


Cloud Virtual Machine

CVM Introduction

Product Introduction



Copyright Notice

©2013-2018 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice

 Tencent Cloud

All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

CVM Introduction

Overview

Regions and Availability Zones

Our Advantages

Use Cases

Instance

Instance Overview

Instance Types

Life Cycle

Spot Instance

Storage

Overview

Cloud Block Storage

Local Storage

Cloud Object Storage (COS)

Image

Overview

Image Types

Network and Security

Overview

Placement Group

Network Environment

Private Network Access

Internet Access

Elastic Public IP (EIP)

Elastic Network Interface (ENI)

Login Password

SSH Key

Security Groups

Monitoring and Warning

Access Control

Access Control Overview

Authorizable Resource Type

Authorization Policy Syntax

CVM Introduction

Overview

Last updated : 2018-07-20 17:22:57



Introduction of Tencent Cloud CVM

Tencent Cloud Cloud Virtual Machine (CVM) provides extendable calculating capacity on the cloud, frees you from estimation of resources usage and in-advance investment. With Tencent Cloud CVM, you can enable CVMs and deploy applications immediately. Tencent Cloud CVM allows you to configure your resources according to your actual requirements, such as CPU, memory, storage, network, security and so on.

Glossary:

- **Instance:** Virtual computing resources on the cloud
- **Region and availability zone:** Position of instances and other resources
- **Image:** A preset instance template, including pre-configuration environment of the server (operating system and other installed software)
- **Instance type:** Different configurations of instance: CPU, memory, storage, network, etc.
- **Local disk:** Devices that are located on the same physical server with instances and can be used for persistent storage by instances
- **Cloud Block Storage:** Distributed persistent block storage devices provided, can be used as system disk for instances or expandable data disk
- **SSH Key Login:** a more secured login method, in which the public key is kept by Tencent Cloud, and private key is kept by user); **Password Login:** log into Tencent Cloud using ordinary password.
- **Private IP address:** Internal service address of the instance; **Public IP address:** External service address of the instance
- **Security Group:** Secure access control of instances that specifies rules for IP, protocol and port of incoming and outbound instances
- **Virtual Private Cloud:** Customized virtual cyberspace, logically isolated from other resources
- **Elastic IP:** Static public IP designed specifically for dynamic network to meet demands for fast troubleshooting
- **Metadata:** Tag for identifying instance sources within the instances
- **Console:** Web-based user interface

If you're new to Tencent Cloud CVM, please refer to [Quick start for Windows CVM](#) and [Quick start for Linux CVM](#).

Related Services

- [Cloud Load Balance](#): may realize automatic distribution of request traffic from clients across multiple CVM instances.
- [Auto Scaling](#): may automatically increase or decrease the quantity of server clusters at regular time or based on specific conditions.
- [Cloud Monitor](#): can be used to monitor operation statistical data for CVM instances.
- [Cloud Database](#): can be used to deploy your relational database on the cloud.

Pricing of CVM

For pricing details of CVMs, see [here](#).

Regions and Availability Zones

Last updated : 2018-09-13 15:38:46

Tencent Cloud CVM's entrusted data centers are located in different regions around the world, which are all made up of region and availability zone.

Each region is an independent geographical location. Within each region, there are multiple mutually isolated locations referred to as availability zones. Each region is entirely independent. Each availability zone is independent. However, availability zones in the same region can be connected via private network links with low latency.

Tencent Cloud supports distribution of cloud resources in different locations by users. Users are advised to consider placing resources in different availability zones when designing system to shield "service unavailable" status caused by single point of failure.

Region

Regions are completely isolated with each other, so as to ensure high stability and fault tolerance. Tencent Cloud now have 3 regions in Mainland China and provides nodes in Hong Kong, Singapore and Toronto. You are recommended to choose nearest region to reduce access latency and improve download speed.

Region attribute is differentiated for all behaviors such as enabling and viewing instances by users. If image of the instance that the users need to enable does not exist in the region, then the image needs to be copied to current region. For more information, please see [Copy Image](#).

- Resources in different availability zones of the same region are connected via private network, and can be accessed directly using [Private IP](#).
- Cloud services of different regions **cannot communicate via private network by default**.
- CVMs cannot access other CVMs, Cloud Database or Cloud Memcached across regions by default.
- When binding Cloud Load Balance to the server, only CVMs in the current region can be chosen;
- Cloud resources on different regions can communicate via [Public IP](#). Cloud Services on VPC can access Internet via [peering connection](#) service.
- [Cloud Load Balance](#) does not support cross-region data forwarding.
- Regions are named by Coverage + City, and availability zones are named by city + serial number.
- Only resources under the same account can communicate via private network.

Notes for Hong Kong region:

- The following cloud services are temporarily unavailable: Cloud Memcached, elastic web engine, Cloud Object Storage, Cloud Block Storage, one-click opening of server and domain binding with separated regions and servers.
- When you need to log in to CVMs in Hong Kong region, log in via jump server is recommended for better operation and maintenance experience.

Notes for North America region:

- The following cloud services are temporarily unavailable: Cloud Memcached, elastic web engine, Cloud Object Storage, mobile acceleration, Cloud Automated Testing, one-click opening of server and domain binding with separated

regions and servers.

- Due to the considerable latency between North America and China, when you need to log in to CVMs in North America region, login via jump server is recommended for better operation and maintenance experience.

Availability Zone

Availability zones (Zone) refer to Tencent Cloud's physical data centers whose power and network are independent from each other within the same region. They are designed to ensure that the failures within an availability zones can be isolated (except for large-scale disaster or major power failure) without spreading to and affecting other zones so that users' businesses can provide continuous online services. By starting an instance in an independent availability zone, users can protect their applications from being affected by the failures occurring in a single location.

When starting an instance, users can choose any availability zone within the specified region. If a user needs to ensure the high reliability of application systems so that the services are still available even when a failure occurs in an instance, the user can use cross-zone deployment scheme (e.g. [Cloud Load Balance](#), [Elastic IP](#), etc.) to allow the instance in another availability zone to handle the relevant requests in replace of the failed instance.

Migrating an instance to another availability zone

Once an instance is started, its availability zone cannot be changed. However you can migrate it to another availability zone.

1. Create a custom image for the instance (see [Create Custom Image](#)).
2. If the instance is using [Virtual Private Cloud](#) and you want to keep the private IP address after the migration, you can first delete the subnet in the current availability zone and then create a subnet in the new availability zone with the same IP address range as that of the original subnet. Please note that a subnet can be deleted only when it contains no available instances. Thus, all the instances in the current subnet should be migrated to the new subnet.
3. Create a new instance in the new availability zone using the custom image you have just created. User can choose the same type and configuration as those of the original instance, or choose new ones.(see [Purchase and Start an Instance](#)).
4. If an elastic IP address is associated with the original instance, then dissociate it from the old instance and associate it with the new instance (see [Elastic IP](#)).
5. (Optional) For [postpaid](#) instances, you can terminate them manually (see [Terminate an Instance](#)). [Prepaid](#) instances will be reclaimed automatically after expiration.

How do I select the region and availability zone?

While purchasing Tencent Cloud services, it is recommended to choose the region that is closest to your customers to minimize the access latency and improve download speed.

Resource Availability

Specify which resources of Tencent Cloud are global, which resources are regional and not specific to any availability zone and which resources are based on availability zones.

--	--	--	--

Resource	Resource ID Format	Availability	Description
User ID	Not limited	Cross-region	Each Tencent Cloud user account is unique and applicable to all resources under the account.
SSH key	skey-xxxxxxx	Cross-region	You can bind an SSH key pair to any CVM in any region under your account.
CVM instance	ins-xxxxxxx	One zone only	A CVM instance is created under a specified availability zone
Custom Image	img-xxxxxxx	Cross-zone	A custom image can be used in different availability zones of the same region. You can copy an image to another region.
Elastic IP	eip-xxxxxxx	Cross-zone	An EIP is created under a specified region, and can only be associated with instance in the same region.
Security Group	sg-xxxxxxx	Cross-zone	A security Group is created under a specified region, and can only be associated with instance in the same region.
(Elastic) Cloud Disk	disk-xxxxxxx	Cross-zone	Elastic cloud disks can be created independently under a certain region and mounted to instances in the same region.
Cloud Load Balance	clb-xxxxxxx	Cross-zone	Cloud Load Balance can bind to CVMs in different availability zones in a single region for forwarding traffic.
Virtual Private Cloud (VPC)	vpc-xxxxxxx	Cross-zone	VPC is created under a certain region. Resources belonging to the same VPC can be created under different availability zones
Subnet	subnet-xxxxxxx	One zone only	Users cannot create a subnet across availability zones.
Routing table	rtb-xxxxxxx	Cross-zone	When creating a routing table, users need to designate a specific VPC, and therefore location attribute of the VPC should be followed.

Our Advantages

Last updated : 2018-08-20 15:59:58



Comprehensive

Tencent Cloud CVM provides comprehensive and extensive service contents for you.

- **Available in multiple regions and availability zones:** South China, East China, North China and Southwest China are covered for Mainland China. Overseas nodes are deployed in Southeast Asia, Asia Pacific, North America, Western U.S. and Europe. Deployment of application in regions close to the location of your clients enables lower delay.
- **Model Configurations:**
 - **Standard:** Suitable for small and medium-sized Web applications and databases.
 - **MEM optimized:** Suitable for applications that require extensive memory operations, searches and computing.
 - **High IO:** Recommended for low latency, I/O-intensive applications.
 - **Computing:** Suitable for large game servers and advertisement service engines, high performance computing and other compute-intensive applications.
 - **Big data:** Suitable for throughput-intensive applications such as Hadoop distributed computing, massive log processing, distributed file systems, and large data warehouses.
 - **Heterogeneous:** Suitable for high-performance applications such as deep learning, scientific computing, video encoding/decoding, and graphics workstations.
 - **Batch-based:** Suitable for compute-intensive applications that use super large computing nodes frequently in a short time, such as rendering, gene analysis, and crystal pharmacy.

Elastic

We are committed to building the most elastic cloud service management platform in the industry with the following capabilities:

- **Hardware configurations:** Instant upgrading/downgrading** of hardware configurations for CVM based on CBS (applicable to both prepaid and postpaid billing methods)
- **Disk change:** Instant capacity expansion** of disk for CVM based on CBS (applicable to both prepaid and postpaid billing methods)
- **Network bandwidth:** Instant upgrading/downgrading** of bandwidth for CVM (limited to postpaid billing method, and prepaid billing method supports instant bandwidth upgrading)
- **Billing method:** CVM supports switch from Bill-by-bandwidth method to Bill-by-traffic method (limited to prepaid billing method)
- **Operating system:** CVM supports switch between Windows and Linux systems at any time (applicable to both prepaid and postpaid billing methods)
- **EIP:** It can be bound to CVMs in various network environments
- **Image type:** Public image (with multiple Linux and Windows operating system types), service marketplace image (integrated with third parties for different application environments), and custom image (image created by users through the image creation feature). Cross-region adjustment and image copying are supported.
- **Custom network structure:** Virtual Private Cloud (VPC) provides independent network space with custom division of network segment as well as IP address and custom router strategy etc. Provide port-level access control for comprehensive logic isolation of network.

Reliable

We are committed to building the most reliable CVM in the industry.

- **CVM reliability:** We guarantee a 99.95% service availability for CVMs with data reliability ranging from 99.99% to 99.999999%. Features including transparent migration during downtime, data snapshot and automatic alarm are supported to ensure the security of your servers.
- **Cloud disk policy:** 3-copy professional storage policy is provided to eliminate single points of failure and ensure reliability of data. Thus, you can rest assured when leaving your data on the cloud without worrying about losing data.
- **Stable network architecture:** CVM comes with well-developed network virtualization technology and ENI binding technology to deliver high network availability. Operation in T3+ IDCs ensures the reliability of the operating environment and unleashes you from the problem concerning network availability.

Fast

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

- **Convenient and rapid operations:** In just a few minutes, you can easily obtain one, hundreds of, or even thousands of server instances, and also purchase, configure, expand or manage your services with just one click.
- **Top-speed public network:** BGP public network with more than 20 lines covers almost all ISPs. No matter which ISP your clients are using, they can always enjoy the same top-speed bandwidth. Optional public network bandwidth of maximum 200 Mbps enables failover in seconds.

- **Top-speed private network:** The private networks of Tencent Cloud's data centers in the same region are interconnected, all with megabyte or gigabyte underlying networks, thus ensuring communication quality within the private networks.

Secure

Tencent Cloud provides multiple solutions to ensure the security of CVMs, as well as backup and rollback mechanisms to protect the security of data.

- **Various methods for remote login to the CVM:** Various login methods are provided, including login with key/password, login via VNC.
- **Diverse security services:** DDoS protection, DNS hijacking detection, intrusion detection, vulnerability scanning, Web Trojan detection, login protection and other security services are provided to safeguard your server.
- **Free cloud monitoring:** Multiple types real-time alerts are supported.
- **Recycle bin protection mechanism:** Cloud services billed on a prepaid basis can be retained in the recycle bin for a certain period of time upon their expiration to avoid data loss and other negative impacts due to immediate termination.
- **Custom access control:** Customize CVM and network access policies through security group and network ACL, and configure different firewalls for different instances in a flexible manner.

Tencent Cloud security service features the following highlights:

- **Comprehensive security protection**
Comprehensive security services are provided for CVM including security check (vulnerability scanning, embedded Trojan detection, web backdoor detection, port security inspection, etc.) and security protection (DDoS protection, intrusion detection, access control), to ensure data security and user privacy.
- **Real-time alarm with regular analysis**
Security services are provided on a 24/7 basis, so that vulnerabilities can be identified in first time and you are notified of this free of charge in real time.
- **Free and easily accessible security protection**
Without having to purchase expensive security equipment for your CVM, your purchase of Cloud services comes with free Cloud Security services. Convenient one-click activation without deployment that can be performed with ease.
- **Professional team giving you reliable protection**
Cloud security developed by Tencent security team with years of expertise and experience in security provides professional and trustworthy security services to cloud services users.

Ease of use

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

- Tencent Cloud CVM provides a Web-based user interface, the Console, allowing you to perform operations such as launching, configuration adjustment and system re-installation on CVM instances just as with physical machines. If you have signed up for a Tencent Cloud account, you can directly [log in to the CVM console](#) to operate on your CVM.

- Tencent Cloud CVM provides an API system, which can be used for easy integration of CVM and your internal monitoring and operation systems to achieve business operation and maintenance system that closely meets business demands and is fully automated. These requests are HTTP/HTTPS requests. For more information about CVM API operations, please see [API Documentation](#).
- If you want to use API to manage your resources, applications and data, you can use SDK (supporting PHP/Python/Java/.NET/Node.js) for programming or use Tencent Cloud CLI to call CVM API. For more information, please see [Use SDK](#) and [Use CLI](#).

Economic

Tencent Cloud provides multiple billing methods at a reasonable price and simplifies traditional OPS tools, thus saving extra IT investment.

- CVM instance and its network deployment both support prepaid and postpaid billing methods to meet the demands of different scenarios.
- You can make the purchase depending on your own demands. Reasonable consumption is realized as pre-purchase or preparation for hardware resources is not required. This helps you effectively reduce investment in infrastructure.

