Cloud Virtual Machine

Operation Guide

Product Documentation

Contents

Operation Guide
Operation Guide Overview

Last updated：2024-01-08 09:32:02

This document provides an overview of CVM instances and their use cases. It also describes how to operate CVM instances.

Purchasing and Using a CVM

If this is the first time you are purchasing and using a CVM instance, we recommend following the instructions below to get started.

1. To learn about CVM instance, see CVM Overview.

2. Select and purchase an appropriate CVM model. See Customizing Linux CVM Configurations.

3. Log in to the CVM instance you purchased: Depending on the instance type purchased, you may choose to either log in to the Windows instance or Linux instance.

Adjusting CVM Configurations

You may need to adjust the disk type, network or other configurations of the CVM instance due to changing demands. See the following documents to make corresponding changes.

Changing Instance Configuration

Adjusting Network Configuration

Adjusting Project Configuration

Reinstalling System

Resetting Password and Key

If you forgot your password or lost your key, refer to the following documents to reset the password or key:

Resetting Instance Password

Managing SSH Keys

Renewing Instances and the Billing

See Renewing Instances

Creating, Importing or Deleting a Custom Image

An Image provides the information required for launching an CVM instances. Tencent Cloud provides three types of images: public image, custom image and shared image. We currently support the following image-related operations.

Creating Custom Images

Deleting Custom Images

Importing Images

Copying Images

Troubleshooting

When you are unable to log in to the CVM instance, or if you are experiencing slow response or other issues, refer to the following for troubleshooting:

CVM Login Failures

CVM Network Latency and Packet Loss

Use Limits

Last updated：2024-05-16 10:55:46

Account-level Limits for Purchasing CVM Instances

You need to sign up for a Tencent Cloud account. For more information, see Signing up for a Tencent Cloud Account. If you create a pay-as-you-go CVM, the system will freeze the cost of one-hour CVM usage. Make sure that your account has sufficient balance for the order.

CVM Instance Use Limits

Virtualized software cannot be installed or re-virtualized (such as installing VMware or Hyper-V).
You cannot use sound cards or mount external hardware devices (such as USB flash drives, external disks, and U-keys).
Only Linux CVMs can act as a public gateway.

CVM Instance Purchase Limits

The **purchase limit** of pay-as-you-go CVM instances for each user in each AZ is between 30 and 60.
For more information, see Purchase Limits.

Image Limits

Public images: No use limits.
Custom images: Each region supports a maximum of 500 custom images.
Shared images: Each custom image can be shared with a maximum of 500 Tencent Cloud users. Custom images can only be shared with accounts in the same region as the source account.
For more information, see Image Types.

ENI Limits

Based on CPU and memory configurations, the number of ENIs bound to a CVM instance differs from the number of private IPs bound to an ENI. The quotes are as shown below:

Note:

 The number of IP addresses bound to a single ENI indicates the maximum number allowed. The EIP quota is not provided based on this upper limit but based on EIP use limits.

ENIs per CVM instance

Private IPs per ENI

Model	Instance Type	Number of ENIs								
		CPU: 1 core	CPU: 2 cores	CPU: 4 cores	CPU: 6 cores	CPU: 8 cores	CPU: 10 cores	CPU: 12 cores	CPU: 14 cores	CPU: 16 core
Standard	Standard S5	2	4	4	-	6	-	-	-	8
	Standard Storage Optimized S5se	-	-	4	-	6	-	-	-	8
	Standard SA2	2	4	4	-	6	-	-	-	8
	Standard S4	2	4	4	-	6	-	-	-	8
	Standard Network-optimized SN3ne	2	4	4	-	6	-	8	-	8
	Standard S3	2	4	4	-	6	-	8	-	8
	Standard SA1	2	2	4	-	6	-	-	-	8
	Standard S2	2	4	4	-	6	-	8	-	8
	Standard S1	2	4	4	-	6	-	8	-	8
High IO	High IO IT5	-	-	-	-	-	-	-	-	8
	High IO IT3	-	-	-	-	-	-	-	-	8

Memory Optimized	Memory Optimized M5	2	4	4	-	6	-	8	-	8
	Memory Optimized M4	2	4	4	-	6	-	8	-	8
	Memory Optimized M3	2	4	4	-	6	-	8	-	8
	Memory Optimized M2	2	4	4	-	6	-	8	-	8
	Memory Optimized M1	2	4	4	-	6	-	8	-	8
Compute	Compute Optimized C4	-	-	4	-	6	-	-	-	8
	Compute Network-optimized CN3	-	-	4	-	6	-	-	-	8
	Compute C3	-	-	4	-	6	-	-	-	8
	Compute C2	-	-	4	-	6	-	-	-	8
GPU-based	GPU Compute GN6	-	-	-	-	-	-	-	-	-
	GPU Compute GN6S	-	-	4	-	6	-	-	-	-
	GPU Compute GN7	-	-	4	-	6	-	-	-	-
	GPU	-	-	-	4	-	-	-	8	-

Model	Instance Type									
	Compute GN8									
	GPU Compute GN10X	-	-	-	-	6	-	-	-	-
	GPU Compute GN10Xp	-	-	-	-	-	6	-	-	-
FPGA-based	FPGA Accelerated FX4	-	-	-	-	-	6	-	-	-
Big Data	Big Data D3	-	-	-	-	6	-	-	-	8
	Big Data D2	-	-	-	-	6	-	-	-	8
	Big Data D1	-	-	-	-	6	-	-	-	-
CPM	Not supported									

Model	Instance Type	Private IPs bound to a single ENI							
		CPU: 1 core	CPU: 2 cores	CPU: 4 cores	CPU: 6 cores	CPU: 8 cores	CPU: 10 cores	CPU: 12 cores	CPU: 14 cores
Standard	Standard S5	6	10	10	-	20	-	-	-
	Standard Storage Optimized S5se	-	-	20	-	20	-	-	-
	Standard SA2	6	10	10	-	20	-	-	-
	Standard S4	6	10	10	-	20	-	-	-

	Standard Network-optimized SN3ne	6	10	10	-	20	-	30	-
	Standard S3	6	10	10	-	20	-	30	-
	Standard SA1	1 GB memory: 2>1 GB memory: 6	10	8 GB memory: 1016 GB memory: 20	-	20	-	-	-
	Standard S2	6	10	10	-	20	-	30	-
	Standard S1	6	10	10	-	20	-	30	-
High IO	High IO IT5	-	-	-	-	-	-	-	-
	High IO IT3	-	-	-	-	-	-	-	-
Memory Optimized	Memory Optimized M5	6	10	10	-	20	-	30	-
	Memory Optimized M4	6	10	10	-	20	-	30	-
	Memory Optimized M3	6	10	10	-	20	-	30	-
	Memory Optimized M2	6	10	10	-	20	-	30	-
	Memory Optimized M1	6	10	10	-	20	-	30	-
Compute	Compute Optimized C4	-	-	10	-	20	-	-	-

	Compute Network-optimized CN3	-	-	10	-	20	-	-	-
	Compute C3	-	-	10	-	20	-	-	-
	Compute C2	-	-	10	-	20	-	-	-
GPU-based	GPU Compute GN2	-	-	-	-	-	-	-	-
	GPU Compute GN6	-	-	-	-	-	-	-	-
	GPU Compute GN6S	-	-	10	-	20	-	-	-
	GPU Compute GN7	-	-	10	-	20	-	-	-
	GPU Compute GN8	-	-	-	10	-	-	-	30
	GPU Compute GN10X	-	-	-	-	20	-	-	-
	GPU Compute GN10Xp	-	-	-	-	-	20	-	-
FPGA-based	FPGA Accelerated FX4	-	-	-	-	-	20	-	-
Big Data	Big Data D3	-	-	-	-	20	-	-	-
	Big Data	-	-	-	-	20	-	-	-

	D2								
	Big Data D1	-	-	-	-	20	-	-	-
CPM	Not supported								

Bandwidth Limits

Maximum outbound bandwidth (downstream bandwidth)

The following rules apply to instances created after 00:00, February 24, 2020:

Network Billing Method	Instance		Maximum Bandwidth Range (Mbps)
	Instance Billing Method	Instance Configuration	
Bill-by-traffic	Pay-as-you-go instances	All	0-100
Bill-by-bandwidth	Pay-as-you-go instances	All	0-100
Bandwidth package	All		0-2000

The following rules apply to instances created before 00:00, February 24, 2020:

Network Billing Method	Instance		Range of Bandwidth Cap (Mbps)
	Instance Billing Method	Instance Configuration	
Bill-by-traffic	Pay-as-you-go instances	All	0-100
Bill-by-bandwidth	Pay-as-you-go instances	All	0-100
Bandwidth package	All		0-2000

Maximum inbound bandwidth (upstream bandwidth)

Purchased fixed bandwidth > 10 Mbps:Tencent Cloud will assign a public network inbound bandwidth equals to the purchased bandwidth.

Purchased fixed bandwidth < 10 Mbps, Tencent Cloud will assign 10-Mbps public network inbound bandwidth.

Disk Limits

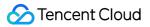| Limitations	Description
Elastic cloud disk capability	Starting from May 2018, all data disks purchased with CVM instances are elastic cloud disks, which can be unmounted from and remounted to CVM instances. This feature is supported in all availability zones.
Cloud disk performance	I/O specification applies to both input and output performance at the same time.For example, if a 1-TB SSD has a maximum random IOPS of 26,000, it means that both its read and write performance can reach this value. Due to performance limits, if the block size in this example is 4 KB or 8 KB, the maximum IOPS can be reached. If the block size is 16 KB, the maximum IOPS cannot be reached (throughput has already reached the limit of 260 MB/s).
Elastic cloud disks per CVM	A maximum of 20
Snapshots per region	64 + Number of cloud disks in the region x 64
Attaching cloud disks to a CVM	The CVM instance and cloud disks must be in the same availability zone.
Snapshot rollback	Snapshot data can only be rolled back to the cloud disk where the snapshot was created.
Creating cloud disks using snapshot - Type limit	Only snapshots of data disks can be used to create new elastic cloud disks.
Creating cloud disks using snapshot - Size limit	The capacity of new cloud disk must be larger than the source disk of the snapshot.

Security Group Limits

Security groups are region-specific. A CVM instance can only be bound to security groups in the same region.

Security groups are applicable to CVM instances in any network environment.

Each user can configure a maximum of 50 security groups for each project in a region.

A maximum of 100 inbound or outbound rules can be configured for a security group.

One CVM instance can be associated with multiple security groups, and a security group can be associated with multiple CVM instances.

Security groups associated with CVM instances on the **classic networkcannot filter packets** from or to TencentDB (MySQL, MariaDB, SQL Server, or PostgreSQL) and NoSQL (Redis or Memcached) databases. Instead, you can use iptables or purchase CFW to filter traffic for such instances.

The quotas are as shown below:

Item	Limit
Security groups	50 per region
Rules in a security group	100 for inbound rules and 100 for outbound rules
CVM instances associated with a security group	2,000
Security groups associated with a CVM instance	5
Security groups referenced by a security group	10

VPC Limits

Resource	Limit
VPCs per region per account	20
Subnets per VPC	100
Classic network-based CVMs associated with each VPC	100
Route tables per VPC	10
Route tables associated with each subnet	1
Routes per route table	50
HAVIPs per VPC	10

Convenience Features
Switching Instance Page View in Console

Last updated：2024-01-08 09:32:02

Overview

The instance list page in the CVM console supports tab and list views. You can switch between them as instructed below.

The tab view has the advantages of the self-service instance detection tool that automatically initiates detection, quick acquisition of the instance information, and shortcuts to frequent operations. We recommend you use the tab view if you have at least 5 CVMs.

Directions

1. Log in to the CVM console and select **Instance** on the left sidebar.

2. On the **Instance** page, you can select **Switch to Tab View** on the right to switch the view as shown below:

3. After the tab view is switched to successfully, the UI is as follows:

In the tab view, you can quickly get instance health status information and instance details, and perform instance management operations.

Instance ID

Availability Zone   Guangzhou Zone 3

IP

Instance Billing Mode   Pay-as-you-go   Modify billing mode

Bandwidth billing   Bill by traffic   Modify billing mode
mode

Instance Configuration Standard S5 - 2C 4G   Adjust Model and Specs

Operating System   TencentOS Server 2.6 (Final)   Reinstall the System

Creation Time   2022-03-03 12:12:06

Basic Information   ENI   Public IP   Monitoring   Security Groups   Operation Logs

Note:

If you have multiple CVM instances, you can select **Switch to List View** on the right to switch to the list view.

Instances

Creating Instances

Guidelines for Creating Instances

Last updated：2024-01-08 09:32:02

This document introduces several methods of creating CVM instances, from basic operations to advanced custom features.

Creating CVM instances via the CVM purchase page is the most commonly used method. It allows you to flexibly select the configurations that meet your business requirements. For more information, see Creating Instances via CVM Purchase Page.

If you want to use a particular operating system, application, or other configuration that you are familiar with, you can first create a custom image and select it when creating an instance to increase efficiency. For more information, see Creating Instances via Images.

If you want to purchase an instance with the same configurations as those of the current instance, you can directly create an instance with the same configurations. For more information, see Purchasing with Same Configurations.

Creating Instances via CVM Purchase Page

Last updated：2024-06-25 15:45:04

Overview

This document guides you through how to create a Tencent Cloud Virtual Machine (CVM) instance using the custom configuration mode as an example.

Preparations

Before creating a CVM instance, you need to complete the following steps:

Sign up for a Tencent Cloud account.

To create a CVM instance whose network type is virtual private cloud (VPC), you need to create a VPC in the target region and create a subnet in the target availability zone under the VPC.

If you do not use the default project, you need to create a project.

If you do not use the default security group, you need to create a security group in the target region and add a security group rule that meets your business requirements.

To bind an SSH key pair when creating a Linux instance, you need to create an SSH key for the target project.

To create a CVM instance with a custom image, you need to create a custom image or import an image.

Directions

1. Log in to Tencent Cloud. Select **Products** > **Compute and Container** > **Compute** > **Cloud Virtual Machine**. Click **Buy Now** to enter the CVM purchase page.

Custom Configuration: It is suitable for specific scenarios and makes it easier for you to purchase CVM instances as needed.

2. Configure the following information as prompted by the page:

Type	Required	Configuration Description
Billing mode	Yes	Select one as needed: **Pay-as-you-go**: It is an elastic billing method of CVM applicable to scenarios such as e-commerce flash sales, where demand will fluctuate significantly in an instant. **Spot instance**: A novel operational mode for instances, aptly suited for scenarios such as big data computing, and load-balanced online services and website services. As market supply and demand dynamics shift, the

		price of spot instances fluctuates accordingly, typically ranging from 3% to 20% of the pay-as-you-go price. For billing details, see Billing Plans.
Region/Availability Zone	Yes	**Region**: We recommend you select the region closest to your end users to minimize the access latency and improve the access speed. **Availability zone**: Select one as needed. If you want to purchase multiple CVM instances, we recommend you select different AZs to implement disaster recovery. For more information on regions and AZs, see Regions and AZs.
Instance	Yes	Tencent Cloud provides different instance types based on the underlying hardware. For more information on instances, see Instance Types.
Image	Yes	Tencent Cloud provides public images, custom images, and shared images. For more information on images, see Image Types.
System disk	Yes	It is used for OS installation and defaults to 50 GB. Available cloud disk types vary by region. Select one as instructed on the page. For more information on cloud disks, see Cloud Disk Types.
Data disk	No	It is used to scale up the storage capacity of the CVM instance to ensure high efficiency and reliability. It is not added by default. For more information on cloud disks, see Cloud Disk Types.
Scheduled Snapshot	No	A scheduled snapshot policy can be set for the system disk or data disk. For more information, see Scheduled Snapshots.
Quantity	Yes	It indicates the quantity of CVM instances to be purchased.

3. Click **Next: Set Network and CVM** to enter the instance settings page.

4. Configure the following information as prompted by the page:

Type	Required	Configuration Description
Network	Yes	It is a logically isolated network space built in Tencent Cloud. A VPC includes at least one subnet. The system provides a default VPC and subnet for each region. If the existing VPC or subnet does not meet your requirements, you can create a VPC or subnet in the VPC console. **Note**: By default, resources in the same VPC are interconnected over the private network. When purchasing a CVM instance, make sure that the CVM instance and its subnet are in the same AZ.
Public IP	No	If your CVM instance needs to access the public network, you need to assign a

		public IP for it. You can assign the public IP when creating the CVM instance or configure an EIP after the creation. **Note**: The dedicated public IP that is assigned free of charge cannot be unbound from the instance. To unbind this IP address, convert it to an EIP first. For more information on EIPs, see Elastic IP (EIP). No dedicated public IP can be assigned in the following two cases, subject to the information on the purchase page: The IP resources have been sold out. Resources are only available in certain regions.
Bill-by-bandwidth mode	Yes	Tencent Cloud provides two network billing modes. Configure a value greater than 0 Mbps as needed. **Bill-by-traffic**: Billing is based on traffic that is actually used. You can specify a peak bandwidth to prevent charges incurred by unexpected traffic. Packet loss will occur when the instantaneous bandwidth exceeds this value. This is applicable to scenarios where the network connection fluctuates significantly. **Bill-by-bandwidth package**: Select this aggregated billing mode when your public network instances have traffic peaks at different times. It is applicable to large-scale businesses where traffic can be staggered between different instances using the public network. BWP is currently in beta test. To try it out, submit a ticket for application. For more information, see Public Network Billing.
Bandwidth value	No	You can set the maximum public network bandwidth of the CVM instance as needed. For more information, see Public Network Bandwidth Cap.
Security group	Yes	If there is no available security group, you can choose New security group. If there are available security groups, you can choose Existing Security Groups. For more information on security groups, see Security Group.
Tag	No	You can add tags for the instance as needed, which can be used to categorize, search for, and aggregate cloud resources. For more information, see Overview.
Instance name	No	You can customize the name of the CVM instance to be created. If no instance name is specified, Unnamed will be used by default. An instance name can contain up to 128 characters. Batch sequential naming or pattern string-based naming is also supported. **Note**: This name is displayed only in the console. It is not the hostname of the CVM instance.
Login Methods	Yes	Configure the method to log in to the CVM as needed. **Set Password**: Customize the password for logging in to the instance. **SSH Key Pair (only for Linux instances)**: Associate the instance with an SSH key to ensure secure login to the CVM instance.If no key is available or existing

		keys are inappropriate, click Create Now to create a key. For more information on SSH keys, see SSH Keys. **Random Password**: A password will be automatically generated and sent to you in Message Center.
Instance Termination Protection	No	It is not enabled by default. You can enable it as needed. Then, you cannot terminate an instance in the console or via the API. For more information, see Enabling Instance Termination Protection.
Security Enhancement	No	By default, Anti-DDoS and Cloud Workload Protection are enabled free of charge to help you build a CVM security system to prevent data leakage.
Tencent Cloud Observability Platform	No	CM is activated by default. You can install add-ons to get CVM monitoring metrics and display them in visual charts. You can also specify custom alarm thresholds. In addition, you can configure three-dimensional CVM data monitoring, smart data analysis, real-time fault alarms, and custom data reports to precisely monitor Tencent Cloud services and the health conditions of CVM instances.
Advanced Settings	No	Configure additional settings for the instance as needed. **Hostname**: You can customize the name of the computer in the CVM operating system. After a CVM instance is created, you can log in to it to view the hostname. **Project**: The default project is selected. You can select an existing project as needed to manage different CVM instances. **CAM Role**: You can set a role and use it to grant a role entity the permissions to access CVM services and resources and perform operations in Tencent Cloud. For detailed directions, see Managing Roles. **Placement Group**: You can add the instances to placement groups to improve your business availability. For detailed directions, see Placement Group. **Custom Data**: You can configure an instance by specifying custom data, and the configured scripts will run when an instance is started. If multiple CVM instances are purchased at a time, the custom data will run on all of them. The Linux operating system supports the Shell format, while the Windows operating system supports the PowerShell format and a maximum of 16 KB of raw data. For more information, see Configuring Custom Data (Linux CVM). **Note**: Custom data configuration applies only to certain public images with the cloud-init service. For more information, see Cloud-Init & Cloudbase-Init.

5. Click **Next: Confirm Configuration** to enter the configuration information confirmation page.

6. Validate the information of the CVM to be purchased and the cost details of each configuration item.

7. Read and indicate your consent to the **Tencent Cloud Terms of Service**.

8. You can perform the following operations as needed:

Select **Save as Launch Template** to save the configuration of this instance as a launch template, based on which you can quickly create instances. For more information, see Managing Instance Startup Template.

Select **Generate API Explorer Reusable Script** to generate the OpenAPI reusable script code for instance creation corresponding to the selected configuration. You can save the code for purchasing CVM instances with the same configuration. For more information, see Generating API Explorer Reusable Scripts to Create Instances.

9. Click **Buy Now** or **Activate** and make the payment.

After making the payment, you can log in to the CVM console to check your CVM instance.

Information such as the instance name, public IP address, private IP address, login username, and initial login password of the CVM will be sent to your account through the Message Center. You can use this information to log in to and manage your instances. To ensure the security of your CVM, please change your CVM login password as soon as possible.

Create Instances via Custom Image

Last updated：2024-01-08 09:32:02

Overview

You can use a custom image to create CVM instances of the same operating system, applications, and data to improve efficiency. This document guides you through how to create an instance using a custom image.

Preparations

You must have a custom image under your account and in the region where you want to create an instance.
If there is no custom image, see the following solutions:

Image Status	Solution
Images on local computers or other platforms	Import the system disk image on local computers or other platforms to the custom image on CVM. For more information, see Overview.
There are template instances but no custom images	For more information, see Creating Custom Images.
Custom images in other regions	Copy the custom image to the target region where you want to create an instance. For more information, see Copying Images.
Custom images under another account	Share the custom image with the account under which you want to create an instance. For more information, see Sharing Custom Images.

Directions

1. Log in to the CVM console.
2. Click **Image** on the left sidebar to enter the image management page.
3. Select a region at the top of the **Image** page.
4. Select a tab based on the image source to view its image list.
Public Image: Go to the public image page.
Custom Image: Go to the custom image page.
Shared Image: Go to the shared image page.
5. In the **Operation** column of the target image, click **Create Instance**.

6. In the pop-up window, click **OK**.

7. Configure and create the instance as prompted by the page.

The **Region** and **Image** fields are automatically filled. Complete the other configurations of the instance as needed.

For more information, see Creating Instances via CVM Purchase Page.

Note:

If you use a custom image that contains one or more data disk snapshots, the system will automatically create the same quantity of cloud disks as data disks and the same capacity as each snapshot. You can increase, but cannot reduce, the cloud disk capacity.

References

You can also call the RunInstances API to create an instance by using a custom image.

Note:

If you use an image of the entire CVM instance to create an instance, call the DescribeImages API to get the snapshot ID associated with the image first and then call the `RunInstances` API to pass in the snapshot ID parameter; otherwise, the created cloud disk cannot match the snapshot ID, the snapshot data cannot be rolled back, the data disk has no data, and mounting cannot be performed.

Purchasing Similar Instances

Last updated：2024-01-08 09:32:02

Overview

You can use the "purchase with same configuration" or "instance startup template" features in the CVM console to create a CVM instance quickly, so as to save your time and improve the horizontal scaling efficiency in certain scenarios.

Directions

Creating instances with same configuration

1. Log in to the CVM console.
2. Select a region at the top of the **Instances** page.
3. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Find the target instance and click **More** > **Purchase with Same Configuration** in the **Operation** column as shown below:

On the page of the target instance, select **More Actions** > **Purchase with Same Configurations** in the top-right corner as shown below:

4. Enter the quantity of CVMs you want to purchase and check the other automatically selected configurations.

You can adjust the parameter configurations based on your actual needs.

5. Read and click **Tencent Cloud Terms of Service** and/or **Refund Policy**.

6. Click **Buy Now** or **Activate** and make the payment.

Using instance startup template to create instance

You can use an existing instance startup template to create an instance quickly. For more information, see Creating Instance from Instance Startup Template.

Generating API Explorer Reusable Scripts to Create Instances

Last updated：2024-01-08 09:32:02

Overview

While purchasing CVMs on the CVM purchase page, you can generate the OpenAPI best practice reusable scripts with the selected configurations. You can then use these codes to purchase CVM instances with the same configurations.

Prerequisites

You have logged in to the Tencent Cloud console and accessed the CVM **Custom Configuration** page.
You have completed CVM configurations and entered the **Confirm Configuration** page. To learn about how to configure parameters, see Creating Instances via CVM Purchase Page.

Directions

1. On the **Confirm Configuration** page, click **Generate API Explorer Reusable Scripts** as shown below:

2. You can view the following information in the pop-up window.

API Workflow: provides the description and actual parameters of the `RunInstances` API based on the selected configurations. The parameters marked with "*" are required for the API. You can hover over the data to display it completely.

API Script: generates codes in Java and Python programing languages. Select the Java or Python tab as needed, click **Copy Script** in the top-right corner, and save the codes to purchase CVM instances that contain the same configurations.

Note:

The instance password will not be displayed on the page or script codes for security reasons. Please modify it by yourself.

The collective expiry date cannot be set in the API Explorer reusable script. You need to set it after creating the CVM.

Enable Model Comparison

Last updated：2024-07-05 20:19:26

Operation Scenarios

When purchasing CVM, you can use the **model comparison** tool to compare configuration parameters, performance indicators, and prices for multiple models. Based on the selected model configuration, it will intelligently recommend the best model to help you efficiently choose the right CVM.

Operation Step

1. Log in to Tencent Cloud official website and go to Custom Configuration Purchase Page of the CVM.
2. Enable the **model comparison function** to add model specifications. For details, you can see the follwing figure.

3. Click **Compare** to open the model comparison page. For details, you can see the follwing figure.

You can add multiple models for horizontal comparison of all parameters.

You can hide common items and highlight differential items.

Based on your selected models, the system intelligently recommends the best model.

You can export the comparison results for easy local archiving and sharing.

4. Click **Purchase** to enter the CVM purchase page to purchase CVMs.

Managing Instance Launch Template

Last updated：2024-05-17 10:55:19

Overview

Instance launch template stores the required configuration information (except the instance password) for creating a CVM instance. You can use the specified instance launch template to quickly create an instance to improve the efficiency and user experience. This document describes how to create, manage, and use an instance launch template in the CVM console to quickly create an instance.

Instructions

After an instance launch template is created successfully, its configuration cannot be modified.
You can create one or multiple versions of an instance launch template and set different configurations for each version. You can also specify the default version and the default configuration of it will be used when you use the template to create an instance.

Directions

Creating and viewing instance template

1. Log in to the CVM console and select **Launch Templates** in the left navigation bar.
2. On the **Instance launch template** page, click **Create template**.
3. On the **Instance Startup Template** page, you can fill in the **Template name** and the **Template description** as needed. For the remaining configuration, please refer to Creating Instances via the Purchase Page.
4. In the Confirm configurations step, read and check **I have read and agree Tencent Cloud Service Terms and Purchasing Channels**, then click **Create Now**.

After successful creation, you can view the instance launch template in the console, which is shown below:

You can click the template ID to enter the template details page and view the specific information.

Creating instance launch template version

1. In the **Instance launch templates** page, click **Create Version** to create a new version for the needed templates on the right of the row, which is shown below:

2. Enter the **Instance Startup Template** page, see Creating Instances via CVM Purchase Page for settings.

3. In the **Confirm configurations** step, read and check **I have read and agree Tencent Cloud Service Terms and Purchasing Channels.**

You can choose to **Compare with original** and confirm the differences between the new version and the original instance launch template in the pop-up **Compare with original** window, which is shown below:

4. After confirmation, click **Create Now**.

After successful creation, in the **Instance Startup Template** page, click

▶

 in front of the row where the template is located to view the version in the expanded list.

Specifying default instance launch template version

1. In the **Instance launch template** page, click

▶

 in front of the row where the template is located.

2. In the expanded list, click **Set as default** on the right of the version you need to set, which is shown below:

3. In the pop-up **Set default template** window, click **OK**.

Using an instance launch template to create an instance

1. In the **Instance launch template** page, select **Create instance** on the right of the row where the template is located.

Note:

The configuration of the **default version** of the instance launch template is used for creating a new instance. You can also click

 in front of the row where the template is located in the expanded list, and select other versions to create an instance.

2. In the **Confirm configurations** step on the **Cloud Virtual Machine (CVM)** creation page, you can select **Compare with original** and confirm the differences between the instance and the instance launch template in the pop-up **Compare with original** window.

3. After confirming it, read and check **I have read and agree Tencent Cloud Service Terms and Purchasing Channels**, and click **Activate**.

Deleting instance launch template

1. In the **Instance launch template** page, select **Delete** on the right side row where the required instance launch template is located.

2. In the pop-up **Delete** window, click **OK**.

Documentation

[Creating Instances via CVM Purchase Page](#)

Batch Sequential Naming or Pattern String-Based Naming

Last updated：2024-01-08 09:32:02

Overview

To allow you to name batch created instances/hosts according to a rule during creation, the features of automatically incrementing suffixed numbers and specifying pattern strings are provided.

When you need to purchase n instances and generate instance/host names in specific forms, such as "CVM+Sequence number" (for example, CVM 1, CVM 2, and CVM 3), you can use the feature of Automatically Ascending Suffixed Numbers.

When you need to create **n** instances and name specific instances/hosts with ascending numbers starting from **x**, you can use the feature of Specifying a Single Pattern String.

When you need to create n instances/hosts with multiple prefixes in their names, each of which contains a specified serial number, you can use the feature of Specifying Multiple Pattern Strings.

Application Scope

This document applies to **setting instance name** and **setting host name**.

Directions

Note:

This document uses setting instance name as an example. The procedure may vary slightly according to the name type.

Automatically incrementing suffixed numbers

This feature allows you to name batch purchased instances with the same prefix and automatically ascending suffixed numbers.

Note:

The created instances are suffixed with numbers starting from 1 by default. You cannot specify the starting number. The following example assumes that you have purchased three instances and want to name these instances in the form of "CVM+Sequence number" (for example, CVM 1, CVM 2, and CVM 3).

Purchase page

API

1. Purchase three instances by referring to Creating Instances via CVM Purchase Page. On the **Configure network and host** tab page, enter the instance name in the form of **Prefix+Sequence number**. In this case, enter `CVM` as the instance name.

2. Follow the prompts on the page and complete payment.

In the RunInstances API, set the relevant fields:

Instance name: set `InstanceName` to `CVM` .

Host name: set `HostName` to `CVM` .

Specifying pattern string

This feature allows you to name batch purchased instances in a complex form with specified serial numbers. You can use one or more pattern strings in instance names as required.

The instance name with a specified pattern string is in the form of **{R:x}**, where **x** indicates the starting number in generated instance names.

Specifying one pattern string

The following example assumes that you want to create three instances and name them with ascending numbers starting from 3.

Purchase page

API

1. Purchase three instances by referring to Creating Instances via CVM Purchase Page. On the **Configure network and host** tab page, enter the instance name in the form of **Prefix+Specified pattern string {R:x}**. In this case, enter `CVM{R:3}` as the instance name.

2. Follow the prompts on the page and complete payment.

In the RunInstances API, set the relevant fields:

Instance name: set `InstanceName` to `CVM{R:3}` .

Host name: set `HostName` to `CVM{R:3}` .

Specifying multiple pattern strings

The following example assumes that you want to create three instances and name them with the **cvm**, **Big**, and **test** prefixes, where **cvm** and **Big** are followed by ascending numbers starting from 13 and 2, respectively. For example, their names are cvm13-Big2-test, cvm14-Big3-test, and cvm15-Big4-test, respectively.

Purchase page

API

1. Purchase three instances by referring to Creating Instances via CVM Purchase Page. On the **Configure network and host** tab page, enter the instance name in the form of **Prefix+Specified pattern string {R:x}-Prefix+Specified pattern string {R:x}-Prefix**. In this case, enter `cvm{R:13}-Big{R:2}-test` as the instance name.

2. Follow the prompts on the page and complete payment.

In the RunInstances API, set the relevant fields:

Instance name: set `InstanceName` to `cvm{R:13}-Big{R:2}-test` .

Host name: set `HostName` to `cvm{R:13}-Big{R:2}-test` .

Feature Verification

After you batch create instances through automatically incrementing suffixed numbers or specifying pattern string, you can verify the feature as follows:

Verifying instance name

Log in to the CVM console and view the newly created instances. You can see that the batch purchased instances are named according to the rule you set as shown below:

ID/Name	Monitoring	Availability Zone ▼	Instance Type ▼	Instance Configuration	Primary IPv6	Instance Billing Mode ▼	Network Billing Mode ▼	Project ▼	Operation
☐ [New] cvm15-Big4-test	📊	Nanjing Zone 1	Standard S5	1-core 2GB 1Mbps System disk: Premium Cloud Storage Network:	-	Pay as you go Created at 2021-03-11 16:33:47	Bill by traffic	Default Project	Log In Mor
☐ [New] cvm14-Big3-test	📊	Nanjing Zone 1	Standard S5	1-core 2GB 1Mbps System disk: Premium Cloud Storage Network	-	Pay as you go Created at 2021-03-11 16:33:44	Bill by traffic	Default Project	Log In Mor
☐ [New] cvm13-Big2-test	📊	Nanjing Zone 1	Standard S5	1-core 2GB 1Mbps System disk: Premium Cloud Storage Network:	-	Pay as you go Created at 2021-03-11 16:33:41	Bill by traffic	Default Project	Log In Mor

Verifying host name

1.

Restart and log in to the CVM instance.

2. Select different steps according to the instance's operating system:

Linux instance

Windows instance

On the operating system UI, run the following commands:

```
hostname
```

Open the command line tool and run the following command:

```
hostname
```

3.

View the returned result of the `hostname` command.

If the returned result is similar to the following, the setting is successful.

```
cvm13-Big2-test
```

4. Repeat step 1–step 3 to verify other batch purchased instances.

Logging In to Linux Instances
Logging In To Linux Instance (Web Shell)

Last updated：2024-01-08 09:32:02

Overview

WebShell is the login method recommended by Tencent Cloud. No matter your local OS is Windows, Linux or Mac OS, as long as you have purchased public IPs for your instances, you can log in via Web Shell. This document describes how to log in to a Linux instance via Web Shell.

Benefits of Web Shell:

Supports copy and paste operations with shortcut keys.

Supports scrolling with mouse wheel.

Supports Chinese input.

Features a high security (password or key is required for each login).

Authentication Method

Password or **Key**

Prerequisites

You already have the admin account and password (or key) to log in to the Linux instance.

If you have chosen to generate a random password when creating an instance, please get it from Message Center.

If you have set a login password, please use it for login. If you forgot it, please reset it.

If a key has been bound to the instance, you can use the key to log in. For more information, see SSH Keys.

You have purchased a public IP for your CVM instance and opened a remote login port (22 by default) for the WebShell proxy IP in the security group associated with the instance.

If you purchase a CVM instance through quick configuration, the port is opened by default.

If you purchase a CVM instance through custom configuration, you can manually open the port as instructed in Security Group Use Cases.

Directions

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Locate the Linux CVM instance you want to log in to and click **Log In** on the right as shown below:

Select the tab of the Linux CVM instance you want to log in to and click **Log In** as shown below:

3. In the **Standard Login** | **Linux Instance** pop-up window, select **Password Login** or **Key Login** as needed:

Refer to the following instructions to enter the required information for login:

Port: the default port is 22. Enter a value as needed.

Username: the default username of Linux instances is `root` , and the default username of Ubuntu instances is `ubuntu` . Enter a value as needed.
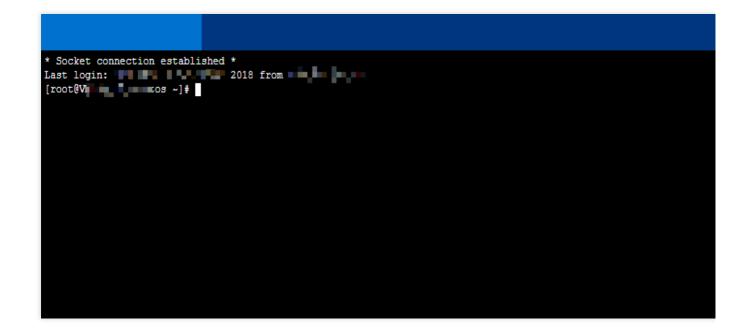
Password: enter the login password obtained in the Prerequisites step.

Key: select the key bound to the instance.

4. Click **Log In** to log in to the Linux instance.

If the login is successful, the following prompt will appear on the WebShell page:

Subsequent Operations

After logging in to the CVM, you can build a personal website or forum or perform other operations. For more information, see:

Manually Building a WordPress Website

Manually Building Discuz! Forum

See Also

Resetting Instance Password

Managing SSH Keys

Logging In To Linux Instances (Remote Login)

Last updated：2024-01-08 09:32:02

Overview

This document takes PuTTY as an example to describe how to log in to a Linux instance from Windows by using remote login software.

Applicable OS

Windows

Authentication Method

Password or **Key**

Prerequisites

You must already have the admin account and password (or key) to log in to the instance.

If you use a system default password to log in to the instance, go to Message Center to obtain the password first.

If you forgot your password, please reset your instance password.

A public IP has been purchased and obtained for your CVM instance, and port 22 is open (this is open by default for CVM purchased with quick configuration).

Directions

Password login

Key login

1. Download the Windows remote login software, PuTTY.

Click here to download PuTTy

2. Double-click **putty.exe** to open the PuTTY client.

3. In the **PuTTY Configuration** window, enter the following content, as shown below:

Configure parameters as follows:

Host Name (or IP address): the public IP of the CVM. Log in to the CVM console to obtain the public IP from the instance list and details pages.

Port: the port of the CVM, which must be "22".

Connection type: select **SSH**.

Saved Sessions: enter the session name, such as `test`.

After configuring **Host Name**, configure and save **Saved Sessions**. You can double-click the session name saved under **Saved Sessions** to log in to CVM.

4. Click **Open** to enter the **PuTTY** interface. The **login as:** command prompt appears.

5. Enter the username after **login as:** and press **Enter**.

6. Enter the password after **Password** and press **Enter**.

The entered password is not displayed by default, as shown below:

Once logged in, you can see the information about the CVM to which you are currently logged in on the left of the command prompt.

1. Download the Windows remote login software, PuTTY. Both putty.exe and puttygen.exe are required.

Download PuTTy

Download PuTTygen

2. Double-click **puttygen.exe** to open the PuTTY Key client.

3. Click **Load**, select and access the path where the downloaded private key is saved. You should download and keep your private key after creating a key pair. For more information, see Managing SSH Keys

For example, select and open the private key file `david` , as shown below:

4.

In the **PuTTY Key Generator** window

, enter the key name and the encrypted private key password (optional), and click **Save private key**, as shown below:

5. In the pop-up window, select the path where the key will be saved. In the **File name** field, enter "Key Name.ppk" and click **Save**. For example, save the private key file `david` as `david.ppk` , as shown below:

6. Double-click **putty.exe** to open the PuTTY client.

7. In the left sidebar, go to **Connection** > **SSH** > **Auth** and enter the **Auth** configuration interface.

8. Click **Browse**, and select and access the path where the key is saved, as shown below:

9. Switch to the **Session** configuration interface. Configure the CVM IP, port, and connection type, as shown below:

Host Name (or IP address): the public IP of the CVM. Log in to the CVM console to obtain the public IP from the instance list and details pages.

Port: the port of the CVM, which must be "22".

Connection type: select **SSH**.

Saved Sessions: enter the session name, such as `test` .

After configuring **Host Name**, configure and save **Saved Sessions**. You can double-click the session name saved under **Saved Sessions** to log in to CVM.

10. Click **Open** to enter the **PuTTY** interface. The **login as:** command prompt appears.

11. Enter the user name after **login as:** and press **Enter**.

12. Enter the password configured in Step 4 after **Passphrase for key "imported-openssh-key":** and press **Enter**. The entered password is not displayed by default, as shown below:

Once logged in, you can see the information about the CVM to which you are currently logged in on the left of the

command prompt.

