

CloudPond

Service Overview

Issue 01
Date 2024-05-16



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. 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 Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property 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 Cloud 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 Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 What Is CloudPond?	1
2 Advantages	7
3 Application Scenarios	11
4 Security	16
5 Networking Requirements	18
6 Deployment Requirements	22
7 Specifications	27
7.1 Compute Resource Specifications	27
7.1.1 General Computing	28
7.2 Storage Resource Specifications	29
7.2.1 Ultra-High I/O Block Storage	29
8 ECS Specifications for	31
9 Constraints	36
10 Pricing	38
11 Basic Concepts	42
12 CloudPond and Other Services	44
A Change History	46

1 What Is CloudPond?

Background Information

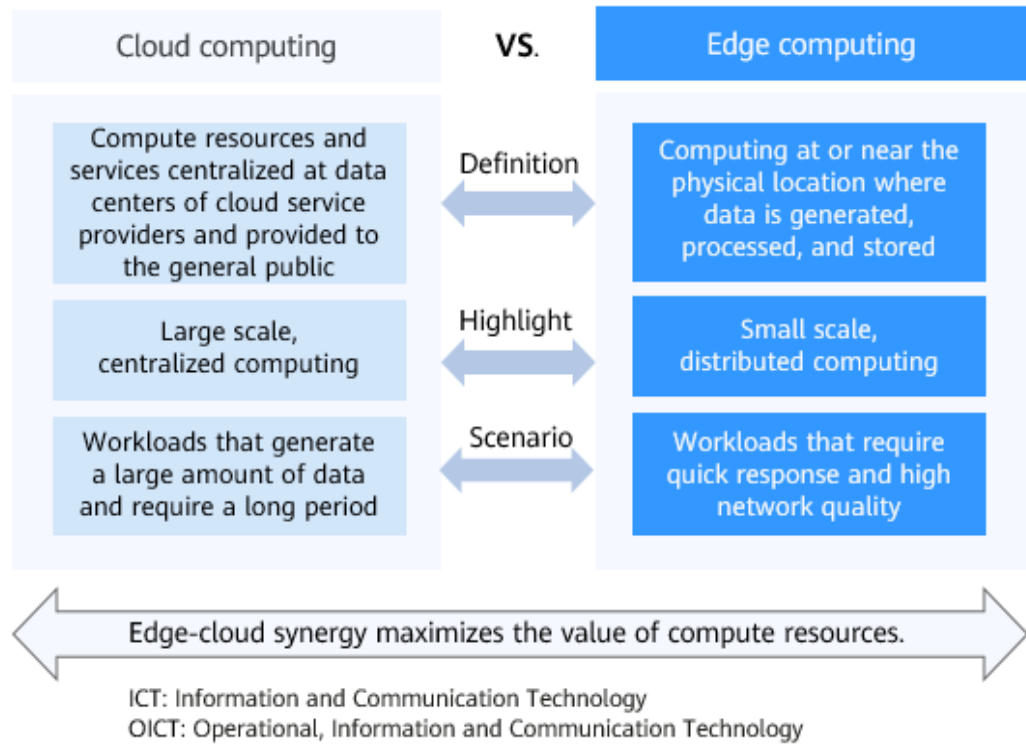
5th generation mobile networks (5G) and artificial intelligence (AI) make innovative services, such as augmented reality (AR), virtual reality (VR), interactive live streaming, automated driving, and smart manufacturing, a reality for the world. These services demand for low latency and high bandwidth. Centralized cloud computing is generally far away from devices and has the following disadvantages:

- High network latency: Centralized cloud computing cannot process and analyze data in a timely manner, resulting in slow application response and poor user experience.
- High bandwidth cost: Data generated by applications needs to be transmitted to the cloud, which consumes large bandwidth. As a result, service providers need to spend a large amount of money on bandwidth.
- Data compliance: Data generated by applications is stored on the cloud, whereas enterprises want sensitive data to be stored on premises to meet the requirement for data compliance.

Edge computing brings computation and data storage closer to where they are needed to improve response speed, save bandwidth, and ensure data compliance.

Figure 1-1 shows the differences between cloud computing and edge computing.

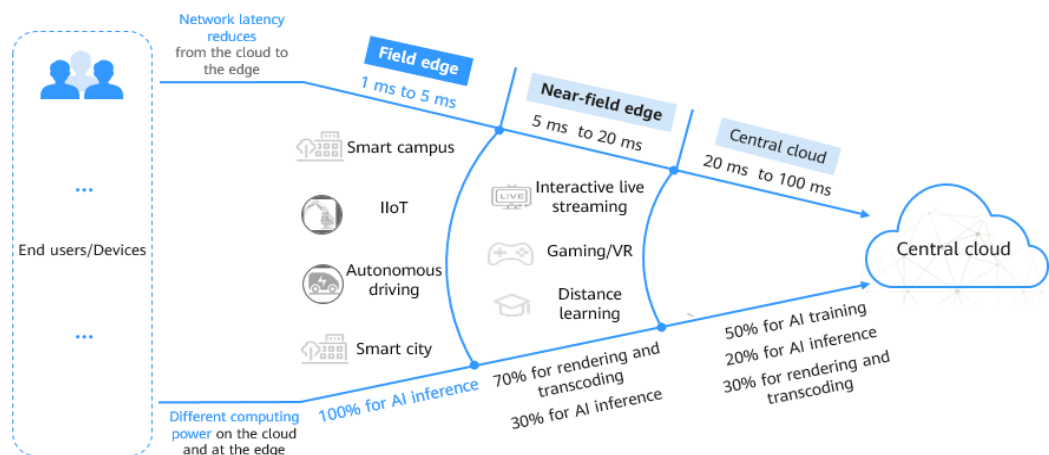
Figure 1-1 Cloud computing and edge computing



Edge computing extends cloud computing to the edge, and the two are mutually complementary. To fit into the requirements in different scenarios, you need to leverage both cloud computing and edge computing to form an edge-cloud synergy.

Figure 1-2 illustrates the scope of edge computing.

Figure 1-2 Scope of edge computing



There are three rough categories of network latency, each defined by their distance from end users. From closest to farthest, these ranges are the field edge, the near-field edge, and the cloud.

- **Field edge:** The network latency ranges from 1 ms to 5 ms, and computing power is primarily used for AI inference. Real-time services, such as autonomous driving and industrial Internet perform well at the field edge.
- **Near-field edge:** The network latency ranges from 5 ms to 20 ms. The computing power is mainly used for rendering, and the rest for AI inference. Video services are positioned to run at the near-field edge.
- **Public cloud (or central cloud):** The network latency is between 20 ms and 100 ms, and workloads that are not delegated to the edge, such as massive data storage, data mining, and training, are generally deployed on the central cloud.

For the field edge scenarios, CloudPond is a good choice that provides integrated hardware and software on your premises.

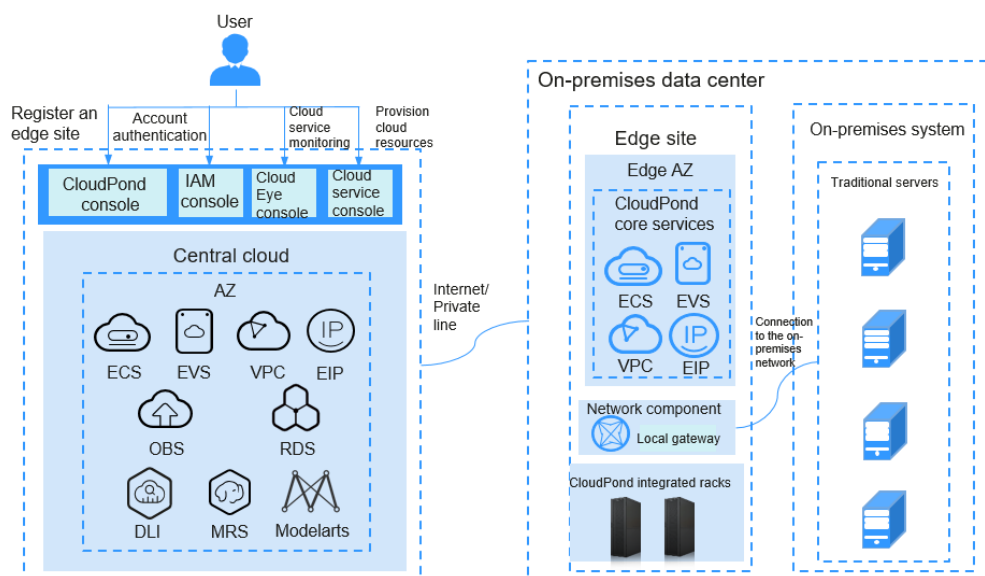
What Is CloudPond?

