

Huawei FusionServer 2288H V5

White Paper

Issue 01

Date 2017-08-07

Copyright © Huawei Technologies Co., Ltd. 2017. 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>

Contents

1 Overview.....	1
2 Features.....	2
3 Logical Structure.....	5
4 Hardware Description.....	7
4.1 Appearance.....	8
4.2 Ports.....	10
4.3 Indicators and Buttons.....	13
4.4 Riser Cards and PCIe Slots.....	20
4.5 Physical Structure.....	24
5 Product Specifications.....	26
5.1 Technical Specifications.....	27
5.2 Physical and Environmental Specifications.....	30
6 Component Compatibility.....	33
6.1 CPU.....	34
6.2 Memory.....	34
6.3 Storage.....	38
6.4 IO Expansion.....	40
6.5 PSU.....	40
6.6 OS and Software Support.....	41
7 System Management.....	42
8 Warranty.....	44
9 Certifications.....	47
10 References.....	48

1 Overview

The Huawei FusionServer 2288H V5 is a 1U 2-socket rack server developed for Internet, Internet data center (IDC), cloud computing, enterprise, and telecom service applications.

Marked H22H-05 on the nameplate, the 2288H V5 is ideal for IT core services, cloud computing virtualization, high-performance computing, distributed storage, big data processing, enterprise or telecom service applications, and other complex workloads. It combines low power consumption with high scalability and reliability, and easy deployment and management.

Figure 1-1 shows the appearance of a server with 12 hard disks.

Figure 1-1 Appearance



2 Features

Performance and Scalability

The 2288H V5 offers the following features to boost performance and improve scalability:

- Intel® Xeon® Scalable series processors ensure high processing performance by providing up to 28 cores, 3.6 GHz frequency, 38.5 MB L3 cache, and two 10.4 GT/s Ultra Path Interconnect (UPI) links between processors.
 - Each 2288H V5 supports two processors with 56 cores and 112 threads to maximize the concurrent execution of multithreaded applications.
 - Intel® Xeon® Scalable series processors support L2 cache. Each core can exclusively use a maximum of 1 MB L2 cache or 1.375 MB L3 cache.
 - Intel® Turbo Boost Technology 2.0 enables processor cores to run at maximum speeds during peak hours by temporarily going beyond the processor thermal design power (TDP).
 - Intel® Hyper-Threading Technology enables each processor core to run up to two threads, improving parallel computation capability.
 - Intel® Virtualization Technology integrates hardware-level virtualization functions to allow OS vendors to better use hardware to address virtualization workloads.
- Twenty-four DDR4 error checking and correcting (ECC) RDIMMs or load-reduced DIMMs (LRDIMMs) provide a maximum memory speed of 2666 MT/s and a maximum memory capacity of 1,536 GB, featuring high speed and availability. The maximum memory bandwidth is 249.9375 GB/s in theory.
- 12 Gbit/s internal SCSI (SAS) storage connection doubles the data transmission rate compared with the 6 Gbit/s SAS storage connection, maximizing the performance of I/O-intensive applications.
- The 2288H V5 supports flexible hard disk configurations and provides elastic and scalable memory capacities to satisfy storage capacity and upgrade requirements.
- Intel® Advanced Vector Extensions 2.0 (AVX 2.0 and AVX-512) improve floating-point computing performance for computing-intensive applications.
- The I/O performance of pure SSDs is higher than that of mixed configuration of SSDs and HDDs and 100 times that of pure HDDs.
- The 2288H V5 supports various integrated network interface cards (NICs) and flexible NICs to provide rich network ports.

- The 2288H V5 supports up to eight PCIe 3.0 slots.
- The Intel® Xeon® Scalable series processors incorporate the PCIe 3.0 controller using the Intel Integrated I/O. This remarkably shortens I/O latency and enhances overall system performance.

Availability and Serviceability

The 2288H V5 provides the following features to improve availability and serviceability:

- The 2288H V5 uses carrier-class components and follows the engineering process, which dramatically improves system reliability.
- The 2288H V5 uses hot-swappable SATA and SAS HDDs. It supports redundant array of independent disks (RAID) 0, 1, 1E, 10, 5, 50, 6, and 60 and offers RAID cache. A supercapacitor is used to protect RAID cache data from power failures.
- SSDs offer better reliability than HDDs, prolonging system uptime.
- The web user interface (WebUI) of the Intelligent Baseboard Management Controller (iBMC) and the UID and HLY indicators on the panel help technical support personnel promptly obtain the status of key components and locate failed (or failing) components. This simplifies servicing, accelerates troubleshooting, and helps improve system availability.
- The iBMC monitors system parameters in real time, triggers alarms, and performs recovery actions in case of failures, minimizing system downtime.
- Huawei provides a three-year warranty for parts replacement and onsite repair for the servers used in China. Huawei provides a 10-hour-a-day, 5-day-a-week support program. Service requests will be handled the next business day. Optional service upgrades are available.
- Huawei provides a three-year warranty for parts replacement and repair for the servers used outside China. Huawei provides a 9-hour-a-day, 5-day-a-week support program. Service requests will be handled the next business day. Huawei delivers the repaired or new parts within 45 calendar days after receiving the defective parts.

Manageability and Security

The 2288H V5 provides the following features to enhance manageability and security:

- The built-in iBMC module monitors server operating status and provides remote management.
- The Network Controller Sideband Interface (NC-SI) feature enables a network port to function as a management network port and a service port, which maximizes the return on investment (ROI) for customers. NC-SI is disabled by default. You can enable it on the iBMC WebUI or in the BIOS.
- The integrated industry-standard Unified Extensible Firmware Interface (UEFI) increases setup, configuration, and update efficiency, and simplifies fault handling.
- The front bezel in the server chassis is locked to ensure local data security and reliability.
- The Intel® Advanced Encryption Standard–New Instructions (AES NI) support faster and stronger encryption.
- The Intel® Execute Disable Bit (EDB) function prevents certain types of malicious buffer overflow attacks when working with a supported OS.

- The Intel® Trusted Execution technology provides enhanced security by using hardware-based defense against malicious software attacks, allowing an application to run in an isolated space from all other applications running on the OS.

 **NOTE**

The service network port supporting NC-SI has the following features:

- The service network port can be bound to a network port (host network port 1 by default) on a flexible or standard NIC.
- The service network port allows you to enable, disable, and configure a VLAN ID. A VLAN ID is disabled by default, and the default VLAN ID is 0.
- The service network port supports IPv4 and IPv6 addresses. You can set an IP address, subnet mask, default gateway, and IPv6 address prefix length for the service network port.

Energy Efficiency

The 2288H V5 offers the following features to save energy:

- The 2288H V5 supports 80 Plus Platinum power supply units (PSUs). The PSUs provide 94% power efficiency at 50% loads.
- The 2288H V5 supports active-standby power supplies and high-voltage DC (HVDC), improving power supply efficiency.
- Efficient VRD PSUs reduce the loss in DC/DC power conversion.
- The 2288H V5 supports area-based and intelligent fan speed adjustment, Proportional-Integral-Derivative (PID) speed adjustment, and intelligent processor frequency adjustment, reducing power consumption.
- The improved thermal design with energy-efficient fans ensures optimal heat dissipation and reduces system power consumption.
- The 2288H V5 supports power capping and power control.
- Hard disks are not powered on simultaneously, which reduces the server startup power consumption.
- The Intel® Intelligent Power Capability allows a single processor to be powered on or off based on site requirements.
- Low-voltage Intel® Xeon® Scalable series processors consume less energy and apply to data centers and telecommunication environments that have power and thermal limitations.
- SSDs consume 80% less power than HDDs.

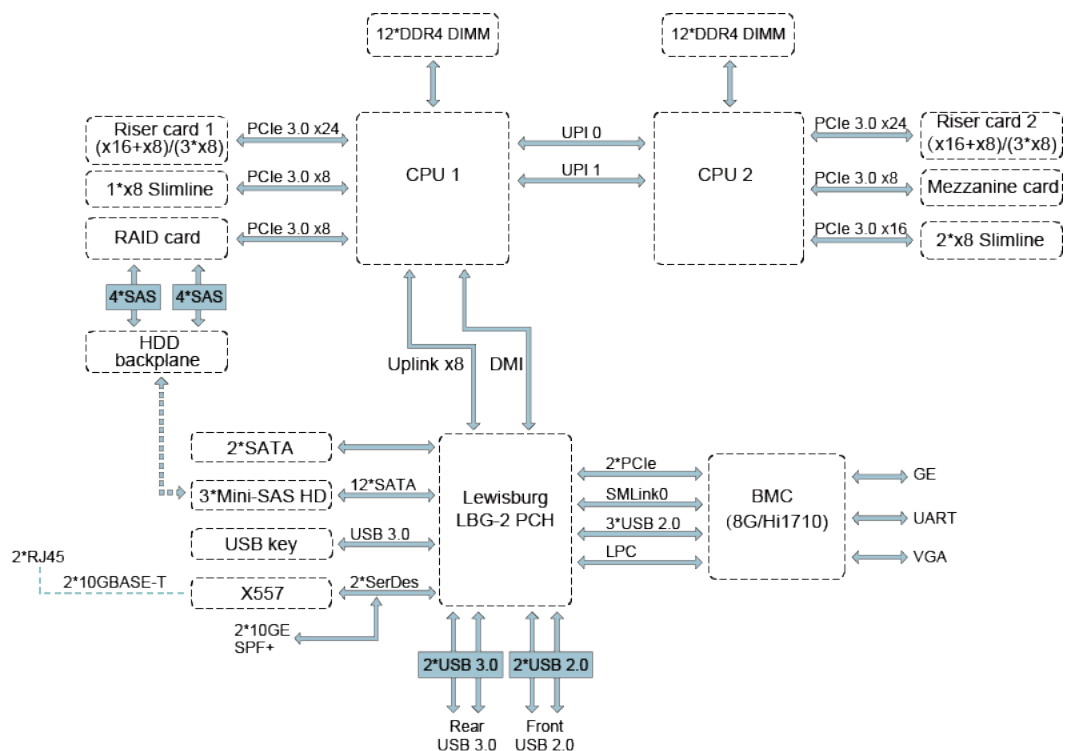
Support for Customization

The 2288 V5 is a Huawei proprietary server. Huawei also provides customized development in a timely manner.

3 Logical Structure

Figure 3-1 shows the logical structure of the 2288H V5.

Figure 3-1 Logical structure of the 2288H V5



- The 2288H V5 supports two Intel® Xeon® Scalable series CPUs and 24 DDR4 DIMMs. The CPUs interconnect with each other through two Ultra Path Interconnect (UPI) buses at speeds of up to 10.4 GT/s.
- The CPUs connect to three PCIe riser cards through PCIe buses and the riser cards provide various PCIe slots.
- The RAID controller card on the mainboard connects to CPU 1 through PCIe buses, and to the hard disk backplanes through SAS signal cables. The hard disk backplanes support various local storage configurations.
- The LBG-2 bridge chip provides two GE electrical ports, and two 10GE optical ports or two 10GE electrical ports (through the X557 PHY).

- The Huawei Hi1710 management chip provides a VGA port, management network port, and debugging serial port.

4 Hardware Description

[4.1 Appearance](#)

[4.2 Ports](#)

[4.3 Indicators and Buttons](#)

[4.4 Riser Cards and PCIe Slots](#)

[4.5 Physical Structure](#)

4.1 Appearance

Front Panel

- **Figure 4-1** shows the front panel of a server with eight hard disks.

Figure 4-1 Front panel of a server with eight hard disks




- | | |
|-------------------------|--|
| 1 Hard disk | 2 Built-in DVD drive or LCD (optional) |
| 3 Label (including ESN) | — — |

For details about how to use the LCD, see the *2288H V5 Server LCD V100R005 User Guide*.

- **Figure 4-2** shows the front panel of a server with 12 hard disks.

Figure 4-2 Front panel of a server with 12 hard disks



- | | |
|-------------|-------------------------|
| 1 Hard disk | 2 Label (including ESN) |
|-------------|-------------------------|

- **Figure 4-3** shows the front panel of a server with 24 hard disks.


Figure 4-3 Front panel of a server with 24 hard disks



- | | |
|-------------|-------------------------|
| 1 Hard disk | 2 Label (including ESN) |
|-------------|-------------------------|

- **Figure 4-4** shows the front panel of a server with 25 hard disks.

Figure 4-4 Front panel of a server with 25 hard disks



ESN

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

Figure 4-5 shows the ESN format by using the example 2102310QPD10F3001263.

Figure 4-5 ESN example





Table 4-1 ESN description

No.	Description
1	ESN ID
2	Material identification code (eight characters)
3	Vendor ID
4	Year and month
5	Serial number

Rear Panel

Figure 4-6 shows the rear panel of a server.

Figure 4-6 Rear panel



NOTE

A hard disk module or riser module can be installed on I/O modules 1 and 2, and only a riser module can be installed on I/O module 3. The preceding figure is for reference only.

4.2 Ports

Front Panel

- **Figure 4-7** shows the ports on the front panel of a server with eight hard disks.

Figure 4-7 Front panel of a server with eight hard disks



- | | | | |
|---|---------------|---|--------------|
| 1 | USB 2.0 ports | 2 | USB 3.0 port |
| 3 | VGA port | — | — |


- **Figure 4-8** shows the ports on the front panel of a server with 12 hard disks.

Figure 4-8 Front panel of a server with 12 hard disks



- **Figure 4-9** shows the ports on the front panel of a server with 24 hard disks.

Figure 4-9 Front panel of a server with 24 hard disks



- **Figure 4-10** shows the ports on the front panel of a server with 25 hard disks.

Figure 4-10 Front panel of a server with 25 hard disks



Table 4-2 describes the ports on the front panel.


Table 4-2 Ports on the front panel

Port	Type	Description
VGA port	DB15	The VGA port is used to connect a terminal, such as a monitor or KVM.
USB port	USB2.0 and USB3.0	The USB ports allow USB devices to be connected to the server. NOTE Before connecting an external USB device, check that the USB device functions properly. A server may operate abnormally if an abnormal USB device is connected.

Rear Panel

Figure 4-11 shows the ports on the rear panel of a server.

Figure 4-11 Rear panel



- | | | | |
|---|-------------------------|---|---------------------|
| 1 | 10GE electrical ports | 2 | 10GE optical ports |
| 3 | VGA port | 4 | GE electrical ports |
| 5 | Management network port | 6 | Serial port |
| 7 | USB 3.0 ports | 8 | PSU sockets |

Table 4-3 describes the ports on the rear panel.

Table 4-3 Ports on the rear panel

Port	Type	Quantity	Description
10GE electrical port	10GBASE-T	2	The mainboard provides two 10GE electrical ports or two GE optical ports for you to select. NOTE
10GE optical port	10GE SFP+	2	<ul style="list-style-type: none"> 10GE electrical ports and 10GE optical ports do not support 10 Mbit/s or 100 Mbit/s networks. The rate of a 10GE electrical port cannot be forcibly set to 1000 Mbit/s.
VGA port	DB15	1	The VGA port is used to connect a terminal, such as a monitor or KVM.
GE electrical port	1000BASE-T	2	Server service network port NOTE This port does not support forcible rates or 10 Mbit/s and 100 Mbit/s networks.

Port	Type	Quantity	Description
Serial port	RJ45	1	The serial port is used as the system serial port by default. You can set it as the iBMC serial port by using the iBMC command. This port is used for debugging.
Management network port	1000BASE-T	1	The 1000 Mbit/s Ethernet port is used to manage the server.
USB port	USB 3.0	2	The USB ports allow USB devices to be connected to the server. NOTICE Before connecting an external USB device, check that the USB device functions properly. A server may operate abnormally if an abnormal USB device is connected.
PSU socket	—	1 or 2	Determine the number of PSUs based on actual requirements, but ensure that the rated power of the PSUs is greater than that of the server. When one PSU is used, Predicted PSU Status cannot be set to Active/Standby .

4.3 Indicators and Buttons

Front Panel

- **Figure 4-12** shows the indicators and buttons on the front panel of a server with eight hard disks.

Figure 4-12 Front panel of a server with eight hard disks



- 1 NMI button
- 2 Fault diagnosis LED
- 3 Health indicator
- 4 UID button/indicator
- 5 Power button/indicator
- 6 Network port connection status indicators (1 to 4 from top to down)

- **Figure 4-13** shows the indicators and buttons on the front panel of a server with 12 hard disks.

Figure 4-13 Front panel of a server with 12 hard disks



- | | |
|---|------------------------|
| 1 Network port connection status indicators (1 to 4 from top to down) | 2 Fault diagnosis LED |
| 3 Health indicator | 4 UID button/indicator |
| 5 Power button/indicator | — — |

- **Figure 4-14** shows the indicators and buttons on the front panel of a server with 24 hard disks.

Figure 4-14 Front panel of a server with 24 hard disks



- | | |
|---|------------------------|
| 1 Network port connection status indicators (1 to 4 from top to down) | 2 Fault diagnosis LED |
| 3 Health indicator | 4 UID button/indicator |
| 5 Power button/indicator | 6 NMI button |

- **Figure 4-15** shows the indicators and buttons on the front panel of a server with 25 hard disks.





Figure 4-15 Front panel of a server with 25 hard disks




- | | |
|---|------------------------|
| 1 Network port connection status indicators (1 to 4 from top to down) | 2 Fault diagnosis LED |
| 3 Health indicator | 4 UID button/indicator |
| 5 Power button/indicator | — — |

Table 4-4 describes the indicators and buttons on the front panel of a server.

Table 4-4 Indicators and buttons on the front panel

Silk Screen	Indicator and Button	State Description
	Fault diagnosis LED	<ul style="list-style-type: none"> ● ---: The server is operating normally. ● Fault code: A server component is faulty. For details about fault code, see the HUAWEI Rack Server Alarm Handling (iBMC).
	Power button/indicator	<p>Power indicator</p> <ul style="list-style-type: none"> ● Steady yellow: The server is ready to power on. ● Steady green: The server is properly powered on. ● Blinking yellow: The iBMC is starting. ● Off: The server is not connected to a power source. <p>Power button</p> <ul style="list-style-type: none"> ● When the server is powered on, you can press this button to shut down the OS. <p>NOTE For RHEL 7.0, CentOS 7.0, and their later versions, set the response action of the power button to shutdown.</p> <ul style="list-style-type: none"> ● When the server is powered on, holding down this button for 6 seconds will power off the server. ● When the server is ready to be powered on, you can press this button to start the server.
	UID button/indicator	<p>The UID button/indicator helps identify and locate a server. You can turn on or off the UID indicator by pressing the UID button or remotely running a command on the iBMC CLI.</p> <p>UID indicator</p> <ul style="list-style-type: none"> ● Steady blue/Blinking blue: The server is being located. ● Off: The server is not being located. <p>UID button</p> <ul style="list-style-type: none"> ● You can press this button to turn on or off the UID indicator. ● You can press and hold down this button for 4 to 6 seconds to reset the iBMC.
	Health indicator	<ul style="list-style-type: none"> ● Steady green: The server is operating properly. ● Blinking red at 1 Hz: A major alarm has been generated on the server. ● Blinking red at 5 Hz: A critical alarm has been generated on the server.

Silk Screen	Indicator and Button	State Description
○	NMI button	<p>The NMI button triggers a server to generate a non-maskable interrupt. You can press this button or control it remotely through the iBMC WebUI.</p> <p>NOTICE</p> <ul style="list-style-type: none"> ● Press the NMI button only when the OS is abnormal. Do not press this button when the server is operating properly. ● Press the NMI button only for internal commissioning. Before pressing this button, ensure that the OS has the handler for NMI interrupts. Otherwise, the OS may crash. Exercise caution when pressing this button.
	Network port connection status indicator	<p>Each indicator shows the status of an Ethernet port on the network interface card (NIC).</p> <ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Off: The network port is not in use or has failed. <p>NOTE The indicators correspond to two 10GE and two GE network ports on the mainboard.</p>

Rear Panel

Figure 4-16 shows the indicators on the rear panel of a server.

Figure 4-16 Indicators on the rear panel

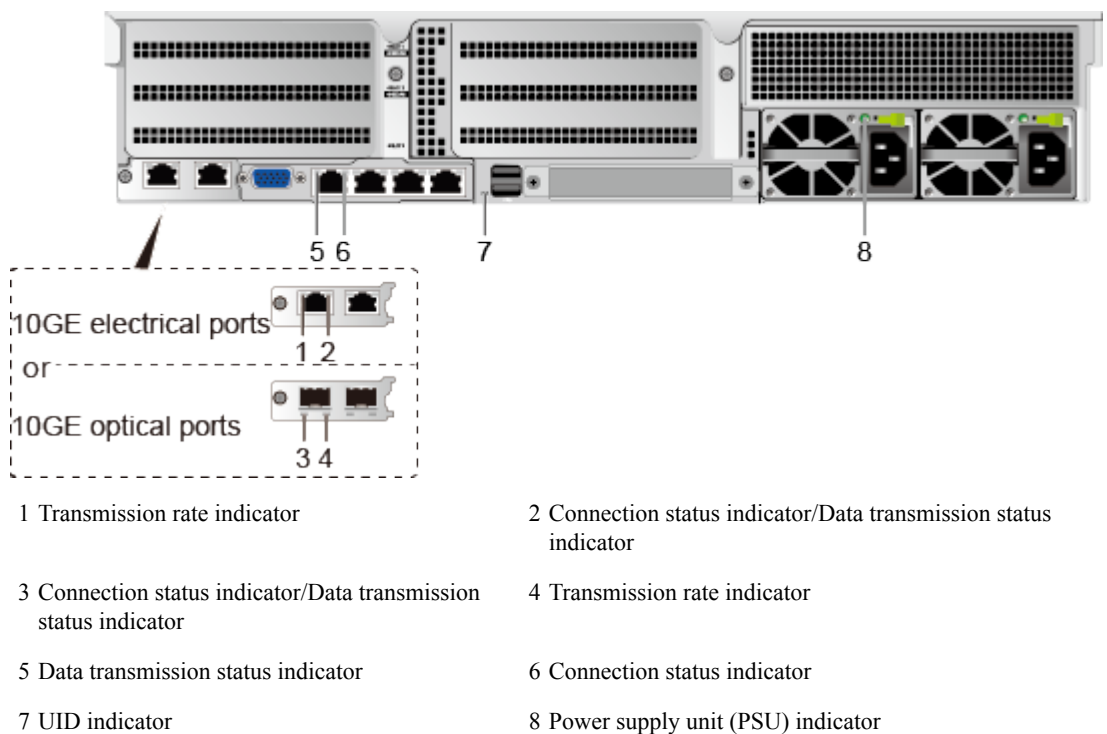


Table 4-5 describes the indicators on the rear panel of a server.

Table 4-5 Indicators on the rear panel

Indicator		State Description
10GE electrical port	Transmission rate indicator	<ul style="list-style-type: none"> ● Steady green: The data transmission rate is 10 Gbit/s. ● Steady yellow: The data transmission rate is 1 Gbit/s. ● Off: The network port is not connected.
	Connection status indicator/Data transmission status indicator	<ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Blinking green: Data is being transmitted. ● Off: The network port is not connected.
10GE optical port	Connection status indicator/Data transmission status indicator	<ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Blinking green: Data is being transmitted. ● Off: The network port is not connected.
	Transmission rate indicator	<ul style="list-style-type: none"> ● Steady green: The data transmission rate is 10 Gbit/s. ● Steady yellow: The data transmission rate is lower than 10 Gbit/s. ● Off: The network port is not connected.
GE electrical port	Connection status indicator	<ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Off: The network port is not connected.
	Data transmission status indicator	<ul style="list-style-type: none"> ● Blinking yellow: Data is being transmitted. ● Off: No data is being transmitted.
UID indicator		<p>The UID indicator helps identify and locate a server. You can turn on or off the UID indicator by pressing the UID button or remotely running a command on the iBMC CLI.</p> <ul style="list-style-type: none"> ● Steady blue/Blinking blue: The server is being located. ● Off: The server is not being located.

Indicator	State Description
PSU indicator	<ul style="list-style-type: none"> ● Steady green: The power input and output are normal. ● Blinking green at 1 Hz: The power input is normal, but the power output is stopped due to power-on or installation detection. An input overvoltage or undervoltage fault occurs. ● Blinking green at 4 Hz: under online upgrade. ● Steady orange: The input is normal, but no power output is supplied due to overheat protection, overcurrent protection, short circuit protection, output overvoltage protection, or some component failures. ● Off: No AC power is supplied.

SAS/SATA Hard Disk Indicators

Figure 4-17 shows the SAS/SATA hard disk indicators.

Figure 4-17 SAS/SATA hard disk indicators

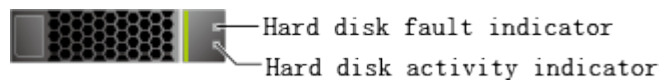


Table 4-6 describes the SAS/SATA hard disk indicators.

Table 4-6 SAS/SATA hard disk indicators

Indicator	State Description
Hard disk fault indicator	<ul style="list-style-type: none"> ● Steady yellow: The hard disk is not detected, the hard disk is faulty, or the RAID array status of the hard disk is abnormal. ● Blinking yellow: The server is locating the hard disk or rebuilding RAID. ● Off: The hard disk is operating normally or is not detected in the RAID array. <p>NOTE If the fault indicator is steady yellow, run a command to check the RAID status to determine whether the RAID array status is abnormal or whether the hard disk is faulty. For details about command description, see the HUAWEI V5 Server RAID Controller Card User Guide.</p>

Indicator	State Description
Hard disk activity indicator	<ul style="list-style-type: none"> ● Steady green: The hard disk is inactive. ● Blinking green: Data is being read from or written to the hard disk, or synchronized between hard disks. ● Off: The hard disk is faulty or not detected.

Flexible NIC Indicators


Supported flexible NICs include:

- SM210 with four GE electrical ports
- SM233 with two 10GE electrical ports

Use the [Huawei Server Compatibility Checker](#) to obtain the latest model list.

- **Figure 4-18** shows the indicators on SM210 with four GE electrical ports.

Figure 4-18 SM210 with four GE electrical ports



- **Figure 4-19** shows the indicators on SM233 with two 10GE electrical ports.

Figure 4-19 SM233 with two 10GE electrical ports




Table 4-7 Flexible NIC indicators


NIC Type	Indicator	State Description
Flexible NIC with four GE electrical ports	Data transmission status indicator	<ul style="list-style-type: none"> ● Blinking yellow: Data is being transmitted. ● Off: No data is being transmitted.

NIC Type	Indicator	State Description
	Connection status indicator	<ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Off: The network port is not connected.
Flexible NIC with two 10GE electrical ports	Transmission rate indicator	<ul style="list-style-type: none"> ● Steady green: The data transmission rate is 10 Gbit/s. ● Steady yellow: The data transmission rate is 1 Gbit/s. ● Off: The data transmission rate is 10/100 Mbit/s.
	Connection status indicator/Data transmission status indicator	<ul style="list-style-type: none"> ● Steady green: The network port is properly connected. ● Blinking green: Data is being transmitted. ● Off: No data is being transmitted or the network port is not connected.

4.4 Riser Cards and PCIe Slots

Figure 4-20, Figure 4-21, Figure 4-22, and Figure 4-23 show the riser cards supported by the server.


Figure 4-20 Riser card (1)



 **NOTE**

When installed in I/O module 1, this riser card provides slots 1 to 3. When installed in I/O module 2, this riser card provides slots 4 to 6.

Figure 4-21 Riser card (2)



 **NOTE**

When installed in I/O module 1, this riser card provides slots 1 and 3. When installed in I/O module 2, this riser card provides slots 4 and 6.

Figure 4-22 Riser card (3)




Figure 4-23 Riser card (4)





Figure 4-24 shows the PCIe slots on the rear panel.

Figure 4-24 PCIe slots



I/O module 1 provides slots 1 to 3, I/O module 2 provides slots 4 to 6, and I/O module 3 provides slots 7 and 8.

- If I/O module 1 uses a 2-slot PCIe riser module (PRM), slot 2 is unavailable.
- If I/O module 2 uses a 2-slot PRM, slot 5 is unavailable.
- If I/O module 3 uses a 1-slot PRM, slot 8 is unavailable.

Table 4-8 describes the PCIe slots.

NOTE

If CPU 2 is not installed, the PCIe slots provided by I/O modules 2 and 3 and the flexible NIC are unavailable.