Subsequent Operations

After logging in to the CVM, you can build a personal website or forum or perform other operations. For more information, see:

Setting up WordPress

Building Discuz! Forum

Logging In To Linux Instance (SSH Key)

Last updated：2024-01-08 09:32:02

Overview

This document describes how to use an SSH key to log in to a Linux instance from a local Linux, Mac OS or Windows server.

Supported Systems

Linux, Mac OS or Windows (including Windows 10 and Windows Server 2019)

Authentication Method

Password or **Key**

Prerequisites

You already have the admin account and password (or key) to log in to the instance.

The default admin account is usually `root` for the Linux instance, and is `ubuntu` for the Ubuntu system. You can modify it according to the actual situation.

If you use a system default password to log in to the instance, go to the Message Center to obtain the password first.

If you use a key to log in, you must have created a key and bound it to this CVM. For more information, see Managing SSH Keys.

If you forgot your password, please reset your instance password.

A public IP has been purchased for your CVM instance, and the port 22 is open. It is open by default for a CVM instance purchased with quick configuration.

Directions

Using the password

Using a key

1. Execute the following command to connect to the Linux CVM.

Note:

If your local computer uses Mac OS, you need to open the terminal that comes with the system before executing the following command.

If your local computer uses Linux, you can directly execute the following command.

If your local computer uses Windows 10 or Windows Server 2019, you need to open the command prompt CMD before executing the following command.

```
ssh <username>@<hostname or IP address>
```

`username` refers to the default account as mentioned in "Prerequisites".

`hostname or IP address` refers to the public IP address or custom domain name of your Linux instance.

2. Enter the password you have obtained, and click **Enter** to log in.

1. Execute the following command to set the private key file readable only to you.

If your local computer uses Mac OS, you need to open the terminal that comes with the system before executing the following command.

If your local computer uses Linux, you can directly execute the following command.

```
chmod 400 <The absolute path of the private key downloaded to be associated
with the CVM>
```

If your local computer uses Windows 10, you need to open the command prompt CMD before executing the following commands in sequence.

```
icacls <The absolute path of the private key downloaded to be associated with
the CVM> /grant <Windows user account>:F
```

```
icacls <The absolute path of the private key downloaded to be associated with
the CVM> /inheritancelevel:r
```

2. Execute the following command for remote login.

```
ssh -i <The absolute path of the private key downloaded to be associated with
the CVM> <username>@<hostname or IP address>
```

`username` refers to the default account as mentioned in "Prerequisites".

`hostname or IP address` refers to the public IP address or custom domain name of your Linux instance.

For example, execute the `ssh -i "Mac/Downloads/shawn_qcloud_stable.pem"` `ubuntu@192.168.11.123` command to remotely log in to the Linux CVM.

Subsequent Operations

After logging in to the CVM, you can build a personal website or forum or perform other operations. For more information, please see:

Manually Building WordPress Website

Manually Building Discuz! Forum

Logging In To Linux Instances (VNC)

Last updated：2024-01-08 09:32:02

Overview

VNC login provided by Tencent Cloud allows users to remotely log in to CVM via a web browser. If a client does not have remote login installed or it cannot be used, user can log in to the CVM using VNC login to check the CVM status and perform basic management operations using the CVM account.

Use Limits

VNC login currently does not support copy and paste, Chinese input method, and file upload or download.
When you use VNC to log in to CVM, mainstream browsers must be used, such as Chrome, Firefox, IE 10 and above.
VNC login is a dedicated terminal, meaning only one user can use VNC login at a time.

Prerequisites

You already have the admin account and password to log in to the instance.
If you have chosen to generate a random password when creating an instance, please get it from Message Center.
If you have set a login password, please use it for login. If you forgot it, please reset it.

Directions

1. Log in to the CVM console.
2. On the **Instances** page, locate the Linux CVM instance you want to log in to and click **Log In** as shown below:

3. In the **Standard Login | Linux Instance** window that is opened, select **login with VNC** as shown below:

4. In the opened window, enter the username after **login** and press **Enter**.
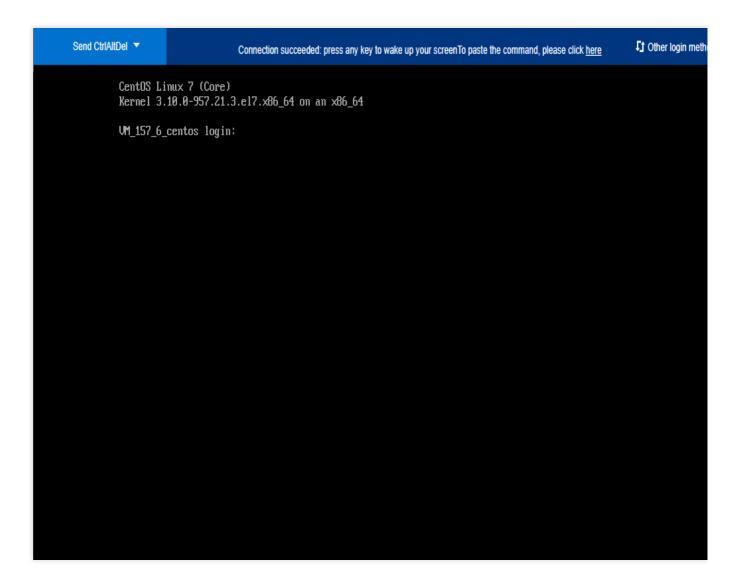
The default username of Linux instances is `root` , and the default username of Ubuntu instances is `ubuntu` .
Please enter as needed.

5. Enter the password after **Password** and press **Enter**.

The entered password is not displayed by default. After login, the information of the CVM that you are currently logged

in to will appear on the left of the command prompt as shown below:

Operations

After logging in to the CVM, you can build a personal website or forum or perform other operations. For more information, please see:

Common Operations and Commands

Manually Building WordPress Website

Manually Building Discuz! Forum

Logging In To Linux Instances (Mobile Devices)

Last updated：2024-01-08 09:32:02

Overview

This document describes how to log in to a Linux instance from different mobile devices. The following tools are used as an example.

iOS device: Termius-SSH client

Android device: JuiceSSH

Applicable Mobile Devices

iOS and Android devices

Prerequisites

The CVM instance is in the **Running** status.

You already have the admin account and password (or key) to log in to the instance.

If you use a system default password to log in to the instance, go to Message Center to obtain the password first.

If you've forgotten your password, you can reset the instance password.

A public IP has been purchased for your CVM instance, and the port 22 is open. It is open by default for a CVM instance purchased with quick configuration.

Directions

Log in to the instance from the mobile device you are using:

iOS device

Android device

1. Download the Termius-SSH client from the App Store, and register as instructed.

2. Tap **New Host** on the home screen.

3. Access the **New Host** page and configure the login information as follows:

Hostname: the public IP address of your CVM instance. For more information, see Getting Public IP Addresses.

Use SSH: enabled by default.

Username: enter the admin account `root`, or `ubuntu` if your instance uses the Ubuntu operating system.

Password: enter the login password of the instance.

4. Tap **Save** in the upper-right corner to save the login configuration.

5. Select the login information on the **Hosts** page and tap **Continue** in the prompt box at the bottom of the page.

6. Login succeeds if you see the following.

Creating an identity

1. Download and install JuiceSSH.

2. From the home screen, tap **Connections** to reach the **Identities** tab.

3. Tap **+** in the lower-right corner.

4. Configure the account name and password on the **Identity** page.

Nickname: enter a custom name for the identity, optional.

Username: enter the admin account `root` , or `ubuntu` if your instance uses the Ubuntu operating system.

Password: tap **Set (optional)** and enter the instance login password in the pop-up window.

5. Tap ✔ in the upper-right corner of the page.

Creating a connection

1. From the home screen, tap **Connections**, then tap **+** in the lower-right corner of the **Connections** page.

2. Configure the login information for the new connection.

Nickname: enter a custom connection name, optional.

Type: select **SSH**.

Address: the public IP address of your CVM instance. For more information, see Getting Public IP Addresses.

Identity: select the identity created in Creating an identity.

Port: enter the port 22.

 Retain the default settings for other parameters.
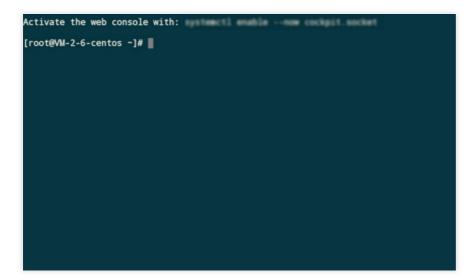
3. Tap **Add to team** in the bottom of the page to save the login configuration.

Logging in to the instance

1. On the **Connections** page, select the instance to log in and tap **Accept**.

2. Login succeeds if you see the following.

Logging in to Windows instance Logging in Using Standard Method (Recommended)

Last updated：2024-01-08 09:32:02

Overview

This document describes how to log in to a Windows instance using the standard login method (WebRDP).
Note:
This method does not vary by the local operating system and supports direct login to the Windows instance in the console.

Prerequisites

You must have the admin account and password for logging in to a Windows instance remotely.

If you have set a login password, please use it for login. If you forgot it, please reset it.

If you have chosen to generate a random password when creating an instance, please get it from Message Center.

You have purchased a public IP for your CVM instance and opened a remote login port (3389 by default) for the WebRDP proxy IP in the security group associated with the instance.

If you purchase a CVM instance through quick configuration, the port is opened by default.

If you purchase a CVM instance through custom configuration, you can manually open the port as instructed in Security Group Use Cases.

Make sure that the public network bandwidth of your instance is ≥ 5 Mbit/s; otherwise, the remote desktop may lag. To adjust the network bandwidth, please see Adjusting Network Configuration.

Directions

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Locate the Windows CVM instance you want to log in to and click **Log In** on the right as shown below:

Select the tab of the Windows CVM instance you want to log in to and click **Log In** as shown below:

3. In the **Standard Login | Windows Instance** window that is opened, enter the login information according to the actual situation.

Port: the default port is 3389. Enter a value as needed.

Username: the default username of Windows instances is `Administrator` . Enter a value as needed.

Password: enter the login password obtained in the Prerequisites step.

4. Click **Log In** to log in to the Windows instance.

This document uses logging in to a CVM instance on Windows Server 2016 Datacenter Edition 64-bit as an example.

If the login is successful, a page similar to the following will appear:

Relevant Documentation

Resetting Instance Password

Adjusting Network Configuration

Logging In to Windows Instance Using RDP

Last updated：2024-01-08 09:32:02

Note:

Currently, the **standard login method (WebRDP)** is used for Windows instances by default. It allows you to log in to a Windows instance in the console without downloading a local login client. For the login method, see Logging in to Windows Instance Using Standard Login Method.

Overview

Remote Desktop Protocol (RDP) is a multiple-channel protocol developed by Microsoft that allows a local computer to connect to a remote computer. We recommend you use RDP to log in to your Windows CVMs. This document describes how to log in to Windows instances using RDP files.

Supported Systems

You can log in to your CVMs from Windows, Linux, and MacOS using RDP.

Prerequisites

You must have the admin account and password for logging in to a Windows instance remotely.

If you have chosen to generate a random password when creating an instance, please get it from Message Center.

If you have set a login password, please use it for login. If you forgot it, please reset it.

You have purchased a public IP for your CVM instance and opened a remote login port (3389 by default) for the WebRDP proxy IP in the security group associated with the instance.

If you purchase a CVM instance through quick configuration, the port is opened by default.

If you purchase a CVM instance through custom configuration, you can manually open the port as instructed in Security Group Use Cases.

Make sure that the public network bandwidth of your instance is ≥ 5 Mbit/s; otherwise, the remote desktop may lag. To adjust the network bandwidth, please see Adjusting Network Configuration.

Directions

Logging in to your Windows CVM using RDP

Logging in to your Linux CVM using RDP

Logging into you MacOS CVM using RDP

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List mode: locate the Windows CVM instance you want to log in to and click **Log In** on the right as shown below:

Tab mode: select the tab of the Windows CVM instance you want to log in to and click **Log In** as shown below:

3. In the **Standard Login | Windows Instance** window that is opened, select **Download RDP File**.

Note:

If you have changed the remote login port, append the IP address with `:port` in the RDP file.

4. Double-Click the downloaded RDP file, enter the password, and click **OK** to remotely connect to your Windows CVM.

If you use a system default password to log in to the instance, you can obtain the password at the Message Center.

If you forgot your password, please reset the instance password.

Note:

We recommend you use rdesktop as the remote desktop client. For more information, see the official introduction to rdesktop.

1. Run the following command to check whether rdesktop has been installed.

```
rdesktop
```

If yes, perform step 4.

If no, you will be prompted with "command not found". In this case, perform step 2.

2.

Open a terminal

 window and run the following command to download rdesktop. This step uses rdesktop v1.8.3 as an example.

```
wget https://github.com/rdesktop/rdesktop/releases/download/v1.8.3/rdesktop-
1.8.3.tar.gz
```

If you want to install the latest version, visit the rdesktop page on GitHub to find it. Then replace the path in the command with that of the latest version.

3. In the directory where rdesktop will be installed, run the following commands to decompress and install rdesktop.

```
tar xvzf rdesktop-<x.x.x>.tar.gz ## Replace x.x.x with the version number of
the downloaded rdesktop.
cd rdesktop-1.8.3
./configure
make
make install
```

4.

Run the following command to connect to the remote Windows instance.

Note:

Replace the parameters in the example with your own parameters.

```
rdesktop -u Administrator -p <your-password> <hostname or IP address>
```

`Administrator` refers to the admin account mentioned in the prerequisites section.

`<your-password>` refers to the login password that you set.

If you use a system default password to log in to the instance, you can obtain the password at the Message Center. If you forgot your password, please reset the instance password.

`<hostname or IP address>` is the public IP or custom domain name of your Windows instance. For more information on how to get the public IP, please see Getting Public IP Addresses.

Note:
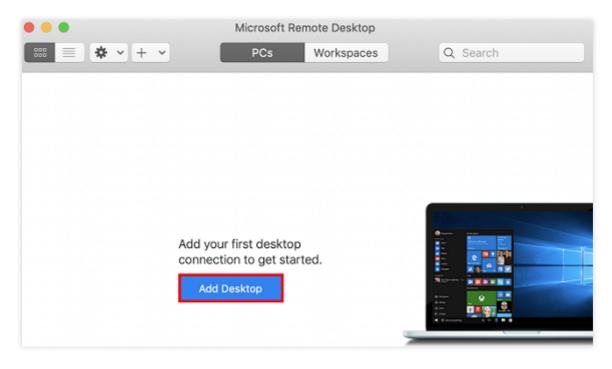
The following operations use Microsoft Remote Desktop for Mac as an example. Microsoft stopped providing a link to download the Remote Desktop client in 2017. Currently, its subsidiary HockeyApp is responsible for releasing the beta client. Go to Microsoft Remote Desktop Beta to download a Beta version.

The following operations use a CVM on Windows Server 2012 R2 as an example.

1. Download and install Microsoft Remote Desktop for Mac on your local computer.

2. Start MRD and click **Add Desktop**, as shown below:

3. In the **Add PC** pop-up window, follow the steps illustrated in the following image to establish a connection to your Windows CVM.

3.1 In the **PC name** text file, enter the public IP address of your CVM instance. For more information on how to obtain the public IP address, see Getting Public IP Addresses.

3.2 Click **Add**.

3.3 Retain the default settings for the other options and establish the connection.
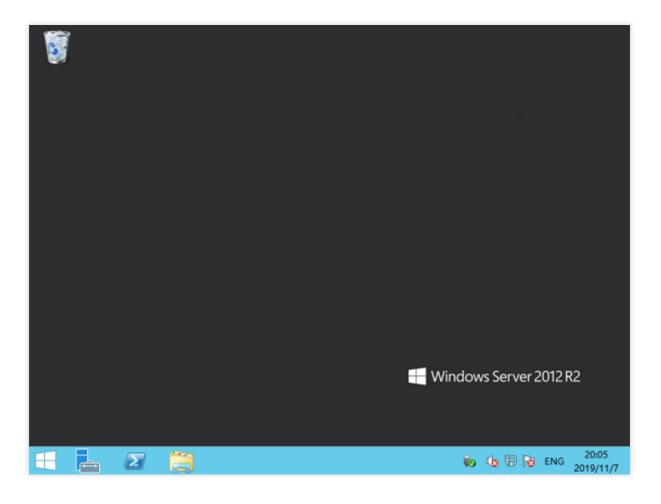
 Your entry has now been saved, as shown below:

4. Double-click the new entry. Input your username and password for CVM and click **Continue**.

5. If you use a system default password to log in to the instance, you can obtain the password at the Message Center.

6. If you forgot your password, please reset the instance password.

7. In the pop-up window, click **Continue** to establish the connection, as shown below:

If the connection is successful, the following page will appear:

RDP Bandwidth Limit Description

The available network bandwidth directly affects the experience of logging in to and using CVM instances over RDP, and different applications and display resolutions require different network configurations. Microsoft has laid down the minimum bandwidth requirements for instances when using RDP in different application scenarios. Please check out the following table to make sure that the network configuration of your instance can meet your business needs; otherwise, issues such as lag may occur.

Note:

To adjust the bandwidth of your instance, please see Adjusting Network Configuration.

These numbers apply to a single monitor configuration with 1920x1080 resolution and with both default graphics mode and H.264/AVC 444 graphics mode.

Scenario	Default Mode	H.264/AVC 444 Mode	Description
Idle	0.3 Kbps	0.3 Kbps	User has paused their work, and there's no active screen updates.
Microsoft Word	100–150	200–300 Kbps	User is actively working with Microsoft Word,

	Kbps		typing, pasting graphics, and switching between documents.
Microsoft Excel	150–200 Kbps	400–500 Kbps	User is actively working with Microsoft Excel and updating multiple cells with formulas and charts simultaneously.
Microsoft PowerPoint	4–4.5 Mbps	1.6–1.8 Mbps	User is actively working with Microsoft PowerPoint, typing, and pasting. User is also modifying rich graphics and using slide transition effects.
Web browsing	6–6.5 Mbps	0.9–1 Mbps	User is actively working with a graphically rich website that contains multiple static and animated images. User scrolls the pages both horizontally and vertically.
Image gallery	3.3–3.6 Mbps	0.7–0.8 Mbps	User is actively working with the image gallery application, browsing, zooming, resizing, and rotating images.
Video playback	8.5–9.5 Mbps	2.5–2.8 Mbps	User is watching a 30 FPS video that consumes 1/2 of the screen.
Fullscreen video playback	7.5–8.5 Mbps	2.5–3.1 Mbps	User is watching a 30 FPS video that is maximized to a fullscreen.

Logging into Windows Instance via Remote Desktop

Last updated：2024-01-08 09:32:02

Overview

This document describes how to log in to a Windows instance through remote desktop on a local computer.

Supported Systems

Windows

Prerequisites

You must have the admin account and password for logging in to a Windows instance remotely.

If you use a system default password to log in to the instance, you can obtain the password at the Message Center.

If you forgot your password, please reset the instance password.

You have purchased public IPs for your CVM instance and port 3389 is open (this port is open by default for a CVM purchased with quick configuration).

Directions

Note:

The following takes the Windows 7 operating system as an example.

1. On the local Window server, click

, enter **mstsc** in **Search programs and files**, and press **Enter** to open the **Remote Desktop Connection** window as shown below:

2. Enter the public IP of the Windows server after **Computer** and click **Connect**. You can get the server public IP as instructed in Getting Public IP Address.

3. Enter the instance's admin account and password in the **Windows Security** pop-up window as shown below:

Note:

If the **Do you trust this remote connection?** window pops up, you can select **Don't ask me again for connections to this computer** and click **Connect**.

The default admin account of the Windows CVM instance is `Administrator` , and the password can be obtained as instructed in Prerequisites.

4. Click **OK**.

Logging into Windows Instance via VNC

Last updated：2024-01-08 09:32:02

Overview

VNC login provided by Tencent Cloud allows users to remotely log in to CVM via a web browser. If a client does not have remote login installed or it cannot be used, user can log in to the CVM using VNC login to check the CVM status and perform basic management operations using the CVM account.

Use Limits

VNC login currently does not support copy and paste, Chinese input method, and file upload or download.
When you use VNC to log in to CVM, mainstream browsers must be used, such as Chrome, Firefox, IE 10 and above.
VNC login is a dedicated terminal, meaning only one user can use VNC login at a time.

Prerequisites

You must already have admin account/password for logging into Windows instance remotely.
If you have chosen to generate a random password when creating an instance, please get it from Message Center.
If you have set a login password, please use it for login. If you forgot it, please reset it.

Directions

1. Log in to the CVM console.
2. On the **Instances** page, locate the Windows CVM instance you want to log in to and click **Log In** as shown below:

3. In the **Standard Login | Windows Instance** window that is opened, select **login with VNC** as shown below:

4. In the login window that pops up, select **Send CtrlAltDel** in the upper-left corner and press **Ctrl-Alt-Delete** to open the system login window as shown below:

5. Enter the login password and press **Enter** to log in to the Windows CVM instance.

Logging in to a Windows Instance from Mobile Devices

Last updated：2024-01-08 09:32:02

Overview

This document describes how to log in to a Windows instance from different mobile devices using Microsoft Remote Desktop.

Applicable Mobile Devices

iOS and Android devices

Prerequisites

The CVM instance is in the **Running** status.

You already have the administrator account and password to log in to the instance.

If you use a system default password to log in to the instance, go to Message Center to obtain the password first.

If you've forgotten your password, you can reset the instance password.

A public IP has been purchased for your CVM instance, and the port 3389 is open. It is open by default for a CVM instance purchased with quick configuration.

Directions

Note:

 This document uses the iOS device as an example. Steps for Android devices are almost the same.

1. Download Microsoft Remote Desktop and start it.

2. In the **PCs** page, tap **+** in the upper-right corner, then tap **Add PC**.

3. Configure the login information to add a PC.

PC name: the public IP address of your CVM instance. For more information, see Getting Public IP Addresses.

User account: by default, **Ask when required** is selected.

4. Tap **Save**.

5. In the **PCs** page, select the instance to log in and enter its administrator account and password.

User name: enter the administrator account `Administrator` .

Password: enter the instance login password.

6. Tap **Continue**. If the page shown in the following figure is displayed, the login succeeds.

Adjusting Configuration

Changing Instance Configuration

Last updated：2024-01-08 09:32:02

Overview

Hardware devices of Tencent Cloud CVM instances can be adjusted quickly and flexibly. This document describes the operation methods for configuration upgrade, downgrade, and cross-model adjustment.

Prerequisites

You can adjust the configuration of an instance when it is in shutdown or running status. If the instance is running, the adjustment takes effect after it is forcibly shut down and restarted.

Note:

If the instance has been **shut down**, you can adjust its configuration directly via the console.

If the instance is **running**, you can adjust its configuration online and confirm to forcibly shut down the instance. The adjustment takes effect after the instance is restarted.

You can adjust the configurations of instances online **in batches**. If an instance in the batch operation is **running**, you need to force the instance to shut down. The adjustment takes effects after the instance is restarted.

Limits and Impacts

Configuration adjustment limits

Only instance **whose system and data disks are both CBS cloud disks** supports configuration adjustment.

Configuration upgrade:

The number of configuration upgrades is unlimited and the upgrade takes effect immediately.

Configuration downgrade:

Pay-as-you-go instances can be downgraded any number of times at any time.

Adjustment across instance families: configurations can be adjusted between instance families without the need for data migration.

During configuration adjustment, instance specifications that can be adjusted are related to the target specifications available in the current AZ. Pay attention to the following restrictions:

Spot instances do not support cross-model configuration adjustment.

Dedicated instances do not support cross-model configuration adjustment. The adjustment scope is subject to the remaining resources of the dedicated host where the instance is located.

Heterogeneous instances such as GPU and FPGA instances cannot be used as the source or target instance type for configuration adjustment across instance families.

Instances configured with a classic network cannot be adjusted to instances that only support VPC.

If the target instance type does not support the CBS disk type configured for the current instance type, the configuration cannot be adjusted.

If the target instance type does not support the image type configured for the current instance type, the configuration cannot be adjusted.

If the target instance type does not support the ENI or ENI quantity configured for the current instance type, the configuration cannot be adjusted. For more information, see Use Limits.

If the target instance type does not support the public network bandwidth cap configured for the current instance type, the configuration cannot be adjusted. For more information, see Public Network Bandwidth Cap.

Impacts

The private IP addresses of few instances may change after adjustment. If any private IP address changes, the relevant information will be displayed on the adjustment page. If no such information is displayed, no private IP address has changed.

Directions

Note:

If your business changes, you can adjust the instance configuration.

During configuration upgrade, upgrade your CVM instance accordingly and pay for fees that may be incurred.

During configuration downgrade, confirm the refund detail and forcibly shut down and restart your CVM instance for the new configuration to take effect immediately.

Via console

Via API

Adjusting the configuration of a single instance

1. Log in to the CVM console and click **Instances** to view the CVM instance list.

2. Proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Resource Adjustment** > **Adjust Model and Specs** as shown below:

Tab view: on the page of the target instance, select **More Actions** > **Resource Adjustment** > **Adjust Model and Specs** in the top-right corner as shown below:

3. In the "Select target configuration" step, confirm the instance status and operation, **select the required model and specifications, confirm the performance parameters**, and click **Next**, as shown in the following figure:

4. Based on the instance billing method, confirm the fees and click **Next**.

Pay-as-you-go instances: confirm the amount to be frozen for the new instance type. After configuration adjustment, pay-as-you-go instances are charged starting from the tier-1 price. Confirm the billing rules, as shown in the following figure:

5. In the "Shutdown CVM" step, read the prompt carefully based on the instance running status.

If the current instance is running, read the prompt carefully and select "Agree to a forced shutdown", as shown in the following figure:

If the current instance is shut down, the following prompt will appear:

6. Click **Adjust Now** to go to the order page and complete the payment.

You can use the ResetInstancesType API to adjust the instance configuration. For more information, see the ResetInstancesType API documentation.

Adjusting Network Configuration

Last updated：2024-01-08 09:32:02

Overview

Tencent Cloud allows you to change the public network billing mode or public network bandwidth as needed. The change takes effect immediately. To learn more about the restrictions and price, see Adjusting Public Network Billing.

Directions

1. Log in to the CVM console. At the top of the **Instances** page, select the region where the target CVM instance resides.

2. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Select **More** > **Resource Adjustment** > **Adjust Network** on the right of the target CVM instance as shown below:

Select **More Actions** > **Resource Adjustment** > **Adjust Network** in the top-right corner of the page of the target instance as shown below:

3. In the **Adjust Network** pop-up window, adjust the public network billing mode or public network bandwidth as needed:

Network billing mode: Tencent Cloud provides two network billing modes: **bill-by-traffic** and **bill-by-bandwidth**. The bill-by-bandwidth mode is **hourly postpaid**.

Target bandwidth cap: Tencent Cloud provides two network configurations: **dedicated public network** and **shared public network** (billed by bandwidth package and currently in beta test). This document takes adjusting the configuration of the dedicated public network as an example, i.e., adjusting the bandwidth cap of a single CVM instance.

Note:

For more information about the bandwidth cap, see Public Network Bandwidth Cap.

4. Select the target billing mode or set the target bandwidth value and click **OK**.

Relevant Documentation

Adjusting Public Network Billing

Public Network Billing

Billing Modes

Public Network Bandwidth Cap

Adjusting Project Configuration

Last updated：2024-01-08 09:32:02

Overview

This project feature is used to manage cloud resources by project. When a CVM instance is created, it must be assigned to a project. Tencent Cloud allows users to reassign an instance to a new project after the instance is created.

Note:

To assign an instance to a new project, create a project first.For more information on how to create a project, refer to New Project.

Directions

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

In the instance list, select the target CVM instance and click **More** > **Instance Settings** > **Assign to Project** on the right as shown below:

Note:

If you want to reassign multiple CVM instances to a new project, select them and click **More Actions** > **Instance Settings** > **Assign to Project** on the top of the page.

On the page of the target CVM instance, select **More Actions** > **Instance Settings** > **Assign to Project** in the top-right corner as shown below:

3. In the **Assign to Project** pop-up window, select the name of the new project and click **Submit**.

Assign to Project

You've selected 1 instance. Collapse

ID/Name	Instance Type	Instance Configuration
	GPU Compute GN6S	4-core 20GB 1Mbps System disk: Premium Cloud Storage Network:vpc-mzg9lleo

Search by project name/description 🔍

Project Name	Description

OK   Close

Suggestions on Adjusting Instance Configuration

Last updated：2024-01-08 09:32:02

Overview

Tencent Cloud can analyze your CVM instance load based on the monitoring metrics including CPU and memory utilization collected by CM in the past three days and give suggestions on how to adjust the instance configuration. You can determine whether to adjust the instance configuration based on the actual conditions.

Notes

The instance configuration adjustment suggestions are made based on the average load data (collected once every 5 minutes) for the last three days and are applicable to instances with a stable load rather than those with CPU or memory utilization spikes.

This feature isn't supported for heterogeneous models such as GPU and FPGA as well as CPM. You can create alarms to actively monitor the instance usage.

The suggestions are for reference only. If you have high requirements for instance usage monitoring, we recommend you use CM for active monitoring.

Directions

1. Log in to the CVM console and enter the instance list page.

2. On the instance list page, if the

warning icon is displayed in the monitoring column of an instance, configuration adjustment suggestions have been given for it.

3. Click the

warning icon, and the **Configuration Adjustment Suggestion** window pops up.

4. In the **Configuration Adjustment Suggestions** window, you can view the target model recommended based on the current instance usage and select **Show More** to view other recommended models.

5. If you want to adjust the instance configuration according to the suggestions, select **I have read and agree to the description of instance configuration fees** and click **Start Adjustment**.

View Instance Details

Viewing Instance Information

Last updated：2024-01-08 09:32:02

Overview

Tencent Cloud provides the following three options for you to view the information of a CVM instance:

View the total number of CVM instances under your account and their status, as well as the quantity and quota of resources in each region on the Overview page of the CVM console.

View the information of all CVM instances in a region on the Instances page on the CVM console.

View the details of a CVM instance on the instance details page.

Prerequisites

You have logged in to CVM console.

Directions

Viewing the CVM instance overview

Select **Overview** on the left sidebar to enter the CVM overview page.

In this page, you can view the following information and perform the following operations:

CVM status: the total number of CVMs, the number of instances that expire within the next 7 days, the number of instances in Recycle Bin, and the number of normal CVMs.

List of CVMs to be renewed (you can renew them on this page).

Resource quantity and quota: you can view the quotas of pay-as-you-go CVMs, custom images and snapshots. You can also apply for quotas on this page.

Perform cross-region search for cloud resources.

Viewing the CVM instance list

Select **Instances** on the left sidebar to enter the instance list page, as shown below:

The information available on this page includes CVM ID and name, monitoring information, status, availability zone, instance type, instance configuration, primary IPv4, primary IPv6, instance billing, network billing, and the project to which the CVM belongs.

Note:

You can adapt to your actual needs Switching Instance Page View in Console。

You can click

in the top-right corner to configure in the pop-up "Display Settings" window the details you want to display, as shown below:

Viewing instance details

1. Go to the Instances page to select the region at the top.

2. Find the instance for which you want to view its details, and click the instance ID or name to enter the details page, as shown below:

On the instance details page, you can view information such as CVM information, architecture, network information, specifications, image information, billing information, ENI, monitoring, security groups, operation logs, and more.

Running

The initial login name for this CVM is root. You can check the initial login password in the Message Center, Reset the password if you forgot it.

Log In Shutdown Restart

Instance ID		Instance Configuration	Ex
Availability Zone	Chengdu Zone 1	Operating System	
IP		Creation Time	2019-12-18 09:44:15
Instance Billing Mode	CDH Billing		
Bandwidth billing mode	Bill by traffic Modify billing mode		

Basic Information   ENI   Public IP   Monitoring   Security Groups   Operation Logs

Instance Information

Name		Project	Default Project
Instance ID		Tags	None
UUID		Key	None
Instance Specification		Placement Group	None
Region	Chengdu	Role	None
Availability Zone	Chengdu Zone 1		

Network Information

Architecture

ins-4
South
8lsikfz

1 secu

1 ENI

Te

Querying Instance Metadata

Last updated：2024-11-08 16:02:50

Instance metadata refers to data relevant to an instance. It can be used for configuring or managing a running instance.
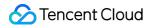
Note:

Although instance metadata can only be accessed after login, the data has not been encrypted. Anyone who accesses the instance can view its metadata. Therefore, you should take proper actions to protect sensitive data.

Overview

Tencent Cloud provides the following metadata:

Name	Description	Version
instance-id	Instance ID	1.0
instance-name	Instance name	1.0
uuid	Instance ID	1.0
local-ipv4	Instance private IP address	1.0
public-ipv4	Instance public IP address	1.0
mac	MAC address of the instance's eth0 device	1.0
placement/region	Instance region	Updated on September 19, 2017
placement/zone	Instance availability zone	Updated on September 19, 2017
network/interfaces/macs/${mac}/mac	MAC address of the instance's network interface	1.0
network/interfaces/macs/${mac}/primary-local-ipv4	Primary private IP of the instance's network interface	1.0
network/interfaces/macs/${mac}/public-ipv4s	Public IP address of the instance's network interface	1.0

network/interfaces/macs/${mac}/vpc-id	VPC ID of the instance's network interface	Updated on September 19, 2017
network/interfaces/macs/${mac}/subnet-id	Subnet ID of the instance's network interface	Updated on September 19, 2017
network/interfaces/macs/${mac}/local-ipv4s/${local-ipv4}/gateway	Gateway address of the instance's network interface	1.0
network/interfaces/macs/${mac}/local-ipv4s/${local-ipv4}/local-ipv4	Private IP address of the instance's network interface	1.0
network/interfaces/macs/${mac}/local-ipv4s/${local-ipv4}/public-ipv4	Public IP address of the instance's network interface	1.0
network/interfaces/macs/${mac}/local-ipv4s/${local-ipv4}/public-ipv4-mode	Public network mode of the instance's network interface	1.0
network/interfaces/macs/${mac}/local-ipv4s/${local-ipv4}/subnet-mask	Subnet mask of the instance's network interface	1.0
payment/charge-type	Instance billing plan	Updated on September 19, 2017
payment/create-time	Instance creation time	Updated on September 19, 2017
payment/termination-time	Instance termination time	Updated on September 19, 2017
app-id	AppID of the user to which the instance belongs	Updated on September 19, 2017
as-group-id	Auto scaling group ID of the instance	Updated on

		September 19, 2017
spot/termination-time	Spot instance termination time	Updated on September 19, 2017
instance/instance-type	Instance type	Updated on September 19, 2017
instance/image-id	Instance image ID	Updated on September 19, 2017
instance/security-group	Information of the security group bound to the instance	Updated on September 19, 2017
instance/bandwidth-limit-egress	Instance private network outbound bandwidth limit, in Kbit/s	Updated on 9/29/2019
instance/bandwidth-limit-ingress	Instance private network inbound bandwidth limit, in Kbit/s	Updated on 9/29/2019
cam/security-credentials/${role-name}	Temporary credential generated by the CAM role policy, which can be obtained only when the instance is associated with the CAM role. Change `${role-name}` to the actual CAM role name; otherwise, `404` will be returned	Updated on 12/11/2019
volumes	Instance storage	1.0

Note:

Field `${mac}` and `${local-ipv4}` in the above table indicate the MAC address and private IP address of the network interface specified for the instance, respectively.

The destination URL address of the request is case-sensitive. You must construct the destination URL address of a new request according to the returned result of the request.

In the current version, the returned data of placement has been changed. To use the data in the previous version, specify the previous version path or leave the version path empty to access the data of version 1.0. For more information on the returned data of placement, see Region and Availability Zone.

Querying Instance Metadata

After logging in to an instance, you can access the metadata such as its local IP address and public IP address to manage connections with external applications.

To view all the instance metadata within a running instance, use the following URI:

```
http://metadata.tencentyun.com/latest/meta-data/
```

You can access the metadata by using the cURL tool or an HTTP GET request, for example:

```
curl http://metadata.tencentyun.com/latest/meta-data/
```

For resources that do not exist, the HTTP error code "404 - Not Found" will be returned.

All metadata-related operations can only be taken place **after you logging in to the instance**. For more information, see Logging In To Windows Instance and Logging In To Linux Instance.

Sample metadata query

The following example shows how to obtain the metadata version.

Note:

When Tencent Cloud modifies the metadata access path or returned data, a new metadata version is released. If your application or script depends on the structure or returned data of the previous version, you can access metadata using the specified previous version. If no version is specified, version 1.0 is accessed by default.

```
[qcloud-user]# curl http://metadata.tencentyun.com/
1.0
9/19/2017
latest
meta-data
```

The following example shows how to view the metadata root directory. The lines ending with / represent directories and other lines represent the accessed data. For the description of accessed data, see the **Overview** section described above.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/
instance-id
instance-name
local-ipv4
mac
```

```
network/
placement/
public-ipv4
uuid
```

The following example shows how to obtain the physical location information of an instance. For the relationship between the returned data and the physical location, see Regions and Availability Zones.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/placement/regio
ap-guangzhou

[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/placement/zone
ap-guangzhou-3
```

The following example shows how to obtain the private IP address of an instance. If an instance has multiple ENIs, the network address of the eth0 device is returned.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/local-ipv4
10.104.13.59
```

The following example shows how to obtain the public IP address of an instance.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/public-ipv4
139.199.11.29
```

The following example shows how to obtain an instance ID. The instance ID is used to uniquely identify an instance.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/instance-id
ins-3g445roi
```

The following example shows how to query the instance UUID. The instance UUID can also be used as the unique identifier of an instance, but we recommend that you use instance ID to identify instances.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/uuid
cfac763a-7094-446b-a8a9-b995e638471a
```

The following example shows how to obtain the MAC address of an instance's eth0 device.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/mac
52:54:00:BF:B3:51
```

The following example shows how to obtain the ENI information of an instance. In case of multiple ENIs, multiple lines of data are returned, with each line indicating the data directory of an ENI.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
52:54:00:BF:B3:51/
```

The following example shows how to obtain the information of a specified ENI.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
local-ipv4s/
mac
vpc-id
subnet-id
owner-id
primary-local-ipv4
public-ipv4s
local-ipv4s/
```

The following example shows how to obtain the VPC information of a specified ENI.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
vpc-ja82n9op

[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
subnet-ja82n9op
```

The following example shows how to obtain the list of private IP addresses bound to the specified ENI. If the ENI is bound with multiple private IP addresses, multiple lines of data are returned.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
10.104.13.59/
```

The following example shows how to obtain the information of a private IP address.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
gateway
local-ipv4
public-ipv4
public-ipv4-mode
subnet-mask
```

The following example shows how to obtain the gateway of a private IP address. This data can be queried only for VPC-based CVMs. For more information about VPC-based CVMs, please see Virtual Private Cloud (VPC).

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
10.15.1.1
```

The following example shows how to obtain the access mode used by a private IP address to access the public network. This data can be queried only for VPC-based CVMs. A classic network-based CVM accesses the public network through the public gateway.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
NAT
```

The following example shows how to obtain the public IP address bound to a private IP address.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
139.199.11.29
```

The following example shows how to obtain the subnet mask of a private IP address.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/network/interfa
255.255.192.0
```

The following example shows how to obtain the billing type of an instance.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/payment/charge-
POSTPAID_BY_HOUR
```

The following example shows how to obtain the creation time of an instance.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/payment/create-
2018-09-18 11:27:33
```

The following example shows how to obtain the termination time for spot instances.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/spot/terminatio
2018-08-18 12:05:33
```

The following example shows how to obtain the account AppId to which the CVM belongs.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/app-id
123456789
```

The following example shows how to obtain the temporary credential generated by the CAM role to which the instance belongs. In this example, the role name is `CVMas` .