CloudPond extends the cloud infrastructure and cloud services to your on-premises facility to allow you to enjoy ultra-low latency and faster access to applications on the cloud. It is suitable for scenarios that have high requirements for network latency, local data storage, and interaction with the on-premises system.

For more information, see [How Is CloudPond Related to Huawei Cloud?](#)

Figure 3 shows the CloudPond product architecture. You are advised to familiarize yourself with the [product concepts](#).

Figure 1-3 CloudPond product architecture



- Pre-integrated racks are delivered to your on-premises data center, allowing you to use cloud services locally.
- An **edge site** is deployed in an edge AZ and is fully managed by Huawei Cloud. You can enjoy the same experience as in general AZs. Edge sites, edge AZs, and cloud resources used by the edge sites are dedicated for your

exclusive use. For more information about edge AZs and regions, see [What Are the Relationships Between Edge Sites, Regions, and AZs?](#)

You can deploy multiple edge sites in one on-premises data center or an edge site in each data center, depending on your requirements.

- The CloudPond console is deployed on the central cloud. You can view edge site information and cloud service resource usage on the console.
- CloudPond allows you to run core cloud services on premises, including [Elastic Cloud Server \(ECS\)](#), [Elastic Volume Service \(EVS\)](#), [Virtual Private Cloud \(VPC\)](#), and [Elastic IP](#). In addition to core cloud services, you can deploy additional cloud services as required. For details about cloud services supported by CloudPond, see [Cloud Services Interacting with CloudPond](#). You can also use Identity and Access Management (IAM) to authenticate accounts and Cloud Eye to monitor cloud resources used by edge site.
- An edge site can connect to both your on-premises network and the cloud.
 - You can connect the edge site to your on-premises system. In this way, the edge site can communicate with the on-premises over a private network.
 - The edge AZ where the edge site resides shares the same VPC with the central AZ on the cloud, so that the edge site can access other cloud services on the cloud over a private network.

For more information, see [Network Connectivity](#) and [Networking Requirements](#).

Why CloudPond?

Considering the difficulties of building IT systems, such as long deployment period, difficult O&M, and high operation costs, many enterprises choose to migrate their workloads to the cloud. Choosing an appropriate type of cloud computing depends on legacy systems of enterprises and their requirements for the cloud. There are three main types of cloud computing, including private clouds, public clouds, and hybrid clouds. Each type has their own strengths and shortcomings.

CloudPond provides a brand new form of cloud computing. It inherits the advantages of mainstream cloud computing and features low network latency, low cost, and data localization, addressing the pain points in some scenarios where only one type of cloud computing is used.

For details about differences between CloudPond and private cloud, public cloud, and hybrid cloud, see [Table 1-1](#).

Table 1-1 Comparisons between CloudPond and private clouds, public clouds, and hybrid clouds

Type	CloudPond	Private Cloud	Public Cloud	Hybrid Cloud
------	-----------	---------------	--------------	--------------

Outline	A cloud service that extends the cloud infrastructure and provides compute and storage resources for the edge.	A cloud computing environment that is dedicated to a single customer. The underlying IT infrastructure is dedicated to this customer with completely isolated access.	A cloud computing environment that is typically created from the IT infrastructure owned by the cloud service provider and provides services to the general public over the Internet in pay-per-use or yearly/monthly mode.	A blend of private clouds and public clouds. Customers keep confidential information secured on their private clouds, but make more general, customer-facing content on a public cloud.
Installation and Maintenance	Customers lease pre-integrated racks from Huawei Cloud, who is responsible for the installation and O&M of racks on premises. They only need to ensure the availability and stability of their equipment rooms.	<ul style="list-style-type: none"> • Building a private cloud: Customers purchase the IT infrastructure by themselves and are responsible for installation and O&M. • Hosted private cloud: Customers lease or buy the IT infrastructure from by third-party cloud providers and work with them to install, operate, and maintain the infrastructure . 	Customers can directly use cloud services provided by the cloud providers, who are responsible for the installation and O&M of the IT infrastructure.	Customers lease or buy the IT infrastructure from both the private and public cloud providers and work with them to install and maintain the infrastructure.

Exclusive Use of Cloud Resources	Exclusive use	Exclusive use	Shared with other users	Blend
Data Localization	Data is stored on customer premises, meeting data compliance requirements.	Data is stored in customer's or cloud provider's data center.	Data is stored in the cloud provider's data center.	Partial data is stored locally, and partial data is stored in the cloud provider's data center, depending on customer requirements.
Network Latency	Low (workloads running on customer premises)	Low (workloads running on customer premises or in the cloud provider's data center)	High (workloads running on the public cloud)	Depends on where workloads are running
Cost	Low. The cost mainly consists of the equipment room, energy, cloud services, service support, and personnel (for on-premises system O&M).	High. The cost mainly consists of the equipment room, energy, devices, private cloud software, service support, and personnel (for on-premises system O&M and platform O&M).	Low. The cost consists of cloud services, network, service support, and personnel costs (for on-premises system O&M).	Medium. The cost depends on how much each cloud is used.

For more details, see [Product Advantages](#).

Access Methods

You can access CloudPond using the management console.

For details about operations that can be performed on the management console, see [Getting Started](#).

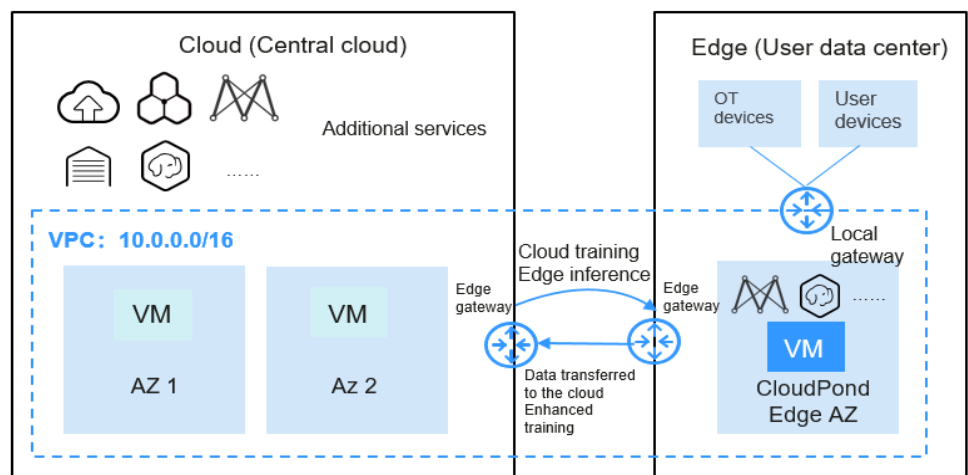
2 Advantages

Homogeneous Architecture, Consistent Experience

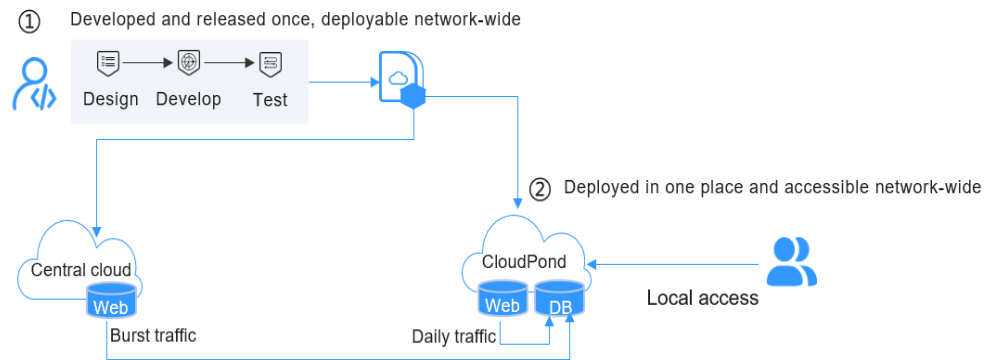
Based on the Huawei Cloud QingTian architecture, CloudPond provides you with a truly consistent experience. The running environment, network, and management system are consistent at the edge and on the cloud. In addition, edge sites share the same VPCs with the cloud, making it easier for cloud resource provisioning and data transmission.

- A distributed network: The edge-cloud distributed network seamlessly connects the on-premises network to the cloud network.

