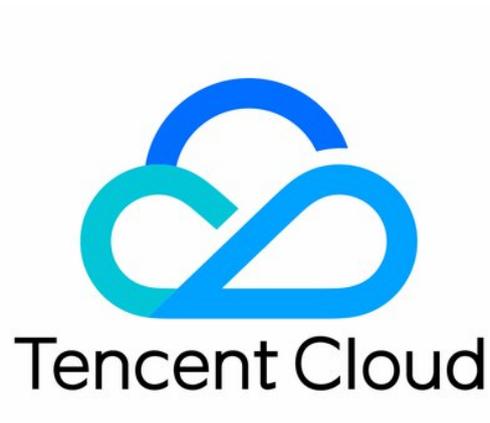


# Cloud Virtual Machine Product Introduction



## Copyright Notice

©2013–2025 Tencent Cloud. All rights reserved.

The complete copyright of this document, including all text, data, images, and other content, is solely and exclusively owned by Tencent Cloud Computing (Beijing) Co., Ltd. ("Tencent Cloud"); Without prior explicit written permission from Tencent Cloud, no entity shall reproduce, modify, use, plagiarize, or disseminate the entire or partial content of this document in any form. Such actions constitute an infringement of Tencent Cloud's copyright, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

## Trademark Notice



This trademark and its related service trademarks are owned by Tencent Cloud Computing (Beijing) Co., Ltd. and its affiliated companies ("Tencent Cloud"). The trademarks of third parties mentioned in this document are the property of their respective owners under the applicable laws. Without the written permission of Tencent Cloud and the relevant trademark rights owners, no entity shall use, reproduce, modify, disseminate, or copy the trademarks as mentioned above in any way. Any such actions will constitute an infringement of Tencent Cloud's and the relevant owners' trademark rights, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

## Service Notice

This document provides an overview of the as-is details of Tencent Cloud's products and services in their entirety or part. The descriptions of certain products and services may be subject to adjustments from time to time.

The commercial contract concluded by you and Tencent Cloud will provide the specific types of Tencent Cloud products and services you purchase and the service standards. Unless otherwise agreed upon by both parties, Tencent Cloud does not make any explicit or implied commitments or warranties regarding the content of this document.

## Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

# Contents

## Product Introduction

CVM Overview

Advantages

Regions and AZs

Instance

Instance Overview

Instance Types

Lifecycle

Spot Instance

Storage

Storage Overview

Cloud Block Storage

COS

Image

Image Overview

Image Types

Public Images

Public Images List

Official Maintenance End Plans of Operating Systems

TencentOS Server

OpenCloudOS

OpenCloudOS Overview

OpenCloudOS V8 Image Update Log

Network and Security

Overview

Placement Group

Private Network Access

Internet Access

Elastic Public IP

Login Password

SSH Keys

Cloud Workload Protection Platform

Anti-DDoS Basic

Ops and Monitoring

Health Check

Monitoring and Alarming

## Access Management

[CAM Overview](#)

[Authorizable Resource Type](#)

[Authorization Policy Syntax](#)

[CAM-Supported APIs](#)

# Product Introduction

## CVM Overview

Last updated: 2024-06-02 16:33:31

### CVM Overview

Tencent Cloud's Cloud Virtual Machine (CVM) is a scalable computing service provided by Tencent Cloud. Utilizing CVMs eliminates the need to estimate resource usage and upfront investment associated with traditional servers, enabling you to quickly launch any number of cloud servers and promptly deploy applications.

CVMs allow users to customize all resources, including CPU, memory, storage, network, and security, and easily adjust them as needs change.

### Using CVM instances

Tencent Cloud allows you to configure and manage CVM instances in the following ways:

- **Console:** A web service interface provided by Tencent Cloud for configuring and managing CVM instances.
- **API:** Tencent Cloud also offers API interfaces for managing CVM instances. For API documentation, please refer to [API Overview](#).
- **SDK:** You can use [SDK programming](#) or Tencent Cloud's [Command Line Tool TCCLI](#) to call CVM APIs.

### Relevant Concepts

Before using Tencent Cloud CVM, you should familiarize yourself with the following concepts:

Concept	Note
<a href="#">Instance</a>	Virtual computing resources in the cloud, including fundamental components such as CPU, memory, operating system, network, and storage. Tencent Cloud offers various CPU, memory, storage, and network configurations for CVM instances. For more information, see <a href="#">Instance Types</a> .
<a href="#">Image</a>	Pre-built templates for CVM instances include pre-configured operating systems and pre-installed software. CVM offers a variety of pre-built images, such as Windows and Linux.
<a href="#">Cloud Disk</a>	A distributed and persistent block storage device provided by Tencent Cloud that can serve as the system disk or an expandable data disk of an instance.

<b>Virtual Private Cloud</b>	Tencent Cloud offers a virtual and isolated network space, logically separated from other resources.
<b>IP Addresses</b>	Tencent Cloud provides <a href="#">Private IPs</a> and <a href="#">Public IPs</a> . In simple terms, private IPs offer Local Area Network (LAN) services for communication between CVM instances, while public IPs are used when users need to access Internet services on their CVM instances.
<b>Elastic IP</b>	Static public network IP addresses designed especially for dynamic networks to meet the demands for fast troubleshooting.
<b>Security Group</b>	A virtual firewall that can check the status and filter data packets. As an important means for network security isolation, it can be used to set network access controls for one or more CVM instances.

## Purchasing and Customizing CVM Instances

If you are using a personal account and purchasing CVM instances for the first time, we recommend the following to get you started:

- [Quickly Customizing Windows CVM Configurations](#)
- [Quick Configuration for Linux CVM instances](#)

If you have specific needs that our standard CVM specifications cannot meet, use the following guides to learn how to obtain custom configurations:

- [Custom Configuration for Windows CVM](#)
- [Custom Configuration for Linux CVM](#)

## CVM Prices

CVM offers both monthly subscription and pay-as-you-go billing options. For more information, see [Pricing Overview](#).

For pricing details of CVM and related resources, see [Product Pricing](#).

## Other related products

- You can use auto-scaling to automatically increase or decrease the number of server clusters based on a schedule or conditions. For more information, please refer to the [Auto-Scaling Product Documentation](#).
- You can use CLB to automatically allocate the request traffic of clients across CVM instances. For more information, see [Cloud Load Balancer](#).
- You can use Tencent Kubernetes Engine (TKE) to manage the application lifecycle on a group of CVM instances. For more information, see [TKE Product Documentation](#).

- You can use Tencent Cloud Observability Platform to monitor CVM instances and their system disks. For more information, please refer to [Tencent Cloud Observability Platform Documentation](#).
- You can deploy relational databases in the cloud or use Tencent Cloud's TencentDB. For more information, please see [TencentDB for MySQL Product Documentation](#).

# Advantages

Last updated: 2024-03-26 15:46:11

## Comprehensive

Tencent Cloud CVM offers a wide range of options

- **Multiple Regions and Availability Zones:** Covering various regions in China, Asia-Pacific, Europe, and the Americas. Deploying applications in regions close to your users can result in lower latency.
- **Various Model Configurations:**
  - **Standard Type:** Suitable for small and medium-sized web applications and databases.
  - **Memory Type:** Suitable for applications requiring extensive memory operations, searches, and calculations.
  - **High I/O Type:** Suitable for low-latency, I/O-intensive applications.
  - **Compute Type:** Ideal for large game servers, advertising service engines, high-performance computing, and other compute-intensive applications.
  - **Big Data Type:** Ideal for throughput-intensive applications such as Hadoop distributed computing, massive log processing, distributed file systems, and large-scale data warehouses.
  - **Heterogeneous Type:** Ideal for high-performance applications such as deep learning, scientific computing, video encoding/decoding, and graphics workstations.
  - **Batch Type:** Ideal for compute-intensive applications requiring frequent use of large-scale computing nodes, such as rendering, gene analysis, and crystallography.

## Elasticity

Dedicated to creating a cloud server management platform with industry-leading elasticity, offering the following capabilities:

- **Hardware Configuration:** Cloud Virtual Machines (CVMs) based on cloud disks can **instantly increase or decrease** hardware configurations (regardless of whether they are billed monthly, annually, or pay-as-you-go).
- **Disk Changes:** Cloud Virtual Machines (CVMs) based on cloud disks can **instantly expand** their disk capacity (regardless of whether they are billed monthly or pay-as-you-go).
- **Network Bandwidth:** CVM instances support **instant upgrade/downgrade** of bandwidth.
- **Billing Mode:** CVMs support switching between bandwidth billing mode and traffic billing mode.
- **Operating System:** Cloud servers in Mainland China can switch between Windows and Linux systems at any time (regardless of monthly or pay-as-you-go billing), while other regions

do not currently support switching between the two.

- **Elastic IP:** Supports binding to servers in various network environments.
- **Image Types:** Public images, service marketplace images, and custom images, with support for cross-region adjustments and image replication. For more information on image types, please refer to [Image Types](#).
- **Custom Network Architecture:** Virtual Private Cloud (VPC) offers an independent network space for users, allowing customization of IP ranges, IP addresses, and routing policies. It provides port-level access control and comprehensive logical network isolation. For more information, see [Virtual Private Cloud](#).

## Reliable

Dedicated to providing industry-leading reliability in cloud virtual machines.

- **CVM Reliability:** Single instance service availability of 99.975% and data reliability of 99.9999999%. Features such as seamless failover migration, data snapshots, and automatic alerts safeguard your server.
- **Cloud Disk Policy:** Offers a three-copy Professional storage strategy to eliminate single points of failure, ensuring data reliability, allowing you to confidently store data in the cloud without worrying about data loss.
- **Stable Network Architecture:** Mature network virtualization and NIC binding technologies ensure high network availability. Running in T3+ data centers guarantees the reliability of the operating environment, freeing you from concerns about network availability.

## Fast

We are committed to providing fast and convenient services for both user operations and CVM performance.

- **Convenient and Fast Operation:** In just a few minutes, you can easily obtain one, hundreds, or even thousands of server instances. You can purchase, configure, manage, and scale your services with a single click.
- **High-Speed Public Network Quality:** Over 20 BGP public networks (including major ISPs, education networks, and broadcasting networks), covering almost all network operators. Regardless of the ISP used by your customers, they can enjoy the same high-speed bandwidth and near-instantaneous failover experience.
- **High-Speed Intranet Quality:** Tencent Cloud provides interconnectivity within the same region data centers, with underlying Gigabit or 10-Gigabit networks ensuring high-quality intranet communication.

## Secure service

Tencent Cloud provides multiple solutions to ensure the CVM security. In addition, the backup and rollback mechanisms can protect data security.

- **Multiple Remote Login Methods for Cloud Servers:** Various login methods are available, including SSH key login, password login, and VNC login.
- **Robust Security Services:** Offering DDoS protection, DNS hijacking detection, intrusion detection, vulnerability scanning, web trojan detection, and login protection to safeguard your servers.
- **Free Access to Tencent Cloud Observability Platform:** Supporting various real-time alerts.
- **Recycle Bin Protection Mechanism:** Supports monthly subscription cloud services entering the recycle bin for a period of time after expiration, preventing significant data loss and other impacts caused by immediate termination.
- **Custom Access Control:** Configure access policies for hosts and networks using security groups and network ACLs, allowing for flexible and customizable firewall settings for different instances.

Tencent Cloud security highlights:

- **Comprehensive Security Protection**  
Integrated security services for cloud servers, including security checks (vulnerability scanning, webshell detection, backdoor detection, port security checks, etc.) and security defenses (DDoS protection, intrusion detection, access control to ensure data security and user privacy).
- **Real-Time Alerts and Periodic Analysis**  
24/7 security service to promptly detect vulnerabilities and provide real-time, free notifications.
- **Free, Convenient, and Secure**  
No need to purchase expensive security equipment for your cloud services. Enjoy free cloud security services with the purchase of cloud services. One-click activation, zero deployment, and easy to use.
- **Professional Team, Reliable Protection**  
Cloud Security is crafted by Tencent's experienced security team, providing professional security services for cloud users, ensuring your trust and confidence.

## Easy-to-Use

A wide range of officially verified application software and Ops tools make operation and maintenance more convenient, so that you no longer have to worry about choosing the right tools.

- Tencent Cloud CVM provides a web-based user interface, known as the console, which allows you to start, adjust configurations, and reinstall the system for cloud server instances, just like working with physical machines. If you have already registered a Tencent Cloud account, you can directly [log in to the CVM console](#) to manage your CVM.

- Tencent Cloud CVM provides an API system, allowing you to conveniently integrate cloud servers with your internal monitoring and operational systems, achieving a fully automated business operation system tailored to your needs. These requests are HTTP or HTTPS requests. For more information on CVM API operations, please refer to the [API Documentation](#).
- If you prefer managing your resources, applications, and data using APIs, you can use SDKs (supporting PHP/Python/Java/.NET/Node.js/Go) for programming or call CVM APIs using Tencent Cloud Command Line Tools. For more information, see [Using SDKs](#) and [Using Command Line Tools](#).

## Cost-effective

Tencent Cloud lets you save on IT investment, and simplify the operating process.

- CVM instances and network deployments support both monthly subscription and pay-as-you-go billing options, catering to various application scenario requirements.
- You only pay for what you need. No up-front investment, no infrastructure, no preparation. Tencent Cloud is the perfect cloud platform for you.

# Regions and AZs

Last updated: 2025-08-18 09:44:50

## Regions

### Feature Overview

A region is the physical location of an IDC. Data centers in different Tencent Cloud regions are isolated to ensure maximal stability and fault tolerance. We recommend that you choose the region closest to your end users to minimize access latency and improve access speed.

You can view the following table or use the [DescribeRegions](#) API to get a complete region list.

### Related features

- The networks of different regions are fully isolated from each other, and Tencent Cloud services in different regions **cannot communicate with each other over the private network by default**.
- Tencent Cloud services in different regions can communicate with each other via [public IPs](#) over the Internet. Services in different Virtual Private Clouds can communicate through [Cloud Connect Network](#), which offers faster and more stable communication.
- [Cloud Load Balancer](#) currently supports traffic forwarding within the same region by default, binding to Cloud Virtual Machines in the local region. If the [cross-region binding](#) feature is enabled, Cloud Load Balancer can bind Cloud Virtual Machines across regions.

#### Note

**Shenzhen/Shanghai Finance Zones** are tailor-made AZs for compliance with regulatory requirements in the finance industry and feature high security and isolation. Currently, they are available for Cloud Virtual Machine, finance database, Redis storage, and face recognition services. They can be activated by verified financial customers by [contacting customer service](#). For more information, see [Introduction to Finance Zones](#).

## Availability Zones

### Feature Overview

An availability zone (AZ) is a physical IDC of Tencent Cloud with independent power supply and network in the same region. It can ensure business stability, as failures (except for major disasters or power failures) in one AZ are isolated without affecting other AZs in the same region. By starting an instance in an independent AZ, users can protect their applications from being affected by a single point of failure.

You can use the [DescribeZones](#) API to get a complete availability zone list.

## Related features

Tencent Cloud services in the same VPC are interconnected via the private network, which means they can communicate using [private IPs](#), even if they are in different AZs of the same region.

### Note

Private network interconnection refers to the interconnection of resources under the same account. Resources under different accounts are completely isolated on the private network.

## China

Regions	Availability Zones
South China (Guangzhou) ap-guangzhou	Guangzhou Zone 1 (Sold out) ap-guangzhou-1
	Guangzhou Zone 3 ap-guangzhou-3
	Guangzhou Zone 4 ap-guangzhou-4
	Guangzhou Zone 6 ap-guangzhou-6
	Guangzhou Zone 7 ap-guangzhou-7
South China (Shenzhen Finance) ap-shenzhen-fsi	Shenzhen Finance Zone 1 (available exclusively for financial institutions and enterprises through <a href="#">online consultation application</a> ) ap-shenzhen-fsi-1
	Shenzhen Finance Zone 2 (exclusive to financial institutions and enterprises, accessible through <a href="#">online consultation application</a> ) ap-shenzhen-fsi-2
	Shenzhen Finance Zone 3 (exclusive to financial institutions and enterprises, accessible through <a href="#">online consultation application</a> ) ap-shenzhen-fsi-3
East China (Shanghai)	Shanghai Zone 1 (Sold out)

ap-shanghai	ap-shanghai-1
	Shanghai Zone 2 ap-shanghai-2
	Shanghai Zone 3 ap-shanghai-3
	Shanghai Zone 4 ap-shanghai-4
	Shanghai Zone 5 ap-shanghai-5
	Shanghai Zone 8 ap-shanghai-8
East China (Shanghai Finance) ap-shanghai-fsi	Shanghai Finance Zone 1 (available exclusively for financial institutions and enterprises through <a href="#">online consultation application</a> ) ap-shanghai-fsi-1
	Shanghai Finance Zone 2 (exclusive access for financial institutions and enterprises through <a href="#">online consultation application</a> ) ap-shanghai-fsi-2
	Shanghai Finance Zone 3 (available exclusively for financial institutions and enterprises through <a href="#">online consultation application</a> ) ap-shanghai-fsi-3
East China (Nanjing) ap-nanjing	Nanjing Zone 1 ap-nanjing-1
	Nanjing Zone 2 ap-nanjing-2
	Nanjing Zone 3 ap-nanjing-3
North China (Beijing) ap-beijing	Beijing Zone 1 (sold out) ap-beijing-1
	Beijing Zone 2 ap-beijing-2
	Beijing Zone 3

	ap-beijing-3
	Beijing Zone 4 ap-beijing-4
	Beijing Zone 5 ap-beijing-5
	Beijing Zone 6 ap-beijing-6
	Beijing Zone 7 ap-beijing-7
Southwest China (Chengdu) ap-chengdu	Chengdu Zone 1 ap-chengdu-1
	Chengdu Zone 2 ap-chengdu-2
Southwest China (Chongqing) ap-chongqing	Chongqing Zone 1 ap-chongqing-1
Hong Kong, Macao and Taiwan, China (Hong Kong) ap-hongkong	Hong Kong Zone 1 (Nodes in Hong Kong, China can cover services in Hong Kong/Macao/Taiwan regions) (Sold out) ap-hongkong-1
	Hong Kong Zone 2 (Hong Kong nodes cover services in the China regions of Hong Kong, Macao, and Taiwan) ap-hongkong-2
	Hong Kong Zone 3 (Hong Kong nodes cover services in the China regions of Hong Kong, Macao, and Taiwan) ap-hongkong-3

**Note**

The product is in beta test for Jinan, Hangzhou, Fuzhou, Wuhan, Changsha, and Shijiazhuang regions. To try it out, contact the sales rep for application.

## Other Countries and Regions

Regions	Availability Zones
---------	--------------------

Southeast Asia (Singapore) ap-singapore	Singapore Zone 1 (Nodes in Singapore can cover Southeast Asia) ap-singapore-1
	Singapore Zone 2 (Nodes in Singapore can cover Southeast Asia) ap-singapore-2
	Singapore Zone 3 (Nodes in Singapore can cover Southeast Asia) ap-singapore-3
	Singapore Zone 4 (Nodes in Singapore can cover Southeast Asia) ap-singapore-4
Southeast Asia (Jakarta) ap-jakarta	Jakarta Zone 1 (Nodes in Jakarta can cover Southeast Asia) ap-jakarta-1
	Jakarta Zone 2 (Jakarta nodes cover services in Southeast Asia) ap-jakarta-2
Northeast Asia (Seoul) ap-seoul	Seoul Zone 1 (Seoul nodes cover services in Northeast Asia) ap-seoul-1
	Seoul Zone 2 (Seoul nodes cover services in Northeast Asia) ap-seoul-2
Northeast Asia (Tokyo) ap-tokyo	Tokyo Zone 1 (Tokyo nodes cover services in Northeast Asia) ap-tokyo-1
	Tokyo Zone 2 (Tokyo node AZs cover services in Northeast Asia) ap-tokyo-2
South Asia (Mumbai) ap-mumbai	Mumbai Zone 1 (Mumbai nodes cover services in South Asia) ap-mumbai-1
	Mumbai Zone 2 (Mumbai nodes cover services in South Asia) ap-mumbai-2

Southeast Asia (Bangkok) ap-bangkok	Bangkok Zone 1 (Bangkok nodes cover services in Southeast Asia) ap-bangkok-1
	Bangkok Zone 2 (Bangkok nodes cover services in Southeast Asia) ap-bangkok-2
North America (Toronto) na-toronto	Toronto Zone 1 (Toronto nodes cover services in North America) na-toronto-1
South America (Sao Paulo) sa-saopaulo	São Paulo Zone 1 (São Paulo nodes cover services in South America) sa-saopaulo-1
US West (Silicon Valley) na-siliconvalley	Silicon Valley Zone 1 (Silicon Valley nodes cover services in West US) na-siliconvalley-1
	Silicon Valley Zone 2 (Silicon Valley nodes cover services in West US) na-siliconvalley-2
US East (Virginia) na-ashburn	Virginia Zone 1 (Virginia nodes cover services in East US) na-ashburn-1
	Virginia Zone 2 (Virginia nodes cover services in East US) na-ashburn-2
Europe (Frankfurt) eu-frankfurt	Frankfurt Zone 1 (Frankfurt nodes cover services in Europe) eu-frankfurt-1
	Frankfurt Zone 2 (Frankfurt nodes cover services in Europe) eu-frankfurt-2

## Selection of Regions and AZs

When selecting a region and AZ, take the following into consideration:

- The region where the CVM is located, your location, and the location of your users.  
We recommend that you choose the region closest to your end users when purchasing CVM instances to minimize access latency and improve access speed.
- The relationship between CVMs and other Tencent Cloud products.  
When you select other Tencent Cloud services, we recommend that you try to locate them all in the same region and AZ to allow them to communicate with each other through the private network, reducing access latency and increasing access speed.

- High reliability and disaster recovery.  
Even if you have just one VPC, we still recommend that you deploy your businesses in different availability zones to prevent a single point of failure and enable cross-AZ disaster recovery.
- There may be network latency among different availability zones. We recommend that you assess your business requirements and find the optimal balance between high availability and low latency.
- If you need access to servers in other countries or regions, we recommend that you select a Cloud Virtual Machine in those other countries or regions. If you use a Cloud Virtual Machine in [China](#) to access [servers in other countries or regions](#), you may encounter higher network latency, which is not recommended.

## Resource Location Description

The following table describes which Tencent Cloud resources are global, which are regional, and which are specific to availability zones.

Resources	Resource ID Format (<Resource Abbreviation>-8-Digit String of Numbers and Letters)	Local Disk Types	Note
Account	No limit	Globally unique	Users can use the same account to access Tencent Cloud resources around the world.
<a href="#">SSH key</a>	skey-xxxxxxxx	Globally available	Users can use the SSH key to bind a CVM in any regions under the account.
<a href="#">Cloud Virtual Machine instance</a>	ins-xxxxxxxx	Instances are AZ-specific.	Users can only create a CVM instance in a specific AZ.
<a href="#">Custom Images</a>	img-xxxxxxxx	Cross AZ in one region	Users can create custom images of instances and use them in different availability zones within the same region. To use a custom

			image in another region, please use the image replication feature to copy the custom image to the desired region.
Elastic IP	eip-xxxxxxxx	Cross AZ in one region	Elastic IP is created under a certain region, and can only be associated with instances in the same region.
Security Group	sg-xxxxxxxx	Cross AZ in one region	Security groups are created within a specific region and can only be associated with instances in the same region. Tencent Cloud automatically creates three default security groups for users.
Cloud Block Storage	disk-xxxxxxxx	Instances are AZ-specific.	Users can only create a Cloud Block Storage disk in a specific AZ and mount it to instances in the same availability zone.
Snapshot	snap-xxxxxxxx	Cross AZ in one region	After creating a snapshot for a specific cloud disk, users can use this snapshot in this region for other operations (such as creating cloud disks).
Cloud Load Balancer	clb-xxxxxxxx	Cross AZ in one region	Cloud Load Balancer can be bound with CVMs in different availability zones of a single region for traffic forwarding.
Virtual Private Cloud	vpc-xxxxxxxx	Cross AZ in one region	A VPC in one region can have resources created in different availability zones.
Subnet	subnet-xxxxxxxx	Instances are AZ-specific.	Users cannot create subnets across availability zones.
Route Table	rtb-xxxxxxxx	Cross AZ in one region	When creating a route table, users need to specify a VPC, so the route table and VPC

share the same location attribute.

## Related Actions

### Migrate an instance to another availability zone

Once launched, an instance cannot be migrated to another availability zone. However, you can create a custom image of the CVM instance and use the image to launch or update an instance in a different availability zone.

1. Create a custom image for the current instance. For more information, see [Create Custom Images](#).
2. If the network environment of the current instance is Virtual Private Cloud and you want to retain the current private IP address after migration, you can first delete the subnet in the current availability zone and then create a subnet with the same IP address range in the new availability zone. Note that only subnets without available instances can be deleted. Therefore, you should move all instances in the current subnet to the new subnet.
3. Use the newly created custom image to create a new instance in a new availability zone. You can choose the same instance type and configuration as the original instance or select a new instance type and configuration. For more information, see [Create Instance](#).
4. If an elastic IP is associated with the original instance, dissociate it from the old instance and associate it with the new instance. For more information, see [Elastic IPs](#).
5. (Optional) If the original instance is [pay-as-you-go](#), you can choose to terminate it. For more information, see [Terminating Instances](#). If the original instance is [monthly subscription](#), you can choose to wait for its expiration and recycling.

### Copy images to other regions

Actions such as launching instances and viewing instances are region-specific. If the image required to launch an instance is not available in the current region, you need to copy the image to the desired region. For more information, see [Copying Images](#).

# Instance

## Instance Overview

Last updated: 2024-09-22 16:53:26

### Overview of a CVM Instance

**Instance** can be understood as a Cloud Virtual Machine (CVM), which includes basic computing components such as CPU, memory, operating system, network, and disk.

CVM instances provide elastic computing services in the cloud in a secure and reliable way to meet computing requirements. As business demands change, computing resources can be scaled in real time to lower your software and hardware costs and simplify IT OPS work.

Different instance types offer varying computing and storage capabilities, suitable for different application scenarios. Users can choose the computing power, storage space, and network access method based on the required service scale. For more information on instance types and their applicable scenarios, see [Instance Types](#). Once an instance is launched, users can use it like a traditional computer, having full control over the started instance.

### Instance Image

**Image** is a template for CVM software configurations, including the operating system and pre-installed programs. Tencent Cloud image provides all the information needed to launch a CVM instance. Users are required to launch instances through an image. An image can launch multiple instances for users to use repeatedly. In simple terms, an image is the "installation disk" for a CVM.

Tencent Cloud provides the following types of images:

- **Public image:** available to all users and suitable for major operating systems.
- **Marketplace images:** Available to all users, these images integrate specific applications in addition to the operating system.
- **Custom image:** only available to the creator and the users with whom the image is shared. A custom image is created from running instances or imported from external sources.
- **Shared image:** shared by other users. They can only be used to create instances.

For more information on images, see [Image Overview](#) and [Image Types](#).

### Instance Storage

The storage of an instance is similar to that of a regular cloud server, divided into **system disk** and **data disk**:

- **System Disk:** Similar to the C drive in a Windows system. The system disk contains a complete copy of the image used to launch the instance, as well as the instance's runtime

environment. When starting, you must choose a system disk size larger than the image size.

- **Data Disk:** Similar to other drives like D and E in the Windows system. Data disks store user data and support flexible resizing, mounting, and unmounting.

Both system and data disks can use different storage types provided by Tencent Cloud. For more information, see [Storage Overview](#).

## Instance Security

Tencent Cloud provides the following instance security protection methods:

- **Policy control:** When a group of cloud resources needs to be controlled by multiple different accounts, users can use policy control to manage access permissions to these resources.
- **Security group:** Control access to instances by allowing trusted addresses to access them using security groups.
- **Login Control:** It is recommended to use [SSH key](#) for logging into Linux instances, while instances using [password login](#) should have their passwords changed periodically.

# Instance Types

Last updated: 2023-09-07 15:42:55

When creating a Tencent Cloud CVM, the instance type specified by the user determines the hardware configuration of the instance. Each instance type offers varying compute, memory, and storage capabilities. Users can choose an appropriate instance type based on the scale of the application they need to deploy and run. These instance families consist of different combinations of CPU, memory, storage, heterogeneous hardware, and network bandwidth, allowing you to flexibly select the right resources for your application.

For stable businesses, it is recommended to choose the monthly subscription billing mode, with longer purchase durations being more cost-effective. For sudden business peaks, you can opt for the pay-as-you-go billing mode, allowing you to create or terminate instances at any time and pay based on the actual usage of the instances. Billing is accurate to the second, minimizing costs to the greatest extent.

## Instance Type

Tencent Cloud instance families are categorized as follows:

Local Disk Types	Subtype	Description
Standard Instance Family	<ul style="list-style-type: none"> <li>Standard S6</li> <li>Standard SA3</li> <li>Standard Type SR1</li> <li>Standard S5</li> <li>Standard Storage Optimized S5se</li> <li>Standard SA2</li> <li>Standard S4</li> <li>Standard Network Optimized Type SN3ne</li> <li>Standard S3</li> <li>Standard SA1</li> <li>Standard Network Optimized S2ne</li> <li>Standard S2</li> </ul>	Standard instances provide a balance of compute, memory, and network resources to accommodate most applications.