Table 4-8 PCIe slot description


PCIe Device	CPU	PCIe Standard	Connector Bandwidth	Bus Bandwidth	Port Number	ROOT PORT (B/D/F)	Device (B/D/F)	Slot Size
Slot 1	CP U 1	PCI e 3.0	x16	<ul style="list-style-type: none"> ● 2-slot PRM: x16 ● 3-slot PRM: x8 	Port 2A	3a/00/0	3b/00/0	Full-height full-length
Slot 2	CP U 1	PCI e 3.0	x16	<ul style="list-style-type: none"> ● 2-slot PRM: N/A ● 3-slot PRM: x8 	Port2C	3a/02/0	3c/00/0	Full-height full-length
Slot 3	CP U 1	PCI e 3.0	x16	x8	Port 3A	5d/00/0	5e/00/0	Full-height half-length
Slot 4	CP U 2	PCI e 3.0	x16	<ul style="list-style-type: none"> ● 2-slot PRM: x16 ● 3-slot PRM: x8 	Port 1A	85/00/0	86/00/0	Full-height full-length
Slot 5	CP U 2	PCI e 3.0	x16	<ul style="list-style-type: none"> ● 2-slot PRM: N/A ● 3-slot PRM: x8 	Port1C	85/02/0	87/00/0	Full-height full-length
Slot 6	CP U 2	PCI e 3.0	x16	x8	Port 2C	ae/02/0	b2/00/0	Full-height half-length
Slot 7	CP U 2	PCI e 3.0	x16	<ul style="list-style-type: none"> ● 1-slot PRM: x16 ● 2-slot PRM: x8 	Port 3A	D7/00/0	D8/00/0	Full-height half-length
Slot 8	CP U 2	PCI e 3.0	x8	<ul style="list-style-type: none"> ● 1-slot PRM: N/A ● 2-slot PRM: x8 	Port 3C	D7/02/0	Db/00/0	Full-height half-length
RAID controller card	CP U 1	PCI e 3.0	x8	x8	Port1C	17/02/0	1B/00/0	—

PCI e De vic e	CP U	PCI e Sta nda rd	Co nne ctor Ban dwi dth	Bus Bandwidth	Port Num ber	ROOT PORT (B/D/F)	Device (B/D/F)	Slot Size
Flex ible NIC	CP U 2	PCI e 3.0	x8	x8	Port1A	17/00/0	1a/00/0	—
<ul style="list-style-type: none"> ● Note 1: B/D/F, Bus/Device/Function Number. ● Note 2: ROOT PORT (B/D/F) indicates the bus number of a CPU internal PCIe root port. Device (B/D/F) indicates the bus number (displayed on the OS) of an onboard or external PCIe port. ● Note 3: The PCIe slots that support full-height full-length PCIe cards are backwards compatible with full-height half-length or half-height half-length PCIe cards. ● Note 4: The PCIe slots that support PCIe x16 cards are backwards compatible with PCIe x8, PCIe x4, and PCIe x1 cards. ● Note 5: Table 4-8 lists the default values of B/D/F. If CPUs are not in full configuration or a PCIe card with a PCI bridge is configured, the values of B/D/F may differ. 								

4.5 Physical Structure

[Figure 4-25](#) shows the components of a server.

Figure 4-25 2288H V5 components



- | | | | |
|----|-------------------------|----|---------------------------|
| 1 | I/O module 1 | 2 | I/O module 2 |
| 3 | Power supply unit (PSU) | 4 | I/O module 3 |
| 5 | Chassis | 6 | Supercapacitor tray |
| 7 | Air duct | 8 | Front hard disk backplane |
| 9 | Fan module bracket | 10 | Fan module |
| 11 | Front hard disk | 12 | Mainboard |
| 13 | Flexible NIC | 14 | RAID controller card |
| 15 | DIMM | 16 | CPU |
| 17 | Heat sink | — | — |

5 Product Specifications

[5.1 Technical Specifications](#)

[5.2 Physical and Environmental Specifications](#)

5.1 Technical Specifications

Table 5-1 Technical specifications

Item	Specifications
Form factor	2U rack server
CPU	<p>One or two Intel® Xeon® Scalable 3100, 4100, 5100, 6100, and 8100 series processors</p> <ul style="list-style-type: none"> ● A maximum of 28 cores (2.5 GHz) ● Maximum frequency: 3.6 GHz (four cores) ● Two UPI links: 10.4 GT/s ● Maximum L3 cache capacity per core: 1.375 MB ● Maximum thermal design power (TDP): 205 W
Chipset	Intel C622
Memory	<ul style="list-style-type: none"> ● Supports a maximum of 24 DDR4 RDIMMs or LRDIMMs. ● RDIMM: twenty-four 32 GB RDIMMs for two CPUs, with a maximum memory capacity of 768 GB ● LRDIMM: twenty-four 64 GB LRDIMMs for two CPUs, with a maximum memory capacity of 1536 GB ● Maximum memory speed: 2666 MT/s ● Memory protection: Error Checking and Correcting (ECC), Mirrored Channel Mode, Single Device Data Correction (SDDC), Rank Sparring Mode, and Lockstep <p>NOTE DIMMs of different types (RDIMMs and LRDIMMs) and specifications (such as the capacity, bit width, rank, and height) cannot be installed on one server. The DIMMs on one server must have the same BOM number. For details about BOM numbers, see the Huawei Server Compatibility Checker.</p>
Storage	<ul style="list-style-type: none"> ● The server supports various hard disk configurations. For details, see Table 6-5. ● Supports hard disk hot swap. ● Supports RAID 0, 1, 10, 1E, 5, 50, 6, and 60, provides a supercapacitor to protect cache data from power failures, and supports RAID state migration, RAID configuration memory, self-diagnosis, and web-based configuration. ● Allows a SAS card or SAS RAID controller card (with a maximum of 4 GB cache) to be configured on the mainboard to improve hard disk storage performance and ensure user data security. <p>NOTE If the BIOS is in Legacy mode, the server does not support 4Kn hard disks.</p>

Item	Specifications
Network port	<ul style="list-style-type: none"> ● Two onboard 10GE electrical or optical ports (provided by the PCH), supporting NC-SI, WOL, and PXE ● Two onboard GE electrical ports (provided by the PCH), supporting NC-SI, WOL, and PXE ● Network ports provided by the flexible NIC: <ul style="list-style-type: none"> – Four GE electrical ports, supporting NC-SI, WOL, and PXE – Two 10GE electrical ports, supporting NC-SI, WOL, and PXE
RAID controller card	<p>The RAID controller card does not occupy a standard PCIe slot, which improves system scalability.</p> <p>The RAID controller card supports RAID level migration and RAID configuration memory. The server supports the following RAID controller cards:</p> <ul style="list-style-type: none"> ● LSI SAS3008 (SR130-M): supports the IT mode and does not support RAID arrays. ● LSI SAS3108 (SR430C-M): supports RAID 0, 1, 10, 5, 50, 6, and 60, and supports a supercapacitor for power-off protection.
PCIe slot	<ul style="list-style-type: none"> ● I/O modules 1 and 2 provide the following PCIe slots: <ul style="list-style-type: none"> - Two standard full-height full-length PCIe 3.0 x16 slots (bandwidth: PCIe 3.0 x8) and one standard full-height half-length PCIe 3.0 x16 slot (bandwidth: PCIe 3.0 x8) - One standard full-height full-length PCIe 3.0 x16 slot and one standard full-height half-length PCIe 3.0 x16 slot (bandwidth: PCIe 3.0 x8) ● I/O module 3 provides the following PCIe slots: <ul style="list-style-type: none"> - One standard full-height half-length PCIe 3.0 x16 slot (bandwidth: PCIe 3.0 x8) and one standard full-height half-length PCIe 3.0 x8 slot - One standard full-height half-length PCIe 3.0 x16 slot ● One Huawei SSD card, which greatly improves I/O performance for search, cache, and download services <p>NOTE Use the Huawei Server Compatibility Checker to check the PCIe cards supported by the server. For PCIe cards not listed in the Huawei Server Compatibility Checker, contact your local Huawei sales representative or Huawei technical support.</p>

Item	Specifications
Port	<ul style="list-style-type: none"> ● Two USB 2.0 ports, one USB 3.0 port, and one DB15 VGA port on the front panel <p>NOTE Only servers with eight 2.5-inch hard disks and twenty-four 2.5-inch hard disks provide the VGA port and USB 3.0 port on the front panel.</p> <ul style="list-style-type: none"> ● Two USB 3.0 ports, one DB15 VGA port, one RJ45 serial port, one RJ45 system management port, two GE electrical ports, and two 10GE electrical ports or two 10GE optical ports on the rear panel ● One built-in USB 3.0 port and two built-in SATA ports <p>NOTE The built-in USB port cannot be used for installing Windows OSs.</p>
Fan module	Four hot-swappable fan modules, allowing one-fan failures
System management	<ul style="list-style-type: none"> ● UEFI ● iBMC <p>Uses an independent port. Supports Simple Network Management Protocol (SNMP) v1/v2c/v3 and IMPI 2.0, and provides the GUI, virtual KVM, virtual media, Serial Over LAN (SOL), intelligent power supply, remote control, hardware monitoring, and LCD features.</p> <ul style="list-style-type: none"> ● NC-SI ● Supports Huawei eSight management software and integration with third-party management systems, such as VMware vCenter, Microsoft SystemCenter, and Nagios.
Security	<ul style="list-style-type: none"> ● Power-on password ● Administrator password ● Front bezel
GPU	<p>Integrates an SM750 graphics card chip to the mainboard to provide a memory capacity of 32 MB and support a maximum resolution of 1920 x 1200 at 60 Hz with 16 M colors.</p> <p>NOTE If a resolution higher than 1280 x 1024 is required, you need to install a dedicated graphics card driver.</p>

5.2 Physical and Environmental Specifications

Physical Specifications

Table 5-2 Physical specifications

Item	Specifications
Dimensions (H x W x D)	<ul style="list-style-type: none"> ● Chassis with 3.5-inch hard disks: 43 mm x 436 mm x 748 mm (1.69 in. x 17.17 in. x 29.45 in.) ● Chassis with 2.5-inch hard disks: 43 mm x 436 mm x 708 mm (1.69 in. x 17.17 in. x 27.87 in.)
Installation space	<p>The server fits into a universal cabinet that complies with the IEC 297 standard.</p> <ul style="list-style-type: none"> ● Cabinet width: 19 in. ● Cabinet depth: > 1000 mm (39.37 in.) <p>Requirements for guide rail installation:</p> <ul style="list-style-type: none"> ● L-shaped guide rails: apply only to a Huawei cabinet. ● Adjustable guide rails: apply to a cabinet with a distance of 543.5 mm to 848.5 mm (21.40 in. to 33.41 in.) between the front and rear mounting bars. ● Holding rails: apply to a cabinet with a distance of 610 mm to 914 mm (24.02 in. to 35.98 in.) between the front and rear mounting bars.
Weight in full configuration	<p>Net weight:</p> <ul style="list-style-type: none"> ● With eight 2.5-inch hard disks: 17.3 kg (38.15 lb) ● With four 3.5-inch hard disks: 18.4 kg (40.57 lb) <p>Packaging materials: 5 kg (11.03 lb)</p>
Power consumption	<p>The power consumption varies with the server configuration. Use the Huawei Server Power Calculator to calculate the power consumption.</p>

Item	Specifications
Acoustic noise	<p>The data listed in the following is the declared A-weighted sound power levels (LWAd) and declared average bystander position A-weighted sound pressure levels (LpAm) when the server is operating in a 23°C (73.4°F) ambient environment. Noise emissions are measured in accordance with ISO 7999 (ECMA 74) and declared in accordance with ISO 9296 (ECMA 109).</p> <ul style="list-style-type: none"> ● Idle: <ul style="list-style-type: none"> - LWAd: 5.92 Bels - LpAm: 42.3 dBA ● Operating: <ul style="list-style-type: none"> - LWAd: 6.14 Bels - LpAm: 46.3 dBA <p>NOTE The actual sound levels generated when the server is operating vary depending on the server configuration, workload, and ambient temperature.</p>

Environmental Specifications

- **Temperature**
 - Operating temperature: 5°C to 45°C (41°F to 113°F) (meeting the ASHRAE CLASS A3 and A4 standards). For details, see [Table 5-3](#).
 - Storage temperature: -40°C to +65°C (-104°F to +149°F).
 - Temperature change rate: < 20°C/h (36°F/h).

Table 5-3 Operating temperature specifications

Model	5°C to 35°C (41°F to 95°F)	5°C to 40°C (41°F to 104°F) (ASHRAE CLASS A3 compliant)	5°C to 45°C (41°F to 113°F) (ASHRAE CLASS A4 compliant)
Server with eight 2.5-inch hard disks	Supports all configurations.	● Does not support PCIe SSD cards.	● Does not support PCIe SSD cards.
Server with four 3.5-inch hard disks		● Supports only Intel® Xeon® Platinum 8180/8168 and Gold 6154 CPUs.	● Supports only Intel® Xeon® Platinum 8153 and Gold 6152/6140/6126/5118 CPUs, and CPUs lower than 105 W.

Model	5°C to 35°C (41°F to 95°F)	5°C to 40°C (41°F to 104°F) (ASHRAE CLASS A3 compliant)	5°C to 45°C (41°F to 113°F) (ASHRAE CLASS A4 compliant)
Note: If one fan fails, the highest operating temperature of the server is 5°C (9°F) lower than that in normal cases.			

- **Humidity**
 - Operating humidity: 8% RH to 90% RH
 - Storage humidity: 5% to 95% RH
 - Humidity change rate: < 20% RH/h
- **Altitude**
 - ≤ 3048 m (9999.88 ft). For altitudes above 900 m (2952.72 ft), the operating temperature decreases by 1°C (1.8°F) every 300 m (984.24 ft).
 - HDDs are not supported when the altitude is higher than 3000 m (9842.4 ft).
- **Corrosive gaseous contaminant**
 - Copper corrosion rate test requirements: The corrosion product thickness growth rate is lower than 300 Å/month (meeting level G1 requirements of the ANSI/ISA-71.04-2013 standard on gaseous corrosion).
 - Silver corrosion rate test requirements: The corrosion product thickness growth rate is lower than 200 Å/month.
- **Particle contaminant**
 - The equipment room environment meets the requirements of ISO 14664-1 Class 8. You are advised to hire a professional organization to monitor particle contaminants in the equipment room.
 - The equipment room is free from explosive, conductive, magnetic conductive, and corrosive dust.

6 Component Compatibility

Use the [Huawei Server Compatibility Checker](#) to check the software and hardware supported by the server.

[6.1 CPU](#)

[6.2 Memory](#)

[6.3 Storage](#)

[6.4 IO Expansion](#)

[6.5 PSU](#)

[6.6 OS and Software Support](#)

6.1 CPU

One or two Intel® Xeon® Scalable 3100, 4100, 5100, 6100, and 8100 series processors

- A maximum of 28 cores (2.5 GHz)
- Maximum frequency: 3.6 GHz (four cores)
- Two UPI links: 10.4 GT/s
- Maximum L3 cache capacity per core: 1.375 MB
- Maximum thermal design power (TDP): 205 W

NOTE

- For details about component options, consult the local Huawei sales representatives.
- CPUs on the same server must be of the same model.
- For details about CPUs, visit https://www.intel.com/content/www/us/en/homepage.html?_ga=2.147254654.87461959.1499832773-883924506.1499832773.

6.2 Memory

Memory Configuration Rules

A server provides 24 DDR4 DIMM slots. Each processor supports six memory channels and each memory channel supports two DDR4 DIMMs.

Observe the following rules when configuring DIMMs:

1. DIMMs of different types (RDIMMs and LRDIMMs) cannot be installed on one server.
2. Each channel supports a maximum of eight ranks.

NOTE

A channel supports more than eight ranks for LRDIMMs, because a quad-rank LRDIMM generates the same electrical load as a single-rank RDIMM on a memory bus.

3. The maximum number of DIMMs to be installed on the server varies with the processor type, DIMM type, number of ranks, and operating voltage. For details, see **Maximum number of DIMMs** in the following tables.

NOTE

Restriction of the number of ranks supported by each channel on the maximum number of DIMMs supported by each channel:

Number of DIMMs supported by each channel ≤ Number of ranks supported by each memory channel/Number of ranks supported by each DIMM

4. All DIMMs operate at the same speed, which is the smaller value of:
 - Memory speed supported by a CPU
 - Lowest maximum operating speed for the selected memory configuration. This speed varies with the rated speed, operating voltage, and number of DIMMs for each memory channel. For details, see **Maximum operating speed** in the following tables.

Table 6-1 RDIMM configuration

Item		Specifications
Rank		Dual-rank
Rated speed (MT/s)		2666
Operating voltage (V)		1.2
Maximum number of DIMMs		24
Maximum capacity per DIMM (GB)		32
Maximum memory capacity (GB)		768
Maximum memory capacity at maximum operating speed (GB)		768
Maximum operating speed (MT/s)	One DIMM per channel	2666
	Two DIMMs per channel	2666
Note 1: The maximum number of DIMMs listed in this table is based on dual-processor configuration. These values are halved for a server with only one processor.		


Table 6-2 LRDIMM configuration

Item		Specifications
Rank		Quad-rank
Rated speed (MT/s)		2666
Operating voltage (V)		1.2
Maximum number of DIMMs		24
Maximum capacity per DIMM (GB)		64
Maximum memory capacity (GB)		1536
Maximum memory capacity at maximum operating speed (GB)		1536
Maximum operating speed (MT/s)	One DIMM per channel	2666
	Two DIMMs per channel	2666
Note 1: The maximum number of DIMMs listed in this table is based on dual-processor configuration. These values are halved for a server with only one processor.		

Memory Slot Configuration Rules

Figure 6-1 shows the DIMM slot numbers and positions.

Figure 6-1 DIMM slots



- The server supports DIMMs of 16 GB, 32 GB, and 64 GB. A server fully configured with DIMMs has up to 1536 GB of memory.
- The server provides 24 DDR4 DIMM slots. Each CPU supports six channels. [Table 6-3](#) lists channels for each CPU.

Table 6-3 Channels

CPU Socket	Channel	DIMM
CPU 1	A	DIMM000(A)
		DIMM001(G)
	B	DIMM010(B)
		DIMM011(H)
	C	DIMM020(C)
		DIMM021(I)
	D	DIMM030(D)
		DIMM031(J)
	E	DIMM040(E)
		DIMM041(K)
	F	DIMM050(F)

CPU Socket	Channel	DIMM
		DIMM051(L)
CPU2	A	DIMM100(A)
		DIMM101(G)
	B	DIMM110(B)
		DIMM111(H)
	C	DIMM120(C)
		DIMM121(I)
	D	DIMM130(D)
		DIMM131(J)
	E	DIMM140(E)
		DIMM141(K)
	F	DIMM150(F)
		DIMM151(L)

Table 6-4 lists the order in which the DIMMs are installed.

Table 6-4 Installation order

CPU	DIMM Installation Order
CPU 1	000(A), 010(B), 020(C), 030(D), 040(E), 050(F), 001(G), 011(H), 021(I), 031(J), 041(K), 051(L)
CPU 1 and CPU 2	000(A), 100(A), 010(B), 110(B), 020(C), 120(C), 030(D), 130(D), 040(E), 140(E), 050(F), 150(F), 001(G), 101(G), 011(H), 111(H), 021(I), 121(I), 031(J), 131(J), 041(K), 141(K), 051(L), 151(L)

Memory Protection Technologies

The server supports the following memory protection technologies:

- ECC
- Mirrored Channel Mode
- SDDC
- Rank Sparing Mode
- Lockstep

Supported DIMMs

 **NOTE**

- For details about component options, consult the local Huawei sales representatives.
- DIMMs on one server must have the same BOM number.

6.3 Storage

Table 6-5 lists hard disk configurations supported by the server.

 **NOTE**

The following table is for reference only. For details about component options, consult the local Huawei sales representatives.


Table 6-5 Hard disk configurations

Configuration	Maximum Front Hard Disks	Maximum Rear Hard Disks	Hard Disk Management Mode
8-disk configuration ^[1]	8 (SAS/SATA hard disk)	—	One RAID controller card
12-disk configuration (expander)	12 (SAS/SATA hard disk)	1. I/O module 1: 2 (SAS/SATA hard disk) 2. I/O module 2: 2 (SAS/SATA hard disk)	One RAID controller card
24-disk configuration (pass-through) ^[1]	24 (SAS/SATA hard disk)	—	Three RAID controller cards (one for eight hard disks)
25-disk configuration (expander) ^[1]	25 (SAS/SATA hard disk)	I/O module 1: 2 (SAS/SATA hard disk)	One RAID controller card
[1]: The front hard disks of an 8-disk, 24-disk (pass-through), or 25-disk (expander) server can only be 2.5-inch. The other disks can be 2.5-inch or 3.5-inch.			

The following figures show the hard disk slots for these configurations.


- 8-disk configuration

Figure 6-2 8-disk configuration



- 12-disk configuration (expander)

Figure 6-3 12-disk configuration (expander)



- 24-disk configuration (pass-through)

Figure 6-4 24-disk configuration (pass-through)



- 25-disk configuration

Figure 6-5 25-disk configuration




Table 6-6 lists the comparison between RAID levels in performance and disk usage.

Table 6-6 RAID level comparison

RAID Level	Reliability	Read Performance	Write Performance	Hard Disk Usage
RAID 0	Low	High	High	100%
RAID 1	High	High	Low	50%
RAID 5	Relatively high	High	Medium	$(N - 1)/N$
RAID 6	Relatively high	High	Medium	$(N - 2)/N$
RAID 10	High	High	Medium	50%
RAID 1E	High	Medium	Medium	50%
RAID 50	High	High	Relatively high	$(N - M)/N$
RAID 60	High	High	Relatively high	$(N - M \times 2)/N$
Note: N indicates the number of member disks in a RAID group, and M indicates the number of subgroups in a RAID group.				

6.4 IO Expansion

The server supports a wide range of PCIe cards for you to choose based on the card type and transmission speed:

- Fiber Channel (FC) host bus adapter (HBA)
- Converged network adapter (CNA)
- Network expansion card
- SSD card

NOTE

For details about component options, consult the local Huawei sales representatives.

6.5 PSU

Table 6-7 lists the PSU supported by the server.

Table 6-7 Supported PSU

PSU Model	Rated Power	Rated Input Voltage
550 W AC Platinum PSU	550 W	100 V AC to 240 V AC, and 240 V DC

 NOTE

- The PSUs are hot-swappable and work in 1+1 redundancy mode.
- [Table 6-7](#) is for reference only. For details about component options, see the [Huawei Server Compatibility Checker](#).
- PSUs on one server must have the same BOM number.
- The PSUs provide short-circuit protection. The PSUs that support dual input live wires provide double-pole fuse.

6.6 OS and Software Support

For details about component options, see the [Huawei Server Compatibility Checker](#).

7 System Management

The server uses Huawei's proprietary Intelligent Baseboard Management Controller (iBMC) to implement remote server management. The iBMC complies with Intelligent Platform Management Interface (IPMI) 2.0 and provides highly reliable hardware monitoring and management.

The iBMC supports the following features and protocols:

- KVM and text console redirection
- Remote virtual media
- IPMI
- Simple Network Management Protocol (SNMP)
- Redfish 1.0
- Login using a web browser

Table 7-1 describes the features of the iBMC.

Table 7-1 iBMC features

Feature	Description
Management interface	Integrates with any standard management system through the following interfaces: <ul style="list-style-type: none">● IPMI● CLI● HTTPS● SNMP● Redfish
Fault detection	Detects faults and accurately locates faults in hardware, for example, an FRU.
Alarm management	Supports alarm management and reports alarms using the SNMP trap, Simple Mail Transfer Protocol (SMTP), and syslog service to ensure 24/7 continuous operation.

Feature	Description
Integrated virtual KVM	Provides remote maintenance measures for troubleshooting the system, and supports a maximum resolution of 1920 x 1200.
Integrated virtual media	Virtualizes local media devices, images, USB keys, and folders into media devices on a remote server, simplifying OS installation. (The virtual DVD-ROM drive supports a maximum transmission rate of 8 MB/s.)
WebUI	Provides a user-friendly graphical user interface (GUI), which simplifies users' configuration and query operations.
Fault reproduction	Reproduces faults to facilitate fault diagnosis.
Screen snapshots and screen videos	Allows you to view screenshots and videos without login, which facilitates routine preventive maintenance inspection (PMI)
Domain Name Service (DNS)/Active Directory (AD)	Supports the DNS and AD, significantly simplifying network and configuration management.
Dual-image backup	Starts software from a backup image if the software fails.
Asset management	Supports intelligent asset management.
Intelligent power management	Uses the power capping technology to increase deployment density, and uses dynamic energy saving to lower operating expenses.
IPv6	Supports IPv6 to ensure sufficient IP addresses.
Network Controller Sideband Interface (NC-SI)	Supports NC-SI, which allows you to access the iBMC through the service network port.

8 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-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 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 the local Huawei representative office.

Table 8-1 describes the warranty service response time.

Table 8-1 Response time

Service	Response Time	Description	Remarks
Help Desk	24/7	Available 24 hours a day, 7 days a week (00:00 to 24:00, Monday to Sunday)	None
Remote troubleshooting	24/7	Available 24 hours a day, 7 days a week (00:00 to 24:00, Monday to Sunday)	The response time is the period between the end of a customer's service request and beginning of the response by the technical support to offer troubleshooting services.

Service	Response Time		Description	Remarks
Online technical support	24/7		Huawei enterprise support website (http://e.huawei.com): available 24 hours a day, 7 days a week (00:00 to 24:00, Monday to Sunday)	None
Licensing of software updates	24/7		Available 9 hours a day, 5 days a week (09:00 to 18:00, Monday to Friday), excluding official holidays.	The repaired or replacement parts will be shipped within 45 calendar days after Huawei receives the defective parts.
Return for repair	Outside China	9/5 hours, shipment within 45 calendar days	Available 9 hours a day, 5 days a week (09:00 to 18:00, Monday to Friday), excluding official holidays.	The repaired or replacement parts will be shipped within 45 calendar days after Huawei receives the defective parts.
	In China	10/5 hours, next business day	Available 10 hours a day, 5 days a week (08:00 to 18:00, Monday to Friday), excluding official holidays. Arrival: NBD	<ol style="list-style-type: none"> Service requests submitted after 15:30 will be handled the next business day. The response time starts when a remote decision is made to appoint an onsite engineer.

Table 8-2 describes warranty services provided by Huawei.

Table 8-2 Warranty services

Service	Description
Help Desk	Huawei provides 24-hour after-sales technical support (such as handling requests for troubleshooting and hardware repair), receives and handles customer inquiries, complaints, and suggestions through a dedicated hotline.
Remote troubleshooting	After receiving a service request for rectifying a network or system fault, Huawei engineers will analyze and handle the fault remotely and proceed to rectify it in the shortest possible time. There are two methods for remote troubleshooting: telephone support and remote access.
Online technical support	Huawei enterprise support website (http://e.huawei.com) provides product and technical materials, such as product manuals, configuration guides, networking case study, and maintenance experience collections. Registered users can access the website and download required documents.
Licensing of software updates	To ensure that the devices operate stably, Huawei provides software patches whenever necessary
Return for repair	Huawei provides repair or replacement services for customers within the promised time to meet customer needs for spare parts. You can return defective parts to the designated Huawei site after submitting a service request. Huawei provides a three-year warranty for parts replacement and limited onsite repair for the servers used in China. Huawei provides a 10-hour-a-day, 5-day-a-week NBD support program. Huawei provides a three-year warranty for parts replacement and repair for the servers used outside China. Huawei provides a 9-hour-a-day, 5-day-a-week NBD support program. Service requests will be handled the next business day. Huawei delivers the repaired or new parts within 45 calendar days after receiving the defective parts.

9 Certifications

No.	Country/ Region	Certification	Standards
1	China	CCC	GB4943.1-2011 GB9254-2008 (Class A) GB17625.1-2012

10 References

For more information about the product, see the following documents:

- [Server Data Sheet](#)
- [Server-related Documentation](#)
- [Huawei Server Information Self-Service Platform](#)

Huawei FusionServer Pro

V5 Rack Server

Bring Pervasive Intelligent Computing
with Chip Innovation



HUAWEI TECHNOLOGIES CO., LTD.



| High-Density Deployment with Lower OPEX |



1288H V5 (4-drive)




1288H V5 (8-drive)



1288H V5 (10-drive)

- 2 Intel® Xeon® Scalable processors in 1U space, with 24 DDR4 DIMMs
- Up to 4 3.5-inch or 10 2.5-inch hard drives for local storage, or 4/8 NVMe SSDs
- 2 10GE and 2 GE LAN on motherboard (LOM) ports, and 5 PCIe expansion slots
- Leverages intelligent energy saving to improve performance per watt by 15%; combines intelligent management features to enable up to 93% accuracy for fault locating

Delivers high computing density, which is especially suitable for scenarios such as virtualization, high-performance computing (HPC), and big data analytics.