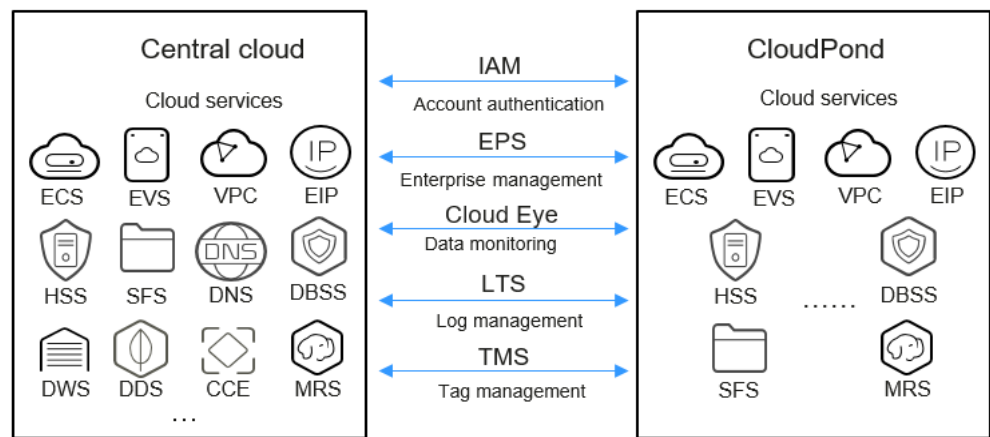
The edge AZ where the edge site resides is connected to the central AZ through edge gateways and shares the same VPC as the central AZ.



- Consistent running environment: The APIs, instances, and service specifications are consistent as those used on the cloud.

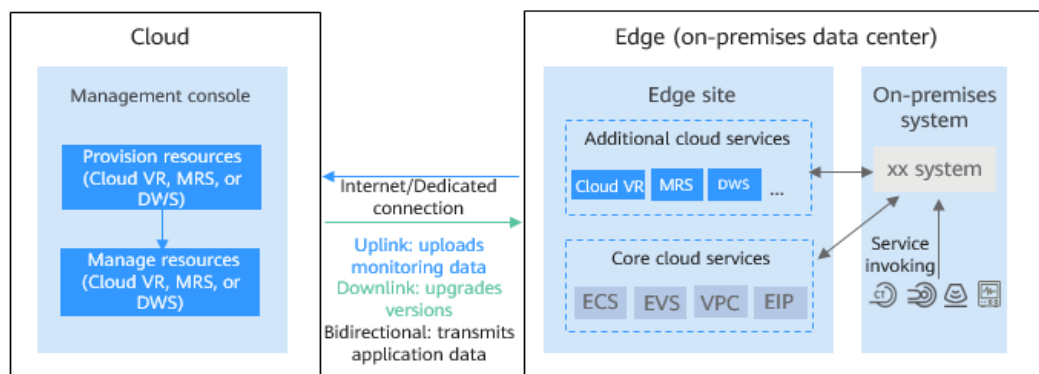


- A management system: The edge uses the same management system as the cloud for multi-grain permission control and resource management. The edge and the cloud use the unified management console.



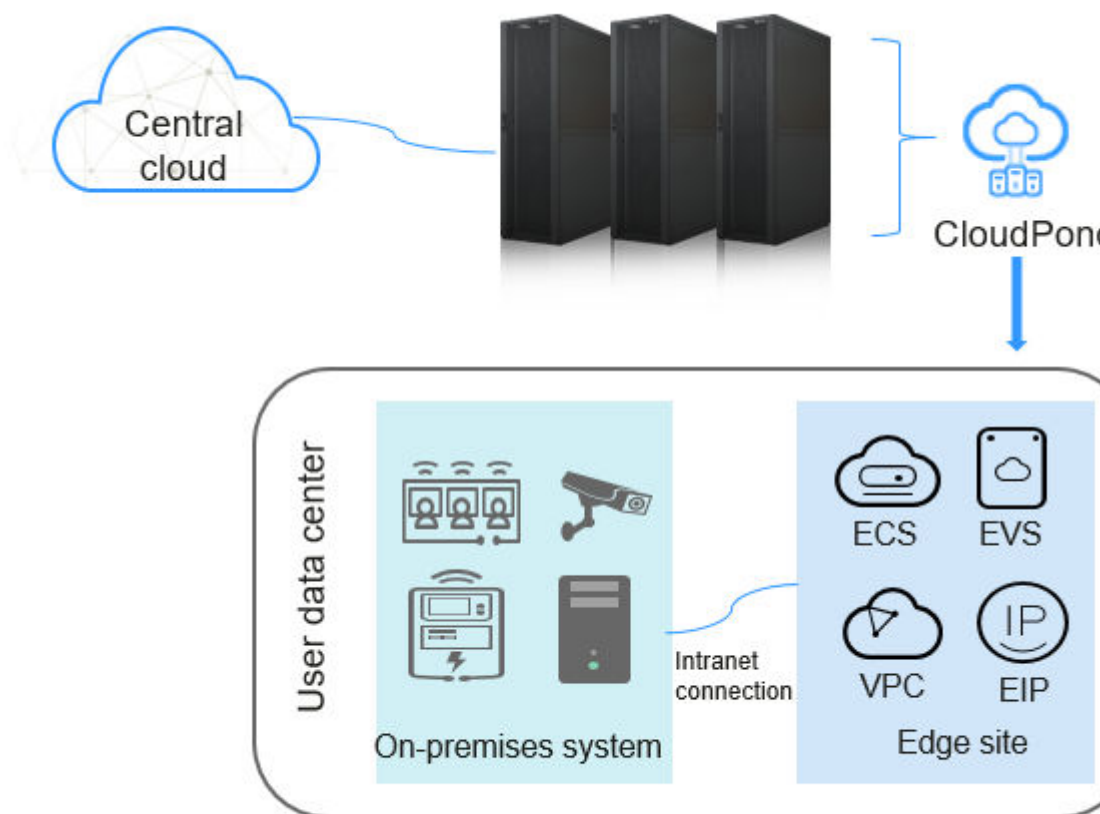
Pay-per-Use Cloud Services

With CloudPond, you can run cloud services locally, such as Cloud VR, MRS, and DWS, and pay for only what you use, meeting your requirements for service innovation at the minimal cost.



Local Data Processing and Storage

CloudPond provides you with a pool of compute and storage resources and allows you to run cloud services locally. With CloudPond, you can process and store data that needs to be retained locally. You can access on-premises resources over a private network. This makes data interconnection easier.