	<ul style="list-style-type: none"> <li>• <a href="#">Standard S1</a></li> </ul>	
<a href="#">Memory-Optimized Instance Family</a>	<ul style="list-style-type: none"> <li>• <a href="#">Memory-optimized MA3</a></li> <li>• <a href="#">Memory-optimized M6</a></li> <li>• <a href="#">Security-enhanced Memory Edition M6ce</a></li> <li>• <a href="#">Memory-optimized M6p</a></li> <li>• <a href="#">Memory-optimized M5</a></li> <li>• <a href="#">Memory-optimized MA2</a></li> <li>• <a href="#">Memory-optimized M4</a></li> <li>• <a href="#">Memory-optimized M3</a></li> <li>• <a href="#">Memory-optimized M2</a></li> <li>• <a href="#">Memory-optimized M1</a></li> </ul>	<p>Memory optimized instances feature large memory and are suitable for applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching.</p>
<a href="#">High I/O Instance Family</a>	<ul style="list-style-type: none"> <li>• <a href="#">High I/O: IT5</a></li> <li>• <a href="#">High I/O: IT3</a></li> </ul>	<p>This instance family features high random IOPS, high throughput and low latency, making it suitable for I/O-intensive applications that require high disk read/write performance and low latency, such as high-performance databases.</p>
<a href="#">Big Data Instance Family</a>	<ul style="list-style-type: none"> <li>• <a href="#">Big Data Type D3</a></li> <li>• <a href="#">Big Data Type D2</a></li> </ul>	<p>This instance family is equipped with massive storage resources and features high throughput, making it suitable for throughput-intensive applications such as Hadoop distributed computing, massive log processing, distributed file systems, and large data warehouses.</p>
<a href="#">Compute-optimized Instance Family</a>	<ul style="list-style-type: none"> <li>• <a href="#">Compute C6</a></li> <li>• <a href="#">Compute C5</a></li> <li>• <a href="#">Compute C4</a></li> <li>• <a href="#">Compute CN3</a></li> <li>• <a href="#">Compute C3</a></li> </ul>	<p>Featuring up to 3.8GHz turbo frequency, these instances provide the highest single-core computing performance. They are suitable for compute-intensive applications such as batch processing, high-performance computing, and large-scale game servers.</p>

	<ul style="list-style-type: none"> <li>• <a href="#">Compute C2</a></li> </ul>	
<a href="#">Heterogeneous Computing Instance Family</a>	–	This family is equipped with heterogeneous hardware such as GPU and FPGA to deliver real-time, fast parallel computing and floating-point computing capabilities. It is suitable for high-performance applications such as deep learning, scientific computing, video encoding/decoding, and graphics workstations.
<a href="#">Batch Instance Family</a>	<ul style="list-style-type: none"> <li>• <a href="#">Batch Computing Type BC1</a></li> <li>• <a href="#">Batch Computing BS1</a></li> </ul>	With the lowest per core-hour cost, this family is suitable for compute-intensive applications that frequently deploy ultra-large amounts of computing nodes for short periods of time. Examples include rendering, genetic analysis, and solid-state pharmaceutical research.
<a href="#">Bare Metal Cloud Server</a>	–	Bare Metal Cloud Servers are high-performance bare metal cloud servers developed based on Tencent Cloud's latest virtualization technology. They combine the flexibility and elasticity of cloud servers with the high stability and robust computing power of physical machines. These servers seamlessly integrate with the full range of Tencent Cloud products.
<a href="#">High-performance Computing Cluster</a>	–	High-performance computing clusters use bare metal cloud servers as nodes, interconnected via RDMA, providing high-bandwidth and ultra-low latency network services. These clusters can meet the parallel computing requirements of large-scale high-performance computing, artificial intelligence, and big data recommendation applications.

**Note:**

Some instances have been discontinued. For more information, see [Instance Specifications \(Discontinued\)](#). Glossary

- **vCPU:** Each vCPU corresponds to a hyperthread of an Intel Xeon processor core or a hyperthread of an AMD EPYC processor core.
- **Send/Receive Packets:** the maximum number of data packets that the instance can process per second for both sending and receiving, without differentiating private or public network traffic.

- **Private Network Bandwidth Capacity:** The maximum amount of data (bit) that can be transmitted per second within the instance's private network.
- **Multi-queue Technology:** multiple send/receive packet queues are supported on a virtual ENI, with the sending and receiving of the data packets of each queue capable of being simultaneously processed by different CPUs

**Note:**

- If you purchase a larger instance specification, it is recommended to use multi-queue technology to achieve higher performance. If you use a single queue, some network metrics may not reach their maximum values.
- To adjust the ENI queues, please refer to the [Steps](#) in the **ENI Multi-Queue Configuration Error** documentation.

- **Number of Queues:** the quantity of send/receive packet queues supported by each virtual ENI (N number of queues indicates that N number of receiving queues and N number of sending queues can be simultaneously supported)

## Instance Limits

- There is a limit to the total number of instances that can be launched within a region. For more information on these limitations, please refer to [CVM Instance Purchase Limits](#).
- **Instance system disk and data disk mounting limitations:** To ensure optimal disk IO performance, Tencent Cloud imposes restrictions on the size and type of data disks purchased with instances. Please refer to the disk configuration module for the corresponding instance family. If necessary, you can bypass the instance disk limitations by purchasing separate cloud disks.
- Note that the private network bandwidth capacity of an instance specification is the maximum private network bandwidth limit of the corresponding instance. If the CVM private network traffic exceeds this limit, random packet loss may happen within the private network for your instances.
- The availability of instance specifications may vary from region to region. Some configurations may be sold out. Please see the purchase page for the latest information.
- The network packet sending and receiving data in this document are based on pure network forwarding test scenarios. For testing methods, please refer to [Network Performance Testing](#). For performance data in actual business scenarios, it is essential to conduct tests before estimating capacity.

Based on your business scenarios, the instance families can be divided into the following categories:

# Standard Instance Family

Standard instances provide a balance of compute, memory, and network resources to accommodate most applications.

## Standard S6

As the latest generation of standard instances, S6 instances are powered by an all-new optimized virtualization platform, and provide balanced and stable compute, memory, and network resources. These instances are ideal for many applications.

S6 instances come with the new Intel® Xeon® processors with the latest DDR4 memory. They are network optimized by default, offering a private network throughput of up to 19 million pps and a private network bandwidth of up to 100 Gbps.

## Features

- A new-generation Star Lake two-way server developed by Tencent Cloud, with the 3rd generation Intel® Xeon® scalable processor
- Intel® Xeon® Ice Lake processor, with a CPU clock rate of 2.7 GHz and a turbo frequency of 3.3 GHz
- Support the processor-to-memory ratios of 1:2 and 1:4
- Up to 100 Gbps private network bandwidth for ultra high packet throughput to meet extremely high private network transfer requirements
- The network performance of an instance corresponds to its specifications; the higher the specifications, the stronger the network forwarding performance, and the higher the internal network bandwidth cap.
- Supports hyper-threading

## Use Cases

Standard S6 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- S6 instances are pay-as-you-go.
- Supports only VPC.

- We recommend you use the instance together with the TencentOS Server operating system to enjoy the best application performance.
- S6 instances are available for purchase with configurations as shown in the table below. Ensure that the selected S6 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.
- Instances can achieve up to 100 Gbps network bandwidth, depending on the support of the instance's operating system kernel version and runtime environment. When the PPS exceeds 10 million and the bandwidth is greater than 50 Gbps, the kernel protocol stack may cause significant network performance degradation. In this case, the netperf test bandwidth value may not meet expectations. You can use the DPDK method to mask the differences in the cloud server kernel protocol stack and obtain the true network performance of the instance. For testing methods, please refer to [High-Throughput Network Performance Testing](#).

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S6.MEDIUM2	2	2	300,000	250,000	2	2	2.7 GHz	-
S6.MEDIUM4	2	4	300,000	250,000	2	2	2.7 GHz	-
S6.MEDIUM8	2	8	300,000	250,000	2	2	2.7 GHz	-
S6.LARGE8	4	8	0.6 million	0.5 million	4	4	2.7 GHz	-
S6.LARGE16	4	16	0.6 million	0.5 million	4	4	2.7 GHz	-

S6.2XLA RGE16	8	16	1.2 million	1 million	8	7	2.7 GHz	–
S6.2XLA RGE32	8	32	1.2 million	1 million	8	7	2.7 GHz	–
S6.4XLA RGE32	16	32	2.5 million	2,000,000	16	13	2.7 GHz	–
S6.4XLA RGE64	16	64	2.5 million	2,000,000	16	13	2.7 GHz	–
S6.8XLA RGE64	32	64	5,000,000	4,000,000	32	26	2.7 GHz	–
S6.8XLA RGE128	32	128	5,000,000	4,000,000	32	26	2.7 GHz	–
S6.12XL ARGE96	48	96	7,000,000	6 million	48	39	2.7 GHz	–
S6.12XL ARGE192	48	192	7,000,000	6 million	48	39	2.7 GHz	–
S6.16XL ARGE216	64	216	10 million	8 million	48	50	2.7 GHz	–
S6.32XL ARGE432	128	432	15 million	16,000,000	48	100	2.7 GHz	–

## Standard SA3

Standard SA3 instances utilize the latest generation of AMD EPYC™ processors, offering a wide range of specifications and exceptional cost–performance ratio. These instances provide a balanced combination of compute, storage, and network resources, making them the optimal choice for numerous applications.

Standard SA3 instances utilize the latest AMD EPYC™ Milan processors, feature the latest DDR4 memory, and come with default network optimization. They offer a maximum intranet transmission capacity of 19 million pps and support up to 100 Gbps intranet bandwidth.

## Features

- Provides high performance with high reliability, security, and stability based on the Star Lake servers developed by Tencent Cloud
- 2.55GHz AMD EPYC™ Milan processor, with a turbo frequency of 3.5GHz, equipped with the latest generation of eight-channel DDR4, ensuring stable memory computing performance.
- Higher specification, SA3.58XLARGE940, offering 232 vCPUs and 940 GB memory
- Support the processor-to-memory ratios of 1:2 and 1:4
- Supports up to 100 Gbps of private network bandwidth for ultra-high packet throughput to meet extremely high private network transmission requirements
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard SA3 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Search and other computing clusters
- Video encoding/decoding, video rendering, and other applications sensitive to single-core performance

## Requirements

- SA3 instances are pay-as-you-go.
- Supports only VPC.
- We recommend you use the instance together with the TencentOS Server operating system to enjoy the best application performance.
- SA3 instances are available for purchase with various configurations. Please refer to the instance specifications below. Ensure that the SA3 instance size you choose meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.
- Instances can achieve up to 100 Gbps network bandwidth, depending on the support of the instance's operating system kernel version and runtime environment. When the PPS

exceeds 10 million and the bandwidth is greater than 50 Gbps, the kernel protocol stack may cause significant network performance degradation. In this case, the netperf test bandwidth value may not meet expectations. You can use the DPDK method to mask the differences in the cloud server kernel protocol stack and obtain the true network performance of the instance. For testing methods, please refer to [High-Throughput Network Performance Testing](#).

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
SA3.MEDIUM2	2	2	300,000	250,000	2	1.5	2.5 5G Hz	-
SA3.MEDIUM4	2	4	300,000	250,000	2	1.5	2.5 5G Hz	-
SA3.MEDIUM8	2	8	300,000	250,000	2	1.5	2.5 5G Hz	-
SA3.LARGE8	4	8	0.5 million	250,000	4	2	2.5 5G Hz	-
SA3.LARGE16	4	16	0.5 million	250,000	4	2	2.5 5G Hz	-
SA3.2XLARGE16	8	16	8 million	0.5 million	8	4	2.5 5G Hz	-
SA3.2XLARGE32	8	32	8 million	0.5 million	8	4	2.5 5G Hz	-
SA3.4XLARGE32	16	32	1.5 million	1.1 million	16	7	2.5 5G Hz	-

SA3.4XL ARGE64	1 6	64	1.5 million	1.1 million	16	7	2.5 5G Hz	-
SA3.8XL ARGE64	3 2	64	2.5 million	2.2 million	32	14	2.5 5G Hz	-
SA3.8XL ARGE128	3 2	128	2.5 million	2.2 million	32	14	2.5 5G Hz	-
SA3.12X LARGE9 6	4 8	96	4,000,00 0	3,300, 000	48	21	2.5 5G Hz	-
SA3.12X LARGE1 92	4 8	192	4,000,00 0	3,300, 000	48	21	2.5 5G Hz	-
SA3.16X LARGE12 8	6 4	128	5.2 million	4.4 million	48	28	2.5 5G Hz	-
SA3.16X LARGE2 56	6 4	256	5.2 million	4.4 million	48	28	2.5 5G Hz	-
SA3.20X LARGE1 60	8 0	160	6.5 million	5,500, 000	48	35	2.5 5G Hz	-
SA3.20X LARGE3 20	8 0	320	6.5 million	5,500, 000	48	35	2.5 5G Hz	-
SA3.24X LARGE1 92	9 6	192	7.8 million	6,600, 000	48	42	2.5 5G Hz	-
SA3.24X LARGE3 84	9 6	384	7.8 million	6,600, 000	48	42	2.5 5G Hz	-
SA3.29X LARGE21 6	1 1 6	216	9.5 million	8 million	48	50	2.5 5G Hz	-

SA3.29X LARGE4 70	1 1 6	470	9.5 million	8 million	48	50	2.5 5G Hz	–
SA3.40X LARGE3 20	1 6 0	320	13,000,0 00	11,000 ,000	48	69	2.5 5G Hz	–
SA3.40X LARGE6 40	1 6 0	640	13,000,0 00	11,000 ,000	48	69	2.5 5G Hz	–
SA3.58X LARGE4 32	2 3 2	432	19 million	16,00 0,000	48	100	2.5 5G Hz	–
SA3.58X LARGE9 40	2 3 2	940	19 million	16,00 0,000	48	100	2.5 5G Hz	–

## Standard SR1

Standard SR1 instances are the latest generation of ARM-based standard instances, built on a newly optimized virtualization platform. They provide balanced, stable compute, memory, and network resources, making them the optimal choice for a wide range of applications.

### Note:

Currently, this model is available for sale in Guangzhou Zone 6. You can visit the [purchase page](#) to buy it. Other regions will be gradually opened, so please stay tuned for product updates.

## Features

- Featuring Ampere® Altra® processors, with a base frequency of 2.8GHz and a consistent maximum frequency of 3.0GHz across all cores.
- Support the processor-to-memory ratios of 1:2 and 1:4
- Data disks can support Premium Cloud Disk, SSD Cloud Disk, and Enhanced SSD Cloud Disk.
- Support up to 25 Gbps of private network bandwidth
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard SR1 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- ARM Simulation Test
- Embedded development and testing
- CPU-based Machine Learning Inference

## Requirements

- SR1 instances are pay-as-you-go.
- The network of SR1 instances must be based on a VPC.
- SR1 instances are available for purchase with configurations as detailed in the table below. Ensure that the SR1 instance size you choose meets the minimum CPU and memory requirements for your operating system and applications. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or opt for other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
SR1.SMALL1	1	1	250,000	1	1.5	2.8 GHz	-
SR1.SMALL2	1	2	250,000	1	1.5	2.8 GHz	-
SR1.MEDIUM4	2	4	250,000	2	1.5	2.8 GHz	-
SR1.MEDIUM8	2	8	250,000	2	1.5	2.8 GHz	-
SR1.LARGE8	4	8	250,000	2	1.5	2.8 GHz	-

SR1.LARGE16	4	16	250,000	2	1.5	2.8 GHz	–
SR1.2XLRGE16	8	16	0.5 million	2	3	2.8 GHz	–
SR1.2XLRGE32	8	32	0.5 million	2	3	2.8 GHz	–
SR1.4XLRGE32	16	32	1.1 million	4	6	2.8 GHz	–
SR1.4XLRGE64	16	64	1.1 million	4	6	2.8 GHz	–
SR1.8XLRGE64	32	64	2.2 million	8	12	2.8 GHz	–
SR1.8XLRGE128	32	128	2.2 million	8	12	2.8 GHz	–
SR1.16XLRGE128	64	128	4.5 million	16	16	2.8 GHz	–

## Standard S5

Standard S5 instances are the next-generation standard instances, based on a newly optimized virtualization platform. They provide balanced, stable compute, memory, and network resources, making them the optimal choice for many applications.

### Features

- Intel® Xeon® Cascade Lake or Intel® Xeon® Cooper Lake processor, with a CPU clock rate of 2.5 GHz and turbo frequency of 3.1 GHz
- New Intel Advanced Vector Extension (AVX-512) instruction set
- 1:2 or 1:4 processor to memory ratio
- It can support up to 29 Gbps intranet bandwidth and ultra-high network packet transmission and reception capabilities, meeting the extremely high intranet transmission demands.

- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard S5 instances are applicable to the following scenarios:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- S5 instances are pay-as-you-go.
- The network of S5 instances must be based on a VPC.
- S5 instances support purchase configurations; please refer to the instance specifications below. Ensure that the selected S5 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S5.SMALL1	1	1	250,000	250,000	1	1.5	2.5 GHz	-
S5.SMALL2	1	2	250,000	250,000	1	1.5	2.5 GHz	-

S5.SMALL4	1	4	250,000	250,000	1	1.5	2.5 GHz	-
S5.MEDIUM2	2	2	250,000	250,000	1	1.5	2.5 GHz	-
S5.MEDIUM4	2	4	300,000	250,000	2	1.5	2.5 GHz	-
S5.MEDIUM8	2	8	300,000	250,000	2	1.5	2.5 GHz	-
S5.LARGE4	4	4	0.5 million	250,000	2	1.5	2.5 GHz	-
S5.LARGE8	4	8	0.5 million	250,000	2	1.5	2.5 GHz	-
S5.LARGE16	4	16	0.5 million	250,000	2	1.5	2.5 GHz	-
S5.2XLARGE16	8	16	8 million	250,000	2	3.0	2.5 GHz	-

S5.2XL ARGE3 2	8	32	8 million	250,0 00	2	3.0	2. 5 G H Z	-
S5.4XL ARGE3 2	16	32	1.5 million	300,0 00	4	6.0	2. 5 G H Z	-
S5.4XL ARGE6 4	16	64	1.5 million	300,0 00	4	6.0	2. 5 G H Z	-
S5.6XL ARGE4 8	24	48	2,000,00 0	400,0 00	6	9.0	2. 5 G H Z	-
S5.6XL ARGE9 6	24	96	2,000,00 0	400,0 00	6	9.0	2. 5 G H Z	-
S5.8XL ARGE6 4	32	64	2.5 million	0.6 million	8	12	2. 5 G H Z	-
S5.8XL ARGE12 8	32	128	2.5 million	0.6 million	8	12	2. 5 G H Z	-
S5.12XL ARGE9 6	48	96	4,000,00 0	900,0 00	12	17.0	2. 5 G H Z	-

S5.12XL ARGE19 2	48	192	4,000,000	900,000	12	17.0	2.5 G H Z	–
S5.16XL ARGE2 56	64	256	5,000,000	1.2 million	16	23.0	2.5 G H Z	–
S5.21XL ARGE3 20	84	320	6 million	1,600,000	16	29.0	2.5 G H Z	–

## Standard Storage Optimized S5se

Standard Storage Optimized S5se instances are based on a newly optimized virtualization platform, comprehensively enhancing cloud disk performance. They utilize the latest Intel® Xeon® Cascade Lake processors and DDR4 memory, providing exceptional cloud disk performance.

### Features

- Supports Premium Cloud Disk and SSD as the system disk
- Supports ulTra SSD as the data disk
- Instance cloud disk storage performance reaches up to 1.1 million IOPS and 4 GB/s read/write bandwidth.
- 2.5 GHz Intel® Xeon® Cascade Lake processor, with a turbo frequency 3.1 GHz and the latest generation 6-channel DDR4 for stable memory computing performance
- New Intel Advanced Vector Extension (AVX-512) instruction set
- 1:2 or 1:4 processor to memory ratio
- Support up to 25 Gbps of private network bandwidth
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

### Use Cases

Standard Storage Optimized S5se instances are applicable to the following scenarios:

- Enterprise applications of different types and sizes
- IO-intensive applications including large databases, NoSQL databases, audio/video processing services, and Elasticsearch cluster

## Requirements

- S5se instances are pay-as-you-go.
- Supports only VPC.
- See below for available S5se specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S5se.LARGE16	4	16	0.5 million	250,000	2	1.5	2.5 GHz	-
S5se.2XLARGE16	8	16	8 million	250,000	2	3.0	2.5 GHz	-
S5se.2XLARGE32	8	32	8 million	250,000	2	3.0	2.5 GHz	-
S5se.4XLARGE32	16	32	1.5 million	300,000	4	6.0	2.5 GHz	-
S5se.4XLARGE64	16	64	1.5 million	300,000	4	6.0	2.5 GHz	-
S5se.8XLARGE64	32	64	2.5 million	0.6 million	8	12.0	2.5	-

							G Hz	
S5se.8XL ARGE128	3 2	128	2.5 million	0.6 millio n	8	12.0	2. 5 G Hz	–
S5se.17X LARGE31 6	6 8	316	5,000,00 0	1.2 millio n	16	25.0	2. 5 G Hz	–

## Instance Types

Standard SA2 instances are the latest generation of instances featuring AMD EPYC™ processors, offering a wide range of specifications with exceptional cost–performance ratios, ensuring your workloads receive industry–leading value. These instances provide a balanced mix of compute, memory, and network resources, making them the optimal choice for many applications.

SA2 instances come with the new AMD EPYC™ ROME processors with the latest DDR4 memory. They are network optimized by default, offering a private network throughput up to 7,500,000 pps and support up to 25 Gbps of private network bandwidth.

## Features

- Provides high performance with high reliability, security, and stability based on the Star Lake servers developed by Tencent Cloud
- 2.6 GHz AMD EPYC™ ROME processor, with a turbo frequency 3.3 GHz and the latest generation 8–channel DDR4 for stable memory computing performance
- Higher specification, SA2.45XLARGE464, offering 180 vCPUs and 464 GB memory
- Supported CPU/RAM ratios: 1:2 and 1:4
- Support up to 25 Gbps of private network bandwidth
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard SA2 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Search and other computing clusters

- Video encoding/decoding, video rendering, and other applications sensitive to single-core performance

## Requirements

- SA2 instances are pay-as-you-go.
- Supports only VPC.
- SA2 instances support purchasing configurations; please refer to the instance specifications below. Ensure that the SA2 instance size you choose meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
SA2.SMALL1	1	1	250,000	1	1.5	2.6 GHz	–
SA2.SMALL2	1	2	250,000	1	1.5	2.6 GHz	–
SA2.SMALL4	1	4	250,000	1	1.5	2.6 GHz	–
SA2.MEDIUM2	2	2	250,000	1	1.5	2.6 GHz	–
SA2.MEDIUM4	2	4	300,000	2	1.5	2.6 GHz	–
SA2.MEDIUM8	2	8	300,000	2	1.5	2.6 GHz	–
SA2.LARGE4	4	4	0.5 million	2	1.5	2.6 GHz	–
SA2.LARGE8	4	8	0.5 million	2	1.5	2.6 GHz	–

SA2.LARGE 16	4	16	0.5 million	2	1.5	2.6 GHz	–
SA2.2XLAR GE16	8	16	0.7 million	2	1.5	2.6 GHz	–
SA2.2XLAR GE32	8	32	0.7 million	2	1.5	2.6 GHz	–
SA2.4XLAR GE32	16	32	1 million	4	3.0	2.6 GHz	–
SA2.4XLAR GE64	16	64	1 million	4	3.0	2.6 GHz	–
SA2.8XLAR GE64	32	64	1.4 million	8	5.0	2.6 GHz	–
SA2.12XLA RGE96	48	96	2,100,000	12	7.0	2.6 GHz	–
SA2.16XLA RGE128	64	128	2.8 million	16	9.0	2.6 GHz	–
SA2.20XLA RGE160	80	160	3.5 million	16	12.0	2.6 GHz	–
SA2.22XLA RGE224	90	224	3.75 million	16	13.0	2.6 GHz	–
SA2.24XLA RGE192	96	192	4.2 million	16	14.0	2.6 GHz	–
SA2.32XLA RGE256	128	256	5.6 million	32	18.0	2.6 GHz	–
SA2.40XLA RGE320	160	320	7.1 million	32	23.0	2.6 GHz	–
SA2.45XLA RGE464	180	464	7.5 million	32	25.0	2.6 GHz	–

## Standard S4

S4 instances provide a balance of compute, memory, and network resources, making them ideal for many applications.

S4 instances come with new Xeon® Skylake processors with the latest DDR4 memory. They are network optimized by default, offering a private network throughput up to 6,000,000 pps,

and support up to 25 Gbps of private network bandwidth.

## Features

- 2.4 GHz Intel® Xeon® Skylake 6148 processor, offering stable computing performance
- New Intel Advanced Vector Extension (AVX-512) instruction set
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- Higher specification, S4.18XLARGE228, offering 72 vCPUs and 228 GB memory
- 1:2 or 1:4 processor to memory ratio
- Support up to 25 Gbps of private network bandwidth
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard S4 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- S4 instances are pay-as-you-go.
- Supports only VPC.
- S4 instances are available for purchase and configuration; please refer to the instance specifications below. Ensure that the selected S4 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks

S4.SMALL 1	1	1	250,000	1	1.5	2.4 GH z	-
S4.SMALL 2	1	2	250,000	1	1.5	2.4 GH z	-
S4.SMALL 4	1	4	250,000	1	1.5	2.4 GH z	-
S4.MEDIU M2	2	2	250,000	1	1.5	2.4 GH z	-
S4.MEDIU M4	2	4	300,000	2	1.5	2.4 GH z	-
S4.MEDIU M8	2	8	300,000	2	1.5	2.4 GH z	-
S4.LARGE 8	4	8	0.5 million	2	1.5	2.4 GH z	-
S4.LARGE 16	4	16	0.5 million	2	1.5	2.4 GH z	-
S4.2XLAR GE16	8	16	8 million	2	3.0	2.4 GH z	-
S4.2XLAR GE32	8	32	8 million	2	3.0	2.4 GH z	-
S4.4XLAR GE32	16	32	1.5 million	4	6.0	2.4 GH z	-
S4.4XLAR GE64	16	64	1.5 million	4	6.0	2.4 GH z	-

S4.6XLAR GE48	24	48	2,000,000	6	8.0	2.4 GH z	–
S4.6XLAR GE96	24	96	2,000,000	6	8.0	2.4 GH z	–
S4.8XLAR GE64	32	64	2.5 million	8	11.0	2.4 GH z	–
S4.8XLAR GE128	32	128	2.5 million	8	11.0	2.4 GH z	–
S4.12XLAR GE96	48	96	4,000,000	12	16.0	2.4 GH z	–
S4.12XLAR GE192	48	192	4,000,000	12	16.0	2.4 GH z	–
S4.16XLAR GE256	64	256	5,000,000	16	22.0	2.4 GH z	–
S4.18XLAR GE288	72	288	6 million	16	24.0	2.4 GH z	–

## Standard Network-optimized SN3ne

Standard SN3ne instances are a relatively new generation of network optimized instances. This family provides a balance of computing, memory, and network resources, with outstanding network throughput, and it is a good choice for many applications.

Standard Network Optimized SN3ne instances are equipped with the latest Xeon® Skylake processors with the latest DDR4 memory, default network optimization, and private network throughput up to 6,000,000 pps, with performance is nearly 8 times greater than the Standard S3 family. It can support up to 25 Gbps of private network bandwidth, with performance 2.5 times greater than the Standard S3 family.

### Features

- 2.5 GHz Intel Xeon® Skylake 6133 processors with stable computing performance

- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- Higher specification, SN3ne.18XLARGE228, offering 72 vCPUs and 228 GB memory
- 1:2 or 1:4 processor to memory ratio
- Support up to 25 Gbps of private network bandwidth
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard Network Optimized SN3ne instances are applicable to the following scenarios:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- SN3ne instances are pay-as-you-go.
- The network of SN3ne instances must be based on a VPC.
- SN3ne instances are available for purchase with various configurations. Please refer to the instance specifications below. Ensure that the selected SN3ne instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
SN3ne.SMALL2	1	2	250,000	1	1.5	2.5 GHz	-

SN3ne.MEDIUM2	2	2	250,000	1	1.5	2.5 GHz	-
SN3ne.MEDIUM4	2	4	300,000	2	1.5	2.5 GHz	-
SN3ne.LARGE4	4	4	0.5 million	2	1.5	2.5 GHz	-
SN3ne.LARGE8	4	8	0.5 million	2	1.5	2.5 GHz	-
SN3ne.LARGE16	4	16	0.5 million	2	1.5	2.5 GHz	-
SN3ne.2XLARGE16	8	16	8 million	2	3.0	2.5 GHz	-
SN3ne.2XLARGE32	8	32	8 million	2	3.0	2.5 GHz	-
SN3ne.3XLARGE24	12	24	1 million	3	4.0	2.5 GHz	-
SN3ne.4XLARGE32	16	32	1.5 million	4	6.0	2.5 GHz	-
SN3ne.4XLARGE64	16	64	1.5 million	4	6.0	2.5 GHz	-
SN3ne.6XLARGE48	24	48	2,000,000	6	8.0	2.5 GHz	-
SN3ne.6XLARGE96	24	96	2,000,000	6	8.0	2.5 GHz	-

SN3ne.8XL ARGE64	32	64	2.5 million	8	11.0	2.5 GH z	–
SN3ne.8XL ARGE128	32	128	2.5 million	8	11.0	2.5 GH z	–
SN3ne.12XL ARGE96	48	96	4,000,000	12	16.0	2.5 GH z	–
SN3ne.12XL ARGE192	48	192	4,000,000	12	16.0	2.5 GH z	–
SN3ne.16XL ARGE128	64	128	5,000,000	16	22.0	2.5 GH z	–
SN3ne.16XL ARGE256	64	256	5,000,000	16	22.0	2.5 GH z	–
SN3ne.18XL ARGE288	72	288	6 million	16	24.0	2.5 GH z	–

## Standard S3

Standard S3 instances are a relatively new generation of standard instances. This family provides a balance of computing, memory, and network resources, and it is a good choice for many applications.