Superior Performance, Ultra-high Density

- Supports 2 Intel® Xeon® Scalable Processors in a 1U space. Its Ultra Path Interconnect (UPI) bus supports rates of up to 10.4 GT/s, and a single CPU supports up to 28 cores. The server supports the Intel® Turbo Boost, hyper-threading, and Advanced Vector Extensions (AVX-512). A single processor delivers up to 40% higher compute power than the previous-generation processor.
- Supports 24 DDR4 DIMMs with a memory capacity of up to 3 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- Supports 12 DCPMMs as volatile or non-volatile storage, which can be used together with 12 DDR4 DIMMs, offering up to 7.5 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- Supports heterogeneous computing acceleration, configurable with 2 single-slot half-height half-length (HHHL) GPU or FPGA accelerator cards.
- Supports two GE and two 10GE LAN on motherboard (LOM) ports, meeting networking requirements of 98% scenarios with streamlined configuration.



Smart Power Saving and Better Energy Efficiency

- Leverages patented Dynamic Energy Management Technology (DEMT), and multiple power-saving measures such as component hibernation, proportional-integral-derivative (PID) algorithm based fan speed tuning, and active-standby power supplies, driving down overall equipment power consumption by up to 15% without compromising workload performance.
- Supports 80 Plus® Titanium power supply units (PSUs), with up to 96% conversion efficiency and compliant with ENERGY STAR and China Environmental Labelling.
- Supports 550 W, 900 W, 1,200 W, and 1,500 W PSU options, flexibly adapting to different power requirements. The 1,200 W and 1,500 W PSUs support DC and high-voltage DC (HVDC) technologies, enabling better energy utilization.



Unmatched Intelligent Manageability, Integration, and Openness

- Uses patented intelligent Fault Diagnosis & Management (FDM) technology, delivering up to 93% accuracy in diagnosing core component faults.
- Integrates eSight or FusionDirector for smart entire-lifecycle O&M, boosting deployment and O&M efficiency.
 - » Supports batch OS installation, slashing the average OS installation time of each server from hours to minutes.
 - » Supports automated firmware upgrade, with flexible and configurable upgrade policies for different components and drivers.
 - » Supports stateless computing, allowing for rapid replication of live-network configuration and swift failover.
- Integrates fault diagnosis LEDs on the front panel to display error codes in real time, enabling maintenance personnel to rapidly locate a fault.
- Provides standardized open interfaces and development guides, facilitating seamless integration with third-party management software.

Huawei FusionServer Pro 1288H V5 Server



Form factor	1U rack server
Processors	1 or 2 1st Generation Intel® Xeon® Scalable processors (3100/4100/5100/6100/8100 series), up to 205 W 1 or 2 2nd Generation Intel® Xeon® Scalable processors (3200/4200/5200/6200/8200 series), up to 205 W
Chipset platform	Intel C622
Memory	24 DDR4 DIMM slots, 2933 MT/s; up to 12 DCPMMs, 2666 MT/s
Internal storage	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none"> • 10 x 2.5-inch hard drives (6–8 NVMe SSDs and 2–4 SAS/SATA drives, with a total number of 10 or less) • 10 x 2.5-inch SAS/SATA/SSDs (2–4 NVMe SSDs and 6–8 SAS/SATA HDDs, with a total number of 10 or less) • 10 x 2.5-inch SAS/SATA/SSDs • 8 x 2.5-inch SAS/SATA hard drives • 4 x 3.5-inch SAS/SATA hard drives Flash storage: <ul style="list-style-type: none"> • 2 M.2 SSDs
RAID support	<ul style="list-style-type: none"> • RAID 0, 1, 10, 5, 50, 6, or 60 • Configured with a supercapacitor for cache power-off protection • Supports RAID level migration, drive roaming, self-diagnosis, and web-based remote configuration
Network ports	LOM: 2 x 10GE + 2 x GE ports Flexible NIC: 2 x GE, 4 x GE, 2 x 10GE, 2 x 25GE, or 1/2 x 56G FDR IB ports
PCIe expansion	Up to 5 PCIe 3.0 slots, including 1 for a RAID controller card and 1 for a flexible NIC
Heterogeneous accelerator cards	2 single-slot HHHH GPU or FPGA heterogeneous accelerator cards For details, visit http://support.huawei.com/onlinetoolsweb/ftca/indexEn?serie=2 .
Fan modules	7 hot-swappable counter-rotating fan modules with support for N+1 redundancy
Power supply units	2 hot-swappable PSUs with support for 1+1 redundancy and the following configuration options: <ul style="list-style-type: none"> • 550 W AC Platinum PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) • 900 W AC Platinum/Titanium PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) • 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC) • 1500 W 380 V HVDC PSUs (input: 260 V to 400 V DC) • 1200 W -48 V to -60 V DC PSUs (input: -38.4 V to -72 V DC)
Management	<ul style="list-style-type: none"> • Huawei iBMC integrates one dedicated management GE network port to provide comprehensive management features such as fault diagnosis, automated O&M, and hardware security hardening. • iBMC supports standard interfaces such as Redfish, SNMP, and IPMI 2.0; provides a remote management interface based on HTML5/VNC KVM; supports CD-free deployment and the Agentless feature, simplifying management. • (Optional) Configured with the Huawei eSight or FusionDirector management software to provide advanced management features such as stateless computing, batch OS deployment, and automated firmware upgrade, enabling smart and automatic entire-lifecycle management.
Operating Systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, CentOS, Citrix XenServer, VMware ESXi For details, visit http://support.huawei.com/onlinetoolsweb/ftca/indexEn?serie=2 .
Security	<ul style="list-style-type: none"> • Power-on password • Administrator password • Trusted Platform Module (TPM) • Security front panel
Operating temperature	5°C to 45°C (41°F to 113°F), compliant with ASHRAE A3 and A4
Certification	CE, UL, FCC, CCC, and RoHS
Installation suite	L-shaped guide rails, adjustable guide rails, and holding rails
Dimensions (H x W x D)	Chassis with 3.5-inch hard drives: 43 mm x 447 mm x 748 mm (1.70 in. x 17.60 in. x 29.45 in.) Chassis with 2.5-inch hard drives: 43 mm x 447 mm x 708 mm (1.70 in. x 17.60 in. x 27.87 in.)

*Last update: July 16, 2019

For more information

To learn more about Huawei servers, contact Huawei sales representatives or business partners, or visit:

<http://e.huawei.com/en/products/cloud-computing-dc/servers>



Scan for an
electronic copy



Scan to learn more about
Huawei servers


| Flexible Configurations for Diverse Workloads |



The Huawei FusionServer Pro 2288/2288H V5 is a 2U 2-socket rack server that supports various configurations and can be widely used in scenarios such as cloud computing virtualization, databases, and big data. The server can be configured with two Intel® Xeon® Scalable processors, 16/24¹ DDR4 DIMM slots, 8/10² PCIe slots, and large-capacity local storage resources.

It incorporates patented technologies such as Dynamic Energy Management Technology (DEMT) and Fault Diagnosis & Management (FDM), and integrates Huawei's eSight or FusionDirector software for entire-lifecycle management, helping customers drive down OPEX and improve ROI.

¹ The 2288 V5 supports 16 DIMM slots, and the 2288H V5 supports 24. ² The 2288 V5 supports 8 PCIe slots, and the 2288H V5 supports 10.



Supreme Performance with Flexible Configurations


- Supports 2 Intel® Xeon® Scalable processors in a 2U space, with an Ultra Path Interconnect (UPI) bus speed of up to 10.4 GT/s. Each CPU supports up to 20/28³ cores. The server supports the Intel® Turbo Boost, hyper-threading, and Advanced Vector Extensions (AVX-512). A single processor delivers up to 40% higher compute power than its predecessor.
- The 2288 V5 supports 16 DDR4 DIMMs with a memory capacity of up to 2 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- The 2288H V5 supports 24 DDR4 DIMMs with a memory capacity of up to 3 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- The 2288H V5 supports 12 DCPMMs as volatile or non-volatile storage, which can be used together with 12 DDR4 DIMMs, offering up to 7.5 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- The 2288H V5 supports heterogeneous computing acceleration. It can be configured with two dual-slot full-height full-length GPU or FPGA accelerator cards.
- The 2288 V5 supports 16 x 3.5-inch or 27 x 2.5-inch local storage drives. The 2288H V5 supports 20 x 3.5-inch or 31 x 2.5-inch (4/8/12/24/28 NVMe SSDs) local storage drives.
- The 2288 V5 supports 2 GE LOM ports. The 2288H V5 supports 2 GE and 2 x 10GE LOM ports, meeting the networking requirements of 98% scenarios with streamlined configuration.

³ For the 2288 V5, a single CPU supports up to 20 computing cores. For the 2288H V5, a single CPU supports up to 28 computing cores.



Smart Power Saving and Better Energy Efficiency

- Leverages patented Dynamic Energy Management Technology (DEMT), and multiple power-saving measures such as component hibernation, proportional-integral-derivative (PID) algorithm based fan speed tuning, and active-standby power supplies, driving down overall equipment power consumption by up to 15% without compromising workload performance.
- Supports 80 Plus® Titanium power supply units (PSUs), with up to 96% conversion efficiency and compliant with ENERGY STAR and China Environmental Labelling.
- Supports 550 W, 900 W, 1,200 W, and 1,500 W PSU options, flexibly adapting to different power requirements. The 1,200 W and 1,500 W PSUs support DC and high-voltage DC (HVDC) technologies, enabling better energy utilization.



Unmatched Intelligent Manageability, Integration, and Openness

- Uses patented intelligent Fault Diagnosis & Management (FDM) technology, delivering up to 93% accuracy in diagnosing core component faults.
- Integrates eSight or FusionDirector for smart entire-lifecycle O&M, boosting deployment and O&M efficiency.
 - » Supports batch OS installation, slashing the average OS installation time of each server from hours to minutes.
 - » Supports automated firmware upgrade, with flexible and configurable upgrade policies for different components and drivers.
 - » Supports stateless computing, allowing for rapid replication of live-network configuration and swift failover.
- Integrates a touchscreen LCD panel for fault diagnosis, allowing O&M personnel to quickly locate faults (supported only by the 2288H V5 8-drive models).
- Provides standardized open interfaces and development guides, facilitating seamless integration with third-party management software.

Huawei FusionServer Pro 2288/2288H V5 Server



	2288 V5	2288H V5
Form factor	2U rack server	
Processors	1 or 2 1st Generation Intel® Xeon® Scalable processors (3100/4100/5100/6100/8100 series), up to 125 W 1 or 2 2nd Generation Intel® Xeon® Scalable processors (3200/4200/5200/6200/8200 series), up to 125 W	1 or 2 1st Generation Intel® Xeon® Scalable processors (3100/4100/5100/6100/8100 series), up to 205 W 1 or 2 2nd Generation Intel® Xeon® Scalable processors (3200/4200/5200/6200/8200 series), up to 205 W
Chipset platform	Intel C621	Intel C622
Memory	16 DDR4 DIMM slots, up to 2,933 MT/s	24 DDR4 DIMM slots, 2933 MT/s; up to 12 DCPMMs, 2666 MT/s
Internal storage	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none"> 8 x 2.5-inch SAS/SATA hard drives or SSDs 12/16 x 3.5-inch SAS/SATA hard drives 27 x 2.5-inch SAS/SATA hard drives or SSDs NVMe SSD card: <ul style="list-style-type: none"> 4 x NVMe SSD cards Flash storage: <ul style="list-style-type: none"> 2 M.2 SSDs 	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none"> 8 x 2.5-inch SAS/SATA hard drives 12/16/20 x 3.5-inch SAS/SATA hard drives 4, 8, 12, 24, or 28 NVMe SSDs 31 x 2.5-inch SAS/SATA hard drives Flash storage: <ul style="list-style-type: none"> 2 M.2 SSDs
RAID support	<ul style="list-style-type: none"> RAID 0, 1, 10, 1E, 5, 50, 6, or 60 Configured with a supercapacitor for cache power-off protection Supports RAID level migration, drive roaming, self-diagnosis, and web-based remote configuration 	
Network ports	LOM: 2 GE ports Flexible NIC: 2 x GE, 4 x GE, 2 x 10GE, or 1/2 x 56G FDR IB ports	LOM: 2 x 10GE + 2 x GE ports Flexible NIC: 2 x GE, 4 x GE, 2 x 10GE, 2 x 25GE, or 1/2 x 56G FDR IB ports
PCIe expansion	Up to 8 PCIe 3.0 slots, including 1 for a RAID controller card and 1 for a flexible NIC.	Up to 10 PCIe 3.0 slots, including 1 for a RAID controller card and 1 for a flexible NIC.
Heterogeneous accelerator cards	/	2 dual-slot FHFL GPU or FPGA heterogeneous accelerator cards For details, visit http://support.huawei.com/online/toolsweb/ftca/indexEn?serie=2 .
Fan modules	4 hot-swappable counter-rotating fan modules with support for N+1 redundancy	
Power supply units	2 hot-swappable PSUs with support for 1+1 redundancy and the following configuration options: <ul style="list-style-type: none"> 550 W AC Platinum PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) 900 W AC Platinum PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC) 	2 hot-swappable PSUs with support for 1+1 redundancy and the following configuration options: <ul style="list-style-type: none"> 550 W AC Platinum PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) 900 W AC Platinum/Titanium PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC) 1500 W 380 V HVDC PSUs (input: 260 V to 400 V DC) 1200 W -48 V to -60 V DC PSUs (input: -38.4 V to -72 V DC)
Management	<ul style="list-style-type: none"> Huawei iBMC integrates one dedicated management GE network port to provide comprehensive management features such as fault diagnosis, automated O&M, and hardware security hardening. iBMC supports standard interfaces such as Redfish, SNMP, and IPMI 2.0; provides a remote management interface based on HTML5/VNC KVM; supports CD-free deployment and the Agentless feature, simplifying management. (Optional) Configured with the Huawei eSight or FusionDirector management software to provide advanced management features such as stateless computing, batch OS deployment, and automated firmware upgrade, enabling smart and automatic entire-lifecycle management. 	
Operating Systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, CentOS, Citrix XenServer, VMware ESXi For details, visit http://support.huawei.com/online/toolsweb/ftca/indexEn?serie=2 .	
Security	Power-on password, Administrator password, Trusted Platform Module (TPM), Security front panel	
Operating temperature	5°C to 40°C (41°F to 104°F), compliant with ASHRAE A2 and A3	5°C to 45°C (41°F to 113°F), compliant with ASHRAE A3 and A4
Certification	CCC	CE, UL, FCC, CCC, and RoHS
Installation suite	L-shaped guide rails, adjustable guide rails, and holding rails	
Dimensions (H x W x D)	Chassis with 3.5-inch hard drives: 86.1 mm x 447 mm x 748 mm (3.39 in. x 29.45 in.) Chassis with 2.5-inch hard drives: 86.1 mm x 447 mm x 708 mm (3.39 in. x 27.87 in.)	

*Last update: July 16, 2019

For more information

To learn more about Huawei servers, contact Huawei sales representatives or business partners, or visit:

<http://e.huawei.com/en/products/cloud-computing-dc/servers>



Scan for an electronic copy



Scan to learn more about Huawei servers

Huawei FusionServer Pro 2488/2488H V5 Server




2488/2488H V5

| New Option for Distributed Deployment with Higher Computing Efficiency |

- 4 Intel® Xeon® Scalable processors in 2U space, with 32/48 DDR4 DIMMs
- Up to 25 2.5-inch hard drives for local storage, or 8 NVMe SSDs
- 2 10GE and 2 GE LOM ports, and 9/11 PCIe expansion slots
- Delivers lower OPEX than 2U 2-socket servers; leverages intelligent energy saving to improve performance per watt by 15%; combines intelligent management features to enable up to 93% accuracy for fault locating

Unlocks high computing efficiency for scenarios such as virtualization, HPC, database, and SAP HANA in-memory computing.



Superior Performance with Higher Efficiency

- Supports 4 Intel® Xeon® Scalable Processors (Platinum or Gold series) in a 2U space. Its Ultra Path Interconnect (UPI) bus supports rates of up to 10.4 GT/s, and a single processor supports up to 28 cores. The server supports Intel® Turbo Boost, hyper-threading, and Advanced Vector Extensions (AVX-512). A single processor delivers up to 40% higher compute power than the previous-generation processor.
- Supports 32/ 48¹ DDR4 DIMMs with a memory capacity of up to 4 TB / 6 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- The 2488 V5 supports 8 DCPMMs as volatile or non-volatile storage, which can be used together with 24 DDR4 DIMMs, offering up to 7 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- The 2488H V5 supports 24 DCPMMs as volatile or non-volatile storage, which can be used together with 24 DDR4 DIMMs, offering up to 15 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- Supports two GE and two 10GE LAN on motherboard (LOM) ports, meeting networking requirements of 98% scenarios with streamlined configuration.
- Supports up to 25 2.5-inch local hard drives (configurable with 8 NVMe SSDs).
- One 2488/2488H V5 saves up to 32%² OPEX in the virtualization scenario compared with two traditional 2U 2S servers.


¹ The 2488 V5 supports up to 32 DIMMs and 2488H V5 supports 48 DIMMs.

² Data is derived from Huawei lab tests; actual improvement depends on the real-world scenario.



Smart Power Saving and Better Energy Efficiency

- Leverages patented Dynamic Energy Management Technology (DEMT), and multiple power-saving measures such as component hibernation, proportional-integral-derivative (PID) algorithm based fan speed tuning, and active-standby power supplies, driving down overall equipment power consumption by up to 15% without compromising workload performance.
- Supports 2,000 W Platinum AC power supply unit (PSU), meeting ultra-high performance requirements; leverages the DC and high-voltage DC (HVDC) technologies to improve energy utilization.
- PSUs meet the requirements of ENERGY STAR and China Environmental Labelling.



Unmatched Intelligent Manageability, Integration, and Openness

- Uses patented intelligent Fault Diagnosis & Management (FDM) technology, delivering up to 93% accuracy in diagnosing core component faults.
- Integrates eSight or FusionDirector for smart entire-lifecycle O&M, driving a leap in deployment and O&M efficiency.
 - » Supports batch OS installation, slashing the average OS installation time of each server from hours to minutes.
 - » Supports automated firmware upgrade, with flexible and configurable upgrade policies for different components and drivers.
 - » Supports stateless computing, allowing for rapid replication of live-network configuration and swift failover.
- Integrates a touchscreen LCD panel for fault diagnosis, allowing O&M personnel to quickly locate faults (supported only by the 2488/2488H V5 8-drive models).
- Integrates fault diagnosis LEDs on the front panel to display error codes in real time, enabling maintenance personnel to rapidly locate a fault.
- Provides standardized open interfaces and development guides, facilitating seamless integration with third-party management software.

Huawei FusionServer Pro 2488/2488H V5 Server



	2488 V5	2488H V5
Form factor	2U rack server	
Processors	2 or 4 1st Generation Intel® Xeon® Scalable processors (5100/6100/8100 series), up to 205 W 2 or 4 2nd Generation Intel® Xeon® Scalable processors (5200/6200/8200 series), up to 205 W	
Chipset platform	Intel C622	
Memory	32 DDR4 DIMM slots, 2933 MT/s; up to 8 DCPMMs, 2666 MT/s	48 DDR4 DIMM slots, 2933MT/s; up to 24 DCPMMs, 2666 MT/s
Internal storage	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none"> • 8 x 2.5-inch SAS/SATA hard drives • 25 x 2.5-inch SAS/SATA hard drives • 8 x 2.5-inch NVMe SSDs and 16 x 2.5-inch SAS/SATA hard drives Flash storage: <ul style="list-style-type: none"> • 2 M.2 SSDs 	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none"> • 8 x 2.5-inch SAS/SATA hard drives • 25 x 2.5-inch SAS/SATA hard drives • 24 x 2.5-inch SAS/SATA hard drives • 8 x 2.5-inch NVMe SSDs and 16 x 2.5-inch SAS/SATA hard drives Flash storage: <ul style="list-style-type: none"> • 2 M.2 SSDs
RAID support	<ul style="list-style-type: none"> • RAID 0, 1, 10, 1E, 5, 50, 6, or 60 • Configured with a supercapacitor for cache power-off protection • Supports RAID level migration, drive roaming 	
LOM network ports	2 x GE + 2 x 10GE ports	
PCIe expansion	Up to 9 PCIe 3.0 slots	Up to 11 PCIe 3.0 slots
Fan modules	4 hot-swappable fan modules, providing protection against single-fan failures	
Power supply units	2 hot-swappable PSUs, with support for 1+1 redundancy. The following PSUs are supported: <ul style="list-style-type: none"> • 2000 W AC Platinum PSUs 1800 W (input: 200 V to 220 V AC, or 192 V to 200 V DC) 2000 W (input: 200 V to 240 V AC, or 200 V to 288 V DC) • 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC) • 900 W AC Platinum PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC) • 1200 W DC PSUs (input: -38.4 V to -72 V DC) 	
Management	<ul style="list-style-type: none"> • Huawei iBMC integrates one dedicated management GE network port to provide comprehensive management features such as fault diagnosis, automated O&M, and hardware security hardening. • iBMC supports standard interfaces such as Redfish, SNMP, and IPMI 2.0; provides a remote management interface based on HTML5/VNC KVM; supports CD-free deployment and the Agentless feature, simplifying management. • (Optional) Configured with the Huawei eSight or FusionDirector management software to provide advanced management features such as stateless computing, batch OS deployment, and automated firmware upgrade, enabling smart and automatic entire-lifecycle management. 	
Operating Systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, CentOS, Citrix XenServer, VMware ESXi For details, visit http://support.huawei.com/onlinetoolsweb/ftca/indexEn?serise=2 .	
Security	<ul style="list-style-type: none"> • Power-on password • Administrator password • Trusted Platform Module (TPM) • Secure startup • Security front panel 	
Operating temperature	5°C to 45°C (41°F to 113°F), compliant with ASHRAE Classes A3 and A4	
Certification	CE, ENERGY STAR, FCC, CCC, RoHS	
Installation suite	L-shaped guide rails, adjustable guide rails, and holding rails	
Dimensions (H x W x D)	86.1 mm x 447 mm x 748 mm (3.39 in. x 17.60 in. x 29.45 in.)	

*Last update: July 16, 2019

For more information

To learn more about Huawei servers, contact Huawei sales representatives or business partners, or visit:
<http://e.huawei.com/en/products/cloud-computing-dc/servers>



Scan for an
electronic copy



Scan to learn more about
Huawei servers



5288 V5

| Hybrid Storage Architecture, Tiered Data Storage |

- 2 Intel® Xeon® Scalable processors in 4U space, with 24 DDR4 DIMMs
- Up to 44 3.5-inch hard drives for local storage, or 4/8 NVMe SSDs
- 2 10GE and 2 GE LOM ports, and 10 PCIe expansion slots
- Leverages intelligent energy saving to improve performance per watt by 15%; combines intelligent management features to enable up to 93% accuracy for fault locating

Offers ultra-large storage capacities, which is ideal for hot, warm, and cold data tiered storage in scenarios such as Content Delivery Network (CDN), video cloud, and massive data archiving.




Ultralarge Capacity, Tiered Storage

- Supports 2 Intel® Xeon® Scalable Processors in a 4U space. Its Ultra Path Interconnect (UPI) bus supports rates of up to 10.4 GT/s, and a single CPU supports up to 28 cores. The server supports Intel® Turbo Boost, hyper-threading, and Advanced Vector Extensions (AVX-512). A single processor delivers up to 40% higher compute power than the previous-generation processor.
- Supports 24 DDR4 DIMMs with a memory capacity of up to 3 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- Supports 12 DCPMMs as volatile or non-volatile storage, which can be used together with 12 DDR4 DIMMs, offering up to 7.5 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- Ultra-large storage space with 44 3.5-inch and 4 2.5-inch hard drives (up to 8 NVMe SSDs), ideal for tiered storage of hot, warm, and cold data.
- Supports two GE and two 10GE LAN on motherboard (LOM) ports, meeting networking requirements of 98% scenarios with streamlined configuration.



Smart Power Saving and Better Energy Efficiency

- Leverages patented DEMA, and multiple power-saving measures such as component hibernation, proportional-integral-derivative (PID) algorithm based fan speed tuning, and active-standby power supplies, driving down overall equipment power consumption by up to 15% without compromising workload performance.
- Fitted with 80 Plus® Platinum power supply units (PSUs), up to 94% conversion efficiency; complies with ENERGY STAR standards and has passed the China Energy Conservation and Environmentally-Friendly Certification.
- PSUs with 900 W, 1500 W and more power options to flexibly adapt to different power requirements, improving energy utilization.



Unmatched Intelligent Manageability, Integration, and Openness

- Uses patented intelligent Fault Diagnosis & Management (FDM) technology, delivering up to 93% accuracy in diagnosing core component faults.
- Integrates eSight or FusionDirector for smart entire-lifecycle O&M, boosting deployment and O&M efficiency.
 - » Supports batch OS installation, slashing the average OS installation time of each server from hours to minutes.
 - » Supports automated firmware upgrade, with flexible and configurable upgrade policies for different components and drivers.
 - » Supports stateless computing, allowing for rapid replication of live-network configuration and swift failover.
- Integrates fault diagnosis LEDs to display error codes in real time, enabling maintenance personnel to rapidly locate a fault.
- Provides standardized open interfaces and development guides, facilitating seamless integration with third-party management software.

Huawei FusionServer Pro 5288 V5 Server



Form Factor	4U rack server
Processors	1 or 2 1st Generation Intel® Xeon® Scalable processors (3100/4100/5100/6100/8100 series), up to 205 W 1 or 2 2nd Generation Intel® Xeon® Scalable processors (3200/4200/5200/6200/8200 series), up to 205 W
Chipset	Intel C622
Memory	24 DDR4 DIMM slots, 2933 MT/s; up to 12 DCPMMs, 2666 MT/s
Internal Storage	Supports hot-swappable hard drives with the following configuration options: Front: <ul style="list-style-type: none">• 24 x 3.5-inch SAS/SATA hard drives Embedded: <ul style="list-style-type: none">• 4 x 3.5-inch SAS/SATA hard drives Rear: <ul style="list-style-type: none">• 16 x 3.5-inch SAS/SATA hard drives• 16 x 3.5-inch SAS/SATA hard drives + 4 x 2.5-inch SAS/SATA hard drives or NVMe SSDs• 14 x 3.5-inch SAS/SATA hard drives (configurable with 4 NVMe SSDs) + 4 x 2.5-inch SAS/SATA hard drives or NVMe SSDs (This configuration does not support internal hard disk and I/O module 1) Flash storage: <ul style="list-style-type: none">• Two M.2 SSDs
RAID	RAID 0, 1, 5, 50, 6, or 60; optional supercapacitor to protect cache data from power failures; RAID level migration, drive roaming, self-diagnosis, and web-based remote configuration
Network Ports	LAN on motherboard (LOM): 2 x 10GE + 2 x GE ports Flexible NIC: 2 x GE, 4 x GE, 2 x 10GE, 2 x 25GE, or 1/2 x 56G FDR IB ports
PCIe Expansion	Up to 10 PCIe 3.0 slots, including 1 for a RAID controller card and 1 for a flexible NIC
Fan Modules	4 hot-swappable counter-rotating fan modules with optional N+1 redundancy
Power Supply	2 hot-swappable PSUs with optional 1+1 redundancy. Supported options include: <ul style="list-style-type: none">• 900 W AC Platinum/Titanium PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC)• 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC)
Management	<ul style="list-style-type: none">• Huawei iBMC integrates one dedicated management GE network port to provide comprehensive management features such as fault diagnosis, automated O&M, and hardware security hardening.• iBMC supports standard interfaces such as Redfish, SNMP, and IPMI 2.0; provides a remote management interface based on HTML5/VNC KVM; supports CD-free deployment and the Agentless feature, simplifying management.• (Optional) Configured with the Huawei eSight or FusionDirector management software to provide advanced management features such as stateless computing, batch OS deployment, and automated firmware upgrade, enabling smart and automatic entire-lifecycle management.
Operating Systems	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, CentOS, Citrix XenServer, VMware ESXi For details, see http://support.huawei.com/online/toolsweb/ftca/index?serise=2 .
Security	<ul style="list-style-type: none">• Power-on password• Administrator password• Trusted Platform Module (TPM)• Security front panel
Operating Temperature	5°C to 40°C (41°F to 104°F) (ASHRAE Class A3 compliant)
Certification	CE, UL, FCC, CCC, and RoHS
Installation Suite	L-shaped guide rails, adjustable guide rails, and holding rails
Dimensions (H x W x D)	175 mm x 447 mm x 748 mm (6.89 in. x 17.60 in. x 29.45 in.)