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/cam/security-cr
{
  TmpSecretId": "**********************************",
  "TmpSecretKey": "******************************",
  "ExpiredTime": 1615590047,
  "Expiration": "2021-03-12T23:00:47Z",
  "Token": "******",
  "Code": "Success"
}
```

The following example shows how to query the instance storage.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/meta-data/volumes
disk-xxxxxxxx/
```

Querying Instance User Data

You can specify instance user data when creating an instance. CVM instances having cloud-init configured can access the data.

Searching user data

After login, you can access user data by using the following method.

```
[qcloud-user]# curl http://metadata.tencentyun.com/latest/user-data
179, client, shanghai
```

Renaming Instances

Last updated：2024-01-08 09:32:02

Overview

To help users manage CVM instances on the console and locate CVMs quickly by name, Tencent Cloud allows users to rename an instance at any time and the new name takes effect instantly.

Directions

On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Modifying the name of an instance

1. Log in to the CVM console.

2. In the row of the target instance in the instance list, select **More** > **Instance Settings** > **Rename** on the right as shown below:

3. In the **Rename** window that pops up, enter the new instance name and click **OK**.

Modifying the names of multiple instances

1. Log in to the CVM console.

2. In the instance list, select the target instances and click **More Actions** > **Instance Settings** > **Rename** above the list as shown below.

3. In the **Rename** window that pops up, enter the new instance name and click **OK**.

Note:

CVMs modified using this method will have the same instance name.

1. Log in to the CVM console.

2. Select the tab of the target instance and select **More Actions** > **Instance Settings** > **Rename** in the top-right corner as shown below:

3. In the **Rename** window that pops up, enter the new instance name and click **OK**.

Resetting Instance Password

Last updated：2024-01-08 09:32:02

Overview

If you forget your CVM instance login password, you can reset it on the console. This document describes how to reset your instance login password on the console.

Note:

When the instance is shut down, you can directly reset the login password.

If the instance is still running, resetting the login password will force shut down it. To avoid affecting your business, please plan ahead and reset passwords during off-peak hours.

Directions

Resetting the password of a single instance

Resetting the passwords of multiple instances

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Password/Key** > **Reset Password** on the right as shown below:

Tab view: on the page of the target instance, click **Reset Password** as shown below:

3. Select **Username** and enter the username of the selected instance. Enter the **New password**, re-enter the new password in the **Confirm Password** field, and click **Next**.

Note:

The **Username** defaults to **System default**, and the default system username is used, such as `Administrator` for Windows, `ubuntu` for Ubuntu, and `root` for other Linux distributions. You can select **Specified user name** and enter the username.

4. Reset the password according to the instance status:

To reset the password of **Running** instances, select **Agree to a forced shutdown** and click **Reset Password**, as shown in the following figure:

To reset the password of **Shutdown** instances, click **Reset Password**, as shown in the following figure.

1. Log in to the CVM console.

2. On the **Instances** page, select the CVM instances to reset password, and click **Reset Password** at the top of the instance list, as shown in the following figure:

3. Select **Username** and enter the username of the selected instance. Enter the **New password**, re-enter the new password in the **Confirm Password** field, and click **Next**.

Note:

The **Username** defaults to **System default**, and the default system username is used, such as `Administrator` for Windows, `ubuntu` for Ubuntu, and `root` for other Linux distributions. You can select **Specified user name** and enter the username.

4. Reset the password according to the instance status:

To reset the password of **Running** instances, select **Agree to a forced shutdown** and click **Reset Password**, as shown in the following figure:

To reset the password of **Shutdown** instances, click **Reset Password**, as shown in the following figure.

FAQs

If you fail to reset the password for Windows CVM instances, refer to Failed to Reset the CVM Password or the CVM Password Is Invalid for troubleshooting.

Managing Instance IPs
Getting Private IP Addresses and Setting DNS

Last updated：2024-01-08 09:32:02

Overview

This document describes how to obtain the private IP address of the instance and configure the private DNS.

Directions

Obtaining the private IP address of an instance

Obtain in console

Obtain via API

Obtain via instance metadata

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

2. On the instance management page, proceed according to the actually used view mode:

List view: select the target instance, move the cursor to the primary IP column, and click

to copy the private IP as shown below:

Tab view: on the instance page, click

after the private network address in "IP Address" to copy the private IP as shown below:

See DescribeInstances.

1. Log in to your CVM.

2. Access the instance metadata by using the cURL tool or an HTTP GET request.

Note:

The following operations use the cURL tool as an example.

Execute the following command to obtain the private IP.

```
curl http://metadata.tencentyun.com/meta-data/local-ipv4
```

The returned information is the private IP address, as shown below:

For more information about instance metadata, see Querying Instance Metadata.

Configuring private network DNS

When a network resolution error occurs, you can manually configure the private network DNS based on your CVM operating system.

Linux

Windows

1. Log into the Linux CVM.

2. Execute the following command to open the `/etc/grub.conf` file.

```
vi /etc/resolv.conf
```
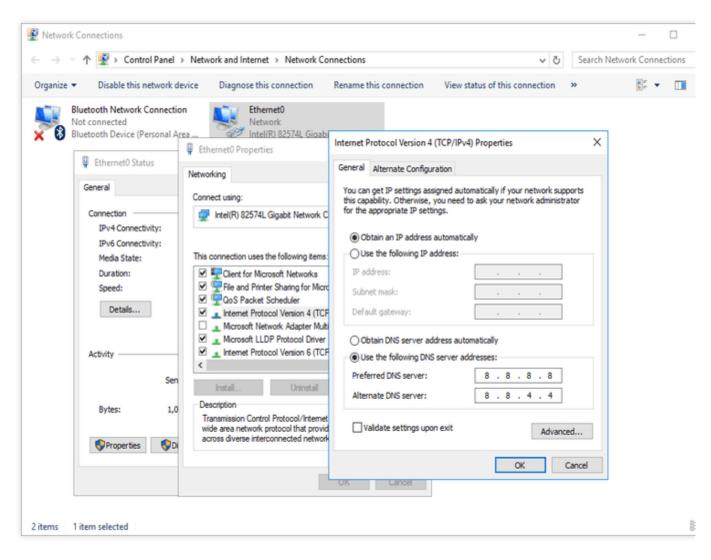
3. Press **i** to switch to the edit mode, and modify the DNS IP according to the corresponding region in the Private Network DNS list.

For example, change the private network DNS IP to an private network DNS server in the Beijing region.

```
nameserver 10.53.216.182
nameserver 10.53.216.198
options timeout:1 rotate
```

4. Press **Esc**, enter **:wq**, save the file and return.

1. Log in to the Windows CVM.

2. On the operating system UI, open **Control Panel** > **Network and Sharing Center** > **Change adapter settings**.

3. Right-click **Ethernet** and select **Properties** to open the "Ethernet Properties" window.

4. In the "Ethernet Properties" window, double-click **Internet Protocol Version 4 (TCP/IPv4)** as shown below:

5. Select **Use the following DNS server addresses** and modify the DNS IP according to the corresponding region in the Private Network DNS list.

6. Click **OK**.

Modifying Private IP Addresses

Last updated：2024-07-10 09:55:57

Scenario

You can directly modify the internal IP of Cloud Virtual Machine (CVM) instances within a private network via the console. This document will guide you on how to modify the internal IP of CVM instances within a private network through the CVM console.

Limits

Modifying the primary IP of a primary ENI may cause the CVM to restart.
The primary IP of a secondary ENI cannot be modified.

Directions

1. Log in to the CVM Console.
2. Select the region of the instance whose private IP you want to modify, and click the instance ID/name to enter its details page.
3. On the instance details page, select the [ENI] tab and click

 to expand the primary ENI.
4. In the primary ENI operation list, click **Modify Primary IP**.
5. In the "Modify Primary IP" window that pops up, enter the new IP and then click **OK**. It takes effect after the instance is restarted.
Note:
You can only enter private IP in the current subnet CIDR.

Getting Public IP Addresses

Last updated：2024-01-08 09:32:02

Overview

This document describes how to obtain the public IP address of a CVM instance.

Directions

Console

API

Instance metadata

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: In the primary IP column, click

to copy the IP.

Tab view: On the instance page, click

after the public network address in "IP Address" to copy the public IP.

Note:

The public IP address is mapped to the private IP address through NAT. Therefore if you view the network interface attributes from within the instance (such as by using `ifconfig (Linux)` or `ipconfig (Windows)` commands), the public IP address is not displayed. To obtain the public IP from within the instance, you need to check the instance metadata.

See DescribeInstances.

1. Log in to the CVM instance.

For more information, see Logging in to Linux Instance Using Standard Login Method and Logging in to Windows Instance.

2. Use the cURL tool or an HTTP GET request to access the metadata and obtain the public IP address.

```
curl http://metadata.tencentyun.com/meta-data/public-ipv4
```

Check the public IP in the result:

For more information, see Instance Metadata.

Changing Public IP Addresses

Last updated：2024-01-08 09:32:02

Overview

This document describes how to change the public IP address of a CVM instance.

Considerations

Each account can change the public IPs of CVMs in the same region a maximum of 3 times per day.

The public IP of each instance can be changed **once only**.

The old public IP will be released after the change.

Prerequisites

You have logged in to CVM console.

Directions

On the instance management page, proceed according to the actually used view mode:

List view

Tab view

1. Locate the target instance and select **More** > **IP/ENI** > **Change Public IP**.

2. In the **Change IP** dialog box, click **Confirm** to complete the change.

1. Go to the page of the target instance and select **More Actions** > **IP/ENI** > **Change Public IP** in the top-right corner.

2. In the **Change IP** dialog box, click **Confirm** to complete the change.

Retrieving Public IP Address

Last updated：2024-01-08 09:32:02

Scenario

This document describes how to retrieve a public IP address that has been used before but not yet assigned to other users.

Notes

The retrieved IP address is an EIP, and the total number of EIPs must not exceed the total quota.

Each account can apply for a specific IP address up to three times per month in each region.

Directions

1. Log in to CVM Console.
2. In the left sidebar, click **EIP** to access the EIP management page.
3. Click **Retrieve IP**, as shown in the following figure:

4. In the **Retrieve IP** pop-up window, enter the public IP address and click **Check** to query whether the IP address can be retrieved, as shown in the following figure.

If yes, click **Apply Now**.

If no, the IP address that you applied for cannot be retrieved for reasons such as it has already been assigned. In this case, try to apply for another IP address or click **Cancel** to exit.

Changing Security Group

Last updated：2024-01-08 09:32:02

Overview

Security group is a virtual firewall for filtering packets and is used to set the network access controls for one or multiple CVMs. It is an important network security isolation method provided by Tencent Cloud. When creating a CVM instance, you must configure a security group for it. Tencent Cloud allows you to configure a new security group for the CVM instance after it is created.

Note:

To configure a new security group for the instance, create a security group first. For more information, please see Creating a Security Group.

Prerequisites

You have logged in to CVM Console.

Directions

Change the configured security group

On the instance management page, proceed according to the actually used view mode:

List mode

Tab mode

1. On the instance management page, select a CVM instance for which a new security group needs to be configured.

Click **More** > **Security Group** > **Configure Security Group**, as shown below:

2. In the pop-up window, check the name of the new security group (multiple names can be selected) and click **Confirm** to change the security group.

1. On the instance management page, select a tab of the CVM instance for which a new security group needs to be configured.

2. On the instance details page, click **More** > **Security Group** > **Configure Security Group**, as shown below:

3. In the pop-up window, check the name of the new security group (multiple names can be selected) and click **Confirm**.

Change the bound security group

1. On the instance management page, click the CVM instance ID/name for which you want to bind the security group and enter the instance details page.

2. On the instance details page, select the **Security Groups** tab and click **Bind** on the "Bound to security group" column, as shown below:

3. In the pop-up window, check the name of the security group (multiple names can be selected) to be bound based on your actual needs and click **OK** to bind the security group, as shown below:

Security Groups   ✕

Projects   [All projects ▾]   Selected (4)

Select a security group   ID/Name   Notes

[Enter the security group name or ID   🔍]

☐   ID/Name   Notes   ⬍   ✖

☑   ⬍   ✖

☑   ↔   ⬍   ✖

☑   ⬍   ✖

☐

☐

☐

[**OK**]   [Cancel]

Conversion from Pay-As-You-Go to Monthly Subscription

Last updated：2024-03-08 17:11:25

Operation scenarios

To make it more convenient for you to use CVM, Tencent Cloud has launched a feature that allows you to convert pay-as-you-go instances to monthly subscription instances. This enables you to convert temporary pay-as-you-go instances into long-term and stable monthly subscription instances. You can execute this conversion operation through the CVM console and Application Programming Interface. This document will guide you on how to convert pay-as-you-go instances to monthly subscription instances by using the CVM console.

Conversion Rules

We provide a billing mode conversion feature on the CVM console, and the specific rules are as follows:
Supports converting single or multiple pay-as-you-go instances into monthly subscription instances.
When a pay-as-you-go instance is converted to a monthly subscription, a renewal order is generated. The payment process for this order must be completed before the change in billing method can take effect.
If the payment is not made or failed, this order can be viewed and handled on your Order Center page.
CVMs that convert the billing mode from pay-as-you-go to monthly subscription do not support the seven-day unconditional refund policy.
After the billing method conversion is successful and the payment is made, the instance will be billed as a monthly subscription immediately. The start time of the new monthly subscription instance is the successful conversion time. Before successful payment, you cannot repeat the conversion of the billing mode for this instance.
Before successful payment, if the configuration information of a instance changes (such as adjusting configuration, reinstalling the system, adjusting bandwidth, adjusting disk, etc.), the amount of the new purchase order doesn't match with the instance, and the unpaid order will be prohibited from being paid. You need to cancel the current unpaid order in Order Center first, then carry out the new conversion operation.
The feature of converting pay-as-you-go to monthly package supports the synchronous conversion of the billing mode of the instance and disk. After the billing mode of the instance is converted, except for the network bandwidth billing mode of the hourly bandwidth of the common public IP for standard account types and the network bandwidth billing mode of hourly bandwidth for traditional account types, which support automatic conversion to monthly package billing by bandwidth, the remaining network bandwidth billing modes remain unchanged.

Use Limits

The conversion is not supported when the remaining quota of the monthly package in the availability zone is less than the number of the instances to be converted from pay-as-you-go.

Instances not billed by the pay-as-you-go billing are not supported for conversion.

Spot instances are not supported for conversion.

Instance network billing mode is billed by bandwidth usage duration. Temporary conversion not supported.

Instances using cloud market images do not support conversion.

Batch instances BC1, BS1 do not support conversion.

Pay-as-you-go instances with unfinished conversion orders are not supported for conversion.

Pay-as-you-go instances that have been set to destruct at a specific time do not support conversion. If you need to convert, please cancel the timed destruction and convert again.

Operation Step

1. Sign in CVM console.

2. Depending on the actual needs, on the instance management page, choose different convertion instance operations.

Convert a Single Instance

Convert Multiple Instances

On the instance management page, proceed according to the actually used view mode:

List view: In the operation column on the right, select **More** on the top > **Instance settings** > **Switch from Pay-as-you-go to Monthly Subscription**, as shown below:

Note:

You can also check the instances that need to be converted, click **More actions** on the top > **Instance setting** > **Switch from Pay-as-you-go to Monthly Subscription**.

Tab view: On the instance page, select **More actions** on the top > **Instance settings** > **Switch from Pay-as-you-go to Monthly Subscription**, as shown below:

Check all the instances that need to be converted, and click **More actions** on the top > **Instance settings** > **Switch from Pay-as-you-go to Monthly Subscription**. This can change the billing mode of multiple instances in batches, as shown below:

The reasons will be displayed for instances that cannot be operated.

3. In the pop-up **Switch to Monthly Subscription** window, according to actual needs, set the renewal period and the auto-renewal, as shown below:

Renewal period: choose the purchase period after the conversion to monthly subscription. If multiple instances are batch converted, the same purchase duration must be set.

Auto-renewal: Choose auto-renewal according to your needs.

4. Check **I have read and agreed to Rules on Switching from Pay-as-you-go to Monthly Subscription**, then click **Change now**.

If there are no unfinished conversion orders for the instance, you will automatically be redirected to the payment page.

5. Follow the prompts on the page to complete the payment and finish the conversion operation.

FAQs

If you encounter any problems during the conversion process, please refer to the purchasing and renewing document.

Searching for Instances

Last updated：2024-01-08 09:32:02

Scenario

By default, the CVM console displays the instances for all projects in the current region. To help you quickly search instances in the current region, Tencent Cloud provides a CVM search feature. You can filter out instances by resource attributes such as project, instance billing method, instance type, availability zone, IP, instance ID, and instance name.

Directions

1. Log in to the CVM Console.
2. Enter the content you wish to search based on your needs, and click

to search.

Enter the keyword in the search text box, and click

, as shown below:

Choose a specific dimension to search (such as project, project, instance billing method, instance type, etc..) and click

, as shown below:

3. To learn more about search syntax, click

to view the relevant syntax of search instances.

For more search instance syntax, please see the following figure.

	Enter Format	Example	Display in Search Box	Description
Single key-word	[Keyword]	10.0.0.1	10.0.0.1　Use '|' to split more than one keywor	List all instances including the keyword "10.0.0
Multiple key-words	[Keyword] [Enter key ↵] [Keyword]	10.0.0.1 www.123.com 192.169.23.54	10.0.0.1　www.123.com　192.169.23.45	List all instances that include all the three key-words"10.0.0.1,""www.123.com"and"192.169.2:
Single re-source type	[Resource type]: [Keyword]	IP: 10.0.0.1	IP: 10.0.0.1　Use '|' to split more than one key	List all instances whose IP is "10.0.0.1"
Multiple re-source types	[Resource type]: [Keyword][Enter key ↵][Resource type]: [Keyword]	Availability Zone: Hong Kong Zone 2 Project: Default	Availability Zone: Hongkon...　Project: Defau	List all instances whose"Availability Zone"is"H Kong Zone 2" and "Project " is "Default"
Single re-source type and multiple keywords	[Resource type]: [Keyword] | [Key-word]	CVM Status: Creating | Shut-down	CVM Status: Creating | Shu...　Use '|' to split	List all instances whose "CVM Status " is "Creati or "Shutdown"
Pasted con-tents	{pasted contents}	112.11.22.33 112.11.22.34 112.11.22.53	112.11.22.33 | 112.11.22.3...　Use '|' to split	List all instances include the keywords "112.11.22.33", "112.11.22.34" or "112.11.22.53

Exporting Instance List

Last updated：2024-01-08 09:32:02

Overview

You can export the CVM instance list of a region in the console, and customize the fields to be exported. You can select a maximum of 27 fields, including ID, instance name, status, region, availability zone, instance type, operating system, image ID, CPU, MEM, bandwidth, public IP, private IP, system disk type, system disk size, data disk type, data disk size, network type, subnet ID, VPC name, creation time, expiry time, instance billing mode, network billing mode, project, dedicated host ID, and tag.

Directions

1. Log in to the CVM console.
2. On the instance management page, select a region and proceed according to the actually used view mode:

List view

Tab view

Click

in the top-right corner of the instance list, as shown below:

Click

in the top-right corner of the instance page, as shown below:

3. In the pop-up "Export instances" window, select the fields you want to export and click "OK", as shown below:

Export instances   ×

☑ Select All

☑ ID   ☑ Bandwidth (Mbps)
☑ Instance Name   ☑ Primary public IPv4
☑ Status   ☑ Primary private IPv4
☑ Region   ☑ Primary IPv6
☑ Availability Zone   ☑ System Disk Type
☑ Instance Type   ☑ System disk size (GB)
☑ CPU (core)   ☑ Data Disk Type
☑ MEM (GB)   ☑ Data disk size (GB)
☑ Operating System   ☑ Network type
☑ Image ID   ☑ VpcId
☑ VPC name
☑ Subnet ID
☑ Subnet name
☑ Creation Time
☑ Expiry Time
☑ Instance Billing Mode
☑ Network billing mode
☑ Project
☑ Dedicated Host ID
☑ Tag

Export range   ⦿ All Instance
                ○ Only export search result
                ○ Selected Instance

            OK   Close

Renewing Instances

Last updated：2024-01-08 09:32:02

This document introduces how to renew **Postpaid instance**.

Postpaid instance: Postpaid instances can be automatically activated with sufficient balance in your account. For more information, please see Online Top-up.

Starting Up Instances

Last updated：2024-01-08 09:32:02

Overview

This document describes how to start up an instance via the console or an API.

Directions

Starting up an instance via the console

Starting up instances via API

Starting up one instance

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Instance Status** > **Start Up** in the **Operation** column on the right as shown below:

Tab view: on the page of the target instance, select **Start Up** in the top-right corner as shown below:

Starting up multiple instances

Select the instances you want to start up, and click **Start up** at the top of the list to start the selected instances, as shown below:

Use the StartInstances API to start up an instance.

Subsequent Operations

Once the instance starts up, you can perform the following operations:

Logging in to the instance: depending on the instance type, log in to the Linux instance or the Windows instance.

Initializing cloud disks: initialize the cloud disks mounted to the instance by formatting, partitioning, and creating a file system.

Shutting Down Instances

Last updated：2024-01-08 09:32:02

Overview

The instance can be shut down when you need to stop the service, or modify configurations that can be done only in the shutdown state. Shutting down an instance is like shutting down a local computer.

Notes

You can shut down an instance using system commands (such as the shutdown command under Windows system and Linux system) or through the Tencent Cloud console. We recommend you view the shutdown process on the console to check whether any problem occurs.

The instance will no longer provide services after the shutdown. Before the shutdown, make sure the CVM has stopped receiving service requests.

During the shutdown, the status of the instance will change from "shutting down" to "shutdown". If the shutdown process takes too long, there may be an exception. For more information, please see Close an CVM to avoid forced shutdown.

After an instance is shut down, all storage is still connected to the instance, and all disk data are retained. Data in the memory will be lost.

Shutting down an instance does not change its physical attributes. The public and private IPs of the instance remain unchanged. Elastic Public IP is still bound to the instance. Due to service interruption, however, you will receive an error response when accessing these IPs. Classiclink relationship remains unchanged.

If the instance belongs to the real server cluster of the CLB instance, it can no longer provide services after the shutdown.

If the health check policy has been configured, the instance that has been shut down will be automatically blocked and requests will no longer be forwarded to it. Otherwise, the client may receive a 502 error code. For more information, please see Health Check.

If the instance that has been shut down is in an auto scaling group, the auto scaling service will mark the instance as having poor performance, and may replace and move it out of the auto scaling group. For more information, please see Auto Scaling.

Directions

Shutting down instance via console

Shut down instance via API

Shutting down a single instance

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: select the target instance and click **More** > **Instance Status** > **Shutdown** in the **Operation** column on the right as shown below:

Tab view: on the page of the target instance, select **More Actions** > **Instance Status** > **Shutdown** in the top-right corner as shown below:

Shutting down multiple instances

1. Log in to the CVM console.

2. Select all the instances you want to shut down and click **Shutdown** at the top of the list to shut down instances in batches as shown below:

Note:

Reasons are given for instances that cannot be shut down.

For more information, see the StopInstances API.

Subsequent Operations

You can modify the following attributes only if the instance has been shut down.

Instance configuration (CPU, memory): To change the instance type, see Change Instance Configuration.

Change password: see Login Password.

Load SSH key: see SSH Key.

Restarting Instances

Last updated：2024-01-08 09:32:02

Overview

Restarting the CVM instance is a common method to maintain it. It is equivalent to restarting the operating system of the local computer. This document describes how to restart instances.

Notes

Preparing to restart instances: The instance cannot provide services during restart. Make sure before restarting the CVM that it has stopped receiving service requests.
How to restart instances: We recommended you restart an instance using the restart operations provided by Tencent Cloud instead of running the restart command in the instance (such as the relaunch command under Windows and the reboot command under Linux).
Restart time: Generally, it takes only a few minutes to restart an instance.
Physical features of instances: Restarting an instance does not change its physical features. Its public and private IP addresses as well as stored data will not be changed.
Billing: Restarting an instance will not start a new instance billing period.

Directions

You can restart instances via the following methods:
Restarting instance in console
Restarting instances via API

Restarting a single instance

1. Log in to the CVM console.
2. On the instance management page, proceed according to the actually used view mode:
List view: in the row of the target instance, select **More** > **Instance Status** > **Restart** as shown below:

Tab view: on the page of the target instance, select **Restart** in the top-right corner as shown below:

Restarting multiple instances

1. Log in to the CVM console.

2. Select all instances you want to restart and click **Restart** at the top of the list to batch restart the instances. If they cannot be restarted, the reason will be displayed as shown below:

Note:

A single instance can also be restarted in this method.

For more information, see the RebootInstances API.

Reinstalling System

Last updated：2024-01-08 09:32:02

Scenarios

System reinstallation allows you to restore an instance to its initial status at launch, which is an important recovery method if the instance has a system failure. This document describes how to reinstall the operating system.

CVM supports the following two reinstallation types:

Reinstall on the same platform: CVMs in all regions can be reinstalled to the OS of the same platform.
For example, you can always reinstall a Linux instance on a Linux OS, and Windows instance on a Windows OS.

Reinstall on different platform: only CVMs in the Chinese mainland can be reinstalled to an OS of different platform.
For example, a Linux instance can be reinstalled on a Windows OS, and a Windows instance on a Linux OS.

Note:

All newly added cloud disk and local disk instances support reinstallation on different platforms. Some existing 20 GB local disk instances do not support cross-platform reinstallation in the console. If you use such instances, you need to submit a ticket for application.

Spot instances do not support system reinstallation.

Notes

Preparation: A reinstallation clears all data in the system disk. Therefore, you must back up important data in the system disk in advance. If you want to retain your system operating data, we recommend you create a custom image and use this image to reinstall the operating system.

Image selection: we recommend that you use the image provided by Tencent Cloud or your custom image instead of those from unknown or other sources. Do not perform other operations while the system disk is being reinstalled.

Instance physical features: the public IP of the instance will not change.

Specification limits: if you want to use an image with Windows 2016, or 2019 versions, the instance memory should be greater than 2 GB.

Billing: if you adjust the size of the system disk (for cloud disks only), you will be charged according to the pricing standards of CBS. For more information, see Pricing List.

Subsequent operations: after the system disk is reinstalled, the data in the data disks will not be affected and will be available for use only after the data disks are reattached.

Directions

You can use either of the following methods to reinstall the operating system:

Reinstalling the system in the console

Reinstalling the system via an API

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Reinstall System** as shown below:

Tab view: on the page of the target instance, select **More Actions** > **Reinstall the System** in the top-right corner as shown below:

3. In the pop-up window that appears, read notes and click **Next**.

4. Select the image that is used by the current instance or another image, set the instance login method, and click **OK**, as shown below:

Note:

Only when the image type is **custom image** or **shared image** can you select **Follow image** as the login method.

Tencent Cloud

Reinstall system

✓ **Heads up** > ② **Configuration**

You've selected 1 instance. Collapse

ID/name	Instance type	Operating system	System disk
	S6 2-core 4GB	OpenCloudOS Server 8 ⚡	Balanced SSD 50GB

ⓘ • Create a snapshot or image to back up your data before continuing, so as to avoid data loss.**Operation guide** ↗

• Data in the instance data disk will not be cleared. But you need to mount the disk manually after reinstallation to use it.**Operation guide** ↗

• If the current system disk size is too small to meet the requirements, please expand the capacity.**Disk capacity expansion** ↗

Image type	**Current image**	Public image	Custom image	Shared image	Market image

Target image   OpenCloudOS Server 8

Login methods   **Set password**   Bind key   Follow image

Username   root

New password   [Please enter the instance   Ø]

Security reinforcement   ☑ Activate Anti-DDoS Protection and Cloud Workload Protection for free. About Security Reinforcement ↗

Cloud monitoring   ☑ FREE cloud monitoring, analysis, alarming, and server monitoring metrics (component installation required) About Cloud Monitor ↗

Fees

Back   **OK**

For details, see ResetInstance.

Related Operations

If the CVM has attached a data disk and you need to reinstall it on a different platform, please see the following documents about how to read data from the data disks of the original operating system:

Read/Write EXT Data Disks after Reinstalling a Linux CVM to Windows CVM

Read/Write NTFS Data Disks after Reinstalling a Windows CVM to Linux CVM

Using Tencent Cloud Automation Tools to execute commands

Last updated：2024-01-08 09:32:02

Overview

TencentCloud Automation Tools (TAT) is a native operations and deployment tool for CVM and Lighthouse instances. You don't need to connect to the instance remotely. TAT can automatically execute shell commands in batch to complete tasks such as running automation scripts, polling processes, installing/uninstalling software, updating applications and installing patches. For more information, see Overview.
This document describes how to use TAT to execute commands for instance management.

Prerequisites

The TAT agent is installed on the CVM instance. See Installing TAT Agent.
Note:
Some existing CVM instances do not currently support the use of TAT. It's expected that TAT will be supported on all instances later.

Directions

Refer to the following documents to create, execute, and view command execution status:
Create a command
Execute the command or execute the command without logging in to the instance, see Executing Without Logging In.
Query the command execution status as instructed in Querying Command Execution Status

Terminating/Returning Instances Overview

Last updated：2024-04-10 10:14:25

This document describes how to terminate and release a Cloud Virtual Machine (CVM) instance. For more information on expiration, see Payment Overdue.

Overview

You can terminate an instance if you no longer need it. The terminated instance will be put into the recycle bin. For instances in the recycle bin, you can restore or release them as needed based on different scenarios.
Note:
If your account is in arrears, then for pay-as-you-go instances, you need to renew the instances first before restoring them.

Methods for Termination/Release

For pay-as-you-go instances, the methods for instance termination and release are as follows:
Manual termination: You can manually terminate a pay-as-you-go instance that is not in arrears. A pay-as-you-go instance is released after it remains in the recycle bin for over 2 hours.
Timed termination: Timed termination is supported for pay-as-you-go instances. You can select a future time to terminate resources. The set termination time is precise to the second. Instance resources for which timed termination is set will be released immediately as scheduled, instead of going into the recycle bin. You can cancel timed termination at any time before the set termination time.
Expiry/arrears auto termination: Pay-as-you-go instance will be automatically terminated when its balance drops below 0 for 2 hours and 15 days. Billing will continue for the first 2 hours, then the instance will shut down and no longer be billed. The pay-as-you-go instance in arrears will not enter the recycle bin and can be viewed on the instance list. You can continue to use the instance if you renew it within the specified time. For more information, see Renewing Instances.

type	Manual termination (not in arrears)	Timed termination (not in arrears)	Automatic termination upon expiration or when in arrears
Pay-as-you-go instances	After termination, the instance is stored in the recycle bin for 2 hours, and	Instances for which timed termination is set will be released immediately as	After an instance enters into arrears, for the first 2 hours, billing will continue and the instance can still be used normally. In the next 15 days, however, the instance will be

	if it is not restored within these 2 hours, it will be released.	scheduled, instead of going into the recycle bin.	shut down, and billing will stop. Pay-as-you-go instances in arrears will not be put into the recycle bin. If the instance is not renewed within the aforementioned period, the instance will be released.
Monthly-Subscribed Instance	Early Destruction Not Supported	Scheduled Termination Not Involved	After expiration, terminated instances enter the recycle bin and are retained for a maximum of 7 days; if not restored upon expiration, the instance will be released.

Relevant Impact

When an instance is terminated, the relevant impact on instance data, EIPs, and billing is as follows:

Billing: when an instance is being terminated or has been released, no expenses related to this instance are incurred.

Instance data: local disks and non-elastic cloud disks attached to the instance are all released, and the data on these disks will be lost. Back up the data in advance. Elastic cloud disks follow their own lifecycle.

EIP: EIPs (including IP addresses on the secondary ENI) of a terminated instance are retained, and idle IP addresses may incur expenses. If you don't need them anymore, release them as soon as possible.

Directions

You can manually terminate/release instances through the following ways:

Terminate/release instances in the console. For more information, see Terminating/Returning Instance in Console.

Terminate/release instances by calling the API. For more information, see TerminateInstances.

Terminating/Returning Instance in Console

Last updated：2024-01-08 09:32:02

Overview

This document describes how to terminate/return a pay-as-you-go CVM instance in the console.

Note:

For the impact of terminating/returning a CVM instance, see Impacts.

Directions

Terminating and releasing pay-as-you-go instances

For pay-as-you-go instances, you can choose immediate termination or timed termination.

1. Log in to the CVM console.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Instance Status** > **Terminate**/**Return** on the right as shown below:

If you need to terminate multiple instances at the same time, select the instances and click **More Actions** > **Terminate**/**Return** at the top of the list.

Tab view: on the details page of the target instance, click **Terminate**/**Return** in the top-right corner of the page as shown below:

3. In the **Terminate**/**Return** pop-up window, choose **Immediate Termination** or **Timed Termination**.

Immediate Termination: if you choose immediate termination, you can choose whether to release resources now or 2 hours later. If you choose to release resources now, the instance data will be cleared and cannot be restored.

Timed Termination: if you choose timed termination, you need to specify the termination time. The instance will be terminated and released upon expiration, and the data cannot be restored.

4. After choose a termination option, click **Next** to confirm the actual resources to be terminated or retained.

5. After confirming the resources to be terminated, click **Start Termination**.

Related Operations

Canceling timed termination

1. Log in to the CVM console.

2. In the instance list, find the instance for which you want to cancel timed termination. In the "**Instance Billing Mode**" column, find "**Timed Termination**" and move the mouse cursor to

ⓘ

to display the timed termination dialog box, as shown below:

3. Click **Cancel**. A dialog box is displayed prompting you to confirm the cancellation.

4. In the dialog box, confirm the information of the instance for which you want to cancel timed termination and click **OK**. The cancellation takes effect immediately, as shown below:

Enabling Instance Termination Protection

Last updated：2024-01-08 09:32:02

Overview

To prevent instances from being terminated unexpectedly, you can enable Termination Protection.
When Termination Protection is enabled, the instance cannot be terminated in the console or by using APIs. You can disable this setting any time as necessary.

Notes

Instance termination protection is disabled by default.
Instance termination protection does not take effect at the system level; for example, when a pay-as-you-go instance is to be terminated due to overdue payment, the protection does not apply.

Directions

Enabling termination protection

Existing instances

Newly purchased instances

1. Log in to the CVM console.

2. You can enable instance termination protection for one or multiple instances as needed:

One instance:

On the **Instances** page, find the target instance and click **More** > **Instance Settings** > **Instance Termination Protection**.

Multiple instances:

On the **Instances** page, select the target instances and select **More** > **Instance Settings** > **Instance Termination Protection**.

3. In the **Instance Termination Protection** pop-up window, select **Enable** and click *OK*.

When purchasing an instance, select **Custom Configuration** and select **Instance termination protection** on the **2. Complete Configuration** tab.

Note:

For more information on other parameters, see Creating Instances via CVM Purchase Page.

Disabling instance termination protection

If you are sure that an instance can be terminated, follow the steps below to disable instance termination protection.

1. Log in to the CVM console.

2. You can disable instance termination protection for one or multiple instances as needed:

One instance:

On the **Instances** page, find the target instance and click **More** > **Instance Settings** > **Instance Termination Protection**.

Multiple instances:

On the **Instances** page, select the target instances and select **More** > **Instance Settings** > **Instance Termination Protection**.

3. In the **Instance Termination Protection** pop-up window, select **Disable** and click *OK*.

References

[Creating Instances via CVM Purchase Page](#)

[Terminating Instances](#)

Instance Repossession or Recovering

Last updated：2024-01-08 09:32:02

This document describes how to repossess a Cloud Virtual Machine (CVM) instance from the recycle bin. For more information, please see Arrears Reminder.

Instance Repossession Description

Tencent Cloud recycle bin is a cloud service repossession mechanism as detailed below:

Pay-as-you-go instances that are manually terminated or terminated at a scheduled time will be put into the recycle bin. If the account is in arrears, the repossession mechanism does not apply to pay-as-you-go instances, and these instances are directly released when the account is in arrears for 2 hours + 15 days.

Instance status in the recycle bin are as follows:

Pay-as-You-Go instances in the recycle bin

Retention period: if your account has no overdue payments, terminated instances will be retained in the recycle bin for 2 hours.

Expiry processing: if instances are not renewed before the retention period ends, the system will release instance resources and automatically terminate instances, which cannot be recovered. EIPs bound to these instances will be retained. If you don't need such EIPs, release them promptly.

Mounting relationship: after the instance enters the recycle bin, its mounting relationship with Cloud Load Balancer, Cloud Block Storage, and Classiclink will **not be automatically terminated**.

Operation restrictions: for instances in the recycle bin, you can only perform the following operations: **renew and recover**, **terminate/return** and **create image** (except for special models).

Note:

You cannot repossess pay-as-you-go instances in the recycle bin if your account is in arrears. Please renew the payment first.

Pay-as-you-go instances are stored in the recycle bin for a maximum of 2 hours. Please note the release time and renew the payment in time to repossess the instances.

Pay-as-you-go instances cannot enter the recycle bin if your account is in arrears. You can view them on the CVM instance list page. The instances will be released after your account has been in arrears for 2 hours + 15 days.

Recovering Instances

1. Log in to the CVM console and select **Recycle Bin** > **Instance Recycle Bin** on the left sidebar.

2. On the **Instance Recycle Bin** page, perform different operations as needed.

Recovering one instance

Batch recovering instances

Find the instance to be recovered in the list, click **Recover** in the **Operation** column, and complete the renewal

payment.

Select all instances to be recovered in the list, click **Batch Recover** at the top, and complete the renewal payment.

Spot Instances

Last updated：2024-01-08 09:32:02

Overview

This document provides guidance on managing and purchasing spot instances. Currently, spot instances are available through the following channels:

CVM console: **Spot Instances** has been added as an option to **Billing Mode** on the CVM purchase page.

BatchCompute console: Spot instances can be selected when users submit jobs and create computing environments in the BatchCompute console.

TencentCloud API: spot instance parameters have been added to the RunInstance API.

Directions

CVM console

BatchCompute console

TencentCloud API

1. Log in to the CVM instance purchase page.

2. On the **Select Model** tab, set **Billing Mode** to **Spot Instances** as shown below:

3. Select region, availability zone, network type, instance and other configuration information as needed and prompted by the page.

4. Check the information of the spot instance to be purchased and the cost details of each configuration item.

5. Click **Activate** and make the payment.

After completing payment, you can log in to the CVM console to check your spot instance.

BatchCompute feature description

Async API

When you submit a job, create a computing environment, or modify the expected number of instances in a computing environment, your BatchCompute instance will process your request asynchronously. When it cannot fulfill the current request due to inventory or price reasons, the BatchCompute instance will continuously apply for spot instance resources until the current request is fulfilled.

If you need to release an instance, you need to adjust the expected number of instances in the computing environment via the BatchCompute console. If you release instances via the CVM console, the BatchCompute console will automatically create instances until the expected number of instances is met.

Cluster mode

The computing environment of a BatchCompute instance can maintain a batch of spot instances as a cluster. You only need to submit the desired quantity, configuration, and maximum price of the spot instances, and the computing environment will automatically and continuously apply for spot instances until the expected quantity is reached. Even if spot instances go offline, the computing environment will automatically apply for spot instances again to reach the expected quantity.

Fixed price

The fixed discount mode is used currently, so you must set the parameter to a value greater than or equal to the current market price. For more information on the market prices, see Spot Instance.

Directions

1. Log in to the BatchCompute console.

2. On the computing environment management page, randomly select a region, such as Guangzhou, and then click **Create**.

The **Create Computing Environment** page appears.

3. On the **Create Computing Environment** page, set **Billing Type** to **Spot Instance** and then configure information such as **Model**, **Image**, **Name**, and **Expected Quantity** as needed as shown below:

4. Click **OK**.

Then you can view the new computing environment in the BatchCompute console. To view the creation progress of CVM instances that are being created in the computing environment, click **Activity Log** and **Instance List** for the computing environment.

In the `RunInstance` API, you can specify the InstanceMarketOptionsRequest parameter to enable or disable the spot instance mode and configure the information about spot instances.