Standard S3 instances are equipped with the latest Xeon® Skylake processors with the latest DDR4 memory and support for up to 10 Gbps of private network bandwidth.

### Features

- 2.5 GHz Intel Xeon® Skylake 6133 processors with stable computing performance
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- Larger instance size, S3.20XLARGE320, offering 80 vCPUs and 320 GB of memory
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

Standard S3 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing

## Requirements

- S3 instances are pay-as-you-go.
- S3 instances can be launched in classic networks and VPCs.
- S3 instances are available for purchase and configuration. Please refer to the instance specifications below. Ensure that the selected S3 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	V C P U	Mem ory (GB)	Through put (pps) (received and sent packets)	Num ber of conn ectio ns	Num ber of Queu es	Private network bandwidth (Gbps) (outbound and inbound)	Cl oc k R at e	Re mar ks
S3.SMALL1	1	1	200,000	250,000	1	1.5	2.5 GHz	-
S3.SMALL2	1	2	200,000	250,000	1	1.5	2.5 GHz	-
S3.SMALL4	1	4	200,000	250,000	1	1.5	2.5 GHz	-

S3.MEDI UM2	2	2	200,000	250, 000	1	1.5	2. 5 G H Z	–
S3.MEDI UM8	2	8	250,000	250, 000	2	1.5	2. 5 G H Z	–
S3.LARG E4	4	4	0.45 million	350, 000	4	1.5	2. 5 G H Z	–
S3.LARG E8	4	8	0.45 million	350, 000	4	1.5	2. 5 G H Z	–
S3.LARG E16	4	16	0.45 million	350, 000	4	1.5	2. 5 G H Z	–
S3.2XLA RGE16	8	16	850,000	0.7 millio n	8	1.5	2. 5 G H Z	–
S3.2XLA RGE32	8	32	850,000	0.7 millio n	8	1.5	2. 5 G H Z	–
S3.3XLA RGE24	1 2	24	850,000	1 millio n	12	1.5	2. 5 G H Z	–

S3.3XLA RGE48	1 2	48	850,000	1 millio n	12	1.5	2. 5 G H Z	–
S3.4XLA RGE32	1 6	32	850,000	1,300 ,000	16	2.0	2. 5 G H Z	–
S3.4XLA RGE64	1 6	64	850,000	1,300 ,000	16	2.0	2. 5 G H Z	–
S3.6XLA RGE48	2 4	48	850,000	2,00 0,00 0	16	3.0	2. 5 G H Z	–
S3.6XLA RGE96	2 4	96	850,000	2,00 0,00 0	16	3.0	2. 5 G H Z	–
S3.8XLA RGE64	3 2	64	850,000	2,60 0,00 0	16	4.0	2. 5 G H Z	–
S3.8XLA RGE128	3 2	128	850,000	2,60 0,00 0	16	4.0	2. 5 G H Z	–
S3.12XL ARGE96	4 8	96	850,000	3.9 millio n	16	6.0	2. 5 G H Z	–

S3.12XL ARGE192	4 8	192	850,000	3.9 millio n	16	6.0	2. 5 G H Z	–
S3.16XL ARGE25 6	6 4	256	850,000	5.2 millio n	16	8.0	2. 5 G H Z	–
S3.20XL ARGE32 0	8 0	320	850,000	6.5 millio n	16	10.0	2. 5 G H Z	–

## Standard SA1

Standard SA1 instances are powered by AMD EPYC™ processors and offer a variety of specifications with exceptional cost-effectiveness, ensuring your workloads receive industry-leading price-performance. These instances provide a balanced mix of compute, memory, and network resources, making them the optimal choice for many applications.

### Features

- 2.0 GHz AMD EPYC™ 7551 processors with a Turbo Boost clock of 2.55 GHz, offering stable computing performance
- The latest generation of 8-channel DDR4 memory with a memory bandwidth of 2,666 MT/s
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

### Use Cases

Standard SA1 instances are applicable to the following use cases:

- Enterprise applications of different types and sizes
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing

### Requirements

- SA1 instances are pay-as-you-go.

- SA1 instances can be launched in classic networks and VPCs.
- SA1 instances are available for purchase with various configurations. Please refer to the instance specifications below. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	V C P U	Mem o r y (GB)	Throug h p u t (pps) (received and sent packets)	Num b e r o f c o n n e c t i o n s	Num b e r o f Q u e u e s	Private network bandwidth (Gbps) (outbound and inbound)	Cl o c k R a t e	Re m a r k s
SA1.SM ALL1	1	1	–	250,000	1	1.5	2.0 G Hz	–
SA1.SM ALL2	1	2	–	250,000	1	1.5	2.0 G Hz	–
SA1.SM ALL4	1	4	–	250,000	1	1.5	2.0 G Hz	–
SA1.ME DIUM2	2	2	–	250,000	1	1.5	2.0 G Hz	–
SA1.ME DIUM4	2	4	–	250,000	2	1.5	2.0 G Hz	–
SA1.ME DIUM8	2	8	–	250,000	2	1.5	2.0 G Hz	–
SA1.LA RGE4	4	4	–	350,000	4	1.5	2.0	–

							G Hz	
SA1.LARGE8	4	8	-	350,000	4	1.5	2.0 G Hz	-
SA1.LARGE16	4	16	-	350,000	4	1.5	2.0 G Hz	-
SA1.2XLARGE16	8	16	-	0.7 millio n	8	1.5	2.0 G Hz	-
SA1.2XLARGE32	8	32	-	0.7 millio n	8	1.5	2.0 G Hz	-
SA1.4XLARGE32	16	32	-	1,300,000	16	1.5	2.0 G Hz	-
SA1.4XLARGE64	16	64	-	1,300,000	16	1.5	2.0 G Hz	-

**Note:**

In the SA1 instance specification list, the "-" in the network packet transmission and reception column indicates that there is no fixed commitment to network packet performance and processor frequency for that specification. If you have a strong demand for consistent performance, it is recommended to choose a model with a performance commitment.

## Standard Network-optimized S2ne

The Standard Network Optimized S2ne instances are the optimal choice for applications with high network packet transmission capabilities, offering up to millions of network transmissions

per second. They are recommended for large-scale gaming servers, video streaming, live broadcasting, and other scenarios with high network PPS requirements.

**Note:**

This instance type is currently only available to whitelisted users. Contact your account manager if you wish to purchase these instance type.

## Features

- Intel Xeon E5-2680 Broadwell (v4) processors with base frequency of 2.4 GHz and DDR4 memory, offering stable computing performance
- Up to 48 cores and 192GB are available for sale.
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

- Scenarios that require sending and receiving massive network packets, such as game services, video services, and financial analysis
- Enterprise applications of different types and sizes

## Requirements

- S2ne instances are pay-as-you-go.
- S2ne instances can only be launched in a VPC.
- See below for available S2ne specifications:

Specifications	V C P U	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S2ne.SMALL2	1	2	120,000	1	1.5	2.4G Hz	-
S2ne.MEDIUM2	2	2	120,000	1	1.5	2.4G Hz	-

S2ne.MEDI UM4	2	4	150,000	2	1.5	2.4G Hz	–
S2ne.LAR GE4	4	4	300,000	2	1.5	2.4G Hz	–
S2ne.LAR GE8	4	8	300,000	2	1.5	2.4G Hz	–
S2ne.LAR GE16	4	16	300,000	2	1.5	2.4G Hz	–
S2ne.2XLA RGE16	8	16	0.6 million	2	2.0	2.4G Hz	–
S2ne.2XLA RGE32	8	32	0.6 million	2	2.0	2.4G Hz	–
S2ne.3XLA RGE24	1 2	24	900,000	3	2.5	2.4G Hz	–
S2ne.3XLA RGE48	1 2	48	900,000	3	2.5	2.4G Hz	–
S2ne.4XLA RGE32	1 6	32	1.2 million	4	3.5	2.4G Hz	–
S2ne.4XLA RGE64	1 6	64	1.2 million	4	3.5	2.4G Hz	–
S2ne.6XLA RGE48	2 4	48	1,800,000	6	5.0	2.4G Hz	–
S2ne.6XLA RGE96	2 4	96	1,800,000	6	5.0	2.4G Hz	–
S2ne.8XLA RGE64	3 2	64	2.4 million	8	6.5	2.4G Hz	–
S2ne.8XLA RGE128	3 2	128	2.4 million	8	6.5	2.4G Hz	–
S2ne.12XL ARGE192	4 8	192	3,600,000	12	9.5	2.4G Hz	–

## Standard S2

Standard S2 instances are a relatively new generation of instances. This family provides a balance of computing, memory, and network resources, and it is a good choice for many applications.

Standard S2 instances are equipped with Intel® Xeon® Broadwell processors with DDR4 memory.

## Features

- Intel Xeon E5–2680 Broadwell (v4) processors with base frequency of 2.4 GHz and DDR4 memory
- Up to 56 cores and 224 GB are available for sale.
- 1:2 or 1:4 processor to memory ratio
- Balance of computing, memory, and network resources

## Use Cases

This family is used for small and mid-size databases, data processing tasks that require additional memory and cache fleets, and for running backend servers for SAP, Microsoft SharePoint, cluster computing and other enterprise applications.

## Requirements

- S2 instances are pay-as-you-go, and can also be used as production instances of standard host HS20 in CDHs.
- S2 instances can be launched in classic networks and VPCs.
- S2 instances are available for purchase with configurations detailed below. Ensure that the selected S2 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S2.SMALL1	1	1	200,000	250,000	1	1.5	2.4 GHz	–

								H z	
S2.SMALL2	1	2	200,000	250,000	1	1.5		2.4 G H z	-
S2.SMALL4	1	4	200,000	250,000	1	1.5		2.4 G H z	-
S2.MEDIUM2	2	2	200,000	250,000	2	1.5		2.4 G H z	-
S2.MEDIUM4	2	4	250,000	250,000	2	1.5		2.4 G H z	-
S2.MEDIUM8	2	8	250,000	250,000	2	1.5		2.4 G H z	-
S2.LARGE4	4	4	0.45 million	350,000	4	1.5		2.4 G H z	-
S2.LARGE8	4	8	0.45 million	350,000	4	1.5		2.4 G H z	-
S2.LARGE16	4	16	0.45 million	350,000	4	1.5		2.4 G	-

								H z	
S2.2XL ARGE16	8	16	0.5 million	0.7 millio n	8	1.5		2. 4 G H z	-
S2.2XL ARGE3 2	8	32	0.5 million	0.7 millio n	8	1.5		2. 4 G H z	-
S2.3XL ARGE2 4	1 2	24	0.5 million	1 millio n	8	2.5		2. 4 G H z	-
S2.3XL ARGE4 8	1 2	48	0.5 million	1 millio n	8	2.5		2. 4 G H z	-
S2.4XL ARGE3 2	1 6	32	0.5 million	1,300 ,000	8	3.0		2. 4 G H z	-
S2.4XL ARGE6 4	1 6	64	0.5 million	1,300 ,000	8	3.0		2. 4 G H z	-
S2.6XL ARGE4 8	2 4	48	0.7 million	2,00 0,00 0	8	4.5		2. 4 G H z	-

S2.6XL ARGE9 6	2 4	96	0.7 million	2,00 0,00 0	8	4.5	2. 4 G H Z	–
S2.8XL ARGE6 4	3 2	64	0.7 million	2,60 0,00 0	8	6.0	2. 4 G H Z	–
S2.8XL ARGE12 8	3 2	128	0.7 million	2,60 0,00 0	8	6.0	2. 4 G H Z	–
S2.14XL ARGE2 24	5 6	224	0.7 million	4.5 millio n	8	10.0	2. 4 G H Z	–

## Standard S1

Series 1 Standard S1 instances are cloud server types that cover a range of low to high core configurations, meeting users' flexible configuration needs and featuring moderate pricing and customizable options. Data disks can be paired with different types, such as local disks and SSD cloud disks, depending on the hardware specifications.

### Features

The Standard S1 instance type has the following features:

- Offer a range of CPUs from low to high core count for flexible configuration options
- Intel Xeon CPUs and DDR3 memories
- Storage can be configured with local disks and SSD cloud disks.
- Balance of computing, memory, and network resources

### Use Cases

Standard S1 instances are suitable for large, medium, and small-sized applications and databases.

### Requirements

- S1 instances are pay-as-you-go, and can also be used as production instances of standard host in CDHs.
- S1 instances can be launched in basic networks and VPCs.
- In the instance specification list, a "-" indicates that there is no fixed commitment for network packet transmission and processor frequency performance for that specification. If you have a strong demand for consistent performance, it is recommended to choose models with performance commitments.
- S1 instances are available for purchase with configurations detailed below. Ensure that the selected S1 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	v CPU	Memory (GB)	Through put (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
S1.SM ALL1	1	1	-	250,000	1	1.5	-	-
S1.SM ALL2	1	2	-	250,000	1	1.5	-	-
S1.SM ALL4	1	4	-	250,000	1	1.5	-	-
S1.ME DIUM2	2	2	-	250,000	1	1.5	-	-
S1.ME DIUM4	2	4	-	250,000	2	1.5	-	-
S1.ME DIUM8	2	8	-	250,000	2	1.5	-	-
S1.ME DIUM12	2	12	-	250,000	2	1.5	-	-

S1.LARGE4	4	4	-	350,000	4	1.5	-	-
S1.LARGE8	4	8	-	350,000	4	1.5	-	-
S1.LARGE16	4	16	-	350,000	4	1.5	-	-
S1.2XLARGE8	8	8	-	0.7 million	8	2.0	-	-
S1.2XLARGE16	8	16	-	0.7 million	8	2.0	-	-
S1.2XLARGE32	8	32	-	0.7 million	8	2.0	-	-
S1.3XLARGE24	12	24	-	1 million	8	2.5	-	-
S1.3XLARGE48	12	48	-	1 million	8	2.5	-	-
S1.4XLARGE16	16	16	-	1.4 million	8	3.5	-	-
S1.4XLARGE32	16	32	-	1.4 million	8	3.5	-	-
S1.4XLARGE48	16	48	-	1.4 million	8	3.5	-	-
S1.4XLARGE64	16	64	-	1.4 million	8	3.5	-	-
S1.6XLARGE48	24	48	-	2,000,000	8	5.0	-	-

S1.8XL ARGE6 4	3 2	64	–	2,700, 000	8	7.0	–	–
S1.12X LARGE 96	4 8	96	–	4,000, 000	8	10.0	–	–

**Note:**

In the S1 Instance Specification list, a "-" in the Network Packets and Base Frequency columns indicates that there is no fixed commitment for network packet transmission and processor base frequency performance. If you have a strong demand for consistent performance, it is recommended to choose a model with performance commitments.

## Memory Optimized Instance Family

Memory Optimized instances feature large memory and are suitable for applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching.

### MEM optimized MA3

As the latest generation of memory optimized instances, MA3 instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- Provides high performance with high reliability, security, and stability based on the Star Lake servers developed by Tencent Cloud
- Equipped with AMD EPYC™ Milan processor, with a base frequency of 2.55GHz and a turbo frequency of 3.5GHz.
- The processor-to-memory ratio is 1:8.
- Purchase the lowest price for the same memory size
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

### Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- MA3 instances are pay-as-you-go.
- The network of MA3 instances must be based on a VPC.
- We recommend you use the instance together with the TencentOS Server operating system to enjoy the best application performance.
- MA3 instances are available for purchase with various configurations. Please refer to the instance specifications below. Ensure that the selected MA3 instance size meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
MA3.SMALL8	1	8	250,000	1	1.5	2.55 GHz	–
MA3.MEDIUM16	2	16	300,000	2	1.5	2.55 GHz	–
MA3.LARGE32	4	32	0.5 million	4	2	2.55 GHz	–
MA3.2XLARGE64	8	64	8 million	8	4	2.55 GHz	–
MA3.4XLARGE128	16	128	1.5 million	16	7	2.55 GHz	–
MA3.8XLARGE256	32	256	2.5 million	32	14	2.55 GHz	–

## MEM optimized M6

As the latest generation of memory optimized instances, M6 instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- A new-generation Star Lake two-way server developed by Tencent Cloud, with the 3rd generation Intel® Xeon® scalable processor
- Intel® Xeon® Ice Lake processor, with a CPU clock rate of 2.7 GHz and a turbo frequency of 3.3 GHz
- The processor-to-memory ratio is 1:8.
- Lowest price per GB of memory among instance types
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

### Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

### Requirements

- M6 instances are pay-as-you-go.
- Supports only VPC.
- We recommend you use the instance together with the TencentOS Server operating system to enjoy the best application performance.
- M6 instances are available for purchase and configuration. Please refer to the instance specifications below. Ensure that the M6 instance size you choose meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.

- Instances can achieve up to 100 Gbps network bandwidth, depending on the support of the instance's operating system kernel version and runtime environment. When the PPS exceeds 10 million and the bandwidth is greater than 50 Gbps, the kernel protocol stack may cause significant network performance degradation. In this case, the netperf test bandwidth value may not meet expectations. You can use the DPDK method to mask the differences in the cloud server kernel protocol stack and obtain the true network performance of the instance. For testing methods, please refer to [High-Throughput Network Performance Testing](#).

Specifications	V C P U	Mem ory (GB)	Through put (pps) (received and sent packets)	Num ber of conn ectio ns	Num ber of Queu es	Private network bandwidth (Gbps) (outbound and inbound)	Cl oc k R ate	Rem arks
M6.SMALL8	1	8	250,000	250,000	1	1.5	2.7 GHz	-
M6.MEDIUM16	2	16	300,000	250,000	2	2	2.7 GHz	-
M6.LARGE32	4	32	0.6 million	0.5 million	4	4	2.7 GHz	-
M6.2XLARGE64	8	64	1.2 million	1 million	8	7	2.7 GHz	-
M6.4XLARGE128	16	128	2.5 million	2,000,000	16	13	2.7 GHz	-

							H z	
M6.8XL ARGE25 6	3 2	256	5,000,00 0	4,00 0,00 0	32	26	2. 7 G H z	–
M6.31M EDIUM4 70	6 2	470	10 million	8 millio n	48	50	2. 7 G H z	–
M6.31X LARGE 940	1 2 4	940	19 million	16,00 0,00 0	48	100	2. 7 G H z	–

## Security-enhanced Memory Optimized M6ce

As the latest generation of memory optimized instances, M6ce instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- A new-generation Star Lake two-way server developed by Tencent Cloud, with the 3rd generation Intel® Xeon® scalable processor
- Intel® Xeon® Ice Lake processor, with a CPU clock rate of 2.7 GHz and a turbo frequency of 3.3 GHz
- Intel® SGX encrypted computing, with a single instance containing up to 428 GB encrypted memory, ensuring the confidentiality and security of critical data and codes
- The processor-to-memory ratio is 1: 8. The encrypted memory accounts for about 50% of the memory
- Lowest price per GB of memory among instance types
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

### Use Cases

These instances are ideal for:

- Data sharing and computing, protecting confidential data shared between different users or vendors
- Blockchain applications, enhancing the privacy and security of transactions and key storage, etc.
- Scenarios with high security and reliability requirements, such as finance, government agencies, medical care, etc.
- Confidential computing scenarios and data encryption applications

## Requirements

- M6ce instances are pay-as-you-go.
- Supports only VPC.
- See below for available M6ce specifications:

### Note:

If you have used hardware-bound keys to encrypt data in Intel SGX Enclave (such as SGX Sealing), you will not be able to decrypt the corresponding data after the host changes in the following scenarios. It is recommended to back up your data at the application layer to ensure application reliability. Scenarios that may cause host changes include host failure, adjusting instance specifications, using the shutdown without charge feature, and modifying instance placement groups.

For more information on how to use security-enhanced memory M6ce instances, please see [Building an SGX Confidential Computing Environment](#).

Specifications	V C P U	Mem ory (GB)	En cry pt ed me mo ry (G B)	Netw ork packe ts trans mitted and receiv ed (pps) (out + in)	Nu mbe r of con nect ions	Nu mbe r of Que ues	Private Networ k Bandwi dth Capabi lity (Gbps) (Outbo und + Inbound)	Clo ck Rat e	Re ma rks
M6ce.ME DIUM16	2	16	8	300,0 00	250, 000	2	2	2.7 /3.	-

								3G Hz	
M6ce.LARGE32	4	32	16	0.6 million	0.5 milli on	4	4	2.7 /3. 3G Hz	–
M6ce.2XLARGE64	8	64	32	1.2 million	1 milli on	8	7	2.7 /3. 3G Hz	–
M6ce.4XLARGE128	1 6	128	64	2.5 million	2,00 0,00 0	16	13	2.7 /3. 3G Hz	–
M6ce.8XLARGE256	3 2	256	128	5,000, 000	4,00 0,00 0	32	26	2.7 /3. 3G Hz	–
M6ce.31XMEDIUM428	6 2	428	214	10 million	8 milli on	48	50	2.7 /3. 3G Hz	–
M6ce.31XLARGE856	1 2 4	856	42 8	19 million	8 milli on	48	100	2.7 /3. 3G Hz	–

## MEM optimized M6p

As the latest generation of memory optimized instances, M6p instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- A new-generation Star Lake two-way server developed by Tencent Cloud, with the 3rd generation Intel® Xeon® scalable processor
- Intel® Xeon® Ice Lake processor, with a CPU clock rate of 2.7 GHz and a turbo frequency of 3.3 GHz

- Paired with Intel® 2nd Gen Optane™ Persistent Memory (BPS), offering a cost-effective memory solution with high performance.
- Lowest price per GB of memory among instance types

## Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Large memory application scenarios such as Hadoop or Redis clusters
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- M6p instances are pay-as-you-go.
- The network of M6p instances must be based on a VPC.
- M6p instances do not support configuration adjustments and the "No Charge when Shut Down" feature.
- See below for available M6p specifications:

### Note:

The reliability of data in persistent memory depends on the reliability of the physical server and the persistent memory device, posing a risk of data loss (e.g., during host downtime). It is recommended that you store data with long-term storage reliability requirements on cloud disks.

For more information on how to use memory-optimized M6p instances, see [Configuring Persistent Memory for M6p Instances](#).

Specifications	vCPU	Memory (GB)	Persistent Memory (GB)	Network packets transmitted and received (pps)	Number of connections	Number of Queues	Private Network Bandwidth Capability (Outbound + Inbound)	Cloud Network Rate	Remarks

				(out + in)					
M6p.LARGE16	4	16	62	0.6 million	0.5 million	4	4	2.7 GHz	-
M6p.2XLARGE32	8	32	124	1.2 million	1 million	8	7	2.7 GHz	-
M6p.4XLARGE53	16	53	248	2.5 million	2,000,000	16	13	2.7 GHz	-
M6p.8XLARGE106	32	106	496	5,000,000	4,000,000	32	26	2.7 GHz	-
M6p.31XMEDIUM212	62	212	992	10 million	8 million	48	50	2.7 GHz	-
M6p.31XLARGE424	124	424	1984	19 million	16,000,000	48	100	2.7 GHz	-

## Memory Optimized M5

As the latest generation of memory optimized instances, M5 instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- Intel® Xeon® Cascade Lake or Intel® Xeon® Cooper Lake processor, with a CPU clock rate of 2.5 GHz and turbo frequency of 3.1 GHz
- New Intel Advanced Vector Extension (AVX-512) instruction set
- CPU/RAM ratio: 1:8
- Lowest price per GB of memory among instance types
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- M5 instances are pay-as-you-go.
- M5 instances can only be launched in a VPC.
- Configuration purchase is available for M5 instances. Please see the instance specifications below.

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
M5.SMALL8	1	8	250,000	250,000	1	1.5	2.5 GHz	-
M5.MEDIUM16	2	16	300,000	250,000	2	1.5	2.5 GHz	-

M5.LARGE32	4	32	0.5 million	250,000	2	1.5	2.5 GHz	–
M5.2XLARGE64	8	64	8 million	250,000	2	3.0	2.5 GHz	–
M5.3XLARGE96	12	96	1 million	250,000	3	5.0	2.5 GHz	–
M5.4XLARGE128	16	128	1.5 million	300,000	4	6.0	2.5 GHz	–
M5.8XLARGE256	32	256	2.5 million	0.6 million	8	12.0	2.5 GHz	–
M5.16XLARGE512	64	512	5,000,000	1.2 million	16	23.0	2.5 GHz	–

## MEM optimized MA2

As the latest generation of memory optimized instances, MA2 instances are powered by an all-new optimized virtualization platform and designed to deliver fast and stable performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

### Features

- Provides high performance with high reliability, security, and stability based on the Star Lake servers developed by Tencent Cloud
- Equipped with AMD EPYC™ ROME processor, with a base frequency of 2.6GHz and a turbo frequency of 3.3GHz.
- CPU/RAM ratio: 1:8
- For instances with the same memory size, the lowest price is offered.

- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Requirements

- MA2 instances are pay-as-you-go.
- The network of MA2 instances must be based on a VPC.
- See below for available MA2 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
MA2.SMALL8	1	8	250,000	250,000	1	1.5	2.6 GHz	-
MA2.LARGE32	4	32	0.5 million	250,000	4	1.5	2.6 GHz	-
MA2.2XLARGE64	8	64	0.7 million	250,000	8	1.5	2.6 GHz	-

MA2.4XL ARGE128	1 6	128	1 million	300,000	16	3	2.6 GHz	–
MA2.8XL ARGE256	3 2	256	1.4 million	0.7 million	32	5	2.6 GHz	–
MA2.12XL LARGE384	4 8	384	2,100,000	1 million	48	7	2.6 GHz	–

## MEM optimized M4

M4 instances are designed to deliver fast performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

As a new generation of memory optimized instances, M4 instances come with new Xeon® Skylake processors with the latest DDR4 memory. They are network optimized by default, offering a private network throughput up to 6,000,000 pps, and support for ultra-high private network bandwidth.

## Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics
- Scenarios that require sending and receiving massive network packets, such as on-screen video comments, live video broadcasting, and gaming

## Features

- 2.4 GHz Intel® Xeon® Skylake 6148 processor, offering stable computing performance
- New Intel Advanced Vector Extension (AVX-512) instruction set
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- Higher specification, M4.16XLARGE512, offering 64 vCPUs and 512 GB memory
- 1:8 or 1:12 processor to memory ratio

- Lowest price per GB of memory among instance types
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- M4 instances are pay-as-you-go.
- Supports only VPC.
- See below for available M4 specifications:

Specifications	V C P U	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
M4.SMALL8	1	8	250,000	1	1.5	2.4 GHz	–
M4.MEDIUM16	2	16	300,000	2	1.5	2.4 GHz	–
M4.LARGE32	4	32	0.5 million	2	1.5	2.4 GHz	–
M4.2XLARGE64	8	64	8 million	2	3.0	2.4 GHz	–
M4.3XLARGE96	12	96	1 million	3	4.0	2.4 GHz	–
M4.3XLARGE144	12	144	1 million	3	4.0	2.4 GHz	–
M4.4XLARGE128	16	128	1.5 million	4	6.0	2.4 GHz	–
M4.4XLARGE192	16	192	1.5 million	4	6.0	2.4 GHz	–
M4.8XLARGE256	32	256	2.5 million	8	11.0	2.4 GHz	–

M4.8XLA RGE384	32	384	2.5 million	8	11.0	2.4 GHz	–
M4.16XLA RGE512	64	512	5,000,000	16	22.0	2.4 GHz	–
M4.18XLA RGE648	72	648	6 million	16	24.0	2.4 GHz	–

## MEM optimized M3

M3 instances are designed to deliver fast performance for workloads that process large data sets in memory. They are ideal for applications that require high in-memory computing performance.

M3 instances are powered by new Intel® Xeon® Skylake processors with the latest DDR4 memory, and support up to 10 Gbps of private network bandwidth.

## Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics

## Features

- 2.5 GHz Intel Xeon® Skylake 6133 processors with stable computing performance
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- Higher specification, M3.16XLARGE512, offering 64 vCPUs and 512 GB memory
- 1:8 or 1:12 processor to memory ratio
- Lowest price per GB of memory among instance types
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- M3 instances are pay-as-you-go.
- Supports both classic network and VPC.
- See below for available M3 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
M3.SMALL8	1	8	300,000	250,000	1	1.5	2.5 GHz	-
M3.MEDIUM16	2	16	400,000	250,000	2	1.5	2.5 GHz	-
M3.LARGE32	4	32	8 million	300,000	4	1.5	2.5 GHz	-
M3.2XLARGE64	8	64	1.2 million	0.6 million	8	1.5	2.5 GHz	-
M3.3XLARGE96	12	96	1.2 million	900,000	12	1.5	2.5 GHz	-
M3.3XLARGE144	12	144	1.2 million	900,000	12	1.5	2.5 GHz	-
M3.4XLARGE128	16	128	1.2 million	1.2 million	16	2.0	2.5	-

8				n			G H Z	
M3.4XL ARGE19 2	1 6	192	1.2 million	1.2 millio n	16	2.0	2. 5 G H Z	–
M3.8XL ARGE2 56	3 2	256	1.2 million	2.4 millio n	16	4.0	2. 5 G H Z	–
M3.8XL ARGE3 84	3 2	384	1.2 million	2.4 millio n	16	4.0	2. 5 G H Z	–
M3.16X LARGE 512	6 4	512	1.2 million	4.8 millio n	16	8.0	2. 5 G H Z	–

## MEM optimized M2

Memory-optimized M2 instances are designed to deliver fast performance for workloads that process large in-memory data sets. Characterized by their large memory capacity, they are the optimal choice for high-memory computing applications.

### Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics

### Features

- 2.4 GHz Intel® Xeon® E5-2680v4 processor with DDR4 memory
- Memory of up to 448 GB is available. The M2.14XLARGE448 model offers 56 vCPUs and 448 GB memory

- CPU/RAM ratio: 1:8
- Lowest price per GB of memory among instance types

## Requirements

- M2 instances are pay-as-you-go, and can also be used as production instances of memory optimized host HM20 in CDHs.
- Supports only VPC.
- See below for available M2 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
M2.SMALL8	1	8	200,000	250,000	1	1.5	2.4 GHz	-
M2.MEDIUM16	2	16	250,000	250,000	2	1.5	2.4 GHz	-
M2.LARGE32	4	32	0.45 million	350,000	4	1.5	2.4 GHz	-
M2.2XLARGE64	8	64	0.5 million	0.7 million	8	1.5	2.4 GHz	-
M2.3XLARGE96	12	96	0.5 million	1 million	8	2.5	2.4 GHz	-

							H z	
M2.4XL ARGE12 8	1 6	128	0.5 million	1,300, 000	8	3.0	2. 4 G H z	–
M2.6XL ARGE19 2	2 4	192	0.7 million	2,000, 000	8	4.5	2. 4 G H z	–
M2.8XL ARGE25 6	3 2	256	0.7 million	2,600, 000	8	6.0	2. 4 G H z	–
M2.12X LARGE 384	4 8	384	0.7 million	3.9 million	8	9.0	2. 4 G H z	–
M2.14X LARGE 448	5 6	448	0.7 million	4.5 million	8	10.0	2. 4 G H z	–

## MEM optimized M1

M1 instances come with approximately 1:8 CPU to memory ratio. They are well suited for applications that require memory-intensive operations, searches, and computations, such as high-performance databases and distributed in-memory caching.

### Features

- 2.3 GHz Intel Xeon® E5-2670 v3 processor with DDR3 memory, providing more instance sizes with stronger computing capacities
- Golden CPU to memory ratio for memory-intensive applications, satisfying the needs for large-scale business deployment

### Use Cases

These instances are ideal for:

- Applications that require extensive memory operations, searches, and computing, such as high-performance databases and distributed memory caching
- External Hadoop clusters or Redis for fields such as computational genomics

## Requirements

- M1 instances are pay-as-you-go.
- Supports both classic network and VPC.
- In the instance specification list, a "-" indicates that there is no fixed commitment for network packet transmission and processor frequency performance for that specification. If you have a strong demand for consistent performance, it is recommended to choose models with performance commitments.
- See below for available M1 specifications:

Specifications	V C P U	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
M1.SMALL8	1	8	-	250,000	1	1.5	2.3 GHz	-
M1.MEDIUM16	2	16	-	250,000	2	1.5	2.3 GHz	-
M1.LARGE32	4	32	-	350,000	4	1.5	2.3 GHz	-
M1.2XLARGE64	8	64	-	0.7 million	8	2.0	2.3 GHz	-

							H z	
M1.3XL ARGE9 6	1 2	96	-	1 millio n	8	2.5	2. 3 G H z	-
M1.4XL ARGE12 8	1 6	128	-	1.4 millio n	8	3.5	2. 3 G H z	-
M1.6XL ARGE19 2	2 4	192	-	2,00 0,00 0	8	5.0	2. 3 G H z	-
M1.8XL ARGE25 6	3 2	256	-	2,70 0,00 0	8	7.0	2. 3 G H z	-
M1.12XL ARGE36 8	4 8	368	-	4,00 0,00 0	8	10.0	2. 3 G H z	-

**Note:**

In the M1 instance specification list, a "-" in the network packet transmission and reception column indicates that there is no fixed commitment for network packet transmission and reception for that specification. If you have a strong demand for consistent performance, it is recommended to choose a model with a performance commitment.

## High I/O Family

This family features high random IOPS, high throughput and low latency, making it suitable for I/O-intensive applications that require high disk read/write performance and low latency, such as high-performance databases.

**Note:**

The data disks of IT3 and IT5 instances are local storage, which are at risk of data loss (e.g., during host downtime). If your application cannot ensure data reliability, it is not recommended to use IT3 or IT5 instances. We strongly advise using instances that allow you to choose cloud disks as data disks.

## High IO IT5

High I/O IT5 instances are the latest generation of high I/O instances designed specifically for I/O-intensive workloads. Based on NVMe SSD instance storage, they provide low-latency, ultra-high IOPS, and high-throughput storage resources, offering up to 2.05 million IOPS and up to 11GB/s throughput capacity at a relatively low cost. They are suitable for high-performance relational databases, Elasticsearch, and other I/O-intensive businesses.

### Use Cases

- High-performance databases, NoSQL databases (e.g. MongoDB), and clustered databases
- I/O intensive applications that require low latency, such as online transaction processing (OLTP) systems, and Elasticsearch

### Features

- 2.5 GHz Intel® Xeon® Cascade Lake processor, offering stable computing performance
- The latest generation of 6-channel DDR4 memory
- Supports up to 23 Gbps of private network bandwidth to meet extremely high private network transmission requirements
- NVMe SSD is used for instance storage, offering low latency and ultra-high IOPS
  - Up to 650,000 random read IOPS (block size of 4 KB), and up to 2.8 GB/s sequential read throughput (block size of 128 KB) for a single disk
  - Up to 2.05 million random read IOPS (block size of 4 KB), and up to 11 GB/s sequential read throughput (block size of 128 KB) for a CVM

### Requirements

- If the CVM instance crashes, we will notify you and perform fixes.
- IT5 instances without monitoring components installed may result in the platform being unable to perform more detailed monitoring. If an instance encounters a failure, proper notifications may not be sent, posing a high risk. Please refer to [Installing CVM Monitoring Components](#) to complete the installation of monitoring components.
- IT5 instances are pay-as-you-go.
- Supports only VPC.

- IT5 instances do not support configuration adjustment and the "No Charge when Shut Down" feature.
- See below for available IT5 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
IT5.4XLRARGE64	16	64	1.5 million	1.5 million	4	6	2.5 GHz	1 x 3570 GB local NVMe SSD
IT5.8XLRARGE128	32	128	2.5 million	300K	8	12	2.5 GHz	2 x 3570 GB local NVMe SSD
IT5.16XLARGE256	64	256	5,000,000	6 million	16	23	2.5 GHz	4 x 3570 GB local NVMe SSD

## High I/O IT3

High I/O IT3 instances are the latest generation of high I/O instances designed specifically for I/O-intensive workloads. Based on NVMe SSD instance storage, they provide low-latency, extremely high IOPS, and high-throughput storage resources, offering up to 1.8 million IOPS and up to 11 GB/s throughput capacity at a relatively low cost. They are suitable for high-performance relational databases, Elasticsearch, and other I/O-intensive businesses. IT3 instances feature inexpensive and intensive storage. They can provide each CPU core with higher IOPS and larger private network bandwidth.

## Use Cases

- High-performance databases, NoSQL databases (e.g. MongoDB), and clustered databases

- I/O intensive applications that require low latency, such as online transaction processing (OLTP) systems, and Elasticsearch

## Features

- 2.5 GHz Intel Xeon® Skylake 6133 processors with stable computing performance
- The latest generation of 6-channel DDR4 memory
- Supports up to 23 Gbps of private network bandwidth to meet extremely high private network transmission requirements
- NVMe SSD is used for instance storage, offering low latency and ultra-high IOPS
  - Up to 620,000 random read IOPS (block size of 4 KB), and up to 2.9 GB/s sequential read throughput (block size of 128 KB) for a single disk
  - Up to 1.8 million random read IOPS (block size of 4 KB), and up to 11 GB/s sequential read throughput (block size of 128 KB) for a CVM

## Requirements

- If the CVM instance crashes, we will notify you and perform fixes.
- IT3 instances without monitoring components installed may result in the platform being unable to perform more detailed monitoring. If an instance encounters a failure, proper notifications may not be sent, posing a high risk. Please refer to [Installing CVM Monitoring Components](#) to complete the installation of monitoring components.
- IT3 instances are pay-as-you-go.
- Supports only VPC.
- IT3 instances do not support configuration adjustment and the no charges when shut down feature.
- See below for available IT3 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
IT3.4XLA RGE64	16	64	1.5 million	4	6	2.5 GHz	1 × 3720 GB local NVMe SSD
IT3.8XLA RGE128	32	128	2.5 million	8	12	2.5 GHz	2 × 3720 GB local

							NVMe SSDs
IT3.16XL ARGE25 6	6 4	25 6	5,000,000	16	23	2.5 GHz	4 × 3720 GB local NVMe SSD

## Big Data Family

The big data family is equipped with massive storage resources, features high throughput, and is suitable for throughput-intensive applications such as Hadoop distributed computing, massive log processing, distributed file systems, and large data warehouses.

### Note:

D3 and D2 instances use local disk as the data disk, which may lose data (e.g., when the host crashes). If your application cannot guarantee data reliability, we recommend you choose an instance that can use cloud disks as the data disk.

## D3

As the latest generation of big data instances, D3 instances are equipped with massive storage resources, and carry up to 94 TB SATA HDD local storage, making them suitable for throughput-intensive services, such as Hadoop distributed computing and parallel data processing.

## Use Cases

- Distributed computing services such as Hadoop MapReduce/HDFS/Hive/HBase
- Workloads such as Elasticsearch, log processing, and large data warehouse
- Customers in the Internet, finance, and industries that require big data computing and storage analysis, as well as workloads that require massive data storage and computing

## Features

- 2.5 GHz Intel® Xeon® Cascade Lake processor with DDR4 memory
- Instances carry up to 24 local disks of 4 TB and up to 94 TB of HDD local storage.
  - At least 190 MB/s sequential read throughput performance for a single disk
  - Millisecond read/write latency
- Processor to memory ratio 1:4, tailored for the big data scenarios

## Requirements

- If the CVM instance crashes, we will notify you and perform fixes.

- D3 instances without monitoring components installed may result in the platform being unable to perform more detailed monitoring. If an instance encounters a failure, proper notifications may not be sent, posing a high risk. Please refer to [Installing CVM Monitoring Components](#) to complete the installation of monitoring components.
- Big Data D3 instances are pay-as-you-go.
- Supports only VPC.
- The configuration of D3 instances cannot be changed.
- See below for available D3 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
D3.2XLARGE32	8	32	8 million	250,000	2	4.0	2.5 GHz	4 x 3720 GB local SATA HDDs
D3.4XLARGE64	16	64	1.5 million	300,000	4	7.0	2.5 GHz	8 x 3720 GB local SATA HDDs
D3.8XLARGE128	32	128	2.5 million	0.6 million	8	14.0	2.5 GHz	12 x 3720 GB local SATA HDDs
D3.16XLARGE256	64	256	5,000,000	1.2 million	12	27.0	2.5 GHz	24 x 3720 GB local SATA HDDs
D3.21XLARGE320	84	320	6 million	1,600,000	16	32.0	2.5 GHz	24 x 3720 GB local

								H z	SATA HDDs
--	--	--	--	--	--	--	--	--------	--------------

## Big Data D2

D2 instances are equipped with massive storage resources, and carry up to 144 TB SATA HDD local storage, making them suitable for throughput-intensive services, such as Hadoop distributed computing and parallel data processing.

### Use Cases

- Distributed computing services such as Hadoop MapReduce/HDFS/Hive/HBase
- Workloads such as Elasticsearch, log processing, and large data warehouse
- Customers in the Internet, finance, and industries that require big data computing and storage analysis, as well as workloads that require massive data storage and computing

### Features

- 2.4 GHz Intel® Xeon® Skylake 6148 processor with DDR4 memory
- Instances carry up to 12 local disks of 12 TB and up to 144 TB of HDD local storage
  - For a single disk, the sequential read throughput is 220+ MB/s and sequential write throughput is 220+ MB/s (128 KB block size and depth of 32)
  - Overall instance throughput can reach up to 2.8 GB/s (128 KB block size and depth of 32)
  - Read/write latency is as low as 2–5 ms
- Higher specification, D2.19XLARGE320, offering 76 vCPUs and 320 GB of memory
- Local storage with a unit price as low as 1/10 of S2, making its total cost close to that of a self-built Hadoop clusters in IDC
- Processor to memory ratio 1:4, tailored for the big data scenarios

### Requirements

- If the CVM instance crashes, we will notify you and perform fixes.
- Big Data D2 instances are pay-as-you-go.
- Supports only VPC.
- The configuration of D2 instances cannot be changed.
- See below for available D2 specifications:

Specifications	v C P U	Me mo ry (G B)	Through put (pps) (receive d and	Numb er of Queu es	Private network bandwidth (Gbps) (outbound	Cloc k Rat e	Remarks

			sent packets )		and inbound)		
D2.2XLAR GE32	8	32	8 million	2	3.0	2.4 GHz	Equipped with one 11176 GB SATA HDD local disk
D2.4XLAR GE64	16	64	1.5 million	4	6.0	2.4 GHz	Equipped with two 11176 GB SATA HDD local disks
D2.6XLAR GE96	24	96	2,000,0 00	6	8.0	2.4 GHz	Equipped with three 11176 GB SATA HDD local disks
D2.8XLAR GE128	32	128	2.5 million	8	11.0	2.4 GHz	Equipped with four 11176 GB SATA HDD local disks
D2.16XLA RGE256	64	256	5,000,0 00	16	22.0	2.4 GHz	Equipped with eight 11176 GB SATA HDD local disks
D2.19XLA RGE320	76	320	6 million	16	25.0	2.4 GHz	Equipped with twelve 11176 GB SATA HDD local disks

## Computing Instance Family

Compute-optimized instances offer CPU frequencies of up to 3.8GHz, providing the highest single-core computing performance. They are suitable for compute-intensive applications such as batch processing, high-performance computing, and large-scale game servers.

### C6 for computing tasks

The C6 Compute instances are the latest generation of compute instances, providing stable, reliable, and high-performance computing services based on a new-generation virtualization platform. They offer greater bandwidth and lower latency. With the highest base frequency processors and the best price-performance ratio in CVM, they are the ideal choice for applications with high computational performance and high concurrent read/write requirements that are constrained by computing limitations.

C6 instances come with new Intel® Xeon® processors, and support up to 100 Gbps of private network bandwidth.

## Use Cases

These instances are ideal for:

- Batch processing workloads and high performance computing (HPC)
- High-traffic Web frontend server
- Other compute-intensive services such as massively multiplayer online (MMO) game servers

## Features

- A new-generation Star Lake two-way server developed by Tencent Cloud, with the 3rd generation Intel® Xeon® scalable processor
- Intel® Xeon® Ice Lake processor, with a CPU clock rate of 3.2 GHz and a turbo frequency of 3.5 GHz
- Offers various processor-to-memory ratios, such as 1:2 and 1:4.
- Up to 100 Gbps private network bandwidth for ultra high packet throughput to meet extremely high private network transfer requirements

## Requirements

- C6 instances are pay-as-you-go.
- We recommend you use the instance together with the TencentOS Server operating system to enjoy the best application performance.
- C6 instances are available for purchase with various configurations. Please refer to the instance specifications below. Ensure that the C6 instance size you choose meets the minimum CPU and memory requirements for your operating system and applications. In many use cases, operating systems with resource-intensive graphical user interfaces (such as Windows) may require larger instance sizes. As your workload's demand for memory and CPU increases over time, you can scale up to higher configurations or choose other instance types.
- Instances can achieve up to 100 Gbps network bandwidth, depending on the support of the instance's operating system kernel version and runtime environment. When the PPS exceeds 10 million and the bandwidth is greater than 50 Gbps, the kernel protocol stack may

cause significant network performance degradation. In this case, the netperf test bandwidth value may not meet expectations. You can use the DPDK method to mask the differences in the cloud server kernel protocol stack and obtain the true network performance of the instance. For testing methods, please refer to [High-Throughput Network Performance Testing](#).

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
C6.LARGE8	4	8	8 million	0.6 million	4	5	3.2 GHz	-
C6.LARGE16	4	16	8 million	0.6 million	4	5	3.2 GHz	-
C6.2XLARGE16	8	16	1,600,000	1,300,000	8	9	3.2 GHz	-
C6.2XLARGE32	8	32	1,600,000	1,300,000	8	9	3.2 GHz	-
C6.4XLARGE32	16	32	3,300,000	2,700,000	16	18	3.2 GHz	-
C6.4XLARGE64	16	64	3,300,000	2,700,000	16	18	3.2 GHz	-

C6.8XL ARGE12 8	3 2	12 8	6,600,00 0	5,500, 000	32	35	3. 2 G Hz	–
C6.23M EDIUM2 16	4 6	21 6	9.5 million	8 million	46	50	3. 2 G Hz	–
C6.16XL ARGE25 6	6 4	25 6	13,000,0 00	11,000 ,000	48	70	3. 2 G Hz	–
C6.20XL ARGE32 0	8 0	32 0	16,000,0 00	13,900 ,000	48	87	3. 2 G Hz	–
C6.23XL ARGE43 2	9 2	43 2	19 million	16,000 ,000	48	100	3. 2 G Hz	–

## C5 for computing tasks

Compute-optimized C5 instances are the latest generation of compute instances, providing stable and reliable high-performance computing services based on a cutting-edge virtualization platform. They offer greater bandwidth and lower latency. With the highest base frequency processors and the best price-performance ratio in CVM, they are the ideal choice for compute-intensive and high-concurrency read-write applications that are constrained by computing limitations.

C5 instances come with the latest Xeon® Cooper Lake processors.

### Features

- Intel® Xeon® Cooper Lake processor, with a base frequency of 3.4 GHz and turbo frequency of 3.8 GHz
- New Intel Brain Floating Point 16-bit (bfloat16) instruction set
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Use Cases

These instances are ideal for:

- Batch processing workloads and high performance computing (HPC)
- High-traffic Web frontend server
- Other compute-intensive services such as massively multiplayer online (MMO) game servers

## Requirements

- C5 instances are pay-as-you-go.
- The network of C5 instances must be based on a VPC.
- See below for available C5 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
C5.LARGE8	4	8	0.5 million	250,000	4	1.5	3.4 GHz	-
C5.LARGE16	4	16	0.5 million	250,000	4	1.5	3.4 GHz	-
C5.2XLARGE16	8	16	8 million	250,000	8	3.0	3.4 GHz	-
C5.2XLARGE32	8	32	8 million	250,000	8	3.0	3.4 GHz	-
C5.4XLARGE32	16	32	1.5 million	300,000	8	6.0	3.4 GHz	-

C5.4XL ARGE64	1 6	64	1.5 million	300,0 00	8	6.0	3. 4 G Hz	–
C5.8XL ARGE64	3 2	64	2.5 million	0.6 million	16	12.0	3. 4 G Hz	–
C5.8XL ARGE12 8	3 2	12 8	2.5 million	0.6 million	16	12.0	3. 4 G Hz	–
C5.12XL ARGE96	4 8	96	4,000,00 0	900,0 00	24	18.0	3. 4 G Hz	–
C5.12XL ARGE19 2	4 8	19 2	4,000,00 0	900,0 00	24	18.0	3. 4 G Hz	–
C5.13XL ARGE18 4	5 2	18 4	4,000,00 0	1 million	32	19.0	3. 4 G Hz	–
C5.16XL ARGE25 6	6 4	25 6	5,000,00 0	1.2 million	32	24.0	3. 4 G Hz	–
C5.26XL ARGE36 8	1 0 4	36 8	6 million	2,000, 000	32	36.0	3. 4 G Hz	–

## Compute Optimized C4

Compute-optimized C4 instances are the next-generation compute instances with up to 25 Gbps internal network bandwidth, offering greater bandwidth and lower latency. They provide the highest base frequency processors and the best price-performance ratio in CVM, making

them an ideal choice for applications with high computational performance and high concurrent read/write operations that are compute-bound.

C4 instances are equipped with the latest Xeon® Cascade Lake processors, and support up to 25 Gbps of private network bandwidth.

## Use Cases

These instances are ideal for:

- Batch processing workloads and high performance computing (HPC)
- High-traffic Web frontend server
- Other compute-intensive services such as massively multiplayer online (MMO) game servers

## Features

- 3.2 GHz Intel® Xeon® Cascade Lake processor, with a turbo frequency 3.7 GHz
- New Intel Advanced Vector Extension (AVX-512) instruction set
- Supports up to 25 Gbps of private network bandwidth to meet extremely high private network transmission requirements
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- Supports pay-as-you-go.
- C4 instances can be launched in classic networks and VPCs.
- See below for available C4 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
C4.LARGE8	4	8	0.6 million	350,000	4	2.0	3.2 GHz	-

C4.LARGE16	4	16	0.6 million	350,000	4	2.0	3.2 GHz	–
C4.2XLARGE16	8	16	1 million	0.7 million	8	4.0	3.2 GHz	–
C4.2XLARGE32	8	32	1 million	0.7 million	8	4.0	3.2 GHz	–
C4.4XLARGE64	16	64	1 million	1,300,000	16	7.0	3.2 GHz	–
C4.8XLARGE174	32	174	1 million	2.5 million	16	13.0	3.2 GHz	–
C4.16XLARGE348	64	348	1 million	5,000,000	16	25.0	3.2 GHz	–

## Compute Network-optimized CN3

The CN3 Compute Optimized instances are the next-generation compute-optimized instances, with a maximum internal network bandwidth of up to 25 Gbps, offering greater bandwidth and lower latency. They provide the highest base clock frequency processors and the best price-performance ratio in CVM, making them an ideal choice for compute-intensive and high-concurrency read/write applications with computational constraints.

CN3 instances are equipped with the latest Xeon® Skylake processors, and support up to 25 Gbps of private network bandwidth, which is 2.5 times greater than Compute C3.

## Use Cases

These instances are ideal for:

- Batch processing workloads and high performance computing (HPC)
- High-traffic Web frontend server
- Other compute-intensive services such as massively multiplayer online (MMO) game servers

## Features

- 3.2 GHz Intel® Xeon® Skylake 6146 processor
- New Intel Advanced Vector Extension (AVX-512) instruction set
- Supports up to 25 Gbps of private network bandwidth to meet extremely high private network transmission requirements
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- CN3 instances are pay-as-you-go.
- Supports both classic network and VPC.
- See below for available CN3 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
CN3.LARGE8	4	8	0.6 million	350,000	4	3.0	3.2 GHz	-
CN3.LARGE16	4	16	0.6 million	350,000	4	3.0	3.2 GHz	-

								H z	
CN3.2X LARGE1 6	8	16	1 million	0.7 million	8	5.0		3. 2 G H z	-
CN3.2X LARGE 32	8	32	1 million	0.7 million	8	5.0		3. 2 G H z	-
CN3.4X LARGE 32	1 6	32	1 million	1.4 million	16	9.0		3. 2 G H z	-
CN3.4X LARGE 64	1 6	64	1 million	1.4 million	16	9.0		3. 2 G H z	-
CN3.8X LARGE 64	3 2	64	1 million	2,700, 000	16	17.0		3. 2 G H z	-
CN3.8X LARGE1 28	3 2	128	1 million	270	16	17.0		3. 2 G H z	-
CN3.12 XLARG E160	4 8	160	1 million	4,000, 000	16	25.0		3. 2 G H z	-

## C3 for computing tasks

C3 instances have the most powerful processors and the best cost performance among CVMs. They are ideal for applications that require high computing performance and high-concurrency reads/writes.

C3 instances are equipped with new Xeon® Skylake processors with the latest DDR4 memory, offering larger private network bandwidth.

## Use Cases

These instances are ideal for:

- Batch processing workloads and high performance computing (HPC)
- High-traffic Web frontend server
- Other compute-intensive services such as massively multiplayer online (MMO) game servers

## Features

- 3.2 GHz Intel® Xeon® Skylake 6146 processor
- New Intel Advanced Vector Extension (AVX-512) instruction set
- The latest generation of 6-channel DDR4 memory with a bandwidth of 2,666 MT/s
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- C3 instances are pay-as-you-go.
- Supports both classic network and VPC.
- See below for available C3 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
C3.LARGE8	4	8	0.6 million	350,000	2	2.5	3.2 G	-

								H z	
C3.LA RGE16	4	16	0.6 million	350,0 00	2	2.5		3. 2 G H z	-
C3.2XL ARGE1 6	8	16	1 million	0.7 million	2	3.0		3. 2 G H z	-
C3.2XL ARGE3 2	8	32	1 million	0.7 million	2	3.0		3. 2 G H z	-
C3.4XL ARGE3 2	1 6	32	1 million	1.4 million	4	4.5		3. 2 G H z	-
C3.4XL ARGE6 4	1 6	64	1 million	1.4 million	4	4.5		3. 2 G H z	-
C3.8XL ARGE6 4	3 2	64	1 million	2,700, 000	8	8.0		3. 2 G H z	-
C3.8XL ARGE1 28	3 2	128	1 million	2,700, 000	8	8.0		3. 2 G H z	-

## C2 for computing tasks

C2 instances are optimized with high-performance processors and deliver high cost performance. They are ideal for applications that require high computing performance and high-concurrency reads/writes.

## Use Cases

These instances are ideal for:

- Batch processing workloads
- High-traffic Web servers, and massively multiplayer online (MMO) game servers
- High-performance computing (HPC) and other compute-intensive applications.

## Features

- 3.2 GHz Intel® Xeon® E5-2667v4 processor with a turbo frequency of up to 3.6 GHz and DDR4 memory
- 1:2 or 1:4 processor to memory ratio
- The network performance of an instance depends on its specification. A higher specification means a greater network forwarding performance and a higher private network bandwidth cap.

## Requirements

- C2 instances are pay-as-you-go, and can also be used as production instances of computing host HC20 in CDHs.
- Supports both classic network and VPC.
- See below for available C2 specifications:

Specifications	VCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of connections	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
C2.LARGE8	4	8	550,000	400,000	4	2.5	3.2 GHz	-
C2.LARGE16	4	16	550,000	400,000	4	2.5	3.2 GHz	-

							H z	
C2.LARGE32	4	32	550,000	400,000	4	2.5	3.2 G H z	–
C2.2XLARGE16	8	16	550,000	8 million	8	3.5	3.2 G H z	–
C2.2XLARGE32	8	32	550,000	8 million	8	3.5	3.2 G H z	–
C2.4XLARGE32	16	32	850,000	1.5 million	8	6.0	3.2 G H z	–
C2.4XLARGE64	16	64	850,000	1.5 million	8	6.0	3.2 G H z	–
C2.8XLARGE96	32	96	850,000	300K	8	10.0	3.2 G H z	–

## Heterogeneous Computing Instance Family

This family is equipped with heterogeneous hardware such as GPU and FPGA to deliver real-time, fast parallel computing and floating-point computing capabilities. It is suitable for high-performance applications such as deep learning, scientific computing, video encoding/decoding, and graphics workstations.

AMD GPU series instances utilize AMD FirePro S7150, offering outstanding graphics processing capabilities. They are the top choice for application scenarios such as remote desktops, 3D rendering, and cloud gaming. [AMD GPU series instances >](#)

NVIDIA GPU series instances utilize NVIDIA Tesla series GPUs, including the previous generation M40, the current mainstream P4/P40, and the latest generation V100. They offer outstanding general computing capabilities, making them the top choice for deep learning training/inference, scientific computing, and other application scenarios. [NVIDIA GPU series instances >](#)

## Batch-based Instance Family

With the lowest per core-hour cost, this family is suitable for compute-intensive applications that frequently deploy ultra-large amounts of computing nodes for short periods of time. Examples include rendering, genetic analysis, and solid-state pharmaceutical research.

### Note:

This instance type is currently only available to whitelisted users. Contact your account manager if you wish to purchase these instance type.

## Batch-based Computing BC1

Batch Compute BC1 instances are cost-effective instances billed by core-hour, with billing accuracy down to the second and prices as low as 0.1 CNY/core-hour. They offer flexibility with immediate availability and termination upon completion, supporting various specifications to meet the demands of compute-intensive users requiring frequent short-term usage of large-scale compute nodes, such as rendering, gene analysis, and crystallography.

### Use Cases

- Video/Film rendering
- Genetic analysis, solid-state pharmaceutical research etc.
- HPC compute-intensive applications such as weather forecasting and astronomy

### Features

- Cost-effective with the lowest price per hour among all instances with the same specification
- CPU/RAM ratio: 1:4

### Requirements

- Batch-based computing BC1 instances are pay-as-you-go.
- BC1 instances can be launched in classic networks and VPCs.

