0 port	-	-

1.6 Logical Structure

Figure 1-7 Logical structure of the CH121 V5

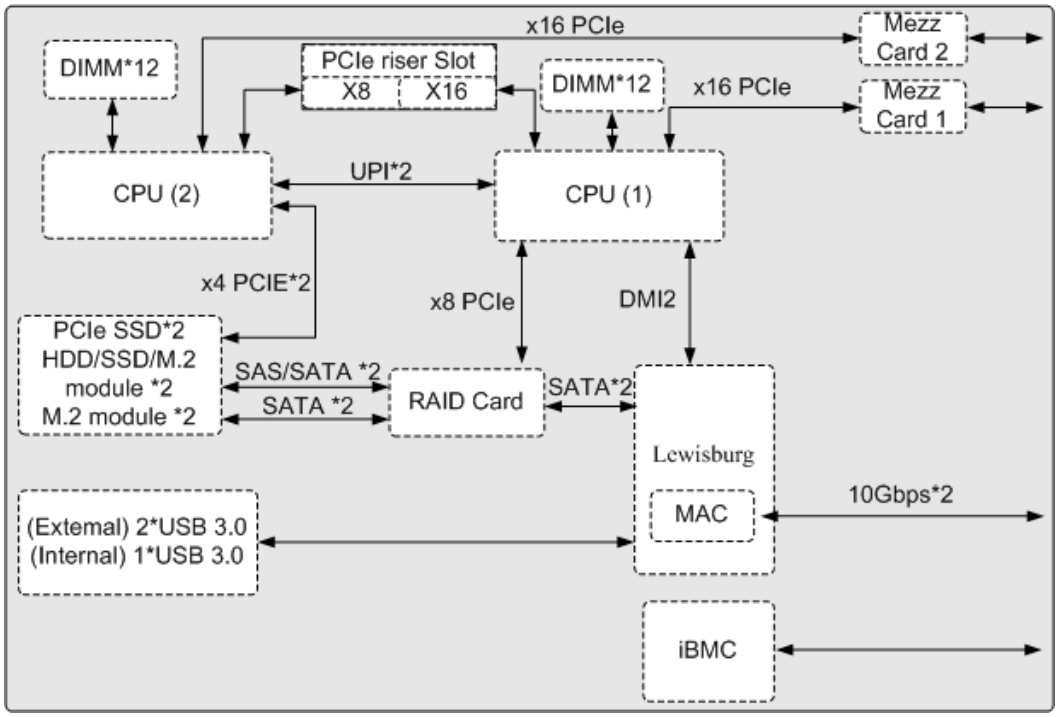

Intel® Xeon® Scalable new-generation CPUs are adopted. Each CPU supports 12 DIMMs. The CPUs are interconnected through UPI links at a speed of up to 10.4 GT/s. CPU 1 connects to the Lewisburg chip through the DMI2 link at a speed of 8 GT/s. Through PCIe links, both CPUs connect to mezzanine cards that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

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

The hard disk interface module consists of a RAID card and a hard disk backplane ("2 x HDD/SSD/M.2 module" in the preceding figure). The hard disk interface module connects to the CPUs through PCIe.

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
Physical specifications	Dimensions (H x W x D)	60.46 mm x 210 mm x 537.2 mm (2.4 in. x 8.26 in. x 21.14 in.)
	Color	● Front panel: black ● Cover: silver
	Weight	● Net: 6.5 kg (14.33 lb) ● Packing materials: 2.3 kg (5.07 lb)
Environmental specifications	Temperature	● Operating temperature: 5°C to 40°C (41°F to 104°F) (ASHRAE Class A3 compliant) ● Storage temperature: -40°C to +65°C (-40°F to +149°F) ● Temperature change rate: < 20°C/h (36°F/h) **NOTE** If Intel® Xeon® Platinum 8180, Platinum 8168, or Gold 6154 CPUs are used, the maximum operating temperature is 35°C (95°F).
	Humidity	● Operating humidity: 5% RH to 85% RH (non-condensing) ● Storage humidity: 5% RH to 95% RH (non-condensing) ● Humidity change rate: < 20% RH/h
	Altitude	At an altitude of 900 m (2952.72 ft), the highest operating temperature is 40°C (104°F). When the compute node is used at an altitude of 900 m to 5000 m (2952.8 ft to 16404.2 ft), the highest operating temperature decreases by 1°C (1.8°F) for every increase of 300 m (984.24 ft). HDDs cannot be configured if the altitude is higher than 3000 m (9842.5 ft).
	Corrosive air pollutant	● Corrosion rate of the copper test piece: < 300 Å/month (in compliance with the ANSI/ISA-71.04-2013 gaseous corrosion level G1) ● Corrosion rate of the silver test piece: < 200 Å/month