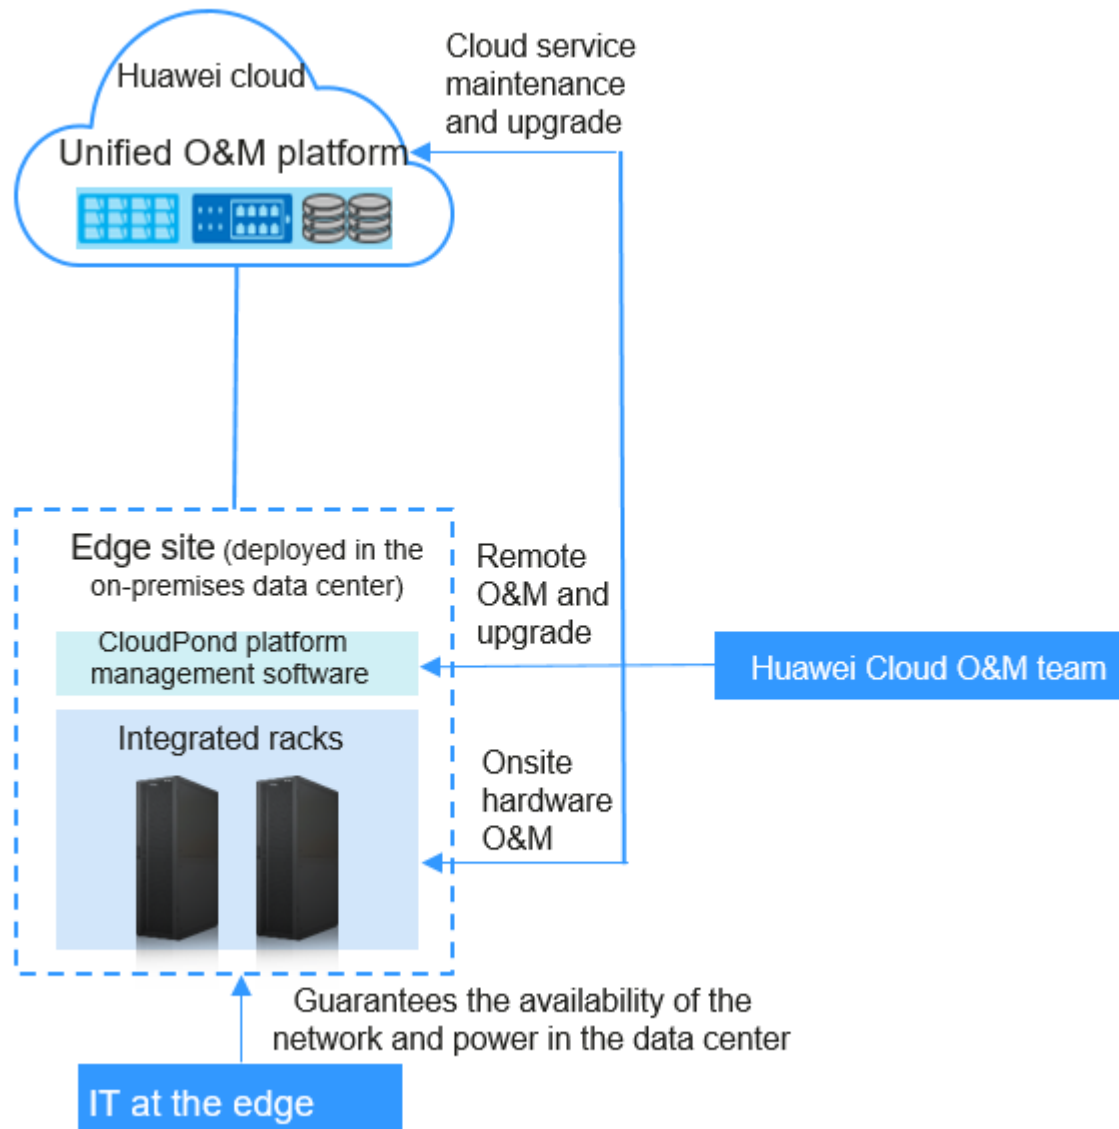


Fully Managed Infrastructure

The racks delivered to your premises are managed and maintained by Huawei Cloud. You do not need to concern about the operation and maintenance of the IT infrastructure and can spend more time and money on service innovation.

- Simplified architecture: There is no independent local management node. The architecture is more lightweight than traditional software products.
- Integrated delivery: Compute and storage resources are delivered to your premises using standard racks and are ready to use within 2 days after being delivered to your premises.
- Periodic upgrades: Software and cloud services used at edge sites are managed using a unified O&M platform and upgraded on a regular basis.
- Managed O&M: The Huawei Cloud O&M team provides onsite O&M for the racks delivered on your premises and replaces faulty hardware on site.

You only need to undertake the O&M of your equipment room, and you do not need a staffed and skilled IT team.



For more information about Huawei Cloud enterprise-class technical support, see [Support Plans](#).

3 Application Scenarios

Low Network Latency

Requirements

Innovative services such as VR, industrial control systems, and IoT have high requirements on network latency. However, centralized cloud computing is generally far away from on-premises devices and cannot meet the requirements for network latency.

To address these issues, enterprises need a new infrastructure that can:

- Provide cloud services locally and consistent experience as the public cloud.
- Allow on-premises devices to access cloud resources over a private network at ultra-low latency.
- Free them from complex O&M to minimize the manpower and cost.

Solution

Deploy an edge site in your on-premises data center so that you can run cloud services locally and develop innovative services at ultra-low latency. See [Figure 3-1](#).

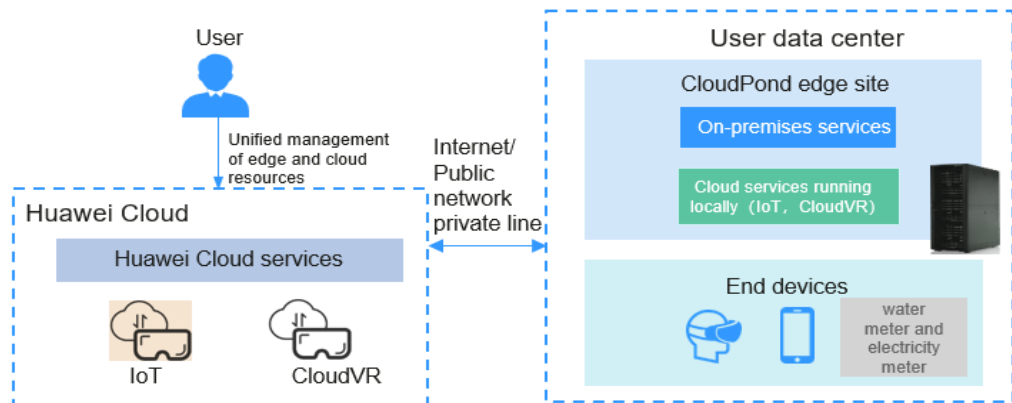
Highlights

- In addition to core cloud services, CloudPond supports a number of additional cloud services to meet your requirements for business innovation.
- Your data is processed and stored locally and can be transmitted to other local devices or applications over your on-premises network.
- Huawei Cloud provides managed infrastructure O&M, enabling you to focus on your own services and innovation.

NOTE

"Managed infrastructure O&M" means that Huawei Cloud is responsible for CloudPond O&M.

Figure 3-1 Low network latency



Local Data Processing and Storage

Requirements

The public cloud provides scalable cloud resources. However, due to restrictions of policies and regulations, data security, service scenarios, and architecture reconstruction, many applications cannot be migrated to the public cloud.

A new infrastructure that has the following features is required:

- It is deployed locally, and the core service data is stored locally.
- It is standardized and can be seamlessly upgraded, batch deployed, and easily managed.
- It is compatible with the legacy IT system and supports cloud migration.

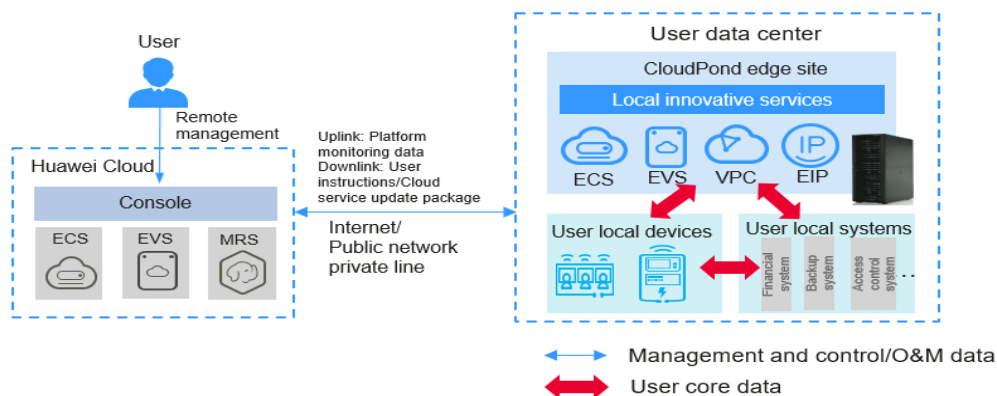
Solution

With pre-integrated racks deployed in your on-premises data center, you can store the service data locally and build a new platform for workloads that cannot be migrated to the cloud.

Highlights

- You can run cloud services and store your data locally.
- Standard racks are delivered as an integrated unit, and cloud services can be used on demand.
- Your on-premises network connects to the edge site using a local gateway, allowing you to easily migrate some challenging applications to the cloud.

Figure 3-2 Local data retention



Enterprise Branch Management

Requirements

Large enterprises often need to set up branches for office, production, warehousing, or R&D. These branches need different applications, such as attendance, finance, logistics, and sales applications.

However, many branches face problems in application running and maintenance due to different underlying infrastructure from the headquarters and complex O&M.

Branches need an infrastructure that allows them to:

- Communicate with the headquarters and transmit the data to the headquarters or other branches more easily.
- Run standard cloud services locally and enjoy consistent experience as the headquarters.
- Free them from burdensome infrastructure O&M.

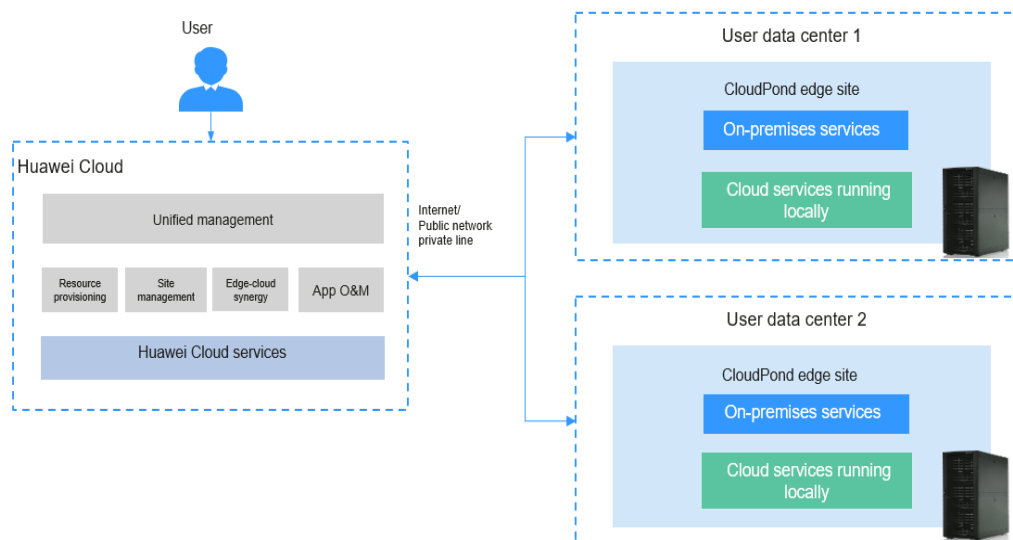
Solution

Deploy an edge site at each branch so that these branches can use the unified infrastructure as the headquarters. See [Figure 3-3](#).