*Last update: July 16, 2019

For more information

To learn more about Huawei servers, contact Huawei sales representatives or business partners, or visit:

<http://e.huawei.com/en/products/cloud-computing-dc/servers>



Scan for an
electronic copy



Scan to learn more about
Huawei servers




5885H V5

| Deliver Excellent Performance and Scalability to Enable Fast and Stable Mission-Critical Services |

- 4 Intel® Xeon® Scalable processors in 4U space, with 48 DDR4 DIMMs
- Up to 25 2.5-inch hard drives for local storage, or 8 NVMe SSDs
- 2 10GE and 2 GE LOM ports, and 15 PCIe expansion slots
- Leverages intelligent energy saving to improve performance per watt by 15%; combines intelligent management features to enable up to 93% accuracy for fault locating

Delivers excellent stability and reliability for scenarios such as virtualization, HPC, and database.




High Efficiency, Stability, and Expandability

- Supports four Intel® Xeon® Scalable processors in a 4U space. Its Ultra Path Interconnect (UPI) bus supports a speed of up to 10.4 GT/s, and a single CPU supports up to 28 computing cores. The server supports Intel® Turbo Boost, hyper-threading, and Advanced Vector Extensions (AVX-512). A single processor delivers up to 65% higher compute performance than the previous-generation processor.
- Supports 48 DDR4 DIMMs with a memory capacity of up to 6 TB (configured with 128 GB DIMMs) to meet large-capacity memory application requirements.
- Supports 24 DCPMMs as volatile or non-volatile storage, which can be used together with 24 DDR4 DIMMs, offering up to 15 TB memory capacity (configured with 512 GB DCPMMs and 128 GB DDR4 DIMMs) to meet various workload requirements.
- 15 PCIe slots, providing excellent scalability.
- Two GE and two 10GE LAN on motherboard (LOM) ports with streamlined configuration, meeting networking requirements of 98% scenarios.
- Supports up to 25 x 2.5-inch local hard drives (configurable with 8 NVMe SSDs).



Smart Power Saving and Better Energy Efficiency

- Leverages patented DEMA, and multiple power-saving measures such as component hibernation, proportional-integral-derivative (PID) algorithm based fan speed tuning, and active-standby power supplies, driving down overall equipment power consumption by up to 15% without compromising workload performance.
- Fitted with 80 Plus® Platinum power supply units (PSUs), up to 94% conversion efficiency; complies with ENERGY STAR standards and has passed the China Energy Conservation and Environmentally-Friendly Certification.
- PSUs with 900 W, 1200 W, 1500 W, and more power options to flexibly adapt to different power requirements, improving energy utilization



Unmatched Intelligent Manageability, Integration, and Openness

- Uses patented intelligent Fault Diagnosis & Management (FDM) technology, delivering up to 93% accuracy in diagnosing core component faults.
- Integrates eSight or FusionDirector for smart entire-lifecycle O&M, driving a leap in deployment and O&M efficiency.
 - » Supports batch OS installation, slashing the average OS installation time of each server from hours to minutes.
 - » Supports automated firmware upgrade, with flexible and configurable upgrade policies for different components and drivers.
 - » Supports stateless computing, allowing for rapid replication of live-network configuration and swift failover.
- Integrates a touchscreen LCD panel for fault diagnosis, allowing O&M personnel to quickly locate faults (supported only by the 5885H V5 8-drive models).
- Integrates fault diagnosis LEDs on the front panel to display error codes in real time, enabling maintenance personnel to rapidly locate a fault.
- Provides standardized open interfaces and development guides, facilitating seamless integration with third-party management software.

Huawei FusionServer Pro 5885H V5 Server



Form Factor	4U rack server
Processors	2 or 4 1st Generation Intel® Xeon® Scalable processors (5100/6100/8100 series), up to 205 W 2 or 4 2nd Generation Intel® Xeon® Scalable processors (5200/6200/8200 series), up to 205 W
Chipset	Intel C622
Memory	48 DDR4 DIMM slots, 2933 MT/s; up to 24 DCPMMs, 2666 MT/s
Internal Storage	Supports hot-swappable hard drives with the following configuration options: <ul style="list-style-type: none">• 8 x 2.5-inch SAS/SATA hard drives (front)• 24 x 2.5-inch SAS/SATA hard drives (front)• 25 x 2.5-inch SAS/SATA hard drives (front)• 16 x 2.5-inch SAS/SATA hard drives and 8 x 2.5-inch NVMe SSDs (front) Flash storage: <ul style="list-style-type: none">• Two M.2 SSDs
RAID	RAID 0, 1, 5, 50, 6, or 60; optional supercapacitor to protect cache data from power failures; RAID level migration, drive roaming, lself-diagnosis, and web-based remote configuration
Network Ports	LAN on motherboard (LOM): 2 x 10GE + 2 x GE ports
PCIe Expansion	Up to 15 PCIe slots for 15 PCIe 3.0 slots 2 dual-width full-height full-length GPU cards (PCIe 3.0 x16) or 4 full-height half-length GPU cards (PCIe 3.0 x16)
Fan Modules	5 hot-swappable counter-rotating fan modules with optional N+1 redundancy
Power Supply	4 hot-swappable PSUs with optional 2+2 redundancy. Supported options include: <ul style="list-style-type: none">• 1500 W AC Platinum PSUs 1000 W (input: 100 V to 127 V AC) 1500 W (input: 200 V to 240 V AC, or 192 V to 288 V DC)• 900 W AC PSUs (input: 100 V to 240 V AC, or 192 V to 288 V DC)• 1200 W DC PSUs (input: -38.4 V to -72 V DC)
Management	<ul style="list-style-type: none">• Provides management features such as fault diagnosis, dynamic energy management technology (DEMT), and hardware security hardening based on Huawei iBMC chips; provides mainstream interfaces, such as Redfish interfaces, enabling easy integration.• Optionally configured with the Huawei eSight or FusionDirector management software to provide advanced management features such as batch OS deployment and automated firmware upgrade, enabling automated entire-lifecycle management.
Operating Systems	<ul style="list-style-type: none">• Microsoft Windows Server• Red Hat Enterprise Linux• SUSE Linux Enterprise Server• CentOS• Citrix XenServer• VMware ESXi For details, see http://support.huawei.com/online/toolsweb/ftca/index?serie=2 .
Security	<ul style="list-style-type: none">• Power-on password• Administrator password• Trusted Platform Module (TPM)• Security front panel
Operating Temperature	5°C to 45°C (41°F to 113°F) (ASHRAE A3 and A4 compliant)
Certification	CE, UL, FCC, CCC, and RoHS
Installation Suite	L-shaped guide rails, adjustable guide rails, and holding rails
Dimensions (H x W x D)	175 mm x 447 mm x 790 mm (6.89 in. x 17.60 in. x 31.10 in.)

*Last update: July 16, 2019

For more information

To learn more about Huawei servers, contact Huawei sales representatives or business partners, or visit:
<http://e.huawei.com/en/products/cloud-computing-dc/servers>

Copyright © Huawei Technologies Co., Ltd. 2019. All rights reserved.
THIS DOCUMENT IS FOR INFORMATION PURPOSE ONLY, AND DOES NOT
CONSTITUTE ANY KIND OF WARRANTIES.



Scan for an
electronic copy



Scan to learn more about
Huawei servers


Copyright © Huawei Technologies Co., Ltd. 2019.

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.

Trademark Notice



HUAWEI, and  are trademarks or registered trademarks of Huawei Technologies Co., Ltd.

Other trademarks, product, service and company names mentioned are the property of their respective owners.

General Disclaimer

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.

HUAWEI TECHNOLOGIES CO., LTD.

Huawei Industrial Base
Bantian Longgang
Shenzhen 518129, P.R.China
Tel: +86-755-28780808

www.huawei.com

Why Huawei servers?

Huawei is a world-leading server provider with a broad spectrum of server offerings including rack, high-density, blade servers and KunLun Mission Critical Servers. Huawei is the industry's only vendor that has the integrated capabilities of server R&D, manufacture, and delivery. Huawei servers have been recognized for their superior quality, rock-solid reliability, extraordinary performance, ease of management, energy efficiency, and security. Huawei servers have served over 5,000 customer accounts across various industries around the globe, including government, finance, electric power, Internet, telecom, energy, transportation, and education.



JERASOFT



Billing Solutions for Telecom Industry

About Us

JeraSoft is a passionate team of engineers focused on development and integration of highly productive Billing Solutions for Telecom (Mobile VoIP, Business Telephony, OTT, Wholesale VoIP and SMS) and IoT/M2M Providers around the world.

Since 2005 JeraSoft has been putting together key insight and successful billing management practices, combining them with state-of-the-art web and telecom technologies. The result is a powerful platform the Telecom Providers can use to streamline Billing Processes and add value to their business.

Why JeraSoft?



Providing Stable Billing Solution for 12 Years



Industry Leading 24/7 Support Team



Serving 200+ Clients in 35 Countries



100% Integration with Third Party Systems



Configured Specially for Your Business



Full Training and Dedicated Initial Help



Customized License Options and Components



New Feature-Rich Releases Throughout the Year


Retail VCS

Billing Platform for Mobile & Business Telecoms, MVNOs and OTT Providers

- ✔ Powerful Subscription Plans
- ✔ DIDs Management
- ✔ Discounts and Taxes Management
- ✔ Flexible Billing Schemas
- ✔ Extended API and Provisioning



20m Subscribers
per system



1 bil Billed Units
per month



250 RPS
per node

Wholesale VCS

Billing Platform for Wholesale VoIP and SMS Business

- ✔ Powerful Rates Management
- ✔ Easy Real-time Statistics
- ✔ Flexible Dynamic Routing and LCR
- ✔ Multi-level Resellers and Agents Tree
- ✔ Traffic Quality Control (ACD, ASR, PDD)



800 Mln Minutes
per month



38k Calls
per node



250 CPS
per node

Care & Support

Your 24/7 Support Team

From simple questions to complicated technical issues - our Support Team will take care of you any time of the day or night!

Multi-channel Support Center

JeraSoft offers live trainings, videos, step-by-step manuals, tutorials, best practices and how-to articles to master the platform, while our support team is always there to help

Live Trainings & Online Resources

Comfortable communication environment based on your preferences - contact our Support Engineers via Phone, Chat or Ticketing System

"The JeraSoft team works accurately and cohesively, and they never leave any of our questions or tasks in "non-reply" status. We are a really satisfied customer!"

Nik Liepins
Bistalk Telecom AG

Clients

With over 200 strong family of clients, we serve a wide range of small and large businesses over 12 years. Here are a few of them:





JERASOFT

Let us help you make the right
choice for your Billing needs

UK +44 (203) 129-9126

US +1 (415) 520-7883

info@jerasoft.net

www.jerasoft.net



MTG2000 Trunk Gateway

MTG2000 is a new-generation intelligent VoIP gateway, which is designed for enterprises, telecom operators and various industries. Focusing on a concept of maintainable, manageable and operable, MTG2000 features high integration and large capacity. It provides carrier-grade VoIP and FoIP services, as well as value-added functions such as modem and voice recognition. Thus it constructs a flexible, high-efficient, future-oriented communication network for users.

MTG2000 supports a range of signaling protocols, realizing the interconnection between SIP and traditional signals like SS7 and PRI. It supports multiple codec methods, offers signal encryption technology and smart voice recognition technology, and improves the utilizing efficiency of trucking resources while ensuring voice quality. The trunk gateway is ideally fit for various access networks of SMEs, call centers, telecom operators and large-scale enterprises.

Product Picture



Key Features

Multi-port and high-integrated structure; up to 20 E1/T1 with 1U size

Provide various services, such as VoIP, FoIP, Modem and POS

Support flexible dialing rules and operations, allowing users to customize dialing rules according to different working environments

Support multiple coding standards: G.711A/U, G.723.1, G.729A/B and iLBC

High compatibility, interoperable with PBX of Avaya, NEC and Alcatel, and also leading soft-switch of Huawei, Cisco and ZTE etc.

Honesty Concentration

Physical Interfaces

- **E1/T1 Ports**
4/8/12/16/20 E1/T1
- **DTU Module :**
4 E1/T1
- **Interface Type**
RJ48(Impedance 120Ω)
- **Ethernet Interface**
GE1: 10/100/1000 BaseT Adaptive Ethernet
GE0: 10/100/1000 BaseT Adaptive Ethernet
- **Serial Port**
1* RS232, 115200bps

Voice Capabilities

- Codecs:G.711a/μ law,G.723.1, G.729A/B, iLBC, AMR
- Silence Suppression
- Comfort Noise
- Voice Activity Detection
- Echo Cancellation (G.168),with up to 128ms
- Adaptive Dynamic Buffer
- Voice ,Fax Gain Control
- FAX:T.38 and Pass-through
- Support Modem/POS
- DTMF Mode: RFC2833/Signal/Inband
- Clear Channel/Clear Mode

PSTN

- ISDN PRI
23B+D(T1),30B+D(E1),NT or TE
ITU-T Q.921, ITU-T Q.931, Q.Sig
- Signal 7/SS7
ITU-T, ANSI , ITU-CHINA
MTP1/MTP2/MTP3, TUP/ISUP
- E1 Frame Type : DF,CRC-4,CRC_ITU
- T1 Frame Type :
4-Frame Multi-frame (F4,FT),2-Frame Multi-frame (F12, D3/4),Extended Super-frame (F24, ESF) ,Remote Switch Mode (F72, SLC96)
- Line Codes:
E1:NRZ,CMI,AMI,HDB3
T1:NRZ,CMI,AMI,B8ZS
- Clock
Local/Remote Clock Source

VoIP Protocol

- SIP v2.0 (UDP/TCP),RFC3261
SDP,RTP(RFC2833), RFC3262, 3263,3264,3265,3515,2976,3311
- RTP/RTCP, RFC2198, 1889
- SIP-T,RFC3372, RFC3204, RFC3398
- SIP Trunk Work Mode :Peer/Access
- SIP/IMS Registration :with up to 256 SIP Accounts
- NAT: Dynamic NAT, Rport

Innovation Open-mind

Software Features

- Local/Transparent Ring Back Tone
- Overlapping Dialing
- Dialing Rules , with up to 2000
- PSTN group by E1 port or E1 Timeslot
- IP Trunk Group Configuration
- Voice Codecs Group
- Caller and Called Number White Lists
- Caller and Called Number Black Lists
- Access Rule Lists
- IP Trunk Priority
- RTP and Signaling Encryption (VOS RC4)

Call Features

- Flexible Route Methods
PSTN-PSTN, PSTN-IP, IP-IP, IP-PSTN
- Intelligent Routing Rules
- Call Routing base on Time
- Call Routing base on Caller/Called Prefixes
- 256 Route Rules for each Direction
- Caller and Called Number Manipulation

Maintenance

- Web GUI Configuration
- Data Backup/Restore
- PSTN Call Statistics
- SIP Trunk Call Statistics
- Firmware Upgrade via TFTP/FTP/Web
- Network Capture
- SNMP v2
- Syslog: Debug, Info, Error, Warning , Notice
- Call History Records via Syslog
- NTP Synchronization
- Centralized Management System

Environmental

- Redundant Power
- Power Supply: 100-240VAC, 50-60 Hz
- Power Consumption:45W
- Operating Temperature:0 °C ~ 45 °C
Storage Temperature: -20 °C ~80 °C
- Humidity:10%-90% Non-Condensing
- Dimensions(W/D/H): 436*300*44.5mm(1U)
- Unit Weight: 3.8kg
- Compliance: CE, FCC

Shenzhen Dinstar Technologies Co., Ltd.

9th Floor, Guoxing Building, Changxing Road, Nanshan District, Shenzhen,
Guangdong, P.R.China, 518052

Tel.: +86 755 6191 9966 FAX: +86 755 2645 6659

Email: sales@dinstar.com support@dinstar.com

Website: www.dinstar.com



the new breed of communication engine

[About](#) | [Downloads](#) | [Documentation](#) | [Community](#) | [Development](#) | [Training](#)

About -> About

1. What is OpenSIPS

OpenSIPS (Open SIP Server) is a mature *Open Source* implementation of a SIP server. OpenSIPS is more than a SIP proxy/router as it includes application-level functionalities. OpenSIPS, as a SIP server, is the core component of any SIP-based VoIP solution. With a very flexible and customizable routing engine, OpenSIPS unifies voice, video, IM and presence services in a highly efficient way, thanks to its scalable (modular) design.

What OpenSIPS has to offer, comes in a reliable and high-performance flavour - OpenSIPS is one of the fastest SIP servers, with a throughput that confirms it as a solution up to enterprise or carrier-grade class.

2. Why OpenSIPS ?

2.1 Performance

What OpenSIPS has to offer, comes in a reliable and high-performance flavour - OpenSIPS is probably the fastest SIP server, with a throughput that confirms it as a solution up to enterprise or carrier-grade class.

Over the course of time, many tests (benchmarks) were done, resulting in a dense collection of various performance tests and measurements (database interaction, memory/cpu usage, etc.) OpenSIPS. The collection of benchmarks can be accessed [here](#), and it showcases OpenSIPS as one of the fastest SIP Proxies on the market.

2.2 Ease of Use

OpenSIPS is an easy to deploy software. It has :

- A well documented, per OpenSIPS version, manual available [here](#), guiding you through all the steps needed for installation, configuration and OpenSIPS scripting
- A graphical, ncurses based, console interface (called menuconfig) which can be used by the user for configuring OpenSIPS compile time options, for including modules with dependencies for compilation and also for generating OpenSIPS scripts with pre-configured capabilities (currently supported are Residential-type script, Trunking and Load-balancing)
- A series of recorded [webinars](#), ranging from the very basics like compiling and installing OpenSIPS, going to more advanced topics like SIP routing and OpenSIPS scripting.

2.3 Flexibility

Due to its very flexible custom scripting language (quite similar in syntax with the C language) and to its modular architecture, OpenSIPS can be easily used in a very wide range of scenarios. You can plug in the needed modules for your desired functionalities and build on top of them by using OpenSIPS's scripting language.

2.4 Monitoring & Control

OpenSIPS can be easily monitored and controlled from the outside, thanks to the MI interface, Event Interface and Statistics interface. Via the MI interface combined with the Stat interface, the administrator can pull various live stats (like consumed memory, UDP queue size, overall load, number of ongoing calls, etc) that can be used to monitor the status of an OpenSIPS instance. Via the Event Interface, the administrator has the ability to get notified by OpenSIPS when certain events get triggered on the proxy side (like DB access is slow, a gateway is down, the MI interface is filled up, etc) and thus can take quick action when such an event happens.

2.5 WEB Interface

OpenSIPS also has a Web Control Panel Application, [OpenSIPS-CP](#). The interface can be used for provisioning OpenSIPS, but it can also display and monitor various statistics, as well as perform actions on OpenSIPS (like instructing OpenSIPS to reload some database tables, etc).

3. OpenSIPS Features

OpenSIPS has to offer many important and interesting features. To mention some of the most important ones:

- SIP registrar server
- SIP router / proxy (lcr, dynamic routing, dialplan features)
- SIP redirect server
- SIP presence agent
- SIP back-to-back User Agent
- SIP IM server (chat and end-2-end IM)
- SIP to SMS gateway (bidirectional)
- SIP to XMPP gateway for presence and IM (bidirectional)
- SIP load-balancer or dispatcher
- SIP front end for gateways/asterisk
- SIP NAT traversal unit
- SIP application server

4. Pre-requisites

Since OpenSIPS is known to work on most *nix based systems, installing and running OpenSIPS requires some very basic Linux knowledge.

Also, since OpenSIPS is a SIP proxy, you will most likely need some basic SIP knowledge in order to deploy / debug OpenSIPS in some more advanced scenarios.

Last but not least, you will be needing some basic programming logic knowledge - The OpenSIPS [configuration file](#) is text-based, written in an OpenSIPS custom language, very similar to C.

5. Getting Started with OpenSIPS

If you want to get started building platforms/solutions with OpenSIPS, here are a few steps that should help :

- Go through the [Getting Started](#) tutorial, which will show you how to get a fully functional platform in a matter of minutes
- Walk your way through the [OpenSIPS Manual](#) for your version, which will guide you through Configuring and using the OpenSIPS Scripting modules in order to achieve your desired functionalities
- If you prefer the classic book versus online material, see [Building Telephony Systems with OpenSIPS 1.6](#)
- If you have a question / issue / bug to report, please see our [Support Page](#)

Table of Content (hide)

- 1. What is OpenSIPS
- 2. Why OpenSIPS ?
 - 2.1 Performance
 - 2.2 Ease of Use
 - 2.3 Flexibility
 - 2.4 Monitoring & Control
 - 2.5 WEB Interface
- 3. OpenSIPS Features
- 4. Pre-requisites
- 5. Getting Started with OpenSIPS

Page last modified on May 25, 2016, at 12:36 PM



Datasheet

Proxmox Virtual Environment 6

AT A GLANCE

- Complete open-source platform for enterprise virtualization
- Easy management of compute, network, and storage with the central web interface
- 100% software-defined architecture
- Two virtualization technologies supported: KVM hypervisor & Linux Container (LXC)
- Web-based UI, CLI, API
- High-Availability (HA) Cluster Manager
- Live Migration
- Built-in services: firewall, backup/restore, storage replication, etc.
- Open-source license GNU AGPL, v3
- Enterprise support agreements

OVERVIEW

Proxmox VE is a complete open-source solution for enterprise virtualization that tightly integrates KVM hypervisor and LXC containers, software-defined storage and networking functionality on a single platform. With the central user interface you can easily run VMs and containers, manage software-defined storage resources and networking functionality, highly available clusters, and multiple out-of-the-box tools like backup/restore, live migration, storage replication, or the integrated firewall. Proxmox VE enables you to virtualize even the most demanding Linux and Windows application workloads.

By combining two virtualization technologies on a single platform, Proxmox VE provides maximum flexibility to your data center. It includes strong high-availability (HA) support and—thanks to the unique multi-master design—you don't need an additional management server thus saving resources and allowing HA without single point of failure (SPOF).

ENTERPRISE-READY

Enterprises use the powerful Proxmox VE platform to easily install, manage and monitor hyperconverged infrastructures (HCI) in their data center. Multiple authentication sources combined with role-based user and permission management enable full control of the HA clusters. The REST API enables easy integration of third-party management tools like for example custom hosting environments.

The future-proof and open-source development model of Proxmox VE guarantees full access to the products source code as well as maximum flexibility and security.

KEY FEATURES

INDUSTRY-LEADING

ENTERPRISE VIRTUALIZATION

- Linux and Windows servers, 32 and 64 bit operation systems.
- Support for the latest Intel and AMD server chip sets – for great VM performance.
- Performance relative to bare-metal for real-world enterprise workloads.
- Management layer containing all the capabilities to manage and monitor an open-source software-defined data center.

OPEN-SOURCE SOFTWARE

- Published under the free and open-source GNU Affero General Public License, version 3 (AGPL, V3: <http://www.gnu.org/licenses/agpl-3.0.html>)
- Designed to ensure cooperation with community
- Public code repository (Git)
- Bugtracker for issue tracking
- Community support forum
- Documentation, Wiki, video tutorials, how-tos,...

ENTERPRISE SUPPORT AGREEMENT

- Subscriptions to ensure business continuity.
- Exclusive access to the stable Enterprise Repository.
- Updates and version upgrades via GUI.
- Professional support from the Proxmox developers.

HIGH-AVAILABILITY CLUSTER

- No single point of failure (no SPOF).
- Multi-master cluster.
- Manage the HA settings for KVM and LXC via GUI.
- pmxcfs—unique Proxmox VE Cluster File System: database-driven file system for storing configuration files replicated in real-time on all nodes using Corosync.
- Based on proven Linux HA technologies, providing stable and reliable HA service.
- Resource agents for KVM and containers (LXC).
- Watchdog-based fencing.

FENCING

- The Proxmox VE HA Manager uses self fencing provided by hardware Watchdog or kernel Softdog.
- No simultaneous data access or corruption.
- Works „out-of-the-box“.

- Included Proxmox VE HA Simulator for testing.

WEB-BASED MANAGEMENT INTERFACE

- Integrated - no need to install a separate management tool or any additional management node.
- Fast, search-driven interface, able to handle thousands of VMs.
- Based on the Ext JS JavaScript framework.
- Secure HTML5 console, supporting SSL.
- Fast and easy creation of VMs and containers.
- Seamless integration and easy management of a whole cluster.
- Subscription management via GUI.
- AJAX technologies for dynamic updates of resources.

COMMAND LINE (CLI)

- Manage all components of your virtual environment.
- CLI with intelligent tab completion.
- Full UNIX man page documentation.

REST API

- Easy integration for third-party management tools.
- REST like API (JSON as primary data format).
- Easy and human readable data format (native web browser format).
- Automatic parameter verification (verification of return values).
- Automatic generation of the API documentation.
- Easy way to create command line tools (use the same API).
- Resource Oriented Architecture (ROA).
- Declarative API definition using JSON Schema.

SOFTWARE-DEFINED STORAGE

- Local storage such as ZFS (encryption possible), LVM, LVMthin, ext4, and XFS.
- Shared storage such as FC, iSCSI or NFS.
- Distributed storage such as Ceph RBD, CephFS or GlusterFS
- Encryption support for Ceph OSD and ZFS
- Unlimited number of storage definitions (cluster-wide).

KEY FEATURES

LIVE MIGRATION

- Moving QEMU virtual servers from one physical host to another with zero downtime

STORAGE REPLICATION STACK (ZFS)

- Built-in open-source storage replication framework.
- Redundancy for guests using local storage.
- Data availability without using shared storage.
- Asynchronous replication.
- Minimize data loss in the case of a failure.
- Improve reliability, fault-tolerance, and accessibility of your data.

HYPER-CONVERGED INFRASTRUCTURE (HCI) WITH CEPH

- Integrated Ceph, a distributed object store and file system.
- Run Ceph RBD and CephFS directly on the Proxmox VE cluster nodes and manage Ceph via GUI.
- Easy-to-use installation wizard for Ceph.
- Proxmox delivers its own Ceph packages.
- Ceph support is included in the support agreement.

VIRTUALIZED NETWORKING

- Bridged networking model
- Each host with up to 4094 bridges
- TCP/IP configuration
- IPv4 and IPv6 support
- VLANs
- Open vSwitch

PROXMOX VE FIREWALL

- Supporting IPv4 and IPv6.
- Linux-based netfilter technology. Stateful firewall, provides high bandwidth.
- Distributed: main configuration in Proxmox VE cluster file system, iptable rules are stored in nodes.
- Cluster-wide settings
- 3 levels of configuration (data center, host, VM/CT)

BACKUP AND RESTORE

- Full backups of VMs and containers
- Live snapshot backups
- Multiple schedules and backup storage
- GUI integrations, but also via CLI
- "Backup Now" and restore via GUI
- All jobs from all nodes can be monitored via the GUI tab "Tasks"

TWO-FACTOR AUTHENTICATION

- Providing high security.
- 2 types: Time-based One Time Passwords (TOTP) and YubiKey.

MULTIPLE AUTHENTICATION SOURCES

- Proxmox VE supports multiple authentication sources.
- Linux PAM standard authentication (e.g., 'root' and other local users)
- Built-in Proxmox VE authentication server
- Microsoft Active Directory (MS ADS)
- LDAP

ROLE-BASED ADMINISTRATION

- User and permission management for all objects (VMs, storage systems, nodes, etc.)
- Proxmox VE comes with a number of predefined roles (lists of privileges) which satisfies most needs. The GUI provides an overview the whole set of predefined roles.
- Permissions to control access to objects (access control lists). In technical terms they are simply a triple containing <path,user,role>. Each permission specifies a subject (user or group) and a role (set of privileges) on a specific path.

VM TEMPLATES AND CLONES

- Deploying VMs from templates is blazing fast, very comfortable, and if you use linked clones you can optimize your storage by using base images and thin-provisioning.