Sync API: currently, `RunInstance` provides a one-time sync request API. This means that if the application fails because the inventory is insufficient or the requested price is lower than the market price, the `RunInstance` API will immediately return a failure code and no longer apply for the spot instance again.

Fixed price: the fixed discount mode is used currently, so you must set the parameter to a value greater than or equal to the current market price. For more information on the market prices, see Spot Instance.

Sample scenario description

You have an instance in Guangzhou Zone 3, and the billing mode of the instance is pay-as-you-go on an hourly basis and in spot mode. The specific configurations of the billing mode are as follows:

MaxPrice: 0.0923 USD/hour

SpotInstanceType: one-time

ImageId: img-pmqg1cw7

InstanceType: S2.MEDIUM4 (Standard 2, 2-core, 4GB)

InstanceCount: 1

Request parameters

```
https://cvm.tencentcloudapi.com/?Action=RunInstances
&Placement.Zone=ap-guangzhou-3
&InstanceChargeType=SPOTPAID
&InstanceMarketOptions.MarketType=spot
&InstanceMarketOptions.SpotOptions.MaxPrice=0.0923
&InstanceMarketOptions.SpotOptions.SpotInstanceType=one-time
&ImageId=img-pmqg1cw7
&InstanceType=S2.MEDIUM4
&InstanceCount=1
&<Common request parameters>
```

Response parameters

```
{
  "Response": {
    "InstanceIdSet": [
      "ins-1vogaxgk"
    ],
    "RequestId": "3c140219-cfe9-470e-b241-907877d6fb03"
  }
}
```

Querying the Repossession Status of a Spot Instance

Last updated：2024-01-08 09:32:02

Spot instances may be repossessed by Tencent Cloud due to price or inventory reasons. To enable users to perform custom operations before instance repossession, we provide an API for obtaining information about repossession status via an internal metadata mechanism.

Metadata

Instance metadata refers to data relevant to an instance. It can be used for configuring or managing a running instance. You can log in to the instance to access and obtain instance metadata. For more information, see Querying Instance Metadata.

Querying the Termination Information of a Spot Instance using Metadata

Run the following command using the cURL tool. You can also send an HTTP GET request.

```
curl metadata.tencentyun.com/latest/meta-data/spot/termination-time
```

If the instance has been terminated, the termination time of the spot instance is returned, as shown below.

Note: The termination time refers to the OS time of the spot instance when it's terminated (in UTC+8).

```
2018-08-18 12:05:33
```

If the error code 404 is returned, the instance is not a spot instance or it's not terminated.

For more information, see Querying Instance Metadata.

No Charges When Shut Down for Pay-as-You-Go Instances

Last updated：2024-01-08 09:32:02

Overview

If you enable "No Charge when Shut Down" when shutting down a pay-as-you-go instance, the billing of CPU and memory resources of this instance stops. However the Cloud disks (system disk and data disk), public network bandwidth, images, and other key components of the CVM instance are still billed.
Note:
When this feature is enabled, the instance's CPU and memory resources **will not be retained**, and the public IP address **will be automatically released** after shutdown. For more information about the feature, its use limits, and impacts, see No Charges When Shut down for Pay-as-You-Go Instances.

Directions

Shutting down an instance via console

1. Log in to the CVM console.
2. Choose the appropriate operation method based on your actual needs.
Shutting down a single instance:
2.1.1 Select the instance you want to shut down, and click **More** > **Instance Status** > **Shut down** under the **Operation** column on the right.
2.1.2 Tick **CVM No Charge when Shut down** and click **OK**.
If the instance does not support this feature, **"No Charge when Shut Down" is not supported** will be displayed in the instance list.
Shutting down multiple instances:
2.1.1 Select all the instances you want to shut down and click **Shut down** at the top of the list to shut down instances in batches.
Reasons are given for instances that cannot be shut down.
2.1.2 Tick **CVM No Charge when Shut down** and click **OK**.
If the instance does not support this feature, **"No Charge when Shut Down" is not supported** will be displayed in the instance list.

Shutting down an instance via API

You can use the `StopInstances` API to shut down an instance. For details, please see StopInstances. To enable this feature via API, please add the following parameter:

Parameter Name	Required	Type	Description
StoppedMode	No	String	The "No Charge when Shut down" feature is only available for pay-as-you-go instances. **Valid values:** KEEP_CHARGING: the instance incurs fees after shutdown STOP_CHARGING: no charges when shut down **Default value:** KEEP_CHARGING

Managing Roles

Last updated：2024-01-08 09:37:00

Overview

A Cloud Access Management (CAM) role is a virtual identity with a collection of permissions. It is used to grant the role entity the permissions to access services and resources and perform operations in Tencent Cloud. You can associate the CAM role with a CVM instance to call other Tencent Cloud APIs from the instance using the periodically updated temporary Security Token Service (STS) key. This ensures the security of your SecretKey and helps you implement refined permission control, avoiding the security risks from using persistent keys.
This document describes how to bind, modify, and delete a role.

Advantages

Binding a CAM role to instances comes with the following features and advantages.
You can use the STS temporary key to access other Tencent Cloud services.
You can grant roles associated with different access policies to instances so that the instances are given different access permissions to Tencent Cloud resources, which helps you implement refined permission control.
You don't need to save SecretKey in an instance. Instead, you can easily control the access permissions of the instance by changing the role authorization.

Notes

The instance only allows the role entity that contains `cvm.qcloud.com` to assume the role. For more information, see Concepts.
The instance must reside in a VPC.
An instance can only bind one CAM role at a time.
You can bind, modify or delete a role without paying extra fees.

Directions

Bind/modifying roles

Binding/Modifying one role

Batch binding/modifying roles

1. Log in to the CVM console and click **Instances** on the left sidebar.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Instance Settings** > **Bind/Modify a Role** on the right as shown below:

Tab view: on the page of the target instance, select **More** > **Instance Settings** > **Bind/Modify a Role** in the top-right corner.

3. In the pop-up window, select the role you want to bind, and click **OK**.

1. On the **Instances** page, select the CVM instances for which you want to bind or modify the roles, click **More Actions** > **Instance Settings** > **Bind/Modify a Role** at the top of list, as shown below:

2. In the pop-up window, select the role you want to bind, and click **OK**.

Note:

CVMs modified using this method will have the same role name.

Deleting roles

Deleting one role

Batch deleting roles

1. Log in to the CVM console and click **Instances** on the left sidebar.

2. On the instance management page, proceed according to the actually used view mode:

List view: in the row of the target instance, select **More** > **Instance Settings** > **Delete a Role** on the right as shown below:

Tab view: on the page of the target instance, select **More Actions** > **Instance Settings** > **Delete a Role** in the top-right corner .

3. Click **OK** in the pop-up window.

1. On the **Instances** page, select the CVM instances for which you want to delete the roles, click **More Actions** > **Instance Settings** > **Delete a Role** above the list, as shown below:

2. Click **OK** in the pop-up window.

Enabling and Disabling Hyper-Threading

Last updated：2024-03-26 09:46:28

Operation Scenarios

Hyper-Threading (HT) technology allows the CPU to publicly run two threads per physical core. This means that one physical core now works like two threads that can handle different software threads. By default, HT is generally enabled for Tencent Cloud CVM instances, and disabled only for a few specific CVM instances. In general, HT does not need to be set.

Note:

Enabling HT: It is suitable for scenarios where the cores need to process more information and background tasks in parallel. Enabling HT can significantly improve the computing experience.

Disabling HT: It is suitable for scenarios where performance is better with HT disabled than enabled, such as compute-intensive scenarios.

When purchasing instances or modifying instance specifications, you can set CPU options (determined by the number of threads per core) for some instance specifications as needed. You can adjust the number of threads per core of a CVM instance (that is, the vCPU of the instance) to enable or disable HT as needed.

Instance Limits

For instance families that support HT enabling and disabling, see Instance Types.
The cost does not change when you enable and disable HT.

Directions

Creating an Instance on the Purchase Page

1. Log in to the Tencent Cloud CVM purchase page.

2. Choose **Custom configuration** > **Advanced settings**. On the page that appears, set **CPU options**.

3. Check the box for **Set the threads bound to CPU** and select the number of threads per core, as shown in the figure below.

When you set the number of threads per core to **1**, HT is disabled.

When you set the number of threads per core to **2**, HT is enabled.

When you do not set the number of threads per core, the default HT policy is used for the instance.

4. Click **Next** to create the instance.

Adjusting the Instance Configuration in the Console

1. Log in to the CVM console.

2. Proceed according to the view mode in use.

List View

Tab View

Locate the target instance to be adjusted, click **More**, and select > **Resource adjustment** > **Adjust model and specs** in the Operation column on the right, as shown in the figure below.

On the page of the target instance, click **More actions** at the top right of the page and select **Resource adjustment** > **Adjust model and specs**, as shown in the figure below.

3. Select the target configuration to be modified. If HT adjustment is allowed after the modification, **CPU options** will appear.

4. Check the box for **Specify the CPU binding thread count** and select the number of threads per core, as shown in the figure below.

When you set the number of threads per core to **1**, HT is disabled.

When you set the number of threads per core to **2**, HT is enabled.

When you do not set the number of threads per core, the default HT policy is used for the instance.

5. Click **Next** to complete configuration adjustment.

Calling APIs to Set HT

Creating an instance: You can call the RunInstances API to enable or disable HT when creating an instance. For more information, see Creating an Instance.

Adjusting configuration: You can call the ResetInstancesType API to enable or disable HT when adjusting the configuration. For more information, see ResetInstancesType.

Instance Quota Management Overview

Last updated：2025-02-13 17:21:20

Instance Quota Management Overview

Tencent Cloud imposes certain limits on the quota of the CVM instance purchase quantity for each AZ, billing mode and instance type. CVM quota management provides you with centralized management of instance quotas. Through CVM quota management, you can query the current and used quotas of each instance or apply for an increase in instance quotas based on business needs.

Basic Concepts

Concept	Description
Instance Quota	A Tencent Cloud account can apply for the maximum value of instance purchases.
Quota Unit	Quota units are distinguished by instance specifications: Instance specifications with vCPU < 64 cores, quota unit is **number of cores**. Instance specifications with vCPU ≥ 64 cores, quota unit is **number of CVM instances**.
Current Quota	Current maximum value of instances that can be applied for purchase.
Quota Used	Quota of purchased and used instances.
Apply for Quota	You can apply for a quota increase on the quota management page.
Target Quota	Target quota value applied for. Once the quota application is approved, it will override the original quota value.

Instance Quota Rules

The quota unit for CVM varies by instance specifications. Specific details are as follows:

For instance specifications with vCPU < 64 cores, the quota unit is counted by number of cores.

For example, all SA5 specifications with vCPU less than 64 cores, including SA5.MEDIUM2, SA5.MEDIUM4, SA5.LARGE8, etc., share a total SA5 quota, which is calculated by the number of cores.

For instance specifications with vCPU ≥ 64 cores, the quota unit is counted by number of CVM instances.
For example, the SA5.20XLARGE160 has 80 vCPU cores, so SA5.20XLARGE160 has a separate quota, which will be calculated by the number of CVM instances.

Features and Usage

If you need to understand	Please Refer to
Query Each Instance's Current Quota and Usage	Query Instance Purchase Quota
Process for Applying for Instance Quota Increase	Increase Instance Purchase Quota
Query Quota Request Records	Query Quota Request Records

Query Instance Purchase Quota

Last updated：2025-02-08 15:15:14

Overview

This article explains how to view the quota details of each instance through the Quota Management Console.

Operation Steps

1. Log in to the Cloud Virtual Machine Console, select Quota Management > Instance Quota in the left navigation bar.

2. On the **Instance Quota** page, you can view the current quota and used quota of specific instance types in each region, each AZ, and specific billing mode by selecting region, AZ, billing mode and quota object to search instances.

Instance quota

	Region	Singapore ▾	Availability zone	Singapore Zone 3 ▾	Billing mode	Monthly subscription ▾	Quota object	All selected, and keyword search ▾

Quota object	Availability zone	Instance configurati...	Quota unit ⓘ	Current quota	Used quota	Instance billing mode	Operation
Standard SA5	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		4 used/ in total	Monthly subscription	Increase quota Purchase instance
Standard S8	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		0 used/ in total	Monthly subscription	Increase quota Purchase instance
Standard SA	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		0 used/ in total	Monthly subscription	Increase quota Purchase instance
Standard S8	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		0 used/ in total	Monthly subscription	Increase quota Purchase instance
Standard S5	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		0 used/ in total	Monthly subscription	Increase quota Purchase instance
Standard SA2	Singapore Zone 3	Instance specifications with vCPU less than 64 cores	Number of Core		0 used/ in total	Monthly subscription	Increase quota Purchase instance
SA5.	Singapore Zone 3	512-core 2304GB	Number of Instances		0 used/ in total	Monthly subscription	Increase quota Purchase instance

Increasing Instance Purchase Quota

Last updated：2025-02-08 15:19:39

Overview

If the existing instance purchase quota cannot meet your actual demand, you can apply for an increase in the CVM instance purchase quota as instructed in this document.

Operation Steps

1. Log in to the Cloud Virtual Machine Console, and select Quota Management > Instance Quota from the left navigation bar.

2. On the Instance Quota page, filter by selecting **Region**, **AZ**, **Billing Mode**, and **Quota object** to quickly locate quota information.

3. In the instance quota list, click **Increase quota** in the corresponding row of the quota object to be increased.

4. In the **Instance quota request** popup window, fill in the **Target quota** and **Reason for request**, and choose whether to notify the application result.

5. After confirming that the information is correct,   click **Submit a request**.

You can view the request status in the Quota Request Records page. For more information about quota request records, please refer to Query Instance Purchase Quota.

Querying Quota Application Records

Last updated：2025-02-08 15:21:54

Overview

This article describes how to view the progress and results of quota applications on the Quota Management page.

Operation Steps

1. Log in to the Cloud Virtual Machine Console, and select Quota Management > Quota Request Records from the left navigation bar.

2. On the **Quota Request Records** page, filter by selecting **Region**, **Availability Zone**, and **Billing Mode** to quickly locate quota request records and view the application progress and results.

application status description is as follows:

Approved: it means the quota increase is successful, you can go to the purchase page immediately to choose.

Under approval: it means it is under approval. If you set a notification of the need for results when increasing instance purchase quota, we will inform you immediately after the result is determined. Thank you for your patience.

Unapproved: it means your application was not approved. It is recommended to consider lowering the target quota and reapplying. If you have any questions, you can submit a ticket to provid.

Reserved Instances
Splitting a Reserved Instance

Last updated：2024-01-08 09:37:00

Overview

You can split a reserved instance into multiple reserved instances with smaller normalization factors, and apply them to smaller on-demand instances.

Rules and Limits

1C1G, 1C2G, and 2C4G reserved instances cannot be split.

You can change the instance size, but not the instance type.

The region and availability zone cannot be changed.

The sum of normalization factors cannot be changed.

Make sure the CPU:MEM ratio is the same before splitting.

One reserved instance can be split into a maximum of 100 reserved instances.

Directions

1. Log in to the CVM console;

2. Click **Reserved Instances** in the left sidebar;

3. On the **Reserved Instances** page, find the original reserved instance and click **Split** under the **Operation** column;

4. In the pop-up window, configure the name, CPU, memory and quantity of new reserved instances.

Execution Result

After splitting, the original reserved instance goes to the **Expired** status, while the new reserved instances are in **Activated** status.

Definitions

Normalization factor: It indicates the CPU performance of a For the splitting and merging of reserved instances, the normalization factor is calculated based on the number of vCPUs.

Sum of normalization factor = Normalization factor of the instance size × Instance quantity. The sum of normalization factors must keep the same as the one before splitting.

Merging Reserved Instances

Last updated：2024-01-08 09:37:00

Overview

You can merge multiple reserved instances into a single reserved instance with a larger normalization factor, so as to apply it to larger on-demand instances.

Rules and Limits

The original reserved instances are in the **Activated** status.

The original reserved instances are in the same availability zone.

The payment methods for the original reserved instances are the same.

The original reserved instances have the same operating system (Windows/Linux) and expiration time.

Merging the 1C1G, 1C2G, and 2C4G reserved instances is not supported.

You can modify the instance size but not the instance family.

Modifying the region or availability zone of a reserved instance is not supported.

The total normalization factors of destination reserved instances must be equal to that of the original reserved instance.

Make sure the CPU:MEM ratio is the same before merging.

Directions

1. Log in to the CVM console.

2. Click **Reserved Instances** in the left sidebar.

3. On the **Reserved Instances** page, find the original reserved instances and click **Merge** under the **Operation** column.

4. In the pop-up window, select the reserved instances to be merged, and enter a new name for the destination reserved instance.

Execution Result

After merging, the original reserved instances enter the **Expired** status, while the new reserved instance is in **Activated** status.

Definitions

Normalization factor: It indicates the CPU performance of an instance. For the splitting and merging of reserved instances, the normalization factor is calculated based on the number of vCPUs.

Sum of normalization factor = Normalization factor of the instance size × Instance quantity. The sum of normalization factors must keep the same as the one before splitting.

Images

Creating a Custom Image

Last updated：2024-05-16 11:03:38

Overview

Besides public images, you can also create custom images, with which you can create CVM instances with the same configurations.
Note:
Images use the CBS snapshot service for data storage:
Upon the creation of a custom image, a snapshot is automatically created and associated with the image. As a result, retaining custom images incurs costs. For more information, please see Billing Overview.

Note

Each region supports a maximum of 500 custom images.
When the Linux instance has a data disk attached, and you only create an image on the system disk, make sure
`/etc/fstab` does not include data disk configuration. Otherwise, instances created with this image cannot be started normally.
The creation process takes ten minutes or longer, which depends on the data size of the instance. Please prepare in advance to avoid business impacts.
You can not create an image by using a CBM instance in the console or via API. You can use CVM to create them.
If your Windows instance needs to enter a domain and uses a domain account, please execute Sysprep before creating a custom image to ensure that the SID is unique after the instance enters the domain. For more information, please see Ensuring Unique SIDs for CVMs Using Sysprep.

Directions

Console
API

Shut down an instance (optional)

1. Log in to the CVM console and check whether the corresponding instance needs to be shut down.
Note:

For CVMs created based on public images after July 2018, you can create images without shutting down the instance. For other CVMs, shut down the instance before creating a custom image to ensure that the image has the same environment deployment as the current instance.

If the instance needs to be shut down, proceed to the next step.

If the instance doesn't need to be shut down, please proceed to Create a custom image.

2. On the instance management page, proceed according to the actually used view mode:

List view: In the row of the target instance, select **More** > **Instance status** > **Shut down** on the right as shown below:

Tab view: Select **Shut down** on the instance details page.

Create a custom image

1. On the instance management page, proceed according to the actually used view mode:

List view: Select **More** > **Create image**.

Tab view: Select **More actions** > **Create image** in the top-right corner.

2. In the **Create a custom image** popup window, complete the configuration.

Image name and **Image description**: Custom name and description.

Tag: You can add tags for the instance as needed, which are used to categorize, search for, and aggregate cloud resources. For more information, see Overview.

Note:

To create custom images that include system disks and data disks, please submit a ticket.

3. Click **Create image**.

You can click **Image** on the left sidebar to view the creation progress in the **Image** page.

Use the custom image to create an instance (optional)

Select the image you created in the image list, and click **Create instance** on the right side to purchase a server with the same configuration as the image, as shown in the following figure:

You can use the `CreateImage` API to create a custom image. For more information, see CreateImage.

Best Practices

Migrating the data on a data disk

If you need to keep the data on the data disk of the original instance when launching a new instance, you can first take a snapshot of the data disk, and then use this data disk snapshot to create a new CBS data disk.

For more information, please see Creating Cloud Disks Using Snapshots.

Sharing Custom Images

Last updated：2024-05-16 10:47:13

Overview

A **shared image** is a **custom image** that a user shared to other users. With a shared image, you can get the necessary components from other users and add your custom contents.
Caution:
Tencent Cloud does not guarantee the integrity or security of shared images. Please only use shared images from trusted sources.

Limits

Each image can be shared with a maximum of 500 Tencent Cloud accounts.
You can not change the name and description of the images shared from others. They can only be used to create or reinstall CVM instances.
When you share an image to others, the shared replicas do not count against your image quota.
If you need to delete a custom image that is shared with others, you need to cancel all the sharing relations first. For more information, see Cancelling Image Sharing. You can not delete an image shared from others.
Custom images can only be shared with accounts in the same region as the source account. To share an image with users in another region, you need to copy it to the target region before sharing.
The shared images that you obtain from others cannot be re-shared.

Directions

Obtaining the ID of the root account to which you want to share the image

To share an image to another user, you need to know their root account ID. They can check their root account ID as instructed below:

1. Log in to the CVM console.
2. Click the account name in the top-right corner and select **Account Information**.
3. View and note down the account ID.
4. Notify the other party to send the obtained account ID to itself.

Sharing images

Sharing images in the console

Sharing images via an API

1. Log in to the CVM console and select **Images** on the left sidebar.

2. Click the **Custom Image** tab to enter the custom image management page.

3. In the custom image list, select the custom image you want to share and click **Share** in the **Operation** column.

4. In the **Shared Image** pop-up window, enter the ID of the account with which you want to share the selected image, and click **Share**.

5. Tell the other user to log in to the CVM console, and select **Images** > **Shared Image** to view the image you have shared.

Repeat the steps above to share an image with multiple users.

You can use the ModifyImageSharePermission API to share images.For more information, please see ModifyImageSharePermission.

Related Operations

Sharing image with Lighthouse

You can share a custom image between Lighthouse and CVM to implement fast offline service migration. You can also use a shared image to quickly create instances and then get the needed components from them or add custom content to them.

For more information, please see ModifyImageSharePermission.

Cancelling Image Sharing

Last updated：2024-05-16 10:46:51

Scenario

This document describes how to cancel custom image sharing. You can cancel your image sharing status with other users at any time. This does not affect instances created by other users using this shared image, but they can no longer see the image nor create new instances using this image.

Directions

Cancel image sharing through the console

Cancel image sharing through API

1. Log in to the CVM Console.On the left sidebar, click Images.

2. Select **Custom Image** tab to enter the custom image management page.

3. In the custom image list, select the custom images you want to cancel sharing and click **More** > **Cancel Sharing**.

4. On the new page, select the unique ID of the account from which you want to cancel the image sharing and click **Cancel Sharing**.

5. In the pop-up window, click **OK** to cancel image sharing.

You can use the ModifyImageSharePermission API to cancel image sharing. For more information, see ModifyImageSharePermission.

Deleting Custom Images

Last updated：2024-01-08 09:37:00

Scenario

This document describes how to delete custom images.

Notes

Before deleting custom images, please note the following items:

After a custom image is deleted, it can no longer be used to start a new CVM instance, but will not affect instances that have already been started. If you want to delete all instances started from this image, see Reclaiming Instances or Terminate Instances.

A custom image that has been shared with others cannot be deleted. To delete it, you need to cancel image sharing first. For more information, see Cancel Image Sharing.

You can only delete the custom image, not common image or shared image.

Directions

Delete images through the console

Delete images through API

1. Log in to the CVM Console, on the left sidebar, click **Images**.

2. Select the **Custom Image** tab to enter the custom image management page.

3. Select the method to delete custom images based on actual needs.

Deleting a single image: locate the custom image to be deleted in the image list and click **More** > **Delete**.

Deleting multiple images: select all custom images to be deleted in the image list and click **Delete** on the top.

4. In the pop-up window, click **OK**.

If the deletion fails, possible reasons will be prompted.

You can use the DeleteImages API to delete images. For details, see Delete Images.

Image Replication

Last updated：2024-12-03 14:54:29

Overview

Features

Image replication offers two options, **Cross-region replication of custom images** and **Intra-region replication of shared images**.

Type	Use Case	Description
Cross-region replication of custom image	Deploy the same CVM instance **across regions** quickly.	Copy a custom image to another region, and use the created copy to create a CVM in the new region.
Intra-region replication of shared image	Make a copy of a shared image and use it as a custom image.	The created custom image is not subject to the limits of shared images.

Limits

Custom images can be replicated across regions, and shared images support intra-region replication.

Due to the data security limitations on financial zone, it supports image duplication within a financial zone, but not supports image duplication cross financial zone and non-financial zone.

Image replication is free of charge. But you need to pay for snapshot service for store the copied custom images.

Image replication takes 10 to 30 minutes.

Cross-region replication is not available for full images.

Methods

Cross-region Replication of Custom Images

Console

API

1. Log in to the CVM console.

2. In the left sidebar, click **Images** to enter the image management page.

3. Select the region where the original image you want to copy resides, and click the **Custom image** tab.

For example, select Guangzhou region.

4. Find the instance whose image needs to be copied, click **More** > **Cross-region replication**.

5. In the pop-up window, select the regions where the image will be copied to and click **OK**.

After the copying is completed, the image list in the destination regions will display images with the same name and different IDs.

6. Switch to a destination region. Select the copied image in the image list, and click **Create instance** to create the same CVM instance.

You can use the `SyncImages` API to copy an image. For more information, see SyncImages.

Intra-region Replication of Shared Images

Console

API

1. Log in to the CVM console.

2. In the left sidebar, click **Images** to enter the image management page.

3. Select the region of the source image, and click the **Shared image** tab.

For example, select Guangzhou region.

4. Find the instance whose image needs to be copied, click **More** > **Intra-region replication**.

5. In the pop-up window, select the regions where the image will be copied to and click **OK**.

After the copying is completed, the image list in the destination regions will display images with the same name and different IDs.

6. Switch to the **Custom image** tab. Select the successfully copied image, and click **Create instance** to create the same CVM instance. The copied image has features like other custom images.

You can use the `SyncImages` API to copy an image. For more information, see SyncImages.

Importing Images Overview

Last updated：2024-10-25 15:36:38

In addition to creating a custom image, Tencent Cloud allows you to import images. You can import an image file of the system disk on a local or a different server into CVM custom images. You can use the imported image to create a CVM or reinstall the operating system for an existing CVM.

Import Preparation

Prepare an image file that meets the import requirements.

Linux System Type Image Limit

Windows System Type Image Limit

Image Attributes	Condition
Operating system	Images based on CentOS, CentOS Stream, Ubuntu, Debian, RedHat, OpenSUSE, CoreOS, FreeBSD, Kylin (Kirin), UnionTech, TencentOS, Fedora, AlmaLinux, Rocky Linux, OpenCloudOS, and Other Linux release version Support 32-bit, 64-bit, and arm_64-bit.
Image format	Support image formats of RAW, VHD, QCOW2, and VMDK. Use qemu-img info imageName | grep 'file format' to check the image format. For image files of other formats, you may refer to Converting Image Format to convert and then import them.
File system and partition	Support xfs, ext3, and ext4 file systems, and support MBR and GPT partitions. **Note:** Support the ext4 file system, but cannot include the feature `64bit` . The features `project` and `quota` cannot appear in pairs. You can run the following command to view the list of features included in the ext4 file system. ```tune2fs -l <ext4 file system disk path>	grep features``` It is not recommended for the system disk to have multiple partitions. We recommend a single root partition only. If there are multiple partitions, it is not recommended to have other partitions after the one where the system resides. Otherwise, disk expansion will fail. It is not recommended to use LVM to create system disk partitions (root partition), as it may lead to issues such as instance startup, password modification, and SSH login via password.
Image size	The actual image size must not exceed 1024 GB. Use `qemu-img info imageName |`

	`grep 'disk size'` to check the actual image size; The image vsize must not exceed 1024 GB. Use `qemu-img info imageName	grep 'virtual size'` to check the image vsize. **Note**: When importing an image, review the size of the image after converting it to QCOW2 format.
Network	Tencent Cloud provides the `eth0` network interface by default for instances. Users can query the instance's network configuration through the metadata service within the instance. For details, see Instance Metadata.
Driver	The image must install the Virtio driver for the virtualization platform KVM. For details, see Linux Import Image Check Virtio Driver. The image needs to install cloud-init. For details, see Linux Import Image Install cloud-init. If the image cannot install cloud-init for other reasons, refer to Force Importing Image to configure the instance manually.
File system	To ensure that the Linux system can correctly identify the disk when the file system is started, check and correctly configure the GRUB file disk identification method. For details, see Configure the GRUB file disk identification method to UUID. To ensure that the Linux system can correctly identify the disk when the file system is mounted, check and correctly configure the fstab file disk identification method. For details, see Configure the fstab file disk identification method to UUID.
Kernel	We recommend a native kernel for the image. Kernel modification may cause the CVM to fail to import.
Region limit	The image import service outside the Chinese mainland only supports COS files in the same region. That is, a COS link in the same region needs to be used for import.

Image Attributes	Condition
Operating system	Windows Server 2022, Windows Server 2019, Windows Server 2016, Windows Server 2012, Windows Server 2008, and Other Windows-related versions Support 32-bit, 64-bit, and arm_64-bit.
Image format	Support RAW, VHD, QCOW2, and VMDK image formats. Use `qemu-img info imageName	grep 'file format'` to check the image format.
File system and partition	Support NTFS file system, and support MBR and GPT partitions. Support creation of multiple partitions on the system disk. Ensure that no other partitions exist after the boot partition (C Drive) on the disk where the system resides, otherwise the instance may fail to start or the disk expansion may fail.
Image size	The actual image size must not exceed 1024 GB. Use qemu-img info imageName	grep 'disk size' to check the actual image size.

	The image vsize must not exceed 1024 GB. Use qemu-img info imageName | grep 'virtual size' to check the image vsize. **Note**: When importing an image, review the size of the image after converting it to qcow2 format.
Network	Tencent Cloud provides the local connection network interface for instances by default. Users can query the instance's network configuration through the metadata service within the instance. For details, see Instance Metadata.

Driver	Images must install the Virtio driver for the virtualization platform KVM. Windows system by default does not have the Virtio driver installed. Users can install the Windows Virtio driver and then export the local image. The download address for the Windows Virtio driver is as follows (download according to your actual network environment):
	Public network download address:
	`http://mirrors.tencent.com/install/windows/virtio_64_1.0.9.exe`
	Private network download address:
	`http://mirrors.tencentyun.com/install/windows/virtio_64_1.0.9.exe`

Region limit	The image import service outside the Chinese mainland only supports COS files in the same region. That is, a COS link in the same region needs to be used for import.
Others	The imported Windows system image does not provide Windows activation service.

Directions

1. Log in to the CVM console and click **Images** on the left sidebar.

2. Select **Custom image** and click **Importing an image**.

3. As prompted in the operation interface, first enable COS, and then create a bucket. Uploading an Object the image file to the bucket and get the image file URL.

4. Click **Next**.

5. Complete the configurations and click **Import**.

Parameter	Required	Description
Region	Yes	Ensure that the region of the custom image you need to import matches the region where you need to create an instance. Custom images from this region cannot be directly used to create instances in other regions. If you need to use the current image in another region, you can replicate your custom image across regions through image replication.
System Disk Files	<u>Yes</u>	The system disk files contain the critical components such as the kernel, library files, and drivers needed for the server operating system. By using the system disk files, the server can start and run the operating system to provide basic services and features. Import method:

		Select the storage from the COS list. In the drop-down list, select the COS bucket where the file is contained, and then select the corresponding image file. The system will automatically obtain the URL of the image file. Note that only COS files within the local domain can be selected in this way. Enter the COS object address. Go to the COS console, find the bucket from **Bucket List** where the image file is located, and find the image file step by step. On the image file details page, click **Copy Temporary Link** to copy the URL of the image file.
Operating System	Yes	1. **Linux and Windows operating systems are supported.** Ensure that it matches the operating system type of the image file. 2. **System platform** Identifies the operating system platform to which the image file is imported. Ensure that it matches the operating system platform of the image file. Linux operating system: CentOS, CentOS Stream,Ubuntu, Debian, RedHat, and other commonly used system platforms can be selected. If the system platform is not listed, select Other Linux. Windows operating system: Windows Server 2022, Windows Server 2019, and other commonly used system platforms can be selected. If the system platform is not listed, select Other Windows. 3. **System Version** Identifies the version of the Linux operating system to which the image file is imported. For example, CentOS supports multiple versions, including 8, 7, 6, and 5. Ensure that it matches the operating system version of the image file. 4. **System Architecture** 32-bit, 64-bit, and ARM 64-bit are supported. Ensure that it matches the system architecture of the image file.
Data Disk Files (only available in the open beta after you contact customer service and apply for that)	No	The data disk files are used to store user data such as application data, user documents and database files. If your image has data disk files, add a data disk by configuring this. Import method: Select the storage from the COS list. In the drop-down list, select the COS bucket where the file is contained, and then select the corresponding image file. The system will automatically obtain the URL of the image file. Note that only COS files within the local domain can be selected in this way. Enter the COS object address. Go to the COS console, find the bucket from **Bucket List** where the image file is located, and find the image file step by step. On the image file details page, click **Copy Temporary Link** to copy the URL of the image file.
Image Name	Yes	The image name displayed after the image file is imported. Note that the image name only supports Chinese characters, letters, digits, hyphens (-), underscores (_), and periods (.), with a maximum of 60 characters.

Image Description	No	Add image description for easy management.
Tag	No	Set a tag to facilitate search and management.

More configuration items:

Parameter	Required	Description
Import Method	No	If your image cannot be imported correctly, you can select the **Enable Forcible Import** option. This method only checks the integrity of the file and will not block the import process due to driver or configuration issues. For details, see Forcibly Importing Image.
Boot Mode	Yes	Ensure that the selected boot mode matches the boot mode of the image file. Otherwise, using this image may affect the normal startup of your instance. For details, see Best Practices for Boot Mode Legacy BIOS and UEFI.

You will be notified about the import result via Message Center.

Import Failed

After you import images in the console, the task may fail for some reasons. In case of task failure, you can check the error code in the Console or Message Center and troubleshoot according to the following content.

Error Code	Reason	Recommended Solution
InvalidUrl	The entered COS link is invalid. Possible reasons are as follows: The entered image file link is not a Tencent Cloud COS link. The object address of the COS has no permission for public read and private write. The access permission of COS files is private read, but the signature has expired. When an image is imported outside the Chinese mainland, a COS link not in the same region was used; the image import service outside the Chinese mainland only supports COS servers in the same region.	Check whether the COS link matches the imported image link.

	The current COS file has been deleted.	
InvalidFormatSize	The format or size of the imported image does not conform to the restrictions of Tencent Cloud image import feature. Imported images support image file formats of qcow2, vhd, vmdk, and raw. The size of the system disk image should not exceed 1024 GB, and a single data disk should not exceed 2048 GiB (based on the image file converted to qcow2 format)	The image must meet the restrictions on `image format` and `image size` specified in Import Preparation. According to the image format conversion content of Linux Image Creation, convert the image file into an appropriate file format and simplify the image content to meet the size requirements, and then re-import the image. For files exceeding the size limit, images can also be migrated using the Offline Instance Migration feature.
VirtioNotInstall	Virtio driver not installed: Tencent Cloud uses KVM virtualization technology and requires users to install the Virtio driver in the image to be imported. Except for a few customized Linux operating systems, most Linux operating systems have the Virtio driver installed; Windows operating systems require users to manually install the Virtio driver.	For Linux image import, see Checking Virtio Drivers in Linux. For Windows image import, see Creating Windows Images to install the Virtio driver.
CloudInitNotInstalled	cloud-init not installed: Tencent Cloud uses the open-source program cloud-init for CVM initialization. Therefore, CVM initialization will fail if the cloud-init program is not installed.	For Linux image import, see Installing Cloud-Init on Linux. For Windows image import, see Installing Cloudbase-Init on Windows.
PartitionNotPresent	Partition information is not found, and the imported image is incomplete.	The image is corrupted. Check if the boot partition was included when creating the image. It may be caused due to an incorrect image creation method.
RootPartitionNotFound	The imported image does not contain a root partition. Possible reasons are as follows: The installation package was uploaded; The data disk image was uploaded; The boot partition image was uploaded;	The image is corrupted. It may be caused due to an incorrect image creation method.

	An incorrect file was uploaded.	
InternalError	Other errors	Contact Customer Service for assistance.

Test Result Explanation and Suggestions for Fix

Image import will automatically detect image availability. You can view the image import results in the console. We recommend refining the image according to the test results to ensure feature availability.

Linux Operating System Inspection Items and Instructions

Inspection Items	Inspection Instructions	Test Results	Suggestions for Fix
Virtio	Whether the Virtio driver is installed in the image	Supported - Satisfied KernelNotSupported - Kernel not supported BlkDriverNotFound - Disk driver not found NetDriverNotFound - Network driver not found	**Strongly recommended to fix:** For details about how to fix, refer to Checking Virtio Drivers in Linux.
CloudInit	Whether cloud-init is installed in the image	Supported - Satisfied NotSupported - Not supported	**Strongly recommended to fix:** For details about how to fix, refer to Installing cloud-init on Linux.
CloudinitConfig	Whether cloud-init is configured correctly	Supported - Satisfied Invalid - cloud.cfg content cannot be loaded InvalidCloudFinalModules - scripts-user not found in the cloud_final_modules configuration InvalidDatasourceList - ConfigDrive not found in datasource_list	Strongly recommended to fix: 1. Modify the cloud-init configuration file: Download cloud.cfg based on different operating systems. Click here to download cloud.cfg for Ubuntu operating system. Click here to download cloud.cfg for CentOS operating system. Click here to download cloud.cfg for OpenSUSE / SUSE operating system. 2. Replace the content of `/etc/cloud/cloud.cfg` with that of the downloaded cloud.cfg file.

| Fstab | The file `/etc/fstab` in the image is configured with system disk device mounting information. Incorrect configurations may cause system startup exceptions, such as device mounting information not existing or incorrect device UUID. | Supported - Satisfied DeviceNotFound - The configuration contains a non-existent device DeviceConfig - `/dev/vd*` configuration is used | **Recommended fix:**
1. If your operating system is not FreeBSD, refer to Configure the fstab file disk identification method to UUID to fix.
2. If your operating system is FreeBSD.
Add Label.

```# Set Label
tunefs -L root /dev/da3
View results
ls /dev/ufs/root```

Modify /etc/fstab.

```# File backup
cp /etc/fstab
/etc/fstab.bak
Change to
vim /etc/fstab```

Modify the result.

```/dev/ufs/home   /
ufs   rw
2   2``` |
| Grub | Whether the grub boot file in the image is normal. For example, if device configurations are abnormal, related devices do not exist, or there are UUID errors, it may cause a system startup anomaly. Do not check this item when the image is FreeBSD. | Supported - Satisfied NotSupported - Not configured using UUID | **Recommended fix:**
For details about how to fix, see Configure the GRUB file disk identification method to UUID. |

Selinux	Whether SELinux is disabled for the image. It is not recommended to enable SELinux for cloud images, as it may cause a system startup anomaly. Do not check this item when the image is FreeBSD.	Supported - Satisfied ConfigNotFound - The /etc/selinux/config file does not exist when SELinux is enabled Enforcing - Highest level enabled	**Recommended fix:** Suggestions for fix
			
			```# Switch to the root user vim /etc/selinux/config```
			
			Change to
			
			```# This file controls the state of SELinux on the system. # SELINUX= can take one of these three values: # enforcing – SELinux security policy is enforced. # permissive – SELinux prints warnings instead of enforcing. # disabled – No SELinux policy is loaded. SELINUX=permissive # SELINUXTYPE= can take one of these two values: # targeted – Targeted processes are protected, # mls – Multi Level Security protection. SELINUXTYPE=targeted```

OnlineResizeFS	Whether the image supports root partition automatic scaling. The image has `cloud-init,` `growpart` `(gpart,` `parted,` `growpart)`	Supported - Satisfied NotSupported - Not supported	**Recommended fix:** When the operating system is CentOS 6, CentOS 7, or TencentOS Server 2 version. ```# root user yum install cloud-utils- growpart``` When the operating system is CentOS 8, CentOS Stream, Rocky Linux, or TencentOS Server 3 version.

	commands installed. After the image is used to create an instance, the root partition will automatically scale. For example, if your image space size is 20 GB, and the system disk size is 100 GB during instance creation, the root partition will automatically scale to 100 GB after the instance is created.		```# root user dnf -y install cloud-utils-growpart``` When the operating system is Ubuntu ```# root user apt-get install cloud-guest-utils```
Network	Whether the network configuration in the image complies with the cloud-init standard. Only Debian and Ubuntu operating systems need to checked.	Supported - Satisfied Unsupported - Not supported	**Recommended fix:** `/etc/network/interfaces` must contain `source` `/etc/network/interfaces.d/*`
SupportBareMachine	Whether the image supports Bare Metal.	Supported - Satisfied Unsupported - Not supported	**Optional fix:** If you are not using Bare Metal instances, do not need to pay attention to; if the image does not meet the requirements for Bare Metal instances, you can contact customer service for assistance.
TimeSync	Whether NTP is installed to	Supported - Satisfied	**Optional fix:**

| | maintain time synchronization. | Unsupported - Not supported | When the operating system is CentOS 6, CentOS 7, or TencentOS Server 2 version. |

```
# root user
yum install ntpd
```

When the operating system is CentOS 8, CentOS Stream, Rocky Linux, or TencentOS Server 3 version.