Highlights

- The headquarters' IT team manages cloud resources for branches. The IT team at each branch only needs to operate and maintain the local equipment room. Huawei Cloud is responsible for the O&M of the infrastructure and cloud services.
- Branches use the same cloud services as the headquarters and can have consistent user experience.
- With edge-cloud synergy and private network communications, the data can be quickly transmitted between the headquarters and branches.

Figure 3-3 Enterprise branch management



SaaS Application Distribution

Requirements

SaaS applications are favored by increasingly more government agencies and enterprises due to their advantages such as rich functions and on-demand use. However, the data generated by SaaS applications is stored on the cloud by default, and all application functions depend on the Internet bandwidth, resulting in poor user experience. In addition, deployment solutions provided by some SaaS application vendors are often complex.

To address these issues, a new infrastructure is required to:

- Provide the same cloud services as those on the cloud, so that SaaS applications can be quickly deployed in multiple regions after being tested on the cloud.
- Provide a cross-region communication mechanism to enable SaaS applications to be easily managed and upgraded.
- Minimize the manpower on infrastructure O&M.

Solution

Deploy edge sites in on-premises data centers to enable SaaS software vendors to distribute their applications across regions.

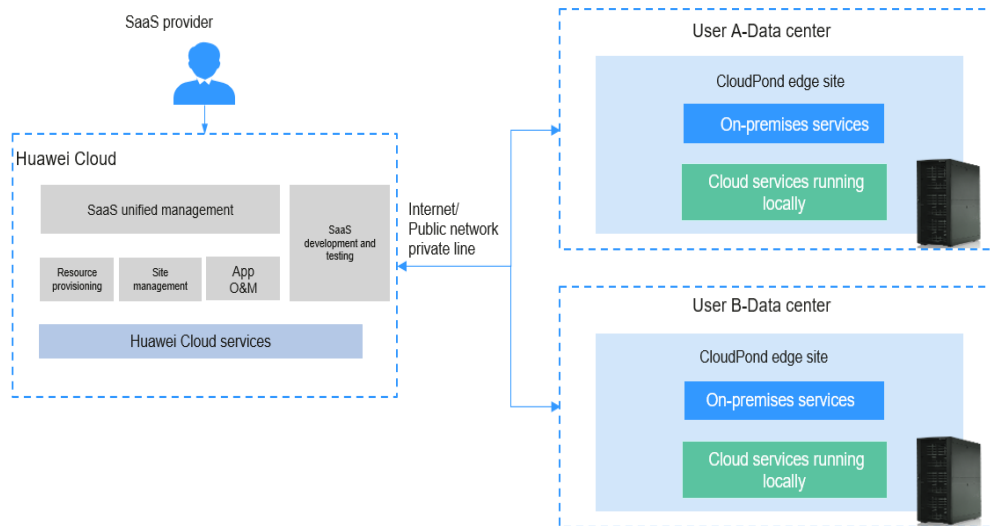
Highlights

- The edge uses the same cloud services as the cloud, including their functions and performance indicators. SaaS application providers can directly deploy SaaS applications at the edge after the applications are tested on the cloud.
- With edge-cloud synergy and private network communication, you can run SaaS applications that are deployed in multiple regions on the cloud for unified management while storing your service data on premises.
- The infrastructure O&M is undertaken by Huawei Cloud, and the infrastructure can be scaled seamlessly.

 NOTE

"Managed infrastructure O&M" means that Huawei Cloud is responsible for CloudPond O&M.

Figure 3-4 SaaS application distribution



4 Security

Shared Responsibility Model

Huawei Cloud is responsible for the infrastructure security. You are responsible only for the security of applications running at the edge sites, the security of the locations of the edge sites, and the security and availability of the networks between edge sites and the cloud.

Communication Security

O&M and service connectivity between edge sites and the cloud is secured as follows:

- O&M connectivity: CloudPond provides a built-in Virtual Private Network (VPN) for ensuring secure transmission of O&M data and management data over the network.
- Service connectivity: Service data is securely transmitted between CloudPond sites and the cloud over the VPC.

Data Security

CloudPond inherits the data security governance requirements of Huawei Cloud. For details, see the [Huawei Cloud Data Security White Paper](#).

- Secure data transmission: Data is encrypted before being transmitted between edge sites and the cloud.
- Secure data use: It is recommended that you identify sensitive data, encrypt it, and use security groups or network access control lists (ACLs) to control access to resources. In addition, you can use [Cloud Trace Service \(CTS\)](#) for auditing.
- Secure data destruction: Before you stop any edge site, dump the data for later use. All data about the edge sites will then be unrecoverably deleted in compliance with industry standards.

Access Control

You can create IAM users to grant minimum permissions required for completing specific tasks, and periodically review the granted permissions. For details, see [IAM Best Practices](#).

CloudPond Infrastructure Security

Huawei Cloud is committed to building a **trusted cloud platform**. CloudPond inherits all security specifications of Huawei Cloud. For details, see the [Huawei Cloud Security White Paper](#).

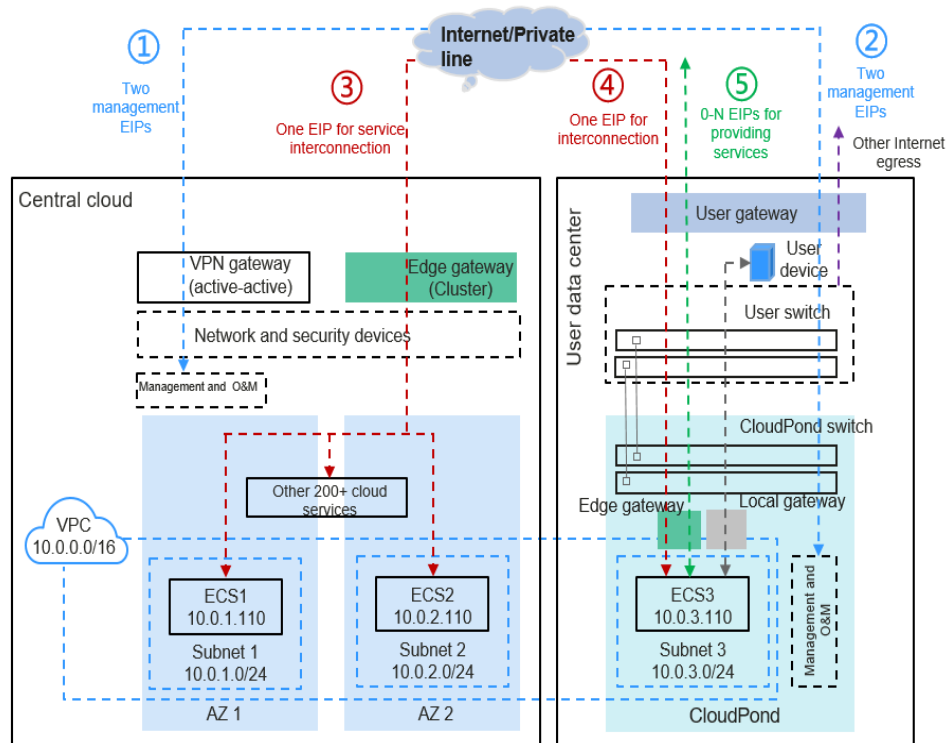
- CloudPond adopts strict security isolation policies at both the user and platform layers. Huawei Cloud provides 24/7 security monitoring and operation services as well as the O&M of the CloudPond platform.
- Host Security Service (HSS) is used to periodically scan for vulnerabilities, and patches are installed in a timely manner based on the Service Level Agreement (SLA) to fix vulnerabilities.
- Huawei Cloud manages CloudPond account permissions in a centralized manner and changes access keys for each account periodically to prevent unauthorized access.

5 Networking Requirements

Network Topology

Learn the definitions of edge gateway, local gateway, and VPN gateway in [Basic Concepts](#) before you proceed with the CloudPond network topology.

Figure 5-1 Network topology



An edge site connects to the cloud through the Internet or a private line. There are four network connections between the edge site and the cloud as shown in [Table 5-1](#).