- See below for available BC1 specifications:

Specifications	vCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
BC1.LARGE16	4	16	–	4	1.5	–	–
BC1.3XLARGE48	12	48	–	8	2.5	–	–
BC1.6XLARGE96	24	96	–	8	5.0	–	–

**Note:**

In the BC1 instance specification list, a "-" in the network packet transmission and CPU frequency columns indicates that there is no fixed commitment to network packet transmission and processor frequency performance for that specification. If you have a strong demand for consistent performance, it is recommended to choose a model with performance commitments.

## Batch-based Computing BS1

Batch computing BS1 instances are cost-effective instances billed by core-hour, with billing accuracy down to the second and prices as low as 0.09 CNY/core-hour. They offer flexibility with immediate availability and termination upon completion, supporting various specifications to meet the demands of compute-intensive users requiring frequent short-term usage of large-scale computing nodes, such as rendering, gene analysis, and crystallography.

### Features

- Cost-effective with the lowest price per hour among all instances with the same specification
- CPU/RAM ratio: 1:2

### Use Cases

- Video/Film rendering
- Genetic analysis, solid-state pharmaceutical research etc.
- HPC compute-intensive applications such as weather forecasting and astronomy

## Requirements

- Batch-based computing BS1 instances are pay-as-you-go.
- BS1 instances can be launched in classic networks and VPCs.
- See below for available BS1 specifications:

Specifications	VCPU	Memory (GB)	Throughput (pps) (received and sent packets)	Number of Queues	Private network bandwidth (Gbps) (outbound and inbound)	Clock Rate	Remarks
BS1.LARGE8	4	8	–	4	1.5	–	–
BS1.3XLARGE24	12	24	–	8	2.5	–	–
BS1.6XLARGE48	24	48	–	8	5.0	–	–

### Note:

In the BS1 Instance Specification list, a "-" in the Network Packets and CPU Frequency columns indicates that there is no fixed commitment to network packet transmission and processor frequency performance for that specification. If you have a strong demand for consistent performance, it is recommended to choose a model with performance commitments.

## Cloud Bare Metal

Tencent Cloud Bare Metal (CBM) is a bare metal cloud service that combines the elasticity of cloud servers with the performance of physical machines. It seamlessly integrates with Tencent Cloud's full range of products (such as networking, storage, and databases) to provide you with high-performance, securely isolated physical server clusters in the cloud. Your business applications can directly access the processor and memory of the bare metal cloud server without any virtualization overhead. With this service, you can elastically scale the number of physical servers based on your business needs, reducing the time to acquire physical servers to minutes. You can entrust capacity management and operations to Tencent Cloud's professional team, allowing you to focus on business innovation.

For instance specifications, please refer to [Bare Metal Server Instance Specifications](#).

## HPC Cluster

High-performance computing clusters use bare-metal cloud servers as nodes, interconnected via RDMA, providing high-bandwidth and ultra-low-latency network services. These clusters can meet the parallel computing requirements of large-scale high-performance computing, artificial intelligence, and big data recommendation applications.

For instance specifications, please refer to [High-Performance Computing Cluster Instance Specifications](#).

# Lifecycle

Last updated: 2025-08-05 19:14:29

The lifecycle of a Tencent Cloud CVM instance refers to the various stages it undergoes from initiation to termination. By effectively managing the instance throughout its lifecycle, you can ensure that applications running on it deliver services efficiently.

## Instance Status

- An instance has the following statuses:

Status	Status Attribute	Status Description
Creating	Interim status	The instance has been created but is not running yet.
Running	Stable status	The instance is in a normal operating state, and you can run your business on instances in this status.
Restarting	Interim status	A restart operation has been performed for the instance via the console or APIs, but the instance is not running yet. If this status lasts for a long time, there may be an exception.
Resetting	Interim status	After an instance undergoes a system reinstallation or disk reset operation via the console or APIs, it enters a state prior to becoming "running."
Shutting down	Interim status	After a shutdown operation has been performed on the instance via the console or APIs, the instance enters a pre-shutdown state. If this status persists for an extended period, it may indicate an exception. It is not recommended to force a shutdown.
Shut down	Stable status	The instance is stopped normally, and in the shutdown state, it cannot provide external services. Some attributes of the instance can only be modified when it is shut down.
Terminating	Interim status	The status when an instance has not yet been terminated after expiring for 15 days or when the user actively initiates the termination process.

To be repossessed	Stable status	Monthly subscribed instances within 15 days of expiration or pay-as-you-go instances within 2 hours of active termination are in the recycle bin status. This status does not provide external services.
Released	Stable status	Upon completion of the termination operation, the original instance no longer exists, cannot provide services, and all data has been completely erased.

## Launching an instance

- After initiating the instance launch operation, the instance enters the "Creating" status. During the "Creating" status, the instance's hardware specifications are configured according to the specified [instance type](#), and the system uses the image specified at startup to launch the instance.
- Upon completion of creation, the instance enters the "Running" status. Instances in the "Running" status have normal connections and access services enabled.

For additional information on instance startup, please refer to [Creating an Instance](#), [Logging In to a Windows Instance](#), and [Logging In to a Linux Instance](#).

## Restarting an instance

We recommend you restart an instance via the Tencent Cloud Console or Tencent Cloud APIs instead of running the OS restart command in the instance.

- Upon executing the restart operation, the instance will transition into the "Restarting" status.
- Restarting an instance is like restarting a computer. After restarted, the instance will retain its public IP address, private IP address and all data on its disk.
- Depending on its configuration, restarting the instance normally takes dozens of seconds to several minutes.

For more information about instance restarts, see [Restarting Instances](#).

## Shutting down an instance

Users can shut down instances using methods such as the Tencent Cloud console or Tencent Cloud API.

- Shutting down an instance is like shutting down a computer.
- A shutdown instance no longer provides external services, but the billing of the instance continues.
- A shutdown instance will still be displayed in the console.
- Shutdown is required for some configuration operations such as adjusting hardware configurations and resetting passwords.

- The shutdown operation does not change the CVM's public IP, private IP, or any data on its disk.

For more information about shutting down instances, please see [Shutting Down Instances](#).

## Terminating and releasing an instance

When a user no longer needs a CVM instance, they can terminate and release it. This can be achieved through the Tencent Cloud console or Tencent Cloud API.

- **Manual Termination:** For monthly subscribed instances that have not expired or pay-as-you-go instances without overdue payments, you can choose to manually terminate them. Monthly subscribed instances are released after being kept in the recycle bin for up to 15 days, while pay-as-you-go instances are released after being kept in the recycle bin for up to 2 hours.
- **Automatic termination due to expiration or overdue payment:** Monthly subscribed instances are retained in the recycle bin for up to 15 natural days after expiration and will be automatically released if not restored. Pay-as-you-go instances with a balance less than 0 for 2 hours + 15 days will be automatically released (fees will be deducted for the first 2 hours, and the instance will be shut down and stop charging for the next 15 days. Overdue pay-as-you-go instances do not enter the recycle bin, and you can view them in the instance list). Renew within the specified time to continue using the instance.

Upon instance termination, the system disk and specified data disks purchased with the instance will be released, but the elastic cloud disks mounted on it will remain unaffected. For more information on instance termination, please refer to [Terminate/Return Instance](#).

# Spot Instance

Last updated: 2024-06-02 15:49:41

## What is a spot instance?

Spot instances are a billing mode for cloud server CVMs, characterized by discounted pricing and a system interruption mechanism. You can purchase instances at a discount, but the system may automatically reclaim these discounted instances. Once you purchase and obtain a spot instance, its usage is virtually identical to that of a pay-as-you-go CVM instance, including console operations, remote login, service deployment, and VPC association.

- Related link: [Common Questions > Instance Category > Spot Instances](#)
- Related link: [How to Purchase Spot Instances](#)

## Special policies for the current stage

- **System-initiated Interruption (Inventory Fluctuation):** At the current stage, the system will not interrupt due to market price reasons, but only due to insufficient spot instance resource inventory. When the inventory is insufficient, the system will randomly reclaim allocated spot instances, and instance data will not be retained.
- **Available in All Regions:** Spot instances have been launched in most Tencent Cloud regions, supporting instance types in sync with the pay-as-you-go billing mode. For the latest regions and instance types, please refer to [Spot Instances – Supported Regions and Types](#).

## Product Features

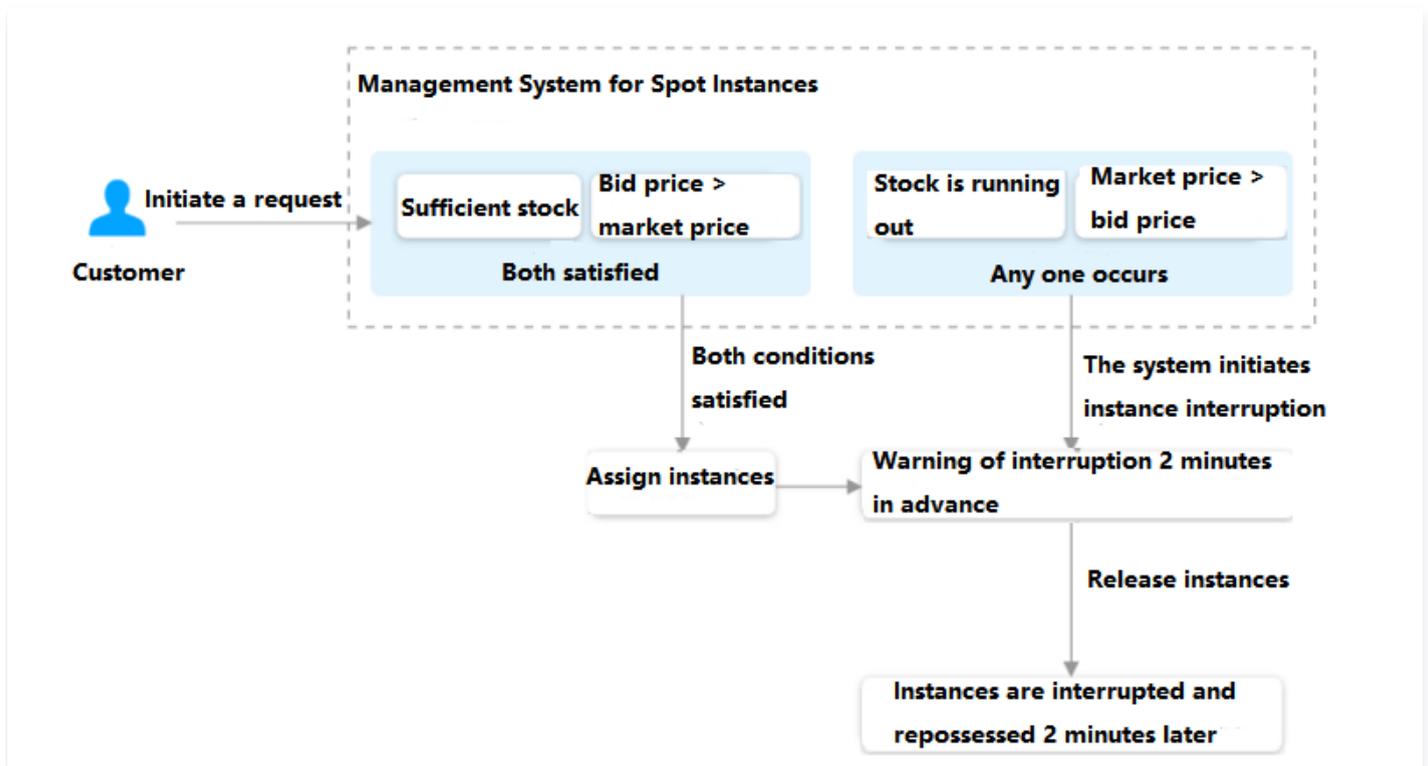
### Cost-effectiveness

CVM	Guangzhou Zone 3	Pay-as-you-go	Spot Instances
	Take S2.2XLARGE16 as an example Unit: USD/hour	Instances	(Take maximal discount as an example)
CPU MEM Items eligible for discount	CPU (8-core) MEM (16 GB)	0.38	0.076
System disk Data disk	System Disk (50 GB HDD disk)	0.01	0.01
Bandwidth/Traffic	Bandwidth (1 Mbps, bill-by-bandwidth)	0.06	0.06
Paid Image			
.....	Total Fee	0.45	0.146

Spot instances are sold at a discount of up to 95% off the prices of pay-as-you-go instances.

- **Discount Range:** The price of spot instances is based on the pay-as-you-go instances with the same specifications, and is sold at a discount ranging from 5% to the original price.
- **Discounted Components:** The discount only applies to the CPU and memory portions of the cloud server. Other components, such as system disks, data disks, bandwidth, and paid images, are not affected by the spot discount.
- **Price Fluctuation:** The discount rate remains stable over a period of time, but when large-scale purchasing occurs in an availability zone, the price may fluctuate.

## System interruption mechanism



Unlike pay-as-you-go instances which can only be released by users, spot instances may be interrupted by the system due to price or resource availability reasons.

**System-initiated Interruption (Inventory Fluctuation):** At the current stage, the system will not interrupt due to market price reasons, but only due to insufficient spot instance resource inventory. When the inventory is insufficient, the system will randomly reclaim allocated spot instances, and instance data will not be retained.

## Non-applicable scenarios

As spot instances may be interrupted, their lifecycle is not under your control. Therefore, it is not recommended to run services with high stability requirement on a spot instance. For example:

- Database services
- Online and website services without load balancers
- Core control nodes in a distributed architecture
- Prolonged big data computing job lasting over 10 hours

## Applicable scenarios and industries

### Scenarios

- Big data computing
- Online and website services with load balancers
- Web crawler service

- Other computing scenarios with fine granularity or support for checkpoint restart

## Applicable Industries

- Gene sequencing and analysis
- Drug crystal form analysis
- Video transcoding and rendering
- Financial and transaction data analysis
- Image and multimedia processing
- Science calculations, such as in geography and hydromechanics

## Restrictions

- **Quota Limits:** Unlike the number of cloud servers, the quota limit for spot instances is based on the total number of vCPU cores for all spot instances owned by a user in an availability zone. At this stage, each account can have up to 50 vCPU cores in total for spot instances in each availability zone. To increase the quota, please [submit a ticket](#).
- **Operation restriction 1:** You cannot upgrade and degrade the configuration of spot instances.
- **Operation restriction 2:** Spot instances cannot be converted to monthly subscriptions.
- **Operation restriction 3:** Spot instances do not support the "No Charge when Shut Down" feature.
- **Operation restriction 4:** System reinstallation is not supported for spot instances.
- **Operation restriction 5:** You cannot expand the system disks and data disks of spot instances.

## Best practices

### Splitting tasks

- Split a prolonged task into fine-grained subtasks for lower possibility of interruption.
- Utilize big data suites with a natural partitioning mindset, similar to EMR.

### Using load balancers to ensure the stability of online and website services

- Use load balancers, such as CLB, at the access layer.
- Use a combination of some pay-as-you-go instances and many spot instances for backend resources.
- Monitor the interruptions of spot instances and remove instances that are about to be interrupted from the CLB.

### Using a computing scheduling mode that supports checkpoint restart

- Store intermediate computing results on permanent storage products such as COS, CFS, and NAS.
- Be aware of the instance metadata to monitor which instances are about to be interrupted and save the computing results within the retention period of 2 minutes.
- Upon recreating a spot instance, continue the previous computation.

# Storage

## Storage Overview

Last updated: 2024-09-22 16:52:23

Tencent Cloud offers an array of flexible, cost-effective, and user-friendly data storage devices for cloud server instances. Each storage device boasts distinct performance and pricing, catering to various use cases.

### Storage Devices

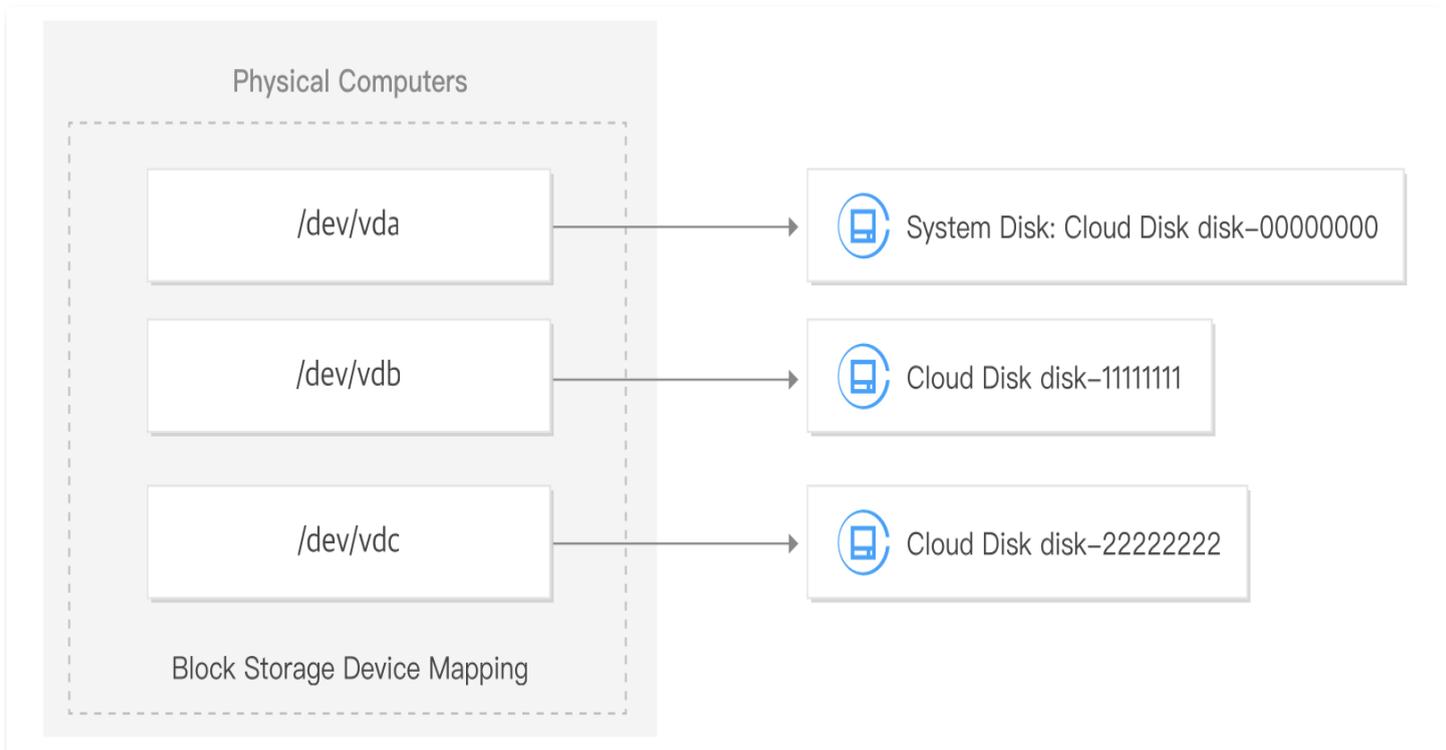
Storage devices can be divided into the following categories by dimension:

Dimension	Type	Note
Use Cases	System Disk	Used to store the collection of systems that control and schedule the operation of CVM. It uses images.
	Data Disk	Used to store all the user data.
Architecture	Cloud Disk	Cloud Block Storage (CBS) is an elastic, highly available, highly reliable, low-cost, and customizable network block device that can be used as a standalone scalable disk for CVM. It provides data storage at a data block level and employs a 3-copy distributed mechanism to ensure data reliability. <b>Cloud server instances with cloud disks can adjust hardware, disk, and network configurations.</b>
	Cloud Object Storage	Object storage is an internet-based data storage device that allows data retrieval from cloud server instances or any location on the internet, thus reducing storage costs. It is not suitable as a storage medium for low-latency, high I/O scenarios.

### Block storage device mapping

Each instance has a system disk to ensure basic operational data and can mount additional data disks. Instances use block storage device mapping (device-mapping) to map these storage devices to recognizable locations.

Block storage is a storage device divided into blocks in bytes and supports random access. Tencent Cloud supports two types of block storage devices: local disks and cloud disks.



This diagram illustrates how CBS maps block storage devices to cloud servers: mapping `/dev/vda` to the system disk and mapping two data disks to `/dev/vdb` and `/dev/vdc`, respectively.

Cloud server instances automatically create block storage device mappings for local disks and cloud disks mounted on them.

# Cloud Block Storage

Last updated: 2024-05-15 15:25:31

Tencent Cloud Block Storage (CBS) provides a persistent block storage service for CVM instances.

- CBS automatically stores data in multiple redundant copies in an availability zone to eliminate the risk of single points of data failure, providing up to 99.9999999% data reliability.
- CBS offers cloud disks of multiple types and specifications to achieve stable and low-latency storage performance.
- CBS can be mounted to or unmounted from instances in the same availability zone, and its storage capacity can be altered in a few minutes, meeting elastic storage requirements. Payment is based on actual usage, and allowing you to enjoy all the above features at a low price point.

## Typical Use Cases

- You can purchase and mount one or more cloud disks to meet the requirements for storage capacity when the disk space on your CVM is insufficient.
- You don't need to purchase extra storage capacity while purchasing a CVM. You can purchase cloud disks later when it's necessary.
- When you need to transfer data from one CVM to another, just detach the cloud data disk from the source CVM and attach it to the target CVM.
- Use multiple cloud disks to form a Logical Volume Manager (LVM) to go beyond the physical limit of a single cloud disk.
- Use multiple cloud disks to form a Redundant Array of Independent Disks (RAID) configuration to go beyond the I/O performance limit of a single cloud disk.

## Lifecycle

- **Non-elastic cloud disks** have a lifecycle that completely follows the CVM. They are purchased together with the CVM and used as system disks, without support for mounting or unmounting.
- The lifecycle of an **elastic cloud disk** is independent of the CVM instance and is not affected by its operation. You can attach multiple cloud disks to a single instance or detach a cloud disk from one instance and attach it to another as a data disk.

## Types

Tencent Cloud offers five types of cloud disks: **Premium Cloud Disk**, **SSD Cloud Disk**, **Balanced SSD Cloud Disk**, **Enhanced SSD Cloud Disk**, and **ulTra SSD Cloud Disk**. Each type has different

performance, features, and pricing. You can refer to the [Cloud Disk Types](#) and [Cloud Disk Pricing Overview](#) documents to choose the most suitable option based on your application requirements.

## Related Actions

- For information on CVM instance and cloud disk configurations, see [Creating Cloud Disks](#) and [Attaching Cloud Disks](#).
- For cloud disk expansion, unmounting, deletion, and other best practices, please refer to [Cloud Disk Product Documentation](#).

# COS

Last updated: 2024-03-26 17:37:02

Cloud Object Storage (COS) provided by Tencent Cloud is a distributed storage service designed to handle a massive amount of files. Users can store and access data at any time through the internet.

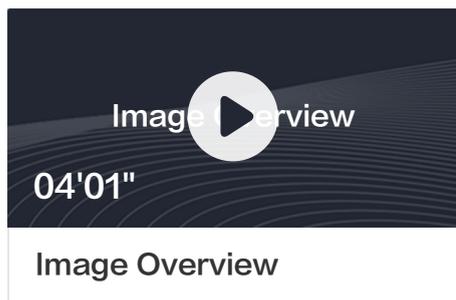
Cloud server users can store and retrieve data from instances or any location on the internet. COS stores user data redundantly across multiple regions, allowing simultaneous read or write operations by various client or application threads. Tencent Cloud COS offers a highly scalable, low-cost, reliable, and secure data storage solution for cloud server users.

For more information about Tencent Cloud Object Storage, please refer to the [COS Product Documentation](#).

# Image

## Image Overview

Last updated: 2024-09-24 09:48:11



## Images

A Tencent Cloud image provides all the information needed to launch a CVM instance. With an image, you can launch instances with similar configurations easily. Generally speaking, an image is the "installation disk" of a CVM instance.

## Image Type

Tencent Cloud provides the following types of images:

- **Public Images:** Available to all users, covering most mainstream operating systems.
- **Service Marketplace Images:** Available to all users, integrating specific applications in addition to the operating system.
- **Custom Images:** Accessible only to the creator and shared users, created from existing running instances or imported from external sources.
- **Shared Images:** Images shared by other users, can only be used for creating instances.

For more information, see [Image Types](#).

## Image Billing

For billing details of images, see [Billing Overview](#).

## Deployment with Images vs Manual Deployment

Model Item	Image Deployment	Manual Deployment
Deployment duration	3 to 5 minutes	1 to 2 days

Deployment process	Quickly create a suitable CVM based on mature marketplace solutions or the already used solutions.	Select the appropriate operating system, database, application software and plugins to create a CVM, and installation and debugging is required.
Security	Apart from shared images, which users need to evaluate themselves, public images, custom images, and marketplace images have all undergone testing and review by Tencent Cloud.	It depends on the development and deployment personnel.
Applicable scenarios	Public images: Genuine operating systems, which contain the initialized add-ons provided by Tencent Cloud. Marketplace images: Quickly build personalized application environments based on mature solutions. Custom images: Quickly create the same software environment as an existing CVM, or perform environment backup. Shared images: Quickly create the same software environment as other users' CVMs.	Configure a CVM completely by yourself, without basic settings provided.

## Use Cases

- **Deploying Specific Software Environments**

Using shared images, custom images, and marketplace images can help quickly set up specific software environments, eliminating the need for time-consuming tasks such as configuring environments and installing software. This meets various personalized needs, such as website building, application development, and visual management, making CVMs ready-to-use and time-saving.

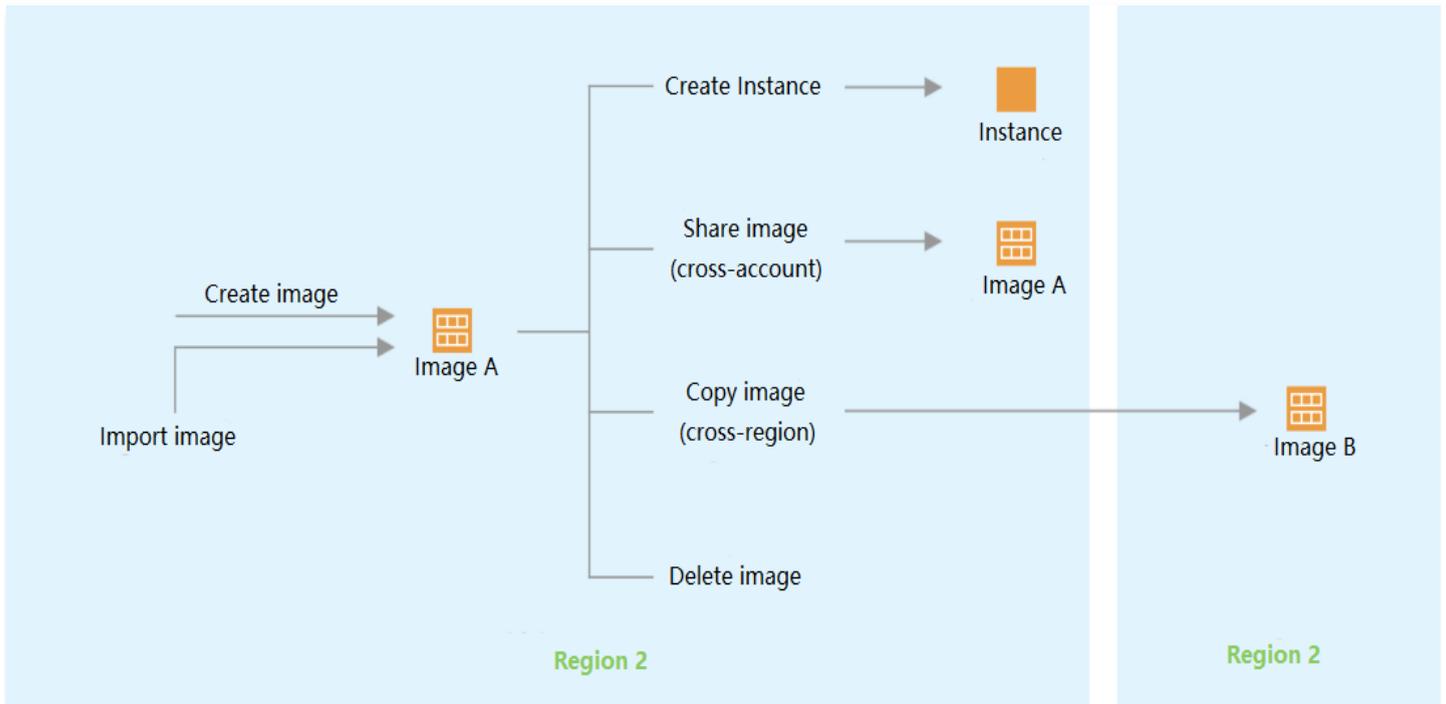
- **Batch Deployment of Software Environment:** Create an image from a CVM instance with a pre-configured environment, and use this image as the operating system when creating multiple CVM instances. This ensures that the new instances have the same software environment as the original instance, achieving the purpose of batch deployment.

- **Server Runtime Environment Backup**

Create an image backup of a CVM instance's runtime environment. If the instance becomes inoperable due to software environment damage, the image can be used for recovery.