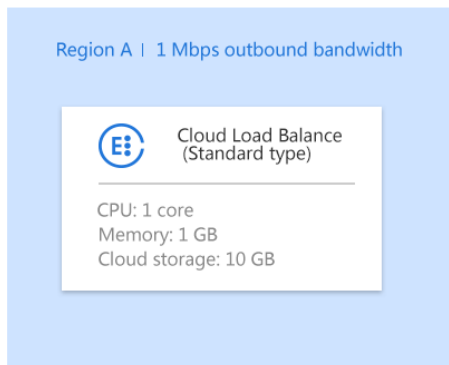
Use Cases

Last updated : 2018-07-17 11:32:32



Static exhibition websites such as small Web/APP and company official websites

Static exhibition websites such as small Web/APP and company official websites

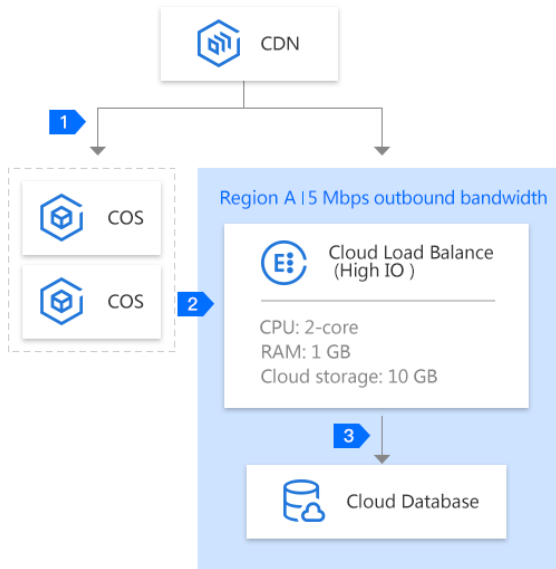


For small-scale websites and Apps, the incoming traffic keeps low during the initial period. It's recommended to purchase a basic-level configuration server for the deployment of application/codes, configuration files, static files, databases and other resources. Tencent Cloud Service Market provides rich images, Apps and OPS tools to satisfy all your needs.

With elastic computing, you can scale up your system when your business traffic increases rapidly. On Tencent Cloud Console, you can easily adjust your CVM configurations (such as CPU, memory, disk capacity, bandwidth), or launch more CVMs with CLB to establish a horizontal scaling traffic distribution system.

Streaming media application for large files such as massive images/videos

Streaming media applications



Basic concept of architecture : Static and dynamic separation

When a user frequently reads big data, I/O may become a bottleneck of cloud virtual machine (CVM). By using COS and CDN service, this problem can be solved easily:

- 1 Forward requests for static files to CDN nodes, and acquire files stored on COS
- 2 Forward others requests to multiple high-IO CVMs backend via load balancers
- 3 Dynamic data including user info can be written/read in CDB of backend server

Instance

Instance Overview

Last updated : 2018-07-30 11:57:35

What is CVM Instance?

Instance can be understood as a CVM, which contains the most basic computing components such as CPU, memory, operating system, network, and disk.

CVM instances can 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 the procurement costs of your software/hardware, and simplify IT OPS work.

Different types of instances provide different computing and storage capacities, and they are applicable to different application scenarios. You can choose the computing capacity, storage, and network access method of the instances based on the service scope as needed. For more instance types and application scenarios, please see [CVM Instance Configuration](#) and [Model Recommendation](#). After an instance is started, you can use it in the same way as you use any other traditional computer, and you have full control over the started instance.

Instance Image

Image is the template for CVM software configuration (operating system, pre-installed programs, etc.). A Tencent Cloud image provides all information required to launch a CVM instance. Tencent Cloud requires users to start the instance via image. An image can launch multiple instances and you can use the image repeatedly. In another word, an image is an "installation disk" of a CVM.

Tencent Cloud provides the following types of images:

- Public images: Available to all users. They cover most mainstream operating systems.
- Service market images: Available to all users. In addition to operating systems, some applications are also integrated into these images.
- Custom images: Only available to the creator and users with whom these images are shared. They are created from running instances or imported from external sources.
- Shared images: Shared by other users. They can only be used to create instances.

For more information about images, please see [Image Overview](#) and [Image Type](#).

Instance Storage

The storage of instances is similar to normal CVMs and is divided into **system disk** and **data disk**:

- System disk: Similar to the C drive in the Windows system. The system disk contains a full copy of the image used to launch an instance, and the running environment for the instance. A larger system disk than the used image is required

when starting.

- **Data disk:** Similar to D and E drives in the Windows system. Data disk is used to save user data, and supports free expansion, mounting and unmounting.

Both system disk and data disk support different types of storage provided by Tencent Cloud. For more information, please see [Storage Overview](#).

Instance Security

Tencent Cloud provides the following instance security protection methods:

- **Policy Control:** When multiple accounts are used to control the same set of cloud resources, users can use Policy Control to manage their access to cloud resources.
- **Security Group:** Users can use a security group to control access by allowing trusted addresses to access the instance.
- **Login control:** Log in to your Linux instances by using [SSH Key](#) whenever possible. For the instances that you [Log in with Password](#), the password needs to be changed from time to time.

Instance Types

Last updated : 2018-09-11 19:07:19

When a Tencent Cloud CVM is created, the instance type specified by the user determines the host hardware configuration of the instance. Each instance type provides different computing, memory and storage functions. Users can choose a proper instance type according to the scale of application to be deployed. These instance families comprise varying combinations of CPU, memory, storage, heterogeneous hardware, and network bandwidth and give you the flexibility to choose the appropriate mix of resources for your applications.

If your demand is stable, we recommend that you choose the prepaid billing method, so your savings will increase with the length of usage. To react to spikes in demand, you can choose the postpaid billing method, which allows you to activate/terminate computing instances at any time and only pay for the actually consumed resources. CVM usage is billed in one-second increments to maximize your savings.

Instance Type

Tencent Cloud instance families are divided into the following types:

Standard Instance Family

This family provides a balance of computing, memory, and network resources to accommodate most applications.

[Standard S3](#) [Standard Network Optimized S2ne](#) [Standard Network Enhanced SN2](#) [Standard S2](#) [Standard S1](#)

Memory Optimized Instance Family

This family features large memory and is suitable for applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching.

[Memory Optimized M3](#) [Memory Optimized M2](#) [Memory Optimized M1](#)

High I/O Instance Family

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

[High I/O I2](#) [High I/O I1](#)

Big Data Instance Family

This instance 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.

Big Data D1

Compute Instance Family

This family comes with a base clock rate of 3.2GHz to provide the highest single-core computing performance. It is suitable for compute-intensive applications such as batch processing, high performance computing, and large game servers.

[Computing CN3](#) [Computing C3](#) [Computing Network Enhanced CN2](#) [Computing C2](#)

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.

[FPGA FX2](#) [GPU Computing GN8](#) [GPU Computing GN2](#) [GPU Rendering GA2](#)

Batch-based Instance Family

With the lowest per core-hour cost, this family is suitable for compute-intensive applications that use super large computing nodes frequently in a short time, such as rendering, gene analysis, and crystal pharmacy.

[Batch Computing BC1](#) [Batch General Purpose BS1](#)

Restrictions on Instances

- The total number of instances that can be started in one zone is limited. For more information, please see [Restrictions on CVM Instance Purchase](#).
- Restrictions on system and data disks mounted on an instance: To ensure premium disk I/O performance, Tencent Cloud sets limits on the size and type of data disks purchased with an instance. For more information, please see the disk matching module of the respective instance family. You can also purchase separate cloud disks if you have higher disk requirements.
- 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 have been sold out. Please see the purchase page for accurate information.

Here is a full list of instance families for different applications.

Standard Instance Family

This family provides a balance of computing, memory, and network resources to support most applications.

Standard S3

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

Standard S3 instances are equipped with the new Skylake Xeon® processors that perform 30% better and the latest DDR4 memories that perform 60% better than Standard S2 instances. Standard S3 instances support up to 10Gbps private network bandwidth.

Features

- 2.5GHz Intel Xeon® Skylake 6133 processors with stable computing performance
- The latest generation of 6-channel DDR4 memories with a memory 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 is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.
- Support all kinds of cloud disks

Application Scenarios

Standard S3 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

Requirements

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S3.SMALL1	1	1	1.5
S3.SMALL2	1	2	1.5
S3.SMALL4	1	4	1.5
S3.MEDIUM4	2	4	1.5
S3.MEDIUM8	2	8	1.5
S3.LARGE8	4	8	1.5
S3.LARGE16	4	16	1.5
S3.2XLARGE16	8	16	1.5
S3.2XLARGE32	8	32	1.5
S3.3XLARGE24	12	24	1.5
S3.3XLARGE48	12	48	1.5
S3.4XLARGE32	16	32	2.0
S3.4XLARGE64	16	64	2.0
S3.6XLARGE48	24	48	3.0
S3.6XLARGE96	24	96	3.0
S3.8XLARGE64	32	64	4.0
S3.8XLARGE128	32	128	4.0
S3.12XLARGE96	48	96	6.0
S3.12XLARGE192	48	192	6.0
S3.16XLARGE128	64	128	8.0

- Both prepaid and postpaid billing methods are available for S3 instances.
- S3 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for S3 instances. Please see the instance specifications on the right. Please make sure that the size of S3 instance you choose can meet the minimum CPU memory requirements of your operating system and applications. In many cases, operating systems with graphical user interface (such as Windows) consuming extensive memory and CPU resources may need larger instances. As the needs of your workload for memory and CPU increase, you can expand to higher configurations or choose other instance types.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S3.16XLARGE256	64	256	8.0
S3.20XLARGE320	80	320	10.0

Standard Network Optimized S2ne

Standard Network Optimized S2ne instances are the best choice for applications that require sending and receiving massive network packets. They support sending and receiving up to millions of network packets per second. Standard Network Optimized S2ne instances are recommended for scenarios with high network PPS requirements, such as large game servers, videos, and live broadcasts.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Features

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)	Network packet forwarding (pps)	Number of queues
S2ne.LARGE8	4	8	1.5	500,000	4
S2ne.2LARGE16	8	16	2.0	800,000	8
S2ne.3LARGE24	12	24	2.5	1 million	8
S2ne.4LARGE32	16	32	3.5	1.6 million	8
S2ne.6LARGE48	24	48	5.0	2 million	8
S2ne.8LARGE64	32	64	7.0	2.5 million	8
S2ne.8LARGE64	48	192	10.0	4.5 million	8

- Intel Xeon E5-2680 Broadwell (v4) processors with a base clock rate of 2.4 GHz and DDR4 memories, offering stable computing performance
- Up to 48 cores and 192 GB are available for sale.
- 1:2 or 1:4 processor to memory ratio.
- The network performance of an instance is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.
- Support all kinds of cloud disks

Application Scenarios

- 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

- Both prepaid and postpaid billing methods are available for S2ne instances.
- S2ne instances can only be started in a VPC.
- Configuration purchase is available for S2ne

instances. Please see the instance specifications on the right.

[Back to the top](#)

Standard Network Enhanced SN2

Standard Network Enhanced SN2 instances come with a 25 GB network environment, featuring greater bandwidth, lower latency, stable computing performance, and the ability to send and receive up to 700,000 packets per second. They are suitable for scenarios that require sending and receiving massive network packets.

Features

- 2.4 GHz Intel Xeon E5-2680 Broadwell (v4) processors, DDR4 memories
- Up to 56 cores and 224 GB are available to meet the need for ultra-large CPUs/memories.
- Support up to 25 Gbps of private network bandwidth to meet extremely high private network transmission requirements.
- Support forwarding up to 700,000 network packets per second to allow for a greater number of concurrent users.
- Support storage options of local disks, HDD cloud disks and SSD cloud disks.

Application Scenarios

- 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
- Small and medium-sized database systems, caches, and search clusters
- Computing clusters, memory-intensive data processing

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
SN2.LARGE8	4	8	2.0
SN2.7XLARGE112	28	112	13.0
SN2.14XLARGE224	56	224	25.0

Requirements

- The prepaid billing method is available for SN2 instances.
- S2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for SN2 instances. Please see the instance specifications on the right.

[Back to the top](#)

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 that bring 40% better integer and floating-point computing performance and DDR4 memories that perform 30% better.

Features

- Intel Xeon E5-2680 Broadwell (v4) processors with a base clock rate of 2.4 GHz and DDR4 memories
- CPU performance is 20% higher than Standard S1
- 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

Application Scenarios

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

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S2.SMALL1	1	1	1.5
S2.SMALL2	1	2	1.5
S2.SMALL4	1	4	1.5
S2.MEDIUM2	2	2	1.5
S2.MEDIUM4	2	4	1.5
S2.MEDIUM8	2	8	1.5
S2.LARGE4	4	4	1.5
S2.LARGE8	4	8	1.5
S2.LARGE16	4	16	1.5
S2.2XLARGE16	8	16	1.5
S2.2XLARGE32	8	32	1.5
S2.3XLARGE24	12	24	2.5
S2.3XLARGE48	12	48	2.5

- S2 instances support both prepaid and postpaid billing methods, and can also be used as production instances of standard host HS20 in CDHs.
- S2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for S2 instances. Please see the instance specifications on the right. Please make sure that the size of S2 instance you choose can meet the minimum CPU memory requirements of your operating system and applications. In many cases, operating systems with graphical user interface (such as Windows) consuming extensive memory and CPU resources may need larger instances. As the needs of your workload for memory and CPU increase, you can expand to higher configurations or choose other instance types.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S2.4XLARGE32	16	32	3.0
S2.4XLARGE64	16	64	3.0
S2.6XLARGE48	24	48	4.5
S2.6XLARGE96	24	96	4.5
S2.8XLARGE64	32	64	6.0
S2.8XLARGE128	32	128	6.0
S2.14XLARGE224	56	224	10.0

Standard S1

The Standard S1 of Series 1 is a type of virtual machines with CPUs ranging from low to high core count. It features moderate prices and flexible options of configurations to meet different needs of users. It also offers the options of local disks, common cloud disks and SSD cloud disks for the data disks (The option may vary with different hardware specifications).

Features

Standard S1 instances have the following features:

- With CPUs ranging from low to high core count, offering you flexible options to configure your CVM
- Intel Xeon CPUs and DDR3 memories
- Support storage options of local disks, HDD cloud disks and SSD cloud disks.
- Balance of computing, memory, and network resources

Application Scenarios

Standard S1 instances are applicable to large-, medium-, and small-sized applications and databases.

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S1.SMALL1	1	1	1.5
S1.SMALL2	1	2	1.5
S1.SMALL4	1	4	1.5
S1.MEDIUM2	2	2	1.5
S1.MEDIUM4	2	4	1.5
S1.MEDIUM8	2	8	1.5
S1.MEDIUM12	2	12	1.5
S1.LARGE4	4	4	1.5
S1.LARGE8	4	8	1.5
S1.LARGE16	4	16	1.5

Requirements

- S1 instances support both prepaid and postpaid billing methods, and can also be used as production instances of standard host in CDHs.
- S1 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for S1 instances. Please see the instance specifications on the right. Please make sure that the size of S1 instance you choose can meet the minimum CPU memory requirements of your operating system and applications. In many cases, operating systems with graphical user interface (such as Windows) consuming extensive memory and CPU resources may need larger instances. As the needs of your workload for memory and CPU increase, you can expand to higher configurations or choose other instance types.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
S1.2XLARGE8	8	8	2.0
S1.2XLARGE16	8	16	2.0
S1.2XLARGE32	8	32	2.0
S1.3XLARGE24	12	24	2.5
S1.3XLARGE48	12	48	2.5
S1.4XLARGE16	16	16	3.5
S1.4XLARGE32	16	32	3.5
S1.4XLARGE64	16	64	3.5
S1.6XLARGE48	24	48	5.0
S1.8XLARGE64	32	64	7.0
S1.12XLARGE96	48	96	10.0

Memory Optimized Instance Family

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.

Memory Optimized M3

As the latest generation of memory optimized instances, Memory Optimized M3 instances are designed to deliver high performance for handling the workloads of large data sets in memory, and are the best choice for applications demanding high-memory computing.