```
# root user
dnf install chrony
```

When the operating system is Ubuntu

```
# root user
apt-get install ntp
```

Windows Operating System Inspection Items and Instructions

Inspection Items	Inspection Instructions	Test Results	Suggestions for Fix
Virtio	Whether the Virtio driver is installed in the image	Supported - Satisfied BlkDriverNotFound - Disk driver not found	**Strongly recommended to fix:** For details about how to fix, refer to section Checking or Installing Virtio Drivers in Creating Windows Images.
Cloudbase	Whether cloudbase is installed in the image	Supported - Satisfied NotSupported - Not supported	**Strongly recommended to fix:** For details about how to fix, refer to Installing Cloudbase-Init on Windows.
CloudbaseConfig	Whether cloudbase is configured correctly	Supported - Satisfied NotSupported - Not supported	**Strongly recommended to fix:** For details about how to fix, refer to Modify the cloudbase-init configuration file.
SupportBareMachine	Whether the image supports Bare Metal	Supported - Satisfied NotSupported - Not supported	**Optional fix:** If you are not using Bare Metal instances, do not need to pay

			attention to; if the image does not meet the requirements for Bare Metal instances, you can contact customer service for assistance. Self-service driver download address: Windows Server 2016 Windows Server 2019 The downloaded file is a ZIP file. After extraction, right-click the bnxtnd.inf file and select **Install.**

Forcibly Importing Image

Last updated：2024-01-08 09:37:01

Scenario

If you cannot install cloudinit in your Linux image, use **Forced Image Import** to import the image. If you use this image for import, which does not have cloudinit installed, Tencent Cloud cannot initialize your CVM. In this case, you need to set up the script on your own to configure the CVM based on the configuration file provided by Tencent Cloud. This document describes how to configure the CVM if the image is forcibly imported.

Tencent Cloud provides the user with CDROM device containing the configuration information. The user needs to mount CDROM and read the information of `mount_point/qcloud_action/os.conf` for configuration. If other configuration data or UserData needs to be used, the user can directly read files under `mount_point/` .

os.conf Configuration File

The content of os.conf is as follows.

```
hostname=VM_10_20_xxxx
password=GRSgae1fw9frsG.rfrF
eth0_ip_addr=10.104.62.201
eth0_mac_addr=52:54:00:E1:96:EB
eth0_netmask=255.255.192.0
eth0_gateway=10.104.0.1
dns_nameserver="10.138.224.65 10.182.20.26 10.182.24.12"
```

Note:

The parameter names above are for reference, and the values are used as examples only.

The description of each parameter in the os.conf configuration file is as follows:

Parameter Name	Description
hostname	CVM name
password	Encrypted password
eth0_ip_addr	LAN IP of eth0
eth0_mac_addr	MAC address of eth0
eth0_netmask	Subnet mask of eth0

eth0_gateway	Gateway of eth0
dns_nameserver	DNS resolution server

Limits

The image must meet the limits on Linux images as outlined in Import Images, except for cloudinit.

The system partition for importing the image is not full.

The imported image contains no vulnerability that can be exploited remotely.

We recommend you change the password immediately after the instance is created successfully with the forcibly imported image.

Notes

Note the following when configuring script parsing:

The script is executed automatically at startup. Please implement this requirement based on your operating system.

Mount `/dev/cdrom` and read `qcloud_action/os.conf` file under the mount point to obtain the configuration information.

The password placed in CDROM by Tencent Cloud is encrypted. You can set new password with `chpasswd -e` . **Note that the encrypted password may contain special characters. We recommend you place it in a file and then set the password with** `chpasswd -e < passwd_file` .

When you use the forcibly imported image to create an instance and then create an image, you need to ensure that the script will still be executed to ensure that the instance is configured correctly. You can also install cloudinit in this instance.

Directions

Note:

Tencent Cloud provides a script sample based on CentOS. You can refer to it to create script for your images. During the creation, note that:

The script must be properly placed in the system before image import.

The script is not applicable to all operating systems. You need to modify it according to your own operating systems.

1. Create an `os_config` script based on the following script sample.

You can modify the script as needed.

```
#!/bin/bash
```

```
### BEGIN INIT INFO
# Provides:        os-config
# Required-Start:   $local_fs $network $named $remote_fs
# Required-Stop:
# Should-Stop:
# Default-Start:    2 3 4 5
# Default-Stop:     0 1 6
# Short-Description: config of os-init job
# Description: run the config phase without cloud-init
### END INIT INFO
##################user settings###################
cdrom_path=`blkid -L config-2`
load_os_config() {
    mount_path=$(mktemp -d /mnt/tmp.XXXX)
    mount /dev/cdrom $mount_path
    if [[ -f $mount_path/qcloud_action/os.conf ]]; then
        . $mount_path/qcloud_action/os.conf
        if [[ -n $password ]]; then
            passwd_file=$(mktemp /mnt/pass.XXXX)
            passwd_line=$(grep password $mount_path/qcloud_action/os.conf)
            echo root:${passwd_line#*=} > $passwd_file
        fi
        return 0
    else
        return 1
    fi
}
cleanup() {
    umount /dev/cdrom
    if [[ -f $passwd_file ]]; then
        echo $passwd_file
        rm -f $passwd_file
    fi
    if [[ -d $mount_path ]]; then
        echo $mount_path
        rm -rf $mount_path
    fi
}
config_password() {
    if [[ -f $passwd_file ]]; then
        chpasswd -e < $passwd_file
    fi
}
config_hostname(){
    if [[ -n $hostname ]]; then
        sed -i "/^HOSTNAME=.*/d" /etc/sysconfig/network
        echo "HOSTNAME=$hostname" >> /etc/sysconfig/network
```

```
        fi
}
config_dns() {
    if [[ -n $dns_nameserver ]]; then
        dns_conf=/etc/resolv.conf
        sed -i '/^nameserver.*/d' $dns_conf
        for i in $dns_nameserver; do
            echo "nameserver $i" >> $dns_conf
        done
    fi
}
config_network() {
    /etc/init.d/network stop
    cat << EOF > /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
IPADDR=$eth0_ip_addr
NETMASK=$eth0_netmask
HWADDR=$eth0_mac_addr
ONBOOT=yes
GATEWAY=$eth0_gateway
BOOTPROTO=static
EOF
    if [[ -n $hostname ]]; then
        sed -i "/^${eth0_ip_addr}.*/d" /etc/hosts
        echo "${eth0_ip_addr} $hostname" >> /etc/hosts
    fi
    /etc/init.d/network start
}
config_gateway() {
    sed -i "s/^GATEWAY=.*/GATEWAY=$eth0_gateway" /etc/sysconfig/network
}
#################init####################
start() {
    if load_os_config ; then
        config_password
        config_hostname
        config_dns
        config_network
        cleanup
        exit 0
    else
        echo "mount ${cdrom_path} failed"
        exit 1
    fi
}
RETVAL=0
case "$1" in
```

```
    start)
        start
        RETVAL=$?
    ;;
    *)
        echo "Usage: $0 {start}"
        RETVAL=3
    ;;
esac
exit $RETVAL
```

2. Place the `os_config` script in the `/etc/init.d/` directory and execute the following command.

```
chmod +x /etc/init.d/os_config
chkconfig --add os_config
```

3. Execute the following command to check whether `os_config` has been added to the startup service.

```
chkconfig --list
```

Note:

You must ensure that the script is correctly executed. If you fail to connect to the instance via SSH or network exception occurs after the image import, try to connect to the instance via the console to execute the script again. If such problems remain, contact the customer service.

Creating an Image
Preparing a Linux Image

Last updated：2024-01-08 09:37:01

Overview

This document describes how to create an image of the system disk of a Linux server.

Directions

Preparations

Check the following before you start:

Note:

If you need to prepare and export a data disk image, skip this operation.

Checking the Partitioning and Starting Mode of the OS

1. Run the following command to check whether the OS partition is an MBR partition.

```
sudo parted -l /dev/sda | grep 'Partition Table'
```

If `msdos` is returned, it's an MBR partition and you can proceed to the next step.

If `gpt` is returned, it's a GPT partition.

2. Run the following command to check whether the OS starts in EFI mode.

```
sudo ls /sys/firmware/efi
```

If files exist, select UEFI as the boot mode when importing images. A discrepancy between the currently selected boot mode and the boot mode of the image file may affect the normal startup of your instance.

If no files exist, select Legacy BIOS as the boot mode when importing images. For details, refer to Best Practices for Boot Modes Legacy BIOS/UEFI.

Checking System-Critical Files

Check system-critical files, including but not limited to the following:

Note:

Follow the distribution standards to ensure that the paths and permissions of the system-critical files are correct and the files can be read and written normally.

`/etc/grub2.cfg` : It's recommended to use uuid in the `kernel` parameter for root mounting. Other methods (such as root= `/dev/sda`) may cause a system startup failure. The mounting steps are as follows:

1.1 Run the following command to get the file system name of `/root` .

```
df -TH
```

Obtain the file system name in the result as shown below. In this document, the file system name of the `/root` is `/dev/vda1` .

```
[root@VM-5-56-centos ~]# df -TH
Filesystem      Type       Size  Used Avail Use% Mounted on
devtmpfs        devtmpfs   938M     0  938M   0% /dev
tmpfs           tmpfs      953M   25k  953M   1% /dev/shm
tmpfs           tmpfs      953M  418k  953M   1% /run
tmpfs           tmpfs      953M     0  953M   0% /sys/fs/cgroup
/dev/vda1       ext4        22G  2.6G   18G  13% /
tmpfs           tmpfs      191M     0  191M   0% /run/user/0
```

1.2 Run the following command to get the UUID.

```
sudo blkid /dev/vda1
```

Note:

The file system UUID is not fixed. Confirm and update it regularly. For example, after the file system is formatted, its UUID will change.

1.3 Run the following command to use VI editor to open the `/etc/fstab` file.

```
vi /etc/fstab
```

1.4 Press **i** to enter edit mode.

1.5 Move the cursor to the end of the file, press **Enter**, and add the following content according to the example above:

```
UUID=d489ca1c-xxxx-4536-81cb-ceb2847f9954 / ext4 defaults    0    0
```

1.6 Press **ESC**, enter **:wq**, and press **Enter** to save the configuration and exit the editor.

`/etc/fstab` : Do not attach other disks here, which may cause the system startup failure after migration because the disk is not found.

`/etc/shadow` : Granted with the read-write permissions.

Uninstalling Software

Uninstall the conflicting drivers and software (including VMware tools, Xen tools, Virtualbox GuestAdditions, and other software that comes with underlying drivers).

Checking the Virtio Driver

For more information, see Checking Virtio Drivers in Linux.

Installing cloud-init

For more information, see Installing Cloud-Init on Linux.

Checking Other Hardware Configurations

After the migration to the cloud, hardware changes include but are not limited to:

The graphics card changes to Cirrus VGA.

The disk changes to Virtio Disk. The device name is vda or vdb.

The ENI changes to Virtio Nic. By default, only eth0 is available.

Querying Partitions and Their Sizes

Run the following command to query the current OS partition format and determine the partitions to be copied and their sizes.

```
mount
```

A result similar to the following is returned:

```
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
sys on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
dev on /dev type devtmpfs
(rw,nosuid,relatime,size=4080220k,nr_inodes=1020055,mode=755)
run on /run type tmpfs (rw,nosuid,nodev,relatime,mode=755)
/dev/sda1 on / type ext4 (rw,relatime,data=ordered)
securityfs on /sys/kernel/security type securityfs
(rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
devpts on /dev/pts type devpts
(rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
cgroup on /sys/fs/cgroup/unified type cgroup2
(rw,nosuid,nodev,noexec,relatime,nsdelegate)
cgroup on /sys/fs/cgroup/systemd type cgroup
(rw,nosuid,nodev,noexec,relatime,xattr,name=systemd)
pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup
(rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
cgroup on /sys/fs/cgroup/cpuset type cgroup
(rw,nosuid,nodev,noexec,relatime,cpuset)
cgroup on /sys/fs/cgroup/rdma type cgroup
(rw,nosuid,nodev,noexec,relatime,rdma)
cgroup on /sys/fs/cgroup/blkio type cgroup
(rw,nosuid,nodev,noexec,relatime,blkio)
```

```
cgroup on /sys/fs/cgroup/hugetlb type cgroup
(rw,nosuid,nodev,noexec,relatime,hugetlb)
cgroup on /sys/fs/cgroup/memory type cgroup
(rw,nosuid,nodev,noexec,relatime,memory)
cgroup on /sys/fs/cgroup/devices type cgroup
(rw,nosuid,nodev,noexec,relatime,devices)
cgroup on /sys/fs/cgroup/pids type cgroup
(rw,nosuid,nodev,noexec,relatime,pids)
cgroup on /sys/fs/cgroup/freezer type cgroup
(rw,nosuid,nodev,noexec,relatime,freezer)
cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup
(rw,nosuid,nodev,noexec,relatime,net_cls,net_prio)
cgroup on /sys/fs/cgroup/perf_event type cgroup
(rw,nosuid,nodev,noexec,relatime,perf_event)
systemd-1 on /home/libin/work_doc type autofs
(rw,relatime,fd=33,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=12692
)
systemd-1 on /proc/sys/fs/binfmt_misc type autofs
(rw,relatime,fd=39,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=12709
)
debugfs on /sys/kernel/debug type debugfs (rw,relatime)
mqueue on /dev/mqueue type mqueue (rw,relatime)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,pagesize=2M)
tmpfs on /tmp type tmpfs (rw,nosuid,nodev)
configfs on /sys/kernel/config type configfs (rw,relatime)
tmpfs on /run/user/1000 type tmpfs
(rw,nosuid,nodev,relatime,size=817176k,mode=700,uid=1000,gid=100)
gvfsd-fuse on /run/user/1000/gvfs type fuse.gvfsd-fuse
(rw,nosuid,nodev,relatime,user_id=1000,group_id=100)
```

According to the result, the root partition resides in `/dev/sda1` , no independent partitions reside in `/boot` or `/home` , sda1 contains the boot partition, and mbr is missing. Therefore, we only need to copy the entire sda.

Note:

The exported image should contain at least the root partition and mbr. If mbr is missing, the operating system cannot be started.

If `/boot` and `/home` are independent partitions in the current operating system, the exported image should also contain them.

Exporting an Image

Choose the appropriate image export method as needed.

Using a platform tool to export an image

Using commands to export an image

For details about how to use the image export tools of virtualization platforms, such as VMWare vCenter Convert and Citrix XenConvert, see the tool documentations on these platforms.

Note:

Tencent Cloud Service Migration supports images in qcow2, vhd, raw, and vmdk formats.

Note:

This method poses higher risks. For example, the file system's metadata may be corrupted when I/O is busy. We recommended that you check the image to make sure that the image is intact and correct after it is exported.

You can use either the qemu-img or dd command to export an image.

Use the `qemu-img` **command**

Run the following command to install the required package. This document uses Debian as an example. The package name may vary by distributions, such as `qemu-img` for CentOS.

```
apt-get install qemu-utils
```

Run the following command to export `/dev/sda` to `/mnt/sdb/test.qcow2` .

```
sudo qemu-img convert -f raw -O qcow2 /dev/sda /mnt/sdb/test.qcow2
```

In this command, `/mnt/sdb` indicates the mounted new disk or another network storage.

To convert its format, modify the value of the `-O` parameter to one of the following:

Parameter Value	Description
qcow2	qcow2 format
vhd	vhd format
vmdk	vmdk format
raw	No format

Using the `dd` **command**

For example, run the following command to export an image in raw format.

```
sudo dd if=/dev/sda of=/mnt/sdb/test.imag bs=1K count=$count
```

The `count` parameter specifies the number of partitions to be copied, which can be queried by running the `fdisk` command. To copy all partitions, ignore `count` .

For example, run the following command to view the number of partitions of `/dev/sda` .

```
fdisk -lu /dev/sda


Disk /dev/sda: 1495.0 GB, 1494996746240 bytes
255 heads, 63 sectors/track, 181756 cylinders, total 2919915520 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 4096 bytes
I/O size (minimum/optimal): 4096 bytes / 4096 bytes
Disk identifier: 0x0008f290
```

According to the result of the fdisk command, the sda1 ends at 41945087 * 512 bytes, so set count to 20481 M.

Note:

The image exported by using the `dd` command is in RAW format. We recommend that you convert it to qcow2, vhd, or other image formats.

Converting the Image Format (Optional)

Refer to **Image Format Conversion** and use `qemu-img` to convert the original image into a supported format.

Checking the Image

Note:

The image file system that you prepare may be corrupted because you prepared the image without stopping the service or due to other reasons. Therefore, we recommend that you check the image after preparing it.

If the image format is supported by the current platform, you can directly open and check the image file system. For example, the Windows platform supports VHD images, the Linux platform allows you to use `qemu-nbd` to open QCOW2 images, and the Xen platform allows you to directly open VHD files. This document uses the Linux platform as an example:

1. Run the following commands in sequence to check whether the nbd component exists.

```
modprobe nbd


lsmod | grep nbd
```

If a result similar to the following is returned, the nbd component exists. If nothing is returned, check whether the kernel compilation option `CONFIG_BLK_DEV_NBD` is enabled. If not, enable it or change the system before compiling the kernel again.

2. Run the following commands in sequence to check the image.

```
qemu-nbd -c /dev/nbd0 xxxx.qcow2


mount /dev/nbd0p1 /mnt
```

After you run the `qemu-nbd` command, `/dev/nbd0` maps to `xxx.qcow2`, and `/dev/nbd0p1` indicates the first partition of the virtual disk. If nbd0p1 does not exist or mount fails, the image may be incorrect. You can also start the CVM to check whether the image file works before uploading the image.

Checking Virtio Drivers in Linux

Last updated：2024-01-08 09:37:01

Overview

To run in Tencent Cloud, a CVM instance must have a kernel supporting virtio drivers, including the block device driver `virtio_blk` and the ENI driver `virtio_net` . To ensure that a CVM instance created with a custom image can start up properly, check whether your image supports virtio drivers in the source server before importing the image. This document uses CentOS as an example to describe how to check whether an image supports virtio drivers.

Directions

Step 1. Check whether the kernel supports virtio drivers

Execute the following command to check whether the current kernel supports virtio drivers:

```
grep -i virtio /boot/config-$(uname -r)
```

A result similar to the following is returned:

If the value of `CONFIG_VIRTIO_BLK` and `CONFIG_VIRTIO_NET` is `m` in the response, please go to Step 2. If the value of `CONFIG_VIRTIO_BLK` and `CONFIG_VIRTIO_NET` is `y` in the response, which means the operating system contains the virtio drivers, you can import the custom image to Tencent Cloud. For detailed directions, see Overview.

If you cannot find `CONFIG_VIRTIO_BLK` and `CONFIG_VIRTIO_NET` in the response, it means that images with the OS **cannot** be imported to Tencent Cloud. Please download and compile kernel.

Step 2. Check whether the temporary file system contains virtio drivers

If the value of the parameters is `m` in Step 1, you need to check whether `initramfs` or `initrd` contains the `virtio` drivers. Please execute the corresponding command according to the operating system:

CentOS Stream Operating System:

```
lsinitrd /boot/initramfs-$(uname -r).img | grep virtio
```

For CentOS 6/CentOS 7/CentOS 8/Red Hat 6/Red Hat 7:

```
lsinitrd /boot/initramfs-$(uname -r).img | grep virtio
```

For RedHat 5/CentOS 5:

```
mkdir -p /tmp/initrd && cd /tmp/initrd
zcat /boot/initrd-$(uname -r).img | cpio -idmv
find . -name "virtio*"
```

For Debian/Ubuntu:

```
lsinitramfs /boot/initrd.img-$(uname -r) | grep virtio
```

OpenSUSE Leap Operating System:

```
lsinitrd /boot/initrd-$(uname -r) | grep virtio
```

A result similar to the following is returned:

It means that initramfs contains the virtio_blk driver and virtio.ko, virtio_pci.ko, and virtio_ring.ko on which the driver depends. In this case, you can import the custom image to Tencent Cloud. For details, see Import Images > Overview. If initramfs or initrd does not contain the virtio drivers, please go to Step 3.

Step 3. Reconfigure the temporary file system

If you find that `initramfs` or `initrd` does not contain the `virtio` drivers in [Step 2](#), you will need to reconfigure the temporary file system to ensure that `initramfs` or `initrd` contains the `virtio` drivers. Run the corresponding command according to the operating system:

CentOS Stream Operating System:

```
mkinitrd -f --allow-missing --with=virtio_blk --preload=virtio_blk --
with=virtio_net --preload=virtio_net --with=virtio_console --
preload=virtio_console /boot/initramfs-$(uname -r).img $(uname -r)
```

For CentOS 8/Red Hat 8:

```
mkinitrd -f --allow-missing --with=virtio_blk --preload=virtio_blk --
with=virtio_net --preload=virtio_net --with=virtio_console --
preload=virtio_console /boot/initramfs-$(uname -r).img $(uname -r)
```

For CentOS 6/CentOS 7/RedHat 6/RedHat 7:

```
mkinitrd -f --allow-missing --with=xen-blkfront --preload=xen-blkfront --
with=virtio_blk --preload=virtio_blk --with=virtio_pci --preload=virtio_pci --
with=virtio_console --preload=virtio_console /boot/initramfs-$(uname -r).img
$(uname -r)
```

For RedHat 5/CentOS 5:

```
mkinitrd -f --allow-missing --with=xen-vbd  --preload=xen-vbd --with=xen-
platform-pci --preload=xen-platform-pci --with=virtio_blk --preload=virtio_blk
--with=virtio_pci --preload=virtio_pci --with=virtio_console --
preload=virtio_console /boot/initrd-$(uname -r).img $(uname -r)
```

For Debian/Ubuntu:

```
echo -e 'xen-blkfront\\nvirtio_blk\\nvirtio_pci\\nvirtio_console' >>
/etc/initramfs-tools/modules
mkinitramfs -o /boot/initrd.img-$(uname -r)
```

OpenSUSE Leap Operating System:

```
mkinitrd -m "virtio_blk virtio_net"
```

Appendix

Downloading and compiling the kernel

Downloading the kernel installation package

1. Execute the following command to install the components necessary for kernel compilation.

```
yum install -y ncurses-devel gcc make wget
```

2. Execute the following command to view the current version of the kernel.

```
uname -r
```

A response similar to the following will be returned, indicating the current kernel version is 2.6.32-642.6.2.el6.x86_64.

```
[root@VM_0_139_centos ~]# uname -r
2.6.32-642.6.2.el6.x86_64
```

3. Download the source code of the corresponding or closest kernel version here.

For example, for the `2.6.32-642.6.2.el6.x86_64` version, you should download `linux-2.6.32.tar.gz` at `https://mirrors.edge.kernel.org/pub/linux/kernel/v2.6/linux-2.6.32.tar.gz`.

4. Execute the following command to switch directory.

```
cd /usr/src/
```

5. Execute the following command to download the installation package.

```
wget https://mirrors.edge.kernel.org/pub/linux/kernel/v2.6/linux-2.6.32.tar.gz
```

6. Execute the following command to decompress the installation package.

```
tar -xzf linux-2.6.32.tar.gz
```

7. Execute the following command to make connection.

```
ln -s linux-2.6.32 linux
```

8. Execute the following command to switch directory.

```
cd /usr/src/linux
```

Compiling the kernel

1. Execute the following commands to compile the kernel.

```
make mrproper
cp /boot/config-$(uname -r) ./.config
make menuconfig
```

Enter the "Linux Kernel vX.X.XX Configuration" interface as shown below:

Note:

If you are not taken to the "Linux Kernel vX.X.XX Configuration" interface, perform Step 18.

"Linux Kernel vX.X.XX Configuration" interface:

Press "Tab" or the "↑"/"↓" key to move the cursor.

Press "Enter" to select or execute the item selected by the cursor.

Press the space bar to select the item selected by the cursor. "*" means compiling to the kernel, and "M" means compiling to a module.

2. Press the "↓" key to move the cursor to "Virtualization" and press the space bar to select "Virtualization".

3. Press "Enter" to enter the Virtualization details interface.

4. In the Virtualization details interface, check whether the Kernel-based Virtual Machine (KVM) support option is selected as shown below:

```
.config - Linux Kernel v2.6.32 Configuration
                              ──── Virtualization ────
  Arrow keys navigate the menu.  <Enter> selects submenus --->.  Highlighted letters are hotkeys.  Pressing <Y>
  includes, <N> excludes, <M> modularizes features.  Press <Esc><Esc> to exit, <?> for Help, </> for Search.
  Legend: [*] built-in  [ ] excluded  <M> module  < > module capable
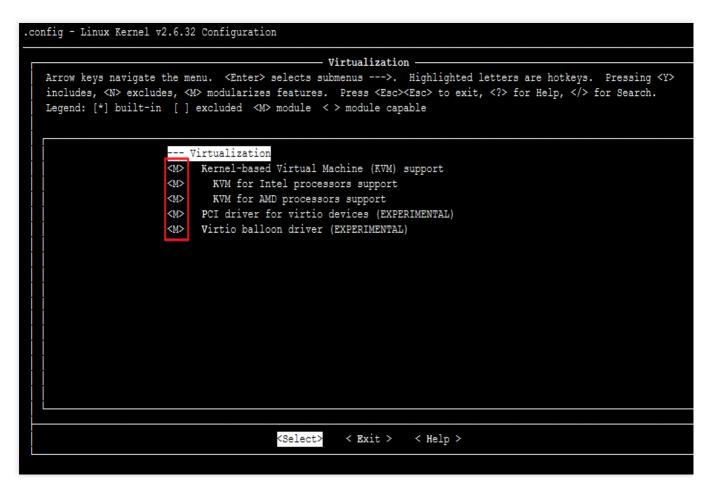
                       --- Virtualization
              <M>     Kernel-based Virtual Machine (KVM) support
              <M>       KVM for Intel processors support
              <M>       KVM for AMD processors support
              <M>     PCI driver for virtio devices (EXPERIMENTAL)
              <M>     Virtio balloon driver (EXPERIMENTAL)




                        <Select>    < Exit >    < Help >
```

If it is not selected, press the space bar to select the "Kernel-based Virtual Machine (KVM) support" option.

5. Press "Esc" to return to the "Linux Kernel vX.X.XX Configuration" main interface.

6. Press the "↓" key to move the cursor to "Processor type and features" and press "Enter" to enter the Processor type and features details interface.

7. Press the "↓" key to move the cursor to "Paravirtualized guest support" and press "Enter" to enter the detailed interface of Paravirtualized guest support.

8. In the Paravirtualized guest support details interface, check whether "KVM paravirtualized clock" and "KVM Guest support" are selected as shown below:

If they are not selected, press the space bar to select the "KVM paravirtualized clock" and "KVM Guest support" options.

9. Press "Esc" to return to the "Linux Kernel vX.X.XX Configuration" main interface.

10. Press the "↓" key to move the cursor to "Device Drivers" and press "Enter" to enter the Device Drivers details interface.

11. Press the "↓" key to move the cursor to "Block devices" and press "Enter" to enter the Block devices details interface.

12. In the Block devices details interface, check whether "Virtio block driver (EXPERIMENTAL)" is selected as shown below:

If it is not selected, press the space bar to select the "Virtio block driver (EXPERIMENTAL)" option.

13. Press "Esc" to return to the Device Drivers details interface.

14. Press the "↓" key to move the cursor to "Network device support" and press "Enter" to enter the Network device support details interface.

15. In the Network device support details interface, check whether "Virtio network driver (EXPERIMENTAL)" is selected as shown below:

If it is not selected, press the space bar to select the "Virtio network driver (EXPERIMENTAL)" option.

16. Press "Esc" to exit the kernel configuration interface, and select "YES" to save the `.config` file.

17. Take Step 1: Checking whether the kernel supports the virtio drivers to verify whether the virtio drivers have been configured correctly.

18. (Optional) Run the following command to manually edit the `.config` file.

Note:

 This step is recommended if either of the following is true:

The kernel still contains no configuration information of the virtio drivers.

When compiling the kernel, you cannot enter the kernel configuration interface or save the `.config` file.

```
make oldconfig
make prepare
make scripts
make
make install
```

19. Execute the following commands to check the installation of the virtio drivers.

```
find /lib/modules/"$(uname -r)"/ -name "virtio.*" | grep -E "virtio.*"
grep -E "virtio.*" < /lib/modules/"$(uname -r)"/modules.builtin
```

If any of the commands returns a list of files such as `virtio_blk`, `virtio_pci.virtio_console`, it indicates that you have installed the virtio drivers correctly.

Installing Cloud-Init on Linux

Last updated：2024-01-08 09:37:01

Overview

This document describes how to install the cloud-init service. Cloud-init allows you to customize configurations during the first initialization of an instance.

Options to install cloud-init:

Download the cloud-init binary package

Download the cloud-init source package

Use the cloud-init package from the software source

Prerequisites

Connect the server that you want to install cloud-init to the public network.

How It Works

Download the cloud-init binary package

Manual download

Using software source

Note:

cloud-init depends on qcloud-python, which is a software package recompiled by Tencent Cloud. qcloud-python is a separate python environment and is only used for cloud-init. It is installed under the directory of `/usr/local/qcloud/python` , and it does not conflict with the default python in the system.

cloud-init is developed by Tencent Cloud based on the community v20.1. It is adapted to Tencent Cloud operation environment.

The cloud-init binary package supports the following operating systems:

Type	OS	Version	x86_64		arm64
			qcloud-python	cloud-init	qcloud
rpm	CentOS	7	qcloud-python-3.7.10-1.el7.x86_64.rpm	cloud-init-20.1.0011-1.el7.x86_64.rpm	qcloud 1.el7.c
		8	qcloud-python-3.7.10-	cloud-init-20.1.0011-	qcloud

			1.el8.x86_64.rpm	1.el8.x86_64.rpm	1.el8.a
	Fedora	36	qcloud-python-3.7.10-2.fc36.x86_64.rpm	cloud-init_20.1.0011-1_arm64.deb	N/A
	Kylin	20sp1	qcloud-python-3.7.10-1.ky10.x86_64.rpm	cloud-init-20.1.0011-2.ky10.x86_64.rpm	qcloud 1.ky1(
	openSUSE	15.4	qcloud-python-3.7.10-2.x86_64.rpm	cloud-init-20.1.0011-2.x86_64.rpm	N/A
deb	Debian	11	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	qcloud 1_arm
		10	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	N/A
		9	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	N/A
		8	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	N/A
	Ubuntu	22.04	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	N/A
		20.04	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	qcloud 1_arm
		18.04	qcloud-python_3.7.10-1%2Bubuntu18.04_amd64.deb	cloud-init_20.1.0011-1%2Bubuntu18.04_amd64.deb	qcloud 1_arm
		16.04	qcloud-python_3.7.10-1_amd64.deb	cloud-init_20.1.0011-1_amd64.deb	N/A

Downloading cloud-init binary package

1. Download the installation package.

2. If cloud-init already exists, run the following command to clear it.

```
rm -rf /var/lib/cloud
rm -rf /etc/cloud
rm -rf /usr/local/bin/cloud*
```

3. Run the following commands based on the OS.

For deb type, run the following command.

```
dpkg -i *.deb
```

For rpm type, run the following command.

```
rpm -ivh *.rpm
```

4. Check whether the version is installed properly.

```
cloud-init qcloud -v
/usr/bin/cloud-init qcloud 0011
```

5. Restart.

Downloading the cloud-init source package

Note:

The cloud-init-20.1.0011 version is most compatible with Tencent Cloud. It ensures that all configuration items of CVMs created through the image can be initialized properly. We recommend that you install **cloud-init-20.1.0011.tar.gz**. You can also click here to download other versions. This document uses cloud-init-20.1.0011 as an example.

Run the following command to download the cloud-init source package:

```
wget https://gerryguan-1306210569.cos.ap-chongqing.myqcloud.com/cloud-init/src/cloud-init-20.1.0011.tar.gz
```

Installing cloud-init

1. Run the following command to decompress the cloud-init installation package.

Note:

If you are using the Ubuntu operating system, run this command with the "root" account.

```
tar -zxvf cloud-init-20.1.0011.tar.gz
```

2. Run the following command to enter the decompressed cloud-init installation package directory, that is, the cloud-init-20.1.0011 directory:

```
cd cloud-init
```

3. Install Python-pip according to the operating system version.

For CentOS 6/7, run the following command:

```
yum install python3-pip -y
```

For Ubuntu, run the following command:

```
apt-get -y install python3-pip
```

During installation, if an error such as "failed to install" or "installation package not found" occurs, see resolving Python-pip installation failure to troubleshoot it.

4. Run the following command to upgrade pip.

```
python3 -m pip install --upgrade pip
```

5. Run the following command to install dependencies.

Note:

Python 2.6 is not supported when cloud-init uses requests 2.20.0 or later. If the Python interpreter installed in the image environment is Python version 2.6 or earlier, run the `pip install 'requests<2.20.0'` command to install requests 2.20.0 or later before installing the cloud-init dependencies.

```
pip3 install -r requirements.txt
```

6. Install the cloud-utils components corresponding to your OS version.

For CentOS 6, run the following command:

```
yum install cloud-utils-growpart dracut-modules-growroot -y
dracut -f
```

For CentOS 7, run the following command:

```
yum install cloud-utils-growpart -y
```

For Ubuntu, run the following command:

```
apt-get install cloud-guest-utils -y
```

7. Run the following command to install cloud-init:

```
python3 setup.py build
```

```
python3 setup.py install --init-system systemd
```

Note:

The `--init-system` can be followed by any of `systemd` , `sysvinit` , `sysvinit_deb` , `sysvinit_freebsd` , `sysvinit_openrc` , `sysvinit_suse` , `upstart` , or `None` (default). Choose one according to the auto-start service management method of the operating system. Otherwise the cloud-init service cannot automatically start upon system startup.

Select `sysvinit` for the CentOS 6 and earlier versions, and select `systemd` for CentOS 7 and later versions. This document uses systemd as an example.

Modifying the cloud-init configuration file

1. Download cloud.cfg for your operating system.

Download cloud.cfg for Ubuntu.

Download cloud.cfg for CentOS.

2. Replace the content of `/etc/cloud/cloud.cfg` with that of the downloaded cloud.cfg file.

Adding syslog user

Run the following command to add a syslog user:

```
useradd syslog
```

Configuring the auto-start of the cloud-init service on boot

If your operating system uses the systemd auto-start service management method, run the following command.

Note:

To check whether the operating system uses systemd, run the `strings /sbin/init | grep "/lib/system"` command, and you will receive a return message.

Run the following command in Ubuntu or Debian.

```
ln -s /usr/local/bin/cloud-init /usr/bin/cloud-init
```

Run the following commands in all operating systems.

```
systemctl enable cloud-init-local.service
systemctl start cloud-init-local.service
systemctl enable cloud-init.service
systemctl start cloud-init.service
systemctl enable cloud-config.service
systemctl start cloud-config.service
systemctl enable cloud-final.service
systemctl start cloud-final.service
systemctl status cloud-init-local.service
systemctl status cloud-init.service
systemctl status cloud-config.service
systemctl status cloud-final.service
```

Run the following commands in CentOS or Redhat.

Replace the content of `/lib/systemd/system/cloud-init-local.service` with the following:

```
[Unit]
Description=Initial cloud-init job (pre-networking)
Wants=network-pre.target
```

```
After=systemd-remount-fs.service
Before=NetworkManager.service
Before=network-pre.target
Before=shutdown.target
Conflicts=shutdown.target
RequiresMountsFor=/var/lib/cloud
[Service]
Type=oneshot
ExecStart=/usr/bin/cloud-init init --local
ExecStart=/bin/touch /run/cloud-init/network-config-ready
RemainAfterExit=yes
TimeoutSec=0
# Output needs to appear in instance console output
StandardOutput=journal+console
[Install]
WantedBy=cloud-init.target
```

Replace the content of `/lib/systemd/system/cloud-init.service` with the following:

```
[Unit]
Description=Initial cloud-init job (metadata service crawler)
Wants=cloud-init-local.service
Wants=sshd-keygen.service
Wants=sshd.service
After=cloud-init-local.service
After=systemd-networkd-wait-online.service
After=networking.service
After=systemd-hostnamed.service
Before=network-online.target
Before=sshd-keygen.service
Before=sshd.service
Before=systemd-user-sessions.service
Conflicts=shutdown.target
[Service]
Type=oneshot
ExecStart=/usr/bin/cloud-init init
RemainAfterExit=yes
TimeoutSec=0
# Output needs to appear in instance console output
StandardOutput=journal+console
[Install]
WantedBy=cloud-init.target
```

If your operating system uses the sysvinit auto-start service management method, run the following commands:

Note:

To check whether the operating system uses sysvinit, run the `strings /sbin/init | grep "sysvinit"` command, and you will receive a return message.

```
chkconfig --add cloud-init-local
chkconfig --add cloud-init
chkconfig --add cloud-config
chkconfig --add cloud-final
chkconfig cloud-init-local on
chkconfig cloud-init on
chkconfig cloud-config on
chkconfig cloud-final on
```

Installing cloud-init

Run the following command to install cloud-init:

```
apt-get/yum install cloud-init
```

Note:

By default, the cloud-init version installed by running `apt-get` or `yum` is the default cloud-init version in the software source configured for the operating system. Some configuration items of instances created by using the image whose cloud-init is installed this way may not be initialized as expected. Therefore, we recommend that you install the service by manually downloading the cloud-init source package.

Modifying the cloud-init configuration file

1. Download cloud.cfg for your operating system.

Download cloud.cfg for Ubuntu.

Download cloud.cfg for CentOS.

2. Replace the content of `/etc/cloud/cloud.cfg` with that of the downloaded cloud.cfg file.

More

Note:

Do not restart the server after performing the following operations. Otherwise, you will need to perform them again.

1. Run the following command to check whether the cloud-init configuration is successful.

```
cloud-init init --local
```

If the following information is returned, it indicates that the cloud-init has been successfully configured.

```
Cloud-init v. 20.1.0011 running 'init-local' at Fri, 01 Apr 2022 01:26:11
+0000. Up 38.70 seconds.
```

2. Run the following command to delete the cache records of cloud-init.

```
rm -rf /var/lib/cloud
```

3. Run the following command in Ubuntu or Debian.

```
rm -rf /etc/network/interfaces.d/50-cloud-init.cfg
```

4. For Ubuntu or Debian, modify the content of `/etc/network/interfaces` to the following:

```
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).
source /etc/network/interfaces.d/*
```

Appendix

Resolving Python-pip installation failure

During installation, if an error such as "failed to install" or "installation package not found" occurs, troubleshoot it based on the operating system as follows:

CentOS 6/7:

Ubuntu:

1. Run the following command to configure the EPEL storage repository.

```
yum install epel-release -y
```

2. Run the following command to install Python-pip.

```
yum install python3-pip -y
```

1. Run the following command to clear the cache.

```
apt-get clean all
```

2. Run the following command to update the software package list.

```
apt-get update -y
```

3. Run the following command to install Python-pip.