## Custom Image Lifecycle

The diagram below summarizes the lifecycle of a custom image. After creating or importing a new custom image, users can use it to launch new instances (users can also launch instances from existing public images or marketplace images). Custom images can be replicated to other regions under the same account, becoming independent images in those regions. Users can also share custom images with other users.



# Image Types

Last updated: 2023-09-07 15:53:30

You can select an image based on the following attributes:

- Location (please refer to [Regions and Availability Zones](#))
- Operating system
- Architecture (32-bit or 64-bit)

Based on different sources, Tencent Cloud offers the following image types: Public images, Custom images, Shared images, and Service Marketplace images.

## Public image

**Public images** are provided by Tencent Cloud, containing basic operating systems and initialization components. They are available to all users. You can view the supported public image versions in the [Public Image List](#) documentation.

### Features:

- **Operating System:** Choose freely (e.g., based on Linux or Windows systems) and update regularly.
- **Software support:** Integrated with software packages provided by Tencent Cloud, and supports multiple versions of common software such as Java, MySQL, SQL Server, Python, Ruby, Tomcat, along with full permissions.
- **Security:** The provided operating systems are fully compliant and use official, genuine versions. Created by Tencent Cloud's internal Professional security operations team, they undergo rigorous testing and can optionally include built-in Tencent Cloud security components.
- **Service Support:**

Public image type	Service support
Tencent Cloud home-grown images	Tencent Cloud's self-developed image <a href="#">TencentOS Server</a> is a Linux operating system specifically designed for cloud scenarios, with technical support and maintenance provided by Tencent Cloud.
Third-party images	For technical support regarding third-party open-source images, please contact the open-source community or the original operating system manufacturer. Tencent Cloud will provide technical assistance in investigating the issue.

When using third-party commercial images on Tencent Cloud, the Tencent Cloud engineering team will assist you in resolving issues related to the operating system. If necessary, the Tencent Cloud team will contact the original manufacturer for help and support.

- **Official End-of-Life (EOL) Plans for Operating Systems**

The maintenance period for public images provided by Tencent Cloud will be consistent with the official maintenance period announced by the image platform. You can view the EOL plans for different operating system platforms and their respective image versions through [Official EOL Plans for Operating Systems](#).

- **Public Image Update Records**

For images that are not discontinued by the official image platform, Tencent Cloud will periodically update public images based on the updates, patches, and fixes released by the official image platform. You can view the image update status through [Public Image Update Records](#).

## Marketplace image

[Service Marketplace Images](#) are images provided by third-party service providers. After being reviewed by Tencent Cloud, they are released for use by all users.

### Service Marketplace Image Features:

- **Operating System:** Choose freely (e.g., based on Linux or Windows systems) and update regularly.
- **Software Support:** Integrated with software packages provided by Tencent Cloud, and supports multiple versions of common software such as Java, MySQL, SQL Server, Python, Ruby, and Tomcat.
- **Security:** The provided operating systems are fully compliant and use official, genuine versions. Created by Tencent Cloud's internal Professional security operations team, they undergo rigorous testing and can optionally include built-in Tencent Cloud security components.
- **Limitations:** Only available for use in Chinese mainland regions.

## Custom Image

**Custom images** are created by users through the image creation feature or imported using the image import feature. They can only be used by the creator and those with whom the image is shared.

### Features:

- **Use case:** Create an image from a CVM instance with a deployed application to quickly launch more instances with the same configuration.

- **Feature support:** Allows users to freely create, copy, share, and delete images.
- **Limitations:** A maximum of 50 custom images are supported in each region.

For more information on operations and limitations, please refer to [Create Custom Images](#), [Copy Custom Images](#), [Share Custom Images](#), [Unshare Custom Images](#), and [Import Custom Images](#).

## Shared Image

**Shared images** are custom images shared with the current user by other Tencent Cloud users through the image sharing feature. The shared images will be displayed in the same region as the original image of the shared user.

### Features:

- **Use case:** Assist other users in quickly creating CVM instances.
- **Functionality support:** Shared images can only be used for creating cloud servers and cannot be modified, copied, shared, or used for other operations.
- **Security:** Shared images are not reviewed by Tencent Cloud and may pose security risks. Therefore, it is strongly recommended not to accept images from unknown sources.
- **Limitations:** Each custom image can be shared with up to 50 Tencent Cloud users. Image sharing only supports sharing within the same region of the recipient's account.

For more operations and limitations, please refer to [Share Custom Images](#) and [Unshare Custom Images](#).

# Public Images

## Public Images List

Last updated: 2023-09-12 16:26:13

**Public images** are images officially provided by Tencent Cloud. Each image contains an operating system and initialization components provided by Tencent Cloud, and is available to all users.

- **Official Operating System End-of-Maintenance Plans**

Tencent Cloud's public image maintenance period aligns with the official maintenance period announced by the image platform. You can refer to the document [Official Operating System End-of-Maintenance Plans](#) to view the end-of-maintenance plans for various image versions across different operating system platforms.

- **Public Image Update Records**

For images that have not reached the end of their maintenance period, Tencent Cloud will periodically update public images based on the updates, patches, and fixes released by the official image platform. You can refer to the document [Public Image Update Records](#) to view the latest updates on images.

## Public Image Characteristics

- **Operating System Type:** Freely choose (e.g., based on Linux or Windows systems) and update regularly.
- **Software Support:** Integrated with software packages provided by Tencent Cloud, and supports multiple versions of common software such as Java, MySQL, SQL Server, Python, Ruby, Tomcat, and more.
- **Security:** The provided operating systems are fully legal and compliant, using official and genuine versions. Created by Tencent Cloud's internal Professional security operations team, these images undergo strict testing and can optionally include built-in Tencent Cloud security components.
- **Service support:**

Public image type	Service support
Tencent Cloud home -	Tencent Cloud's self-developed image <a href="#">TencentOS Server</a> is a Linux operating system specifically designed for cloud scenarios, with technical support and maintenance provided by Tencent Cloud.

grown images	
Third-party images	<p>For technical support regarding third-party open-source images, please contact the open-source community or the original operating system manufacturer. Tencent Cloud will provide technical assistance in investigating the issue.</p> <p>When using third-party commercial images on Tencent Cloud, the Tencent Cloud engineering team will assist you in resolving issues related to the operating system. If necessary, the Tencent Cloud team will contact the original manufacturer for help and support.</p>

## Tencent Cloud home-grown images

TencentOS Server is a self-developed Linux system image by Tencent Cloud.

Image tag	Image name	Image ID
TencentOS Server 3.1	TencentOS Server 3.1 for ARM64 (TK4)	img-9a0eg8mt
	TencentOS Server 3.1 (TK4)	img-eb30mz89
TencentOS Server 2.4	TencentOS Server 2.4 for ARM64 (TK4)	img-ieb3k16j
	TencentOS Server 2.4 (TK4)	img-9axl1k53
	TencentOS Server 2.4 (ARM64)	img-l7xr2ilb
	TencentOS Server 2.4	img-hdt9xxkt

## Third-party images

### Windows Server images

Image tag	Image name	Image ID
Windows Server 2022 Datacenter Edition	Windows Server 2022 Datacenter Edition 64-bit Chinese version	img-m07ny34j
	Windows Server 2022 Datacenter Edition 64-bit English version	img-9tzezztj
Windows Server 2019 Datacenter	Windows Server 2019 Datacenter Edition 64-bit Chinese version	img-bhvhr6pr

	Windows Server 2019 Datacenter Edition 64-bit English Version	img-mmy6qctz
Windows Server 2016 Datacenter	Windows Server 2016 Datacenter Edition 64-bit Chinese version	img-9id7emv7
	Windows Server 2016 Datacenter Edition 64-bit English Version	img-1eckhm4t
Windows Server 2012 R2 Datacenter	Windows Server 2012 R2 Datacenter Edition 64-bit Chinese version	img-29hl923v
	Windows Server 2012 R2 Datacenter Edition 64-bit English version	img-2tddq003

## Linux images

Tencent Cloud public images include a wide range of third-party Linux system images, such as OpenCloudOS, CentOS, Ubuntu, and commercial images like Red Hat Enterprise Linux.

### OpenCloudOS

Image tag	Image name	Image ID
OpenCloudOS Server 9	OpenCloudOS Server 9.0	img-j5e5hadz
OpenCloudOS Server 8	OpenCloudOS Server 8	img-028qly2h

### CentOS Stream

Image tag	Image name	Image ID
CentOS Stream 9	CentOS Stream 9 64-bit	img-9xqekomx
CentOS Stream 8	CentOS Stream 8 64-bit	img-8m9ugrip

### CentOS

Image tag	Image name	Image ID
CentOS 8	CentOS 8.5 64-bit	img-es95t8wj
	CentOS 8.4 64-bit	img-l5eqiljn
	CentOS 8.3 64-bit	img-5w4qozfr

	CentOS 8.2(arm64)	img-n74hgdx
	CentOS 8.2 64-bit	img-n7nyt2d7
	CentOS 8.0 64-bit	img-25szkc8t
CentOS 7	CentOS 7.9(arm64)	img-50v2hx0z
	CentOS 7.9, 64-bit	img-l8og963d
	CentOS 7.8 64-bit	img-3la7wgnt
	CentOS 7.7 64-bit	img-1u6l2i9l
	CentOS 7.6 64-bit	img-9qabwvbn
	CentOS 7.5 64-bit	img-oikl1tzv
	CentOS 7.4(arm64)	img-k4xgkxa5
	CentOS 7.4 64-bit	img-8toqc6s3
	CentOS 7.3 64-bit	img-dkwyg6sr
	CentOS 7.2 64-bit	img-31tjrtph
CentOS 6	CentOS 6.10 64-bit	img-fizif873
	CentOS 6.9 64-bit	img-jhhcsd4h
	CentOS 6.9 32-bit	img-060ov1xz
	CentOS 6.8 64-bit	img-6ns5om13

### Note

CentOS has officially discontinued the maintenance of the CentOS Linux project. Tencent Cloud provides you with alternatives, recommending the use of CentOS-compatible OpenCloudOS or TencentOS Server as replacements. For detailed information and migration guidance, please refer to [CentOS Discontinuation Notice and Recommended Solutions](#).

## Ubuntu

Image tag	Image name	Image ID
Ubuntu Server 22.04 LTS	Ubuntu Server 22.04 LTS 64-bit	img-487zeit5

Ubuntu Server 20.04 LTS	Ubuntu Server 20.04 LTS 64-bit	img-22trbn9x
	Ubuntu 20.04(arm64)	img-2eokmhf5
Ubuntu Server 18.04 LTS	Ubuntu Server 18.04 LTS 64-bit	img-pi0ii46r
	Ubuntu 18.04(arm64)	img-9eh8c1p1
Ubuntu Server 16.04 LTS	Ubuntu Server 16.04 LTS 64-bit	img-pyqx34y1
	Ubuntu Server 16.04 LTS 32-bit	img-8u6dn6p1
Ubuntu Server 14.04 LTS	Ubuntu Server 14.04 LTS 64-bit	img-3wnd9xpl
	Ubuntu Server 14.04 LTS 32-bit	img-qpxvpujt

## Debian

Image tag	Image name	Image ID
Debian 12	Debian 12.0 64-bit	img-7ag0z2jt
Debian 11	Debian 11.5(arm64)	img-mbt4qk91
	Debian 11.4 64-bit	img-btz2mndd
	Debian 11.1 64-bit	img-4cmp1f33
Debian 10	Debian 10.2 64-bit	img-h1yvfw1
	Debian 10.12 64-bit	img-7ay90qj7
	Debian 10.11 64-bit	img-qhtfjw1d
Debian 9	Debian 9.13 64-bit	img-5k0ys7jp
	Debian 9.0 64-bit	img-6rrx0ymd
Debian 8	Debian 8.2 64-bit	img-hi93l4ht
	Debian 8.2 32-bit	img-ez7jwngr
	Debian 8.11 64-bit	img-2lj11q1f

## Red Hat Enterprise Linux

Image tag	Image name	Image ID
-----------	------------	----------

Red Hat Enterprise Linux 9	Red Hat Enterprise Linux 9.0 64-bit	img-ae5vndnh
Red Hat Enterprise Linux 8	Red Hat Enterprise Linux 8.6 64-bit	img-kp3mv36j
	Red Hat Enterprise Linux 8.5 64-bit	img-r5xber0b
Red Hat Enterprise Linux 7	Red Hat Enterprise Linux 7.9, 64-bit	img-0qhxz7dl
	Red Hat Enterprise Linux 7.6 64-bit	img-92meq8tb
	Red Hat Enterprise Linux 7.2 64-bit	img-g7mtf785

### Note

- Red Hat Enterprise Linux image is now in beta. To join the beta, [submit a request](#).
- To use Red Hat Enterprise Linux image, select an instance type that is verified by Red Hat Enterprise Linux when you purchase the CVM. For more details, see [Red Hat Enterprise Linux Image Usage Guide](#).

## AlmaLinux

Image tag	Image name	Image ID
AlmaLinux 9.1	AlmaLinux 9.1 64-bit	img-2cyb0wu1
AlmaLinux 9.0	AlmaLinux 9.0 64-bit	img-f089mf4l
AlmaLinux 8.8	AlmaLinux 8.8 64-bit	img-3cvyk589
AlmaLinux 8.6	AlmaLinux 8.6 64-bit	img-jy2bb29p
AlmaLinux 8.5	AlmaLinux 8.5 64-bit	img-4ogcw28j

## CoreOS

Image tag	Image name	Image ID
CoreOS Container Linux	CoreOS 1745.5.0 64-bit	img-0hvei5hp

## Fedora

Image tag	Image name	Image ID
-----------	------------	----------

Fedora 37	Fedora 37 64-bit	img-d7j9x59z
Fedora 36	Fedora 36 64-bit	img-ge141oql

## FreeBSD

Image tag	Image name	Image ID
FreeBSD 13.1	FreeBSD 13.1 64-bit	img-ng3lehjp
FreeBSD 13.0	FreeBSD 13.0 64-bit	img-1lkqxofp
FreeBSD 12.3	FreeBSD 12.3 64-bit	img-j9m732cx
FreeBSD 12.2	FreeBSD 12.2 64-bit	img-pi37fg9j
FreeBSD 11.4	FreeBSD 11.4 64-bit	img-aif2u6pf
FreeBSD 11.1	FreeBSD 11.1 64-bit	img-c1y1axb9

## Rocky Linux

Image tag	Image name	Image ID
Rocky Linux 9.2	Rocky Linux 9.2 64-bit	img-no59bf11
Rocky Linux 9.1	Rocky Linux 9.1 64-bit	img-ev270u7p
Rocky Linux 9.0	Rocky Linux 9.0 64-bit	img-k1g1wwy9
Rocky Linux 8.6	Rocky Linux 8.6 64-bit	img-no575grb
Rocky Linux 8.5	Rocky Linux 8.5 64-bit	img-qd4bf0jb

## OpenSUSE

Image tag	Image name	Image ID
openSUSE Leap 15.4	openSUSE Leap 15.4	img-aaa4d8d1
openSUSE Leap 15.3	openSUSE Leap 15.3	img-1e4uwwol
openSUSE Leap 15.2	openSUSE Leap 15.2	img-i6u3kbtj

openSUSE Leap 15.1	openSUSE Leap 15.1	img-4orfj3l
openSUSE 42.3	openSUSE 42.3 64-bit	img-0ytr67o7

# Official Maintenance End Plans of Operating Systems

Last updated: 2023-09-12 16:26:38

Tencent Cloud's public image maintenance period aligns with the official maintenance period announced by the image platform. You can refer to this document for the end-of-maintenance plans for different operating system platforms and their respective image versions.

For images that have not been discontinued by the official platform, Tencent Cloud will periodically update the public images provided on the public cloud based on the updates, patches, and fixes released by the image platform. You can view the image update history in the document [Public Image Update Records](#).

## Note

After end of maintenance, any software maintenance and support including bug fixes and feature updates are unavailable. We recommend that you update your operating system soon or choose a more stable image version.

## Self-Developed Operating Systems of Tencent Cloud

### TencentOS Server

Version	Routine Maintenance Time	Extend Maintenance Time
TencentOS Server 3	2029-12-31	2032-12-31
TencentOS Server 2	2024-12-31	2032-12-31

## Third-Party Operating Systems

### Windows Server

Version	Mainstream Update End Date	Extension Update End Date
Windows Server 2022 Datacenter	2026-10-13	2031-10-14

Windows Server 2019 Datacenter	2024-01-09	2029-01-09
Windows Server 2016 Datacenter	2022-01-11	2027-01-12
Windows Server 2012 R2 Datacenter	2018-10-09	2023-10-10

For official maintenance periods, please refer to the [Microsoft Windows Server Official Website](#).

## OpenCloudOS

Version	EOL
OpenCloudOS 8	2029-12

## CentOS

Version	End of Update	End of Maintenance
CentOS Stream 9	2027-05-31	2027-05-31
CentOS Stream 8	2024-05-31	2024-05-31
CentOS 8	2021-12-31	2021-12-31
CentOS 7	2020-08-06	2024-06-30
CentOS 6	2017-05-10	2020-11-30

For official maintenance periods, please refer to the [CentOS Official Website](#).

### Note

CentOS has officially discontinued the maintenance of the CentOS Linux project. Tencent Cloud provides you with alternatives, recommending the use of OpenCloudOS or TencentOS Server, which are compatible with CentOS. For detailed information and migration guidance, please refer to [CentOS Discontinuation Notice and Recommended Solutions](#).

## Ubuntu

Version	Standard Support End Date	Extension Update End Date
---------	---------------------------	---------------------------

Ubuntu 22.04 LTS	2027-04	2032-04
Ubuntu 20.04 LTS	2025-04	2030-04
Ubuntu 18.04 LTS	2023-04	2028-04
Ubuntu 16.04 LTS	2021-04	2026-04

For official maintenance periods, please refer to the [Ubuntu Official Website](#).

## Debian

Version	Support End Date	EOL-Long-Term Support End Date	EOL-Extension Long-Term Support End Date
Debian 11	2024-07	2026-06	/
Debian 10	2022-07	2024-06	2029-06-30
Debian 9	2020-07-06	2022-06-30	2027-06-30

For official maintenance periods, please refer to the [Debian Official Website](#).

## Red Hat Enterprise Linux

Version	End of Full Support	End of Phase One Maintenance Support	End of Second-Stage Maintenance Support	End of Extended Lifecycle Support (ELS)
Red Hat Enterprise Linux 8	2024-05-31	N/A	2029-05-31	2032-05-31
Red Hat Enterprise Linux 7	2019-08-06	2020-08-06	2024-06-30	2028-06-30

For more information, please refer to the [Red Hat Official Website](#).

### Note

- Red Hat Enterprise Linux image is now in beta. To join the beta, [submit a request](#).

- To use Red Hat Enterprise Linux image, select an instance type that is verified by Red Hat Enterprise Linux when you purchase the CVM. For more details, see [Red Hat Enterprise Linux Image Usage Guide](#).

## AlmaLinux

Version	EOL
AlmaLinux 8.5	2031-11

For official maintenance periods, please refer to the [AlmaLinux official website](#).

## CoreOS

Version	EOL
CoreOS Container Linux	2020-05-26

## FreeBSD

Version	End of Maintenance
FreeBSD 13.1	2023-06-30
FreeBSD 13.0	2022-08-31
FreeBSD 12.3	2023-03-31
FreeBSD 12.2	2022-03-31

For the official maintenance period, please refer to the [FreeBSD official website](#).

## Rocky Linux

Version	EOL
Rocky Linux 9.0	2032-05-31
Rocky Linux 8.6	2029-05-31
Rocky Linux 8.5	2029-05-31

For official maintenance schedule, please refer to [Rocky Linux Official Website](#).

## OpenSUSE

Version	End of Maintenance
OpenSUSE Leap 15.4	2023-11
OpenSUSE Leap 15.3	2022-11
OpenSUSE Leap 15.2	2022-01-04
OpenSUSE Leap 15.1	2021-02-02

For official maintenance periods, please refer to the [OpenSUSE Official Website](#).

# TencentOS Server

Last updated: 2024-09-24 16:59:31

## Feature Overview

TencentOS Server is Tencent's Linux OS designed for cloud scenarios. With specific features and optimized performance, TencentOS Server provides a high-performance, secure, and reliable operating environment for applications in CVM instances. TencentOS Server is publicly available currently, and its user-mode environment is compatible with CentOS, so applications developed on CentOS can run directly on TencentOS Server.

For more information about TencentOS Server, see [TencentOS Server Introduction](#).

## Related Actions

You can start using TencentOS Server by referring to [Usage Instructions](#). If you already have a CentOS instance, you can also refer to [CentOS to TencentOS Migration Guide](#) to migrate to TencentOS Server.

## Services and Updates

Tencent Cloud provides more than five years of maintenance and updates for each major version of TencentOS Server, including regular image updates, introduction of new features and optimizations, timely security vulnerability fixes, and bug fixes. Existing servers can be upgraded from the YUM repository for bug fixes in time.

- For more information on TencentOS Server, follow Tencent Cloud Assistant on Weixin Mini Program.
- For more information on TencentOS Server release notes, see [TencentOS Server Image Update Log](#).

# OpenCloudOS

## OpenCloudOS Overview

Last updated: 2023-09-12 16:27:21

The OpenCloudOS open-source operating system community was jointly initiated by Tencent and its partners. It is a completely neutral, fully open, secure, stable, and high-performance Linux operating system and ecosystem. OpenCloudOS consolidates the advantages of multiple vendors in software and open-source ecosystems, inherits Tencent's over 10 years of technical accumulation in operating systems and kernel level, and has solid support in cloud-native, stability, performance, and hardware support, capable of equitably and comprehensively supporting all hardware platforms.

### Use Cases

OpenCloudOS works with the vast majority of production instances on the cloud, including Cloud Virtual Machine (CVM) and Cloud Bare Metal (CBM).

### OpenCloudOS Version

- OpenCloudOS 8.6 is the latest version of OpenCloudOS V8 released by the OpenCloudOS community. It offers CentOS 8-compatible basic libraries and user-mode components, and is equipped with the OpenCloudOS Kernel based on the community's 5.4 LTS. With kernel enhancements and extensive deployment on over 10 million nodes, OpenCloudOS delivers 70% improved stability and 50% better performance in specific scenarios than CentOS 8, providing users with a superior solution compared to CentOS 8.
- OpenCloudOS 9.0 is the latest version of OpenCloudOS V9 released by the OpenCloudOS community. OpenCloudOS V9 utilizes the latest Kernel 6.1 from the upstream community, offering multi-architecture and new hardware support, multi-core performance optimization, and enhanced hybrid isolation features. Simultaneously, features like MGLRU and Mapple Tree enhance memory management efficiency. It also provides comprehensive Cgroup V2 support, multi-architecture hot patch support, and detailed system tuning and adaptation.

### Using OpenCloudOS

- You can choose the public image and the corresponding version of OpenCloudOS when creating a new instance or reinstalling the operating system of an existing instance. For more information, please refer to [Create Instance](#) and [Reinstall System](#).
- If you already have a CentOS instance, you can refer to the [CentOS to OpenCloudOS Migration Guide](#) to migrate to OpenCloudOS.

## Obtaining OpenCloudOS

- Please visit [OpenCloudOS 8.6](#) to obtain OpenCloudOS V8.
- Please visit [OpenCloudOS 9.0](#) to obtain OpenCloudOS V9.

## Update History

For more information, see [OpenCloudOS Image Update Log](#).

## Services and Updates

The OpenCloudOS community provides up to a decade of maintenance and updates for each major version of OpenCloudOS (such as OpenCloudOS V8), encompassing the latest kernel features, security vulnerability remediation, and bug fixes. Existing servers can be upgraded via yum to promptly address vulnerabilities.

For further inquiries, we invite you to scan the QR code and follow the OpenCloudOS WeChat Official Account to join the official community discussion group.



# OpenCloudOS V8 Image Update Log

Last updated: 2023-09-12 16:27:48

## Note

- The image update history is organized by release time.
- Image releases follow a regional rollout process. If the image used when creating a CVM is not the latest version listed in the update history, it may not have been released in that region yet.
- If you cannot find a specific image mentioned in the update history within the console, it may not have been fully released yet. For more information, please consult the [Tencent Cloud Assistant](#).

## 2023

Updated Feature	Updated On
<ul style="list-style-type: none"> <li>• Resolved the issue with the <code>i8042.noaux</code> parameter</li> <li>• Resolved the issue of the keymap being empty</li> <li>• Resolved the issue of certain repo URLs lacking <code>mirrors.tencentyun.com</code></li> </ul>	2022-07-15
<ul style="list-style-type: none"> <li>• Upgraded the kernel to <code>v5.4.119-20.0009.20.oc8</code></li> <li>• User Mode Software Upgrade</li> <li>• Upgraded dracut-related software to rectify the abnormal startup of the <code>kdump</code> service</li> <li>• Launched OpenCloudOS 8.6 on the public cloud</li> </ul>	2023-03-15

## 2022

Updated Feature	Updated On
<ul style="list-style-type: none"> <li>• Disabled the <code>firewalld/sssdrngd</code> service</li> <li>• Uninstalled the <code>microcode_ctl/nss-softokn/avahi</code> package</li> <li>• Set keymap</li> <li>• Set timezone</li> <li>• Set the <code>cloudinit.target</code> dependency for the boot of <code>kdump</code></li> <li>• Configured <code>mirrors.tencentyun.com</code> as the first URL in repo</li> <li>• Modified the <code>/etc/rc.d/rc.local</code> file permission to <code>755</code></li> </ul>	2022-09-16

<ul style="list-style-type: none"><li>• Fixed permission errors of some directories in <code>/var/lib/</code></li></ul>	
<ul style="list-style-type: none"><li>• Updated the kernel version to 5.4.119-19.0010</li><li>• Updated other user mode software</li><li>• Updated the image timestamp</li></ul>	2022-07-27
<ul style="list-style-type: none"><li>• Launched OpenCloudOS 8.5 to the public cloud</li></ul>	2022-03-04

# Network and Security

## Overview

Last updated: 2024-03-26 14:28:42

Tencent Cloud offers network infrastructure and security services to ensure your business runs securely, efficiently, and flexibly.

### Encrypted Login

Tencent Cloud provides two encrypted login methods: [Password Login](#) and [SSH Key Login](#). Users can choose either method to securely connect to their cloud server. Please note that Windows instances do not support SSH Key Login.

### Network access

Tencent Cloud services can be accessed through [Internet Access](#) or [Private Network Access](#).

- **Internet Access:** Internet access is a service provided by Tencent Cloud for public data transfer between instances. Instances are assigned public IP addresses to enable communication with other computers on the network.
- **Private network access:** used to provide Local Area Network (LAN) service. Tencent Cloud assigns resources with private IP addresses to allow a free private network communication in the same region.

### Network Environment

Tencent Cloud's [network environments](#) can be divided into: Classic Network and Virtual Private Cloud (VPC).

- **Classic Network:** The Classic Network is a public network resource pool for all Tencent Cloud users. It is suitable for users who are new to Tencent Cloud and just getting started.
- **Private Network (VPC):** A Private Network is a logically isolated network space in Tencent Cloud that you can customize. Instances within a VPC can be launched in predefined, custom IP ranges, ensuring isolation from other users. This is suitable for users familiar with network management.

### Security Group

[Security Group](#) is a virtual firewall with stateful packet filtering capabilities, used to configure network access control for single or multiple cloud servers. It is an essential network security isolation measure provided by Tencent Cloud.

You can use the following methods to control access permissions for your instances:

- Create multiple security groups and specify different rules for them.

- Associate one or more security groups with each of your CVM instances. These security groups control traffic to your instances and the resources your instances can access.
- Configure your security groups to allow only the specified IP addresses or security groups to access your instances.

## Elastic IP (EIP)

**Elastic IP (EIP)** is a fixed public IP address that can be independently purchased and owned within a specific region. We recommend using Elastic IPs in the following scenarios:

- An instance may crash for unforeseeable reasons, and the failover instance needs to use the same IP address to provide uninterrupted service.
- An instance does not have a public IP address but still needs a static IP address.

## ENIs

The **Elastic Network Interface (ENI)** is a flexible network interface that can be bound to cloud servers within a private network and freely migrated between multiple servers. ENIs are highly beneficial for configuring and managing networks and building highly reliable network solutions.

## Cloud Workload Protection

**Host Security** leverages Tencent Security's extensive threat data and machine learning to provide users with security protection services such as hacker intrusion detection and vulnerability risk alerts. Key features include password cracking interception, remote login reminders, trojan file detection, and high-risk vulnerability detection. These capabilities address the primary network security risks faced by servers, helping businesses build a robust server security defense system and prevent data breaches.

## Anti-DDoS Basic

**Anti-DDoS Basic** is a free basic DDoS protection service provided by Tencent Cloud for resources such as cloud servers and load balancers, meeting daily security operation requirements. Tencent Cloud dynamically adjusts the blocking threshold based on the user's security reputation status. Anti-DDoS Basic is enabled by default, monitoring network traffic in real-time, initiating immediate cleansing upon detecting an attack, and providing second-level protection for public IPs on Tencent Cloud.

# Placement Group

Last updated: 2023-09-07 18:02:50

Placement groups are policies for distributing instances across underlying hardware. Instances created within a placement group provide disaster recovery and high availability upon startup. Tencent Cloud's CVM instance placement policies allow you to distribute instances using specific strategies during creation, reducing the impact of underlying hardware/software failures on your CVM-based business. By utilizing placement groups, you can deploy instances related to your business across different physical servers, ensuring high availability and underlying disaster recovery capabilities.