„Exactly what we needed: High availability features, ease of use and proper integration of a shared storage technology (Ceph) without the need to set up separate storage systems.“

Martin Gollowitz,
Voluntary System Administrator at FSFE

LEARN MORE

Wiki: <https://pve.proxmox.com>

Community Forums:

<https://forum.proxmox.com/>

Bugtracker: <https://bugzilla.proxmox.com>

Code repository: <https://git.proxmox.com>

HOW TO BUY

Visit the Proxmox Online Shop to purchase a subscription: <https://shop.maurer-it.com>

Find an authorized reseller in your area:

www.proxmox.com/partners

SALES INQUIRIES

office@proxmox.com

HELP AND SUPPORT

Proxmox Customer Portal:

<https://my.proxmox.com>

Support Forum: <https://forum.proxmox.com>


TRAINING PROXMOX VE

Learn Proxmox VE easily – Visit a training:

<https://www.proxmox.com/training>

ABOUT PROXMOX

Proxmox Server Solutions GmbH is a software provider dedicated to develop powerful and efficient open-source server solutions. The privately held company is based in Vienna (Europe).



Proxmox Server Solutions GmbH
Bräuhausgasse 37
1050 Vienna
Austria

office@proxmox.com
<https://www.proxmox.com>

RTPENGINE MODULE

Maxim Sobolev

Sippy Software, Inc.

Juha Heinanen

TuTPro, Inc.

Edited by

Maxim Sobolev

Edited by

Bogdan-Andrei Iancu

Edited by

Juha Heinanen

Edited by

Sas Ovidiu

Edited by

Carsten Bock

ng-voice GmbH

Edited by

Richard Fuchs

Sipwise GmbH

Copyright © 2003-2008 Sippy Software, Inc.

Copyright © 2005 Voice Sistem SRL

Copyright © 2009-2014 TuTPro Inc.

Copyright © 2010 VoIPEmbedded Inc.

Copyright © 2013-2014 Sipwise GmbH

Table of Contents**1. Admin Guide****1.1. Overview****1.2. Multiple RTP proxy usage****1.3. Dependencies****1.3.1. OpenSIPS Modules****1.3.2. External Libraries or Applications****1.4. Parameters****1.4.1. rtpengine_sock (string)****1.4.2. rtpengine_disable_tout (integer)****1.4.3. rtpengine_tout (integer)****1.4.4. rtpengine_retr (integer)****1.4.5. extra_id_pv (string)****1.4.6. setid_avp (string)****1.5. Functions****1.5.1. rtpengine_use_set(setid)****1.5.2. rtpengine_offer([flags])****1.5.3. rtpengine_answer([flags])****1.5.4. rtpengine_delete([flags])****1.5.5. rtpengine_manage([flags])****1.5.6. rtpengine_start_recording()****1.6. Exported Pseudo Variables****1.6.1. \$rtpstat****1.7. MI Commands**

1.7.1. [rtengine_enable](#)

1.7.2. [rtengine_show](#)

2. Frequently Asked Questions

List of Examples

- 1.1. [Set rtengine_sock parameter](#)
- 1.2. [Set rtengine_disable_tout parameter](#)
- 1.3. [Set rtengine_tout parameter](#)
- 1.4. [Set rtengine_retr parameter](#)
- 1.5. [Set extra_id_pv parameter](#)
- 1.6. [Set setid_avp parameter](#)
- 1.7. [rtengine_use_set usage](#)
- 1.8. [rtengine_offer usage](#)
- 1.9. [rtengine_answer usage](#)
- 1.10. [rtengine_delete usage](#)
- 1.11. [rtengine_manage usage](#)
- 1.12. [rtengine_start_recording usage](#)
- 1.13. [\\$rtostat Usage](#)
- 1.14. [rtengine_enable usage](#)
- 1.15. [rtengine_show usage](#)

Chapter 1. Admin Guide

1.1. Overview

This is a module that enables media streams to be proxied via an RTP proxy. The only RTP proxy currently known to work with this module is the Sipwise rtengine <https://github.com/sipwise/rtengine>. The rtengine module is a modified version of the original rtpproxy module using a new control protocol. The module is designed to be a drop-in replacement for the old module from a configuration file point of view, however due to the incompatible control protocol, it only works with RTP proxies which specifically support it.

1.2. Multiple RTP proxy usage

The rtengine module can support multiple RTP proxies for balancing/distribution and control/selection purposes.

The module allows definition of several sets of rtengines. Load-balancing will be performed over a set and the admin has the ability to choose what set should be used. The set is selected via its id - the id being defined with the set. Refer to the "rtengine_sock" module parameter definition for syntax description.

The balancing inside a set is done automatically by the module based on the weight of each RTP proxy from the set.

The selection of the set is done from script prior using rtengine_delete(), rtengine_offer() or rtengine_answer() functions - see the rtengine_use_set() function.

Another way to select the set is to define setid_avp module parameter and assign setid to the defined avp before calling rtengine_offer() or rtengine_manage() function. If forwarding of the requests fails and there is another branch to try, remember to unset the avp after calling rtengine_delete() function.

For backward compatibility reasons, a set with no id take by default the id 0. Also if no set is explicitly set before rtengine_delete(), rtengine_offer() or rtengine_answer() the 0 id set will be used.

IMPORTANT: if you use multiple sets, take care and use the same set for both rtengine_offer()/rtengine_answer() and rtengine_delete()!! If the set was selected using setid_avp, the avp needs to be set only once before rtengine_offer() or rtengine_manage() call.

1.3. Dependencies

1.3.1. OpenSIPS Modules

The following modules must be loaded before this module:

- *tm module* - (optional) if you want to have rtengine_manage() fully functional

1.3.2. External Libraries or Applications

The following libraries or applications must be installed before running OpenSIPS with this module loaded:

- *None.*

1.4. Parameters

1.4.1. rtengine_sock (string)

Definition of socket(s) used to connect to (a set) RTP proxy. It may specify a UNIX socket or an IPv4/IPv6 UDP socket.

Default value is "NONE" (disabled).

Example 1.1. Set rtengine_sock parameter

```
...
# single rtproxy
modparam("rtengine", "rtengine_sock", "udp:localhost:12221")
```

```
# multiple rtproxies for LB
modparam("rtpengine", "rtpengine_sock",
         "udp:localhost:12221 udp:localhost:12222")
# multiple sets of multiple rtproxies
modparam("rtpengine", "rtpengine_sock",
         "1 == udp:localhost:12221 udp:localhost:12222")
modparam("rtpengine", "rtpengine_sock",
         "2 == udp:localhost:12225")
...
```

1.4.2. rtpengine_disable_tout (integer)

Once an RTP proxy was found unreachable and marked as disabled, the rtpengine module will not attempt to establish communication to that RTP proxy for rtpengine_disable_tout seconds.

Default value is "60".

Example 1.2. Set rtpengine_disable_tout parameter

```
...
modparam("rtpengine", "rtpengine_disable_tout", 20)
...
```

1.4.3. rtpengine_tout (integer)

Timeout value in waiting for reply from RTP proxy.

Default value is "1".

Example 1.3. Set rtpengine_tout parameter

```
...
modparam("rtpengine", "rtpengine_tout", 2)
...
```

1.4.4. rtpengine_retr (integer)

How many times the module should retry to send and receive after timeout was generated.

Default value is "5".

Example 1.4. Set rtpengine_retr parameter

```
...
modparam("rtpengine", "rtpengine_retr", 2)
...
```

1.4.5. extra_id_pv (string)

The parameter sets the PV definition to use when the "b" parameter is used on rtpengine_delete(), rtpengine_offer(), rtpengine_answer() or rtpengine_manage() command.

Default is empty, the "b" parameter may not be used then.

Example 1.5. Set extra_id_pv parameter

```
...
modparam("rtpengine", "extra_id_pv", "$avp(extra_id)")
...
```

1.4.6. setid_avp (string)

The parameter defines an AVP that, if set, determines which RTP proxy set rtpengine_offer(), rtpengine_answer(), rtpengine_delete(), and rtpengine_manage() functions use.

There is no default value.

Example 1.6. Set setid_avp parameter

```
...
modparam("rtpengine", "setid_avp", "$avp(setid)")
...
```

1.5. Functions

1.5.1. rtpengine_use_set(setid)

Sets the ID of the RTP proxy set to be used for the next `rtpengine_delete()`, `rtpengine_offer()`, `rtpengine_answer()` or `rtpengine_manage()` command. The parameter can be an integer or a config variable holding an integer.

This function can be used from `REQUEST_ROUTE`, `ONREPLY_ROUTE`, `BRANCH_ROUTE`.

Example 1.7. rtpengine_use_set usage

```
...
rtpengine_use_set("2");
rtpengine_offer();
...
```

1.5.2. rtpengine_offer([flags])

Rewrites SDP body to ensure that media is passed through an RTP proxy. To be invoked on `INVITE` for the cases the SDPs are in `INVITE` and `200 OK` and on `200 OK` when SDPs are in `200 OK` and `ACK`.

Meaning of the parameters is as follows:

- `flags` - flags to turn on some features.

The "flags" string is a list of space-separated items. Each item is either an individual token, or a token in "key=value" format. The possible tokens are described below.

- `via-branch=...` - Include the "branch" value of one of the "Via" headers in the request to the RTP proxy. Possible values are: "1" - use the first "Via" header; "2" - use the second "Via" header; "auto" - use the first "Via" header if this is a request, or the second one if this is a reply; "extra" - don't take the value from a header, but instead use the value of the "extra_id_pv" variable. This can be used to create one media session per branch on the RTP proxy. When sending a subsequent "delete" command to the RTP proxy, you can then stop just the session for a specific branch when passing the flag '1' or '2' in the "rtpengine_delete", or stop all sessions for a call when not passing one of those two flags there. This is especially useful if you have serially forked call scenarios where the RTP proxy gets an "offer" command for a new branch, and then a "delete" command for the previous branch, which would otherwise delete the full call, breaking the subsequent "answer" for the new branch. *This flag is only supported by the Sipwise rtpengine RTP proxy at the moment!*
- `asymmetric` - flags that UA from which message is received doesn't support symmetric RTP. (automatically sets the 'r' flag)
- `force-answer` - force "answer", that is, only rewrite SDP when corresponding session already exists in the RTP proxy. By default is on when the session is to be completed.
- `internal`, `external` - these flags specify the direction of the SIP message. These flags only make sense when the RTP proxy is running in bridge mode. "internal" corresponds to the proxy's first interface, "external" corresponds to the RTP proxy's second interface. You always have to specify two flags to define the incoming network and the outgoing network. For example, "internal external" should be used for SIP message received from the local interface and sent out on the external interface, and "external internal" vice versa. Other options are "internal internal" and "external external". So, for example if a SIP requests is processed with "internal external" flags, the corresponding response must be processed with "internal external" flags.
- `auto-bridge` - this flag an alternative to the "internal" and "external" flags in order to do automatic bridging between IPv4 on the "internal network" and IPv6 on the "external network". Instead of explicitly instructing the RTP proxy to select a particular address family, the distinction is done by the given IP in the SDP body by the RTP proxy itself. Not supported by Sipwise rtpengine.
- `address-family=...` - instructs the RTP proxy that the recipient of this SDP body expects to see addresses of a particular family. Possible values are "IP4" and "IP6". For example, if the SDP body contains IPv4 addresses but the recipient only speaks IPv6, you would use "address-family=IP6" to bridge between the two address families.

Sipwise rtpengine remembers the address family preference of each party after it has seen an SDP body from them. This means that normally it is only necessary to explicitly specify the address family in the "offer", but not in the "answer".

Note: Please note, that this will only work properly with non-dual-stack user-agents or with dual-stack clients according to RFC6157 (which suggest ICE for Dual-Stack implementations). This short-cut will not work properly with RFC4091 (ANAT) compatible clients, which suggests having different m-lines with different IP-protocols grouped together.

- `force` - instructs the RTP proxy to ignore marks inserted by another RTP proxy in transit to indicate that the session is already goes through another proxy. Allows creating a chain of proxies. Not supported and ignored by Sipwise rtpengine.
- `trust-address` - flags that IP address in SDP should be trusted. Without this flag, the RTP proxy ignores address in the SDP and uses source address of the SIP message as media address which is passed to the RTP proxy.
- `replace-origin` - flags that IP from the origin description (o=) should be also changed.
- `replace-session-connection` - flags to change the session-level SDP connection (c=) IP if media description also includes connection information.
- `symmetric` - flags that for the UA from which message is received, support symmetric RTP must be forced.
- `repacketize=NN` - requests the RTP proxy to perform re-packetization of RTP traffic coming from the UA which has sent the current message to increase or decrease payload size per each RTP packet forwarded if possible. The NN

is the target payload size in ms, for the most codecs its value should be in 10ms increments, however for some codecs the increment could differ (e.g. 30ms for GSM or 20ms for G.723). The RTP proxy would select the closest value supported by the codec. This feature could be used for significantly reducing bandwidth overhead for low bitrate codecs, for example with G.729 going from 10ms to 100ms saves two thirds of the network bandwidth. Not supported by Sipwise rtppengine.

- *ICE=...* - controls the RTP proxy's behaviour regarding ICE attributes within the SDP body. Possible values are: "force" - discard any ICE attributes already present in the SDP body and then generate and insert new ICE data, leaving itself as the *only* ICE candidates; "remove" instructs the RTP proxy to discard any ICE attributes and not insert any new ones into the SDP. The default (if no "ICE=..." is given at all), new ICE data will only be generated if no ICE was present in the SDP originally; otherwise the RTP proxy will only insert itself as an *additional* ICE candidate. Other SDP substitutions (c=, m=, etc) are unaffected by this flag.
- *RTP, SRTP, AVP, AVPF* - These flags control the RTP transport protocol that should be used towards the recipient of the SDP. If none of them are specified, the protocol given in the SDP is left untouched. Otherwise, the "SRTP" flag indicates that SRTP should be used, while "RTP" indicates that SRTP should not be used. "AVPF" indicates that the advanced RTCP profile with feedback messages should be used, and "AVP" indicates that the regular RTCP profile should be used. See also the next set of flags below.
- *RTP/AVP, RTP/SAVP, RTP/AVPF, RTP/SAVPF* - these serve as an alternative, more explicit way to select between the different RTP protocols and profiles supported by the RTP proxy. For example, giving the flag "RTP/SAVPF" has the same effect as giving the two flags "SRTP AVPF".
- *to-tag* - force inclusion of the "To" tag. Normally, the "To" tag is always included when present, except for "delete" messages. Including the "To" tag in a "delete" messages allows you to be more selective about which dialogues within a call are being torn down.
- *rtcp-mux-demux* - if rtcp-mux (RFC 5761) was offered, make the RTP proxy accept the offer, but not offer it to the recipient of this message.
- *rtcp-mux-reject* - if rtcp-mux was offered, make the RTP proxy reject the offer, but still offer it to the recipient. Can be combined with "rtcp-mux-offer" to always offer it.
- *rtcp-mux-offer* - make the RTP proxy offer rtcp-mux to the recipient of this message, regardless of whether it was offered originally or not.
- *rtcp-mux-accept* - if rtcp-mux was offered, make the RTP proxy accept the offer and also offer it to the recipient of this message. Can be combined with "rtcp-mux-offer" to always offer it.
- *media-address=...* - force a particular media address to be used in the SDP body. Address family is detected automatically.

This function can be used from ANY_ROUTE.

Example 1.8. rtppengine_offer usage

```
route {
...
  if (is_method("INVITE")) {
    if (has_body("application/sdp")) {
      if (rtppengine_offer())
        t_on_reply("1");
    } else {
      t_on_reply("2");
    }
  }
  if (is_method("ACK") && has_body("application/sdp"))
    rtppengine_answer();
...
}

onreply_route[1]
{
...
  if (has_body("application/sdp"))
    rtppengine_answer();
...
}

onreply_route[2]
{
...
  if (has_body("application/sdp"))
    rtppengine_offer();
...
}
```

1.5.3. rtppengine_answer([flags])

Rewrites SDP body to ensure that media is passed through an RTP proxy. To be invoked on 200 OK for the cases the SDPs are in INVITE and 200 OK and on ACK when SDPs are in 200 OK and ACK.

See rtppengine_offer() function description above for the meaning of the parameters.

This function can be used from REQUEST_ROUTE, ONREPLY_ROUTE, FAILURE_ROUTE, BRANCH_ROUTE.

Example 1.9. rtppengine_answer usage

See `rtppengine_offer()` function example above for example.

1.5.4. `rtppengine_delete([flags])`

Tears down the RTPProxy session for the current call.

See `rtppengine_offer()` function description above for the meaning of the parameters. Note that not all flags make sense for a "delete".

This function can be used from ANY_ROUTE.

Example 1.10. `rtppengine_delete` usage

```
...
rtppengine_delete();
...
```

1.5.5. `rtppengine_manage([flags])`

Manage the RTPProxy session - it combines the functionality of `rtppengine_offer()`, `rtppengine_answer()` and `rtppengine_delete()`, detecting internally based on message type and method which one to execute.

It can take the same parameters as `rtppengine_offer()`. The flags parameter to `rtppengine_manage()` can be a configuration variable containing the flags as a string.

Functionality:

- If INVITE with SDP, then do `rtppengine_offer()`
- If INVITE with SDP, when the tm module is loaded, mark transaction with internal flag FL_SDP_BODY to know that the 1xx and 2xx are for `rtppengine_answer()`
- If ACK with SDP, then do `rtppengine_answer()`
- If BYE or CANCEL, or called within a FAILURE_ROUTE[], then do `rtppengine_delete()`
- If reply to INVITE with code >= 300 do `rtppengine_delete()`
- If reply with SDP to INVITE having code 1xx and 2xx, then do `rtppengine_answer()` if the request had SDP or tm is not loaded, otherwise do `rtppengine_offer()`

This function can be used from ANY_ROUTE.

Example 1.11. `rtppengine_manage` usage

```
...
rtppengine_manage();
...
```

1.5.6. `rtppengine_start_recording()`

This function will send a signal to the RTP proxy to record the RTP stream on the RTP proxy. *This function is not supported by Sipwise rtppengine at the moment!*

This function can be used from REQUEST_ROUTE and ONREPLY_ROUTE.

Example 1.12. `rtppengine_start_recording` usage

```
...
rtppengine_start_recording();
...
```

1.6. Exported Pseudo Variables

1.6.1. `$rtpstat`

Returns the RTP statistics from the RTP proxy. The RTP statistics from the RTP proxy are provided as a string and it does contain several packet counters. The statistics must be retrieved before the session is deleted (before `rtppengine_delete()`).

Example 1.13. `$rtpstat` Usage

```
...
append_hf("X-RTP-Statistics: $rtpstat\r\n");
...
```

1.7. MI Commands

1.7.1. rtppengine_enable

Enables a RTP proxy if parameter value is greater than 0. Disables it if a zero value is given.

The first parameter is the RTP proxy url (exactly as defined in the config file).

The second parameter value must be a number in decimal.

NOTE: if a RTP proxy is defined multiple times (in the same or diferente sete), all of its instances will be enables/disabled.

Example 1.14. rtppengine_enable usage

```
...
$ opensslctl fifo rtppengine_enable udp:192.168.2.133:8081 0
...
```

1.7.2. rtppengine_show

Displays all the RTP proxies and their information: set and status (disabled or not, weight and recheck_ticks).

No parameter.

Example 1.15. rtppengine_show usage

```
...
$ opensslctl fifo rtppengine_show
...
```

Chapter 2. Frequently Asked Questions

2.1. How do I migrate from "rtpproxy" or "rtpproxy-ng" to "rtppengine"?

For the most part, only the names of the functions have changed, with "rtpproxy" in each name replaced with "rtppengine". For example, "rtpproxy_manage()" has become "rtppengine_manage()". A few name duplications have also been resolved, for example there is now a single "rtppengine_delete()" instead of "unforce_rtp_proxy()" and the identical "rtpproxy_destroy()".

The largest difference to the old module is how flags are passed to "rtppengine_offer()", "rtppengine_answer()", "rtppengine_manage()" and "rtppengine_delete()". Instead of having a string of single-letter flags, they now take a string of space-separated items, with each item being either a single token (word) or a "key=value" pair.

For example, if you had a call "rtpproxy_offer("FRWOC+PS");", this would then become:

```
rtppengine_offer("force trust-address symmetric replace-origin replace-session-connection ICE=force RTP/SAVPF");
```

Finally, if you were using the second paramater (explicit media address) to any of these functions, this has been replaced by the "media-address=..." option within the first string of flags.

2.2. Where can I find more about OpenSIPS?

Take a look at <http://www.opensips.org/>.

2.3. Where can I post a question about this module?

First at all check if your question was already answered on one of our mailing lists:

- User Mailing List - <http://lists.opensips.org/cgi-bin/mailman/listinfo/users>
- Developer Mailing List - <http://lists.opensips.org/cgi-bin/mailman/listinfo/devel>

E-mails regarding any stable OpenSIPS release should be sent to users@lists.opensips.org and e-mails regarding development versions should be sent to devel@lists.opensips.org.

If you want to keep the mail private, send it to users@lists.opensips.org.

2.4. How can I report a bug?

Please follow the guidelines provided at: <https://github.com/OpenSIPS/opensips/issues>.



Yate stands for **Yet Another Telephony Engine**, and like the name states it is mainly a telephony engine; while currently focused on [Voice over Internet Protocol \(VoIP\)](#) and [PSTN](#), its power lies in its ability to be easily extended. Voice, video, data and instant messaging can all be unified under Yate's flexible routing engine, maximizing communications efficiency and minimizing infrastructure costs for businesses.

Yate usability

- [VoIP Server](#)
- [VoIP client](#)
- [Conference server](#) - with up to 200 channels in a single conference.
- [VoIP to PSTN gateway](#)
- PC2Phone and Phone2PC gateway
- IP Telephony server and/or client
 - [H.323 gatekeeper](#)
 - [H.323 multiple endpoint server](#)
 - [H.323<->SIP Proxy](#)
 - [SIP Router](#)
 - [SIP session border controller](#)
 - [SIP registration server](#)
 - [IAX Server](#)
 - [IAX Client](#)
 - [Jingle client or server](#)
 - MGCP server (Call Agent)
- ISDN passive and active recorder
- ISDN, RBS, analog passive recorder
- SS7 switch
- Call center server
- IVR engine
- Prepaid and/or postpaid cards system

About Yate's software

The software is *written in C++* and it supports scripting in various programming languages (such as those supported by the currently implemented PHP, Python, Perl and Javascript libraries) and even any Unix shell. The

PHP, Python, Perl and Javascript libraries have been developed and made available in order to ease development of external functionalities for Yate.

Yate is production-ready software and is easily extensible.

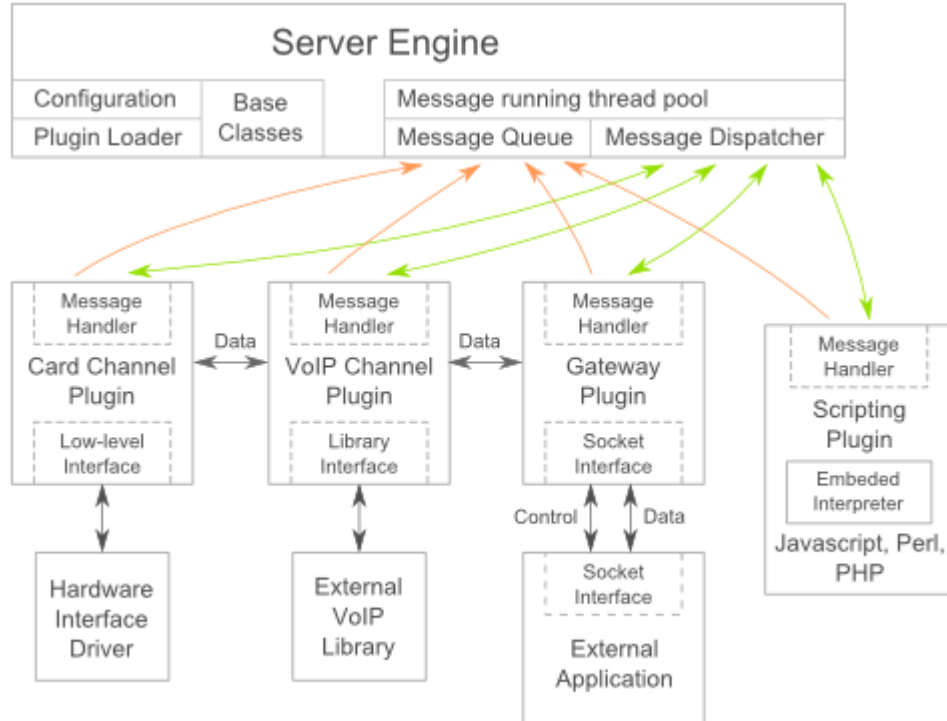
Yate's license

Yate architecture

The most important aspect of Yate is its message-passing system. Modules are passing around messages among them. This allows us to have a bigger flexibility than with plain functions, mainly because messages in Yate can have an arbitrary number of parameters, and can be sent to more than one module by changing the priority.

Yate's four main components are:

- **Core** - generic classes like String, Thread, Socket, Mutex
- **Message Engine** - message related classes Message, Engine, Plugin
- **Telephony Engine** - telephony related classes like Driver, Channel
- **Yate Modules** - modules of Yate are equal no matter if they are telephony or routing or anything else related, because of the message passing system.

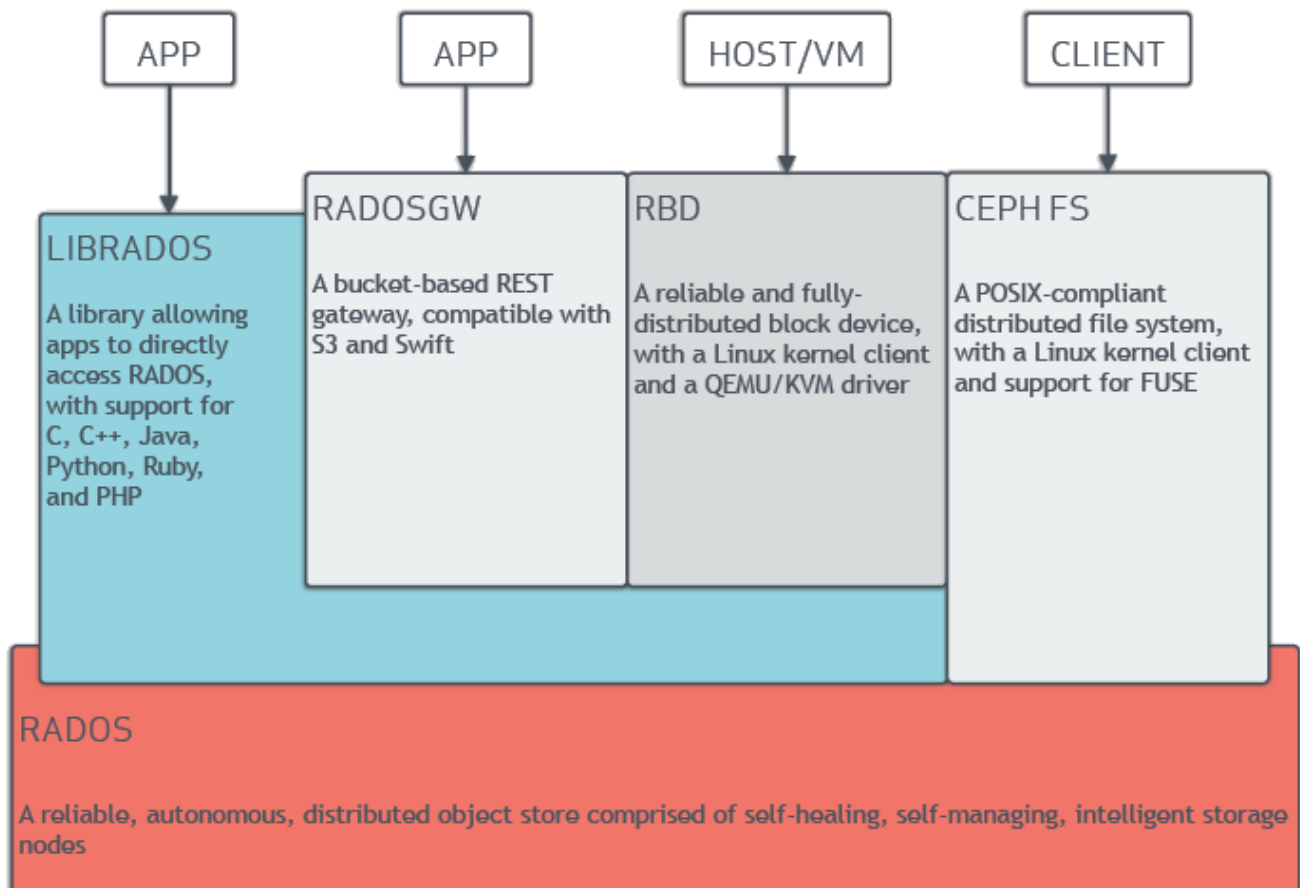


Notice: This document is for a development version of Ceph.