```
apt-get -y install python3-pip
```

Creating Windows Images

Last updated：2024-10-22 12:01:24

Overview

This document uses Windows Server 2012 as an example to guide you through the process of creating a Windows image. If you are using a different version of the Windows Server operating system, you can still refer to this document for image creation.

Prerequisites

Before performing the operation, ensure that you have prepared a virtual machine and have installed the Windows operating system (Windows Server 2012 is used as an example in this document).

Preparations

Checking the Partitioning and Starting Mode of the Operating System

Note:

If you need to create and export a data disk image, skip this step.

Checking the Partitioning and Starting Mode of the Operating System

1. On the desktop, click

to open the **Windows PowerShell** window.

2. In the **Windows PowerShell** window, enter **diskmgmt.msc** and click **Enter** to open the **Disk Management** window.

3. Right-click the disk you need to check and then choose **Properties**.

4. Select the **Volume** tab to view the disk partition.

5. Check whether the disk partition is a GPT partition.

If yes, as GPT partitions are not available for service migration, please submit a ticket.

If no, proceed to the next step.

6. Start CMD as the admin user and run the following command to check whether the operating system starts in EFI mode:

```
bcdedit /enum {current}
```

A result similar to the following will be returned:

```
Windows boot loader
ID                     {current}
device                 partition=C:
path                   \\WINDOWS\\system32\\winload.exe
description            Windows 10
locale                 zh-CN
inherit                {bootloadersettings}
recoverysequence       {f9dbeba1-1935-11e8-88dd-ff37cca2625c}
displaymessageoverride Recovery
recoveryenabled        Yes
flightsigning          Yes
allowedinmemorysettings 0x15000075
osdevice               partition=C:
systemroot             \\WINDOWS
resumeobject           {1bcd0c6f-1935-11e8-8d3e-3464a915af28}
nx                     OptIn
bootmenupolicy         Standard
```

If the **path** parameter contains "efi", the current operating system is started in EFI mode. In this case, please submit a ticket.

If **path** does not contain "efi", proceed to the next step.

Uninstalling Conflicting Drivers and Software

Uninstall the conflicting drivers and software (including VMware tools, Xen tools, Virtualbox GuestAdditions, and other software that comes with underlying drivers).

Directions

Step 1: Installing cloud-base

Install cloud-base as instructed in Installing Cloudbase-Init on Windows.

Step 2: Verifying or Installing VirtIO Drivers

Choose **Control panel** > **Programs and Features** and enter "Virtio" in the search box.

If the result shown in the figure below is returned, the Virtio driver is installed.

If the Virtio driver is not installed, please download an edition based on your needs and install it.

Note:

Tencent Cloud does not support importing Windows Server 2003.

If you are using Windows Server 2008R2/2012R2/2016/2019/2022, please install Tencent Cloud customized VirtIO driver.

If you are using another version of Windows operating system, please try Tencent Cloud customized VirtIO driver first. If the running is not stable, install the community edition instead.

Installing Tencent Cloud customized VirtIO driver (recommended)

Installing community edition of VirtIO driver

The download addresses are as below:

Public network download address:

```
http://mirrors.tencent.com/install/windows/virtio_64_1.0.9.exe
```

Private network download address:

```
http://mirrors.tencentyun.com/install/windows/virtio_64_1.0.9.exe
```

Try Tencent Cloud customized VirtIO driver first. If the running is not stable, install the community edition instead.

Download community edition of VirtIO driver

Checking Other Hardware Configurations

After the migration to the cloud, hardware changes include but are not limited to:

The graphics card has been changed to Cirrus VGA. (**Path: Computer Management > Device Manager > Display adapters**)

The disk has been changed to VirtIO disk. (**Path: Computer Management > Device Manager > Disk Drives**)

The Elastic Network Interface (ENI) has been changed to VirtIO NIC and local area connection is used by default.
(Path: Computer Management > Device Manager > Network adapter)

Step 3: Exporting the Image File Using the Disk2vhd Tool

1. Click here to download Disk2vhd.

2. Install and run Disk2vhd.

Note:

Install and run Disk2vhd on a non-system disk.

Disk2vhd can be started only after the Volume Shadow Copy Service (VSS) is installed in the Windows system. For more information about the VSS features, see Volume Shadow Copy Service.

3. Configure the parameters as below and click **Create** to export the image.

Use Vhdx: Do not select it because the system currently does not support VHDX images.

Use volume Shadow Copy: It is recommended that you select it for higher data integrity.

VHD File name: The location where the .vhd file is stored. Please select a non-system disk.

Volume to include: The entire system disk is required to be exported when you export the image. **Please choose all partitions of your system disk**, otherwise an error will occur when you import the image.

System disk partitions usually include C:\\ partition, boot partition and recovery partition. All partitions need to be chosen.

Configuration samples

Run Disk2vhd in E drive, choose all partitions of the system disk (boot partition and C:\\ partition). Select "Use volume Shadow Copy" and deselect "Use Vhdx". The .vhd file will be stored on E drive after the image is exported.

Step 4: Checking the Image

Note:

The image file system that you create may be corrupted because you created the image without stopping the service or due to other reasons. Therefore, we recommend that you check the image after creating it.

If the image format is supported by the current platform, you can directly open the image to check the file system. For example, the Windows platform supports images in the vhd format; the Linux platform allows you to use qemu-nbd to

open images in the qcow2 format; and the Xen platform allows you to directly open files in the vhd format.

This document takes checking the VHD images through **Attach VHD** in **Disk Management** on Windows as an example.

1. On the desktop, right click

and select **Computer Management** in the pop-up menu.

2. Select **Storage** > **Disk Management** to enter the disk management page.

3. Select **Action** > **Attach VHD** as shown in the figure below.

If the result similar to the following figure appears, the image has been created.

FAQs

How Do I Convert Image Formats?

Note:

Currently, Tencent Cloud Virtual Machine (CVM) migration supports images in the following formats: qcow2, vhd, raw, and vmdk.

If the image exported using other tools does not match the above formats, refer to Converting Image Format and use `qemu-img` to convert the image file format to a supported format.

Installing Cloudbase-Init on Windows

Last updated：2024-01-08 09:37:01

Overview

This document describes how to install Cloudbase-Init on the Windows Server 2012 R2 64-bit operating system.

Required Software

The following table describes the software required for installing Cloudbase-Init.

Software	Download Link	Description
CloudbaseInitSetup_X_X_XX_xXX.msi	Download the Cloudbase-Init installation package based on the operating system used. Stable version (recommended) Windows 64-bit operating system: Click here to download the installation package. Windows 32-bit operating system: Click here to download the installation package. Beta version For details, see the Cloudbase-Init official website.	Used to install Cloudbase-Init
TencentCloudRun.ps1	Click here to download the installation package.	-
localscripts.py	Click here to download the installation package.	Used to ensure that Cloudbase-Init starts properly

Directions

Installing Cloudbase-Init

1. On the desktop, double-click the Cloudbase-Init installation package.

2. In the dialog box, click **Run** to enter the Cloudbase-Init setup wizard, as shown below:

3. Click **Next**.

4. Check "I accept the terms in the License Agreement" and click **Next** for the following two operations.

5. On the **Configuration options** page, set **Serial port for logging** to **COM1**, select **Run Cloudbase-Init service as LocalSystem** and click **Next**, as shown below:

6. Click **Install**.

7. When the installation is completed, click **Finish** to close the Cloudbase-Init setup wizard, as shown below:

Note:

When closing the Cloudbase-Init setup wizard, do not check any checkbox or run Sysprep.

Modifying the Cloudbase-Init configuration file

1. Open the `cloudbase-init.conf` configuration file.

The `cloudbase-init.conf` configuration file is saved in `C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\conf` by default.

2. Replace content in the `cloudbase-init.conf` configuration file with the following:

```
[DEFAULT]
username=Administrator
groups=Administrators
inject_user_password=true
config_drive_raw_hhd=true
config_drive_cdrom=true
config_drive_vfat=true
bsdtar_path=C:\\Program Files\\Cloudbase Solutions\\Cloudbase-
Init\\bin\\bsdtar.exe
mtools_path=C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\bin\\
san_policy=OnlineAll
metadata_services=cloudbaseinit.metadata.services.configdrive.ConfigDriveServic
e,cloudbaseinit.metadata.services.ec2service.EC2Service
#,cloudbaseinit.metadata.services.httpservice.HttpService
#,cloudbaseinit.metadata.services.maasservice.MaaSHttpService
metadata_base_url=http://169.254.0.23/
ec2_metadata_base_url=http://169.254.0.23/
```

```
retry_count=2
retry_count_interval=5
plugins=cloudbaseinit.plugins.windows.extendvolumes.ExtendVolumesPlugin,cloudba
seinit.plugins.common.networkconfig.NetworkConfigPlugin,cloudbaseinit.plugins.c
ommon.sethostname.SetHostNamePlugin,cloudbaseinit.plugins.common.setuserpasswor
d.SetUserPasswordPlugin,cloudbaseinit.plugins.common.localscripts.LocalScriptsP
lugin,cloudbaseinit.plugins.common.userdata.UserDataPlugin
verbose=true
debug=true
logdir=C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\log\\
logfile=cloudbase-init.log
default_log_levels=comtypes=INFO,suds=INFO,iso8601=WARN,requests=WARN
#logging_serial_port_settings=COM1,115200,N,8
mtu_use_dhcp_config=true
ntp_use_dhcp_config=true
first_logon_behaviour=no
netbios_host_name_compatibility=false
allow_reboot=true
activate_windows=true
kms_host="kms.tencentyun.com"
local_scripts_path=C:\\Program Files\\Cloudbase Solutions\\Cloudbase-
Init\\LocalScripts\\
C:\\powershell
PS C:\\Set-ExecutionPolicy Unrestricted
volumes_to_extend=1,2
```

3. Copy the `TencentCloudRun.ps1` script to `C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\LocalScripts`.

4. Right-click the `TencentCloudRun.ps1` script, select **Properties**, and check for its executable permission in the pop-up window, as shown below:

Check **Unblock** and click **OK**.

Skip this step if the **Unblock** option does not exist.

5. Replace `localscripts.py` in `C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\Python\\Lib\\site-packages\\cloudbaseinit\\plugins\\common` with the `localscripts.py` file in Required Software.

Converting Image Format

Last updated：2024-01-08 09:37:01

Overview

This document describes how to use qemu-img to convert image files to VHD or RAW format. Currently, you can import image files in RAW, VHD, QCOW2, or VMDK to Tencent Cloud CVM. Image files in other formats need to be converted before being imported.

Directions

Select the method according to the operating system of the CVM instance:

Windows

Linux

Note:

This document uses Windows 10 as an example to describe how to convert the image format. As the steps may vary by operating system, proceed based on the actual conditions.

Installing qemu-img

Download qemu-img here and install it. This document takes `C:\\Program Files\\qemu` as the installation path as an example.

Configuring environment variable

1. Right-Click **Start** and select **System** in the pop-up menu.
2. In the pop-up window, select **Advanced system settings**.
3. In the **System Properties** pop-up window, select the **Advanced** tab and click **Environment Variables**.
4. In the **Environment Variables** window, select `Path` in **System variables** and click **Edit** as shown below:

5. In the **Edit environment variable** pop-up window, click **Create**, enter the installation path of qemu-img
`C:\\Program Files\\qemu` , and click **OK**.

6. In the **Environment Variables** window, click **OK** again.

Verifying environment variable configuration

1. Press **Win + R** to open the **Run** window.

2. In the **Run** window, enter **cmd** to open the command line.

3. Run the following command to determine whether the environment variable has been configured successfully based on the returned result:

```
qemu-img --help
```

Converting image format

1. Run the following command on the command line to switch to the directory of the image file:

```
cd <directory of the source image file>
```

2. Run the following command to convert the image format:

```
qemu-img convert -f <source image file format> -O <target image format> <source
image filename> <target image filename>
```

The parameters are described as follows:

`-f` : source image file format.

`-O` (in uppercase): target image format and source and target image filenames.

For example, run the following command to convert the `test.qcow2` image file to `test.raw` :

```
qemu-img convert -f qcow2 -O raw test.qcow2 test.raw
```

After conversion, the target file will be displayed in the directory of the source image file.
Note:

This document uses Ubuntu 20.04 and CentOS 7.8 as an example to describe how to convert the image format. As the steps may vary by operating system, proceed based on the actual conditions.

Installing qemu-img

1. Run the following command to install qemu-img:
Ubuntu:

```
apt-get update # Update the package list


apt-get install qemu-utils # Install qemu-img
```

CentOS:

```
yum install qemu-img
```

2. Run the following command to convert the image format:

```
qemu-img convert -f qcow2 -O raw test.qcow2 test.raw
```

The parameters are described as follows:

`-f` : source image file format.

`-O` (in uppercase): target image format and source and target image filenames.

After conversion, the target file will be displayed in the directory of the source image file.

References

Overview

Preparing a Windows Image

Preparing a Linux Image

Setting the GRUB File Disk Identification Method to UUID

Last updated：2024-01-08 09:37:00

Operation Scenario

To ensure the Linux system can correctly identify the disk when launching the file system, please inspect and correctly set the GRUB file disk identification method.

The GRand Unified Bootloader (GRUB) serves as a bootloader for initiating the operating system. GRUB permits the utilization of device names (for instance, `/dev/vda1` , `/dev/vdb1` and so forth) to identify disk partitions. However, these device names may change due to the change in the actual operating environment after an image is imported. To guarantee the correct booting of the system even when the device name changes, you can modify the disk identification method in the GRUB file to the Universally Unique Identifier (UUID).

Setting the GRUB File Disk Identification Method to UUID

Confirming the GRUB File Path

There are two common versions of GRUB: GRUB (GRUB Legacy) and GRUB2. The configuration files for GRUB and GRUB2 are located in different paths.

For GRUB, the configuration file is typically located in `/boot/grub/menu.lst` or `/boot/grub/grub.conf` .
For GRUB2, the configuration file is commonly located in `/boot/grub/grub.cfg` or `/boot/grub2/grub.cfg` .
If you find the `menu.lst` or `grub.conf` file in the `/boot/grub` directory, you are probably using GRUB (GRUB Legacy). If you find the `grub.cfg` file in the `/boot/grub` or `/boot/grub2` directory, you are probably using GRUB2.

Obtaining the UUID

To obtain the UUID of a partition, the `blkid` command can be used. Running the `blkid` command will display the detailed information of all the available partitions, including the UUIDs. Run the following command in the terminal:

```
sudo blkid
```

The output similar to the following one indicates that the associated UUID of the device `/dev/vda1` is `c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b` .

```
/dev/vda1: UUID="c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b" BLOCK_SIZE="4096"
TYPE="ext4" PARTUUID="bcfcb5cb-01"
```

Modifying the GRUB configuration file

This segment highlights an example where the modification of the GRUB2 configuration file located in the `/boot/grub/grub.cfg` directory is made. If you are using GRUB, or if the GRUB2 configuration file for distribution is located in the `/boot/grub2/grub.cfg` directory, you can adjust the configuration according to the actual situation.

1. Back up the current `/boot/grub/grub.cfg` file to the `/home` directory.

```
sudo cp /boot/grub/grub.cfg /home
```

2. Use the **vi** editor to open the `/boot/grub/grub.cfg` file and confirm the root partition marked in the configuration file. In this case, the root partition is located on the `/dev/vda1` device.

```
sudo vi /boot/grub/grub.cfg


# /boot/grub/grub.cfg
...
echo    'Loading Linux 6.1.0-13-amd64 ...'
linux   /boot/vmlinuz-6.1.0-13-amd64 root=/dev/vda1 ro
echo    'Loading initial ramdisk ...'
...
```

3. Edit the configuration starting with a device name in the `grub.cfg` file, and change the `root=/dev/vda1` device name to the `root=UUID=xxx` format. The content after `root=UUID=` is the UUID value corresponding to the device returned by running the `blkid` command. This configuration may appear for multiple times in the `grub.cfg` file. The modifcation is required for each configuration.

```
# Before modification
...
echo    'Loading Linux 6.1.0-13-amd64 ...'
linux   /boot/vmlinuz-6.1.0-13-amd64 root=/dev/vda1 ro
echo    'Loading initial ramdisk ...'
...
# After modification
...
echo    'Loading Linux 6.1.0-13-amd64 ...'
linux   /boot/vmlinuz-6.1.0-13-amd64 root=UUID=c0b9ecd8-f922-4e5d-bccb-
83fbc94ad23b ro
echo    'Loading initial ramdisk ...'
...
```

4. Press **Esc** to enter **:wq**. Press **Enter** to save the configuration and exit the editor.

5. (Optional) Run the following command to ensure the modification has been successfully saved.

```
sudo cat /boot/grub/grub.cfg


...
linux   /boot/vmlinuz-6.1.0-13-amd64 root=UUID=c0b9ecd8-f922-4e5d-bccb-
83fbc94ad23b ro
...
```

6. (Optional) Delete the `grub.cfg` backup file in the `/home` directory.

Setting the Fstab Disk Identification Method to UUID

Last updated：2024-01-08 09:37:01

Operation Scenario

To guarantee correct disk recognition by the Linux system during file system mounting, please inspect and correctly set the fstab file disk identification method.

The file system table (fstab) is a configuration file in the Linux system that stores file system mounting information. Typically, the `/etc/fstab` file supports the use of device names (such as `/dev/vda1`) to identify file systems. However, device names may change due to the change in the actual operating environment after an image is imported, so there may be some problems using device names to identify file systems. To avoid these problems, you can change the file system identification method in the `/etc/fstab` file to UUID. The UUID is a unique characteristic string that identifies a disk partition and won't be affected by the change in device names. Using a UUID as the fstab file disk identification can ensure that the system can still correctly mount the file system when the device name changes.

Setting the Fstab Disk Identification Method to UUID

Confirming the Current Configuration of Fstab

Run the following command to view the current configuration method.

```
sudo cat /etc/fstab
```

If the output resembles the following one, with the first column beginning with UUID, it indicates that the current fstab is configured using the UUID method.

```
UUID=c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b / ext4 defaults 1 1
```

If the output resembles the following one, with the first column beginning with the block device name (such as `/dev/vda1`), it indicates that the current fstab is using a device name. You can refer to the subsequent operation to switch to the UUID method.

```
/dev/vda1 / ext4 defaults 1 1
```

Obtaining the UUID

To obtain the UUID of a partition, the `blkid` command can be used. Running the `blkid` command will display the detailed information of all the available partitions including the UUIDs. Run the following command in the terminal:

```
sudo blkid
```

The output similar to the following one indicates that the associated UUID of the device `/dev/vda1` is `c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b` .

```
/dev/vda1: UUID="c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b" BLOCK_SIZE="4096"
TYPE="ext4" PARTUUID="bcfcb5cb-01"
```

Modifying fstab

1. Backup the current `/etc/fstab` file to the `/home` directory.

```
sudo cp /etc/fstab /home
```

2. Use the **vi** editor to open the `/etc/fstab` file.

```
sudo vi /etc/fstab
```

3. Edit the configurations beginning with device names in the fstab file. Change device names to the `UUID=xxx` format. The content after `UUID=` is the UUID value corresponding to the device returned by running the `blkid` command.

```
# Before modification
/dev/vda1 / ext4 defaults 1 1
# After modification
UUID=c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b / ext4 defaults 1 1
```

4. Press **Esc** to enter **:wq**. Press **Enter** to save the configuration and exit the editor.

5. (Optional) Run the following command to ensure that the modification has been successfully saved.

```
sudo cat /etc/fstab
```

If the following content is returned, it indicates the modification has been saved successfully.

```
UUID=c0b9ecd8-f922-4e5d-bccb-83fbc94ad23b / ext4 defaults 1 1
```

6. (Optional) Run the following command. If no errors are returned, it means that the configuration has been successfully performed in accordance with the UUID method.

```
sudo mount -a
```

7. (Optional) Delete the backup fstab file in the `/home` directory.

If the modification to UUID identification failed, the system can be restored to the original state by restoring the fstab file.

```
sudo mv /home/fstab /etc/fstab
```

Exporting an image

Last updated：2024-01-08 09:37:00

Overview

Tencent Cloud allows you to export created custom images to COS buckets.

Precautions

You have activated the COS service in the COS console.

You have created a bucket in the region where the custom image to export resides. For more information, see Creating Bucket.

Note

You cannot export commercial images such as Windows images.

For a custom image, the capacity of a system disk or data disk cannot be greater than 500 GB.

When the image of an entire CVM instance is exported, the CVM instance cannot contain more than 5 data disks.

Billing Description

If you use other services such as COS when using CVM, fees will be calculated according to the billing rules of the actually used services.

The fees are as described below:

Use Case	Billing	Document
Exporting an image to a COS bucket	Storage usage fees. Storing an image in a COS bucket will incur storage usage fees. COS will calculate the object size and charge fees based on the storage type and region of the target object.	Storage Usage Fees
	Request fees. Exporting an image to a COS bucket will incur write request fees. COS will calculate the number of write requests and charge fees accordingly.	Request Fees
	Traffic fees. Exporting an image to a COS bucket will generate upstream traffic. COS will calculate the traffic volume. Private network upstream traffic and public network upstream traffic are free of charge.	Traffic Fees

| Downloading an image from a COS bucket | Request fees. Downloading an image from a COS bucket will incur write request fees. COS will calculate the number of write requests and charge fees accordingly. | Request Fees |
| | Traffic fees. Downloading an image from a COS bucket will generate downstream traffic. COS will calculate the traffic volume. Private network downstream traffic is free of charge, while public network downstream traffic is not. | Traffic Fees |

Directions

1. Log in to the CVM console and click **Images** in the left sidebar.

2. In the upper part of the **Images** page, select the region where the custom image to export resides and click the **Custom Image** tab.

3. Locate the image to export and choose **More** > **Export Image**.

4. In the pop-up **Export Image** window, set parameters as follows:

COS Bucket: Select the bucket where the image to export resides. Make sure that the bucket is in the same region as the image to export.

Export File Prefix: Customize the prefix of the file to export.

Select to **agree to authorize CVM to access my COS bucket**.

5. Click **OK** to start exporting the image.

6. In the pop-up window, click **OK**.

The export duration depends on the size of the image file and the length of the export task queue. After the export task is completed, the image file will be stored in the destination bucket. You can go to the **Bucket List** page and click the ID of the destination bucket to go to the bucket details page. On the bucket details page, the image file just exported is displayed as `Custom prefix_xvda.raw` .

FAQs

1. How is the public network downstream traffic in COS generated and billed?

Public network downstream traffic is the traffic generated by data transfer from COS to the client over the internet. Traffic generated by downloading an object directly through an object link or by browsing an object at a static website endpoint is public network downstream traffic. For more information about the billing details, see Billable Items and Pricing | Cloud Object Storage.

2. Will I be charged for public network downstream traffic generated by downloading files through the COS console, tools, API, or SDK?

The traffic (private or public network traffic) generated by accessing COS is subject to the use case, and only access to COS from a Tencent Cloud product in the same region will be over the private network by default, with no public network downstream traffic fees incurred. For more information on how to identify private network access, see Overview > Private network access.

3. What is public network traffic in COS?

Public network downstream traffic is the traffic generated by data transfer from COS to the client over the internet. Downloading a file stored in COS in the COS console, accessing or downloading an object through a tool, object address, or custom domain name, and previewing an object in a browser will generate public network downstream traffic. For more information, see Overview > Private network access.

4. Will accessing COS over the private network incur fees?

Accessing COS over the private network will incur **storage usage fees** and **request fees** but not **traffic fees**. For more information about the billing details, see Billable Items.

CentOS Linux Operations

Background

Last updated：2024-01-08 09:37:01

Background

CentOS plans to officially discontinue support for CentOS Linux, with details shown in the table below. For more information, see CentOS's official announcement.

OS Version	Maintenance End Date	Customer Impact
CentOS 8	January 1, 2022	After end of maintenance, any software maintenance and support including bug fixes and feature updates are unavailable.
CentOS 7	June 30, 2024	

Choose the OpenCloudOS Community Stable Edition (free of charge) or TencentOS Server images for newly purchased CVMs

Introduction to OpenCloudOS and TencentOS Server

Initiated by Tencent and its partners, **CloudOpenOS** is a standalone, fully open OS and ecosystem with security, stability, and high performance.
For more information about OpenCloudOS, see OpenCloudOS Overview.
TencentOS Server is Tencent's Linux OS designed for cloud scenarios. With specific features and optimized performance, TencentOS Server provides a high-performance, secure, and reliable operating environment for applications in CVM instances.
For a more comprehensive understanding of TencentOS Server, please refer to the TencentOS Server Documentation.

Linux editions issued in the ecosystem supply chain are classified into four categories here:

L1 Stream Edition, such as OpenCloudOS Stream and the well-known Fedora and Debian.
L2 Commercial Edition, the majority of which are issued by commercial companies, such as TencentOS Server by Tencent, RHEL by Redhat, and Ubuntu by Canonical.
L3 Community Stable Edition, usually a free reissue of a commercial system such as OpenCloudOS and the original CentOS. This edition has few differences from the L2 Commercial Edition.

L4 Community Derived Edition, an optimized and customized edition based on L3.

OpenCloudOS falls under the L3 Community Stable Edition category and TencentOS Server the L2 Commercial Edition category. OpenCloudOS is to TencentOS Server what CentOS is to RHEL.

OpenCloudOS is derived from the commercial stable edition of TencentOS Server and has basically the same source code. The main difference is that the commercial edition provides SLA Guaranteed technical support.

	OpenCloudOS	TencentOS Server
Kernel version	Linux v5.4 kernel	Linux v5.4 kernel
User mode	Compatible with CentOS 8 (OpenCloudOS 8.X)	Compatible with CentOS 7 (TencentOS Server 2.4) and CentOS 8 (TencentOS Server 3.1)
Technical support	From the OpenCloudOS community	From TencentOS Server technical support
Flaw/vulnerability publish	In the community	By TencentOS Server technical support

OpenCloudOS is a community-based OS, available to you for free and maintained by developers in the community.

If you need service and maintenance from a professional OS team, you can purchase the TencentOS Server

subscription service.

Directions for CentOS migration

If you are using CentOS 8, you can migrate it to OpenCloudOS as instructed in Migrating CentOS to OpenCloudOS. If you have CentOS instances, you can migrate them to TencentOS Server as instructed in Migrating CentOS to TencentOS.

Migrating CentOS to TencentOS Server

Last updated：2024-01-08 09:37:01

Overview

CentOS plans to officially discontinue support for CentOS Linux, with details shown in the table below. For more information, see CentOS's official announcement.

OS Version	EOL	Impact
CentOS 8	January 1, 2022	After end of maintenance, any software maintenance and support including bug fixes and feature updates are unavailable.
CentOS 7	June 30, 2024	

Choose TencentOS Server images for newly purchased CVMs. For existing CentOS instances, you can migrate them to TencentOS Server.

Supported versions

OS versions supported for source servers:

CentOS 7 series:

CentOS_7.2_64-bit, CentOS_7.3_64-bit, CentOS_7.4_64-bit, CentOS_7.5_64-bit, CentOS_7.6_64-bit, CentOS_7.7_64-bit, CentOS_7.8_64-bit, and CentOS_7.9_64-bit

CentOS 8 series:

CentOS_8.0_64-bit, CentOS_8.2_64-bit, and CentOS_8.4_64-bit

OS versions recommended for target servers:

CentOS 7 series: TencentOS Server 2.4 (TK4) is recommended.

CentOS 8 series: TencentOS Server 3.1 (TK4) is recommended.

Note:

CentOS 7.2 and CentOS 7.3 public images may contain packages for 32-bit systems by default. Delete these packages before update.

Reminders

OS migration is not supported in the following cases:

A GUI is installed.

An i686 RPM package is installed.

Business may fail to run properly after migration under the following conditions:

The business program is installed with and relies on a third-party RPM package.

The business program relies on a fixed kernel version or has its own kernel module compiled.

The target version after migration is TK14 based on the v5.4 kernel. This version is later than the kernel versions of CentOS 7 and CentOS 8 and may have changes in some old features.

The business program relies on a fixed GCC version.

TencentOS 2.4 is installed with GCC v4.8.5 by default, and TencentOS 3.1 is installed with GCC v8.5 by default.

After migration, you need to restart the instance to enter the TencentOS kernel.

Migration does not affect data disks. Upgrade only in the OS layer does not involve any operation on data disks.

Resource Requirements

500 MB of available memory

10 GB of available space in the system disk of the destination instance

Directions

Preparation

1. Create a snapshot to back up system disk data.

2. Uninstall i686 RPM package (if any).

Migration execution

CentOS 7 series to TencentOS Server 2.4 (TK4)

CentOS 8 series to TencentOS Server 3.1 (TK4)

1. Log in to the target CVM instance. For operation details, see Logging in to Linux Instance Using Standard Login Method.

2. Run the following command to install Python 3:

```
yum install -y python3
```

3. Run the following command to obtain the migration tool:

```
wget
http://mirrors.tencent.com/tencentos/2.4/tlinux/x86_64/RPMS/migrate2tencentos-
1.0-4.tl2.noarch.rpm
```

4. Run the following command to install the migration tool. The command will create `migrate2tencentos.py` in `/usr/sbin` .

```
rpm -ivh migrate2tencentos-1.0-4.tl2.noarch.rpm
```

5. Run the following command to start migration:

```
python3 /usr/sbin/migrate2tencentos.py -v 2.4
```

The migration takes some time. When the script execution is completed, the following information will be displayed:

```
Metadata Cache Created
Switch complete. TencentOS recommends rebooting this system.
[root@VM-2-43-centos ~]#
```

6. Restart the instance. For operation details, see Restarting Instances.

7. Check the migration result.

7.1 Run the following command to check the OS release information:

```
cat /etc/os-release
```

The information shown in the figure below is displayed:

```
[root@VM-2-43-centos ~]# cat /etc/os-release
NAME="TencentOS Server"
VERSION="2.4"
ID="tencentos"
ID_LIKE="rhel fedora centos tlinux"
VERSION_ID="2.4"
PRETTY_NAME="TencentOS Server 2.4"
ANSI_COLOR="0;31"
CPE_NAME="cpe:/o:tencentos:tencentos:2"
HOME_URL="https://cloud.tencent.com/product/ts"
```

7.2 Run the following command to check the kernel:

```
uname -r
```

The information shown in the figure below is displayed:

```
[root@VM-2-43-centos ~]# uname -r
5.4.119-19-0009.1
[root@VM-2-43-centos ~]#
```

Note:

By default, the kernel is the latest version of YUM. The actual result prevails. This document uses the version shown in the figure as an example.

7.3 Run the following command to check YUM:

```
yum makecache
```

The information shown in the figure below is displayed:

```
[root@VM-2-43-centos ~]# yum makecache
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
 * epel: mirrors.tencentyun.com
 * tlinux: mirrors.tencentyun.com
 * tlinux-extras: mirrors.tencentyun.com
 * tlinux-os: mirrors.tencentyun.com
 * tlinux-updates: mirrors.tencentyun.com
epel
tlinux
tlinux-extras
tlinux-os
tlinux-tkernel4
tlinux-updates
Metadata Cache Created
[root@VM-2-43-centos ~]#
```

1. Log in to the target CVM instance. For operation details, see Logging in to Linux Instance Using Standard Login Method.

2. Run the following command to install Python 3:

```
yum install -y python3
```

3. Run the following command to obtain the migration tool:

```
wget
http://mirrors.tencent.com/tlinux/3.1/Updates/x86_64/RPMS/migrate2tencentos-
1.0-4.tl3.noarch.rpm
```

4. Run the following command to install the migration tool. The command will create `migrate2tencentos.py` in `/usr/sbin` .

```
rpm -ivh migrate2tencentos-1.0-4.tl3.noarch.rpm
```

5. Run the following command to start migration:

```
python3 /usr/sbin/migrate2tencentos.py -v 3.1
```

The migration takes some time. When the script execution is completed, the following information will be displayed:

```
Metadata cache created.
Switch complete. TencentOS recommends rebooting this system.
[root@VM-2-2-centos ~]# 
```

6. Restart the instance. For operation details, see Restarting Instances.

7. Check the migration result.

7.1 Run the following command to check the OS release information:

```
cat /etc/os-release
```

The information shown in the figure below is displayed:

```
[root@VM-2-2-centos ~]# cat /etc/os-release
NAME="TencentOS Server"
VERSION="3.1 (Final)"
ID="tencentos"
ID_LIKE="rhel fedora centos"
VERSION_ID="3.1"
PLATFORM_ID="platform:el8"
PRETTY_NAME="TencentOS Server 3.1 (Final)"
ANSI_COLOR="0;31"
CPE_NAME="cpe:/o:tencentos:tencentos:3"
HOME_URL="https://cloud.tencent.com/product/ts"
```

7.2 Run the following command to check the kernel:

```
uname -r
```

The information shown in the figure below is displayed:

```
[root@VM-2-2-centos ~]# uname -r
5.4.119-19-0009.1
[root@VM-2-2-centos ~]#
```

Note:

By default, the kernel is the latest version of YUM. The actual result prevails. This document uses the version shown in the figure as an example.

7.3 Run the following command to check YUM:

```
yum makecache
```

The information shown in the figure below is displayed:

```
[root@VM-2-2-centos ~]# yum makecache
TencentOS Server 3.1 - TencentOS
TencentOS Server 3.1 - Updates
TencentOS Server 3.1 - TencentOS-AppStream
TencentOS Server 3.1 - Base
TencentOS Server 3.1 - AppStream
TencentOS Server 3.1 - Extras
TencentOS Server 3.1 - PowerTools
Extra Packages for TencentOS Server 3.1 - x86_64
Extra Packages for TencentOS Server 3.1 Modular - x86_64
Metadata cache created.
[root@VM-2-2-centos ~]#
```

Should you encounter any issues during the migration process, please reach out to the Contact Us.

Migrating CentOS to OpenCloudOS

Last updated：2024-01-08 09:37:01

Overview

CentOS has officially discontinued support for CentOS 8 (see CentOS's official announcement). This document guides you migrating your servers from CentOS 8 to OpenCloudOS.

OS Version	EOL	Impact
CentOS 8	January 1, 2022	After end of maintenance, any software maintenance and support including bug fixes and feature updates are unavailable.

Version Description

OS versions supported for source servers

OS	Version
CentOS 8 series	CentOS_8.0_64-bit, CentOS_8.2_64-bit, CentOS_8.3_64-bit, CentOS_8.4_64-bit, and CentOS_8.2_ARM64

OS versions recommended for target servers

If you are using CentOS 8 series, migrate it to OpenCloudOS 8.

OS migration is not supported for CentOS Stream 8 public images.

Limits

OS migration is not supported in the following cases:

A GUI is installed.

An i686 RPM package is installed.

Business may fail to run properly after migration under the following conditions:

The business program is installed with and relies on a third-party RPM package.

The business program relies on a fixed kernel version or has its own kernel module compiled. The target version after migration is tkernel4 (TK4) based on the v5.4 kernel. This version is later than the kernel versions of CentOS 8. Some old features may be updated in this new version. If your business program relies heavily on the kernel, we recommend

that you know which features your business program actually relies on. You can also visit the OpenCloudOS community Bugtracker.

The business program relies on a fixed GCC version. Currently, OpenCloudOS 8 is installed with GCC v8.5 by default.

After migration, you need to restart the instance to enter the OpenCloudOS kernel.

Migration does not affect data disks. Upgrade only in the OS layer does not involve any operation on data disks.

Requirements

500 MB of available memory

10 GB of available space in the system disk

Directions

Preparation

1. Create a snapshot to back up system disk data before you start migration.

2. Check whether an i686 RPM package is installed and, if so, uninstall the package.

3. Install Python 3 in your operating environment if you have not installed it. You can install Python 3 using a CentOS Vault repository.

```
# cat <<EOF | sudo tee /tmp/centos8_vault.repo
[c8_vault_baseos]
name=c8_vault - BaseOS
baseurl=https://mirrors.cloud.tencent.com/centos-vault/8.5.2111/BaseOS/\\$basearch/
gpgcheck=0
enabled=1
[c8_vault_appstream]
name=c8_vault - AppStream
baseurl=https://mirrors.cloud.tencent.com/centos-vault/8.5.2111/AppStream/\\$basear
gpgcheck=0
enabled=1
EOF
# yum -y install python3 --disablerepo=* -c /tmp/centos8_vault.repo --enablerepo=c8
```

Migration execution

Do to following to migrate a CentOS 8 instance to OpenCloudOS 8:

1. Log in to the target CVM instance. See Logging In To Linux Instance (Web Shell).

2. Run the following command to install Python 3. If no YUM repository is available, install Python 3 using a CentOS Vault repository. For more information, see item 3 in the **Preparing for the migration** section.

```
yum install -y python3
```

3. Run one of the following commands based on your Python version to download the migration tool:

```
#x86 version
wget https://mirrors.opencloudos.tech/opencloudos/8.6/AppStream/x86_64/os/Packages/
#ARM version
wget https://mirrors.opencloudos.tech/opencloudos/8/AppStream/aarch64/os/Packages/m
```

4. Run the following command to install the migration tool. The command will create the

`migrate2opencloudos.py` file in `/usr/sbin` .