Memory Optimized M3 instances come with the new Intel® Xeon® Skylake processors that bring 30% better performance than Standard S2 instance, and the latest DDR4 memories with an increase in performance

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
M3.SMALL8	1	8	1.5
M3.MEDIUM16	2	16	1.5
M3.LARGE32	4	32	1.5

at 60%. They support up to 10 Gbps private network bandwidth.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Application Scenarios

This type of instances is suitable for the following circumstances:

- Applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching.
- Users who build their own hadoop clusters or redis for gene computing and so on.

Features

- 2.5 GHz Intel Xeon® Skylake 6133 processors with stable computing performance
- The latest generation of 6-channel DDR4 memories with a bandwidth of 2,666 MT/sec
- Larger instance size, M3.16XLARGE512, offering 64 vCPUs and 512 GB of memory
- 1:8 or 1:12 processor to memory ratio
- Price is lower than any other memory-optimized instances with the same memory
- The network performance of an instance is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.

Requirements

- Both prepaid and postpaid billing methods are available for M3 instances.
- M3 instances can be launched in basic networks and VPCs.

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
M3.2XLARGE64	8	64	1.5
M3.3XLARGE96	12	96	1.5
M3.3XLARGE144	12	144	1.5
M3.4XLARGE128	16	128	2.0
M3.4XLARGE192	16	192	2.0
M3.8XLARGE256	32	256	4.0
M3.8XLARGE384	32	384	4.0
M3.16XLARGE512	64	512	8.0

- Configuration purchase is available for M3 instances. Please see the instance specifications on the right.

[Back to the top](#)

Memory Optimized M2

Memory Optimized M2 instances are designed to deliver high performance for handling the workloads of large data sets in memory. They feature large memory, which makes them the perfect choice for applications demanding high memory computing.

Application Scenarios

This type of instances is suitable for the following circumstances: - Applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching. - Users who build their own hadoop clusters or redis for gene computing and so on.

Features

- 2.4 GHz Intel Xeon® E5-2680v4 processors, DDR4 memories
- Up to 448 GB memory available. M3.16XLARGE512, offering 56 vCPUs and 448 GB of memory
- 1:8 processor to memory ratio
- Price is lower than any other memory-optimized instances with the same memory

Requirements

- M2 instances support both prepaid and postpaid billing methods, and can also be used as production instances of memory optimized host HM20 in CDHs.
- M2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for M2 instances. Please see the instance specifications on the right.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
M2.SMALL8	1	8	1.5
M2.MEDIUM16	2	16	1.5
M2.LARGE32	4	32	1.5
M2.2XLARGE64	8	64	1.5
M2.3XLARGE96	12	96	2.5
M2.4XLARGE128	16	128	3.0
M2.6XLARGE192	24	192	4.5
M2.8XLARGE256	32	256	6.0
M2.12XLARGE384	48	384	9.0
M2.14XLARGE448	56	448	10.0

Memory Optimized M1

Memory Optimized M1 instances with about 1:8 CPU to memory ratio are applicable to applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching.

Features

- 2.3 GHz Intel Xeon® E5-2670 v3 processors and DDR3 memories, providing larger-sized instances with stronger computing capacities.
- Memory Intensive with a golden ratio, satisfying the needs for large-scale business deployment.

Application Scenarios

This type of instances is suitable for the following circumstances:

- Applications that require extensive memory operations, searches, and computations, such as high-performance databases and distributed memory caching.
- Users who build their own hadoop clusters or redis for gene computing and so on.

Requirements

- Both prepaid and postpaid billing methods are available for M1 instances.
- M1 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for M1 instances. Please see the instance specifications on the right.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
M1.SMALL8	1	8	1.5
M1.MEDIUM16	2	16	1.5
M1.LARGE32	4	32	1.5
M1.2XLARGE64	8	64	2.0
M1.3XLARGE96	12	96	2.5
M1.4XLARGE128	16	128	3.5
M1.6XLARGE192	24	192	5.0
M1.8XLARGE256	32	256	7.0
M1.12XLARGE368 Model	48 vCPU (core)	368 Memory (GB)	Private 10.0 network bandwidth (Gbps)

High I/O Instance Family

It features high random IOPS, high throughput, low latency, etc. and is suitable for I/O-intensive applications that require fast hard disk read/write and low latency, such as high-performance databases.

High I/O I2

High I/O I2 instances are optimized to provide tens of thousands of low-latency random I/O operations per second (IOPS) to applications, which are the ideal choice for high IOPS scenarios.

High I/O I2 instances come with Intel® Broadwell processors that bring 40% better integer and floating-point computing performance and DDR4 memories that perform 30% better.

Application Scenarios

- 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

- 2.4 GHz Intel Xeon E5-2680 Broadwell (v4) processors, DDR4 memories
- CPU performance is 20% higher than Series 1 High IO I1
- SSD is used for instance storage, and SSD local disk is adopted for all system disks
 - High random IOPS, with up to 75,000 random read IOPS (blocksize = 4k, iodepth = 32) and up to 10,000 random write IOPS (blocksize = 4k, iodepth = 32) in typical scenarios.
 - High throughput, with up to 250 MB/sec random read throughput (blocksize =4k, iodepth =32) in typical scenarios.
 - Low latency, with the access latency in sub-milliseconds (blocksize = 4k, iodepth = 1) in typical scenarios.

Requirements

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
I2.MEDIUM4	2	4	1.5
I2.MEDIUM8	2	8	1.5
I2.LARGE8	4	8	1.5
I2.LARGE16	4	16	1.5
I2.2XLARGE16	8	16	1.5
I2.2XLARGE32	8	32	1.5
I2.3XLARGE24	12	24	2.5
I2.3XLARGE48	12	48	2.5
I2.4XLARGE32	16	32	3.0
I2.4XLARGE64	16	64	3.0
I2.6XLARGE48	24	48	4.5
I2.6XLARGE96	24	96	4.5
I2.8XLARGE64	32	64	6.0
I2.8XLARGE128	32	128	6.0
I2.14XLARGE224	56	224	10.0

- High IO I2 instances support both prepaid and postpaid billing methods, and can also be used as production instances of high IO host HI20 in CDHs.
- I2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for I2 instances. Please see the instance specifications on the right.

[Back to the top](#)

High I/O I1

High I/O I1 is a virtual machine mounted with a high-performance SSD local disk, which can meet the high requirements for fast disk read/write and low latency.

Features

- High random IOPS, with up to 40,000 random read IOPS in typical scenarios (blocksize =4k, iodepth =32);
- High throughput, with up to 250 MB/sec random read throughput (blocksize =4k, iodepth =32) in typical scenarios.
- Low latency, with the access latency in sub-milliseconds.

Application Scenarios

- 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.

Requirements

- Both prepaid and postpaid billing modes are available for High I/O I1 instances.
- I1 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for I1 instances. Please see the instance specifications on the right.

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
I1.MEDIUM4	2	4	1.5
I1.MEDIUM8	2	8	1.5
I1.MEDIUM16	2	16	1.5
I1.LARGE8	4	8	1.5
I1.LARGE16	4	16	1.5
I1.2XLARGE16	8	16	2.0
I1.2XLARGE32	8	32	2.0
I1.3XLARGE24	12	24	2.5
I1.3XLARGE48	12	48	2.5
I1.4XLARGE32	16	32	3.5
I1.4XLARGE64	16	64	3.5
I1.6XLARGE48	24	48	5.0
I1.6XLARGE96	24	96	5.0
I1.8XLARGE64	32	64	7.0

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
I1.8XLARGE128	32	128	7.0
I1.12XLARGE192	48	192	10.0

Big Data Instance Family

The big data instance 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.

Big Data D1

The Big Data D1 instance is equipped with massive storage resources, can carry up to 48 TB SATA HDD local storage, and is suitable for throughput-intensive services, such as Hadoop distributed computing and parallel data processing.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Application Scenarios

- Distributed computing services such as Hadoop MapReduce/HDFS/Hive/HBase.
- Business scenarios such as Elasticsearch, log processing, and large data warehouse.
- Customers in the Internet, finance, and other industries who require big data computing and storage analysis, as well as business scenarios of massive data storage and computing.

Features

- 2.4 GHz Intel Xeon E5-2680v4 processors, DDR4 memories
- Equipped with up to 48 TB local HDD storage

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
D1.2XLARGE32	8	32	1.5
D1.4XLARGE64	16	64	3.0
D1.6XLARGE96	24	96	4.5
D1.8XLARGE128	32	128	6.0
D1.14XLARGE224	56	224	10.0

- For a single disk, the sequential read throughput is 190+ MB/s, and sequential write throughput is 190+ MB/s (128 KB of block size and depth of 32)
- For the machine, the throughput can reach up to 2.3 GB/s (128 KB of block size and depth of 32).
- The read/write delay is as low as 2-5 ms.
- The unit price of local storage is as low as 1/10 of that of S2. The total cost is similar to that of Hadoop clusters built in IDC.
- The 1:4 processor to memory ratio is tailored for the big data scenarios.

Requirements

- Big Data D1 instance uses local disk as its data disk, so there is a risk of data loss (for example, when host crashes). If your application does not have a data reliability architecture, it is highly recommended that you select an instance, which can use a cloud disk as its data disk.
 - After the local disk is damaged, you need to shut down the CVM instance before we can change the local disk.
 - If the CVM instance has crashed, we will inform you and make repairs.
- Both prepaid and postpaid billing modes are available for Big Data D1 instances.
- D1 instances can be launched in basic networks and VPCs.
- D1 instance does not support the adjustment of configuration.
- Configuration purchase is available for D1 instances. Please see the instance specifications on the right.

[Back to the top](#)

Compute Instance Family

This family comes with a base clock rate of 3.2 GHz to provide the highest single-core computing performance. It is suitable for compute-intensive applications such as batch processing, high performance computing, and large game servers.

Computing CN3 Instance

As the latest generation of computing instances, the Computing CN3 instances come with 25 GB ENI that is 2.5 times faster than that of normal computing instances, featuring greater bandwidth and lower latency. Computing CN3 instances provide the processor with the maximum base clock rate and the best cost performance among CVMs. It is an ideal choice for applications subject to computing, such as high computing performance and high concurrent read and write.

Computing CN3 instances are equipped with the new Skylake Xeon® processors, and support up to 10 Gbps private network bandwidth.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Application Scenarios

These instances are an ideal choice for:

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

Features

- 3.2 GHz Intel Xeon® Skylake 6146 processors with up to 3.6 GHz of turbo frequency
- New Intel Advanced Vector Extension (AVX-512) instruction set
- Support up to 25 Gbps of private network bandwidth to meet extremely high private network transmission requirements.
- The latest generation of 6-channel DDR4 memories with a memory bandwidth of 2,666 MT/s

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
CN3.LARGE8	4	8	3.0
CN3.LARGE16	4	16	3.0
CN3.2XLARGE16	8	16	5.0
CN3.2XLARGE32	8	32	5.0
CN3.4XLARGE32	16	32	9.0
CN3.4XLARGE64	16	64	9.0
CN3.8XLARGE64	32	64	17.0
CN3.8XLARGE128	32	128	17.0

- 1:2 or 1:4 processor to memory ratio.
- The network performance of an instance is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.
- Support all kinds of cloud disks

Requirements

- Both prepaid and postpaid billing methods are available for CN3 instances.
- CN3 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for CN3 instances. Please see the instance specifications on the right.

[Back to the top](#)

Computing C3 Instance

As the latest generation of computing instances, Computing CN3 instances provide the processor with the maximum base clock rate and the best cost performance among CVMs. It is an ideal choice for applications subject to computing, such as high computing performance and high concurrent read and write.

Computing C3 instances are equipped with the new Skylake Xeon® processors that perform 30% better and the latest DDR4 memories that perform 60% better than Computing C2 instances. Computing C3 instances support up to 10 Gbps private network bandwidth.

Application Scenarios

These instances are an ideal choice for:

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

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
C3.LARGE8	4	8	2.5
C3.LARGE16	4	16	2.5
C3.LARGE32	4	32	2.5
C3.2XLARGE16	8	16	3.0
C3.2XLARGE32	8	32	3.0
C3.4XLARGE32	16	32	4.5
C3.4XLARGE64	16	64	4.5
C3.8XLARGE64	32	64	8.0
C3.8XLARGE128	32	128	8.0

Features

- 3.2 GHz Intel Xeon® Skylake 6146 processors with up to 3.6 GHz of turbo frequency
- New Intel Advanced Vector Extension (AVX-512) instruction set
- The latest generation of 6-channel DDR4 memories with a memory bandwidth of 2,666 MT/s
- 1:2 or 1:4 processor to memory ratio.
- The network performance of an instance is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.
- Support all kinds of cloud disks

Requirements

- Both prepaid and postpaid billing methods are available for C3 instances.
- C3 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for C3 instances. Please see the instance specifications on the right.

[Back to the top](#)

Computing Network Enhanced CN2

Computing Network Enhanced CN2 instances come with 25 GB ENI, providing a network performance that is 2.5 times faster than that of normal computing instances, featuring greater bandwidth, lower latency, and super-high clock rate. They are suitable for scenarios that require high computing resources and sending and receiving massive network packets.

Features

- 3.2 GHz Intel Xeon® E5-2667v4 processors with up to 3.6 GHz of turbo frequency; DDR4 memories

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
CN2.4XLARGE32	16	32	13.0
CN2.4XLARGE64	16	64	13.0
CN2.8XLARGE96	32	96	25.0

- 1:2 or 1:4 processor to memory ratio.
- Support up to 25 Gbps of private network bandwidth to meet extremely high private network transmission requirements.
- Support forwarding up to 700,000 network packets per second to allow for a greater number of concurrent users.
- Support storage options of SSD local disks, HDD cloud disks and SSD cloud disks.

Application Scenarios

- Scenarios that require sending and receiving massive network packets, such as game services, video services, and financial analysis
- Batch processing workloads
- High-traffic Web server, massively multiplayer online (MMO) game server
- High performance computing (HPC)

Requirements

- The prepaid billing method is available for CN2 instances.
- CN2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for CN2 instances. Please see the instance specifications on the right.

[Back to the top](#)

Computing C2 Instance

Computing CN3 instances provide the processor with the highest performance and the best cost performance among CVMs. It is an ideal choice for applications subject to computing, such as high computing performance and high concurrent read and write.

Application Scenarios

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
C2.LARGE8	4	8	2.5
C2.LARGE16	4	16	2.5

These instances are an ideal choice for:

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

Features

- 3.2 GHz Intel Xeon® E5-2667v4 processors with up to 3.6 GHz of turbo frequency; DDR4 memories
- 1:2 or 1:4 processor to memory ratio.
- The network performance of an instance is determined by its specification. The higher the specification is, the greater the network forwarding performance and the higher the maximum private network bandwidth limit will be.
- Support all kinds of cloud disks

Requirements

- C2 instances support both prepaid and postpaid billing methods, and can also be used as production instances of computing host HC20 in CDHs.
- C2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for C2 instances. Please see the instance specifications on the right.

[Back to the top](#)

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
C2.LARGE32	4	32	2.5
C2.2XLARGE16	8	16	3.5
C2.2XLARGE32	8	32	3.5
C2.4XLARGE32	16	32	6.0
C2.4XLARGE60	16	60	6.0
C2.4XLARGE64	16	64	6.0
C2.8XLARGE64	32	64	10
C2.8XLARGE96	32	96	10
C2.8XLARGE120	32	120	10

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.

GPU Computing GN8

GPU Computing GN8 uses high-performance NVIDIA Tesla P40 and is applicable to generic GPU computing applications with CUDA and OpenCL programming models. The peak computing capacity of a single CVM exceeds 96 TFLOPS in single-precision floating point compute and 376 TOPS in INT8. Based on top-notch TFLOPS and INT8 performance, GPU Computing GN8 is the most cost-effective AI training and inference solution that perfectly meeting the needs of one-stop deep learning training and real-time inference.