When creating instances within a placement group, we will distribute and launch instances in the specified region according to your pre-set deployment strategy. If you have not assigned a placement group for the instances, we will attempt to start instances on different physical machines to ensure service availability.

## Spread Placement Group

Currently, placement groups support spread groups. Spread placement groups distribute a set of instances across different underlying hardware, providing high availability. It is recommended for applications requiring individually placed critical instances, such as master-slave databases and highly available clusters. By launching instances in a spread placement group, the risk of instances with the same underlying hardware failing simultaneously is reduced.

Spread placement groups have regional attributes and can span multiple availability zones. There are quantity limitations for instances in each group; for related quantity restrictions, please refer to the [official website console](#) page.

### Note

If there is not enough hardware to fulfill the request, launching instances in a spread placement group will fail. In general, you can wait for some time and then retry your operation.

## Spread placement group rules and limits

Before using a spread placement group, note the following rules:

- Placement groups cannot be merged.
- An instance cannot be placed across multiple placement groups.
- You can choose from three spread placement levels: physical machines, switches, or racks. For example, if you select the physical machine level, instances in the spread group will not

be placed on the same physical machine, ensuring high availability for your business.

- The maximum number of instances supported by spread placement group on physical machine, switch, or rack is different. For more information, please see the official website.
- After implementing a disaster recovery group strategy, the specified policy will be strictly followed. It is important to note that if there is insufficient underlying hardware to distribute instances, some instances may fail to be created.
- Instances on CDH do not support spread placement groups.

## Operations Guide

For more related operations, please see [Spread Placement Groups](#) and [Managing Instance Placement Groups](#).

# Private Network Access

Last updated: 2024-05-15 15:25:32

Private network services, also known as Local Area Network (LAN) services, enable cloud services to communicate with each other via internal connections. Tencent Cloud services can be accessed through the [Internet](#) or interconnected via Tencent Cloud's private network. Tencent Cloud data centers are interconnected with high-speed, low-latency Gigabit and 10-Gigabit connections, providing high-bandwidth and low-latency communication services to help you flexibly build your network architecture.

## Private IP Address

### Overview

A private IP address cannot be accessed through the Internet and is the implementation method for Tencent Cloud's private network services. Each instance has a default network interface (eth0) assigned with a private IP address, which can be automatically allocated by Tencent Cloud or customized by the user (only in the [Private Network](#) environment).

#### Note

Changing the private IP address within the operating system may cause disruption to internal network communication.

### Properties

- Private network is user-sensitive, and different users are isolated from each other. By default, cloud services of another user cannot be accessed via the private network.
- Private network is region-sensitive, and different regions are isolated from each other. By default, cloud services under the same account in a different region cannot be accessed via the private network.

### Scenarios

Private IP can be used for the communication between CLBs and CVM instances, and between CVM instances and other Tencent Cloud services (such as TencentDB).

### Address Assignment

Each CVM instance is assigned a default private IP address upon startup. The private IP address varies depending on the [network environment](#):

- Basic network: private IP address is automatically assigned by Tencent Cloud and cannot be changed.

- **Private Network:** Currently, Tencent Cloud VPC supports CIDR blocks in any of the three major private IP ranges, as follows:
  - **10.0.0.0 – 10.255.255.255** (subnet mask range must be between 12 and 28)
  - **172.16.0.0 – 172.31.255.255** (subnet mask range must be between 12 and 28)
  - **192.168.0.0 – 192.168.255.255** (subnet mask range must be between 16 and 28)

## Private Network DNS

### DNS Server Address

The private network DNS service is responsible for domain name resolution. Incorrect DNS configuration may result in inaccessible domain names. Tencent Cloud provides reliable private network DNS servers in different regions. The specific configurations are as follows:

Network Environment	Regions	Private DNS Server
Classic network	Guangzhou	Guangzhou Zone 1: 10.112.65.31 10.112.65.32
		Guangzhou Zone 2: 10.112.65.31 10.112.65.32
		Guangzhou Zone 3: 10.59.218.193 10.59.218.194
		Guangzhou Zone 4: 100.121.190.140 100.121.190.141
	Shanghai	10.236.158.114 10.236.158.106
	Beijing	10.53.216.182 10.53.216.198
	Shanghai Finance	10.48.46.77 10.48.46.27
	Shenzhen Finance	100.83.224.91 100.83.224.88
North America	10.116.19.188 10.116.19.185	

	Hong Kong (China)	10.243.28.52 10.164.55.3
	Singapore	100.78.90.19 100.78.90.8
	Guangzhou Open	10.59.218.18 10.112.65.51
	Chengdu	100.88.222.14 100.88.222.16
	Silicon Valley	100.102.22.21 100.102.22.30
	Frankfurt	100.120.52.60 100.120.52.61
	Seoul	10.165.180.53 10.165.180.62
VPCs	All Regions	183.60.83.19 183.60.82.98

## Operations Guide

You can view or modify the private IP address of the instance. For detailed instructions, see:

- [Obtaining the instance's private IP address and configuring DNS](#)
- [Modify Instance Private IP Address](#)

# Internet Access

Last updated: 2024-03-26 16:07:21

When users need to deploy applications on CVM instances that require public services, data must be transmitted via the Internet, and an IP address on the Internet (also known as a public IP address) is required. Tencent Cloud provides Internet access through its high-speed data center connections. The domestic multi-line BGP network covers over twenty network operators, with BGP public network exports switching across domains in seconds, ensuring that your users can enjoy high-speed, secure network quality regardless of the network they use.

## Public IP address

- **Overview:** A public IP address is a non-reserved address on the Internet, allowing CVMs with public IP addresses to access and be accessed by other computers on the Internet.
- **Acquisition:** When creating a CVM, set the bandwidth greater than 0 Mbps in the network settings. Upon completion, the Tencent Cloud system will automatically assign a public IP address from the Tencent Cloud public IP address pool to the instance. This address can be changed. For more information on specific operations, please see [Changing Public IP](#).
- **Configuration:** You can log in to CVM instances with public IP addresses on the Internet and configure them accordingly. For more information on logging in to CVM instances, please see [Logging in to Linux Instances](#) and [Logging in to Windows Instances](#).
- **Conversion:** Public IP addresses are mapped to the instance's [private IP address](#) through Network Address Translation (NAT).
- **Maintenance:** All public network interfaces of Tencent Cloud are uniformly processed by Tencent Gateway (TGW). The public network cards of Tencent Cloud CVM instances are configured on the unified interface layer TGW, making them transparent to the CVMs. Therefore, when users view network interface information in CVMs using `ifconfig (Linux)` or `ipconfig (Windows)` commands, they can only see [private network](#) information. Public network information needs to be viewed by users themselves by logging in to the [CVM Console](#) and checking the CVM list/details page.
- **Cost:** Providing services through public IP addresses incurs associated fees. For specific details, please refer to [Public Network Billing Modes](#).

## Public IP address release

Users cannot actively associate or release public IP addresses associated with instances. Public IP addresses will be released or reassigned in the following situations:

- **When terminating an instance:** If a user actively terminates a pay-as-you-go instance or a monthly/yearly subscribed instance after its expiration, Tencent Cloud will release its public

IP address.

- [Elastic Public IP Address](#) **When associating and disassociating with instances.** When an instance is associated with an elastic public IP address, Tencent Cloud will release the original public IP address of the instance. When disassociating an instance from an elastic IP address, the instance will be automatically assigned a new public IP address, and the original released public IP address will return to the public IP address pool, making it unavailable for reuse.

If you need a fixed, permanent public IP address, you can use an [Elastic Public IP Address](#).

## Operations Guide

For detailed instructions on how to obtain and change your public IP address, refer to:

- [Obtaining the Public IP Address of an Instance](#)
- [Changing Instance Public IP Address](#)

# Elastic Public IP

Last updated: 2024-03-26 16:07:21

## Feature Overview

Elastic IP (EIP) is a fixed public IP address that can be independently purchased and owned within a specific region. With an elastic public IP address, you can quickly remap the address to another instance or NAT gateway instance in your account, effectively masking instance failures.

An Elastic IP can be retained in your account until it is released. Unlike public IPs, which can only be applied for and released with a cloud server, Elastic IPs can be decoupled from the cloud server's lifecycle and operated as independent cloud resources. For example, if you need to retain a public IP related to your business, you can convert it to an Elastic IP and keep it in your account.

## Rules and Limits

Rules/Restrictions	Notes
Elastic IP Usage Rules	For more details, refer to <a href="#">Usage Rules</a> .
Elastic IP quota restrictions	For more information, please refer to <a href="#">EIP Quota Limits</a> .
Limits on public IPs bound to CVM	For more information, please refer to <a href="#">Binding CVM Limits</a> .

## Difference Between an EIP and a Common Public IP

Public IP addresses are non-reserved addresses on the Internet, allowing cloud servers with public IPs to access and be accessed by other computers on the Internet. Both regular public IPs and EIPs are public IP addresses, providing cloud resources with the ability to access the public network and be accessed by it.

- Ordinary public IP: It can only be allocated when purchasing a cloud server and cannot be unbound from the server. If not allocated during purchase, it cannot be obtained.
- EIP: An independently purchasable and owned public IP address resource that can be bound or unbound to cloud resources such as cloud servers, NAT gateways, elastic network interfaces, and highly available virtual IPs at any time.

### Note

Only general BGP IP lines are applicable for the current common public IP addresses.

Compared to ordinary public IPs, EIP offers more flexible management options, as shown in the table below. For more information, please refer to [Public IPv4 Addresses](#).

Comparison Item	Common public IP address	EIP
Public network access	✓	✓
Can be purchased and held	×	✓
Can be bound and unbound at any time	×	✓
Adjust bandwidth in real time <sup>1</sup>	✓	✓
IP resource occupation fee	×	✓

**Note**

[Public IP Console](#) only supports adjusting the bandwidth of EIPs. For specific operations, please refer to [Adjusting Network Configuration](#). For adjusting the bandwidth of regular public IPs, please refer to [Adjusting Regular Public IP Network Configuration](#).

EIP can be decoupled from the lifecycle of cloud resources and operated independently. For example, if you need to retain a public IP address related to your business, you can convert a regular public IP to an EIP and keep it in your account.

# Login Password

Last updated: 2024-03-26 14:28:42

The password is a unique login credential for each cloud server instance. To ensure the security and reliability of instances, Tencent Cloud offers two encrypted login methods:

- Password login
- [SSH key](#) login

When creating the CVM, you can refer to the following documents to choose different encrypted login methods based on your CVM operating system.

- [Custom Configuration for Windows CVM](#)
- [Custom Configuration for Linux CVM](#)

Anyone with the instance login password can remotely access the cloud server instance through a public network address allowed by the security group. We recommend using a secure password, properly storing it, and periodically changing it. For more password-related operations, please refer to [Managing Login Passwords](#).

# SSH Keys

Last updated: 2024-06-02 10:16:51

To ensure the security and reliability of instances, Tencent Cloud offers two encrypted login methods: [Password Login](#) and SSH Key Pair Login. This document provides information on the configuration of SSH Key Pair Login.

When [customizing the configuration of a Linux cloud server](#), you can choose SSH Key as the encrypted login method for the server.

## SSH Key Overview

Tencent Cloud recommends that you **use SSH Key Pairs** to log in to Linux instances. SSH Key Pairs are a pair of keys generated through an encryption algorithm. Tencent Cloud-created SSH Key Pairs use RSA 2048-bit encryption, generating both a public key and a private key:

- **Public Key:** After the SSH Key Pair is successfully generated, Tencent Cloud only stores the public key. For Linux instances, the public key content is stored in the `~/.ssh/authorized_keys` file.
- **Private Key:** You need to download and securely store the private key, which is only available for download once. Tencent Cloud will not save your private key. Anyone with your private key can decrypt your login information, so you must keep the private key in a secure location.

You can securely connect to cloud servers using key pairs, which provide a more secure login method than using regular passwords. To use a key pair to log in to a Linux instance, simply specify the key pair when creating the instance or bind the key pair after the instance is created. This allows you to log in using the private key without entering a password.

## Features and Advantages

Compared to traditional password authentication methods, SSH key pair login has the following advantages:

- SSH key pair login is more complex and difficult to brute-force.
- SSH key pair login is easier to use. You can log in to instances remotely with a few simple configuration steps on the console and your local client, and do not need to enter a password when you log in again.

## Usage Limits

- SSH key pair login is only available for Linux instances.
- Each Tencent Cloud account can have up to 100 SSH key pairs.
- Tencent Cloud will not retain your private key. You need to download the private key after creating an SSH key, and keep it safe.

- To ensure data security, you need to shut the instance down before loading the key.
- To enhance the security of cloud servers, password login will be disabled by default once an instance is bound to a key. If you need to use password login simultaneously, please go to the Cloud Server Console to [reset the instance password](#).

## Use Cases

- For information on how to create, bind/unbind, and delete keys, please refer to [Managing SSH Keys](#).
- To learn about how to log in to CVM instances remotely using an SSH key pair, see:
  - [Logging In to a Linux Instance via a Remote Login Tool](#)
  - [Logging In to a Linux Instance via an SSH Key](#)

# Cloud Workload Protection Platform

Last updated: 2024-06-02 14:20:41

## Feature Overview

Based on Tencent Security's massive amounts of threat data, Tencent Cloud Workload Protection (CWP) leverages machine learning to provide a wide variety of security services ranging from intrusion detection to vulnerability alerting. It offers various security features such as brute force attack prevention, unusual login location reminding, trojan protection, and high-risk vulnerability detection, helping build a security protection system to cope with major network security risks faced by servers and prevent data leakage.

CWP is available in Basic and Pro editions. When creating a CVM instance, you can choose to activate CWP Basic by default.

### Note

For an introduction and comparison of CWP Basic and Pro features, please refer to [Feature Overview and Edition Comparison](#).

## Billing Mode

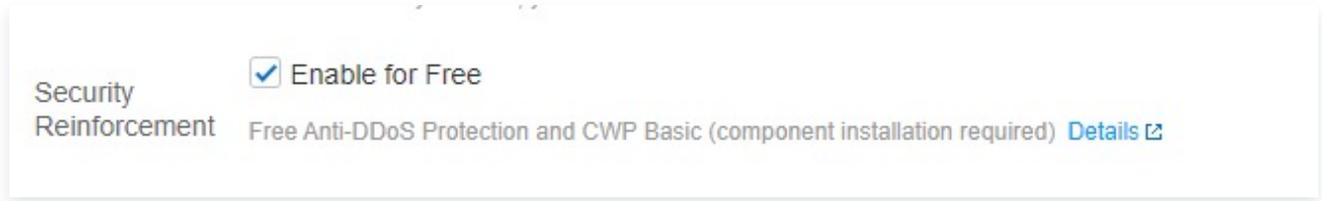
CWP Basic is free of charge. If you wish to upgrade to CWP Pro, please refer to [Purchasing CWP Pro](#).

## Installing CWP Basic

You can install CWP Basic in one of the following methods based on the actual conditions:

### Automatically Installing CWP during CVM Creation

When creating a Tencent Cloud CVM instance, you can choose to activate CWP Basic by default. On the CVM instance purchase page, under "Security Services," select "Enable for Free" to automatically install CWP, as shown in the image below:



Security Reinforcement  Enable for Free  
Free Anti-DDoS Protection and CWP Basic (component installation required) [Details](#)

### Manually Installing CWP on Existing CVMs

You can use one of the following methods to install the CWP agent for an existing instance based on its operating system:

- [Windows CVM Environment](#)
- [Linux CVM Environment](#)

After successful installation, you can view the security status of your CVM instances on the [CVM Console Overview page](#) or the [CWP Console](#).

## Documentation

- CWP Basic and Pro: [Feature Overview and Edition Comparison](#)
- [Security Dashboard](#)

# Anti-DDoS Basic

Last updated: 2024-05-15 17:08:01

## Feature Overview

Anti-DDoS Basic is a free service provided by Tencent Cloud for resources such as cloud servers and load balancers, offering fundamental DDoS protection to meet daily security operation requirements. Tencent Cloud dynamically adjusts the blocking threshold based on the user's security reputation status. Anti-DDoS Basic is enabled by default, monitoring network traffic in real-time, initiating immediate cleansing upon detecting an attack, and providing second-level protection for public IPs on Tencent Cloud.

## Billing Mode

Anti-DDoS Basic is a free service. If you require other DDoS protection solutions, please refer to [DDoS Protection Solution Comparison](#).

## Enabling Anti-DDoS Basic

When creating a Tencent Cloud CVM instance, you can choose to activate Anti-DDoS Basic for free by default. On the CVM instance purchase page, under "Security Services," select "Activate for free" to enable Anti-DDoS Basic.

After activation, you can view the protection configuration of the CVM instance on the [Cloud Server Console Overview](#) page or the [Anti-DDoS Basic](#) page in the Anti-DDoS Console.

## Documentation

Anti-DDoS Basic [Operation Overview](#)

# Ops and Monitoring

## Health Check

Last updated: 2023-09-20 10:54:36

### Overview

Instance Self-Check allows you to monitor the performance, cost, network, and disk status of your CVM instances, helping you understand their operational status. This feature enables you to promptly identify and resolve any issues related to your instances.

### Use Cases

We recommend you use health check in the following two scenarios:

- **Troubleshooting:** if any failure or problem occurs during instance operations, you can use self-service instance detection to locate and troubleshoot it and handle it according to the provided suggestions.
- **All-Around instance check:** during daily Ops, self-service instance detection can help you keep up to date with the overall instance running status and discover and solve problems promptly, guaranteeing normal business operations.

### Check Item Description

The health check items are as detailed below:

#### Local Network Detection

Category	Detection Instructions	Risk level	Solution
Network latency	<p>Local network latency refers to the PING value between your computer and the Tencent Cloud server (the time it takes for your computer to send data to the server and receive feedback), which is used to check the transmission delay between networks.</p> <ul style="list-style-type: none"> <li>• Detecting high network latency in instances by sending HTTP requests. The criteria are as follows:</li> <li>• If the latency is above 600 ms, the network quality is considered poor.</li> </ul>	Abnormal	We recommend that you check your local network, address specific issues accordingly, or switch to another network.

	<ul style="list-style-type: none"> <li>If no response is received within 5s, the request is considered timed out.</li> <li>If all requests time out, the network is considered disconnected.</li> </ul>		
Network jitter	Calculate the difference in latency values between adjacent requests, and the average of these differences represents the network jitter value. If the network jitter value divided by the network latency value is less than or equal to 0.15, it indicates a stable network. If it is greater than 0.15, it signifies network fluctuations.	–	
Upstream bandwidth	Data packets are uploaded to the instance to calculate its upstream bandwidth.	–	
Downstream bandwidth	Data packets are downloaded from the instance to calculate its downstream bandwidth.	–	

### Security Group Rule Detection

Category	Detection Instructions	Risk level	Solution
Common ports	Checks whether requests to common ports such as ports 22 and 3389 used by the TCP protocol for the inbound traffic are blocked in the security group.	Warnings	Inbound (Ingress) rules for TCP port 22 in the instance security group are blocked, which may prevent normal SSH login. You can open the required ports. For more information, please see

[Security Group Use Cases](#).

## Account Cost Detection

Category	Detection Instructions	Risk level	Solution
Cloud disk status	Checks whether cloud disks associated with the instance have expired and whether they can be read/written.	Abnormal	The cloud disk of this instance has expired. Please visit the <a href="#">Cloud Disk Console</a> to renew it as soon as possible.
	For pay-as-you-go instances and non-auto-renewal monthly subscription cloud disks, check if the cloud disk has expired and become unusable.	Warnings	The cloud disk of this instance is not set to auto-renew, which may result in its expiration and unavailability. It is recommended to visit the <a href="#">Cloud Disk Console</a> to enable auto-renewal for the cloud disk.
	For monthly-subscription instances with auto-renewal enabled and monthly-subscription cloud disks without auto-renewal, the cloud disk may become unusable due to expiration.	Warnings	
	Non-auto-renewal monthly-subscription instances and non-auto-renewal monthly-subscription cloud disks may become unusable due to inconsistent expiration dates between the two, causing the cloud disk to expire.	Warnings	The expiry time of the instance and its attached cloud disk are not consistent,

which may result in the cloud disk becoming unavailable due to expiration. It is recommended to visit the [Cloud Disk Console](#) to set up auto-renewal for the cloud disk.

### Instance Storage Check

Category	Detection Instructions	Risk level	Solution
High cloud disk latency	Checks whether the I/O performance metric svctm is abnormal.	Warnings	A cloud disk associated with the instance has a high latency. We recommend you pay attention to the cloud disk usage.
Checks for cloud disk I/O hang	Checks for cloud disk I/O hang	Warnings	A cloud disk associated with the instance has an I/O hang. We recommend you pay attention to the cloud disk usage.
System disk inode utilization	Checks whether the inode utilization of the cloud disk has reached 100%.	Warnings	Please pay attention to the usage of cloud disks. For troubleshooting,

System disk read-only	Checks whether the cloud disk is read-only.	Abnormal	refer to <a href="#">Kernel and I/O Related Issues</a> .
System disk space utilization	Checks whether the utilization of the cloud disk has reached 100%.	Warnings	
Partition I/O utilization	Checks whether the io_util of the cloud disk has reached 100%.	Warnings	

### Instance Network Detection

Category	Detection Instructions	Risk level	Solution
Public EIP connection	Checks whether the EIP is isolated due to overdue payments.	Abnormal	Public IP connectivity may be disrupted due to overdue payments. We recommend you visit the <a href="#">Billing Center</a> to top up and renew your account as soon as possible.
Existence of EIP	Checks whether the instance has an EIP.	Warnings	This instance does not have a public IP. If you need a public IP for external network access, please go to the <a href="#">EIP Console</a> to bind an EIP.
EIP blocked	EIP blocked	Abnormal	The public IP of this instance has been blocked due to a DDoS attack. Please refer to the <a href="#">Unblocking Protected IP</a> documentation for further guidance.

Public network bandwidth utilization	Checks whether the instance has experienced a high public network inbound bandwidth utilization in the last 12 hours.	Warnings	<p>To prevent becoming a bottleneck for your business, it is recommended to monitor network usage. For troubleshooting, please refer to <a href="#">High Bandwidth Usage Preventing Login</a>.</p>	
	Checks whether the instance has experienced a high public network outbound bandwidth utilization in the last 12 hours.	Warnings		
Private network bandwidth utilization	Checks whether the instance has experienced a high private network inbound bandwidth utilization in the last 12 hours.	Warnings		
	Checks whether the instance has experienced a high private network outbound bandwidth utilization in the last 12 hours.	Warnings		
Packet loss	Checks whether the instance has experienced TCP packet loss due to triggering of traffic throttling in the last 12 hours.	Warnings		<p>To prevent bottlenecks, it is recommended to check the health of your services. For more information, please refer to <a href="#">Cloud Server Network Packet Loss</a>.</p>
	Checks whether the instance has experienced UDP packet loss due to triggering of traffic throttling in the last 12 hours.	Warnings		

	Checks whether the instance has experienced packet loss due to a soft interrupt in the last 12 hours.	Warnings
Kernel network conditions	Checks whether the instance has experienced a full UDP send buffer in the last 12 hours.	Warnings
	Checks whether the instance has experienced a full UDP receive buffer in the last 12 hours.	Warnings
	Checks whether the instance has experienced a full TCP complete connection queue in the last 12 hours.	Warnings
	Checks whether the instance has experienced TCP request overflow in the last 12 hours.	Warnings
Connection utilization	Checks whether the number of connections of the instance has reached the upper limit in the last 12 hours.	Warnings

**Internal Detection for Linux Hosts**

Category	Detection Instructions	Risk level	Solution
----------	------------------------	------------	----------

SSH Login Related	Is password-based login disabled for SSHD?	In the <code>/etc/ssh/sshd_config</code> file, is <code>ssh_password_authentication</code> set to <code>yes</code> ?	Warnings	The <code>sshd</code> configuration has disabled password login. If you need to enable password-based login, please refer to the <a href="#">Procedure</a> .
	Whether SSHD prohibits root user login	In the <code>/etc/ssh/sshd_config</code> file, check if <code>ssh_permit_root_login</code> is set to <code>"yes"</code> .	Warnings	SSHD has disabled root user login. To enable root user login, please refer to the <a href="#">Handling Steps</a> .
	SSH Private Key File Permissions	Is the file permission for <code>/etc/ssh/ssh_host_rsa_key</code> incorrect?	Warnings	The <code>ssh_host_rsa_key</code> permission configuration is incorrect. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.
	<code>/var/empty/sshd</code> Permissions	Is the permission for <code>/var/empty/sshd</code> incorrect?	Critical	The <code>/var/empty/sshd</code> permission configuration is incorrect, causing login issues. Please refer to the <a href="#">Troubleshooting Steps</a> for resolution.
	<code>hosts_deny</code> Configuration	Whether the <code>/etc/hosts.deny</code> configuration file contains special	Warnings	The <code>/etc/hosts.deny</code> configuration file contains special restriction rules that may prevent login. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.

		login rules		
	Root User Shell Configuration	Whether the shell configuration in the /etc/passwd file is incorrect	Critical	The root user's bash configuration in the /etc/passwd file is incorrect, causing login failure. Please refer to <a href="#">Steps</a> for repair.
	wtmp or btmp Files	Checking if /var/log/wtmp or /var/log/btmp files are oversized	Warnings	Oversized /var/log/wtmp or /var/log/btmp files may cause slow login. You need to clear the corresponding files. Please refer to <a href="#">Unresponsive VNC Login After Entering Correct Password</a> for the repair process.
	Dynamic libraries required by the SSHD processes	Does the dynamic library required by the SSHD process exist?	Critical	The dynamic library required by the SSHD process is missing, causing login issues. Please refer to <a href="#">troubleshooting steps</a> for resolution.
	/etc/profile invoking /etc/profile	Whether there is a /etc/profile calling /etc/profile infinite loop	Critical	A dead loop caused by /etc/profile invoking itself prevents login. Please refer to <a href="#">Fixing the /etc/profile Dead Loop Invocation Issue</a> for a solution.
	SSHD Processes	SSH process existence	Critical	The sshd process is missing; you need to start the corresponding sshd service. Please refer to the <a href="#">Troubleshooting Steps</a> for resolution.
Network Configur	Number of ENI	Whether all ENI	Warning	If the number of ENI queues is not fully enabled, the instance may

ation	Queue s	queues have not been fully enabled	s	not achieve its maximum network performance. Please refer to <a href="#">Incorrect ENI Multi-Queue Configuration</a> for a solution.
	NAT Enviro nment Kernel Param eters	Is packet loss in the NAT environm ent caused by irregular kernel network configura tions?	War ning s	If <code>tcp_tw_recycle</code> is configured, packet loss may occur in a NAT environment. Please refer to <a href="#">Common Kernel Parameters for Linux Instances</a> for temporary removal.
OS Environ ment	Syste m Limits Config uration	Is the <code>/etc/secur ity/limits.c onf</code> configura tion abnormal ?	War ning s	An abnormal configuration in <code>/etc/security/limits.conf</code> may cause login failures. Please refer to <a href="#">Steps</a> for repair instructions.
	Syste m OOM	Has the system recently experienc ed an OOM event?	Criti cal	If the system has experienced OOM, it is recommended to evaluate whether the memory usage is reasonable or to upgrade the instance configuration. Please refer to <a href="#">High Memory Utilization</a> for troubleshooting and resolution.
	Is SELinu x enable d?	Check if SELinux is enabled on the instance.	War ning s	Enabling SELinux may cause login issues. We recommend referring to <a href="#">Handling Steps</a> to disable SELinux first.
	Has the PID been exhaus ted?	Instance PID exhaustio n check	Criti cal	System PID is nearing exhaustion, which may lead to system anomalies. It is recommended to evaluate whether the system's launched threads are reasonable

				or increase the system's pid_max. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.
	Cloud-init Environment	Checking if the Cloud-Init environment is functioning properly	Critical	Cloud-init environment anomalies may cause issues such as inability to reset passwords or modify hostnames. We recommend reinstalling cloud-init by referring to <a href="#">Installing cloud-init on Linux systems</a> .
	Basic information of a file system	/etc/fstab Configuration and File System Check	Warnings	/etc/fstab contains a non-existent partition, which may cause the instance to fail to start upon reboot. Please refer to <a href="#">Cloud Disk Not Auto-attaching upon Linux CVM Restart</a> for troubleshooting and resolution.
	Entering Emergency Mode	Checking for abnormalities in /etc/fstab configurations	Critical	Abnormal configurations in /etc/fstab may cause the system to enter emergency mode. Please refer to <a href="#">Unable to Log In Due to /etc/fstab Configuration Errors</a> for troubleshooting and resolution.
	Firewall Detection	Checking for abnormal firewall rules	Warnings	If the iptables policy is set to drop, it may cause network connectivity issues. Please refer to <a href="#">iptables policy settings</a> for troubleshooting and resolution.
Launch Configuration-related	bin and lib soft links	Checking for missing soft links in bin and lib directories	Critical	Missing bin or lib soft links may cause system anomalies. Please refer to <a href="#">System bin or lib Soft Link Missing</a> to recreate the corresponding soft links.
	Configuring Huge	Has huge pages memory	Warnings	The system has enabled huge pages memory, which may cause system anomalies. Please refer to

	Pages Memory	been enabled?		<a href="#">Configuring Huge Pages in sysctl.conf</a> to check if it is configured by your application. If not, the instance may have been compromised.
	Dynamic Library Hijacking Configuration	Has dynamic library hijacking been configured?	Warnings	The system is configured with dynamic library hijacking, which may cause system abnormalities. Please refer to <a href="#">ld.so.preload for adding dynamic library hijacking</a> to check if it is configured by the business program. If not, the instance may have been compromised.
System Resource Usage	Is the CPU utilization too high?	Are there any processes with CPU utilization exceeding 80%?	Warnings	High CPU utilization: assess whether it's reasonable or consider upgrading the instance configuration. Please refer to <a href="#">Troubleshooting and resolving login issues due to high CPU or memory usage</a> for guidance.
	Is the memory utilization too high?	Are there any processes with memory utilization exceeding 80%?	Warnings	High memory usage: assess whether it's reasonable or consider upgrading the instance configuration. Please refer to <a href="#">Troubleshooting and Resolving Login Failures Due to High CPU or Memory Usage</a> for more information.
	Is the file system inode usage too high?	Has the file system inode usage exceeded 95%?	Warnings	Excessive disk space usage may lead to system anomalies and prevent data writing. It is recommended to assess whether you can delete some files or expand the disk size. Please refer to <a href="#">Solving the Problem of Full Disk Space</a> for resolution.
	Checking if the file system	Has the disk space usage	Warnings	High disk inode usage may lead to system anomalies and prevent data writing. It is recommended to evaluate whether you can delete

space usage is too high	exceeded 95%?	some files or expand the disk size. Please refer to <a href="#">Resolving File System Inode Full Issues</a> for guidance on fixing the problem.
-------------------------	---------------	---

### Internal Detection for Windows Hosts

Category		Detection Instructions	Risk level	Solution
OS Environment	Windows Operating System Version	Is the Windows operating system Windows Server 2008 R2 or an earlier version?	Warnings	Windows Server 2008 R2 and earlier versions have inferior security, stability, and compatibility, and are no longer maintained by Microsoft or Tencent Cloud. It is recommended to follow the <a href="#">processing steps</a> to back up your data and reinstall Windows Server 2016 or a higher version.
	Memory limit	Has a memory limit been set?	Warnings	The installed system memory is not being fully utilized. To remove the memory limit, please refer to the <a href="#">resolution steps</a> .
	CPU Limit	Has a CPU limit been set?	Warnings	If the allocated CPU is not being fully utilized, refer to <a href="#">Resolution Steps</a> to remove the CPU limit.
	Handle leakage	Are the number of file handles normal?	Warnings	There may be a process handle leak; please refer to the <a href="#">troubleshooting steps</a> to investigate and resolve the issue.