```
rpm -ivh migrate2opencloudos-1.0-1.oc8.noarch.rpm
```

5. Run the following command to start migration:

```
python3 /usr/sbin/migrate2opencloudos.py -v 8
```

The migration takes some time. When the script execution is completed, the following information will be displayed:

```
Metadata cache created.
Switch complete. OpenCloudOS recommends rebooting this system.
[root@VM-64-6-centos ~]#
```

6. Restart the instance. See Restarting Instances.

7. Check the migration result.

Run the following command to check the OS release information:

```
cat /etc/os-release
```

The information shown in the figure below is displayed:

```
[root@VM-64-27-centos ~]# cat /etc/os-release
NAME="OpenCloudOS"
VERSION="8.6"
ID="opencloudos"
ID_LIKE="rhel fedora"
VERSION_ID="8.6"
PLATFORM_ID="platform:oc8"
PRETTY_NAME="OpenCloudOS 8.6"
ANSI_COLOR="0;31"
CPE_NAME="cpe:/o:opencloudos:opencloudos:8"
HOME_URL="https://www.opencloudos.org/"
BUG_REPORT_URL="https://bugs.opencloudos.tech/"
```

Run the following command to check the kernel:

```
uname -r
```

The information shown in the figure below is displayed:

```
[root@VM-64-27-centos ~]# uname -r
5.4.119-19.0010.ocrelease.7
```

By default, the kernel is the latest version of YUM.

Run the following command to check YUM:

```
yum makecache
```

The information shown in the figure below is displayed:

```
[root@VM-64-6-centos ~]# yum makecache
OpenCloudOS 8 - Base
OpenCloudOS 8 - AppStream
OpenCloudOS 8 - Extras
OpenCloudOS 8 - HighAvailability
OpenCloudOS 8 - PowerTools
OpenCloudOS 8 - ResilientStorage
Extra Packages for OpenCloudOS 8 - x86_64
Extra Packages for OpenCloudOS 8 Modular - x86_64
Metadata cache created.
[root@VM-64-6-centos ~]#
```

Migrating Servers
Online Migration
Overview

Last updated：2024-01-08 09:37:00

Online migration supports migrating or synchronizing systems and applications on the source server or virtual machine from your IDCs or other cloud platforms to Tencent Cloud with no system downtime.
With go2tencentcloud, the migration tool provided by Tencent Cloud, you can directly migrate all systems and applications on the source server to the destination CVM, without the need to create, upload, and import images. It can meet enterprises' business requirements for cloud deployment, cross-cloud migration, cross-account or cross-region migration, and hybrid cloud deployment.
Note:
The source server can be a physical server, a virtual machine, or a cloud server on third-party cloud platform, such as AWS, Google Cloud Platform, VMware, Alibaba Cloud, or Huawei Cloud.

Use Cases

Online migration is applicable to the scenarios including but not limited to:
 IT architecture cloudification
Hybrid cloud architecture deployment
Cross-cloud migration
Cross-account or cross-region migration

Differences from Offline Migration

In offline migration, you need to create images for system disks or data disks on source servers, and then migrate images to the Cloud Virtual Machine (CVM) or Cloud Block Storage (CBS). You do not need to create images for online migration. Instead, you can run the migration tool on source servers to migrate them to destination CVMs.

Starting Migration

Two methods for online migration are available. You can select the appropriate one as needed.

Migration	Overview	Use cases	Characteristics

method			
Online Migration - Importing source by using the migration tool	Log in to the source instance, import the migration source with the tool, and create a migration task in the console to implement the migration.	Migrate via the public and private network Migrate from other cloud platforms to Tencent Cloud Migrate from customer IDC to Tencent Cloud	High compatibility
Online Migration - Console	Log in to the console, perform identity authentication, quickly import the migration source and create a migration task.	No need to log in to the source server Migration via the public network Cross-cloud migration: Applicable to the scenarios where the source instances are located in Alibaba Cloud	Quick migration in batch

FAQs

For more information, please see FAQs about Server Migration.

Migration Operation Guide
Online Migration Directions

Last updated：2024-01-08 09:37:01

Online migration supports migrating or synchronizing systems and applications on the source server or virtual machine from your IDCs or other cloud platforms to Tencent Cloud with no system downtime.

Two methods for online migration are available. You can select the appropriate one as needed.

Migration method	Overview	Use cases	Characteristics
Online Migration - Importing Migration Source from the Client	Log in to the source instance, import the migration source with the tool, and create a migration task in the console to implement the migration.	Migration via the public and private networks Cross-cloud migration: Applicable to any source environment Migration from IDCs to cloud	High compatibility
Online Migration - Quick Migration in the Console	Log in to the console, perform identity authentication, quickly import the migration source and create a migration task.	Migration via the public network Cross-cloud migration: Applicable to the scenarios where the source instances are located in Alibaba Cloud	Quick migration in batchOperation in the console

Migrating with a Migration Tool

Last updated：2025-06-18 10:22:11

This document describes how to migrate your source server to Tencent Cloud CVM by importing the migration source from the client.

Migration Workflow

The procedure of importing the migration source from the client is shown below:

Migration Directions

Step 1. Prepare for migration

Go to Manage API Key to create a key and obtain the `SecretId` and `SecretKey` .

Stop applications on the server and back up your data.

Source server: You can use the snapshot feature or other methods to back up data on the source server. The source server is the server to be migrated.

Destination CVM: Create a snapshot of the instance (See Creating Snapshots) to back up the data.

If you are using a sub-account, ask the root account to assign you the `QcloudCSMFullAccess` and `QcloudCVMFullAccess` permissions in the CAM console.

Before the migration, you need to check the following configuration based on the actual conditions:

Migrate to a CVM instance: Check the source server and destination CVM.

Migrate to a CVM image: Check the source server.

Linux source server	1. Check and install Virtio. For more information, see Checking Virtio Drivers in Linux. 2. Run `which rsync` to check whether Rsync is installed. If not, install it as instructed in How do I install Rsync? 3. Check whether SELinux is enabled, If yes, disable it as instructed in How do I disable SELinux? 4. After a migration request is made to the Tencent Cloud API, the API will use the current UNIX time to check the generated token. Make sure that the system time of your server is correct.
Windows source server	1. Check and install Virtio. For more information, see Checking or installing the Virtio driver.

	2. (Optional) Check and install Cloudbase-Init (See Installing Cloudbase-Init on Windows). It's recommended to install it on the source server before the migration. In this case, the network configuration and OS license activation are performed automatically after the migration. Otherwise you need to log in to the instance via VNC, and modify the network configuration manually on the destination server after migration.
Destination CVM	Storage space: The cloud disks (including the system disk and data disks) of the destination CVM must offer sufficient storage space for saving data from the source server. Security group: Port 80, port 443 and port 3389 are opened. Bandwidth: Set the bandwidth cap on both the two ends to the highest possible value. During the process, the traffic consumed is approximately the amount of data migrated. Adjust the billing mode beforehand if necessary. Network: If the source or destination server only supports IPv6 but not IPv4, see Parameters in the client.json file.

Note:

Check the source server by executing `sudo ./go2tencentcloud_x64 --check` .

By default, go2tencentcloud automatically performs checks upon launch. To skip checks, open the `client.json`

file, set `Client.Extra.IgnoreCheck` to `true` .

Step 2. Import the migration source

Import with the migration tool

Linux source server

Windows source server

1. Run the following command on the source server to download the migration tool `go2tencentcloud.zip` , and

go to the corresponding directory.

```
wget https://go2tencentcloud-1251783334.cos.ap-
guangzhou.myqcloud.com/latest/go2tencentcloud.zip


unzip go2tencentcloud.zip


cd go2tencentcloud/go2tencentcloud-linux
```

Note:

The files in the `go2tencentcloud` directory will not be migrated. Do not place the files to be migrated in this

directory.

2. (Optional) Exclude files and directories on the source server that do not need to be migrated.

Add files and directories that don't need to be migrate to the rsync_excludes_linux.txt file.

3. Import the migration source.

3.1 For example, on a 64-bit Linux source server, execute the following commands in sequence as the root user to run the tool.

```
chmod +x go2tencentcloud_x64


sudo ./go2tencentcloud_x64
```

3.2 Enter the `SecretId` and `SecretKey` of the account API access key obtained in Prerequisites and press **Enter** as shown below:

If you see the following message, the source server is imported successfully. You can now see the server in the CVM console.

1. Download or upload `go2tencentcloud.zip` to the source server. Decompress the file to the `go2tencentcloud` folder. Open `go2tencentcloud-windows` , and the directory is shown as below.

2. Run `go2tencentcloud_x64.exe` .
Method 1: Right-click `go2tencentcloud_x64.exe` and run it as admin. Enter `SecretId` and `SecretKey` in the pop-up window.
Method 2: Start cmd or PowerShell command line as admin: cd /d "absolute path of the directory of go2tencentcloud_x64.exe", and run `go2tencentcloud_x64.exe` .
3. Enter Tencent Cloud API key (`SecretId` and `SecretKey`) in the pop-up window.

4. If the following message appears, the source server information is imported. You can now check the source server in the CVM console.

Note:
If "Import source server successfully" does not appear, check the logs in the `logs/log` file under the migration tool directory for troubleshooting.

Check the source server in the console

Log in to the CVM console and check the imported server. Its status should be **Online**, as shown below:

Note:

Importing the source server is the first step of migration. Keep the migration tool alive till the whole migration progress ends.

Step 3. Create a migration task

1. Create a migration task

Log in to the CVM console, go to the online migration page, locate the source server, and click **Create migration task**. In the **Create migration task** pop-up window, configure the task as shown below:

Configuration description:

Basics:

Item	Required	Description
Destination region	Yes	Tencent Cloud region to which the source server is to be migrated. For more information on regions, see Regions and AZs.
Task name	Yes	The migration task name.
Task description	No	Migration task description.
Destination type	Yes	Set the destination type for the source server to be migrated to Tencent Cloud. **CVM image**: Create a CVM image for the source server. Image name: Name of the destination CVM image that will be generated for the migration source. If the name already exists, the migration task ID is appended to the name. **CVM instance**: Select a CVM instance in the destination region as the migration destination. Destination instance: We recommend you use the same operating system for the source server and destination CVM. For example, to migrate a CentOS 7 source server, select a CentOS 7 CVM as the destination.
Network mode	Yes	The network used for transferring data. **Public network**: Transfer over the public network. **Private network**: Transfer over the private network. For details, see Migrating via Private Network.

		VPC: Create the relay instance in a VPC when migrating to a CVM image. Subnet: Create the relay instance in a subnet when migrating to a CVM image.
Migration method	Yes	For Linux instances: **File-level migration**: Higher compatibility and relatively slower transfer speed. **Block-level migration**: Faster transfer speed and relatively lower compatibility.
		Windows block-level migration: Block-level migration, with faster transfer speed and relatively lower compatibility.
Configure incremental sync	No	You can customize the incremental sync duration to continuously sync the data. **Not enable**: The migration tool scans for and migrate the increments. Generally, it is implemented for once. **Enable**: You can select the incremental sync duration. The migration tool will continuously sync the data to Tencent Cloud. You can also manually stop the incremental sync in the task list.
Scheduled execution time	No	Set the time when the migration task will be automatically started after creation. It can be as early as 10 minutes after the current time.

Advanced (Optional):

Item	Required	Description
Data rate (KB/s)	No	The upper limit of data rate during the migration (0 to 25600 KB/s). It's set to 0 by default. This item is not available for migration to Windows.
Checksum verification	No	When it is enabled, data consistency check is enhanced, but the transfer speed may be reduced. This item is not available for migration on Windows.

2. Start the migration task

Note:

You can skip this step if your task is scheduled, which will automatically start running at the scheduled execution time.

After creating a migration task, you can click the **Migration task** tab to view the task as shown below:

You can click **Start**/**Retry** on the right of the task to start it, click **OK** in the pop-up window, and the task status will become **Migrating** as shown below:

Note:

If the migration destination is a CVM instance, the destination CVM enters migration mode after the migration starts. Do not reinstall the system, shut down, terminate, or reset passwords of the destination CVM until the migration ends and the destination CVM exits the migration mode.

If the migration destination is a CVM image, a relay instance `do_not_delete_csm_instance` will be created under your account after the migration starts. Don't reinstall, shut down, or terminate the relay instance or reset its password. It will be automatically terminated by the system after the migration ends.

Step 4. Check after migration

1. View the migration progress in the console

After the migration task status becomes **Successful**, the migration is completed successfully, as shown below:

Note:

The time required for data transfer depends on the size of the data on the source server, network bandwidth, etc. Please wait for the migration process to complete.

After the migration task starts, you can click **Pause** on the row of the task to stop it.

The migration tool supports checkpoint restart. After a task is paused, you can click **Start**/**Retry** again to resume migration from where you paused.

A migration task can be paused during data transfer. After you click **Pause** for it in the console, the migration tool will pause the data transfer in progress.

If the migration process is time-consuming and you need to stop it, you can pause the migration task first and click **Delete** to delete it.

2. Check after migration

Failed migration:

Check the error information in log files (under the migration tool directory by default), operation guides, or FAQs about Server Migration for troubleshooting. After troubleshooting, click **Start**/**Retry** under the operation column to restart the migration task.

Successful migration:

Migrating to a CVM: The destination CVM starts up normally. Data on the CVM is consistent with that on the source server. The network and other system services are normal.

Migrating to a CVM image: Click the **CVM image ID** on the row of the migration task to go to the CVM image page and view the image information. You can use this image to create CVM instances.

If you have any questions or the migration has an exception, see FAQs about Server Migration or contact us.

Migrating via the Console

Last updated：2024-01-08 09:37:01

This document describes how to migrate a source server to Tencent Cloud CVM through quick migration in the Console.

Overview

Quick migration is an agile version of online migration. It eliminates complex operations such as login to the source server and tool download, allowing you to create a batch migration task to migrate data such as the operating system and applications on the source server to Tencent Cloud.
Quick migration supports both the Linux and Windows operating systems. You can query the migration progress on the **Online migration** page of the CVM console.

Limits

Quick migration has certain requirements for the source server environment. Specifically, you need to install the cloud assistant (such as Alibaba Cloud ECS Cloud Assistant), configure the public IP, and use the VPC (classic networks are not supported) on the source server.
Currently, quick migration allows you to migrate only Alibaba Cloud servers to Tencent Cloud.
The quick migration feature is iteratively optimized and is supported only in certain scenarios currently. If your scenarios are not covered, please choose Online Migration - Importing Migration Source from the Client.

Migration Workflow

The procedure of quick migration in the console is shown below:

Migration Directions

Step 1. Prepare for migration

Get the access key in the Tencent Cloud console

Create an API key and get the **SecretId** and **SecretKey** on the **Manage API Key** page in the Cloud Access
Management (CAM) console. For detailed directions, see Root Account Access Key Management.

Note:

If you want to migrate the source server by using a sub-account, ask the root account owner to associate the
QcloudCSMFullAccess and QcloudCVMFullAccess policies with the sub-account in the CAM console.

Get the access key on the source cloud platform

Get the `AccessKeyID` and `AccessKeySecret` of Alibaba Cloud in the following steps:

1.1 Log in to the RAM console and select **Identities** > **Users**.

1.2 Click **Create User** and select **OpenAPI Access** for **Access Mode**. (Do not select **Console Access**, which does
not apply to this scenario.) Then, save the `AccessKeyID` and `AccessKeySecret` . For detailed directions, see
Create a RAM user.

1.3 In the user list, find the target user and click **Add Permissions** to add the ECS read-only permission
(`AliyunECSReadOnlyAccess`) and ECS cloud assistant management permission
(`AliyunECSAssistantFullAccess`). For detailed directions, see Grant permissions to the RAM user.

(Optional) Stop applications on the source server

We recommend you stop all applications on the source server to prevent them from being affected by the migration.

(Optional) Back up data on the source server and destination CVM

We recommend you back up data in the following ways before the migration:

Source server: You can use the snapshot feature or other methods to back up data on the source server.

Destination CVM: Create a snapshot of the instance (See Creating Snapshots) to back up the data.

Check the destination CVM

If the migration destination is a CVM instance, you need to check the destination CVM.

Destination CVM	1. Storage space: The cloud disks (including the system disk and data disks) of the destination CVM must offer sufficient storage space for saving data from the source server. 2. Security group: Port 80, port 443 and port 3389 are opened. 3. Bandwidth: Set the bandwidth cap on both the two ends to the highest possible value. During the process, the traffic consumed is approximately the amount of data migrated. Adjust the billing mode beforehand if necessary. 4. Network: If the source or destination server only supports IPv6 but not IPv4, see Parameters in the client.json file.

Go to the quick migration page

i. Log in to the CVM console and choose **Server Migration** > **Online Migration** in the left sidebar. Click **Import migration source** to go to the Import migration source page.

ii. Select **Quick migration** to batch create migration tasks.

Step 2. Create a migration task

1. **Configure the task**

Enter the task name and description.

2. **Configure the migration source**

The source ISP is set to Alibaba Cloud ECS by default, and you need to enter the `AccessKey` and `SecretKey` of your Alibaba Cloud account (they can be obtained as instructed here). Then, verify that **you have the permission to access the source server information** as shown below:

Note:

Keep your access key confidential. We recommend you delete or disable the access key after the migration.

3. **Configure the migration destination**

The destination ISP is set to CVM by default, and you need to enter the `SecretId` and `SecretKey` of your TencentCloud API (they can be obtained as instructed here) to **get the permissions to use CVM**. You can copy the key information on the **Manage API Key** page. Make sure that the API key is correct. Otherwise, the migration fails.

Note:

Keep your access key confidential. We recommend you delete or disable the access key after the migration.

4. **Configure the migration information**

4.1 After the migration source information is verified, click **Add migration source** to select the target instance in the pop-up window.

4.2 Select the **region** in the top-left corner of the pop-up window to get the **instance list** in the region. The number after the region indicates the number of instances.

4.3 Select the target instance to add it to the **Selected** list on the right.

Note:

You can batch migrate **instances from different regions** and add migration sources multiple times.

Currently, you can batch migrate up to five instances.

4.4 Click **OK**. Then, the information of the target instance is displayed in the migration source list. You can click **Add destination information** in the **Operation** column to configure the migration destination information.

4.5 In the **Add migration destination** pop-up window, select the region and migration destination type:

Item	Required	Description
Destination region	Yes	Tencent Cloud region to which the source server is to be migrated. For more information on regions, see Regions and AZs.
Destination type	Yes	The type of the Tencent Cloud destination to which the migration source is to be migrated. CVM image: The Tencent Cloud destination image to be generated for the migration source after the migration task is completed.Image name: The name of the Tencent Cloud destination image generated for the migration source. If the image name already exists in the target region, the migration task will automatically add the task ID in the image name. CVM instance: The CVM instance in the target region to be used as the migration destination.Destination instance: We recommend you select a destination CVM instance with the same operating system as the source server. For example, to migrate a CentOS 7 source server, select a CentOS 7 CVM instance as the destination. In addition, the system disk and data disk capacity of the destination CVM instance must be larger than those of the source server.

5. **Click Create and start the migration task. The Reminder window will pop up. Note the following:**

You need to wait a minute before the progress can be viewed in the console, as it takes some time to execute the task on the migration source.

If the migration source fails to be imported due to an abnormal source server environment or incorrect information, the failure cause may not be indicated in the Tencent Cloud console. In this case, recreate the task or use online migration instead.

Step 3. Check after migration

1. **View the migration status and progress**

A successfully created migration task will run automatically. You can view the migration source information on the

migration source page and the task progress on the migration task page.

If the migration destination is a CVM instance, the destination CVM enters migration mode after the migration starts.

Do not reinstall the system, shut down, terminate, or reset passwords of the destination CVM until the migration ends and the destination CVM exits the migration mode.

If the migration destination is a CVM image, a relay instance named `do_not_delete_csm_instance` will be created under your account after the migration starts. Do not reinstall, shut down, or terminate the relay instance or reset its password. It will be automatically terminated by the system after the migration ends.

2. **Wait for the migration task to end**

After the migration task status becomes **Successful**, the migration is completed successfully, as shown below:

Note:

The time required for data transfer depends on the size of the data on the source server, network bandwidth, etc. Please wait for the migration process to complete.

After the migration task starts, you can click **Pause** on the row of the task to stop it.

The migration tool supports checkpoint restart. After a task is paused, you can click **Start/Retry** again to resume migration from where you paused.

A migration task can be paused during data transfer. After you click **Pause** for it in the console, the migration tool will pause the data transfer in progress.

If the migration process is time-consuming and you need to stop it, you can pause the migration task first and click **Delete** to delete it.

3. **Check after migration**

Failed migration:

Check the error information in log files (under the migration tool directory by default), operation guides, or FAQs about Server Migration for troubleshooting. After troubleshooting, click **Start/Retry** under the operation column to restart the migration task.

Successful migration:

Migrating to a CVM: The destination CVM starts up normally. Data on the CVM is consistent with that on the source server. The network and other system services are normal.

Migrating to a CVM image: Click the **CVM image ID** on the row of the migration task to enter the CVM image page and view the image information. You can use this image to create CVM instances.

If you have any questions or the migration has an exception, see FAQs about Server Migration or Contact Us.

Migration Tool Compatibility and Tool Configuration Description

Last updated：2024-01-08 09:37:01

Supported Operating Systems

Operating systems supported by the online migration tool include but not limited to the following:

Linux	Windows
CentOS 5/6/7/8	
Ubuntu 10/12/14/16/18/20	
Debian 7/8/9/10	Windows Server 2008
SUSE 11/12/15	Windows Server 2012
openSUSE 42	Windows Server 2016
Amazon Linux AMI	Windows Server 2019
Red Hat 5/6/7/8	Windows Server 2022
Oracle Linux 5/6/7/8	

Supported Migration Modes

Public network migration mode

Private network migration mode

If both your source server and destination CVM can access the public network, you can use the public network migration mode.

In the current public network migration mode, the source server calls Tencent Cloud APIs through the Internet to initiate a migration request, and transfers data to the destination CVM to complete the migration. The public network

migration scenario is shown below:

If your source server or destination CVM is located in a private network or Virtual Private Cloud (VPC), the source server cannot directly establish a connection with the destination CVM through the Internet. In this case, you can use the private network migration mode of the tool. You need to establish a connection between the source server and the destination CVM through VPC peering connection, VPN connections, Cloud Connect Network, or Direct Connect.

Scenario 1: This scenario is applicable to the migration via Online Migration Tool. If your source server or the destination CVM cannot access the public network, you can first access the Tencent Cloud API via the internet through a server with public network access (such as a gateway) to initiate a migration request, and then transfer data and migrate to the destination CVM through the connection. This scenario does not require the source server and destination CVM can access the public network.

Scenario 2: If your source server can access the public network, use the source server to call Tencent Cloud APIs through the Internet to initiate a migration request, and then transfer data to the destination CVM through the connection to complete the migration. This scenario requires the source server, but not the destination CVM, to be able to access the public network.

Scenario 3: If your source server can access the public network through a proxy, use the source server to call Tencent Cloud APIs through the network proxy to initiate a migration request, and then transfer data to the destination CVM through the connection to complete the migration. This scenario requires neither the source server nor the destination CVM to be able to access the public network.

Files in the Compressed Package

After `go2tencentcloud.zip` is decompressed, it contains the following files:

File Name	Description
go2tencentcloud-linux.zip	The migration zip for Linux system.
go2tencentcloud-windows.zip	The migration zip for Windows system.
readme.txt	Directory overview file.
release_notes.txt	Migration tool change log.

After `go2tencentcloud-linux.zip` is decompressed, it contains the following files:

File Name	Description
go2tencentcloud_x64	Executable program of the migration tool for the 64-bit Linux operating

	system.
go2tencentcloud_x32	Executable program of the migration tool for the 32-bit Linux operating system.
user.json	User information in the migration.
client.json	Configuration file of the migration tool.
rsync_excludes_linux.txt	rsync configuration file, which excludes files and directories that do not need to be migrated in the Linux system.

After `go2tencentcloud-windows.zip` is decompressed, it contains the following files:

File Name	Description
go2tencentcloud_x64.exe	Executable program of the migration tool for the 64-bit Windows operating system.
user.json	User information in the migration.
client.json	Configuration file of the migration tool.
client.exe	Executable program of the migration tool for the Windows operating system.

Note:

The configuration files cannot be deleted. You must store them under the same folder as the go2tencentcloud executable program.

Parameters in the user.json file

The user.json configuration file is described as below:

Parameter	Type	Required	Description
SecretId	String	Yes	Secret ID for your account to access APIs. For more information, see Access Key.
SecretKey	String	Yes	Secret key for your account to access APIs. For more information, see Access Key.

Parameters in the client.json file

The client.json configuration file is described as below:

Parameter	API	Required	Description

	Type		
Client.Extra.IgnoreCheck	Bool	No	The default value is `false` . The default value is false. The migration tool automatically checks the source server environment upon startup by default. To skip the check, set this parameter to `true` .
Client.Extra.Daemon	Bool	No	The default value is `false` . If you need the migration tool to run in the background, set this parameter to `true` .
Client.Net.Proxy.Ip	String	No	The default value is empty. In the private network migration Scenario 3, the IP address of the network proxy needs to be configured.
Client.Net.Proxy.IPv6	Bool	No	It defaults to false. Set it to true if you want to transfer data via IPv6. Otherwise, the migration data will be transferred via IPv4.
Client.Net.Proxy.Port	String	No	The default value is empty. In the private network migration Scenario 3, the port of the network proxy needs to be configured.
Client.Net.Proxy.User	String	No	The default value is empty. In the private network migration Scenario 3, if your network proxy needs to be verified, configure the username of the network proxy.
Client.Net.Proxy.Password	String	No	The default value is empty. In the private network migration Scenario 3, if your network proxy needs to be verified, configure the password of the network proxy.

Note:

Except for the above parameters, other configuration items in the `client.json` file usually don't need to be entered.

rsync_excludes_linux.txt file description

This file is used to exclude files on the Linux source server or configuration files under specified directories that do not need to be migrated. By default, the rsync_excludes_linux.txt file already excludes the following directories and files. **Do not delete or modify the existing configurations.**

```
/dev/*
/sys/*
/proc/*
/var/cache/yum/*
/lost+found/*
```

```
/var/lib/lxcfs/*
/var/lib/docker-storage.btrfs/root/.local/share/gvfs-metadata/*
```

To exclude other directories or files, append them to the rsync_excludes_linux.txt file. For example, to exclude all content on the data disk attached to `/mnt/disk1` , configure the rsync_excludes_linux.txt file as follows:

```
/dev/*
/sys/*
/proc/*
/var/cache/yum/*
/lost+found/*
/var/lib/lxcfs/*
/var/lib/docker-storage.btrfs/root/.local/share/gvfs-metadata/*
/mnt/disk1/*
```

Parameters of the Migration Tool

Parameter	Description
--help	Prints help information.
--check	Checks the source server
--log-file	Configures the log file name, which is log by default.
--log-level	Configures the logging level. Valid values: 1(ERROR level), 2 (INFO level) and 3(DEBUG level). Default value: 2.
--version	Prints the version number.
--clean	Ends the migration task.

Migration Time Estimation

Last updated：2024-01-08 09:37:01

This document describes how to estimate the time for online migrating the system and applications from a source server in your IDC or cloud platform to Tencent Cloud CVM.

The migration time is subject to the data transfer speed during migration. You can estimate it by testing the transfer speed between the source server and destination CVM.

Estimating Migration Time in Different Scenarios

Scenario 1

If the destination type of a migration task is CVM instance, the estimated migration time is mainly the actual data transfer time.

For example, if the size of the data on all disks to be migrated on the source server is 50 GB and the outbound bandwidth is 100 Mbps, the estimated total migration time will be 1.14 hours as calculated below:

1. Convert the unit

Convert the unit of the actual bandwidth to MB/s: 100 Mbps = 100 / 8 = 12.5 MB/s

Convert the unit of the actual disk data volume to MB: 50 GB = 50 * 1024 = 51200 MB

2. Estimate the actual data migration time

51200 / 12.5 = 4096 seconds = 1.14 hours

Scenario 2

If the destination type of a migration task is CVM image, the migration time includes the actual data transfer time and image creation time.

For example, if the size of the data on all disks to be migrated on the source server is 50 GB and the outbound bandwidth is 100 Mbps, the estimated total migration time will be 1.23 hours as calculated below:

1. Convert the unit

Convert the unit of the actual bandwidth to MB/s: 100 Mbps = 100 / 8 = 12.5 MB/s

Convert the unit of the actual disk data volume to MB: 50 GB = 50 * 1024 = 51200 MB

2. Estimate the actual data migration time

51200 / 12.5 = 4096 seconds = 1.14 hours

3. Calculate the image creation time at a speed of about 160 MB/s

51200 / 160 = 320 seconds = 0.089 hour

4. Calculate the total migration time

1.14 + 0.089 = 1.23 hours

Relevant Operations: Testing Data Transfer Speed

You can use the **iperf3** tool to test the data transfer speed, such as bandwidth and speed of data transfer from client to server.

Factors affecting transfer speed

Outbound bandwidth of the source server and inbound bandwidth of the destination instance.

For example, if the outbound bandwidth of the source server is 50 Mbps and the inbound bandwidth of the destination instance is 100 Mbps, the actual transfer speed won't exceed 50 Mbps theoretically.

During the migration, the bandwidth isn't always fully used, and you can dynamically adjust the inbound bandwidth of the destination or relay instance.

If the source server and destination instance are in different regions, the transfer speed will be lower than that when they are in the same region.

Note:

When you use online migration in the console, if the migration destination is a CVM image, the relay instance `do_not_delete_csm_instance` with a bandwidth cap of 50 Mbps will be created during migration.

You can dynamically adjust the inbound bandwidth of the destination or relay instance in the console during migration to control the migration speed.

Speed test for migration to Linux CVM instance

For example, if you use the online migration feature in the console to migrate a server to a CentOS 7.5 CVM instance, you can test the transfer speed in the following steps:

1. Create a pay-as-you-go CentOS 7.5 CVM instance in the migration destination region.

Note:

If the migration destination is a CVM image, a CentOS 7.5 relay instance will be created during migration. To test its speed, we recommend you choose an available Standard model with a low CPU and memory configuration, which is more like the actual migration scenario.

The default port of the iperf3 server is TCP 5201. You need to add it to and open it in the inbound traffic configuration in the security group of the CentOS 7.5 instance.

2. Install iperf3 on the source server and in the testing destination instance respectively.

Run the following command to install iperf3 in the destination CentOS 7.5 instance:

```
yum -y install iperf3
```

Install iperf3 on the source server. Use the installation command corresponding to the Linux distribution on the source server for installation.

3. Run the following command to start iperf3 in the testing destination CentOS 7.5 instance as the server:

```
iperf3 -s
```

If "Server listening on 5201" is returned, the start succeeded.

4. Run the following command to start iperf3 on the source server as the client:

```
iperf3 -c [destination instance IP]
```

The returned test result is as shown below, indicating that the transfer speed between the source server and the test CentOS 7.5 instance is 111 Mbps.

Migration billing instructions

Last updated：2024-01-08 09:37:01

Service migration is free to use. However, it may involve fees for **relay instances** and **network traffic**. This document describes the billable items and billing modes that may be involved when you use service migration.

Relay Instances

If the migration destination is a Cloud Virtual Machine (CVM) image, a relay instance named `do_not_delete_csm_instance` will be created under your account after the migration starts. The instance incurs instance fees and cloud disk fees.

Billing mode: pay-as-you-go

Do not reinstall, shut down, or terminate the relay instance or reset its password. It will be automatically terminated by the system after the migration ends.

Network Traffic

Network traffic is generated during online migration, which is billed as follows:

For migration over the public network, if your source server has a bandwidth plan, no additional fees will incur on the source server. The inbound traffic on the destination CVM does not incur fees.

For migration over the public network, if your source server is billed by traffic, traffic fees will incur on the source server. The inbound traffic on the destination CVM does not incur fees.

For migration through another channel, such as a VPC peering connection, a VPN connection, Cloud Connect Network, or Direct Connect, refer to the billing rules of the corresponding network service.

Offline Migration

Last updated：2024-01-08 09:37:01

This document describes how to migrate your instance and data in an offline manner.

Overview

Supported by Tencent Cloud Service Migration (CSM), the service migration feature lets you migrate operating systems, applications, and application data from a source server to a Cloud Virtual Machine (CVM) instance or Cloud Block Storage (CBS) instance. It helps meet enterprise needs for cloudification, cross-cloud migration, cross-account or cross-region migration, and hybrid cloud deployment.

Service migration provides offline migration and online migration. Offline migration includes:

Migration to CVM allows you to migrate a system disk image (or both system disk image and data disk image if necessary) to a specific CVM instance.

Migration to CBS allows you to migrate a data disk image to a specific cloud disk.

Precautions

Activate Tencent Cloud Cloud Object Storage (COS) and make sure COS is available in your region.

For information on regions currently supported by COS, see Regions and Access Endpoints.

Considerations

Note:

Supported image formats: QCOW2, VHD, VMDK, and RAW. We recommend using the compressed image format to shorten the transmission and migration time.

The image to upload must be stored in the COS bucket in the same region as of the destination CVM instance. In addition, the COS bucket must allow public read access.

If you need to import both the system disk image and data disk images, a corresponding number of data disks must be mounted to the target instance.

The capacity of the target disk should be greater than (as recommended) or equal to that of the source disk.

Snapshot files (such as *-00000*.vmdk) are not supported.

Create an image of the source server.

For Windows, see Creating Windows Images.

For Linux, see Preparing a Linux Image.

Upload the created image to COS.

Because images are large in size, upload using the browser may fail. We recommend using the COSCMD tool to upload images. For more information, see COSCMD.

If images exported from other cloud platforms are compressed packages (such as .tar.gz files), you can upload them directly to COS.

Obtain the COS address of the uploaded image.

Go to the COS console, locate the image file you just uploaded and copy the temporary URL on the image file details page.

Prepare the destination CVM or CBS instance.

Click here to purchase a CVM instance.

Click here to view CBS purchase instructions.

Directions

Migration to CVM

Migration to CBS

1. Log in to the CVM console and click **Service Migration** in the left sidebar.

2. Click **Migrate to CVM** on the **Offline migration** page.

3. In the "Migrate to CVM" pop-up window, complete the preparation, and click **Next**.

4. Select the region and enter the task name, COS link and the CVM instance to migrate to.

5. Click **Complete**.

During the migration, you can quit or close the Service Migration page. You can also return to this page anytime to check the migration progress.

1. Log in to the CVM console and click **Service Migration** in the left sidebar.

2. Click **Migrate to CBS** on the **Offline migration** page.

3. In the **Migrate to CBS** pop-up window, complete the preparation, and click **Next**.

4. Select the region and enter the task name, COS link and the destination CBS instance.

5. Click **Complete**.

During the migration, you can quit or close the Service Migration page. You can also return to this page anytime to check the migration progress.

FAQs

For more information, see Service Migration.

Contact Us

Last updated：2024-01-08 09:37:01

If you encounter any issue during service migration, or have any feedback or suggestions, do not hesitate to contact us.

Submitting a Ticket

If you encounter any operation or technical problems when using our product, you can log in to the Tencent Cloud Console and follow the on-screen prompts to submit a ticket. We will get back to you as soon as possible.

Ticket links:

Submitting a ticket: Submit a ticket

Querying ticket status: Ticket list

Maintenance Tasks

Overview

Last updated：2024-01-08 09:37:01

Maintenance Task is designed to provide users with standardized CVM troubleshooting and authorized maintenance services.

To improve the running performance and stability of instances, and ensure the safe and efficient operation of the underlying platform, we regularly maintain and upgrade the underlying host and platform architecture without CVM shutdown. During the upgrade and maintenance, your CVM instances can operate stably without the need to interrupt the business applications.

Maintenance Task helps users learn and handle all kinds of issues of CVM instances in real time, prevent potential downtime risks in advance, improve maintenance efficiency and reduce maintenance costs. You can back up data of abnormal instances to ensure stable operation of your business. Also, you can configure preset authorization policies or use APIs as needed for automatic Ops of CVM failures and risks.

Advantages

Free enablement

Maintenance Task is now fully available for free. After you create and use a CVM instance, you can go to Task List to check all maintenance task records of your CVM instances.

Full coverage of exceptions and risks

All kinds of sudden exceptions (such as sudden abnormal downtime of underlying host, causing the CVM to abnormally restart), running risks (predict the risks of various software and hardware failures of the underlying host), disk exceptions/warnings (instance disk usage exceptions/ early warnings) and scheduled maintenance and upgrade tasks are covered.

Elastic configuration

Multiple preset authorization policies can be configured based on your own business scenarios and Ops requirements. Each policy can be associated with different instance families, and can be quickly bound through CVM tags.

Flexible authorization

Users can authorize for maintenance through the Maintenance Task console, preset authorization policies and APIs.

Use Cases

Real-time awareness of instance exceptions and quick recovery

Users are notified with all kinds of CVM instance exceptions. Corresponding maintenance tasks are created. You can log in to the Maintenance Task console to check the recovery of the affected instances and avoid risks in time.

Real-time monitoring of risks on instances and avoid in advance

When the CVM instances are currently running normally, but the platform detects that there are software and hardware risks on the underlying host, or there are maintenance tasks planned by the platform for the CVM instances, users can receive relevant information in real time, make maintenance plans, and authorize for maintenance during low-peak business periods to avoid failures in advance and eliminate potential downtime risks.

Automatic Ops for CVM exceptions

Users can authorize for automatic Ops through preset authorization policies and APIs. When a new maintenance task or alarm event is triggered, the failure can be healed with automatic Ops to improve Ops efficiency.

Use Limits

Maintenance Task is currently applicable to CVMs, CDHs and CPMs.

Maintenance Task Type and Processing Policy

Last updated：2024-01-08 09:37:01

When exceptions that affect instance availability and performance are detected, the maintenance process is initiated automatically, record the maintenance tasks, and notify users about the affected instances. Users can go to the Maintenance Task console to check details and authorize Tencent Cloud to perform maintenance.

The maintenance tasks of CVM instances are classified into the following types based on the reasons the tasks are triggered. See below for details:

Maintenance Task Types

Task Type	Description	Suggestion	Applicable Authorization Policies
Instance running exception	Sudden software and hardware failures or system errors of the underlying host of the instance, which cause abnormal downtime or restart of the instance.	When a maintenance task of abnormal instance running is triggered, the platform immediately performs relevant maintenance and tries to restart the abnormal instance.It is recommended to wait for the completion of instance restarting, and check the update progress of maintenance task.	Choose the policy based on the current status of the maintenance task: When the task is in "Processing" status, the platform is urgently performing related maintenance on the abnormal instance. After the maintenance is completed, the task status will be updated immediately, and relevant notifications will be pushed to you. When the task is in "Ended" status, the abnormal instance has automatically restarted and restored. You can verify whether the instance and application have been restored to normal mode.
Instance running risk	The instance is currently running normally, but there are risks on software and hardware of the host	To complete the maintenance as soon as possible to avoid risks of the underlying software and hardware and	According to the fixing method of the underlying risks of the instance, the following authorization methods can be selected:

	or the underlying platform, which may cause the fluctuation of the instance performance or the abnormal downtime.	potential downtime, it is recommended to back up your business data in advance and go to the Maintenance Task console to perform the following operations: 1. (Optional) Back up the instance data. 2. Authorize the platform to initiate maintenance immediately, or reserve a planned maintenance within 48 hours in advance. 3. Wait for the system to automatically initiate maintenance at the scheduled maintenance time.	Authorization for migration without CVM shutdown (the instance does not need to be shut down, and the CVM may experience short-term high load or network jitter during the migration). Authorization for shutdown maintenance (the instance is fast restored after restart). **Note**: 1. If the user does not authorize within 48 hours, the system will initiate maintenance at the scheduled maintenance time. 2. Local disk instances do not support fast restoration after restart, and require a longer maintenance period to fix underlying hardware risks. Users can choose to redeploy local disk instances to quickly avoid the risks (local disk data cannot be retained).
Instance disk exception	A sudden failure occurs on the local disk, which may cause reduced I/O performance of the instance or damage to the disk.	To complete the maintenance as soon as possible to restore the disk, it is recommended to back up your business data in advance and go to the Maintenance Task console to perform the following operations: 1. (Optional) Back up the instance data. 2. Authorize the platform to change the abnormal disk immediately, or reserve a planned maintenance within 48 hours in advance. 3. Wait for the platform to replace the abnormal disk, and reattach and use the replaced disk	According to the fixing method of the abnormal disk, the following authorization methods can be selected: Change disk without CVM shutdown (replace the abnormal disk without CVM shutdown. During the maintenance, the I/O of the abnormal disk is temporarily unavailable. After the maintenance is completed, you can attach and use the new disk). Shut down to change disk (the instance needs to be shut down to replace the abnormal disk. The local disk data may be retained. A long maintenance period is required). (Optional) Migrate without the disk: The local disk instance is

		according to the prompts in the restoration notification.	redeployed, and the local disk data cannot be retained. The instance availability can be restored in minutes.
Instance disk warning	The local disk of the instance may be damaged, or its service life is about to end, which may cause instance I/O exceptions or disk offline.	To complete the maintenance as soon as possible to eliminate the potential failure risks of the local disk, it is recommended to back up your business data in advance and go to the Maintenance Task console to perform the following operations: 1. (Optional) Back up the instance data. 2. Authorize the platform to change the disk with potential failure risks immediately, or reserve a planned maintenance within 48 hours in advance. 3. Wait for the platform to replace the abnormal disk, and reattach and use the replaced local disk according to the prompts in the restoration notification.	According to the fixing method of the abnormal disk, the following authorization methods can be selected: Change disk without CVM shutdown (replace the abnormal disk without CVM shutdown. During the maintenance, the I/O of the abnormal disk is temporarily unavailable. After the maintenance is completed, you can attach and use the new disk). Shut down to change disk (the instance needs to be shut down to replace the abnormal disk. The local disk data may be retained. A long maintenance period is required). (Optional) Migrate without the disk: The local disk instance is redeployed, and the local disk data cannot be retained. The instance availability can be restored in minutes.
Instance network connection exception	A sudden failure occurs at the underlying network connection of the instance, which may cause network jitter or abnormal network connection.	When a maintenance task of an abnormal instance network connection is triggered, the platform immediately performs relevant maintenance on the underlying network and tries to restore the network connection of the abnormal instance. It is recommended to wait for the completion of the automatic fixing of the	Choose the policy based on the current status of the maintenance task: When the task is in "Processing" status, the platform is urgently performing related maintenance on the underlying network of the abnormal instance. After the maintenance is completed, the task status will be updated immediately, and relevant notifications will be pushed to you.

		network connection, and check the update progress of the maintenance task.	When the task is in "Ended" status, the network connection of the abnormal instance has been recovered. You can verify whether the instance and application have been restored to normal mode.
Instance maintenance and upgrade	Maintenance without CVM shutdown is initiated due to reasons such as underlying host architecture and software upgrades to improve instance performance and security.	To complete maintenance as soon as possible to improve instance performance and security, it is recommended to back up your business data in advance, and go to the Maintenance Task console to perform the following operations: 1. (Optional) Back up the instance data. 2. Authorize the platform to initiate maintenance immediately, or reserve a planned maintenance within 48 hours in advance. 3. Wait for the system to automatically initiate maintenance at the scheduled maintenance time.	You can choose from the following authorization methods: Maintenance without CVM shutdown (the instance does not need to be shut down, and the CVM may experience short-term high load or network jitter during the maintenance). **Note**: If the user does not authorize within 48 hours, maintenance starts at the next scheduled maintenance time.

Task Status

Task Status	Description
Pending authorization	Wait for user authorization. The user can choose the maintenance method and time. If the user does not authorize for a non-disk maintenance task within 48 hours, the system will initiate maintenance at the scheduled maintenance time and the maintenance task status will be changed into "Processing".
Scheduled	The user has authorized for maintenance and reserved a maintenance time. The default scheduled maintenance time can be modified within 48 hours after the task is created.

Processing	The maintenance task is being processed.
Ended	The maintenance task is completed.
Avoided	If the instance has an unfinished maintenance task, when the user returns or terminates the instance, or adjusts the instance configuration, the avoidance of the maintenance task will be interrupted.
Canceled	The maintenance task is canceled by the system.

Viewing Maintenance Task

Last updated：2024-01-08 09:37:01

Overview

This document describes how to view the lists of pending and historical maintenance tasks and the detailed troubleshooting information in the CVM console.

Directions

1. Log in to the CVM console and select **Maintenance Task** > **Task List** on the left sidebar.
2. On the **Maintenance Task** list page, select the filter conditions above the list to get the list of the required maintenance tasks.
3. Click the ID of a maintenance task to view more information on the task details page.

Authorizing Maintenance Policy and Scheduling Maintenance Time

Last updated：2024-01-08 09:37:01

Overview

This document describes how to select a specific maintenance policy and schedule the maintenance time for a maintenance task in the CVM console.

Directions

1. Log in to the CVM console and select **Maintenance Task** > **Task List** on the left sidebar.
2. Click **Authorize**/**Schedule** on the right of the row of the target maintenance task.
3. In the pop-up window, specify the authorized maintenance method and scheduled maintenance time.
Note:
The authorized maintenance method is determined by the task type. For more information, see Maintenance Task Type and Processing Policy.
If the **scheduled maintenance time** is not specified, maintenance will start immediately by default.
4. Click **OK**.

Configuring Preset Authorization Policy

Last updated：2024-01-08 09:37:01

Overview

This document describes how to configure a preset maintenance authorization policy in the CVM console. You can set such a policy for all CVM instances under a tag. If a maintenance task is generated, it will be processed according to the configured preset policy with no need for your manual authorization.

Directions

1. Log in to the CVM console and select **Maintenance Task** > **Preset Authorization Policy** on the left sidebar.

2. On the **Preset Authorization Policy** page, click **Create**.

3. In the **Create Preset Authorization Policy** pop-up window, select the specific product type, metric, and policy for preset authorization and associate tags.

4. Click **OK**. After an instance associated with a set tag generates a maintenance task, the preset policy will be used by default for maintenance.

Configuring Maintenance Task Alarm Notification

Last updated：2024-01-08 09:37:01

Overview

This document describes how to use EventBridge to set alarm notification for a CVM instance in the EventBridge console. You can set alarm notification for the maintenance tasks of a CVM instance, so that you will be notified immediately that you need to take countermeasures through channels such as email, SMS, and phone call when an exception occurs.

Directions

1. Log in to the EventBridge console and activate the service as instructed in Activating EventBridge.
2. Select **Event Rule** on the left sidebar, select the target region and event bus at the top of the **Event Rule** page, and click **Create Event Rule**.
3. On the **Create Event Rule** page, perform the following operations:
3.1 In **Basic Information**, set the **Rule Name** parameter as shown below:

3.2 In **Event Pattern**, set **Event Matching** parameters as needed as shown below:

Tencent Cloud service: Select **CVM** from the drop-down list.

Event Type: Select an option as needed from the drop-down list.

3.3 Click **Next**.

3.4 In **Delivery target**, select an option as needed from the **Trigger method** drop-down list.

Select **CLS** for **Trigger method**. For more information, see CLS Log Target.

Select **Notification message** for **Trigger method**. For more information, see Message Push Target.

4. Click **Complete**.

Cloud Disks

Expanding Cloud Disks

Last updated：2024-01-08 09:37:01

Scenarios

A cloud disk is an expandable storage device on cloud. You can expand its capacity at any time without losing any data in it.

After expanding the cloud disk, you need to expand the partition and file system. You can allocate the capacity of the expanded part to an existing partition or format it into a new partition.

Note:

MBR partition supports disk with a maximum capacity of 2 TB. When you partition disk with a capacity greater than 2 TB, we recommend that you create and mount a new data disk and use the GPT partition format to copy data.

Expanding Data Disks

If the cloud disk is a data disk, you can expand it using the following three methods.

Note:

If multiple cloud disks of the same capacity and type are attached to the CVM, you can identify them using the method shown in Distinguishing data disks. Select a data disk and expand its capacity as instructed below.

Expand in the CVM console (recommended)

Expand in the CBS console

Expand via an API

1. Log into the CVM Console.

2. Select **More** > **Resource Adjustment** > **Expand Cloud Disks** in the **Operation** column.

3. Select the data disk to be expanded in the pop-up window, and click **Next**.

4. Select a new capacity (it must be greater than or equal to the current capacity) and click **Next**.

5. Read the notes and click **Adjust Now**, as shown below:

6. Assign its expanded capacity to an existing partition, or format it into an independent new partition. Depending on the operating system of the CVM, see Extending Partitions and File Systems (Windows) or Determining the Expansion Method.

1. Log in to the CBS Console.

2. Select **More** > **Expand** for the target cloud disk.

3. Select a new capacity. It must be greater than or equal to the current capacity.

4. Complete the payment.

5. Assign its expanded capacity to an existing partition, or format it into an independent new partition. Depending on the operating system of the CVM, see Extending Partitions and File Systems (Windows) or Determining the Expansion Method.

You can use the `ResizeDisk` API to expand the specified cloud disks. For more information, see ResizeDisk.

Expanding System Disks

1. Log in to the CVM console. Locate the target CVM, and select **More** > **Resource Adjustment** > **Expand Cloud Disks** in the **Operation** column.

2. Select the system disk to expand in the pop-up window, and click **Next**.

3. Select a new capacity (it must be greater than or equal to the current capacity) and click **Next**.

4. Expand the cloud disk as instructed below.

Expand in the CVM console

Expand in the CBS console

Expand via an API

Note:

CVM supports expanding a cloud system disk without shutting down the instance.

1. In the **Expand partition and file system** tab, read the notes and click **Adjust Now**.

2. Complete the capacity expansion in the console and log in to the instance to check whether the file system has been extended automatically. If not, extend the partition and file system as instructed in Extending System Disk and File System Online.

1. Log in to the CBS Console.

2. Select **More** > **Expand** for the target cloud disk.

3. Select a new capacity. It must be greater than or equal to the current capacity.

4. Complete the payment.

5. Assign its expanded capacity to an existing partition, or format it into an independent new partition. Depending on the operating system of the CVM, see Extending Partitions and File Systems (Windows) or Determining the Expansion Method.

You can use the `ResizeInstanceDisks` API to expand the non-elastic disks. For more information, see ResizeInstanceDisks.

Relevant operations

Distinguishing data disks

Check cloud disks according to the operating system of the CVM.

Linux

Windows

1. Log in to a Linux instance.

2. Run the following command to view the relationship between the elastic cloud disks and the device name.