Table 5-1 CloudPond network connections

Connection Name	Position in the Figure	Description
O&M connection	Blue line connected by public IP addresses (1) and (2)	This connection is used for O&M at the edge site. With this connection, the latest version of cloud services that can run on premises can be pushed to the edge site, platform status information can be collected for easy O&M, and image file used by the ECSs can be cached for the edge site. To ensure network reliability, two public IP addresses are required for dual-channel communications. Your data is not transmitted over this connection, ensuring data security and meeting data compliance requirements.
Service communication connection	Red line connected by public IP addresses (3) and (4)	This connection is used to transmit your data and management data between the edge site and the cloud. To ensure data security, network encryption technologies are used to encrypt your data transmitted over this connection.
Connection for publishing Internet-accessible services	Green line connected using IP address (5)	This connection is used to provide services to Internet users. You need to purchase a public IP address from a carrier.
Connection for on-premises network communications	Black line in the on-premises data centers	This connection is used by the edge site to communicate with other on-premises devices. To enable the edge site to communicate with the on-premises network, you need to configure routes.

Public IP Addresses describes the public IP addresses required for network communications between an edge site and the cloud.

With the four connections, the following can be achieved:

- ECS 3 at the edge site can communicate with ECS 1 and ECS 2 on the cloud over a private network.
- ECS 3 at the edge site can communicate with on-premises network devices over a private network.
- ECS 3 can use IP address (5) to provide web services for Internet users.
- The Huawei Cloud O&M team can use the O&M connection to monitor the edge site status. For example, if the disk of a server is damaged, the Huawei

Cloud O&M team can receive a warning and arrange replacement in a timely manner.

Public IP Addresses

Table 5-2 Public IP addresses for network communications between the edge site and the cloud

No.	Category	Quantity	Provider	Description
(1)	Management EIP used by the cloud	2	Huawei Cloud	Used to connect to the active-active VPN gateways on the cloud
(2)	Management public IP address used at the edge site	2	Customer	Purchase public IP addresses from a carrier for interconnecting with the management plane of the cloud. For each public IP address, the minimum downlink bandwidth and the minimum uplink bandwidth are 100 Mbit/s and 30 Mbit/s, respectively.
(3)	EIP used by the cloud for interconnection with the edge site	1	Customer	The EIP is used by an ECS on the cloud to communicate with the edge site. The bandwidth is at least 10 Mbit/s.
(4)	Edge site public IP address for interconnection with the cloud	1	Customer	Purchase a public IP address and the same bandwidth from the same carrier you select when buying the EIP used by the cloud for interconnection with the edge site.

Table 5-3 Public IP addresses for providing Internet-accessible services

No	Category	Quantity	Provider	Description
(5)	Public IP addresses for service release	0-N	Customer	<p>Used to provide services accessible from the Internet.</p> <ol style="list-style-type: none"> 1. You need to purchase public IP addresses from a carrier. 2. Contact the Huawei Cloud O&M team to record the carrier's line information (public IP address and bandwidth) into CloudPond. 3. Buy an EIP for each of your public IP addresses on the EIP console. 4. Locate your public IP addresses and bind them to the ECSs running at the edge site. <p>For details, see Provisioning Cloud Resources to an Edge Site.</p>

See [Networking Requirements](#) for more details.

Helpful Links

- [Configuring Local Network Connectivity](#)
- [Configuring and Testing Network Connectivity Between an Edge Site and the Cloud](#)
- [Managing Local Gateway Route Tables](#)
- [What Can I Do If the Network Between My Edge Site and the Cloud Is Disconnected? Will My Data Be Lost?](#)

6 Deployment Requirements

Evaluate whether your location is suitable for deploying an edge site based on the following tables.

If your location fails to meet the requirements, you can [register an edge site](#) and [place an order](#) on the console, and then contact the Huawei Cloud technical support for a solution.

Equipment Room

Table 6-1 Equipment room environment

Category	Requirement
Temperature	The ambient temperature ranges from 5°C to 35°C when the altitude of the equipment room is from -60 m to 1800 m. The ambient temperature is reduced by 0.45°C per 100 meters when the altitude of the equipment room ranges from 1800 m to 3000 m.
Humidity	The ambient relative humidity range of 20 to 80 percent is optimal for server reliability and operator comfort, and no condensing is allowed.
Rack layout	Racks are at least 1 m away from walls or any obstacles.
Firefighting equipment	Gas fire extinguishers are provided in the equipment room where the racks are deployed.
Water resistance	Racks should not be installed near toilets, kitchens, or wet areas.
Soot	Racks should not be installed near smoke sources or powder layers.

Table 6-2 Electric and electrical power

Category	Requirement
Power supply availability	≥ 99.982% (no more than 1.6 hours of downtime per year)
24/7 power supply	Edge sites should always be on if there is power supply in the equipment room. Do not power off edge sites daily or periodically.
Mains supply	Dual-channel mains power
Rack power supply	<ul style="list-style-type: none"> • Power of each channel: ≥ 6 kW • Current of each channel: ≥ 32 A • Air circuit breakers connected to racks: 220 V power supply, 1 pole, ≥ 32 A current • Power supply: 220 V AC, 230 V AC, or 240 V AC, 1 PH
Industrial connector	<p>Each cable tray at the rear of racks must be equipped with industrial female connectors. It is recommended that the cable tray is installed at the top of racks. If the cable tray is installed on the bottom of racks, record the information for subsequent evaluation.</p> <p>Industrial connector specifications: IP44 connectors (2P+E) that comply with IEC60309 requirements must be used. Female connectors must be installed in the equipment room, and male connectors (with 3 x 6 mm² cables) must be installed on racks. Connectors can be used for devices with 32 A/220 V power supply.</p>
UPS	<ul style="list-style-type: none"> • 2N redundancy is required for UPSs. • UPSs can provide power supply for at least 15 minutes.
Ground cable of racks	<p>There must be ground copper bars at the rear (top or bottom) of racks, and two screw holes must be reserved for each rack.</p> <p>Ground bars cannot exceed 30 m and must be securely connected to the ground grid in the equipment room. You can use galvanized flat steel with a cross-sectional area of 40 x 4 mm² or larger or a cable with a cross-sectional area of 95 mm² or larger. The flat steel or cable must be insulated, corrosion-proof, and free of any damaged points, and its exposed part must be protected against mechanical damage.</p> <p>The equipment room must use joint grounding. The working ground and the protection ground must share the same ground with the lightning-proof ground. The ground resistance must be less than 10 Ω or meet the requirements of local standards. Each rack has a proper ground point and equipotential grounding is used between racks.</p>

Rack Space

Table 6-3 Rack installation

Category	Requirement
Cabling	Overhead cabling is recommended. The distance between the cable tray or fiber trough and the top of racks should be less than or equal to 1 m. If the distance is greater than 1 m, record the detailed cabling length.
Floor load bearing	The bearing capacity of the floor is greater than 1,000 kg/m ² .
Power	The minimum power for a single rack is 6,000 W.
Installation position	The installation position must be able to accommodate a rack with the dimensions of 2000 mm x 600 mm x 1200 mm (H x W x D). The cold aisle must be 1.2 m or larger in width, and the hot aisle must be 0.6 m or larger in width.
Expandable space	Sufficient space should be reserved if capacity expansion is required later.
Cabling between the edge site and on-premises network devices	Check whether there is a cable trough on the top of racks and whether multi-mode optical fiber can be used for cabling between racks and on-premises network devices. Record the cabling length.

Rack Transportation

Table 6-4 Rack unloading

Category	Requirement
Unloading area	<ul style="list-style-type: none"> The unloading area meets height and ramp requirements for an 8-ton freight carrier truck. The working space must be of sufficient width, depth, and height to allow unloading of racks. The depth must be greater than 6.5 m.

Category	Requirement
Unloading route	<p>Any platforms, hallway, elevators, doorways, turns, and ramps from the delivery site to the final resting position must have sufficient clearance.</p> <ul style="list-style-type: none"> • The road surface for trucks passes must be level (roughness: ± 3 cm). • There are no thresholds or steps along the route from the unpacking room to the rack installation position. • There is no height difference along the route. If the ground is not level, the incline must be less than 8 degrees.
Freight elevator	<ul style="list-style-type: none"> • If racks are transported with packaging, the freight elevator door must be at least 2.4 m high and 1.5 m wide. If racks are transported without packaging, the freight elevator door must be at least 2 m high. • The clearance between the car-platform sill and the hoistway edge of any landing sill must be less than 20 mm. • The load-bearing capacity of the freight elevator must be greater than or equal to 2 tons. • There must be at least two freight elevators in the equipment room.

On-Premises Network

Learning about the [networking requirements](#) can help you quickly understand the requirements for on-premises networks.