Application Scenarios

These instances are an ideal choice for:

- Deep learning such as image classification and recognition, speech recognition, and natural language processing
- Scientific computing including computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other server-side GPU computing workloads.

Features

- NVIDIA Tesla P40 GPU compute cards
- 2.4GHz Intel Xeon E5-2680v4 processors, DDR4 memories
- The peak computing capacity of a single CVM: over 96 TFLOPS (single-precision floating point compute); over 376 TOPS (INT8)

Requirements

- Both prepaid and postpaid billing methods are available for GN8 instances.
- GN8 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for GN8 instances. Please see the instance specifications on the right.
- The configuration of GN8 instances cannot be changed.

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
GN8.LARGE56	6	56	1.5
GN8.3XLARGE112	14	112	2.5
GN8.7XLARGE224	28	224	5.0
GN8.14XLARGE448	56	448	10.0

[Back to the top](#)

GPU Computing GN2

GPU Computing GN2 uses high-performance NVIDIA Tesla M40 and is applicable to generic GPU computing applications with CUDA and OpenCL programming models. It offers powerful single/double-precision floating point feature with 6,144 acceleration cores. The single-precision floating point compute capacity reaches 14 TFlops.

Application Scenarios

These instances are an ideal choice for:

- Deep learning such as image classification and recognition, speech recognition, and natural language processing
- Scientific computing including computational fluid dynamics, computational finance, genomics research, environmental analysis, high-performance computing, and other server-side GPU computing workloads.

Features

- NVIDIA Tesla M40 GPU compute cards, with 24GB GDDR5 video memory for each card
- 2.4GHz Intel Xeon E5-2680v4 processors
- Peak computing capacity for single machine: over 14TFlops (single-precision floating point); over 0.4TFlops (double-precision floating point).

Requirements

- Both system disks and data disks use local SSD disks and can be mounted with cloud disks.
- Both prepaid and postpaid billing methods are available for GN2 instances.
- GN2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for GN2 instances. Please see the instance specifications on

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
GN2.7XLARGE56	28	56	5.0
GN2.14XLARGE112	56	112	10.0

the right.

- The configuration of GN2 instances cannot be changed.

[Back to the top](#)

GPU Rendering GA2

GA2 instances are optimized for graphics-intensive applications and are suitable for generic GPU rendering applications. These instances are equipped with the latest AMD S7150 GPUs. A single GPU has 2,048 processor cores and provides a single-precision floating-point compute capacity of 3.77 TFLOPS. With powerful computing capacity and elastic expansion based on your business needs, GA2 instances are optimal for high-performance rendering and computing applications.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Application Scenarios

GA2 Instance is especially suitable for GPU computing workloads requiring high-performance rendering and excellent graphic processing capacities.

- Graphic rendering scenarios such as 3D modeling, rendering, multimedia encoding/decoding, and nonlinear editing
- Business scenarios that require a small amount of virtual GPU resources for optimal graphics performance, such as cloud gaming

Features

- AMD FirePro™ S7150 GPU, with single-precision floating-point compute capacity of up to 3.77 TFlops for a single GPU
- Intel Xeon E5-2680v4 2.5GHz processors; high-speed DDR4 memories

Requirements

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
Model GA2.2XLARGE16	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)

- Both prepaid and postpaid billing methods are available for GA2 instances.
- GA2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for GA2 instances. Please see the instance specifications on the right.

[Back to the top](#)

FPGA FX2

FPGA FX2 is an FPGA-based computing service equipped with a Xilinx KU115 accelerator. It is designed to accelerate compute-intensive algorithms, and achieve high throughput, low latency, and hardware programming. We recommend that you use it for high-performance computing services, such as genomics research, financial analysis, image compression and real-time video processing.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Application Scenarios

Ideal for scenarios that require large amounts of parallel computing and high throughput

- Deep learning and inference scenarios such as natural language processing and image classification
- Scenarios that require a large amount of analysis and computing such as genomics research and financial analysis
- Large-scale image processing scenarios such as image compression and real-time video processing

Features

- Use Xilinx Kintex UltraScale KU115 FPGA
- Intel Xeon E5-2680v4 2.5GHz processors; high-speed DDR4 memories

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
FX2.3XLARGE60	14	60	2.5
FX2.7XLARGE120	28	120	5.0
FX2.14XLARGE240	56	240	10.0

Requirements

- The prepaid billing method is available for FX2 instances.
- FX2 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for FX2 instances. Please see the instance specifications on the right.
- The configuration of FX2 instances cannot be changed.

[Back to the top](#)

Batch-based Instance Family

With the lowest per core-hour cost, this family is suitable for compute-intensive applications that use super large computing nodes frequently in a short time, such as rendering, gene analysis, and crystal pharmacy.

- This instance type is only available to users on the whitelist for now. Contact your pre-sales manager for the permission to purchase these instances.

Batch-based Computing BC1

The batch-based computing BC1 instance is an instance of ultra-high cost performance. It applies the core/hour billing method, which is accurate to seconds. The price is as low as 0.1 CNY/core. It is flexible and ready to use, and can be terminated when you do not need it. It supports a variety of specifications to meet the needs of computing-intensive users who frequently use super large computing nodes in short time, such as rendering, gene analysis, and crystal pharmacy.

Application Scenarios

- Video/film rendering
- Genomics, crystal pharmacy, etc.
- HPC computing-intensive business such as weather forecasting and astronomy

Features

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
BC1.LARGE16	4	16	2.0
BC1.3XLARGE48	12	48	5.0
BC1.6XLARGE96	24	96	10.0

- Cost-effective with the lowest price per hour among all instances with the same specification
- 1:4 processor to memory ratio

Requirements

- Only the postpaid billing method is available for batch-based computing BC1 instances.
- BC1 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for BC1 instances. Please see the instance specifications on the right.

[Back to the top](#)

Batch-based Computing BS1

The batch-based computing BS1 instance is an instance of ultra-high cost performance. It applies the core/hour billing method, which is accurate to seconds. The price is as low as 0.09 CNY/core. It is flexible and ready to use, and can be terminated when you do not need it. It supports a variety of specifications to meet the needs of computing-intensive users who frequently use super large computing nodes in short time, such as rendering, gene analysis, and crystal pharmacy.

Features

- Cost-effective with the lowest price per hour among all instances with the same specification
- 1:2 processor to memory ratio

Application Scenarios

- Video/film rendering
- Genomics, crystal pharmacy, etc.
- HPC computing-intensive business such as weather forecasting and astronomy

Requirements

Model	vCPU (core)	Memory (GB)	Private network bandwidth (Gbps)
BS1.LARGE8	4	8	2.0
BS1.3XLARGE24	12	24	5.0
BS1.6XLARGE48	24	48	10.0

- Only the postpaid billing method is available for batch-based computing BS1 instances.
- BS1 instances can be launched in basic networks and VPCs.
- Configuration purchase is available for BS1 instances. Please see the instance specifications on the right.

[Back to the top](#)

Life Cycle

Last updated : 2018-09-28 16:58:26

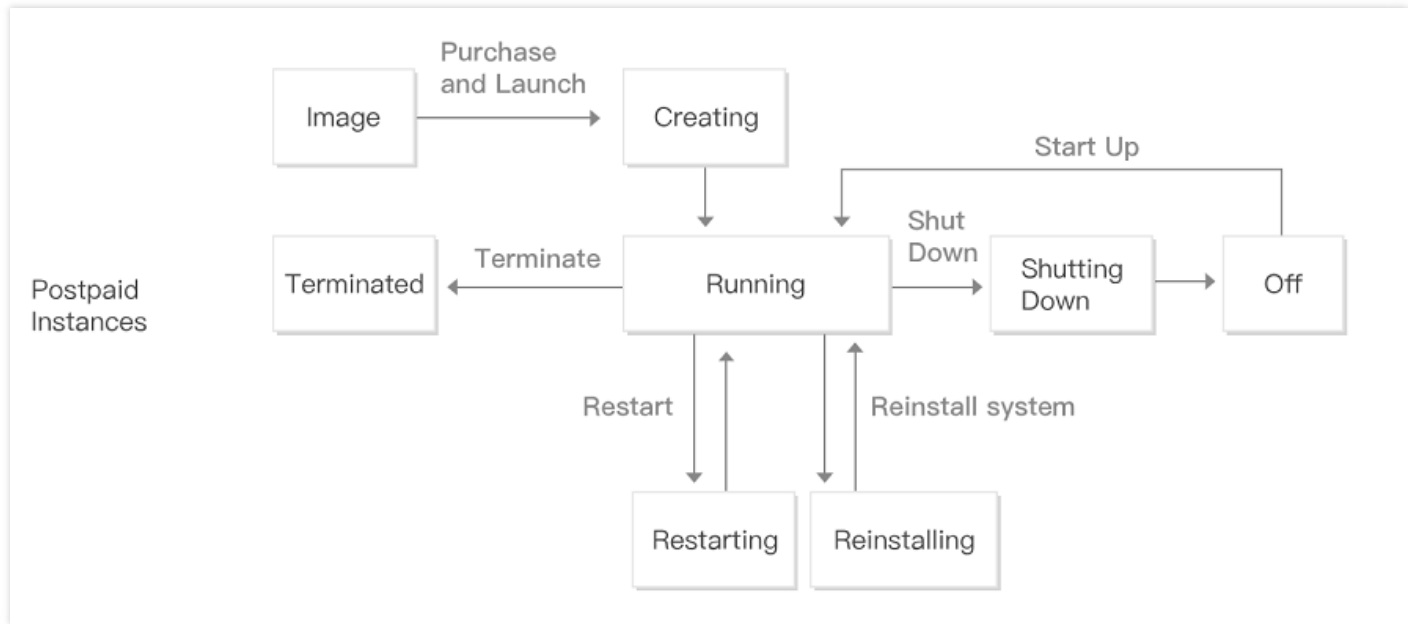
The lifecycle of Tencent Cloud CVM instance refers to all the statuses from the launch of the instance to its termination. Properly managing Tencent Cloud instances in different statuses (from launch to termination) can ensure that the applications running on the instances provide services in an efficient and economical way.

Instance Statuses

- **An instance has the following statuses:**

Status Name	Status Attributes	Description
Creating	Intermediate status	The instance is created, but is not running yet.
Running	Steady status	The instance is running normally. Instances in this status can be used to run your business.
Restarting	Intermediate status	A restart operation is performed for the instance via the console or APIs, but is not running yet. If this status lasts for a long time, there may be an exception.
Resetting	Intermediate	The instance's system is reinstalled or its disk is reset via the console or APIs, but the instance is not running yet.
Shutting down	Intermediate status	A shutdown operation is performed for the instance, but it is not completely shut down yet. If this status lasts for a long time, there may be an exception. Forced shutdown is not recommended.
Shut down	Steady status	The instance is shut down normally. Instances in this status cannot provide external services. Some attributes of instance can only be modified when the instance is in the shutdown status.
Terminating	Intermediate	The user actively performed the termination operation, but the termination has not yet completed.
Terminated	Steady status	The termination operation is completed. The original instance does not exist and cannot provide services. Its data is completely cleared.

- **Instance Status Transition:**



Instance Launching

- After an instance is launched, it enters the creating status. The hardware specifications are configured according to the specified [Instance Specifications](#) for the instance in the creating status, and the system launches the instance using the image specified on launch.
- The instance enters the running status after it is created. An instance in the running status can be connected and accessed normally.

For more information on instance launching, please see [Purchase and Launch an Instance](#), [Log in to Windows Instance](#), and [Log in to Linux Instance](#).

Instance Restarting

You are recommended to restart an instance via Tencent Cloud Console or Tencent Cloud APIs, instead of running the OS restart command in the instance.

- After restarting, it enters restarting status.
- Restarting an instance is like restarting the computer. After restarting, the instance maintains its public IP address, private IP address and all data on the disk.
- Normally, it takes dozens of seconds or even several minutes to restart the instance, depending on the instance configuration.

For more information on instance restarting, please see [Restart Instance](#).

Instance Shutdown

You can shut down the instance via the console or APIs.

- Shutting down an instance is like shutting down the computer.
- An shutdown instance no longer provides external services, but the billing is not stopped.
- The shutdown instance will still be displayed in the console.
- Shutdown is required for some configuration operations, such as adjusting hardware configurations and resetting passwords.
- Shutdown operation does not change the CVM's public IP, private IP, or any data on its disk.

For more information on instance shutdown, please see [Instance Shutdown](#).

Instance Termination

You can terminate an CVM instance if you no longer need it. You can terminate an instance through Tencent Cloud Console or Tencent Cloud APIs.

- Manual termination: you can actively terminate postpaid instances and prepaid instances in the recycle bin through the console.
- Auto termination: postpaid instances cannot be automatically terminated.

When terminating an instance, the system disks and data disks designated when purchasing will be terminated. But elastic cloud disks mounted on it will not be affected.

For more information on instance termination, please see [Terminate Instance](#).

Spot Instance

Last updated : 2018-09-11 11:00:14

What is Spot Instance

Spot instances are a new mode of CVM instances that featured by spot prices at substantial discounts and an interruption mechanism (available at certain discounts and can be automatically interrupted and reclaimed by the system). After you purchase a spot instance, you can use it the same way as you use a postpaid CVM instance, including console operations, remote login, service deployment, associating with VPC, etc.


- Related document: [FAQ About Spot Instances](#)
- Related document: [How to purchase Spot Instances](#)

Special Policies During Public Trial

- Application: During the public trial, you need to submit an application before you can use a spot instance. Tencent Cloud technical engineers will review and reply to your application. [Application Page >>](#)
- Fixed discount: Spot instances of all specifications are available at a fixed discount. Any spot instance is sold at a discount of 20% off compared to the postpaid instance of the same specification. The prices of a small number of instances may be adjusted slightly.
- Interruption by system: The interruption will not be caused by the reason that the market price is higher than the price paid by user, but will be caused by insufficient stock of spot instance pool. In case of insufficient resources, the system will reclaim some allocated spot instances at random.
- Increasing coverage: At the early stage of public trial, spot instances are only available for some availability zones and instance types. More availability zones and instance types will be supported later as the public trial progresses. For the up-to-date list of supported regions and instance types, please see [FAQ About Spot Instances - Supported Regions and Instances>>](#).