Category	Item	Specifications
	Particulate pollutant	● The ISO14664-1 Class 8 requirements are met. You are advised to ask a professional organization to monitor particulate pollutants in the equipment room. ● There is no explosive, conductive, magnetic, or corrosive dust in the equipment room.
Input power supply	Rated input voltage	12 V DC
Power consumption	Maximum power consumption	730 W

Table 1-6 Configurations used for testing the maximum power consumption

Compute Node	CPU	DIMM	Hard Disk	RAID Controller Card	NIC	PCIe Card
CH121 V5	Two 8168 CPUs	24 x 64 G PC4 2666 MT/s	2 x SATA SSDs	LSI SAS3008	● MEZZ 1: MZ710 ● MEZZ 2: MZ620	ES3000 V3

2 Features

Performance and Scalability

- Each Intel® Xeon® Scalable CPU (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 provides the intelligent adaptation function to enable the CPU cores to run at the maximum speed during peak workloads by temporarily going beyond the CPU thermal design power (TDP).
- 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 to allow OS vendors to better use hardware for addressing virtualization workloads.
- Intel® advanced vector extensions (AVX) improves floating-point computing performance for compute-intensive applications.
- A total of 24 load-reduced DIMMs (LRDIMMs) provides quick speed, high availability, and a maximum memory capacity of 3 TB.
- Two Intel® Xeon® Scalable CPUs provide a theoretical maximum memory bandwidth of 256 GB/s (64 bits/8 x 2666 MHz x 6 channels x 2), 66.7% higher than the previous generation.
- The I/O performance of a system configured with only solid-state drives (SSDs) is much higher than that of a system configured with only HDDs or both SSDs and HDDs. An SSD supports up to 100 times I/O operations per second (IOPS) of a typical HDD.
- The compute node supports 96-lane PCIe 3.0 (8 GT/s per lane), which improves the maximum bandwidth by 20% over the previous 80-lane PCIe generation.
- Intel® integrated I/O technology enables the PCIe 3.0 controller to be integrated into the Intel® Xeon® Scalable CPUs, shortening I/O latency and enhancing overall system performance.
- A compute node supports multiple network ports and mezzanine cards to provide a variety of ports.
- A compute node supports standard PCIe SSDs.

- A LOM with two 10GE ports is supported.
- M.2 SSDs are supported. The supported RAID levels depend on the configured RAID controller card.

Availability and Serviceability

- A compute node provides the memory mirroring and memory backup functions to avoid system shutdown caused by uncorrectable memory errors.
- A compute node supports hot-swappable hard disks for configuring RAID properties to protect data and prolong system running time.
- The UID and HLY indicators on the panel and the key component status displayed on the iBMC WebUI help technical support personnel quickly locate faulty components. This simplifies maintenance, shortens troubleshooting time, and improves system availability.
- The compute node supports SSDs to provide higher reliability than HDDs and prolong system operating time.
- The integrated iBMC module continuously monitors system parameters, triggers alarms, and performs recovery actions to minimize system downtime.

Manageability and Security

- The iBMC module monitors the compute node operating status and provides remote management.
- An integrated industry-standard unified extensible firmware interface (UEFI) increases setting, configuring, and updating efficiencies, and simplifies error handling.
- The optional TPM 2.0 provides advanced encryption functions, such as digital signatures and remote authentication.
- The industry-standard advanced encryption standard–new instruction (AES NI) implements faster and stronger encryption.
- Intel® Execute Disable Bit (EDB) works with a supported OS to prevent certain types of malicious buffer overflow attacks.
- Intel® Trusted Execution technology enhances security by using hardware-based resistance against malicious software attacks, allowing applications to run in isolated mode to avoid any interference from the other applications running on the OS.