Table 6-5 Requirements for on-premises networks

Category	Requirement
Brand and model of network switches	The brand and model of network switches connected to racks can be selected as needed.
Number of switches	One or two switches (If there is only one switch, evaluate the risk of high availability. If there are two racks, stack the switches.)
Optical modules on switches	Multi-mode or single-mode optical modules must be used. Electrical ports are not supported.
Rate and number of optical ports on switches	<ul style="list-style-type: none"> • Optical port rate: 100 Gbit/s (recommended), 40 Gbit/s, 25 Gbit/s, or 10 Gbit/s • Number of ports: 2, 4 (recommended), 8, or 14 (The number of ports varies according to the rate of the optical port.)

Category	Requirement
Public IP address	Four public IP addresses are required for connecting the edge site to the cloud (IP addresses (2), (3), and (4) described in Table 5-3) (2) indicates two public IP addresses and (4) indicates one public IP address you purchase from a carrier, and (3) indicates one IP address you purchase on the EIP console.
Public network egress	At least 200 Mbit/s is required (for management IP addresses and interconnection IP addresses). If you use other connections, ensure that they meet the requirements for network communications between edge sites and the cloud. To avoid network quality problems, do not share the existing public network egress with CloudPond. If these IP addresses need to be shared, ensure that there are no network conflicts.
Public IP address availability	Public IP addresses must be 24/7 available. The network interruption cannot exceed 24 hours.
Network latency	The network latency between the edge site and the cloud must be less than or equal to 150 ms.

7 Specifications

Compute Resource Specifications

There are three types of existing CloudPond computing resources. You can determine the computing power based on your service requirements by referring to [Table 7-1](#). For details about each type, see [Compute Resource Specifications](#).

Table 7-1 Compute resource specifications

Resource Type	Computing Power	ECS Specification
General computing units	C6	For details, see ECS Specifications for .
	S6	
	M6	

Storage Resource Specifications

CloudPond provides block storage resources for you to run your services. You can determine storage resources based on your service requirements by referring to [Table 7-2](#). For details about supported types, see [Storage Resource Specifications](#).

Table 7-2 Storage Resource Specifications

Resource Type	Storage Resource	Disk Specification
Ultra-high I/O	Ultra-high I/O EVS disks	Disk Types and Performance

7.1 Compute Resource Specifications

7.1.1 General Computing

General computing units, such as the C, S, M ECS series, are suitable for enterprises to deploy common services. For details, see [Table 7-3](#).

Table 7-3 General computing units

Name	Computing Power	Description
General Computing Unit-BU-2C6 General Computing Unit-EU-2C6 General Computing Unit-CU-2C6	2C6	Workloads running on 2 C6 ECSs
General Computing Unit-BU-2S6 General Computing Unit-EU-2S6 General Computing Unit-CU-2S6	2S6	Workloads running on 2 S6 ECSs
General Computing Unit-BU-2M6 General Computing Unit-EU-2M6 General Computing Unit-CU-2M6	2M6	Workloads running on 2 M6 ECSs
General Computing Unit-BU-4C6 General Computing Unit-EU-4C6 General Computing Unit-CU-4C6	4C6	Workloads running on 4 C6 ECSs
General Computing Unit-BU-4S6 General Computing Unit-EU-4S6 General Computing Unit-CU-4S6	4S6	Workloads running on 4 S6 ECSs
General Computing Unit-BU-4M6 General Computing Unit-EU-4M6 General Computing Unit-CU-4M6	4M6	Workloads running on 4 M6 ECSs

Name	Computing Power	Description
General Computing Unit-EU-6C6 General Computing Unit-CU-6C6	6C6	Workloads running on 6 C6 ECSs
General Computing Unit-EU-6S6 General Computing Unit-CU-6S6	6S6	Workloads running on 6 S6 ECSs
General Computing Unit-EU-6M6 General Computing Unit-CU-6M6	6M6	Workloads running on 6 M6 ECSs
General Computing Unit-EU-8C6 General Computing Unit-CU-8C6	8C6	Workloads running on 8 C6 ECSs
General Computing Unit-EU-8S6 General Computing Unit-CU-8S6	8S6	Workloads running on 8 S6 ECSs
General Computing Unit-EU-8M6 General Computing Unit-CU-8M6	8M6	Workloads running on 8 M6 ECSs
General Computing Unit-CU-10C6	10C6	Workloads running on 10 C6 ECSs
General Computing Unit-CU-10S6	10S6	Workloads running on 10 S6 ECSs
General Computing Unit-CU-10M6	10M6	Workloads running on 10 M6 ECSs

7.2 Storage Resource Specifications

7.2.1 Ultra-High I/O Block Storage

Ultra-high I/O EVS disks are suitable for services that require high performance. For details, see [Table 7-4](#).

Table 7-4 Ultra-high I/O

Name	Type	Minimum Capacity	Description
Ultra-high I/O - 17 TiB	Ultra-high I/O	17 TiB	Services that require small capacity, ultra-high I/O EVS disks
Ultra-high I/O - 34 TiB	Ultra-high I/O	34 TiB	Services that require medium capacity, ultra-high I/O EVS disks
Ultra-high I/O - 51 TiB	Ultra-high I/O	51 TiB	Services that require large capacity, ultra-high I/O EVS disks
Ultra-high I/O - 68 TiB or larger	Ultra-high I/O	68 TiB	Services that require super-large capacity, ultra-high I/O EVS disks

 **NOTE**

1. Scaled-out storage will be billed at the unit price applied when you placed the CloudPond order.
For example, you purchased a 17 TiB high I/O storage unit and then expanded it to 51 TiB. The scale-out storage will be billed at the unit price of a 17 TiB high I/O unit.
2. Block storage hardware models support x86 or arm CPUs.

8 ECS Specifications for

This section describes the specifications supported by .

- [General Computing](#)
- [General Computing-plus](#)
- [Memory-Optimized](#)

For details, see [Specifications](#).

For details about how to purchase an ECS, see [Creating an ECS at an edge site](#).

General Computing

Table 8-1 S6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
s6.small.1	1	1	0.8/0.1	10	1	2	KVM
s6.medium.2	1	2	0.8/0.1	10	1	2	KVM
s6.large.2	2	4	1.5/0.2	15	1	2	KVM
s6.xlarge.2	4	8	2/0.35	25	1	2	KVM
s6.2xlarge.2	8	16	3/0.75	50	2	2	KVM
s6.medium.4	1	4	0.8/0.1	10	1	2	KVM

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
s6.large.4	2	8	1.5/0.2	15	1	2	KVM
s6.xlarge.4	4	16	2/0.35	25	1	2	KVM
s6.2xlarge.4	8	32	3/0.75	50	2	2	KVM

General Computing-plus

Table 8-2 C6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
c6.large.2	2	4	4/1.2	40	2	2	KVM/QingTian
c6.xlarge.2	4	8	8/2.4	80	2	3	KVM/QingTian
c6.2xlarge.2	8	16	15/4.5	150	4	4	KVM/QingTian
c6.3xlarge.2	12	24	17/7	200	4	6	KVM/QingTian
c6.4xlarge.2	16	32	20/9	280	8	8	KVM/QingTian
c6.6xlarge.2	24	48	25/14	400	8	8	KVM/QingTian
c6.8xlarge.2	32	64	30/18	550	16	8	KVM/QingTian

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
c6.12xlarge.2	48	96	35/27	750	16	8	KVM/QingTian
c6.16xlarge.2	64	128	40/36	1000	32	8	KVM/QingTian
c6.22xlarge.2.physical	88	192	44/40	1000	16	33	BMS
c6.large.4	2	8	4/1.2	40	2	2	KVM/QingTian
c6.xlarge.4	4	16	8/2.4	80	2	3	KVM/QingTian
c6.2xlarge.4	8	32	15/4.5	150	4	4	KVM/QingTian
c6.3xlarge.4	12	48	17/7	200	4	6	KVM/QingTian
c6.4xlarge.4	16	64	20/9	280	8	8	KVM/QingTian
c6.6xlarge.4	24	96	25/14	400	8	8	KVM/QingTian
c6.8xlarge.4	32	128	30/18	550	16	8	KVM/QingTian
c6.12xlarge.4	48	192	35/27	750	16	8	KVM/QingTian
c6.16xlarge.4	64	256	40/36	1000	32	8	KVM/QingTian

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
c6.22xlarge.4	88	352	44/40	1200	28	8	KVM/QingTian

Memory-Optimized

Table 8-3 M6 ECS specifications

Flavor	vCPUs	Memory (GiB)	Max./Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
m6.large.8	2	16	4/1.2	40	2	2	KVM
m6.xlarge.8	4	32	8/2.4	80	2	3	KVM
m6.2xlarge.8	8	64	15/4.5	150	4	4	KVM
m6.3xlarge.8	12	96	17/7	200	4	6	KVM
m6.4xlarge.8	16	128	20/9	280	8	8	KVM
m6.6xlarge.8	24	192	25/14	400	8	8	KVM
m6.8xlarge.8	32	256	30/18	550	16	8	KVM
m6.12xlarge.8	48	384	35/27	750	16	8	KVM

Flavor	vCPUs	Memory (GiB)	Max./ Assured Bandwidth (Gbit/s)	Max. PPS (x 10,000)	NIC Multi-Queue	Max. NICs	Virtualization Type
m6.16xlarge.8	64	512	40/36	1000	32	8	KVM

9 Constraints

The following are constraints and limitations of CloudPond and its functions. For details about restrictions on parameter settings, see corresponding sections in the [CloudPond User Guide](#).