[Edit on GitHub](#) | [Report a Documentation Bug](#)

ARCHITECTURE

Ceph uniquely delivers **object, block, and file storage** in one unified system. Ceph is highly reliable, easy to manage, and free. The power of Ceph can transform your company's IT infrastructure and your ability to manage vast amounts of data. Ceph delivers extraordinary scalability—thousands of clients accessing petabytes to exabytes of data. A **Ceph Node** leverages commodity hardware and intelligent daemons, and a **Ceph Storage Cluster** accommodates large numbers of nodes, which communicate with each other to replicate and redistribute data dynamically.




THE CEPH STORAGE CLUSTER

Ceph provides an infinitely scalable **Ceph Storage Cluster** based upon [RADOS](#), which you can read about in [RADOS - A Scalable, Reliable Storage Service for Petabyte-scale Storage Clusters](#).

A Ceph Storage Cluster consists of two types of daemons:

- **Ceph Monitor**
- **Ceph OSD Daemon**




A Ceph Monitor maintains a master copy of the cluster map. A cluster of Ceph monitors ensures high availability should a monitor daemon fail. Storage cluster clients retrieve a copy of the cluster map from the Ceph Monitor.

A Ceph OSD Daemon checks its own state and the state of other OSDs and reports back to monitors.

Storage cluster clients and each **Ceph OSD Daemon** use the CRUSH algorithm to efficiently compute information about data location, instead of having to depend on a central lookup table. Ceph’s high-level features include providing a native interface to the Ceph Storage Cluster via `librados`, and a number of service interfaces built on top of `librados`.

STORING DATA

The Ceph Storage Cluster receives data from **Ceph Clients**—whether it comes through a **Ceph Block Device**, **Ceph Object Storage**, the **Ceph Filesystem** or a custom implementation you create using `librados`—and it stores the data as objects. Each object corresponds to a file in a filesystem, which is stored on an **Object Storage Device**. Ceph OSD Daemons handle the read/write operations on the storage disks.



Ceph OSD Daemons store all data as objects in a flat namespace (e.g., no hierarchy of directories). An object has an identifier, binary data, and metadata consisting of a set of name/value pairs. The semantics are completely up to **Ceph Clients**. For example, CephFS uses metadata to store file attributes such as the file owner, created date, last modified date, and so forth.

ID	Binary Data	Metadata
1234	0101010101010100110101010010 0101100001010100110101010010 0101100001010100110101010010	name1 value1 name2 value2 nameN valueN

Note: An object ID is unique across the entire cluster, not just the local filesystem.

SCALABILITY AND HIGH AVAILABILITY

In traditional architectures, clients talk to a centralized component (e.g., a gateway, broker, API, facade, etc.), which acts as a single point of entry to a complex subsystem. This imposes a limit to both performance and scalability, while introducing a single point of failure (i.e., if the centralized component goes down, the whole system goes down, too).

Ceph eliminates the centralized gateway to enable clients to interact with Ceph OSD Daemons directly. Ceph OSD Daemons create object replicas on other Ceph Nodes to ensure data safety and high availability. Ceph also uses a cluster of monitors to ensure high availability. To eliminate centralization, Ceph uses an algorithm called CRUSH.

CRUSH INTRODUCTION

Ceph Clients and Ceph OSD Daemons both use the [CRUSH](#) algorithm to efficiently compute information about object location, instead of having to depend on a central lookup table. CRUSH provides a better data management mechanism compared to older approaches, and enables massive scale by cleanly distributing the work to all the clients and OSD daemons in the cluster. CRUSH uses intelligent data replication to ensure resiliency, which is better suited to hyper-scale storage. The following sections provide additional details on how CRUSH works. For a detailed discussion of CRUSH, see [CRUSH - Controlled, Scalable, Decentralized Placement of Replicated Data](#).

CLUSTER MAP

Ceph depends upon Ceph Clients and Ceph OSD Daemons having knowledge of the cluster topology, which is inclusive of 5 maps collectively referred to as the “Cluster Map”:

1. **The Monitor Map:** Contains the cluster `fsid`, the position, name address and port of each monitor. It also indicates the current epoch, when the map was created, and the last time it changed. To view a monitor map, execute `ceph mon dump`.
2. **The OSD Map:** Contains the cluster `fsid`, when the map was created and last modified, a list of pools, replica sizes, PG numbers, a list of OSDs and their status (e.g., `up`, `in`). To view an OSD map, execute `ceph osd dump`.
3. **The PG Map:** Contains the PG version, its time stamp, the last OSD map epoch, the full ratios, and details on each placement group such as the PG ID, the *Up Set*, the *Acting Set*, the state of the PG (e.g., `active + clean`), and data usage statistics for each pool.
4. **The CRUSH Map:** Contains a list of storage devices, the failure domain hierarchy (e.g., device, host, rack, row, room, etc.), and rules for traversing the hierarchy when storing data. To view a CRUSH map, execute `ceph osd getcrushmap -o {filename}`; then, decompile it by executing `crushtool -d {comp-crushmap-filename} -o {decomp-crushmap-filename}`. You can view the decompiled map in a text editor or with `cat`.
5. **The MDS Map:** Contains the current MDS map epoch, when the map was created, and the last time it changed. It also contains the pool for storing metadata, a list of metadata servers, and which metadata servers are `up` and `in`. To view an MDS map, execute `ceph fs dump`.

Each map maintains an iterative history of its operating state changes. Ceph Monitors maintain a master copy of the cluster map including the cluster members, state, changes, and the overall health of the Ceph Storage Cluster.

HIGH AVAILABILITY MONITORS

Before Ceph Clients can read or write data, they must contact a Ceph Monitor to obtain the most recent copy of the cluster map. A Ceph Storage Cluster can operate with a single monitor; however, this introduces a single point of failure (i.e., if the monitor goes down, Ceph Clients cannot read or write data).

For added reliability and fault tolerance, Ceph supports a cluster of monitors. In a cluster of monitors, latency and other faults can cause one or more monitors to fall behind the current state of the cluster. For this reason, Ceph must have agreement among various monitor instances regarding the state of the cluster. Ceph always uses a majority of monitors (e.g., 1, 2:3, 3:5, 4:6, etc.) and the **Paxos** algorithm to establish a consensus among the monitors about the current state of the cluster.

For details on configuring monitors, see the [Monitor Config Reference](#).

HIGH AVAILABILITY AUTHENTICATION

To identify users and protect against man-in-the-middle attacks, Ceph provides its `cephx` authentication system to authenticate users and daemons.

Note: The `cephx` protocol does not address data encryption in transport (e.g., SSL/TLS) or encryption at rest.


`Cephx` uses shared secret keys for authentication, meaning both the client and the monitor cluster have a copy of the client's secret key. The authentication protocol is such that both parties are able to prove to each other they have a copy of the key without actually revealing it. This provides mutual authentication, which means the cluster is sure the user possesses the secret key, and the user is sure that the cluster has a copy of the secret key.

A key scalability feature of Ceph is to avoid a centralized interface to the Ceph object store, which means that Ceph clients must be able to interact with OSDs directly. To protect data, Ceph provides its `cephx` authentication system, which authenticates users operating Ceph clients. The `cephx` protocol operates in a manner with behavior similar to **Kerberos**.

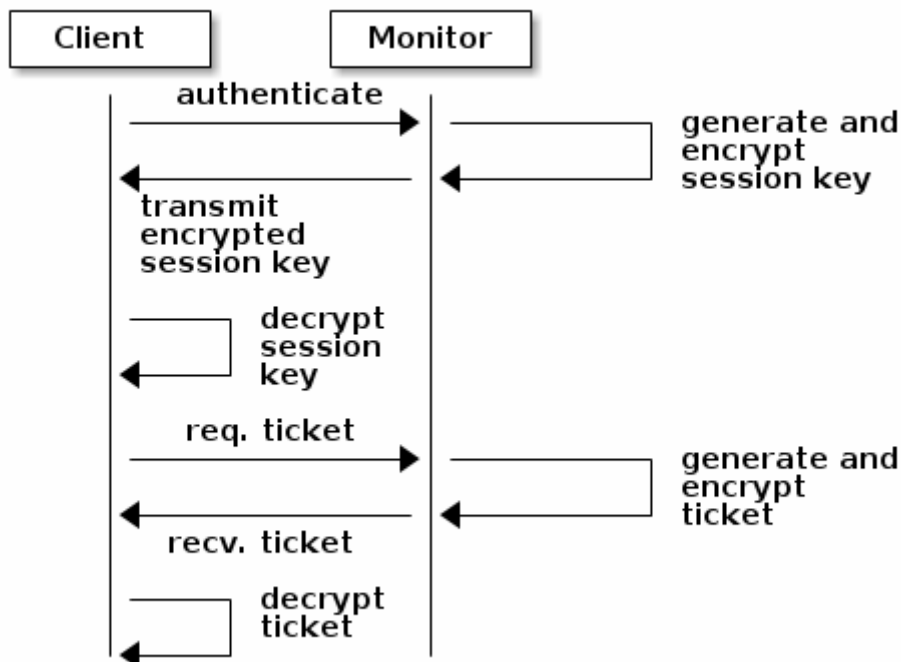
A user/actor invokes a Ceph client to contact a monitor. Unlike Kerberos, each monitor can authenticate users and distribute keys, so there is no single point of failure or bottleneck when using `cephx`. The monitor returns an authentication data structure similar to a Kerberos ticket that contains a session key for use in obtaining Ceph services. This session key is itself encrypted with the user's permanent secret key, so that only the user can request services from the Ceph Monitor(s). The client then uses the session key to request its desired services from the monitor, and the monitor provides the client with a ticket that will authenticate the client to the OSDs that actually handle data. Ceph Monitors and OSDs share a secret, so the client can use the ticket provided by the monitor with any OSD or metadata server in the cluster. Like Kerberos, `cephx` tickets expire, so an attacker cannot use an expired ticket or session key obtained surreptitiously. This form of authentication will prevent attackers with access to the communications medium from either creating bogus messages under another user's identity or altering another user's legitimate messages, as long as the user's secret key is not divulged before it expires.

To use `cephx`, an administrator must set up users first. In the following diagram, the `client.admin` user invokes `ceph auth get-or-create-key` from the command line to generate a username and secret key. Ceph's `auth` subsystem generates the username and key, stores a copy with the monitor(s) and transmits the user's secret back to the `client.admin` user. This means that the client and the monitor share a secret key.

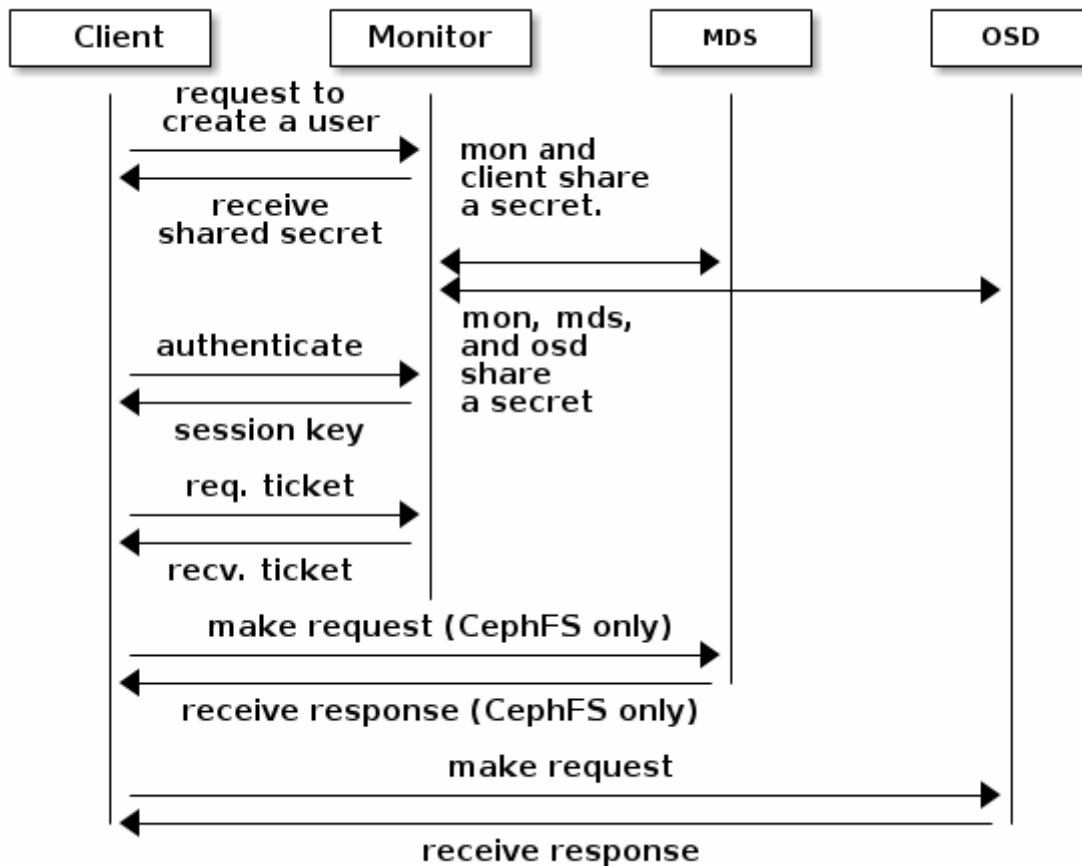
Note: The `client.admin` user must provide the user ID and secret key to the user in a secure manner.



To authenticate with the monitor, the client passes in the user name to the monitor, and the monitor generates a session key and encrypts it with the secret key associated to the user name. Then, the monitor transmits the encrypted ticket back to the client. The client then decrypts the payload with the shared secret key to retrieve the session key. The session key identifies the user for the current session. The client then requests a ticket on behalf of the user signed by the session key. The monitor generates a ticket, encrypts it with the user's secret key and transmits it back to the client. The client decrypts the ticket and uses it to sign requests to OSDs and metadata servers throughout the cluster.



The `cephx` protocol authenticates ongoing communications between the client machine and the Ceph servers. Each message sent between a client and server, subsequent to the initial authentication, is signed using a ticket that the monitors, OSDs and metadata servers can verify with their shared secret.



The protection offered by this authentication is between the Ceph client and the Ceph server hosts. The authentication is not extended beyond the Ceph client. If the user accesses the Ceph client from a remote host, Ceph authentication is not applied to the connection between the user's host and the client host.

For configuration details, see [Cephx Config Guide](#). For user management details, see [User Management](#).

SMART DAEMONS ENABLE HYPERSCALE

In many clustered architectures, the primary purpose of cluster membership is so that a centralized interface knows which nodes it can access. Then the centralized interface provides services to the client through a double dispatch—which is a **huge** bottleneck at the petabyte-to-exabyte scale.

Ceph eliminates the bottleneck: Ceph's OSD Daemons AND Ceph Clients are cluster aware. Like Ceph clients, each Ceph OSD Daemon knows about other Ceph OSD Daemons in the cluster. This enables Ceph OSD Daemons to interact directly with other Ceph OSD Daemons and Ceph Monitors. Additionally, it enables Ceph Clients to interact directly with Ceph OSD Daemons.

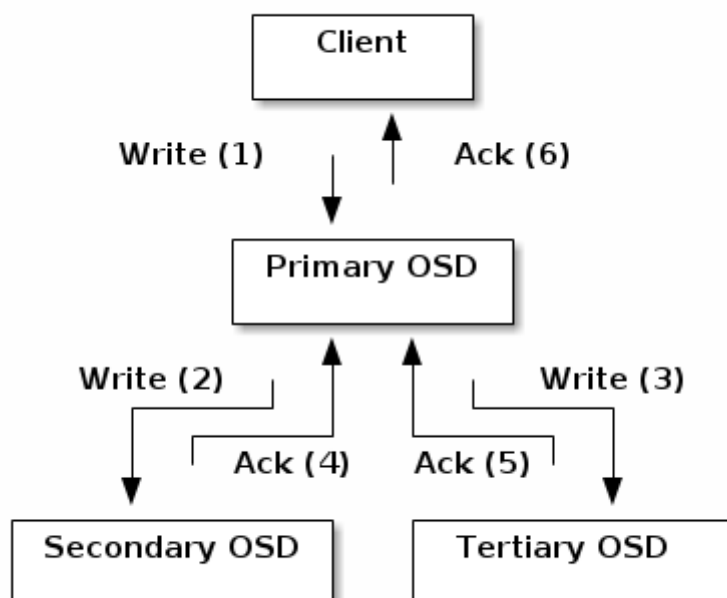
The ability of Ceph Clients, Ceph Monitors and Ceph OSD Daemons to interact with each other means that Ceph OSD Daemons can utilize the CPU and RAM of the Ceph nodes to easily perform tasks that would bog down a centralized server. The ability to leverage this computing power leads to several major benefits:

1. **OSDs Service Clients Directly:** Since any network device has a limit to the number of concurrent connections it can support, a centralized system has a low physical limit at high scales. By enabling Ceph Clients to contact Ceph OSD Daemons directly, Ceph increases both performance and total system capacity simultaneously, while removing a single point of failure. Ceph Clients can maintain

a session when they need to, and with a particular Ceph OSD Daemon instead of a centralized server.

2. **OSD Membership and Status:** Ceph OSD Daemons join a cluster and report on their status. At the lowest level, the Ceph OSD Daemon status is up or down reflecting whether or not it is running and able to service Ceph Client requests. If a Ceph OSD Daemon is down and in the Ceph Storage Cluster, this status may indicate the failure of the Ceph OSD Daemon. If a Ceph OSD Daemon is not running (e.g., it crashes), the Ceph OSD Daemon cannot notify the Ceph Monitor that it is down. The OSDs periodically send messages to the Ceph Monitor (MPGStats pre-luminous, and a new MOSDBeacon in luminous). If the Ceph Monitor doesn't see that message after a configurable period of time then it marks the OSD down. This mechanism is a failsafe, however. Normally, Ceph OSD Daemons will determine if a neighboring OSD is down and report it to the Ceph Monitor(s). This assures that Ceph Monitors are lightweight processes. See [Monitoring OSDs](#) and [Heartbeats](#) for additional details.
3. **Data Scrubbing:** As part of maintaining data consistency and cleanliness, Ceph OSD Daemons can scrub objects within placement groups. That is, Ceph OSD Daemons can compare object metadata in one placement group with its replicas in placement groups stored on other OSDs. Scrubbing (usually performed daily) catches bugs or filesystem errors. Ceph OSD Daemons also perform deeper scrubbing by comparing data in objects bit-for-bit. Deep scrubbing (usually performed weekly) finds bad sectors on a drive that weren't apparent in a light scrub. See [Data Scrubbing](#) for details on configuring scrubbing.
4. **Replication:** Like Ceph Clients, Ceph OSD Daemons use the CRUSH algorithm, but the Ceph OSD Daemon uses it to compute where replicas of objects should be stored (and for rebalancing). In a typical write scenario, a client uses the CRUSH algorithm to compute where to store an object, maps the object to a pool and placement group, then looks at the CRUSH map to identify the primary OSD for the placement group.

The client writes the object to the identified placement group in the primary OSD. Then, the primary OSD with its own copy of the CRUSH map identifies the secondary and tertiary OSDs for replication purposes, and replicates the object to the appropriate placement groups in the secondary and tertiary OSDs (as many OSDs as additional replicas), and responds to the client once it has confirmed the object was stored successfully.



With the ability to perform data replication, Ceph OSD Daemons relieve Ceph clients from that duty, while ensuring high data availability and data safety.


DYNAMIC CLUSTER MANAGEMENT

In the [Scalability and High Availability](#) section, we explained how Ceph uses CRUSH, cluster awareness and intelligent daemons to scale and maintain high availability. Key to Ceph's design is the autonomous, self-healing, and intelligent Ceph OSD Daemon. Let's take a deeper look at how CRUSH works to enable modern cloud storage infrastructures to place data, rebalance the cluster and recover from faults dynamically.

ABOUT POOLS

The Ceph storage system supports the notion of 'Pools', which are logical partitions for storing objects.

Ceph Clients retrieve a [Cluster Map](#) from a Ceph Monitor, and write objects to pools. The pool's size or number of replicas, the CRUSH rule and the number of placement groups determine how Ceph will place the data.



Pools set at least the following parameters:

- Ownership/Access to Objects
- The Number of Placement Groups, and
- The CRUSH Rule to Use.

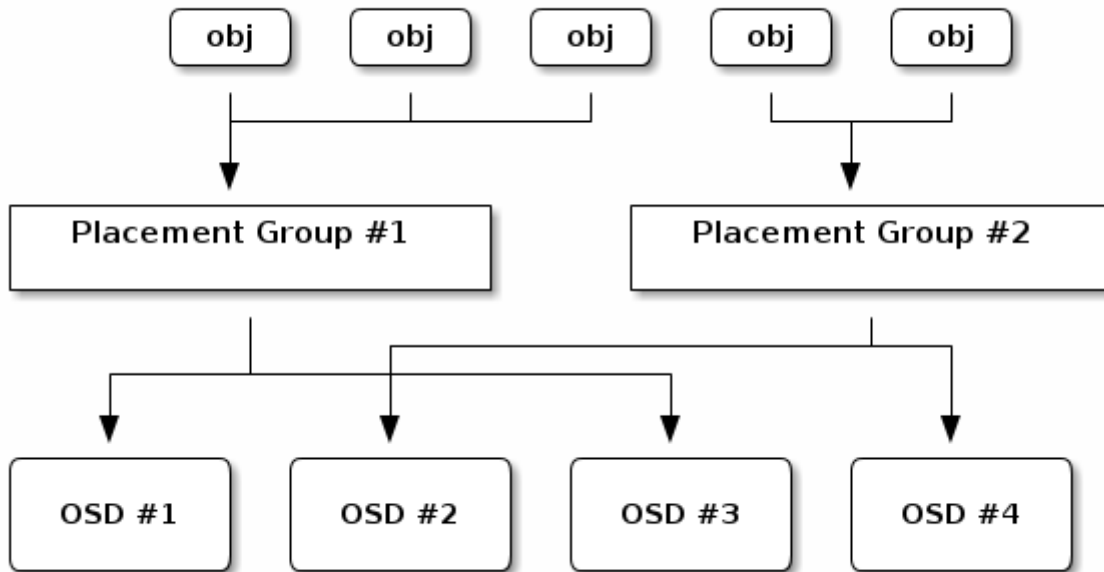
See [Set Pool Values](#) for details.

MAPPING PGS TO OSDS

Each pool has a number of placement groups. CRUSH maps PGs to OSDs dynamically. When a Ceph Client stores objects, CRUSH will map each object to a placement group.

Mapping objects to placement groups creates a layer of indirection between the Ceph OSD Daemon and the Ceph Client. The Ceph Storage Cluster must be able to grow (or shrink) and rebalance where it stores objects dynamically. If the Ceph Client "knew" which Ceph OSD Daemon had which object, that would create a tight coupling between the Ceph Client and the Ceph OSD Daemon. Instead, the CRUSH algorithm maps each object to a placement group and then maps each placement group to one or more

Ceph OSD Daemons. This layer of indirection allows Ceph to rebalance dynamically when new Ceph OSD Daemons and the underlying OSD devices come online. The following diagram depicts how CRUSH maps objects to placement groups, and placement groups to OSDs.



With a copy of the cluster map and the CRUSH algorithm, the client can compute exactly which OSD to use when reading or writing a particular object.

CALCULATING PG IDS

When a Ceph Client binds to a Ceph Monitor, it retrieves the latest copy of the **Cluster Map**. With the cluster map, the client knows about all of the monitors, OSDs, and metadata servers in the cluster.

However, it doesn't know anything about object locations.

Object locations get computed.

The only input required by the client is the object ID and the pool. It's simple: Ceph stores data in named pools (e.g., "liverpool"). When a client wants to store a named object (e.g., "john," "paul," "george," "ringo", etc.) it calculates a placement group using the object name, a hash code, the number of PGs in the pool and the pool name. Ceph clients use the following steps to compute PG IDs.

1. The client inputs the pool name and the object ID. (e.g., pool = "liverpool" and object-id = "john")
2. Ceph takes the object ID and hashes it.
3. Ceph calculates the hash modulo the number of PGs. (e.g., 58) to get a PG ID.
4. Ceph gets the pool ID given the pool name (e.g., "liverpool" = 4)
5. Ceph prepends the pool ID to the PG ID (e.g., 4.58).

Computing object locations is much faster than performing object location query over a chatty session. The **CRUSH** algorithm allows a client to compute where objects *should* be stored, and enables the client to contact the primary OSD to store or retrieve the objects.

PEERING AND SETS

In previous sections, we noted that Ceph OSD Daemons check each others heartbeats and report back to the Ceph Monitor. Another thing Ceph OSD daemons do is called 'peering', which is the process of bringing all of the OSDs that store a Placement Group (PG) into agreement about the state of all of the objects (and their metadata) in that PG. In fact, Ceph OSD Daemons **Report Peering Failure** to the Ceph Monitors. Peering issues usually resolve themselves; however, if the problem persists, you may need to refer to the **Troubleshooting Peering Failure** section.

Note: Agreeing on the state does not mean that the PGs have the latest contents.

The Ceph Storage Cluster was designed to store at least two copies of an object (i.e., `size = 2`), which is the minimum requirement for data safety. For high availability, a Ceph Storage Cluster should store more than two copies of an object (e.g., `size = 3` and `min size = 2`) so that it can continue to run in a degraded state while maintaining data safety.

Referring back to the diagram in **Smart Daemons Enable Hyperscale**, we do not name the Ceph OSD Daemons specifically (e.g., `osd.0`, `osd.1`, etc.), but rather refer to them as *Primary*, *Secondary*, and so forth. By convention, the *Primary* is the first OSD in the *Acting Set*, and is responsible for coordinating the peering process for each placement group where it acts as the *Primary*, and is the **ONLY** OSD that that will accept client-initiated writes to objects for a given placement group where it acts as the *Primary*.

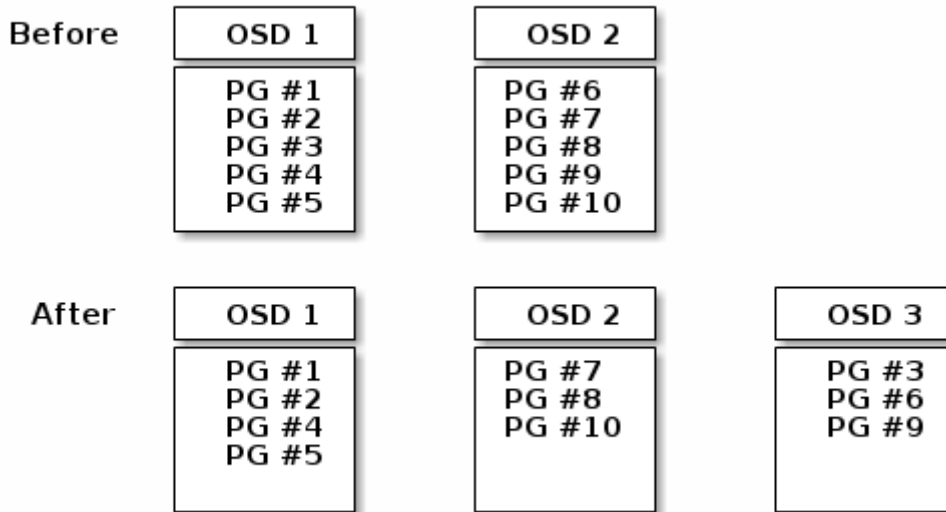
When a series of OSDs are responsible for a placement group, that series of OSDs, we refer to them as an *Acting Set*. An *Acting Set* may refer to the Ceph OSD Daemons that are currently responsible for the placement group, or the Ceph OSD Daemons that were responsible for a particular placement group as of some epoch.

The Ceph OSD daemons that are part of an *Acting Set* may not always be up. When an OSD in the *Acting Set* is up, it is part of the *Up Set*. The *Up Set* is an important distinction, because Ceph can remap PGs to other Ceph OSD Daemons when an OSD fails.

Note: In an *Acting Set* for a PG containing `osd.25`, `osd.32` and `osd.61`, the first OSD, `osd.25`, is the *Primary*. If that OSD fails, the Secondary, `osd.32`, becomes the *Primary*, and `osd.25` will be removed from the *Up Set*.