Energy Efficiency

- The Intel® Xeon® Scalable Platinum 8100 CPUs provide significantly better performance than the previous generation. The maximum CPU TDP supported by the compute node is increased by 60 W.
- Intel® intelligent power capability powers on and off a single CPU based on the site requirements to reduce power consumption.
- Low-voltage Intel® Xeon® CPUs consume less energy to satisfy demands of power and thermally constrained data centers and telecommunication environments.
- Low-voltage 1.2 V DDR4 registered DIMMs (RDIMMs) consume 20% to 30% less energy than 1.35 V DDR3 RDIMMs.
- SSDs consume 80% less power than HDDs.
- A compute node uses hexagonal ventilation holes to enable higher ventilation density than round holes, remarkably increasing the system cooling efficiency.

- The efficient voltage regulator down (VRD) PSUs reduce the loss in the mainboard DC power conversion.
- Power capping and power control are supported.

3 Components

This topic describes the software and hardware supported by the CH121 V5.

For more details, see **Huawei Server Compatibility Checker**.

CPU

The mainboard supports one or two CPUs.

- Intel® Xeon® Scalable full-series CPUs are all supported (Bronze 3100, Silver 4100, Gold 5100/6100, and Platinum 8100), with up to 28 cores.
- Each CPU is integrated with memory controllers and supports six DDR4 memory channels. Each channel supports two DDR4 DIMMs at 2400 or 2666 MT/s frequency.
- Each CPU integrates a PCIe controller with 48 lanes and supports PCIe 3.0.
- Every two CPUs are interconnected through two UPI links, with each link delivering up to 10.4 GT/s.
- The maximum core frequency supported by the CPUs reaches 3.6 GHz.

Memory

Up to 24 DIMM slots for installing DIMMs (12 DIMMs for each CPU). At least one DIMM must be configured.

DIMM Configuration Rules

Observe the following rules to configure DIMMs:

1. The CH121 V5 supports the DIMMs with the capacity of 8 GB, 16 GB, 32 GB, 64 GB, or 128 GB. A compute node provides a maximum memory capacity of 3 TB when DIMMs are fully configured.
2. The maximum number of DIMMs supported by one compute node depends on the CPU type, DIMM type, and rank quantity. See "Maximum number of DIMMs" in **Table 3-1**.

   NOTE

   - CPU 1 must be configured with DIMMs. If DIMMs are configured only for CPU 2, the compute node cannot be powered on. **Figure 1-6** shows the positions of CPUs 1 and 2.
   - Note the following rule:

     Maximum number of DIMMs per channel ≤ Maximum number of ranks per channel/Number of ranks per DIMM

3. A compute node does not support mixed use of RDIMMs and LRDIMMs. Mixed use of RDIMMs or LRDIMMs of different capacities is supported during capacity expansion but may affect the DIMM RAS feature.
4. The speed of any DIMM is the smaller value of the following:
   - Memory speed supported by the connected CPU.
   - Lowest maximum operating speed of a specific memory configuration. See "Maximum operating speed" in **Table 3-1**.

Table 3-1 DIMM configuration rules for Intel® Xeon® Scalable CPUs

Parameter		DIMM		
Rank		Dual rank	Quad rank	Octal rank
Rated speed (MT/s)		2666	2666	2666
Rated voltage (V)		1.2	1.2	1.2
Operating voltage (V)		1.2	1.2	1.2
Maximum number of DIMMs[a]		24	24	24
Maximum capacity per DIMM (GB)		32	64	128
Maximum memory capacity[b] (GB)		768	1536	1536
Maximum memory capacity at the maximum operating speed (GB)		768	1536	1536
Maximum operating speed (MT/s)	One DIMM per channel	2666	2666	2400
	Two DIMMs per channel	2666	2666	2400
a: The maximum numbers of DIMMs are based on the two-CPU configuration. If only one CPU is installed, the maximum numbers of DIMMs are half the values given in this table. b: Octal-rank 128 GB DIMMs cannot be fully configured because a single CPU supports up to 768 GB memory when octal-rank 128 GB DIMMs are configured. 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

For details about DIMM slot configuration rules, see **Huawei Server Product Memory Configuration Assistant**.

Table 3-2 lists the DIMM slot installation sequence. For details about the DIMM slot numbers, see **Figure 1-6**.