General Constraints

- Management mode: You can use only the CloudPond console to manage locations, edge sites, and cloud resources used at edge sites.
- Hardware restrictions: User devices cannot be installed in CloudPond racks.
- O&M and monitoring: O&M and monitoring are completed by the Huawei Cloud O&M team. You can use Cloud Eye to monitor cloud resources but cannot access the O&M and monitoring platforms.
- Hardware types and resource usage restrictions: Hardware devices used to provide EVS services at an edge site include converged models (converged nodes provide EVS services and cloud platform O&M management and control component services) and storage models (storage nodes provide EVS services for you). Up to four converged nodes can be deployed at an edge site.

Hardware devices used to provide ECS services at an edge site include virtualization models and bare metal models based on ECS flavors. Except for necessary management overheads, all resources on the hardware devices are provided for service running.

Compute and storage resources are separated. You can expand compute or storage resources as needed (except for local disks provided by D/I series ECSs)

Core Cloud Services

[Table 9-1](#) describes the constraints on running ECS, EVS, VPC, and EIP at an edge site.

Table 9-1 Constraints on core cloud services at edge sites

Cloud Service	Constraint
ECS	<ul style="list-style-type: none"> • ECSs do not support Virtual Network Console (VNC) login. Log in to the ECSs using SSH (Linux) or mstsc (Windows). • An annual license is required for each graphics card used by a GPU-accelerated P12 ECS. • ECSs cannot be encrypted, backed up, or used for creating full-ECS images, and their billing modes cannot be changed. • To run commercial OSs on an ECS, such as Windows and Red Hat, you need to acquire the license. • ECS supports the UEFI secure boot of local VMs. By default, this feature can be enabled for C7n compute instances in CloudPond. Other resources are not supported.
EVS	<ul style="list-style-type: none"> • Only high I/O and ultra-high I/O EVS disks are available. • Currently, high I/O EVS disks use HDDs, and ultra-high I/O EVS disks use NVMe SSDs. You cannot customize the media type and size. • Currently, EVS disks are deployed in three-copy redundancy. • EVS disks cannot be backed up or encrypted. • Currently, block storage snapshots of the EVS service can be stored only at the edge.
VPC	<ul style="list-style-type: none"> • If two edge sites are in different VPCs, VPC peering cannot be used to connect the two VPCs. • If two edge sites are in the same VPC, the subnet used by each edge site cannot communicate with each other.
EIP	<ul style="list-style-type: none"> • An EIP in a central AZ cannot be bound to an ECS in an edge AZ, or vice versa. • An EIP used at an edge site cannot be bound to the ECS that is running at another edge site. • An EIP used at an edge site cannot be added to a shared bandwidth in a central AZ, and vice versa.
CloudPond encryption mode	By default, the commercial encryption solution is used. In specific configuration and feature scenarios, SM series cryptographic algorithms can be used (evaluated based on projects).

10 Pricing

Overview

CloudPond provides you with various resources of different billing modes.

Billing Items

CloudPond cloud services consist of core services and advanced services. The billing items of core services consist of computing resources, storage resources, network resources (free), and Elastic IP (EIP) used for communication between edge sites and the central cloud.

For details, see [Table 10-1](#).

Table 10-1 Billing items

Category	Billing Item	Description
Core cloud services	Compute resources	You pay only for the compute resources required for running your workloads on premises, including vCPUs and memory. For details, see Compute Resource Specifications .
	Storage resources	You pay only for the disks storing your service data on premises. There are four tiers of fixed storage capacity. For details, see Storage Resource Specifications .

Category	Billing Item	Description
	Network resources (free)	<p>Network resources are free of charge.</p> <ul style="list-style-type: none"> Local gateways: used for internal network communication between edge sites and on-premises systems. Virtual Private Cloud (VPC): provides VPCs and subnets (including edge AZs). EIP: allows you to launch local services using public IP addresses of your data center. <p>NOTE The EIP service does not provide Huawei Cloud public IP addresses. You need to bind public IP addresses of your data center before you use the EIP service.</p>
	EIP (communication between edge sites and the central cloud)	You need to purchase resources in the EIP service for communication between edge sites and the central cloud.

For CloudPond advanced services, their billing modes and prices of each specification vary by services.

Payment Methods

[Table 10-2](#) describes the payment methods of each billing item.

Table 10-2 Payment methods

Billing Item	Payment Method	Description
<p>Compute resources</p> <p>Storage resources</p>	<ul style="list-style-type: none"> • Full upfront: You pay for the three-year contract at a time. • Partial upfront + monthly subscription: You pay about 50% at the time of purchase, and the remaining fee is paid in installments within 36 months after site acceptance. 	<ul style="list-style-type: none"> • You do not need to pay for the compute and storage resources within the capacity. For more information, see Completing the Payment. • After all hardware and software are installed and commissioned, the Huawei Cloud technical support team will notify you to carry out an acceptance within three working days. The contract term starts after Huawei Cloud receives your confirmation in writing.
EIP (communication between edge sites and the central cloud)	Yearly/ Monthly subscription or pay-per-use	For details, see EIP Billing .

Capacity Expansion

You can log in to the console to expand edge site resources as needed.

Unsubscription

Contact Huawei Cloud sales team about the specific requirements and arrange the subsequent process together.

Renewal

- If you need to renew the contract upon expiration, communicate with the Huawei Cloud sales team.
- If you use CloudPond advanced services or other services on the central cloud, renew them based on their renewal mode.

Overdue Payment

- If your account is in arrears, you can view the arrears details. For details, see [Top-Up and Repayment](#).
- If your account is in arrears, the purchased CloudPond resources are still available.

11 Basic Concepts

Edge Site

An edge site is a pool of compute and storage resources that are deployed on premises. An edge site associates a specific Huawei Cloud region with your premises.

Location

A location is where your edge site is deployed.

Edge AZ

An edge AZ is where an edge site resides. Edge AZs are deployed in on-premises data centers and are not shared with the general public. For more information, see [What Are the Relationships Between Edge Sites, Regions, and AZs?](#)

Local Gateway

A local gateway is a virtual gateway deployed at the edge site for connecting the edge site to the on-premises system, so that the on-premises network can access the cloud over a private network. You can view details about local gateways and configure routes on the CloudPond console.

Edge Gateway

Edge gateways are deployed in pairs to connect the edge site to the cloud so that your data can be transmitted over a VPC after being encrypted.

VPN Gateway

A VPN gateway is deployed on the cloud in clusters and used for interconnecting the CloudPond O&M network with the cloud and the edge site.

Customer Gateway

A customer gateway can be a router or firewall that connects to both the Internet and CloudPond pre-integrated racks delivered to your location for network connectivity of your on-premises system.

12 CloudPond and Other Services

Cloud Services on Which CloudPond Depends

[Table 12-1](#) describes the cloud services on which CloudPond depends.

Table 12-1 CloudPond dependent services

Cloud Service	Function	Reference
IMS	Create ECSs using images that are automatically cached for CloudPond. When you use an image provided by IMS to create an ECS at the edge site for the first time, the image will be automatically cached at the edge site. When you want to create an ECS later, the image cached at the edge site will be used. This shortens the time required for creating the ECSs.	Cached Images

Cloud Services Interacting with CloudPond

[Table 12-2](#) describes the cloud services interacting with CloudPond.

Table 12-2 Cloud services interacting with CloudPond

Category	Cloud Service	Interaction	How to Use	Reference
Core cloud services	<ul style="list-style-type: none"> • ECS • EVS • VPC • EIP 	Provision ECSs, EVS disks, VPCs, and EIPs for your edge site and run cloud services on premises for low latency and local data storage.	On the cloud service console, provision the cloud resources and select Edge AZ (ECSs, EVS disks, or VPCs) or Edge path (EIP) when you provision the cloud resources.	<ul style="list-style-type: none"> • Provisioning Cloud Resources to an Edge Site • Managing Cloud Resources at an Edge Site

In addition to the preceding cloud services, CloudPond uses the same monitoring and management system as the cloud and supports cloud services such as Cloud Eye, Identity and Access Management (IAM), Enterprise Project Management Service (EPS), Log Tank Service (LTS), and Tag Management Service (TMS).

A Change History

Released On	Description
2024-05-16	This issue is the first official release.