Product Features

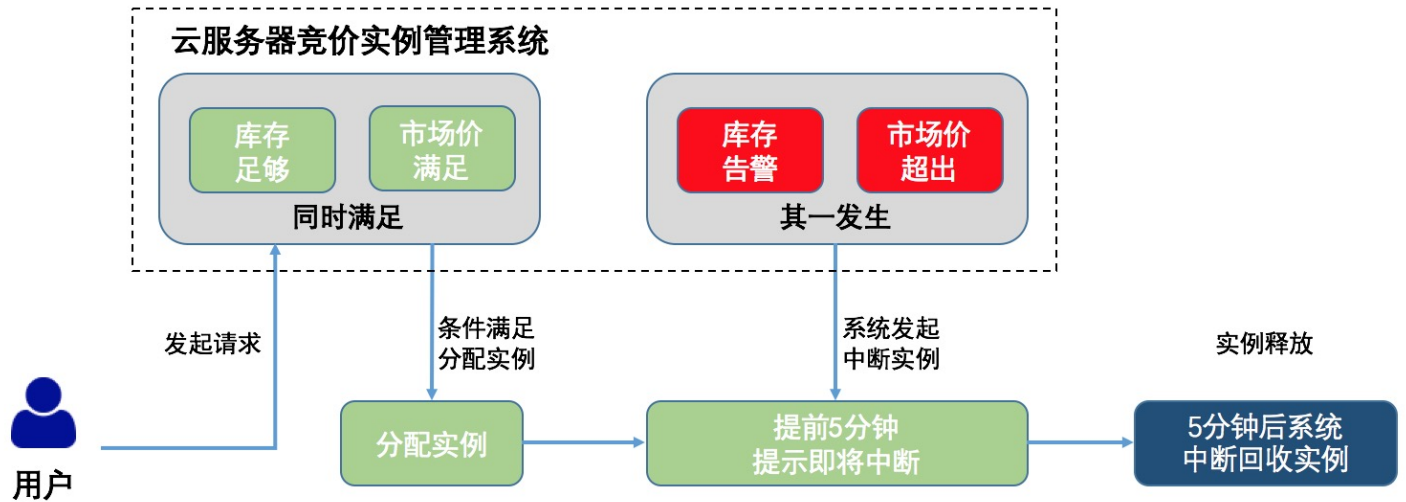
1. Greatly reduced cost

云服务器 (CVM)	广州二区 以S2. 2XLARGE16为例 单位: 元/小时	按量计费	竞价实例 (以最低折扣为例)
	CPU (8 核) 内存 (16 GB)	2.5	0.25
系统盘	系统盘 (50 GB 普通云硬盘)	0.02	0.02
带宽/流量	带宽 (1 Mbps 按带宽计费)	0.06	0.06
收费镜像	收费总计	2.58	0.33



Spot instances are sold at a discount of up to 90% off compared to the postpaid instances.

- **Discount range:** Spot instances are priced at a discount ranging from 10% to 100% off based on the prices of postpaid instances of the same specifications.
- **What are covered by discount:** The discount only involves CVM vCPU and memory. Other CVM-related billed items such as system disk, data disk, bandwidth, and paid images are not affected by the discount.
- **Price fluctuation:** The discount is stable for a period of time, but when there are bulk purchases in an availability zone, the price will fluctuate (during the public trial, the discount is fixed at 80%).

2. Interruption by system



Unlike postpaid and prepaid instances that can only be released by users, spot instances are interrupted by the backend of system, which determines whether to interrupt the allocated spot instances based on the current price and the stock of resource pool.



价格原因回收示意

库存原因回收示意

- **Interruption due to price:** When the market price is higher than your highest bid, the spot instance will be reclaimed (during the public trial, the market price is fixed).
- **Interruption due to insufficient stock:** In case of insufficient stock of resource pool of your spot instance, the system reclaims the spot instance based on the stock shortage; when the pool is fully stocked, you can apply for spot instances again.

Scenarios To Which Spot Instances Do Not Apply

Due to the interruption by system, you cannot bring the lifecycle of a spot instance under your control. Therefore, it is not recommended to run services that require high stability on a spot instance, for example:

- Database services
- Online and website services without load balancer
- Master control nodes in a distributed architecture
- Long-time (over 10 hours) big data computing tasks

Scenarios and Industries To Which Spot Instances Apply

Scenarios

- Big data computing
- Online and website services with load balancer
- Web crawler service
- Other computing scenarios with fine granularity or support for checkpoint restart

Industries

- Gene sequencing and analysis
- Drug crystal form analysis
- Video transcoding and rendering
- Finance and transaction data analysis
- Image and multimedia processing
- Scientific calculations (geography, hydromechanics...)

Limits

- Quota limit: Unlike the quantity quota of non-spot instances, the quota of spot instances depends on the total number of vCPU cores of all spot instances you own in the availability zone. During the public trial period, the number of spot instances in each availability zone under an account should be limited such that the number of vCPU cores is not more than 30 (you can submit a ticket to apply for a bigger quota).
- Operation limit 1: Upgrading/degrading configuration is not supported for spot instances.
- Operation limit 2: Spot instances cannot be changed to prepaid instances.

Best Practices

1. Splitting high-granularity tasks

- Split a long-time task into fine-grained subtasks to reduce the possibility of interruption;
- Use a big data suite like EMR that comes with splitting mechanism.

2. Ensure the stability of online and website services with load balancer

- Use load balancer, such as CLB, for the access layer;

- Use a small proportion of postpaid instances + a large proportion of spot instances for backend resources;
- Listen on the interruptions of spot instances and remove from the CLB the instances that are about to be interrupted.

3. Computing scheduling mode supporting checkpoint restart

- Store the intermediate results of computation on the permanent storage products such as COS/CFS/NAS;
- Keep aware of instances that are about to be interrupted through Metadata and store the computing results for the retention period of 5 minutes;
- When the spot instance is started again, resume the last computation.

Storage Overview

Last updated : 2018-09-28 16:41:21

Tencent Cloud provides a wide range of flexible, economic and user-friendly data storage devices for the CVM instances. Various storage devices are provided to cater for different usage scenarios. The performance and price varies with the category of storage device.

Storage Device Classification

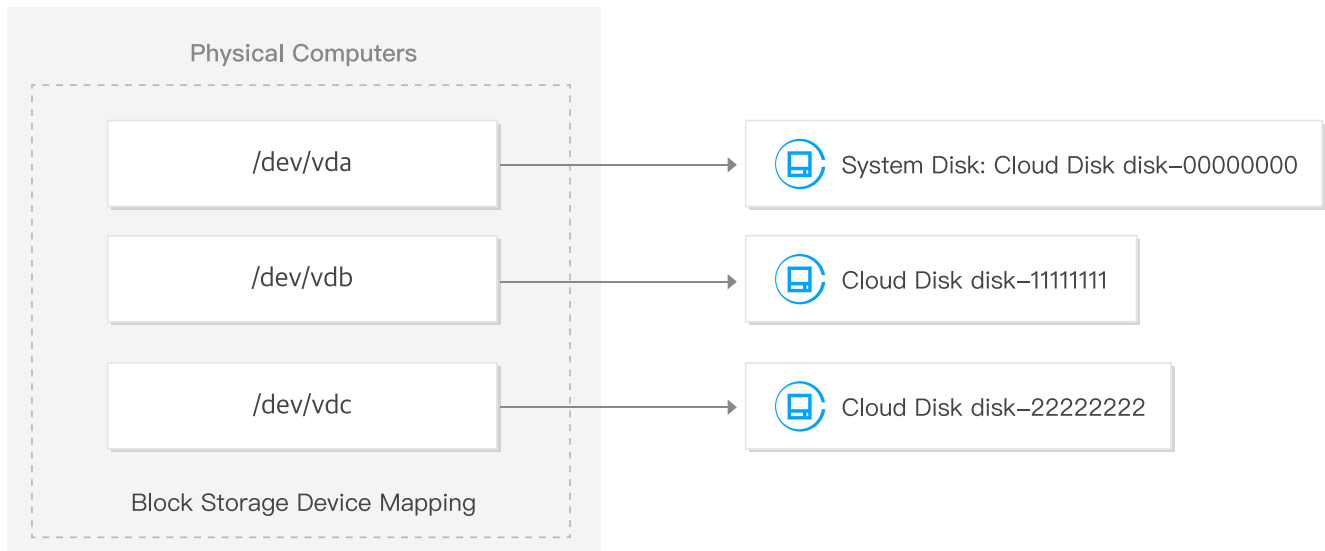
Storage devices can be divided into the following categories by dimensions:

Dimension	Category	Description
Storage medium	HDD disk	Uses mechanical hard disk as the storage medium. It is characterized by a lower price and a good read/write speed.
	SSD disk	Uses Solid State Drive (SSD) as the storage medium. It has an excellent performance in IOPS and read/write speed, with an IOPS and throughput up to 20 times and 16 times higher than those of HDD disk, respectively. It is more expensive than HDD disk.
Application scenarios	System disk	Used to store the collection of systems that control and schedule the operation of CVM. It is operated by using images.
	Data disk	Used to store all the user data.
Architectural pattern	Cloud storage	Cloud storage is an elastic, highly available, highly reliable low-cost and customizable network block device that can be used as a standalone scalable hard disk for CVM. It provides data storage at data block level and employs a 3-copy distributed mechanism to ensure the data reliability for CVM. For a CVM with cloud storage, adjustments of hardware, disks and network are allowed.
	Local disk	It is a local storage reserved on the physical machine where the CVM resides in. It allows a data access with a low latency, but involves a risk of single point failure. For a CVM with local disk, you cannot upgrade hardware (CPU, memory and disk), and can only upgrade bandwidth.
	COS	COS is a data storage device on the Internet. It allows data retrieval from any location on CVM instance or the Internet, thus reducing the storage cost. It is not suitable to be used as a storage medium in the low-latency and high-IO scenarios.

Block Storage Device Mapping

Each instance has a system disk to keep the basic operation data. More data disks can be mounted to an instance. An instance uses block storage device-mapping to map the storage devices to locations it can identify.

Block storage is a storage device that puts data into blocks in bytes and allows random access. Tencent Cloud supports two types of block storage devices: local disk and cloud disk.



This figure shows how CBS maps the block storage device to the CVM and maps `/dev/vda` to the system disk, and how it maps the two data disks to `/dev/vdb` and `/dev/vdc` respectively.

The CVM instance can automatically create block storage device mapping for the local disks and cloud storages mounted to it.

Cloud Block Storage

Last updated : 2017-12-15 11:37:36

Tencent Cloud's Cloud Block Storage provides storage at data block level for CVM instances. As a storage medium with high availability and high reliability, Tencent Cloud's Cloud Block Storage comes with a variety of hard drives to cater to the diversified needs of users in read/write operations. For more information about Cloud Block Storage, see [Cloud Block Storage product documentation](#).

In case of frequent data changes that require faster read/write speed and persistent storage, you're recommended to use Tencent Cloud's Cloud Block Storage. Cloud disks can be mounted to any running instance in the same availability zone, just like an ordinary mobile hard disk. It is particularly suitable to be used as such application like instance's file system and database storage, with data existing independently of instance's life cycle.

For more information, see:

- [Technical architecture of Tencent Cloud CBS](#)
- [Price of Tencent Cloud CBS](#)

Local Storage

Last updated : 2018-10-10 11:37:15

Overview

Local disk is a storage device located on the same physical server as the CVM instance and features high read/write IO and low latency.

It is a local storage reserved on the physical machine where the CVM resides in. You can choose local disks as system disks and data disks for most of Tencent instances.

- **Lifecycle:** Local disk is created with the creation of CVM instance, so its life cycle is same as that of CVM.
- **Purchase:** Local disk can only be enabled with the enabling of CVM, which means you can specify a local disk only when purchasing a CVM instance. For more information about purchasing a CVM, please see [Purchase and Enable an Instance](#).

Note:

For a CVM with local disk, you cannot upgrade hardware (CPU and memory), and can only upgrade bandwidth.

Type

Local disk is a local storages on the physical machine where the CVM resides in, and can be classified into local HDDs and SSDs according to their media.

Local HDD

Type	Available specification	Performance	Price
System disk	Always 50 GB and unchangeable	Peak throughput: 40 - 100 MB/s. IOPS: hundreds - 1000.	Prepaid: 0.3 CNY/GB/month Postpaid: 0.00042 CNY/GB/hour
Data disk	A capacity ranging from 10 GB to 1600 GB (in 10 GB increments) is supported. The maximum capacity supported varies with hardware configuration.		

Local SSD

Local SSD is a local storage on the physical machine where the CVM resides in. It provides instances with full SSD block-level data access capability with a low latency, high random IOPS, and high I/O throughput.

Type	Available specification	Performance	Price
System disk	Always 50 GB and unchangeable	Peak throughput: 250 MB/s Its IOPS for random write reaches up to 10000	Prepaid: 0.8 CNY/GB/month Postpaid:

Data disk	A capacity ranging from 10 GB to 7000 GB (in 10 GB increments) is supported. The maximum capacity supported varies with hardware configuration.	(blocksize = 4k, iodepth = 32) Its IOPS for random read reaches up to 75000 (blocksize = 4k, iodepth = 32) Access latency is less than 3 ms.	0.0033 CNY/GB/hour
-----------	---	--	--------------------

Local SSD is suitable for the following scenarios:

- Low latency: Access latency within microseconds.
- Distributed application: NoSQL, MPP data warehouse, distributed file system and other I/O intensive applications. These applications themselves have distributed data redundancy.
- Logs for large online applications: Large online applications produce a large amount of log data, which require high-performance storage with less demand on storage reliability.
- Single point of failure (SPOF) risk: If SPOF risk exists, it is recommended to implement data redundancy at the application layer to ensure data availability.

Cloud Object Storage (COS)

Last updated : 2018-03-22 13:20:33

Tencent Cloud's object storage COS is a data storage device located on the Internet and provides users with a fast and reliable data storage solution for a low price. Users can store and retrieve data via any location on CVM instance or the Internet. COS stores user data redundantly across multiple regions and allows multiple clients or application threads to perform read/write operations on those data at the same time. The data on the COS can be retrieved in a similar way to Internet domain names and be accessed by using a HTTP URL address.

For more information about Tencent Cloud Object Storage, see [COS product documentation](#).

Image

Overview

Last updated : 2017-11-21 16:48:36

What are images?

Tencent Cloud images provide all the information needed to start CVM instances. After you specify an image, you can start any number of instances from it. You can also start instances from any number of images.

The major roles of images are:

- Batch deployment of software environment
By creating an image for the CVM instance with its software environment deployed and then using the image in batch-creating CVM instances, you can make these CVMs instances created later have the same software environment as the original one, therefore to achieve batch deployment of software environment.
- As a backup of CVM operating environment
After an image is created for a CVM instance, you can use the image to recover the CVM instance if its software environment is corrupted and cannot run normal during subsequent use.

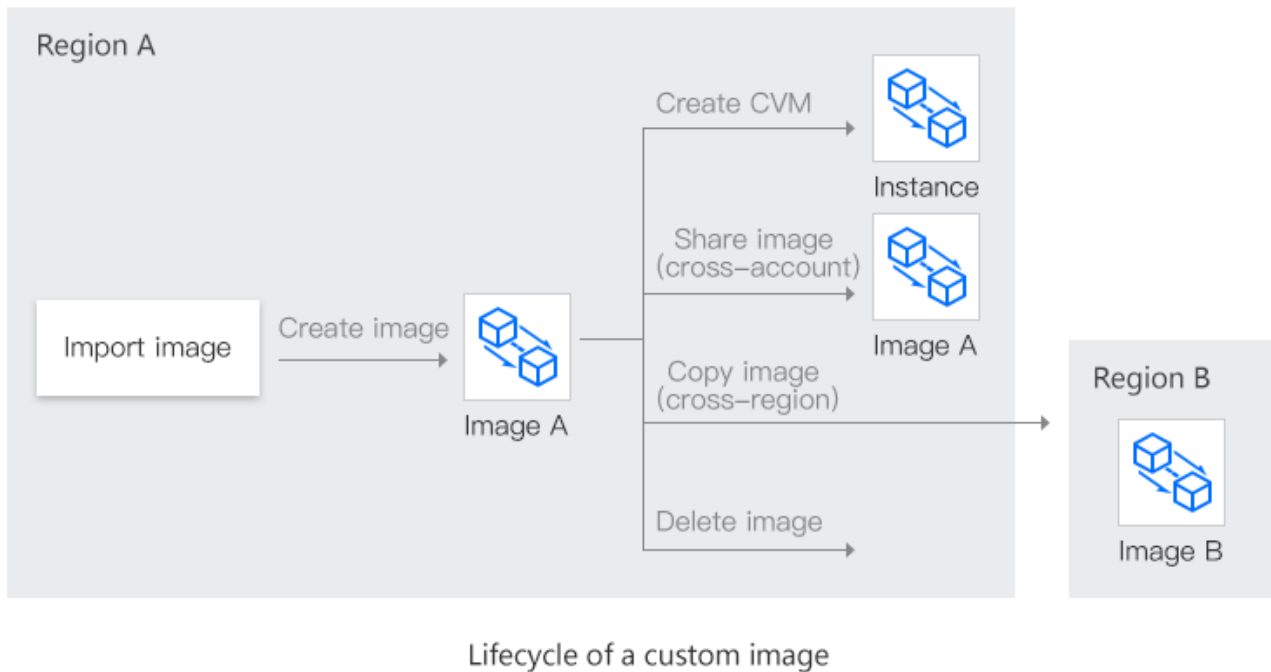
Image types

Tencent Cloud provides the following types of images:

- Common images: All users can use them, covering most of the mainstream operating systems.
- Service market images: All users can use them, which integrate some special applications in addition to the operating system.
- Custom images: Only the creators and shared objects can use them, which are created from existing running instances, or imported externally.
- Shared images: Shared by other users and used only to create instances.