Table 3-2 DIMM installation sequence

CPU	DIMM Slot Installation Sequence
CPU 1	1A1, 1B1, 1C1, 1D1, 1E1, 1F1, 1A2, 1B2, 1C2, 1D2, 1E2, 1F2
CPU 1 and CPU 2	1A1, 2A1, 1B1, 2B1, 1C1, 2C1, 1D1, 2D1, 1E1, 2E1, 1F1, 2F1, 1A2, 2A2, 1B2, 2B2, 1C2, 2C2, 1D2, 2D2, 1E2, 2E2, 1F2, 2F2

The CH121 V5 provides 24 DDR4 DIMM slots. Each CPU integrates six memory channels. Six memory channels for CPU 1 are 1A, 1B, 1C, 1D, 1E and 1F, and those for CPU 2 are 2A, 2B, 2C, 2D, 2E and 2F. **Table 3-3** lists the memory channels for each CPU.

Slots 1A1, 1B1, 1C1, 1D1, 1E1, 1F1, 2A1, 2B1, 2C1, 2D1, 2E1 and 2F1 are the primary slots of channels 1A, 1B, 1C, 1D, 1E. 1F, 2A, 2B, 2C, 2D, 2E and 2F respectively. When installing DIMMs in a channel, choose the primary slot first. If you do not install a DIMM in the primary slot, the DIMM in the standby slot cannot work.

Table 3-3 Memory channels for each CPU

Channel Location	Memory Channel	DIMM Slot
CPU 1	1A	DIMM000 (1A1)
		DIMM001 (1A2)
	1B	DIMM010 (1B1)
		DIMM011 (1B2)
	1C	DIMM020 (1C1)
		DIMM021 (1C2)
	1D	DIMM030 (1D1)
		DIMM031 (1D2)
	1E	DIMM040 (1E1)
		DIMM041 (1E2)
	1F	DIMM050 (1F1)
		DIMM051 (1F2)
CPU 2	2A	DIMM100 (2A1)
		DIMM101 (2A2)
	2B	DIMM110 (2B1)
		DIMM111 (2B2)
	2C	DIMM120 (2C1)
		DIMM121 (2C2)

Channel Location	Memory Channel	DIMM Slot
	2D	DIMM130 (2D1)
		DIMM131 (2D2)
	2E	DIMM140 (2E1)
		DIMM141 (2E2)
	2F	DIMM150 (2F1)
		DIMM151 (2F2)

Storage

The CH121 V5 supports two 2.5-inch HDDs or SSDs and allows mixed configuration of an HDD and an SSD. Each HDD or SSD is hot-swappable and can be independently installed and removed.

NOTE

After the OS is installed on a hard disk, do not move the hard disk to another compute node; otherwise, mounting a virtual flash drive or CD/DVD-ROM on the KVM screen may fail.

The CH121 V5 supports the LSI SAS3008 RAID controller card, which supports RAID 0 and 1.

Table 3-4 lists the performance, minimum number of disks required, and disk utilization of different RAID levels.

Table 3-4 RAID level comparison

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

I/O Expansion

The CH121 V5 supports the following types of PCIe mezzanine cards for connecting to switch modules through the midplane. You can choose a mezzanine card based on the I/O card type and rate requirements.

- GE expansion card
- 10GE expansion card
- 40GE expansion card
- FC or FCoE expansion card

- RoCE expansion card
- IB expansion card

I/O Modules Supported by the LOM

The LOM can connect to I/O modules (switch modules). The following figure shows the connections between the LOM on a half-width compute node and the I/O modules.

Figure 3-1 Connections between the LOM on a half-width compute node and the I/O modules