System brute-force attacks and security breaches	Determining if the system is experiencing a significant number of brute-force attacks and other types of attacks.	Warnings	The system may experience lag or crashes due to brute force attacks or other malicious activities, affecting normal operations and even posing a risk of data loss. Please refer to <a href="#">Security Group Overview</a> and configure security group policies appropriately through the console, allowing only necessary IPs and ports.
System Environment Variables	Are the system environment variables functioning properly?	Warnings	If system environment variables are missing or abnormal, please refer to <a href="#">Resolution Steps</a> for repair.
System Activation	Determining if the system has been activated	Warnings	System not activated, please follow the <a href="#">activation steps</a> to activate the system.
System time	Is the system time accurate?	Warnings	System time anomaly detected. Please follow the <a href="#">resolution steps</a> to fix it.
System Route Table	Is the system default route missing?	Warnings	The system lacks a default route. Please refer to <a href="#">Resolution Steps</a> for repair.
System Internet Explorer Proxy	Is the Internet Explorer proxy configured?	Warnings	The system has set an IE proxy. Please refer to <a href="#">Troubleshooting Steps</a> for further investigation.

	CD-ROM Status	Checking if the system CD-ROM device is functioning properly	Warnings	CD-ROM Anomaly: The console requires a CD-ROM to reset the password. Please refer to the <a href="#">Troubleshooting Steps</a> for repair.
System Resource Usage	Is the memory utilization too high?	Are there any processes with memory utilization exceeding 80%?	Warnings	If the system memory usage is too high, please refer to the <a href="#">Troubleshooting Steps</a> for investigation.
	Is the virtual memory utilization too high?	Is there insufficient virtual memory resources?	Critical	If the system's virtual memory usage is too high, please follow the <a href="#">resolution steps</a> to fix the issue.
	High total CPU usage	Are there any processes with total CPU usage exceeding 80%?	Warnings	If the system CPU usage is too high, please refer to the <a href="#">troubleshooting steps</a> for investigation.
	High usage of a single CPU	Are there any processes with single CPU utilization exceeding 80%?	Warnings	If the utilization of a single logical CPU is too high, please refer to the <a href="#">troubleshooting steps</a> for investigation.

	Insufficient Available Disk Space	Whether the disk usage exceeds 95% or the available disk space is less than 5GB.	Warnings	Insufficient available disk space, please refer to <a href="#">Recovery Steps</a> for resolution.
	NTFS System Metadata Files	Is the NTFS metadata disk usage proportion high?	Warnings	NTFS metadata usage is too high; please follow the <a href="#">Steps</a> to repair it.
Remote Connection	Remote Desktop Service Status	Is the remote desktop service status abnormal?	Warnings	If the remote desktop service status is abnormal, please refer to the <a href="#">troubleshooting steps</a> for further investigation.
	Remote Desktop Service Port	Whether the remote desktop service port is listening on the default port 3389.	Warnings	The Remote Desktop Service port is not listening. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.
	RDP Listener	Is the RDP Listener enabled?	Critical	The RDP listener is not enabled, preventing remote login. Please refer to the <a href="#">resolution steps</a> for fixing this issue.
	Allow Remote Desktop Connection	Whether to allow Remote Desktop Connection	Critical	RDP is disabled, preventing remote login. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.

	RDP Self-Signed Certificate Expiration Date	Has the RDP self-signed certificate expired?	Critical	If the RDP self-signed certificate has expired, you may not be able to remotely log in. Please refer to the <a href="#">resolution steps</a> for repair.
	Remote Desktop Services Role Installation and Authorization	Determining if the Remote Desktop Services role is installed and licenses are imported.	Warnings	Multi-user login is enabled, but the License has not been imported. Please refer to <a href="#">Repair Steps</a> for resolution.
	Network Access Account	Verify if the network access sharing and security model for local accounts is set to "forceguest".	Critical	The network access account is set as a guest and cannot be used for remote login. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.
	Allow Remote Desktop Service Port through Firewall	Check if the firewall allows Remote Desktop Services.	Warnings	The Windows internal firewall has not allowed the Remote Desktop Service port, preventing remote login. Please refer to the <a href="#">troubleshooting steps</a> for resolution.
Network Configuration	Port Exhaustion	Whether the number of TCP and	Critical	All available ports have been exhausted. Please follow the <a href="#">Steps</a> to resolve the issue.

		UDP ports has been exhausted		
	Timewait/Closewait Connection Count	Is the number of Timewait/Closewait connections normal?	Warnings	An abnormal number of Timewait/Closewait connections may prevent remote login. Please refer to the <a href="#">troubleshooting steps</a> for resolution.
	Gateway Status	Is the gateway status normal?	Warnings	Gateway status is abnormal. Please follow the <a href="#">Steps</a> to resolve the issue.
	MAC Addresses	Is it the default system MAC address?	Critical	The MAC address is not the system default. Please follow the <a href="#">resolution steps</a> to fix it.
	Private Network Domain Name Resolution	Can Tencent Cloud's private network domain names be resolved normally?	Warnings	Intranet domain name cannot be resolved. Please refer to <a href="#">Troubleshooting Steps</a> for resolution.

### Instance Status Detection

Category	Detection Instructions	Risk level	Solution
Instance shutdown	Checks whether the instance is shut down.	Warnings	The instance is shut down. You can go to the <a href="#">CVM console</a> to power it on.
Instance	Checks whether the	Warnings	The instance has been restarted in the last 12 hours. Pay attention to the

restart history	instance has been restarted in the last 12 hours.		instance running status.
Instance kernel crash	Checks whether a hung task has occurred in the instance in the last 12 hours.	Abnormal	The instance has experienced hung tasks, panic, or soft lockups within the last 12 hours. Please pay attention to the instance's operational status. For troubleshooting, refer to <a href="#">Kernel and IO Related Issues</a> .
	Checks whether a panic has occurred in the instance in the last 12 hours.	Abnormal	
	Checks whether a soft deadlock occurred in the instance in the last 12 hours.	Abnormal	

### Instance Performance Detection

Category	Detection Instructions	Risk level	Solution
CPU utilization	Checks whether the instance has experienced a high CPU utilization in the last 12 hours.	Warnings	<p>To avoid becoming a bottleneck for your business, it is recommended to monitor CPU usage and adjust configurations accordingly. For troubleshooting, please refer to the corresponding instance operating system documentation:</p> <ul style="list-style-type: none"> <li>Windows Instance: <a href="#">Unable to log in due to high CPU or memory usage</a></li> </ul>

Memory utilization	Has the instance experienced high memory load in the last 12 hours?	Warnings	<ul style="list-style-type: none"><li>For Linux, refer to: <a href="#">High CPU or Memory Usage Causing Login Failures</a></li></ul>
Basic CPU utilization	Checks whether the instance has experienced a high CPU utilization in the last 12 hours.	Warnings	

## Related Actions

You can refer to [Using Instance Self-Check](#) to generate an instance detection report or view historical detection reports.

# Monitoring and Alarming

Last updated: 2024-06-02 16:05:31

**Monitoring and Alarming** are essential components in ensuring high reliability, availability, and performance of cloud servers. When creating a cloud server, Tencent Cloud Observability Platform is activated by default for free. You can analyze and receive real-time alerts, as well as obtain host monitoring metrics through the platform.

This document provides an overview of the monitoring and alarming features available for cloud servers. For more detailed information, please refer to the [Tencent Cloud Observability Platform Product Documentation](#).

## Overview

Cloud server monitoring and alarming is a management tool for real-time monitoring of cloud servers. The monitoring and alarming features provide comprehensive and detailed monitoring data, extracting key metrics from cloud servers and displaying them in the form of monitoring charts. This allows you to gain a thorough understanding of resource utilization, performance, and operational status of your cloud servers. Additionally, it supports setting custom alarm thresholds and sending notifications based on your defined rules.

## Basic Features

You can access the following CVM monitoring and alarms features in the Cloud Monitor console:

Module	Feature	Main Feature
<a href="#">Monitoring Overview</a>	Tencent Cloud Observability Platform Overview	Provides an overview of the overall status, alarm summary, and comprehensive monitoring information.
<a href="#">Alarm Policy</a>	Customizable alarm thresholds are supported for users.	Currently, the cloud server alarm setting service is supported. When an anomaly occurs in the cloud server metrics, you will be promptly notified to address the issue.
<a href="#">Cloud Server</a>	View Cloud Product	Current Cloud Server Monitoring View

<b>Monitoring</b>	Monitoring Dashboard	
<b>Dashboard</b>	Preset Monitoring Dashboard Custom Monitoring Dashboard	Offers flexible and personalized chart features for cloud server monitoring scenarios, including cross-instance aggregated data, real-time/historical data display, similar metric comparisons, and linked charts.
<b>Traffic Monitoring</b>	Monitoring Traffic	View Overall Bandwidth Information for Users
<b>Tencent Managed Service for Prometheus (TMP)</b>	Inherit open-source Prometheus monitoring capabilities	You can monitor the internal status of applications or services, such as the number of requests processed, orders placed, etc. Additionally, you can monitor the processing time of core logic, such as the time consumed when requesting external services. For more information, please refer to <a href="#">Custom Integration in Cloud Server Scenarios</a> .
<b>Grafana Service</b>	Provide open-source visualization Grafana service	Pre-configured cloud server monitoring dashboard, including various commonly used metrics.

For more information, please see [Basic Features of Tencent Cloud Observability Platform](#).

## Use Cases

- **Daily Management Scenarios:** Log in to the Tencent Cloud Observability Platform console to view the operational status of various Tencent Cloud Observability Platforms.
- **Timely handling of abnormal scenarios:** Sends alarm notifications when monitoring data reaches the alarm threshold, allowing you to promptly receive alerts and investigate the cause of the anomaly.
- **Timely Expansion Scenarios:** By setting alarm rules for monitoring items such as bandwidth, connection count, and disk usage, you can conveniently understand the current status of your cloud services and receive timely alarm notifications for service expansion when the business volume increases.

## Monitoring Items

To monitor instance performance benchmarks, you should monitor at least the following items. You can access the relevant monitoring information in the [Cloud Server Console](#) on the instance details page.

Monitoring Metric	Monitored metrics	Note
CPU utilization	cpu_usage	CPU usage ratio. The data is collected and reported by the internal monitoring component of the server, making the data more accurate.
Memory Utilization	mem_usage	The ratio of the actual amount of memory used by the user to the total amount of memory, excluding the memory occupied by buffer and system cache.
Private network bandwidth out	lan_outtraffic	Average outbound traffic per second of private ENI.
Private network bandwidth in	lan_intraffic	Average inbound traffic per second of private ENI.
Public network bandwidth out	wan_outtraffic	Average outbound traffic per second over the public network. The minimum granularity for bandwidth statistics is 10 seconds (bandwidth calculation method: total traffic in 10 seconds divided by 10 seconds).
Public network bandwidth in	wan_intraffic	Average inbound traffic per second of the public network.
Disk utilization	disk_usage	Disk usage.
Disk I/O wait time	disk_io_wait	Average wait time per disk I/O operation.

## Monitoring Data

- Monitoring Interval:** Tencent Cloud Observability Platform currently offers various monitoring data aggregation granularities, including 10 seconds, 1 minute, 5 minutes, 1 hour, and 1 day. Cloud servers support a 1-minute monitoring granularity, meaning data is aggregated every minute. By default, the interval is set to 5 minutes.

- **Data Storage:** Monitoring data at second-level granularity is stored for 1 day; data at 1-minute and 5-minute granularities is stored for 31 days; data at 1-hour granularity is stored for 93 days; and data at 1-day granularity is stored for half a year.
- **Alarm Display:** Data is presented in easy-to-read charts, and the console integrates monitoring data from all products, providing users with a comprehensive overview of system performance.
- **Alarm Settings:** You can set monitoring metric thresholds, and when the conditions are met, timely alarm notifications will be sent to the concerned parties. For more information, refer to [Creating Alarm Policies](#).
- **Dashboard Configuration:** You can set up a dashboard for monitoring metrics, allowing you to dynamically analyze abnormal metric causes. You can also observe real-time metric changes and promptly scale resources as needed. For more information, refer to [Create a Dashboard](#).

# Access Management

## CAM Overview

Last updated: 2024-06-02 10:16:51

If you have multiple users managing different Tencent Cloud services such as CVM, VPC, and database, and they all share your Tencent Cloud account access key, you may face the following problems:

- Your key will be easily compromised because it is shared by several users.
- Your users might introduce security risks from misoperations due to the lack of user access control.

In this case, you can use sub-accounts to assign different services to different users, mitigating the aforementioned issues. By default, sub-accounts do not have the rights to use CVM or access CVM-related resources. Therefore, you need to create policies that grant sub-accounts the necessary permissions to use the resources they require.

Cloud Access Management (CAM) is a web-based service provided by Tencent Cloud, primarily designed to help customers securely manage access permissions to resources within their Tencent Cloud accounts. With CAM, you can create, manage, and terminate users (groups), and control which individuals can access specific Tencent Cloud resources through identity and policy management.

When using CAM, you can associate a policy with a user or user group to allow or deny them access to specified resources for completing specified tasks. For more information on CAM policies, see [Policy Syntax](#). For more information on how to use CAM policies, see [Policy](#).

If you do not need to manage access permissions to CVM-related resources for sub-accounts, you can skip this section. Skipping this part will not affect your understanding and usage of the remaining portions of the documentation.

### Get Started

CAM policies must either grant or deny the use of one or more CVM actions. Additionally, they must specify the resources that can be used for these actions (either all resources or a subset, depending on the operation). Policies can also include conditions set for operating resources. Some CVM API operations do not support resource-level permissions, which means that you cannot specify a particular resource when using such operations; instead, you must specify all resources for usage.

Task scheduling	Document
-----------------	----------

Basic policy structure	<a href="#">Policy Syntax</a>
Operation definition in a policy	<a href="#">Operations of CVM</a>
Resource definition in a policy	<a href="#">CVM Resource Path</a>
Apply conditions to restrict policies.	<a href="#">Conditional Keys for CVM</a>
Resource-level permissions supported by CVM	<a href="#">Resource-level Permissions Supported by CVM</a>
Console Samples	<a href="#">Console Samples</a>

# Authorizable Resource Type

Last updated: 2023-09-12 14:24:29

Resource-level permission refers to the ability to specify which resources a user can perform operations on. Cloud Virtual Machine (CVM) partially supports resource-level permission, meaning that for CVM operations that support resource-level permission, you can control when a user is allowed to perform operations or use specific resources. For example, you can [grant a user permission to operate CVMs in the Guangzhou region](#).

The following resource types can be authorized in Cloud Access Management (CAM):

ResourceType	Resource Description Method in Authorization Policy
<a href="#">Cloud Virtual Machine Instance-related</a>	<code>qcs::cvm:\$region::instance/*</code>
<a href="#">Cloud Virtual Machine Key-related</a>	<code>qcs::cvm:\$region::keypair/*</code>
<a href="#">Cloud Virtual Machine Image-related</a>	<code>qcs::cvm:\$region:\$account:image/*</code>

[Cloud Virtual Machine Instance-related](#), [Cloud Virtual Machine Key-related](#), and [Cloud Virtual Machine Image-related](#) sections introduce the currently supported CVM API operations with resource-level permissions, as well as the resources and condition keys supported by each operation. **When setting the resource path**, you need to replace variables such as `$region` and `$account` with your actual parameter information, and you can also use the \* wildcard in the path. For related operation examples, please refer to [Access Management Examples](#).

## Note

For CVM API operations not listed in the table, it means that they do not support resource-level permissions. For a CVM API operation that does not support authorization at the resource level, you can still authorize a user to perform it, but you must specify "\*" as the resource element in the policy statement.

## CVM Instance

Configuring using API	Resource Path	Condition Key
-----------------------	---------------	---------------

<b>DescribeInstanceInternetBandwidthConfigs</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ModifyInstanceInternetChargeType</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ModifyInstancesAttribute</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ModifyInstancesProject</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ModifyInstancesRenewFlag</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>RebootInstances</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>RenewInstances</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ResetInstance</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId qcs::cvm:\$region:\$account:systemdisk/*</pre>	<pre>cvm:region cvm:zone cvm:instance_type</pre>
<b>ResetInstancesInternetMaxBand</b>	<pre>qcs::cvm:\$region:\$account:instance/*</pre>	<pre>cvm:region cvm:zone</pre>

<b>width</b>	<code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:instance_type</code>
<b>ResetInstancesPassword</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>ResetInstancesType</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>ResizeInstanceDisks</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>RunInstances</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:image/*</code> <code>qcs::cvm:\$region:\$account:image/\$imageId</code> <code>qcs::cvm:\$region:\$account:keypair/*</code> <code>qcs::cvm:\$region:\$account:keypair/\$keyId</code> <code>qcs::cvm:\$region:\$account:sg/*</code> <code>qcs::cvm:\$region:\$account:sg/\$sgId</code> <code>qcs::vpc:\$region:\$account:subnet/*</code> <code>qcs::vpc:\$region:\$account:subnet/\$subnetId</code> <code>qcs::cvm:\$region:\$account:systemdisk/*</code> <code>qcs::cvm:\$region:\$account:datadisk/*</code> <code>qcs::vpc:\$region:\$account:vpc/*</code> <code>qcs::vpc:\$region:\$account:vpc/\$vpcId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>StartInstances</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>StopInstances</b>	<code>qcs::cvm:\$region:\$account:instance/*</code> <code>qcs::cvm:\$region:\$account:instance/\$instanceId</code>	<code>cvm:region</code> <code>cvm:zone</code> <code>cvm:instance_type</code>
<b>TerminateInstances</b>	<code>qcs::cvm:\$region:\$account:instance/*</code>	<code>cvm:region</code> <code>cvm:zone</code>

```
qcs::cvm:$region:$account:instance/$instanceId
```

```
cvm:instance_type
```

## CVM Key

Configuring using API	Resource Path	Condition Key
<a href="#">AssociateInstancesKeyPairs</a>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId</pre>	—
<a href="#">CreateKeyPair</a>	<pre>qcs::cvm:\$region:\$account:keypair/*</pre>	—
<a href="#">DeleteKeyPairs</a>	<pre>qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId</pre>	—
<a href="#">DescribeKeyPairs</a>	<pre>qcs::cvm:\$region:\$account:keypair/*</pre>	—
<a href="#">DescribeKeyPairsAttribute</a>	<pre>qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId</pre>	—
<a href="#">DisassociateInstancesKeyPairs</a>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId</pre>	—
<a href="#">ImportKeyPair</a>	<pre>qcs::cvm:\$region:\$account:keypair/*</pre>	—
<a href="#">ModifyKeyPairAttribute</a>	<pre>qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId</pre>	—

## CVM Image

Configuring using API	Resource Path	Condition Key
<b>CreateImage</b>	<pre>qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:image/*</pre>	<b>cvm:region</b>
<b>DeleteImages</b>	<pre>qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId</pre>	<b>cvm:region</b>
<b>DescribeImages</b>	<pre>qcs::cvm:\$region:\$account:image/*</pre>	<b>cvm:region</b>
<b>DescribeImagesAttribute</b>	<pre>qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId</pre>	<b>cvm:region</b>
<b>DescribeImageSharePermission</b>	<pre>qcs::cvm:\$region:\$account:image/*</pre>	<b>cvm:region</b>
<b>ModifyImageAttribute</b>	<pre>qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId</pre>	<b>cvm:region</b>
<b>ModifyImageSharePermission</b>	<pre>qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId</pre>	<b>cvm:region</b>
<b>SyncImages</b>	<pre>qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId</pre>	<b>cvm:region</b>

# Authorization Policy Syntax

Last updated: 2024-09-24 15:20:41

## Policy Syntax

CAM policy:

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "effect",
      "action": ["action"],
      "resource": ["resource"],
      "condition": {"key": {"value"}}
    }
  ]
}
```

- **Version** is required, and currently, only the value "2.0" is allowed.
- **Statement** is used to describe the details of one or more permissions. It includes a permission or permission set of multiple other elements such as `effect`, `action`, `resource`, and `condition`. A policy has only one `statement` element.

**1.1 Effect** describes the result of a statement, whether it is "allow" or "explicit deny". It includes two cases: "allow" and "deny" (explicit deny). This element is required.

**1.2 Action** is used to describe the allowed or denied operations. Operations can be an API (described with a "name" prefix) or a feature set (a group of specific APIs described with a "permit" prefix). This element is required.

**1.3 Resource** describes the specific data being authorized. A resource is described using a six-segment format. Detailed resource definitions vary by product. For information on specifying resources, refer to the product documentation corresponding to the resource statement you are writing. This element is required.

**1.4 Condition** describes the constraints under which the policy takes effect. Conditions consist of operators, keys, and values. Condition values may include information such as time and IP address. Some services allow you to specify other values in the conditions. This element is optional.

## CVM Operations

In a CAM policy statement, you can specify any API operation from any service that supports CAM. For CVM, use APIs prefixed with `name/cvm:`. For example: `name/cvm:RunInstances` or `name/cvm:ResetInstancesPassword`.

To specify multiple operations in a single statement, separate them with commas as shown below:

```
"action":["name/cvm:action1","name/cvm:action2"]
```

You can also use wildcards to specify multiple actions. For example, you can specify all actions whose names begin with the word "Describe", as follows:

```
"action":["name/cvm:Describe*"]
```

To specify all operations in CVM, use the wildcard "\*" as follows:

```
"action":["name/cvm:*"]
```

## CVM Resource Path

Each CAM policy statement has its own applicable resources. Resource paths are generally in the following format:

```
qcs:project_id:service_type:region:account:resource
```

**project\_id:** Describes project information, only for compatibility with early CAM logic, no need to fill in.

**service\_type:** Product abbreviation, such as CVM.

**region:** Regional information, such as bj.

**account:** Root account information of the resource owner, such as uin/164256472.

**resource:** Detailed resource information of each product, such as instance/instance\_id1 or instance/\*.

For example, you can specify a specific instance (i-15931881scv4) in the statement as follows:

```
"resource":["qcs::cvm:bj:uin/164256472:instance/i-15931881scv4"]
```

You can also use the wildcard "\*" to specify it for all instances that belong to a specific account as shown below:

```
"resource":["qcs::cvm:bj:uin/164256472:instance/*"]
```

If you want to specify all resources or if a specific API operation does not support resource-level permission, you can use the wildcard (\*) in the `resource` element as shown below:

```
"resource": ["*"]
```

To specify multiple resources in one policy, separate them with commas. In the following example, two resources are specified:

```
"resource": ["resource1", "resource2"]
```

The following table describes the resources CVM can use and the corresponding resource description methods.

In the following table, the words prefixed with \$ are all alternative names.

- “project” refers to project ID.
- “region” refers to region.
- “account” refers to account ID.

Resources	Resource Description Method in Authorization Policy
Instance	qcs::cvm:\$region:\$account:instance/\$instanceId
Key	qcs::cvm:\$region:\$account:keypair/\$keyId
VPC	qcs::vpc:\$region:\$account:vpc/\$vpcId
Subnets	qcs::vpc:\$region:\$account:subnet/\$subnetId
Video flipping	qcs::cvm:\$region:\$account:image/*
CBS	qcs::cvm:\$region:\$account:volume/\$diskId
Security Group	qcs::cvm:\$region:\$account:sg/\$sgId
EIP	qcs::cvm:\$region:\$account:eip/*

## CVM Condition Keys

In policy statements, you can optionally specify conditions to control the effective time of the policy. Each condition contains one or more key-value pairs. Condition keys are case-insensitive.

- If you specify multiple conditions or specify multiple keys in one condition, they will be evaluated with the logical operation "AND".

- If you specify a key with multiple values in a single condition, it will be evaluated using the logical OR operation. All conditions must be met to grant permission.  
The table below describes the condition keys used by CVM for specific services.

Condition keys	Reference types	Key value pairs
<code>cvm:instance_type</code>	String	<p><code>cvm:instance_type= instance_type</code></p> <ul style="list-style-type: none"> <li>• Here, <code>instance_type</code> refers to the instance type (e.g., S1.SMALL1).</li> </ul>
<code>cvm:image_type</code>	String	<p><code>cvm:image_type= image_type</code></p> <ul style="list-style-type: none"> <li>• Here, <code>image_type</code> refers to the image type (e.g., IMAGE_PUBLIC).</li> </ul>
<code>vpc:region</code>	String	<p><code>vpc:region= region</code></p> <ul style="list-style-type: none"> <li>• Here, <code>region</code> refers to the region (for example, "ap-guangzhou").</li> </ul>
<code>cvm:disk_size</code>	Integer	<p><code>cvm:disk_size= disk_size</code></p> <ul style="list-style-type: none"> <li>• In this context, <code>disk_size</code> refers to the disk size (e.g., 500).</li> </ul>
<code>cvm:disk_type</code>	String	<p><code>cvm_disk_type= disk_type</code></p> <ul style="list-style-type: none"> <li>• Here, <code>disk_type</code> refers to the disk type (e.g., CLOUD_BASIC).</li> </ul>
<code>cvm:region</code>	String	<p><code>cvm:region= region</code></p> <ul style="list-style-type: none"> <li>• Here, <code>region</code> refers to the region (for example, "ap-guangzhou").</li> </ul>

# CAM-Supported APIs

Last updated: 2023-09-07 16:33:44

## Basic information

Product Name	Abbreviation	Authorization Granularity
CVM	cvm	Resource-level

### Note

Three authorization granularity levels are supported by the Tencent Cloud product: service level, operation level, and resource level.

- **Service Level:** Determines whether a user has overall access to a service, either granting full access to all operations or denying access to all operations. Cloud products with service-level authorization granularity do not support authorization for specific APIs.
- **Operation Level:** Defines whether access to specific service interfaces (APIs) is allowed, for example, granting read-only access to a certain account for CVM services.
- **Resource-level:** Determines whether access to specific resources is allowed, providing the finest level of authorization granularity. For example, granting an account read and write access to a specific CVM instance. Products that support resource-level API authorization are considered to have resource-level authorization granularity.

## API authorization granularity

- **Resource-level API:** This type of API supports authorizing a specific, individual resource.
- **Action-level API:** This type of API does not support authorizing a specific resource.

During authentication with a resource-level API, the cloud product passes the specific 6-segment resource description to CAM for authentication, allowing authorization and authentication of a specific resource.

For operation-level APIs, the cloud product does not pass the specific 6-segment resource description to CAM for authentication, but only passes an arbitrary resource `*`. Therefore, if the policy syntax restricts a specific resource during authorization, and the API does not pass

that resource during authentication, CAM will determine that the API is not within the authorized scope and will deem it as lacking permission.

## Procedure

For more information, see [Write Operations](#) , [List Operations](#) , and [Read Operations](#) .