Image lifecycle

The following figure summarizes the lifecycle of custom images. After you create or import a new custom image, you can use it to start a new instance (you can also start an instance from an existing common image or a service market image). A custom image can be synchronized to the same accounts in another region, and becomes an independent image in that region. You can also share custom images to other users.



Usage restrictions

Common images: No usage restrictions.

Custom images:

- 1) Cross-regional synchronization of custom images is free.
- 2) Each region supports 10 custom images.
- 3) Each user can only carry out five operations every day (including cloud API operations, counted by region).
- 4) Image synchronization is available only in Guangzhou, Beijing, Shanghai, Hong Kong, but not in the North American region.

Shared images:

- 1) Each custom image can be shared up to 50 Tencent Cloud master accounts.
- 2) Shared images do not support modifying names and description, and can only be used to create CVMs.
- 3) Sharing images to the same region as the shared account is supported.

Service market images: No usage restrictions.

Charging standards

Currently, all Tencent Cloud common images are free of charge except that Windows images in overseas regions require some license fees. Service market images are priced by the providers. When you start a CVM instance using an image, the costs consist of two parts: CVM instance costs and image costs. In addition, since your custom images are based on some standard images, they may also incur costs. Specific prices and charge modes will be published on the product page.

Security

Tencent Cloud common images are manufactured after rigorous testing by the professional security O&M team of Tencent Cloud. We also provide optional built-in Tencent Cloud security components, so you can use them confidently.

Service market images are manufactured by third-party providers, who have been selected carefully by Tencent Cloud and have signed settlement agreements. Those images have been tested rigorously by the providers and reviewed by Tencent Cloud, to guarantee the security of their contents, so you can also use them.

Images shared by other users may be subject to security risks because they are not reviewed by Tencent Cloud. Therefore, it is strongly recommended that you do not accept images from unknown sources.

Image Types

Last updated : 2018-08-09 16:19:47

You can select an image based on the following attributes:

- Location (please see [Region and Availability zone](#)).
- Operating System
- Architecture (32-bit or 64-bit)

Based on different sources, images provided by Tencent Cloud are divided into public images, custom images, shared images, and service market images.

Public Images

Public images are provided, supported and maintained by Tencent Cloud. Each image is composed of a basic operating system and initialization components from Tencent, and is available to all users.

Features:

- **Operating system:** You are free to choose a Linux-based or Windows-based operating system, and update it regularly.
- **Software support:** Public images are integrated with software packages (such as APIs) provided by Tencent Cloud, and support multiple versions of common software with full permissions, such as Java, MySQL, SQL Server, Python, Ruby, Tomcat.
- **Security:** The operating system provided officially is completely legitimate and compliant. The images are made by our professional security and OPS teams. They are strictly tested with optional built-in Tencent Cloud security components.
- **Limit:** None.
- **Fees:** All public images are free of charge, except for Windows images in some oversea regions which require a certain amount of license fee.

Service Marketplace Images

Service marketplace images are provided by third-party service providers. They are published and made available to all users after having gone through an audit of Tencent Cloud.

Features:

- **Operating system:** Pre-installed environments including multiple conventional operating systems.
- **Software support:** Integrated with common software environments, such as PHP, Java, FTP, Nginx, Docker, WordPress, and Discuz.
- **Security:** All providers have gone through our strict audit and entered into an agreement with us. The images are strictly tested by the providers to ensure the security of image content.
- **Limit:** None.
- **Fees:** The prices of images are determined by their providers. Some images are free, and some are fee-based. When you use a fee-based image to launch an instance, both instance and image will incur costs.

Custom Images

Custom images are made by users with image creation feature, or imported via image import feature. They are only available to their creators and sharers.

Features:

- **Application scenario:** An image created out of an CVM instance with deployed applications can be used to quickly create more instances that contain the same configuration.
- **Supported features:** You can create, copy, share and terminate images.
- **Limit:** Each region supports a maximum of 10 custom images.
- **Fees:** Creation of images may incur a fee. Refer to the prices displayed on the instance creation page for the actual prices. Copying custom images across regions is free of charge.

For more operations and limits, please see [Create Custom Images](#), [Copy Custom Images](#), [Share Custom Images](#), [Cancel Custom Image Sharing](#), and [Import Custom Images](#).

Shared Images

Shared images are custom images shared by other Tencent Cloud users to the current user via the image sharing feature. These images are displayed in the same region as the original images of the users with whom they are shared.

Features:

- **Application scenario:** A shared image can be used to quickly create a CVM.
- **Supported features:** A shared image can only be used to create CVMs. You cannot modify its name, copy it, share it with other users, or perform any operations on it.
- **Security:** Shared images are not reviewed by Tencent Cloud, and thus may pose security risks. Therefore, we strongly recommend against using images from unknown sources.
- **Limit:** Each custom image can be shared to a maximum of 50 Tencent Cloud users. An image can only be shared to the accounts in the same region as the source account.

For more operations and limits, please see [Share Custom Images](#) and [Cancel Custom Image Sharing](#).

Network and Security Overview

Last updated : 2018-05-04 15:07:24

Tencent Cloud provides the following network and security features:

- Security group
- Encryption login method
- Elastic IP
- Internet access
- Private network access
- Basic network and private network

You can use [Security Group](#) to control access to your instance. These security groups resemble a network firewall, allowing you to specify the protocols, ports, and source/target IP ranges that are allowed access. You can create multiple security groups and assign different rules to each security group. You can then assign one or more security groups to each instance, and we will use these rules to determine what traffic is allowed to access instances and which resources the instance can access. You can configure a security group so that only a specific IP address or a specific security group can access the instance.

Tencent Cloud provides two encryption login methods: [Password Login](#) and [SSH Key Pair Login](#). Users are free to choose two ways to securely connect with the CVM.

Instances may fail because of uncontrollable reasons. If an instance fails and you start a replacement instance, the public IP of the alternate instance will be different from the original instance. However, if your application requires a static IP address, you can use an [Elastic IP Address](#).

Tencent's [Internet Link](#) gives access to more than 20 domestic mainstream network operators to ensure that your customers, regardless of ISP, can enjoy the same high-speed access; [Private Network Link](#) goes through an underlying 10 Gigabit / Gigabit network interoperability to ensure high-speed access, high reliability and low latency.

The user's [Network Environment](#) can be roughly divided into 'basic network' and 'private network'. Under a basic network, your cloud product instance is located in a large resource pool preset by Tencent Cloud; under a private network, your cloud product instance can be activated under your own preset, custom network segments, and isolated from other users.

The above network and security services protect your instances; making them safe, efficient and able to freely provide external services.

Placement Group

Last updated : 2018-09-30 12:06:50

You can launch instances in a placement group which determines how the instances are placed on the underlying hardware. When creating a placement group, you can distribute the instances in the group to different underlying hardware.

This feature is under internal trial. Please click [here](#) to apply for a trial use.

Spread Placement Groups

A spread placement group contains a set of instances that are placed on different underlying hardware.

For applications of important instances that need to be placed separately, such as master/slave databases and high-availability clusters, a spread placement group is recommended. By launching instances in a spread placement group, you can reduce the risk of simultaneous failure that occurs when the instances are placed on the same underlying hardware.

A spread placement group has geographical attributes and can be deployed across multiple availability zones. There is a limit on the number of instances in a group. For more information, please see the Console page on the official website.

When you launch instances in a spread placement group, the request may fail in case of insufficient hardware. We provide more different hardware so that you can retry your request.

Spread Placement Group Rules and Limits

Before using a spread placement group, you should pay attention to the following rules:

1. Placement groups cannot be merged.
2. An instance can be launched in a placement group at a time.
3. An instance cannot be placed across multiple placement groups.
4. Existing instances cannot be automatically added into a placement group.
5. Spread placement layers can be selected, including physical machine, exchange, and rack.
6. The maximum numbers of instances in groups on different placement layers are different. For more information, please see the official website.
7. The disaster recovery group policy you have specified is strictly complied with. Please note that if there is not enough hardware to distribute instances, the creation of some instances will fail.
8. Instances on CDH don't support spread placement groups.

Operation Instructions

For more information on how to work with spread placement groups, please see [Spread Placement Group](#).

Network Environment

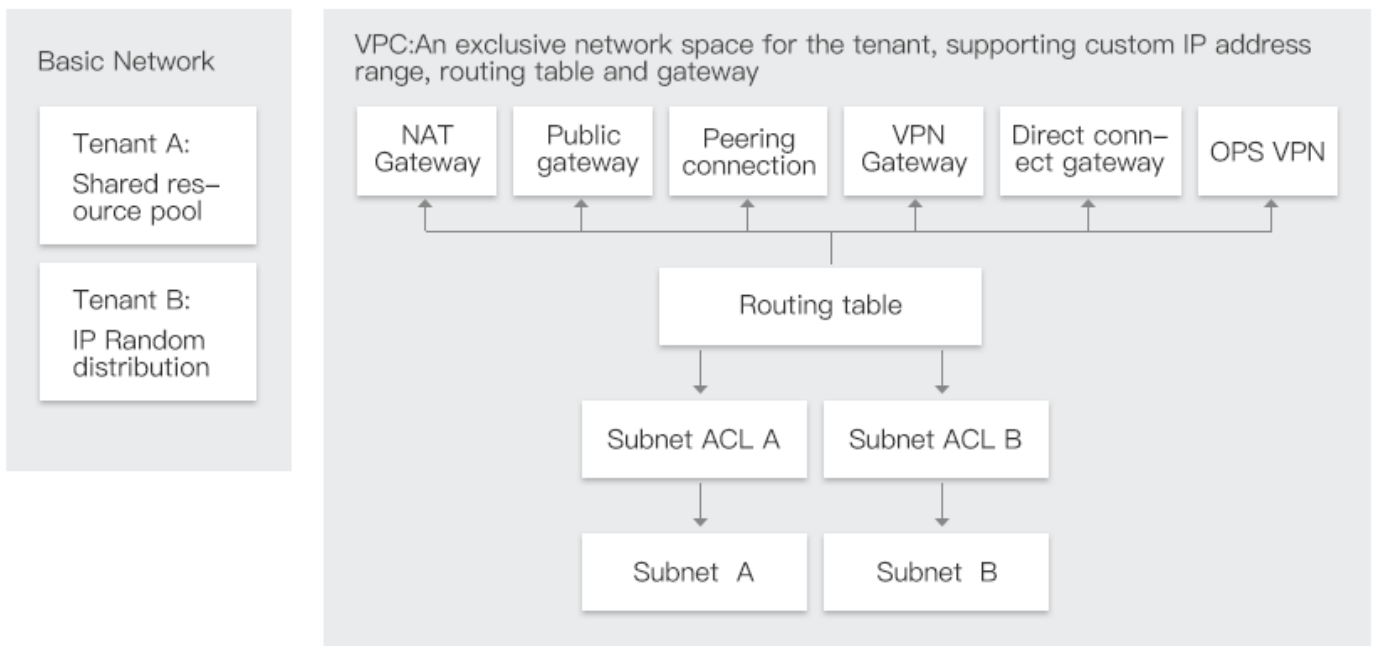
Last updated : 2017-11-21 16:51:03

Tencent Cloud network environment can be divided into either basic network and private network (VPC).

Through the Tencent Cloud basic network, all of the user's resources on the cloud are managed uniformly by Tencent Cloud; relatively speaking, these configurations are simpler and more convenient to use, helping users manage their CVMs faster and easier. All basic network instances in the same geographical area are free to communicate via [Private Network Service](#) under the user account. Most of the user's needs can be met through the basic network and implementation; this is best and most convenient if you are just beginning to understand and use Tencent Cloud.

With the Tencent Cloud Private Network (VPC), you can customize a logical isolated virtual network within the cloud and launch a CVM resource (such as an instance) into the quarantine. Even in the same area, different VPCs cannot communicate with each other by default. VPC is very similar to traditional networks that data centers run, but at the same time, can offer you faster and more extensible infrastructures on the cloud. Users can customize network topology and IP addresses within the network and configure the [Router Table](#), gateway and security settings; supports [Dedicated Connection](#) through your local data center, and rapid expansion of computing resources; freely plan how your VPC communicates with the Internet. Use a variety of control methods (including [Security Group](#) and [Network ACL](#) to protect VPC resources. For more information, see [Private Network Product Documentation](#). Private networks can help users build more complex network architectures, suitable for users who are familiar with network management.

Basic network and private network



The functional differences between the private network and the basic network are shown in the table below:

Function	Private Network	Basic Network
Renter Association	Logical Isolated Network Based on GRE Encapsulation	Renter Association
Network customization	Yes	No
Router customization	Yes	No
Custom IP	Yes	No
Intercommunication Rules	Supports cross-domain, cross-account intercommunication	Interoperable among renters in the same geographical area
Security control	Security group and Network ACL	Security group

Advantages of VPC

By starting the instance within the VPC, you can:

- Assign an instance of **your custom** private static IP.
- Assign multiple IP addresses to your instance (coming soon).
- Controls inbound and outbound traffic for an instance.
- Add an additional access control layer to the instance using the Network Access Control List (ACL).

Share and access resources between basic network and VPC

Some of the resources and functions on Tencent Cloud can support two kinds of network environments, and can be shared or accessed between different networks.

Resources	Instructions
Image	can be used to start a CVM instance in any network environment
Elastic IP	Elastic IPs can be bound to any network environment on a CVM instance
Instances	Instances on a basic network and instances within the private network can be accessed through Public IP or Basic Network Interoperability functions to achieve intercommunication
SSH Key	SSH key supports loading a CVM instance under any network environment
Security Group	Security Groups support binding to CVM instances in any network environment

Note: [Cloud Load Balance](#) cannot be shared between the underlying network and the VPC. That is, Cloud Load Balance does not support binding basic network instances and VPCs at the same time; even though the VPC and

the basic network are connected through the basic network and can intercommunicate.

Migrate instances within the basic network to VPC

- 1) [Create a Custom Image](#) for the CVM instance in the basic network environment.
- 2) (Optional) [Create a Snapshot](#) of the CVM instance data disk in the basic network environment.
- 3) [Create a VPC and Subnet](#).
- 4) [Purchase and Start the CVM Instance](#) in the VPC.

Private Network Access

Last updated : 2018-09-27 17:53:52

Private network services are LAN services. Cloud services can access each other via internal linkages. The cloud products on the Tencent Cloud can access each other via [Internet Access](#) or via the private network of Tencent Cloud.

Interconnected with megabyte/gigabyte underlying networks, Tencent Cloud's data centers can enable communication via private network featured by large bandwidth and low latency, which is free of charge in the same region, giving you the flexibility to build a network architecture.