REBALANCING

When you add a Ceph OSD Daemon to a Ceph Storage Cluster, the cluster map gets updated with the new OSD. Referring back to **Calculating PG IDs**, this changes the cluster map. Consequently, it changes object placement, because it changes an input for the calculations. The following diagram depicts the rebalancing process (albeit rather crudely, since it is substantially less impactful with large clusters) where some, but not all of the PGs migrate from existing OSDs (OSD 1, and OSD 2) to the new OSD (OSD 3). Even when rebalancing, CRUSH is stable. Many of the placement groups remain in their original configuration, and each OSD gets some added capacity, so there are no load spikes on the new OSD after rebalancing is complete.



DATA CONSISTENCY

As part of maintaining data consistency and cleanliness, Ceph OSDs can also scrub objects within placement groups. That is, Ceph OSDs can compare object metadata in one placement group with its replicas in placement groups stored in other OSDs. Scrubbing (usually performed daily) catches OSD bugs or filesystem errors. OSDs can also perform deeper scrubbing by comparing data in objects bit-for-bit. Deep scrubbing (usually performed weekly) finds bad sectors on a disk that weren't apparent in a light scrub.

See [Data Scrubbing](#) for details on configuring scrubbing.

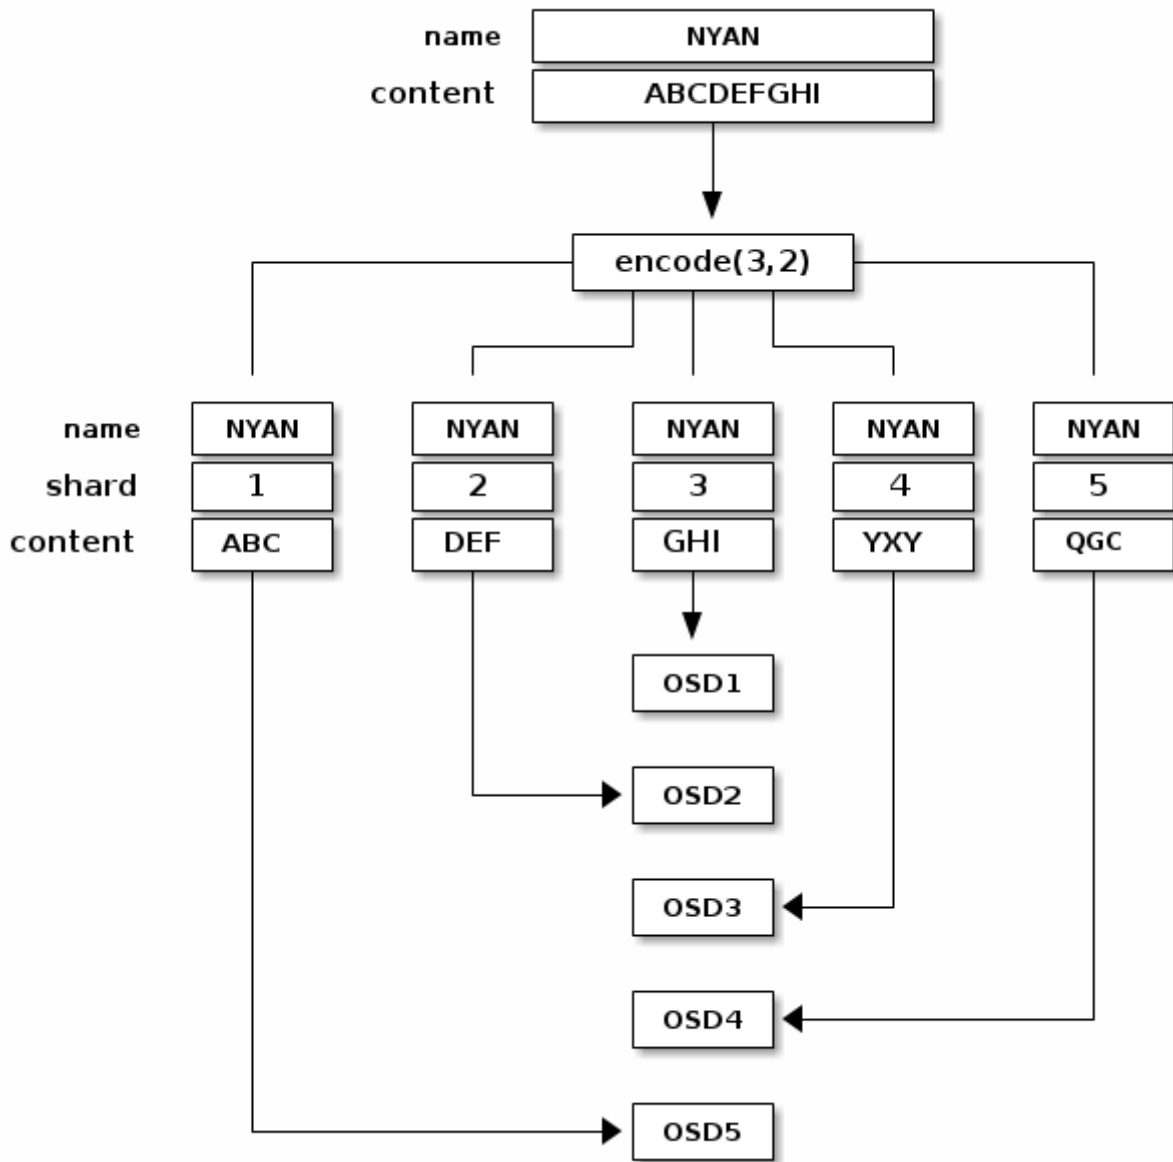
ERASURE CODING

An erasure coded pool stores each object as $K+M$ chunks. It is divided into K data chunks and M coding chunks. The pool is configured to have a size of $K+M$ so that each chunk is stored in an OSD in the acting set. The rank of the chunk is stored as an attribute of the object.

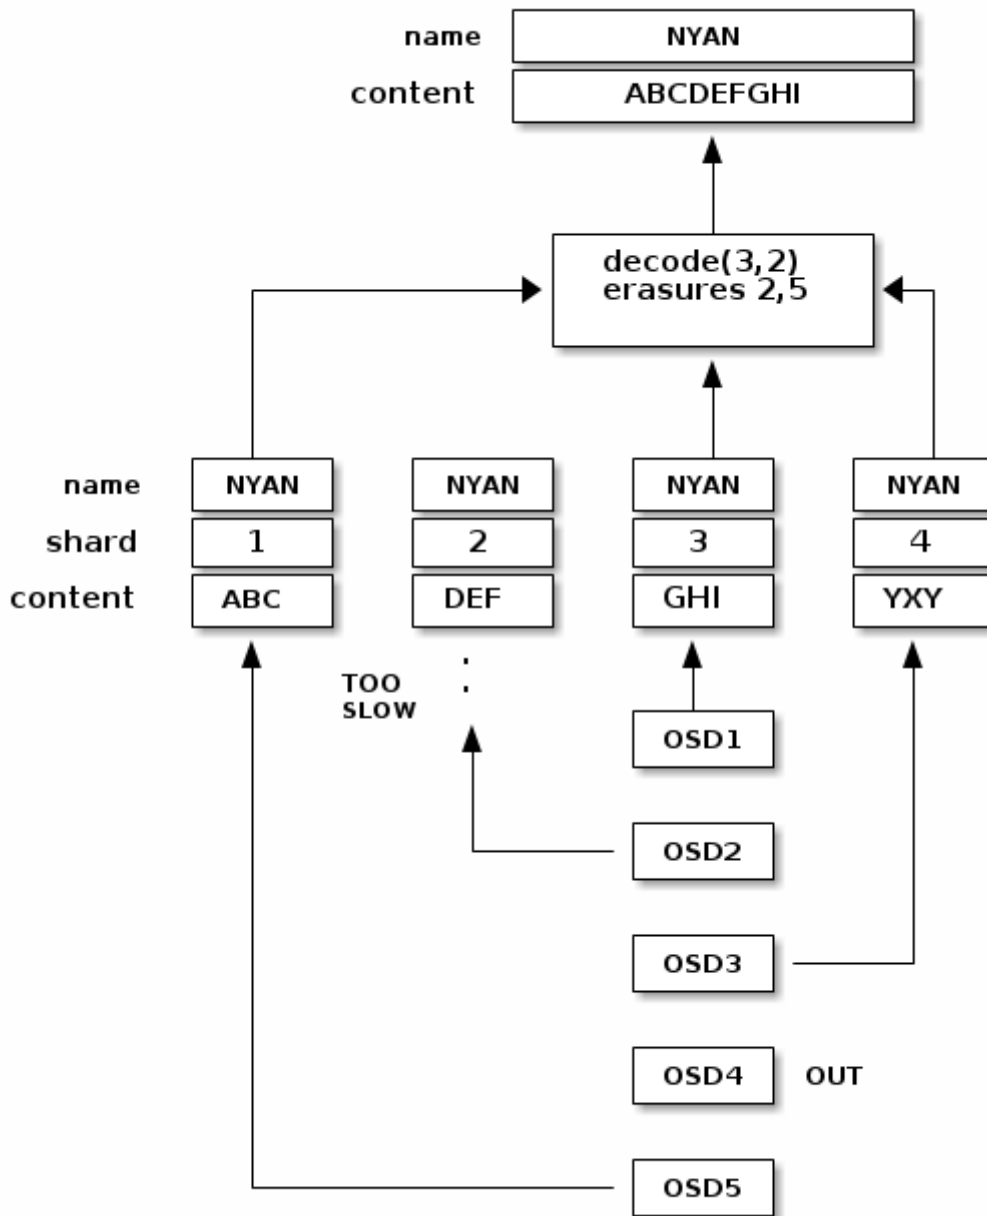
For instance an erasure coded pool is created to use five OSDs ($K+M = 5$) and sustain the loss of two of them ($M = 2$).

READING AND WRITING ENCODED CHUNKS

When the object **NYAN** containing ABCDEFGHI is written to the pool, the erasure encoding function splits the content into three data chunks simply by dividing the content in three: the first contains ABC, the second DEF and the last GHI. The content will be padded if the content length is not a multiple of K . The function also creates two coding chunks: the fourth with YXY and the fifth with QGC. Each chunk is stored in an OSD in the acting set. The chunks are stored in objects that have the same name (**NYAN**) but reside on different OSDs. The order in which the chunks were created must be preserved and is stored as an attribute of the object (`shard_t`), in addition to its name. Chunk 1 contains ABC and is stored on **OSD5** while chunk 4 contains YXY and is stored on **OSD3**.



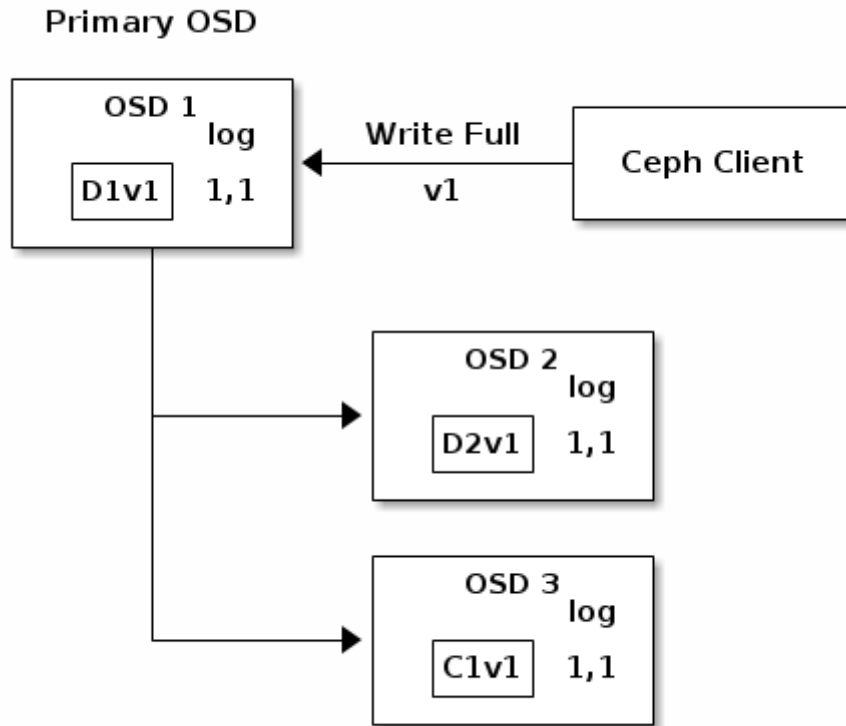
When the object **NYAN** is read from the erasure coded pool, the decoding function reads three chunks: chunk 1 containing ABC, chunk 3 containing GHI and chunk 4 containing YXY. Then, it rebuilds the original content of the object ABCDEFGHI. The decoding function is informed that the chunks 2 and 5 are missing (they are called 'erasures'). The chunk 5 could not be read because the **OSD4** is out. The decoding function can be called as soon as three chunks are read: **OSD2** was the slowest and its chunk was not taken into account.



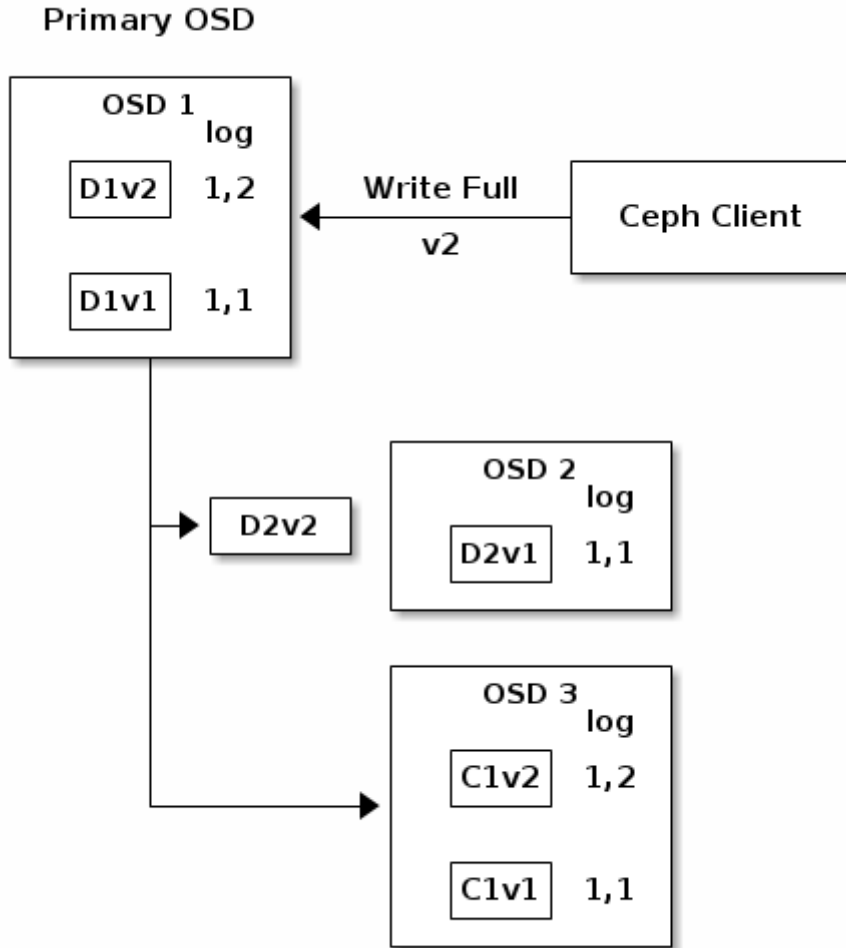
INTERRUPTED FULL WRITES

In an erasure coded pool, the primary OSD in the up set receives all write operations. It is responsible for encoding the payload into $K+M$ chunks and sends them to the other OSDs. It is also responsible for maintaining an authoritative version of the placement group logs.

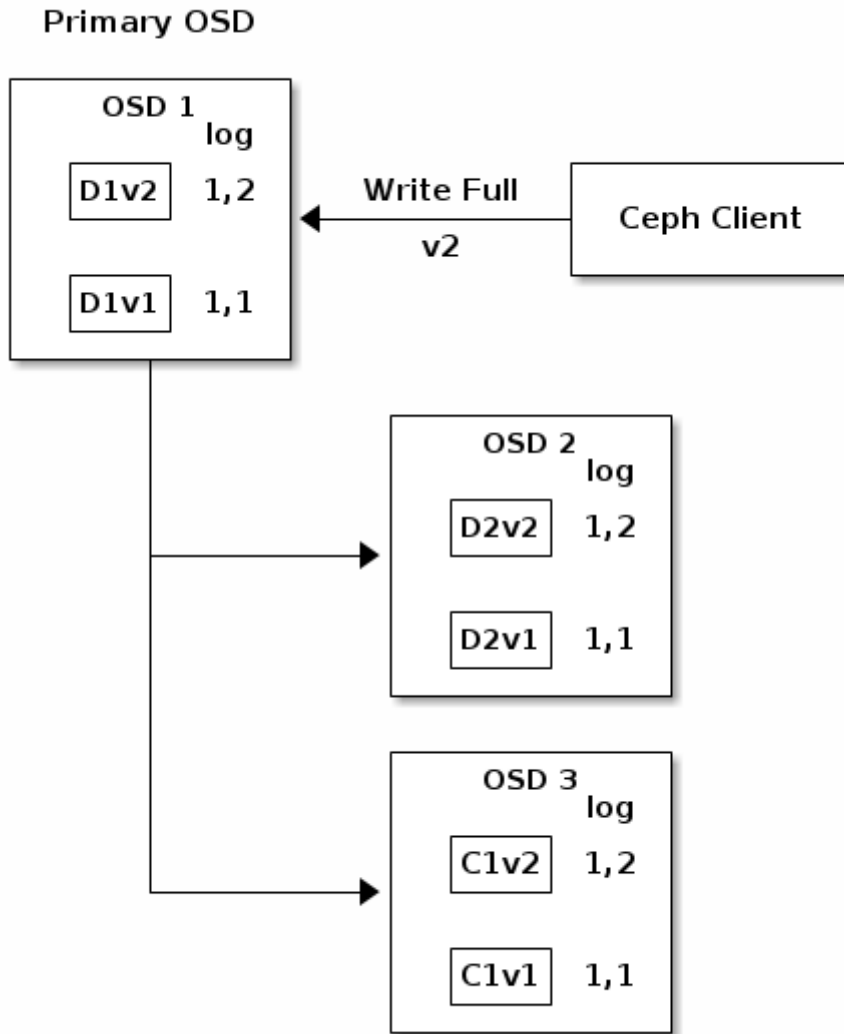
In the following diagram, an erasure coded placement group has been created with $K = 2 + M = 1$ and is supported by three OSDs, two for K and one for M . The acting set of the placement group is made of **OSD 1**, **OSD 2** and **OSD 3**. An object has been encoded and stored in the OSDs : the chunk D1v1 (i.e. Data chunk number 1, version 1) is on **OSD 1**, D2v1 on **OSD 2** and C1v1 (i.e. Coding chunk number 1, version 1) on **OSD 3**. The placement group logs on each OSD are identical (i.e. 1, 1 for epoch 1, version 1).



OSD 1 is the primary and receives a **WRITE FULL** from a client, which means the payload is to replace the object entirely instead of overwriting a portion of it. Version 2 (v2) of the object is created to override version 1 (v1). **OSD 1** encodes the payload into three chunks: D1v2 (i.e. Data chunk number 1 version 2) will be on **OSD 1**, D2v2 on **OSD 2** and C1v2 (i.e. Coding chunk number 1 version 2) on **OSD 3**. Each chunk is sent to the target OSD, including the primary OSD which is responsible for storing chunks in addition to handling write operations and maintaining an authoritative version of the placement group logs. When an OSD receives the message instructing it to write the chunk, it also creates a new entry in the placement group logs to reflect the change. For instance, as soon as **OSD 3** stores C1v2, it adds the entry 1,2 (i.e. epoch 1, version 2) to its logs. Because the OSDs work asynchronously, some chunks may still be in flight (such as D2v2) while others are acknowledged and on disk (such as C1v1 and D1v1).

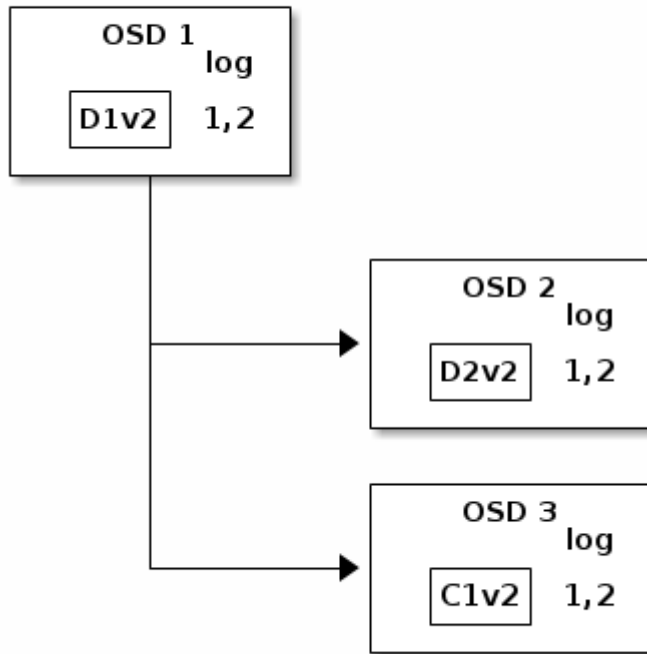


If all goes well, the chunks are acknowledged on each OSD in the acting set and the logs' last_complete pointer can move from 1,1 to 1,2.

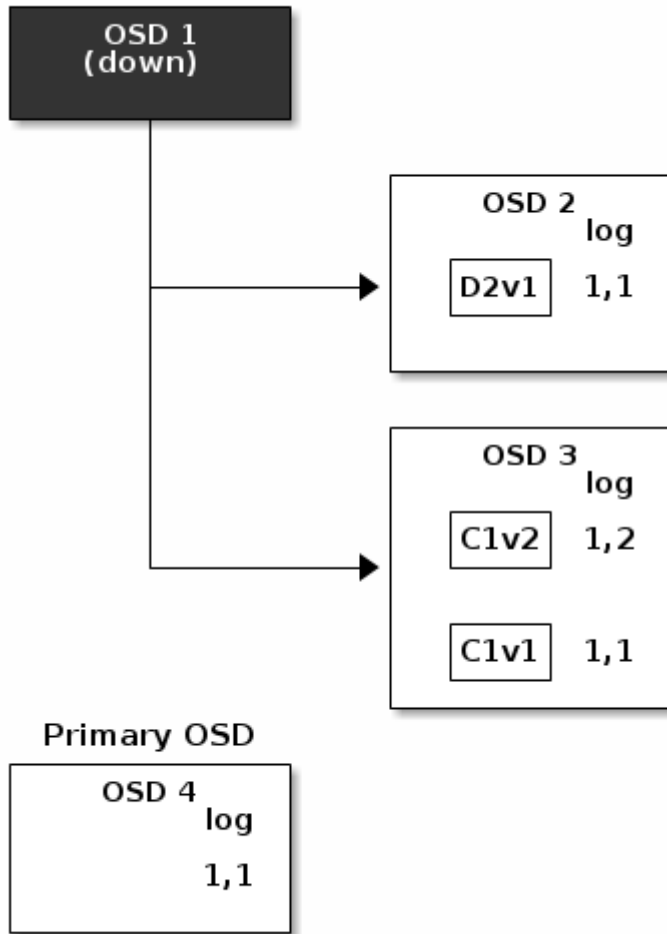


Finally, the files used to store the chunks of the previous version of the object can be removed: D1v1 on **OSD 1**, D2v1 on **OSD 2** and C1v1 on **OSD 3**.

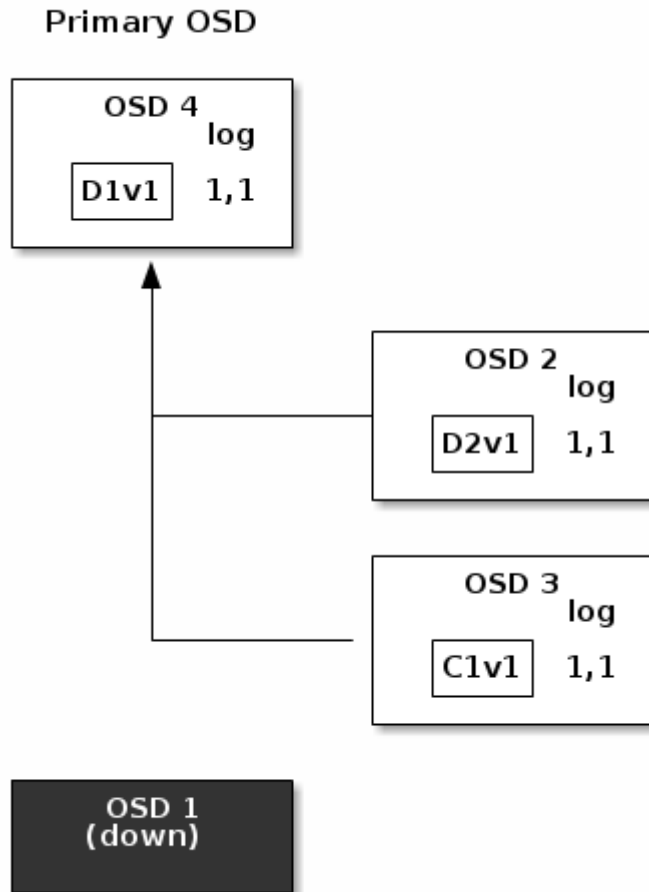
Primary OSD



But accidents happen. If **OSD 1** goes down while D2v2 is still in flight, the object's version 2 is partially written: **OSD 3** has one chunk but that is not enough to recover. It lost two chunks: D1v2 and D2v2 and the erasure coding parameters $K = 2$, $M = 1$ require that at least two chunks are available to rebuild the third. **OSD 4** becomes the new primary and finds that the `last_complete` log entry (i.e., all objects before this entry were known to be available on all OSDs in the previous acting set) is 1,1 and that will be the head of the new authoritative log.




The log entry 1,2 found on **OSD 3** is divergent from the new authoritative log provided by **OSD 4**: it is discarded and the file containing the C1v2 chunk is removed. The D1v1 chunk is rebuilt with the decode function of the erasure coding library during scrubbing and stored on the new primary **OSD 4**.



See [Erasure Code Notes](#) for additional details.

CACHE TIERING

A cache tier provides Ceph Clients with better I/O performance for a subset of the data stored in a backing storage tier. Cache tiering involves creating a pool of relatively fast/expensive storage devices (e.g., solid state drives) configured to act as a cache tier, and a backing pool of either erasure-coded or relatively slower/cheaper devices configured to act as an economical storage tier. The Ceph objecter handles where to place the objects and the tiering agent determines when to flush objects from the cache to the backing storage tier. So the cache tier and the backing storage tier are completely transparent to Ceph clients.



See [Cache Tiering](#) for additional details.

EXTENDING CEPH

You can extend Ceph by creating shared object classes called 'Ceph Classes'. Ceph loads .so classes stored in the `osd class dir` directory dynamically (i.e., `$libdir/rados-classes` by default). When you implement a class, you can create new object methods that have the ability to call the native methods in the Ceph Object Store, or other class methods you incorporate via libraries or create yourself.

On writes, Ceph Classes can call native or class methods, perform any series of operations on the inbound data and generate a resulting write transaction that Ceph will apply atomically.

On reads, Ceph Classes can call native or class methods, perform any series of operations on the outbound data and return the data to the client.

Ceph Class Example

A Ceph class for a content management system that presents pictures of a particular size and aspect ratio could take an inbound bitmap image, crop it to a particular aspect ratio, resize it and embed an invisible copyright or watermark to help protect the intellectual property; then, save the resulting bitmap image to the object store.

See `src/objclass/objclass.h`, `src/fooclass.cc` and `src/barclass` for exemplary implementations.


SUMMARY

Ceph Storage Clusters are dynamic—like a living organism. Whereas, many storage appliances do not fully utilize the CPU and RAM of a typical commodity server, Ceph does. From heartbeats, to peering, to rebalancing the cluster or recovering from faults, Ceph offloads work from clients (and from a centralized gateway which doesn't exist in the Ceph architecture) and uses the computing power of the OSDs to

perform the work. When referring to [Hardware Recommendations](#) and the [Network Config Reference](#), be cognizant of the foregoing concepts to understand how Ceph utilizes computing resources.

CEPH PROTOCOL

Ceph Clients use the native protocol for interacting with the Ceph Storage Cluster. Ceph packages this functionality into the `librados` library so that you can create your own custom Ceph Clients. The following diagram depicts the basic architecture.