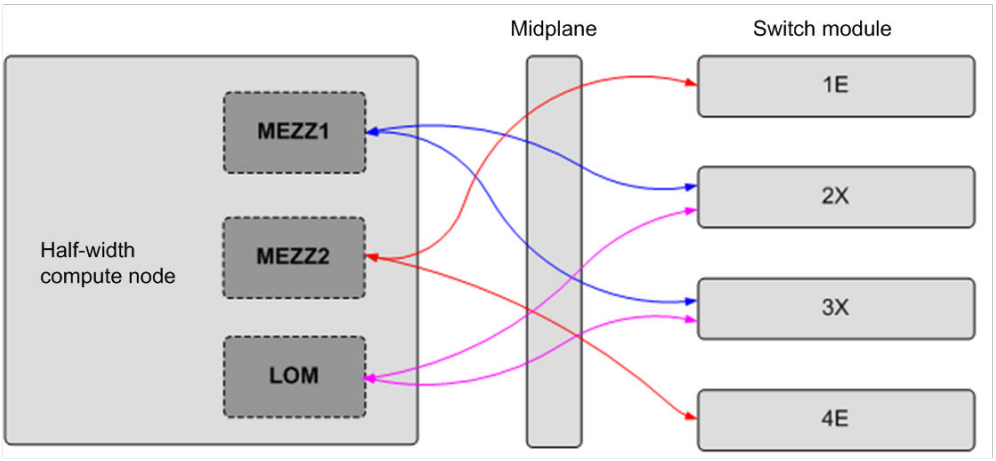

Mezzanine cards on the compute node connect to switch modules. Mezz1 connects to Fabric ports of switch module slots 2X and 3X; Mezz2 connects to Fabric ports of switch module slots 1E and 4E.

The LOM is an NIC integrated in the PCH and provides two 10GE ports to connect to the Base ports of slots 2X and 3X.

Table 3-5 I/O modules supported by the LOM

I/O Module	I/O Slot	LOM	Remarks
CX916	2X/3X	√	N/A
	1E/4E	×	The LOM cannot communicate with I/O modules in slots 1E and 4E.
CX920	2X/3X	√	N/A
	1E/4E	×	The LOM cannot communicate with I/O modules in slots 1E and 4E.

NOTE

Forcibly powering off the compute node will cause the WOL function of the LOM ports to be unavailable.

Power Supply

The CH121 V5 is powered by the power supply units (PSUs) in the E9000 chassis, without any independent power supply.

Peripherals

The CH121 V5 supports peripherals such as a USB DVD-ROM drive.

OSs and Software

For details about the OSs and virtualization software supported by the CH121 V5, see **Huawei Server Compatibility Checker**.

4 Management

Huawei iBMC, a remote management system for servers, is integrated on the E9000 compute nodes. It complies with the IPMI V2.0 standards and provides reliable hardware monitoring and management functions. Huawei iBMC implements seamless communications with chassis management modules. 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

Item	Description
Management interface	iBMC supports various management interfaces to implement system integration. iBMC can be integrated with any standard management systems through the following interfaces: ● IPMI V2.0 ● CLI ● HTTPS ● Redfish ● SNMP
Fault detection	iBMC helps to detect faults and accurately locate hardware faults.
System watchdog	iBMC supports BIOS POST, OS watchdog, and fault timeout automatic system reset. You can enable or disable these functions in the iBMC.
Boot device configuration	iBMC supports out-of-band configuration for boot devices.

Item	Description
Alarm management	iBMC supports alarm management and reports alarms in various ways such as the (SMTP), and syslog service to ensure that the compute node runs properly without interruption.
Integrated KVM	iBMC provides remote maintenance measures, such as KVM and KVM over IP, for troubleshooting. The maximum resolution is 1600 x 1200.
Integrated virtual media	iBMC virtualizes local media devices or images for remote compute nodes to facilitate OS installation. The virtual DVD-ROM drive supports a transmission rate of up to 8 MB/s.
WebUI	iBMC provides a visual WebUI for quick configuration and information queries. The following web browsers are supported: ● Internet Explorer 8.0 ● Firefox 9.0 ● Chrome 13.0 ● Safari
Fault reproduction	iBMC reproduces faults to diagnose the faults quickly.
Screenshots and videos	iBMC allows you to view screenshots and videos without login, which facilitates preventive maintenance inspection (PMI).
Black Box	Allows you to enable or disable the black box function and download black box data.
DNS/LDAP	iBMC supports domain management and directory services, which significantly simplifies network and configuration management.
Dual-image backup	If iBMC software fails, it starts again from a backup image.
Asset management	iBMC provides intelligent asset management to facilitate asset management.
Intelligent power management	iBMC supports power capping to increase deployment density and uses dynamic energy saving technology to lower the operating expense (OPEX).

5 Warranty

According to the *Huawei Warranty Policy for Servers & Storage Products* (*Warranty Policy* for short), Huawei provides a three-year warranty for servers, a one-year warranty for DVD-ROM drives and iBBUs, 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 service and/or parts may not be available in all countries. For more information about warranty services in your country, contact Huawei technical support or your local Huawei representative office.

6 Certifications

This topic 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: ● EN 55022: 2010 ● CISPR 22: 2008 ● EN 55024: 2010 ● CISPR 24: 2010 ● ETSI EN 300 386 V1.6.1: 2012 ● ETSI ES 201 468 V1.3.1: 2005
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