Private IP address

Overview

Private IP addresses are IP addresses that cannot be accessed through Internet. Tencent Cloud's private network services are realized based on them. Each instance has a default network interface (eth0) for the assignment of private IPs. Private IP addresses can be automatically assigned by Tencent Cloud, or defined by users (only under the [VPC](#)).

Note:

Changing the private IP by yourself within the operating system may cause private network interruption.

Attribute

- Private network services are user-sensitive. Different users are isolated from each other, which means that the cloud services of the other user cannot be accessed via the private network by default.
- Private network services are also region-sensitive. Different regions are isolated from each other, which means that the cloud services under the same account in a different region cannot be accessed via the private network by default.

Application Scenarios

A private IP can be used for access between CLBs, CVMs, or access between CVMs and other cloud services (such as CDN and CDB) via a private network.

Address Assignment

Each CVM instance will be given a default private IP address when activated. The private IP varies with [Network Environment](#):

- Basic network: The private IP address is assigned by Tencent Cloud automatically and cannot be changed.
- VPC: The initial private IP address is assigned by Tencent Cloud randomly in VPC IP address range, and users can define the private IP address for the CVM instance within the 10.[0 - 255].0.0/8 , 172.[0 - 31].0.0/16 and 192.168.0.0/16 IP address ranges. The specific value range is determined by the VPC where the instance locates. For more information, please see [VPC and Subnet](#).

Private Network DNS

DNS Server Address

Private network DNS service is used for domain name resolution. If DNS configuration is incorrect, the domain name will become inaccessible.

Tencent Cloud provides reliable private network DNS servers in different regions. Specific configurations are shown below:

Network Environment	Region	Private Network DNS Server
Basic network	Guangzhou	10.225.30.181 10.225.30.223
	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 Zone	100.83.224.91 100.83.224.88
	North America	10.116.19.188 10.116.19.185
	Hong Kong	10.243.28.52 10.225.30.178
	Singapore	100.78.90.19 100.78.90.8
	Guangzhou Open Zone	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	
VPC	All regions	183.60.83.19 183.60.82.98

Internet Access

Last updated : 2018-09-20 16:50:24

To make the applications deployed on CVM instances publicly available, data must be transferred over the Internet and an Internet IP address (also known as public IP address) is required. Tencent Cloud provides Internet access via high-speed connected networks of Tencent Cloud IDC. The domestic multi-line BGP network covers more than 20 network operators and public network egress achieves cross-region switchover in a matter of seconds. This ensures that your users can enjoy high-speed and secure network no matter which network they use.

Public IP Address

- **Overview:** Public IP addresses are non-reserved addresses on the Internet. A CVM with a public IP address can access other computers on the Internet and can be accessed by other computers.
- **How to obtain:** When you create a CVM, set the bandwidth to a value greater than 0 Mbps in the network. Then Tencent Cloud system automatically assigns a public IP address to the instance from Tencent Cloud public IP address pool. This address cannot be changed and it is not associated with your Tencent Cloud account.
- **Configuration:** You can log in to a CVM instance with a public IP address on the Internet and configure it accordingly. For more information, please see [Log in to Linux Instance](#) and [Log in to Windows Instance](#).
- **Translation:** The public IP address is mapped to the [private IP address](#) of the instance through Network Address Translation (NAT).
- **Maintenance:** All the public network interfaces of Tencent Cloud are managed by Tencent Gateway (TGW). Public network ENIs of Tencent CVM instances are configured on the API layer TGW, with CVMs unaffected. Therefore, when you use such commands as `ifconfig` (Linux) or `ipconfig` (Windows) in the CVM to view the network interface information, only the information of [private network](#) is displayed. To view public network information, log in to [CVM Console](#) and go to CVM List/Details page.
- **Fees:** Fees are charged on instances providing services through a public IP address. For more information, please see [Purchase Network Bandwidth](#).

Public IP Address Release

You cannot actively associate or release the public IP address associated with an instance.

In the following cases, the public IP address is released or reallocated:

- **The instance is terminated.** When you actively terminate a postpaid instance or a prepaid instance is terminated upon expiry, Tencent Cloud will release the public IP address of the instance.
- **An [EIP address](#) is associated or dissociated with the instance.** When an instance is associated with an EIP address, Tencent Cloud will release the original public IP address of the instance. When the instance is dissociated with the EIP address, the instance is automatically assigned a new public IP address. The previously released public IP address is returned to the public IP address pool and you cannot reuse it.

If you need a permanent public IP address, use an [EIP address](#).

Elastic Public IP (EIP)

Last updated : 2018-07-30 12:03:27

Elastic IP, is referred to as EIP for short. It is a static IP designed for dynamic cloud computing, and a fixed public IP in a certain region. In case of an instance failure, the EIP can be remapped to another instance in your account (or [NAT gateway instance](#)) quickly to block the failure.

Common Operations

The following describes how to use EIPs.

Applying for EIPs

1. Log in to the [CVM Console](#).
2. In the left navigation pane, click **EIP**.
3. Click the **Apply** button, enter a region and the number of EIPs you want to apply for, and then click **OK**.
4. After this, you can see in the list the new EIP(s) you just applied for, which have/has an unbound status.

Binding EIPs to cloud products

1. Log in to the [CVM Console](#).
2. In the left navigation pane, click **EIP**.
3. In the EIP list, click the **Bind** button next to the EIP to be bound to a cloud product. (If the EIP is already bound to an instance, this button is unavailable. Please unbind it first.)
4. In the popup box, select the cloud product type that you want to bind, and then select the cloud product instance ID. Click the **Bind** button to complete the binding.

Unbinding EIPs from cloud products

1. Log in to the [CVM Console](#).
2. In the left navigation pane, click **EIP**.
3. In the EIP list, click the **Unbind** button next to the EIP that is already bound to a cloud product.
4. Click **OK**.

Note:

After unbinding, the cloud product instance may be assigned a new public IP, which may be different from the one before binding.

Releasing EIPs

1. Log in to the [CVM Console](#).
2. In the left navigation pane, click **EIP**.
3. In the EIP list, click **More** -> **Release** button next to the EIP to be released.
4. Click **OK**.

Adjusting bandwidth

1. Log in to the [CVM Console](#).
2. In the left navigation pane, click **EIP**.
3. In the EIP list, click the **Change Bandwidth** button next to the EIP for which you want to adjust bandwidth.
4. Adjust the target bandwidth value in the Adjust Bandwidth page.
5. Click **OK**.

Converting public IP to EIP

The public IP purchased along with the CVM instance is an ordinary public IP. It does not have elasticity and cannot be mounted and unmounted. Tencent Cloud allows you to convert an ordinary public IP to an EIP by following the steps below:

1. Log in to the [CVM Console](#).

2. Click the Convert icon.

云主机

云服务器新购特惠, 超值活动价 0.73元/天起 立即购买

广州 深圳金融 上海 上海金融 北京 成都 香港 新加坡 孟买 首尔 多伦多 广州Open 硅谷 弗吉尼亚 法兰克福

+ 新建 开机 关机 重启 续费 重置密码 更多操作

所属项目: 默认项目 多个关键字用竖线"|"分隔

ID/主机名	监控/状态	可用区	主机类型	配置	主IP地址	主机计费模式	所属项目	操作
ins- [redacted]	运行中	上海二区	标准型S3	32核 128GB 1Mbps 系统盘: 本地硬盘 网络: 基础网络	[redacted] (公) IP 10.105.98.166 (内) IP	包年包月		<input type="checkbox"/> 转换 IP保持不变, 转为弹性公网IP, 您最多可开通20个弹性公网IP, 已开通14个。

3. Click **Confirm Conversion**.

转换为弹性公网IP

检测到这台CVM已有公网IP([redacted]), 是否将其转化为EIP?

特性: IP将保持不变, 直接转为弹性公网IP保持绑定, 且不中断您的服务!

配额: 您最多可开通 20 个EIP, 已开通 14 个。

计费: 为保证IP资源有效利用, 未绑定云服务器的EIP, 将按小时收取少量费用。

确定转换

Troubleshooting

Network inaccessibility may occur with an EIP. This is generally caused by the following reasons:

- The EIP is not bound to any cloud product. For more information about how to bind an EIP to cloud products, please see [Binding EIP to Cloud Products](#).
- Security policy is invalid. Check if there is a valid security policy (security group or network ACL). If the bound cloud product has a security group policy, for example: access to 8080 port is denied, the port 8080 of the EIP is also inaccessible.

Elastic Network Interface (ENI)

Last updated : 2018-08-23 10:46:45

ENI is an elastic network interface bound to CVMs in a VPC, which can be migrated freely among multiple CVMs. It is very useful for configuring management networks and establishing highly reliable network solutions.

Elastic NIC has VPC, availability zone and subnet attributes. You can only bind it to CVMs under the same availability zone. A CVM can be bound with multiple ENIs. The maximum number of ENIs allowed to be bound to a CVM depends on the CVM's specification.

Concepts

- **Primary ENI or secondary ENI:** The ENI created with the creation of CVM within VPC is the primary ENI, and those created by users are secondary ENIs. The primary ENI does not support binding and unbinding, while secondary ENIs support.
- **Primary private IP:** The primary private IP of an ENI is assigned by the system or specified by users when the ENI is created. You can modify the primary private IP of the primary ENI, but not the ones of secondary ENIs.
- **Secondary private IP:** The secondary private IPs bound to the ENI, other than the primary IP, are configured by users when they create or modify the ENI. You can bind/unbind these IPs.
- **EIP:** Bound with private IPs of an ENI in a one-to-one manner.
- **Security group:** An ENI can be bound with one or more security groups.
- **MAC address:** An ENI has a globally unique MAC address.

Application Scenarios

- **Isolation among private network, public network and management network:**
Isolation among private network, public network and management network for the purpose of secure data transmission is required for the network deployment of important businesses. Data security and network isolation are guaranteed by employing different routing policies and security group policies. Like a physical server, a CVM can be bound with ENIs residing in different subnets to achieve isolation among three networks.
- **Highly reliable application deployment:**
The high availability of key components in the system architecture is ensured through multi-server hot backup. Tencent Cloud provides ENIs and private IPs that can be flexibly bound and unbound. You can configure the disaster recovery of Keepalived to achieve highly available deployment of key components.

Use Limits

The number of ENIs bound to a CVM is quite different from that of private IPs bound to an ENI depending on the CPU and memory configurations. These allowed numbers are shown in the following table:

CVM Configuration	Max. Number of ENIs	Max. Number of IPs Bound to Each ENI
-------------------	---------------------	--------------------------------------

CVM Configuration	Max. Number of ENIs	Max. Number of IPs Bound to Each ENI
CPU: 1-core. Memory: 1 GB	2	2
CPU: 1-core. Memory: > 1 GB	2	6
CPU: 2-core	2	10
CPU: 4-core. Memory: < 16 GB	4	10
CPU: 4-core. Memory: > 16 GB	4	20
CPU: 8- to 12-core	6	20
CPU: > 12-core	8	30

API Overview

CVM-related APIs of ENI are shown below. For more information on how to work with ENIs, please see [Overview of ENI APIs](#).

Feature	Action ID	Description
Create ENI	CreateNetworkInterface	Create an ENI
Apply for private IPs for ENI	AssignPrivateIpAddresses	Apply for private IPs for an ENI
Bind EIP to CVM	AttachNetworkInterface	Bind an EIP to a CVM

Login Password

Last updated : 2018-05-07 18:28:05

The first step in using a CVM instance is to login. To ensure the security and reliability of the instance, Tencent provides two encryption login methods: password login and [SSH key pair login](#). A password is the login credentials specific to each CVM instance, and the SSH key can be used for multiple CVM instances at the same time.

Anyone with an instance login password can log into the CVM instance remotely through a public network address that is allowed by the security group. Therefore, it is recommended that you use a more secure password, keep it safe, and modify the login password for your instance periodically.

A user can specify to use a password or SSH key when purchasing [Purchase and Start Instance](#). When using a password, it can be set by itself or automatically generated. When the password is generated automatically, the initial password will be delivered to the user via [Internal Message](#). Users can learn from the below content on how to set an initial password and how to reset settings in case you forgot your password, etc.

Set the initial password

- 1) [Purchase and Start Instance](#), you can select the login method in the Set Host Name and Login Mode section. The default is [Set Password].
- 2) In accordance with password character limitations, enter the host password and confirmation, then click Buy Now; the initial password will be set successfully when the CVM instance is successfully assigned.
- 3) You can also select [Auto Generate Password], and click [Buy Now] to get the CVM instance initial password via [Internal Message](#) after the CVM instance is successfully assigned.

It is important to note that the character limitations for setting a password is as follows:

- Linux device passwords must be between 8-16 chars, and include 2 of the following items (a-z 、 A-Z 、 0-9 and [、 (、)、 \、 ~、 !、 @、 #、 \$、 %、 ^、 &、 \、 *、 -、 +、 =、 |、 {、 }、 [、]、 :、 ;、 '、 /、 . 、 ?、 /、] and ` or other special symbols)
- Windows device passwords must be between 12-16 chars, and include 3 of the following items (a-z 、 A-Z 、 0-9 and [、 (、)、 \、 ~、 !、 @、 #、 \$、 %、 ^、 &、 \、 *、 -、 +、 =、 |、 {、 }、 [、]、 :、 ;、 '、 /、 . 、 ?、 /、] and ` or other special symbols)

Reset password

Note: You can reset the password for the cloud host only when it is powered off. If the machine is running, please shut down the host first.

- 1) Open [CVM console](#).
- 2) For a single CVM instance that is shut down, in the right-hand action bar, click [More] - [Reset Password].

- 3) For batch CVM instances, select all the hosts that need a password reset; then at the top of the list, click [Reset Password] to modify the host login password in batches. A CVM instance that cannot have its password reset will display the reason why.
- 4) In the Reset password pop-up box, enter the new password, confirm the password and the verification code, then click [Confirm Reset].
- 5) Wait for the reset to succeed, and you will receive a successful reset message in your station inbox; now you can use the new password to start and use your CVM.

SSH Key

Last updated : 2018-07-30 11:59:45

To ensure the security and reliability of instances, Tencent Cloud allows users to log in to instances with: [password](#) and SSH key pair. This document describes how to make configurations for login with SSH key pair. Users of instances with different operating systems can choose an encryption method by referring to the settings section in [Custom Configuration of Windows CVMs](#) and [Custom Configuration of Linux CVMs](#).

SSH Key Overview

Tencent Cloud allows encryption and decryption of the information for logging into Linux instances with public key cryptography. The public key cryptography uses a public key to encrypt data (such as a password), and then recipients can decrypt the data using a private key. The combination of a public key and a private key is also known as key pair. An SSH key can be used to connect to instances in a safe way, making it a more secured login method than normal passwords.

Tencent Cloud only stores the public key, and you must retain the private key. Anyone who has your private key can decrypt your login information. So be sure to keep your private key in a secure location.

Features and Advantages

Compare to the traditional verification method involving user name and password, SSH key has the following advantages:

- Login verification with SSH key offers a higher security to prevent brute force attacks.
- Login with SSH key makes it more convenient for you to log in to your instance by making simple configurations on the console and local client, without the need to enter the password again in case of re-login.

Use Limits

- Only supported for Linux instances.
- Tencent Cloud will not retain your private key information. You need to click **Download** to obtain the private key within 10 minutes after you have created an SSH key, and keep it properly.
- A Linux instance can only bound with an SSH key. A key that has been bound to your instance will be replaced with the new key.
- To ensure data security, you need to shut the instance down before the key can be loaded.

Operation Instructions

For more information, please [SSH Key](#).

Security Groups

Last updated : 2018-08-06 10:10:34

Security Group Overview

Security group is a stateful virtual firewall for filtering packets and is used to set the network access controls for a single or multiple CVMs. It is a logical grouping and an important means of network security isolation.