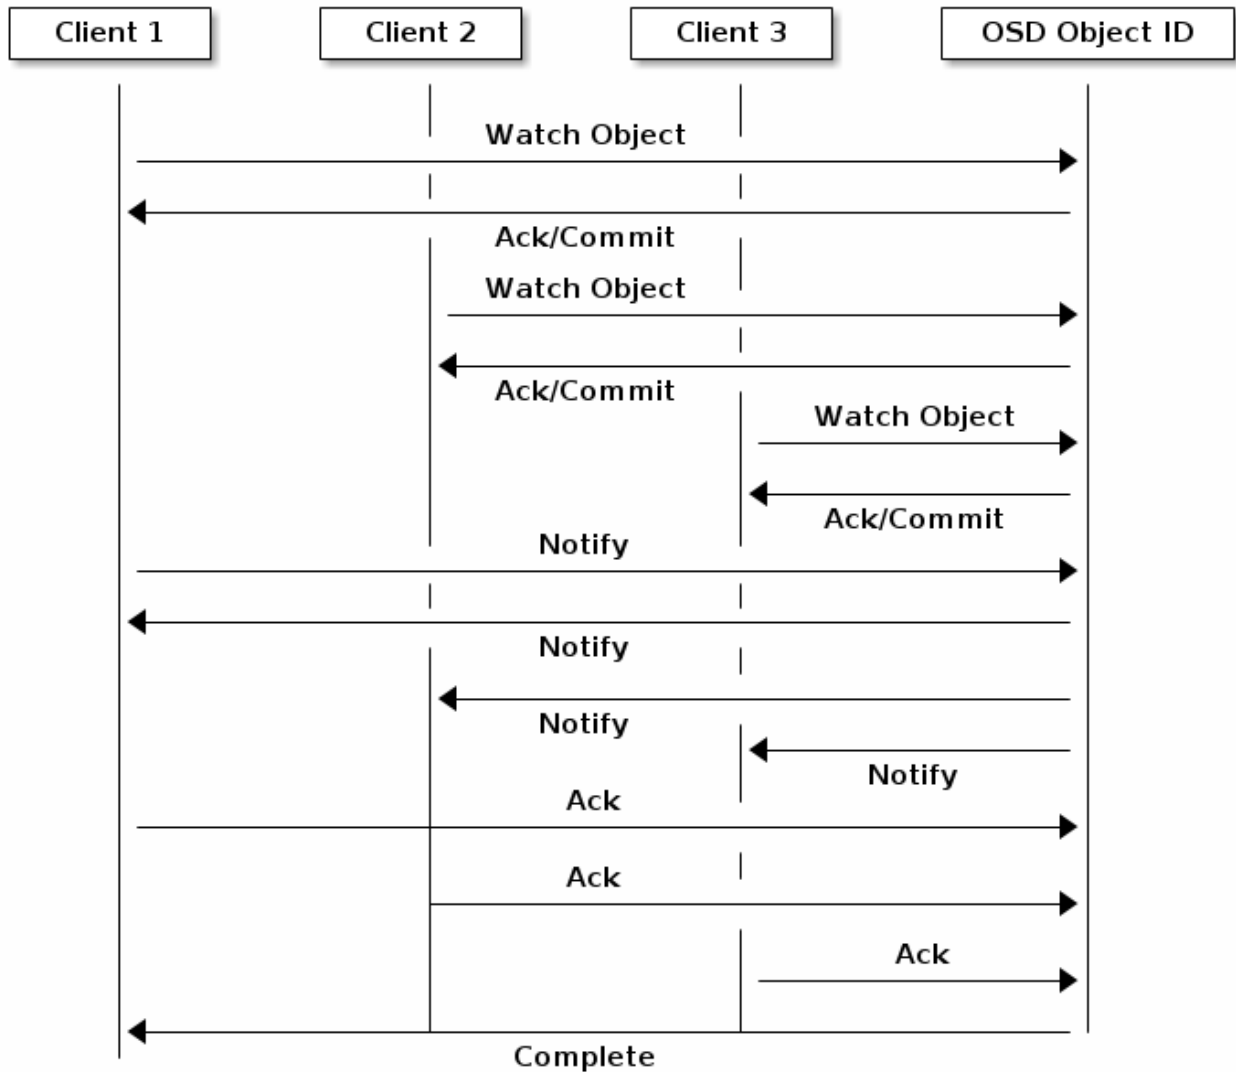
NATIVE PROTOCOL AND LIBRADOS

Modern applications need a simple object storage interface with asynchronous communication capability. The Ceph Storage Cluster provides a simple object storage interface with asynchronous communication capability. The interface provides direct, parallel access to objects throughout the cluster.

- Pool Operations
- Snapshots and Copy-on-write Cloning
- Read/Write Objects - Create or Remove - Entire Object or Byte Range - Append or Truncate
- Create/Set/Get/Remove XATTRs
- Create/Set/Get/Remove Key/Value Pairs
- Compound operations and dual-ack semantics
- Object Classes

OBJECT WATCH/NOTIFY

A client can register a persistent interest with an object and keep a session to the primary OSD open. The client can send a notification message and a payload to all watchers and receive notification when the watchers receive the notification. This enables a client to use any object as a synchronization/communication channel.



DATA STRIPING


Storage devices have throughput limitations, which impact performance and scalability. So storage systems often support **striping**—storing sequential pieces of information across multiple storage devices—to increase throughput and performance. The most common form of data striping comes from **RAID**. The RAID type most similar to Ceph’s striping is **RAID 0**, or a ‘striped volume’. Ceph’s striping offers the throughput of RAID 0 striping, the reliability of n-way RAID mirroring and faster recovery.

Ceph provides three types of clients: Ceph Block Device, Ceph Filesystem, and Ceph Object Storage. A Ceph Client converts its data from the representation format it provides to its users (a block device image, RESTful objects, CephFS filesystem directories) into objects for storage in the Ceph Storage Cluster.

Tip: The objects Ceph stores in the Ceph Storage Cluster are not striped. Ceph Object Storage, Ceph Block Device, and the Ceph Filesystem stripe their data over multiple Ceph Storage Cluster objects. Ceph Clients that write directly to the Ceph Storage Cluster via librados must perform the striping (and parallel I/O) for themselves to obtain these benefits.

The simplest Ceph striping format involves a stripe count of 1 object. Ceph Clients write stripe units to a Ceph Storage Cluster object until the object is at its maximum capacity, and then create another object for additional stripes of data. The simplest form of striping may be sufficient for small block device images, S3

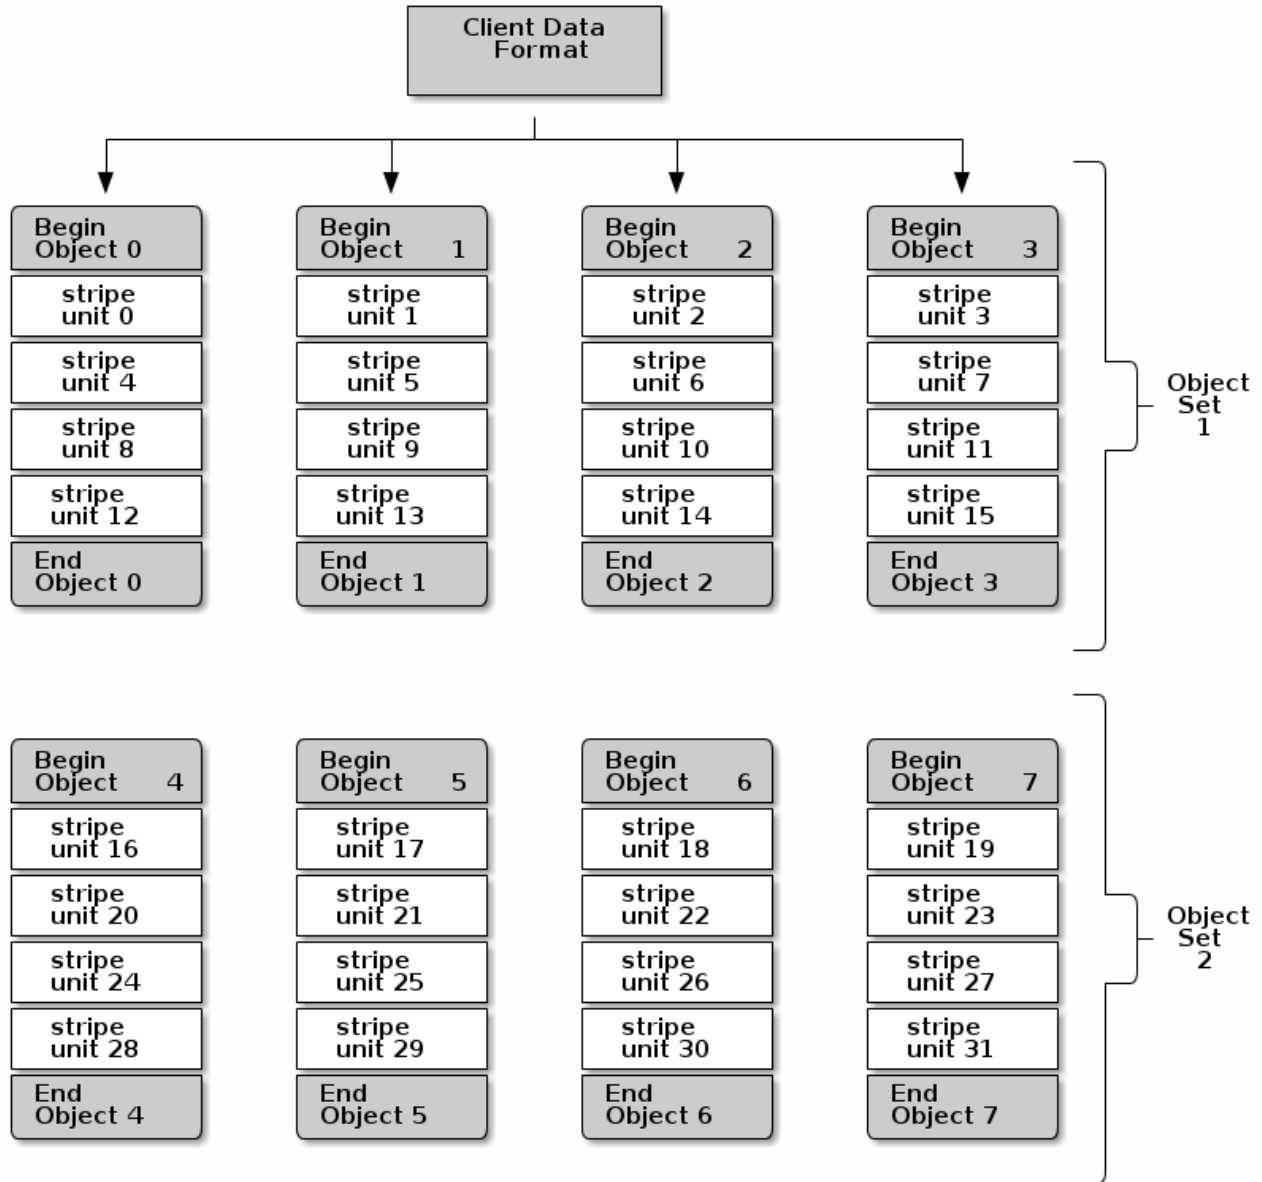
or Swift objects and CephFS files. However, this simple form doesn't take maximum advantage of Ceph's ability to distribute data across placement groups, and consequently doesn't improve performance very much. The following diagram depicts the simplest form of striping:



If you anticipate large images sizes, large S3 or Swift objects (e.g., video), or large CephFS directories, you may see considerable read/write performance improvements by striping client data over multiple objects within an object set. Significant write performance occurs when the client writes the stripe units to their corresponding objects in parallel. Since objects get mapped to different placement groups and further mapped to different OSDs, each write occurs in parallel at the maximum write speed. A write to a single disk would be limited by the head movement (e.g. 6ms per seek) and bandwidth of that one device (e.g. 100MB/s). By spreading that write over multiple objects (which map to different placement groups and OSDs) Ceph can reduce the number of seeks per drive and combine the throughput of multiple drives to achieve much faster write (or read) speeds.

Note: Striping is independent of object replicas. Since CRUSH replicates objects across OSDs, stripes get replicated automatically.

In the following diagram, client data gets striped across an object set (object set 1 in the following diagram) consisting of 4 objects, where the first stripe unit is stripe unit 0 in object 0, and the fourth stripe unit is stripe unit 3 in object 3. After writing the fourth stripe, the client determines if the object set is full. If the object set is not full, the client begins writing a stripe to the first object again (object 0 in the following diagram). If the object set is full, the client creates a new object set (object set 2 in the following diagram), and begins writing to the first stripe (stripe unit 16) in the first object in the new object set (object 4 in the diagram below).



Three important variables determine how Ceph stripes data:

- **Object Size:** Objects in the Ceph Storage Cluster have a maximum configurable size (e.g., 2MB, 4MB, etc.). The object size should be large enough to accommodate many stripe units, and should be a multiple of the stripe unit.
- **Stripe Width:** Stripes have a configurable unit size (e.g., 64kb). The Ceph Client divides the data it will write to objects into equally sized stripe units, except for the last stripe unit. A stripe width, should be a fraction of the Object Size so that an object may contain many stripe units.
- **Stripe Count:** The Ceph Client writes a sequence of stripe units over a series of objects determined by the stripe count. The series of objects is called an object set. After the Ceph Client writes to the last object in the object set, it returns to the first object in the object set.

Important: Test the performance of your striping configuration before putting your cluster into production. You CANNOT change these striping parameters after you stripe the data and write it to objects.

Once the Ceph Client has striped data to stripe units and mapped the stripe units to objects, Ceph's CRUSH algorithm maps the objects to placement groups, and the placement groups to Ceph OSD

Daemons before the objects are stored as files on a storage disk.

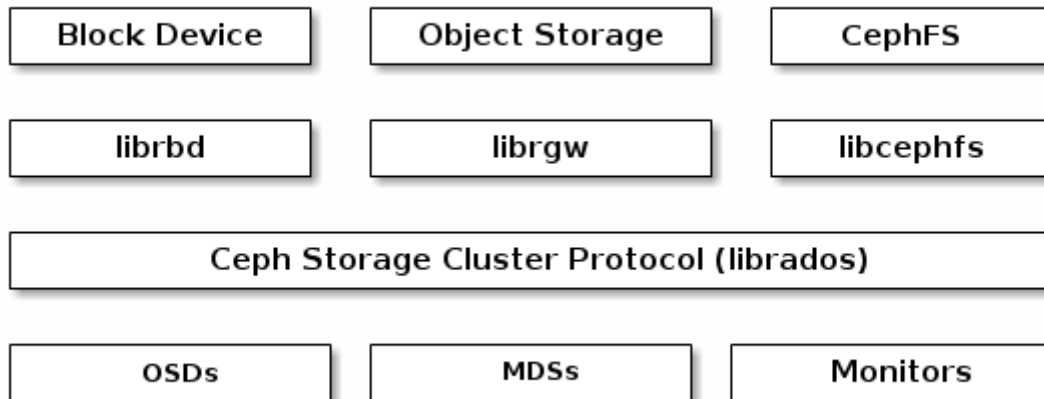
Note: Since a client writes to a single pool, all data striped into objects get mapped to placement groups in the same pool. So they use the same CRUSH map and the same access controls.

CEPH CLIENTS

Ceph Clients include a number of service interfaces. These include:

- **Block Devices:** The **Ceph Block Device** (a.k.a., RBD) service provides resizable, thin-provisioned block devices with snapshotting and cloning. Ceph stripes a block device across the cluster for high performance. Ceph supports both kernel objects (KO) and a QEMU hypervisor that uses `librbd` directly—avoiding the kernel object overhead for virtualized systems.
- **Object Storage:** The **Ceph Object Storage** (a.k.a., RGW) service provides RESTful APIs with interfaces that are compatible with Amazon S3 and OpenStack Swift.
- **Filesystem:** The **Ceph Filesystem** (CephFS) service provides a POSIX compliant filesystem usable with `mount` or as a filesystem in user space (FUSE).

Ceph can run additional instances of OSDs, MDSs, and monitors for scalability and high availability. The following diagram depicts the high-level architecture.



CEPH OBJECT STORAGE

The Ceph Object Storage daemon, `radosgw`, is a FastCGI service that provides a **RESTful** HTTP API to store objects and metadata. It layers on top of the Ceph Storage Cluster with its own data formats, and maintains its own user database, authentication, and access control. The RADOS Gateway uses a unified namespace, which means you can use either the OpenStack Swift-compatible API or the Amazon S3-compatible API. For example, you can write data using the S3-compatible API with one application and then read data using the Swift-compatible API with another application.

S3/Swift Objects and Store Cluster Objects Compared

Ceph's Object Storage uses the term *object* to describe the data it stores. S3 and Swift objects are not the same as the objects that Ceph writes to the Ceph Storage Cluster. Ceph Object Storage objects are

mapped to Ceph Storage Cluster objects. The S3 and Swift objects do not necessarily correspond in a 1:1 manner with an object stored in the storage cluster. It is possible for an S3 or Swift object to map to multiple Ceph objects.

See [Ceph Object Storage](#) for details.

CEPH BLOCK DEVICE

A Ceph Block Device stripes a block device image over multiple objects in the Ceph Storage Cluster, where each object gets mapped to a placement group and distributed, and the placement groups are spread across separate `ceph-osd` daemons throughout the cluster.

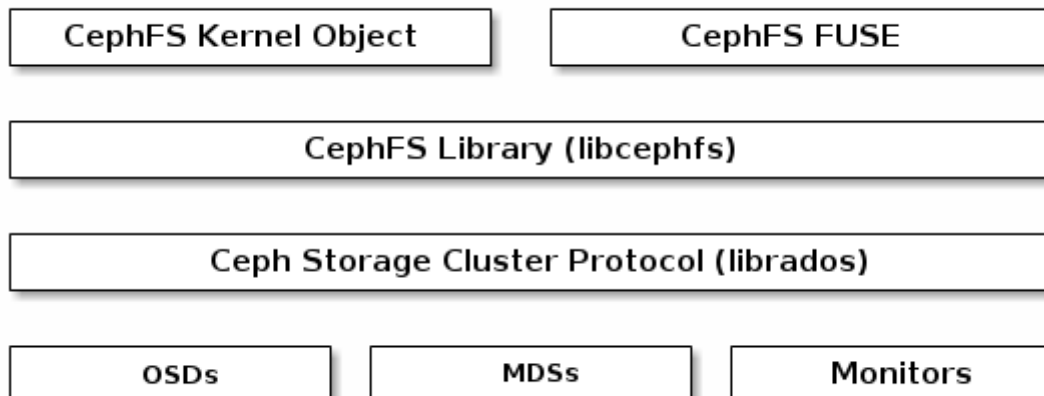
Important: Striping allows RBD block devices to perform better than a single server could!

Thin-provisioned snapshottable Ceph Block Devices are an attractive option for virtualization and cloud computing. In virtual machine scenarios, people typically deploy a Ceph Block Device with the `rbd` network storage driver in QEMU/KVM, where the host machine uses `librbd` to provide a block device service to the guest. Many cloud computing stacks use `libvirt` to integrate with hypervisors. You can use thin-provisioned Ceph Block Devices with QEMU and `libvirt` to support OpenStack and CloudStack among other solutions.

While we do not provide `librbd` support with other hypervisors at this time, you may also use Ceph Block Device kernel objects to provide a block device to a client. Other virtualization technologies such as Xen can access the Ceph Block Device kernel object(s). This is done with the command-line tool `rbd`.

CEPH FILESYSTEM

The Ceph Filesystem (CephFS) provides a POSIX-compliant filesystem as a service that is layered on top of the object-based Ceph Storage Cluster. CephFS files get mapped to objects that Ceph stores in the Ceph Storage Cluster. Ceph Clients mount a CephFS filesystem as a kernel object or as a Filesystem in User Space (FUSE).



The Ceph Filesystem service includes the Ceph Metadata Server (MDS) deployed with the Ceph Storage cluster. The purpose of the MDS is to store all the filesystem metadata (directories, file ownership, access

modes, etc) in high-availability Ceph Metadata Servers where the metadata resides in memory. The reason for the MDS (a daemon called `ceph-mds`) is that simple filesystem operations like listing a directory or changing a directory (`ls, cd`) would tax the Ceph OSD Daemons unnecessarily. So separating the metadata from the data means that the Ceph Filesystem can provide high performance services without taxing the Ceph Storage Cluster.

CephFS separates the metadata from the data, storing the metadata in the MDS, and storing the file data in one or more objects in the Ceph Storage Cluster. The Ceph filesystem aims for POSIX compatibility. `ceph-mds` can run as a single process, or it can be distributed out to multiple physical machines, either for high availability or for scalability.

- **High Availability:** The extra `ceph-mds` instances can be *standby*, ready to take over the duties of any failed `ceph-mds` that was *active*. This is easy because all the data, including the journal, is stored on RADOS. The transition is triggered automatically by `ceph-mon`.
- **Scalability:** Multiple `ceph-mds` instances can be *active*, and they will split the directory tree into subtrees (and shards of a single busy directory), effectively balancing the load amongst all *active* servers.

Combinations of *standby* and *active* etc are possible, for example running 3 *active* `ceph-mds` instances for scaling, and one *standby* instance for high availability.

Mediant™ CE/VE/SE

Mediant CE/VE/SE Software Session Border Controller (SBC)

AudioCodes **Mediant software session border controller (SBC)** is a highly scalable SBC solution supporting broad SIP interoperability, advanced media handling and robust security. AudioCodes Mediant software SBC enables enterprises and service providers to deliver voice services, such as SIP trunking and unified communications, via private or public clouds.



The Mediant software SBC is available in three variants to meet different customer deployment needs:

Mediant CE | a cloud-native SBC delivering high scalability and elasticity in virtualized cloud environments

Mediant VE | built for deployment in virtualized data centers, public clouds and NFV environments

Mediant SE | designed to run on commercial off-the-shelf servers (COTS) in high-scale communications environments



Comprehensive SBC functionality and SIP interoperability

Shared code base with AudioCodes field-proven, hardware-based SBCs



Rapid cloud deployment

Optimized resource consumption for private and public clouds such as Microsoft Azure and Amazon Web Services



NFV-ready

Proven interoperability with leading NFV orchestrators



Enhanced scalability

Easily scale from tens up to tens of thousands of concurrent sessions



High availability

1:1 active-standby configuration for business continuity



High performance and robust security

Built-in software-based media transcoding with support for encryption and protection from attacks



Qualified for leading UC and hosted telephony platforms

Supported environments include Microsoft Teams, Skype for Business and BroadSoft BroadWorks



Integrated WebRTC gateway

Simple and secure WebRTC deployment, supporting both signaling and media

Specifications

Capacities			
	Mediant CE	Mediant VE	Mediant SE
Max. Signaling Sessions	40,000	24,000	70,000
Max. Media Sessions	40,000	24,000	70,000
Max. SRTP-RTP Sessions	40,000	10,000	40,000
Max. Transcoding Sessions	27,000	12,000 *	30,000*
Max. Registered Users	100,000	75,000	500,000
Security			
Access Control	DoS/DDoS line rate protection, bandwidth throttling, dynamic blacklisting		
VoIP Firewall	RTP pinhole management, rogue RTP detection and prevention, SIP message policy, advanced RTP latching		
Encryption and Authentication	TLS, DTLS, SRTP, HTTPS, SSH, client/server SIP Digest authentication, RADIUS Digest		
Privacy	Topology hiding, user privacy		
Traffic Separation	VLAN/physical interface separation for multiple media, control and OAMP interfaces		
Intrusion Detection System	Detection and prevention of VoIP attacks, theft of service and unauthorized access		
Interoperability			
SIP B2BUA	Full SIP transparency, mature and broadly deployed SIP stack, stateful proxy mode		
SIP interworking	3xx redirect, REFER, PRACK, session timer, early media, call hold, delayed offer		
Registration and Authentication	User registration restriction control, registration and authentication on behalf of users, SIP authentication server		
Transport Media	SIP over UDP/TCP/TLS/WebSocket, IPv4 / IPv6, RTP / SRTP (SDES/DTLS)		
Header Manipulation	Ability to add/modify/delete SIP headers and message body using advanced regular expressions (regex)		
URI and Number Manipulations	URI user and host name manipulations, ingress and egress digit manipulation		
Transcoding and Vocoders	Coder normalization including transcoding, coder enforcement and re-prioritization, extensive vocoder support: G.711, G.723.1, G.726, G.729A/B, GSM-FR, AMR-NB, AMR-WB (G.722.2), SILK-NB/WB, Opus-NB/WB		
Signal Conversion	DTMF/RFC 2833/SIP, T.38 fax, packet-time conversion		
WebRTC Controller	Interworking between WebRTC devices and SIP networks. Supports WebSocket, Opus, VP8 video coder, lite ICE, DTLS, RTP multiplexing, secure RTCP with feedback		
NAT	Local and far-end NAT traversal for support of remote workers		
Voice Quality and SLA			
Call Admission Control	Based on bandwidth, session establishment rate, number of connections/registrations		
Packet Marking	802.1p/Q VLAN tagging, DiffServ, TOS		
Standalone Survivability	Maintains local calls in the event of WAN failure.		
Impairment Mitigation	Packet Loss Concealment, Dynamic Programmable Jitter Buffer, Silence Suppression/Comfort Noise Generation, RTP redundancy, broken connection detection		
Voice Enhancement	Transrating, RTPC-XR, Acoustic echo cancellation, replacing voice profile due to impairment detection, Fixed & dynamic voice gain control		
Direct Media	Hair-pinning of local calls to avoid unnecessary media delays and bandwidth consumption while avoiding media anchoring		
Voice Quality Monitoring	RTPC-XR, AudioCodes One Voice Operations Center (OVOC)		
High Availability	SBC high availability with two-box redundancy, active calls preserved		
Quality of Experience	Access control and media quality enhancements based on QoE and bandwidth utilization		
Test agent	Ability to remotely verify connectivity, voice quality and SIP message flow between SIP UAs		
SIP Routing			
Routing Methods	Request URL, IP address, FQDN, ENUM, advanced LDAP, third-party routing control through REST API		
Advanced Routing Criteria	QoE, bandwidth, SIP message (SIP request, coder type, etc.), Layer-3 parameters		
Redundancy	Detection of proxy failures and subsequent routing to alternative proxies		
Routing Features	Least-cost routing, call forking, load balancing, E911 gateway support, emergency call detection and prioritization		
SIPREC	IETF standard SIP recording interface		
Management			
OAM&P	Browser-based GUI, CLI, SNMP, INI Configuration file, REST API, EMS, HTTP reverse proxy		
Multi-tenancy	Advanced multi-tenant SBC partitioning		
Deployment tools	VNF/Stack manager (Mediant CE), HEAT templates, Cloud Formation		
Auto-scaling (CE)	Automatic, REST API, CLI, Web UI		
Cloud Environments			
Public cloud	Azure, AWS		
Private cloud	OpenStack, VMware® vSphere		
Mediant VE SBC Minimum Requirements			
Hypervisor	VMware® vSphere ESXi™ 5.5, & 6.5 and above, Linux KVM, Microsoft Hyper-V	Virtual Resources	1 vCPU; 2 GB RAM; 10 GB Disk; Virtual NICS - 2/3 (Standalone/HA)

* With media transcoding cluster

About AudioCodes

AudioCodes Ltd. (NasdaqGS: AUDC) is a leading vendor of advanced voice networking and media processing solutions for the digital workplace. With a commitment to the human voice deeply embedded in its DNA, AudioCodes enables enterprises and service providers to build and operate all-IP voice networks for unified communications, contact centers and hosted business services. AudioCodes' wide range of innovative products, solutions and services are used by large multinational enterprises and leading tier one operators worldwide.

International Headquarters

1 Hayarden Street,
Airport City
Lod 7019900, Israel
Tel: +972-3-976-4000
Fax: +972-3-976-4040

AudioCodes Inc.

200 Cottontail Lane, Suite A101E
Somerset, NJ 08873
Tel:+1-732-469-0880
Fax:+1-732-469-2298

Contact us: www.audiocodes.com/contact
Website: www.audiocodes.com

©2019 AudioCodes Ltd. All rights reserved. AudioCodes, AC, HD VoIP, HD VoIP Sounds Better, IPmedia, Mediant, MediaPack, What's Inside Matters, OSN, SmartTAP, User Management Pack, VMAS, VoIPerfect, VoIPerfectHD, Your Gateway To VoIP, 3GX, VocaNom, AudioCodes One Voice and CloudBond are trademarks or registered trademarks of AudioCodes Limited. All other products or trademarks are property of their respective owners. Product specifications are subject to change without notice.

Ref. 08/19