- You can add the basic network-based CVM or ENI instances with the same network security isolation requirements in the same region to the same security group.
- You can use the security group policies to filter the inbound and outbound traffic of an instance, which can be a basic network-based CVM or an ENI instance.
- You can modify security group rules at any time, and the new rules take effect immediately.

Security Group Template

You can create a custom security group, or create a security group from template. Three templates are available in the system:

- Open Port 22 on Linux: Only TCP port 22 for SSH login is open to Internet, and all private network ports are open.
- Open Port 3389 on Windows: Only TCP port 3389 for MSTSC login is open to Internet, and all private network ports are open.
- Open all ports: All ports are open to Internet and private networks. This poses certain security risks.

Security Group Rules

Security group rules are used to control the inbound and outbound traffic of instances associated with the security group (filtered based on the rules from top to bottom). By default, a new security group rejects all traffic (All Drop). Therefore, a CVM associated with a security group with no rules will reject all traffic.

Each security group rule involves the following items:

- Type: You can select a system rule template, or create custom rules.
- Source or destination: Traffic source (inbound rules) or destination (outbound rules). Choose one of the following options:
 - Specify a single IP in CIDR notation.
 - Specify an IP address range in CIDR notation, such as 203.0.113.0/24.
 - Reference a security group ID. You can reference one of the following security group IDs:
 - Current security group. (This indicates whether CVMs associated with the security group can be mutually accessed or not)
 - Another security group. Another security group ID of the same project in the same region.
 - Reference an IP address object or an IP address group object in the [Parameter Template](#).

- Protocol port: Enter the protocol type and port range, or reference an IP address object or an IP address group object in the [Parameter Template](#).
- Policy: Allow or Reject.

Note:

- Referencing security group ID is optional. The rules of the referenced security group will not be added to the current security group.
- If you enter the security group ID in Source/Destination when configuring security group rules, the private IPs of the CVM's primary ENI and the ENI that are associated with this security group ID are used as Source/Destination, excluding public IP.

Security Group Priority

- Multiple security groups associated with an instance are **prioritized from the smaller number to the bigger number**.
- The rules in a security group are **prioritized from top to bottom**.

An instance associated with a security group with no rules will reject all traffic by default.

Security Group Restrictions

- Security groups are region and project-specific. CVMs can only be associated with the security groups in the same region and project.
- Security groups apply to any CVM instances in [network environment](#).
- Each user can set a maximum of 50 security groups for each project in each region.
- A maximum of 100 inbound/outbound access policies can be set for a security group.
- A CVM can be associated with multiple security groups, and a security group can be associated with multiple CVMs. No number limit is imposed.
- Security groups associated with CVMs in **basic networks cannot filter** data packets sent from (or to) CDB and cloud cache service (Redis and Memcached) of Tencent Cloud. If necessary, you can use iptables to filter traffic of such instances.

Feature	Count
Security group	50/Region
Access policy	100 (Inbound/Outbound)
Number of security groups associated with an instance	No limit
Number of instances associated with a security group	No limit

Note:

If you have a large number of instances that need to access each other, you can assign them to multiple security groups, and achieve mutual authorization and access by configuring the rules for security group IDs.

Comparison of Security Group and Network ACL

Security Group	Network ACL
Operations at the CVM instance level (the primary defense layer)	Operations at the subnet level (the secondary defense layer)
Allow and Reject rules are supported	Allow and Reject rules are supported
Stateful: Returned traffic is automatically allowed without subjecting to any rules	Stateless: Returned traffic must be explicitly allowed by rules
Operations are only applied to an instance if the security group is specified when the instance is launched or the security group is associated with the instance later	Operations are automatically applied to all CVM instances in the associated subnet

Security Group-related Cloud APIs

Security group-related cloud APIs are Developer Tools of security groups. Cloud APIs can be used to work with security groups and manage the configurations of security groups and CVM instances. For more information, please see [Security Group-related APIs](#).

Monitoring and Warning

Last updated : 2018-05-07 11:21:39

Provides Tencent Cloud CVMs with a good monitoring environment, in order to maintain high reliability, high availability and high performance for the most important parts of CVM instances. Users can easily collect different dimensions of monitoring data from different resources; this allows for an easy grasp of the use of these resources and location of faults.

Collecting monitoring data helps the user establish normal criteria for CVM instance performance. By measuring the performance of an instance and collecting monitoring data history at different times and under different load conditions, the user can clearly understand the normal performance of the instance and quickly determine whether the instance is in an abnormal state according to the current monitoring data; this will allow them to quickly find a solution to the problem. For example, users can monitor instances of CPU utilization, memory usage and disk I/O. If the instance performance is lower than normal at a certain time, you may need to trigger an alarm to notify the user to upgrade the instance configuration or increase the number of instances, in order to reduce single instance loads and keep the system running smoothly.

To establish a normal performance baseline for a CVM instance, you should monitor at least the following:

Item	Metrics
CPU usage	cpu_usage
Memory usage	mem_usage
Private network outgoing traffic	lan_outtraffic
Private network incoming traffic	lan_intraffic
Public network outgoing traffic	wan_outtraffic
WAN incoming traffic	wan_intraffic
Disk usage	disk_usage
Disk IO waiting time	disk_io_wait

For more details on monitoring indicators, refer to [Cloud Monitoring Product Documentation](#).

Cloud monitoring collects raw data from running CVM instances and presents the data as easy-to-read tables. By default, the statistics are saved for one month. You can view the operation status of an instance for one month to get a better understanding of the application service. By default, cloud monitoring collects operating data of an instance every 5 minutes. Some of Tencent Cloud services support a more detailed 1 minute granular monitoring.

The console of different products may display a series of graphs based on the raw data of cloud monitoring. The cloud monitoring console integrates monitoring data of all products, which is beneficial for users when obtaining an overall operations overview. According to the user's requirements, they can choose different access points to obtain instance data.

At the same time, creating an alarm for an instance indicator that you are following, will enable the CVM instance to send alarm information to the concerned group in a timely manner, when the operation status has reached a certain condition.

This allows you to detect anomalies in a timely manner and take the appropriate measures to maintain system stability and reliability. For more information, refer to [Creating Alarms](#).

Access Control

Access Control Overview

Last updated : 2018-06-22 16:08:20

If you use such services as Cloud Virtual Machine (CVM), VPC, and database in Tencent Cloud and these services are managed by different users who share your cloud account key, the following problems may exist:

- Your key is shared by multiple users, which leads to a high risk of leakage.
- You cannot limit the access permission of other users, which is easy to pose a security risk due to misoperation.

To avoid the above problems, you can use sub-accounts to allow different users to manage different services. By default, sub-accounts have no permission to use CVM or related resources. So it is necessary to create a policy to grant sub-accounts the permissions to use the resources they need.

Cloud Access Management (CAM) is a service package provided by Tencent Cloud, which is used to help customers manage the permissions to access resources under Tencent Cloud accounts in a secure way. By using CAM, you can create, manage and terminate users (or user groups), as well as determine which Tencent Cloud resources can be accessed and who can use them through identity management and policy management.

When using CAM, you can associate a policy to a user or a group of users. The policy can authorize or reject users to finish specified tasks using specified resources. For more basic information about CAM policy, please see [Policy Syntax](#). For more information on how to use CAM policy, please see [Policy](#).

If you do not need to manage the access permission to CVM resources for sub-accounts, you can skip this chapter. This will not affect your understanding and usage of other parts in this document.

This feature is in Beta test for now. Submit a [Ticket](#) to apply for it.

Getting Started

CAM policy must authorize or reject users to perform one or more CVM operations. Meanwhile, it must specify the resources to work with (all resources, or part of resources for some operations). Policy can also contain conditions set for resource operations.

Some of CVM-based API operations support resource-level permissions, which means that you must specify all resources to use but cannot specify a specific resource for use when performing such API operations.

Task	Link
Learn about the basic structure of the policy	Policy Syntax
Define operations in the policy	CVM Operations
Define resources in the policy	CVM Resource Path
Limit the policy with conditions	CVM Condition Keys
Resource-level permissions supported by CVM	Resource-level Permissions Supported by CVM
CAM policy example	CAM Policy Example

Authorizable Resource Type

Last updated : 2018-06-22 16:04:42

Resource-level permissions can be used to specify which resources users can work with. CVM supports part of the resource-level permissions. This means that for some CVM operations, you can control the time when users are allowed to perform operations (condition must be met) or to use specified resources. The following table describes the types of resource that can be authorized in CVM.

The type of resource that can be authorized in CAM:

Resource type	Resource description method in authorization policy
CVM instance related	qcs::cvm:\$region::instance/*
CVM key related	qcs::cvm:\$region::keypair/*
CVM image related	qcs:t:cvm:\$region:\$account:image/*

The following table describes the API operations of CVM (Cloud Virtual Machine) that support resource-level permissions, and the resources and condition keys supported for each operation. You can use wildcard "*" in the path when specifying resource path.

Note:

Any CVM API operation not listed here does not support resource-level permission. If any CVM API operation does not support resource-level permission, you can still authorize the permission of this operation to users, but you must specify * for the resource element of policy statement.

CVM Instance Related:

API Operation	Resource Path	Condition Key
DescribeInstanceInternetBandwidthConfigs	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ModifyInstanceInternetChargeType	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ModifyInstancesAttribute	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ModifyInstancesProject	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type

API Operation	Resource Path	Condition Key
ModifyInstancesRenewFlag	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
RebootInstances	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
RenewInstances	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ResetInstance	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/*	cvm:region cvm:zone cvm:instance_type
ResetInstancesInternetMaxBandwidth	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ResetInstancesPassword	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ResetInstancesType	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
ResizeInstanceDisks	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
RunInstances	qcs::cvm:\$region:\$account:instance/* 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:sg/* qcs::cvm:\$region:\$account:sg/\$sgId qcs::vpc:\$region:\$account:subnet/* qcs::vpc:\$region:\$account:subnet/\$subnetId qcs::cvm:\$region:\$account:systemdisk/* qcs::cvm:\$region:\$account:datadisk/* qcs::vpc:\$region:\$account:vpc/* qcs::vpc:\$region:\$account:vpc/\$vpcId	cvm:region cvm:zone cvm:instance_type

API Operation	Resource Path	Condition Key
StartInstances	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
StopInstances	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type
TerminateInstances	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId	cvm:region cvm:zone cvm:instance_type

CVM Key Related:

API Operation	Resource Path	Condition Key
AssociateInstancesKeyPairs	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId	-
CreateKeyPair	qcs::cvm:\$region:\$account:keypair/*	-
DeleteKeyPairs	qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId	-
DescribeKeyPairs	qcs::cvm:\$region:\$account:keypair/*	-
DescribeKeyPairsAttribute	qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId	-
DisassociateInstancesKeyPairs	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId	-
ImportKeyPair	qcs::cvm:\$region:\$account:keypair/*	-
ModifyKeyPairAttribute	qcs::cvm:\$region:\$account:keypair/* qcs::cvm:\$region:\$account:keypair/\$keyId	-

CVM Image Related:

API Operation	Resource Path	Condition Key
CreateImage	qcs::cvm:\$region:\$account:instance/* qcs::cvm:\$region:\$account:instance/\$instanceId qcs::cvm:\$region:\$account:image/*	cvm:region
DeleteImages	qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId	cvm:region

API Operation	Resource Path	Condition Key
DescribeImages	qcs::cvm:\$region:\$account:image/*	cvm:region
DescribeImagesAttribute	qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId	cvm:region
DescribeImageSharePermission	qcs::cvm:\$region:\$account:image/*	cvm:region
ModifyImageAttribute	qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId	cvm:region
ModifyImageSharePermission	qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId	cvm:region
SyncImages	qcs::cvm:\$region:\$account:image/* qcs::cvm:\$region:\$account:image/\$imageId	cvm:region

Authorization Policy Syntax

Last updated : 2018-06-22 16:07:22

Policy Syntax

CAM policy:

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

- **version** is required. Now, the only available value is "2.0".
- **statement** describes details of one or multiple permissions. This element includes permission or permission set for other elements such as effect, action, resource, and condition. One policy can only have one statement element.
 - i. **action** describes allowed or rejected actions. An action can be an API (described using the prefix "name") or a feature set (a set of specific APIs, described using the prefix "permid"). This element is required.
 - ii. **resource** describes the detailed data of authorization. Resource is described using 6-piece format. Detailed resource definition for each product is different. For more information on how to specify resource, please see the product documentation corresponding to your resource declaration. This element is required.
 - iii. **condition** describes the condition for the policy to take effect. Condition consists of operator, operational key and operate value. Condition value may include information such as time and IP address. Some services allow you to specify other values in the condition. This element is optional.
 - iv. **effect** describes whether the result produced by the declaration is "Allow" or "Explicitly Deny". This includes two situations: allow, deny. This element is required.

CVM Operations

In CAM policy statements, you can specify any API operation from any service that supports CAM. For CVM, use API 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 specify multiple operations with a wildcard. For example, you can specify all operations whose name starts 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 resources.

The general form of resource path is as follows:

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

project_id describes the project information, which is only used to be compatible with the earlier CAM logics and is not required.

service_type: the abbreviation of the product, for example, CVM.

region: region information, for example, bj.

account: the root account information of the resource owner, for example, uin/164256472.

resource: detailed resource information of each product, for example, 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 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 any API operation does not support resource-level permission, you can use the wildcard "*" in Resource element as shown below:

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

To specify multiple resources in one instruction, 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" represents project ID.
- "region" represents region.
- "account" represents account ID.

Resource	Resource description method in authorization policy
Instance	qcs:cvm:\$region:\$account:instance/\$instanceId
Key	qcs:cvm:\$region:\$account:keypair/\$keyId

Resource	Resource description method in authorization policy
VPC	qcs::vpc:\$region:\$account:vpc/\$vpclId
Subnet	qcs::vpc:\$region:\$account:vpc/\$vpclId
System disk	qcs::cvm:\$region:\$account:systemdisk/*
Image	qcs::cvm:\$region:\$account:image/*
Subnet	qcs::vpc:\$region:\$account:subnet/\$subnetId
Data disk	qcs::cvm:\$region:\$account:datadisk/*
Security group	qcs::cvm:\$region:\$account:sg/\$sgId
EIP	qcs::cvm:\$region:\$account:eip/*

CVM Condition Keys

In a policy statement, you can selectively specify conditions for controlling the effective time of the policy. Each condition includes one or multiple key value pairs. Condition keys are case insensitive.

- If you specify multiple conditions or specify multiple keys in one condition, we make an evaluation with logical operation "AND".
- If you specify a key with multiple values in one condition, we make an evaluation with logical operation "OR". All conditions must be met to grant a permission.

The following table describes CVM condition keys for specific services.

Condition keys	Types for reference	Key value pairs
cvm:instance_type	String	cvm:instance_type= instance_type <ul style="list-style-type: none"> ◦ Instance_type represents instance type (for example, S1.SMALL1).
cvm:image_type	String	cvm:image_type= image_type <ul style="list-style-type: none"> ◦ Image_type represents image type (for example, IMAGE_PUBLIC).
vpc:region	String	vpc:region= region <ul style="list-style-type: none"> ◦ region represents region (for example, ap-guangzhou).
cvm:disk_size	Integer	cvm:disk_size= disk_size <ul style="list-style-type: none"> ◦ disk_size represents disk size (for example, 500).

Condition keys	Types for reference	Key value pairs
cvm:disk_type	String	<p>cvm_disk_type= disk_type</p> <ul style="list-style-type: none">◦ disk_type represents disk type (for example, CLOUD_BASIC).
cvm:region	String	<p>cvm:region= region</p> <ul style="list-style-type: none">◦ region represents region (for example, ap-guangzhou).