```
ls -l /dev/disk/by-id
```

The following information will appear:

Note that `disk-xxxx` is the ID of a cloud disk. You can use it to view cloud disk details on the CBS console.

1. Log in to a Windows instance.

2. Right-click

, and select **Run**.

3. Enter `cmd` in the pop-up window and press **Enter**.

4. Run the following command to view the relationship between the elastic cloud disks and the device name.

```
wmic diskdrive get caption,deviceid,serialnumber
```

You can also run the following command.

```
wmic path win32_physicalmedia get SerialNumber,Tag
```

The following information will appear:

Note that `disk-xxxx` is the ID of a cloud disk. You can use it to view cloud disk details on the CBS console.

Checking the cloudinit configuration

Check cloud disks according to the operating system of the CVM.

Linux instance

Windows instance

After the expansion, log in to the Linux instance and check whether the `/etc/cloud/cloud.cfg` file contains the `growpart` and `resizefs` configuration items.

If yes, ignore other operations.

```
cloud_init_modules:
 - migrator
 - bootcmd
 - write-files
 - growpart
 - resizefs
 - set_hostname
 - update_hostname
 - ['update_etc_hosts', 'once-per-instance']
 - rsyslog
 - users-groups
 - ssh
```

growpart: Expands the partition to the disk size.

resizefs: Expands or adjusts the file system in the `/` partition to the partition size.

If no, manually extend partitions and file systems (Linux) according to the operating system, and assign its extended capacity to an existing partition, or format it into an independent new partition.

After the system disk is expanded, log in to the Windows instance and check whether the `ExtendVolumesPlugin` configuration item under `plugin` exists in `C:\\Program Files\\Cloudbase Solutions\\Cloudbase-Init\\conf\\cloudbase-init.conf`.

If yes, reboot the machine. `cloudbase-init` will automatically extend the volume, adding the empty space behind the C partition to the C partition. Note that there must be no other partition between partition C and the blank space. If you don't want to reboot the machine, there are other partitions between the C partition and the blank space, or `cloudbase-init` is blocked by third-party security software, you need to execute the following powershell command manually.

```
$DiskOps="@
select disk 0
select volume c
extend
exit
@"
$DiskOps | diskpart.exe | Out-Null
```

If no, manually extend partitions and file systems (Windows) according to the operating system, and assign its extended capacity to an existing partition, or format it into an independent new partition.

Changing Disk Media Type

Last updated：2024-01-08 09:37:01

Overview

Tencent Cloud CVM supports adjusting the storage hardware media, which enables you to flexibly respond to diversified storage needs of different businesses.

Tencent Cloud provides two types of storage media: Cloud Block Storage and Local Storage. A local disk can be converted to a cloud disk. This document describes how to change the disk media type.

The downsides of CVMs with local disks are as follows:

The configuration cannot be customized due to the limit of host resources.

Features such as snapshots and creation acceleration are not supported.

Low data reliability.

Host failures will have a longer impact.

To avoid these downsides, you can convert the local disks attached to your CVMs to cloud disks.

Prerequisites

CVM Status

Make sure that the related CVM is shut down.

Unsupported CVM Types

Spot instances

Big data and high I/O models

Bare metal instances

CVM Configuration

At least one of the CVM's system disk and data disks must be **HDD** or **SSD local disk**.

There are available cloud disks that have matched size with local disks in the availability zone where the CVM resides.

The adjustment will convert **all** of the local disks to cloud disks if both the system disk and data disks of the CVM are local disks. You will also be able to configure the cloud disk type for each disk separately.

In other words, the disk media change of a CVM whose disks are all local disks applies to all of its disks, rather than only system disk or data disks.

Changing the cloud disk media will not resize the disk. You can expand the cloud disk after changing the media type.

This operation will not change the lifecycle of a CVM, instance ID, private/public IP, disk name, and mount point.

Notes

This conversion needs to copy all the data from the local disk to the cloud disk. Depending on the disk size and transmission speed, this could take some time.

You can only convert local disks to cloud disks. The conversion CANNOT be reverted.

After the adjustment, we recommend you to start up and log in to the CVM to check the data integrity.

Directions

1. Log in to the CVM console and access the **Instances** page.

Note:

 If the CVM has already been shut down, go to Step 3.

2. (Optional) Locate the target CVM, and click **More** > **Instance Status** > **Shutdown** under the **Operation** column to shut it down.

3.

Under the **Operation** column, click **More** > **Resource Adjustment** > **Change Disk Media Type**.

4. In the pop-up window, select the target cloud disk type, check **I have read and agreed to Rules for Changing Disk Media Type**, and click **Change Now**.

5. Double-check the information, make a payment if applicable, and wait for the process to complete.

Adjusting Cloud Disk Performance

Last updated : 2025-06-04 16:19:24

The performance of a cloud disk depends on its capacity. You can improve the performance by adjusting its capacity till it reaches the ceiling. When the ceiling is reached, you can purchase extra performance to get even higher performance. Note that the extra performance is only available for enhanced SSD instances. For more information, see Enhanced SSD Performance.

Caution:

Currently, only **Enhanced SSD** supports independent performance adjustment.

The extra performance can be independently adjusted only after the basic performance reaches the ceiling.

The performance adjustment will not affect the running of your cloud disks and businesses.

Performance Adjustment Billing

Upgrading

For pay-as-you-go cloud disks, the performance upgrade takes effect immediately, and the cloud disks are charged by the new configuration right away.

Downgrading

For pay-as-you-go cloud disks, the performance upgrade takes effect immediately, and the cloud disks are charged by the new configuration right away.

Performance Upgrade

Upgrading a disk via the console

When prerequisites are met, you can upgrade a disk as instructed below in the console:

1. Log in to the CBS console.

2. Select the region and the cloud disk that requires performance adjustment.

3. Click **More** > **Adjust Performance** under the **Operation** column of the selected cloud disk.

4. Select a target configuration in the pop-up window.

5. Read and confirm the notes and start the adjustment.

Upgrading a disk via API

You can also use the `ModifyDiskExtraPerformance` API to upgrade a specified cloud disk. For detailed directions, see ModifyDiskExtraPerformance.

Performance Downgrade

Downgrading a disk via the console

When prerequisites are met, you can downgrade a disk as instructed below in the console:

1. Log in to the CBS console.

2. Select the region and the cloud disk that requires performance adjustment.

3. Click **More** > **Adjust Performance** under the **Operation** column of the selected cloud disk.

4. Select a target configuration in the pop-up window.

5. Read and confirm the notes and start the adjustment.

Downgrading a disk via API

You can also use the `ModifyDiskExtraPerformance` API to downgrade a specified cloud disk. For detailed directions, see ModifyDiskExtraPerformance.

Networking
MTU

Last updated：2024-07-31 14:49:33

Maximum Transmission Unit (MTU)

The MTU determines the maximum size of a data packet that can be transmitted at one time on the network, including the IP packet header and payload, but excluding the Ethernet header. In theory, the larger the MTU, the more data can be transferred in a single packet, resulting in more efficient network communication. This document describes features related to Jumbo Frames.

Jumbo Frame

A Jumbo Frame refers to an Ethernet frame with a payload exceeding 1500 bytes. The increased payload percentage helps improve link utilization rate and achieve better network performance. Tencent Cloud supports Jumbo Frames up to 8500 bytes, allowing you to send Ethernet frames with an 8500-byte payload.
The latest generation instances (SA5/S8/IA5se/ITA5/M8/MA5) have enabled the Jumbo Frame capability by default, setting the MTU to 8500. For more information, refer to Instance Specifications.
Note:
Network communication using Jumbo Frames requires considering the network link and peer host's MTU limit. Packets exceeding the MTU limit will be sharded or discarded. It is recommended to test if there are connectivity and performance issues due to MTU mismatches in business scenarios before deploying services within CVM using Jumbo Frames, and manage the packet size in these scenarios to reduce unnecessary performance loss.
Disabling the Jumbo Frame feature is in grayscale. If you need to disable it by definition, Submit a ticket **for support.**

CVM Intercommunication Link MTU Limit

When CVM instances supporting Jumbo Frames are used as the source host, the link maximum MTU limit for the destination host in different communication scenarios is shown in the table below:

Destination Host	Allowed Maximum MTU	Description

Within the same VPC, CVM instances with Jumbo Frames enabled	8500	No
Within the VPCs interconnected through CCN, CVM instances with Jumbo Frames enabled		The Jumbo Frame capability needs to be enabled on the CCN. CCN's support for the Jumbo Frame capability is in beta testing. If needed, you can submit a ticket to apply.
Hosts within the offline IDC connected through dedicated channels		The Jumbo Frame capability needs to be enabled on the Direct Connect Gateway. Direct Connect Gateway's support for the Jumbo Frame capability is in beta testing. For details, see Change Channel.
Within the same VPC, CVM instances without enabling Jumbo Frames	1500	No
Within the VPCs interconnected through CCN, CVM instances without enabling Jumbo Frames		No
Hosts within the offline IDC connected through dedicated channels		No
Access external hosts through EIP		No

Cloud Network Product MTU Limit

When accessing specific network products using non-TCP protocols (UDP and ICMP), pay attention to the MTU limitations of different network products. Tencent Cloud is currently improving and perfecting the scenarios for supporting Jumbo Frames. The maximum MTU limitations of some network products are as follows:

Network Product	Supported Maximum MTU	Description
NAT Gateway	1500	No
CLB (including ALB, NLB, and CLB)	1500	No
Direct Connect Gateway	8500	The Jumbo Frame capability needs to be enabled on the Direct

		Connect Gateway. Direct Connect Gateway's support for the Jumbo Frame capability is in beta testing. For details, refer to [Change Channel](#).
Cloud Connect Network	8500	The Jumbo Frame capability needs to be enabled on the CCN. CCN's support for the Jumbo Frame capability is in beta testing. If needed, you can [Submit a Ticket](#) to apply.
Peering connection	1500	No

Note:

For data packets exceeding the maximum MTU limitations of cloud products, Tencent Cloud's forwarding components will shard or negotiate the message. This ensures connectivity only but not network communication performance. Therefore, in communication scenarios where the maximum MTU of cloud products is 1500, avoid using Jumbo Frames for UDP and ICMP communication.

FAQs

Problem Description: When using Jumbo Frames for UDP and ICMP communication, how to resolve the situation if there is a sharp performance decline?

Solution: Check for data packet sharding and disable Jumbo Frames on CVM instances if necessary.

If the following scenarios are involved, it is recommended to disable the Jumbo Frame capability on CVM instances:

Scenario One: If there is a need for interconnection with other public cloud providers through the public network, it is recommended to disable the Jumbo Frame capability on CVM instances.

Scenario Two: If there is a need for interconnection with other cloud providers through a dedicated line or VPN, it is recommended to disable the Jumbo Frame capability on CVM instances.

If the following scenarios require enabling Jumbo Frames, ensure that all gateway devices along the link are enabled for Jumbo Frames:

Scenario One: When multiple gateway devices are passed, all gateway devices along the link need to enable Jumbo Frames, such as in CCN + dedicated line scenarios.

Scenario Two: When multiple subnets form ECMP routing, all gateway devices need to enable Jumbo Frames, for example, two dedicated line gateways form ECMP.

How to calculate the appropriate MTU value when data is encapsulated in a tunnel?

Scenario: How to calculate a reasonable MTU when establishing a tunnel between the customer's IDC and the cloud CVM?

Calculation method: 1500 - {Tunnel encapsulation overhead} - 36 bytes (GRE encapsulation overhead). The tunnel encapsulation overhead should be calculated based on the actual tunnel type used.

Common Public IP

Last updated：2024-01-08 09:37:01

Overview

This document describes how to use the common public IP address. The common public IP can only be assigned when you purchase a CVM and cannot be unbound from the CVM.

Note:

For traditional accounts, when unbinding the EIP from CVM, each account can reallocate a common public IP 10 times per day for free.

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

Directions

You can use the following common public IP features:

Feature	Overview	Documentation
Recovering public IP addresses	If you release or return a public IP (including EIP and common public IP) by mistake, you can recover it in the console, and the recovered public IP will be an EIP.	-
Converting common public IPs to EIPs	You can convert a common public IP of CVM to an EIP. After conversion, the EIP can be bound to and unbound from CVM at any time, making it easier to manage the public IP.	-
Changing the public IP	Change the common public IP of the CVM and release the original public IP.	Changing the public IP in the CVM console
Adjusting the network bandwidth	Adjust the bandwidth or billing mode as needed. This feature will take effect in real time.	Adjusting Network Configuration

Elastic IP

Last updated：2024-04-10 14:13:09

Scenario

Elastic IP (EIP) is a static IP designed for dynamic cloud computing and a fixed public IP in a certain region. With EIP, you can quickly remap an address to another instance in your account or NAT gateway instance to avoid instance failure. This document describes how to use EIPs.

Prerequisites

You have logged in to the CVM Console.

Directions

Apply for EIPs

1. In the left sidebar, click Public IP to enter the EIP management page.
2. Click **Apply** on the EIP management page.
3. In the pop-up **Apply for EIP** window, select the region, IP address type, billing method and bandwidth limit, and enter the number of EIPs you want to apply for.

 Note:

A standard account is taken as an example below. If you are unsure of your account type, see Account Types.

Parameter	Description
IP Address Type	Tencent Cloud supports various types of EIPs, such as general BGP IP, premium BGP IP, accelerated IP, static single-line IP, and anti-DDoS EIP. General BGP IP: The domestic multi-line BGP network covers more than twenty ISPs (including the three major ISPs, CERNET, and China Broadnet). The BGP public network outbound supports switchover across regions within seconds, providing your users with high-speed and secure networks. Premium BGP IP: Dedicated lines can avoid the use of international ISP services. The latency is lower, which effectively improves the quality of overseas services for users in Chinese Mainland. Accelerated IP: Anycast is used for acceleration to ensure more stable and reliable public network access with a low latency. Static single-line IP: Users can access the public network using services of a single ISP, featuring low cost and convenient scheduling.

	Anti-DDoS EIP: This type of GBP IP provides Tbps-level cloud-native DDoS protection capability and should be used together with Anti-DDoS Pro for Enterprise. After the IP address is bound to business resources and Anti-DDoS Pro resources, uses can enjoy the anti-DDoS capability.
IP Resource Pool	If certain adjacent IP addresses need to be reserved for your services or IP addresses in a specific network segment should be allocated, you can submit a ticket for consultation. A dedicated IP resource pool will be assigned for you. Dedicated resource pools are supported for general and premium BGP IPs and static single-line IPs currently. For the fees, please consult your business manager.
Billing Mode	General BGP IPs support billing by traffic and bandwidth package. For details, see Public Network Fees. Premium BGP IPs support billing by bandwidth package. For details, see Public Network Fees and Bandwidth Packages for Premium BGP IPs in Billing Overview. Accelerated IPs, static single-line IPs, and anti-DDoS EIPs support only billing by bandwidth package. Other billing modes are not supported now.
Bandwidth Cap	Set the bandwidth cap based on your needs and allocate bandwidth resources reasonably.
Amount	Determine the amount of EIPs to be applied as needed and make sure that the amount does not exceed the total quota. For details, see Quota Limit in Usage Restrictions.
Name	Specify the EIP instance name. This parameter is optional.
Tag	You can add a tag for permission management.

4. Click **OK** to complete the EIP application.

5. After the application is completed, you can see in the list the EIP you have applied for, which is in an unbound status.

Bind EIPs to cloud products

1. In the left sidebar, click **Public IP** to enter the EIP management page.

2. In the EIP management page, select the EIP which you want to bind to a cloud product and click **More** > **Bind**.

Note:

If the EIP has been bound to a instance, please unbind it first.

3. In the pop-up "Bind resources" window, select the resource to be bound to the EIP and click **OK**.

4. In the pop-up window, click **OK** to complete binding the EIP to the cloud product.

Unbind EIPs from cloud products

1. In the left sidebar, click **Public IP** to enter the EIP management page.

2. In the EIP management page, select the EIP which you want to unbind from the cloud product and click **More** >
Unbind.

3. In the pop-up "Unbind EIP" window, confirm the unbinding information and click **OK**.

4. In the pop-up window, click **OK** to complete unbinding the EIP from the cloud product.

Note:

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

Release EIPs

1. In the left sidebar, click **Public IP** to enter the EIP management page.

2. In the EIP management page, select the EIP which you want to release from the cloud product and click **More** >
Release.

3. In the pop-up "Are you sure you want to release the selected EIPs?" window, select **Release the above EIPs** and
click **Release**.

Adjust Bandwidth

1. In the left sidebar, click **Public IP** to enter the EIP management page.

2. Select the EIP whose bandwidth needs to be adjusted and click **Adjust network**

3. In the pop-up "Change bandwidth" window, configure the bandwidth value and click **OK** to complete the
adjustment.

Convert a public IP to an EIP

The public IP purchased along with the CVM instance is not elastic and cannot be mounted or unmounted. Tencent
Cloud allows you to convert the public IP to an EIP by the following steps:

1. In the left sidebar, click **instances** to enter the instance management page.

2. Select the instance whose public IP needs to be converted to an EIP and then click

, as shown below:

3. In the pop-up "Convert to EIP" window, click **OK**.

Troubleshoot Exceptions

Network inaccessibility may occur with an EIP due to the following reasons:

The EIP is not bound to any cloud product. For more information about how to bind an EIP to the cloud product, please see bind EIPs 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, such as access to 8080 port is denied, the port 8080 of the EIP is also inaccessible.

ENI

Last updated：2024-01-08 09:37:01

To configure ENIs for your CVM, following these instructions:

1. Create an ENI.

View the ENI you just created.

2. Bind the ENI to your CVM and configure it.

3. Configure the CVM and VPC route table.

4. Assign a private IP.

4.1 Log in to Virtual Private Cloud Console.

4.2 Click **ENI** under **IP and ENI** in the left sidebar. The ENI page appears.

4.3 Click the **ID**/**Name** of an ENI to see its details.

4.4 Click **IP Management** to go to the details page.

4.5 Click **Assign private IP** to assign a private IP to the ENI. If you do this manual, pick a usable private IP. Click **OK**.

5. Manage the ENI.

Releasing private IPs

Unbinding CVMs

Deleting ENIs

Binding EIPs

Unbinding EIPs

Modifying primary private IP

Changing the subnet of an ENI

Configuring a Public Gateway

Last updated：2024-01-08 09:37:01

Warning:

Using a single CVM instance as the public gateway has the risk of single point of failure. You are advised to use a NAT gateway **in the production environment.**

As of December 6, 2019, Tencent Cloud no longer supports configuring a CVM as the public gateway on the CVM purchase page. If you need to configure a gateway, follow the instructions below.

Overview

You can access the internet by using a public gateway CVM with a public IP or EIP when some of your VPC-based CVMs lack the public IPs. The public gateway CVM translates the source IP for outbound traffic. When other CVMs access the internet through the public gateway CVM, the source IPs will be translated into public IP of the public gateway CVM. See the figure below.

Prerequisites

You are logged in to the CVM console.

The public gateway CVM and the CVMs that need to access the internet through the public gateway CVM must be in different subnets because the public gateway CVM can only forward requests from other subnets.

The public gateway CVM must be a Linux CVM. Windows CVMs will not work.

Directions

Step 1: bind an EIP (optional)

Note:

Skip this step if the public gateway CVM already has a public IP.

1. Log in to the CVM console and choose **EIP** on the left sidebar.

2. Locate the EIP to bind the instance, select **More** > **Bind** in the **Operation** column.

3. In the pop-up window, select a CVM to be configured and bind it to the EIP.

Step 2: configure a route table for the gateway subnet

Note:

The gateway subnet and other subnets cannot share the same route table. You need to create a separate route table for the gateway subnet.

1. Create a custom route table.

2. Associate the route table with the subnet where the public gateway CVM resides.

Step 3: configure a route table for the other subnets

This route table directs all traffic from the CVMs without a public IP to the public gateway so these CVMs can access public networks as well.

Add the following routing policies to the route table:

Destination: the public IP you want to access.

Next hop type: CVM.

Next hop: private IP of the CVM instance to which the EIP is bound in Step 1.

For more information, see Manage Route table.

Step 4: configure the public gateway

1. Log in to the public gateway CVM instance and perform the following operations to enable the network forwarding and NAT proxy features:

1.1 Run the following command to create the `vpcGateway.sh` script in `usr/local/sbin`.

```
vim /usr/local/sbin/vpcGateway.sh
```

1.2 Press **i** to switch to the edit mode and add the following code to the script.

```
#!/bin/bash
echo "-------------------------------------------------"
echo " `date`"
echo "(1)ip_forward config......"
file="/etc/sysctl.conf"
grep -i "^net\\.ipv4\\.ip_forward.*" $file &>/dev/null && sed -i \\
's/net\\.ipv4\\.ip_forward.*/net\\.ipv4\\.ip_forward = 1/' $file || \\
echo "net.ipv4.ip_forward = 1" >> $file
echo 1 >/proc/sys/net/ipv4/ip_forward
[ `cat /proc/sys/net/ipv4/ip_forward` -eq 1 ] && echo "-->ip_forward:Success"
|| \\
echo "-->ip_forward:Fail"
echo "(2)Iptables set......"
iptables -t nat -A POSTROUTING -j MASQUERADE && echo "-->nat:Success" || echo
"-->nat:Fail"
iptables -t mangle -A POSTROUTING -p tcp -j TCPOPTSTRIP --strip-options
timestamp && \\
echo "-->mangle:Success" || echo "-->mangle:Fail"
echo "(3)nf_conntrack config......"
echo 262144 > /sys/module/nf_conntrack/parameters/hashsize
[ `cat /sys/module/nf_conntrack/parameters/hashsize` -eq 262144 ] && \\
echo "-->hashsize:Success" || echo "-->hashsize:Fail"
echo 1048576 > /proc/sys/net/netfilter/nf_conntrack_max
```

```
[ `cat /proc/sys/net/netfilter/nf_conntrack_max` -eq 1048576 ] && \\
echo "-->nf_conntrack_max:Success" || echo "-->nf_conntrack_max:Fail"
echo 10800 >/proc/sys/net/netfilter/nf_conntrack_tcp_timeout_established \\
[ `cat /proc/sys/net/netfilter/nf_conntrack_tcp_timeout_established` -eq 10800
] \\
&& echo "-->nf_conntrack_tcp_timeout_established:Success" || \\
echo "-->nf_conntrack_tcp_timeout_established:Fail"
```

1.3 Click **Esc** and enter **:wq** to save and close the file.

1.4 Run the following command to set the script permission.

```
chmod +x /usr/local/sbin/vpcGateway.sh
echo "/usr/local/sbin/vpcGateway.sh >/tmp/vpcGateway.log 2>&1" >> /etc/rc.local
```

2. Set the RPS of the public gateway.

2.1 Run the following command to create the `set_rps.sh` script in `usr/local/sbin`.

```
vim /usr/local/sbin/set_rps.sh
```

2.2 Press **i** to switch to the edit mode and add the following code to the script.

```
# !/bin/bash
echo "------------------------------------------"
date
mask=0
i=0
total_nic_queues=0
get_all_mask() {
local cpu_nums=$1
if [ $cpu_nums -gt 32 ]; then
mask_tail=""
mask_low32="ffffffff"
idx=$((cpu_nums / 32))
cpu_reset=$((cpu_nums - idx * 32))
if [ $cpu_reset -eq 0 ]; then
mask=$mask_low32
for ((i = 2; i <= idx; i++)); do
mask="$mask,$mask_low32"
done
else
for ((i = 1; i <= idx; i++)); do
mask_tail="$mask_tail,$mask_low32"
done
mask_head_num=$((2 ** cpu_reset - 1))
mask=$(printf "%x%s" $mask_head_num $mask_tail)
fi
else
```

```
mask_num=$((2 ** cpu_nums - 1))
mask=$(printf "%x" $mask_num)
fi
echo $mask
}
set_rps() {
if ! command -v ethtool &>/dev/null; then
source /etc/profile
fi
ethtool=$(which ethtool)
cpu_nums=$(cat /proc/cpuinfo | grep processor | wc -l)
if [ $cpu_nums -eq 0 ]; then
exit 0
fi
mask=$(get_all_mask $cpu_nums)
echo "cpu number:$cpu_nums mask:0x$mask"
ethSet=$(ls -d /sys/class/net/eth*)
for entry in $ethSet; do
eth=$(basename $entry)
nic_queues=$(ls -l /sys/class/net/$eth/queues/ | grep rx- | wc -l)
if (($nic_queues == 0)); then
continue
fi
cat /proc/interrupts | grep "LiquidIO.*rxtx" &>/dev/null
if [ $? -ne 0 ]; then # not smartnic
#multi queue don't set rps
max_combined=$(
$ethtool -l $eth 2>/dev/null | grep -i "combined" | head -n 1 | awk '{print
$2}'
)
#if ethtool -l $eth goes wrong.
[[ ! "$max_combined" =~ ^[0-9]+$ ]] && max_combined=1
if [ ${max_combined} -ge ${cpu_nums} ]; then
echo "$eth has equally nic queue as cpu, don't set rps for it..."
continue
fi
else
echo "$eth is smartnic, set rps for it..."
fi
echo "eth:$eth queues:$nic_queues"
total_nic_queues=$(($total_nic_queues + $nic_queues))
i=0
while (($i < $nic_queues)); do
echo $mask >/sys/class/net/$eth/queues/rx-$i/rps_cpus
echo 4096 >/sys/class/net/$eth/queues/rx-$i/rps_flow_cnt
i=$(($i + 1))
done
```

```
done
flow_entries=$((total_nic_queues * 4096))
echo "total_nic_queues:$total_nic_queues flow_entries:$flow_entries"
echo $flow_entries >/proc/sys/net/core/rps_sock_flow_entries
}
set_rps
```

2.3 Click **Esc** and enter **:wq** to save and close the file.

2.4 Run the following command to set the script permission.

```
chmod +x /usr/local/sbin/set_rps.sh
echo "/usr/local/sbin/set_rps.sh >/tmp/setRps.log 2>&1" >> /etc/rc.local
chmod +x /etc/rc.d/rc.local
```

3. Restart the public gateway CVM to apply the configurations. Then, test if a CVM without a public IP can access the Internet through the public gateway CVM.

Switching to VPC

Last updated：2024-03-26 14:58:42

Overview

Tencent Cloud provides the classic network and VPC for different scenarios. Various features are offered to help you flexibly manage your networks.

Switching between networks:

Switching from the classic network to VPC: Tencent Cloud allows you to migrate one or more CVM instances from the classic network to VPC at a time.

Switching between VPCs: Tencent Cloud allows you to migrate one or more CVM instances from VPC A to VPC B at a time.

Specifying a custom IP address.

Choosing to retain the original private IP and `HostName` of the instance.

Preparations

Before migration, unbind the CVM instance from the CLB instances and secondary ENIs in the private and public networks and release the secondary IP of the primary ENI. Rebind them after migration.

Directions

Determining the network attribute of the CVM instance

1. Log in to the CVM console.
2. On the "instance" list page, view the target instance of which the network is pending switched based on the actually used view mode.

List view

Tab view

The instance is on the classic network if **Network: Classic Network** is displayed in the **Instance Configuration** column.

The instance is in the classic network if "Network: Classic Network" appears in **Basic Information**.

Note:

Switching from a classic network to a VPC is irreversible. A CVM instance cannot communicate with CVM instances in classic networks after being migrated from a classic network to a VPC.

Before switching from a classic network to a VPC, you need to create a VPC in the same region and a subnet in the same AZ as the target CVM instance in advance. For detailed directions, see Creating VPC.

After determining the network attribute of the instance, switch to VPC as required.

Switching to VPC

1. Log in to the CVM console.

2. On the **Instances** page, migrate the target instance to VPC.

List view

Tab view

Migrating a single instance to the VPC

Select the target instance and click **More** > **Resource Adjustment** > **Switch VPC** in the **Operation** column on the right.

Batch migrating instances to the VPC

Select the target instances and click **More Actions** > **Resource Adjustment** > **Switch VPC** above the list of instances.

Note:

Batch migration is only supported for CVM instances in the same availability zone.

Select the tab of the target instance and click **More Actions** > **Resource Adjustment** > **Switch VPC** in the top-right corner.

3. In the **Switch VPC** window that appears, read the notes and then click **Next**.

4. Select the destination VPC and the corresponding subnet and then click **Next**.

5. Set the private IP and **HostName Options** of the selected subnet as needed.

Set the main parameters as follows:

Pre-allocate IP: If the original private IP is not retained, you can enter the **Pre-allocate IP**. If it is not entered, the system will automatically assign one.

Retain original private IP: Set it as needed.

HostName Options: Set it as needed.

6. Click **Next**, perform the operations according to the instructions on the **Shutdown CVM** page, and click **Start Migration**. On the **Instances** page, you will see that **Modifying instance VPC attributes** is displayed in the **Status** column of the migrated instances.

Note:

During the migration, the CVM instance or instances need to be restarted. Therefore, do not perform other operations. After the migration, please check whether the CVM instance or instances are running normally and can be accessed via a private network and logged in to remotely.

EIP Direct Connection

Last updated：2024-01-08 09:37:01

EIP Direct Connection is ideal for scenarios where you want to check the public IP in CVM, like when you need to forward private and public traffic to different IP addresses. This document provides instructions on how to configure EIP Direct Connection in both Linux and Windows CVM.

Note:

EIP Direct Connection may cause network interruption. Please consider whether a short interruption to your business operations is acceptable.

Use Cases

When you want to access internet via an EIP, you can choose NAT mode or direct connection mode. The default mode is NAT mode.

In NAT mode, EIP is invisible on the local machine. You need to manually add an EIP address for each configuration.

In direct connection mode, the EIP is visible on the local machine. You do not need to manually add an EIP address for each configuration, which can minimize development cost.

Use Limits

At present, EIP direct connection is under beta test and is only available to allowed users. It only supports devices in a VPC. You can submit a ticket.

If you switch your devices to a VPC, you need to reconfigure EIP Direct Connection.

On CVM, EIP direct connection cannot take effect at the same time as an NAT gateway. If the routing table associated with the subnet where your CVM resides is configured with a routing policy of accessing the public network through the NAT gateway, direct connection cannot be implemented through the EIP on the CVM. You can allow the CVM to access the public network through its EIP by adjusting the priorities of NAT gateways and EIPs. In this case, EIP direct connection can be implemented.

Directions

Note:

To use EIP direct connection, you need to enable it in the console first, then download the script for EIP Direct Connection and run it in your CVM. Otherwise, EIP direct connection might not funtion properly.

We provide a script for configuring the IP so that private network traffic goes through the private IP and public network traffic goes through the public. For other applications, configure the routing accordingly.

Configuring EIP direct connection on Linux CVM

Configuring EIP direct connection on Windows CVM

The script for Linux is applicable to the following scenarios: both the private IP and public IP are bound to the primary ENI (eth0), where the public network address is accessed through the public IP, and the private network address is accessed through the private IP.

Note:

The script for Linux supports CentOS 6 and later, and Ubuntu.

Step one: download the script for EIP direct connection

EIP direct connection may cause network interruption. Therefore, you need to download the script for EIP direct connection and upload it to CVM in advance. You can obtain the script by using one of the following methods:

Method 1: upload the script for EIP direct connection

(1) Download the configuration script for EIP direct connection from Download Script for Linux.

(2) After the script for Linux is downloaded onto the local machine, upload it to the CVM that requires EIP direct connection.

Method 2: directly use a command

Log in to the CVM, and run the following command on the CVM to download the script:

```
wget https://network-data-1255486055.cos.ap-
guangzhou.myqcloud.com/eip_direct.sh
```

```
wget https://network-data-1255486055.cos.ap-
guangzhou.myqcloud.com/eip_direct.sh
```

Step two: configure EIP direct connection in the EIP Console

1. Log in to the EIP Console.

2. Find the EIP that is bound to the primary ENI and choose **More > Direct Connection** in the Operation column on the right.

3. Click on OK on the pop-up window.

Step three: run the script for EIP direct connection

After configuring EIP for the primary ENI (eth0), you need to log in to the CVM and run the script for EIP direct connection.

1. Log in to the CVM that requires EIP direct connection.

2. Run the script for EIP direct connection as follows:

2.1 Run the following command to add the execution permission:

```
chmod +x eip_direct.sh
```

2.2 Execute the `ip addr` command to check the name of the ENI that requires EIP direct connection.

2.3 Execute the following command to run the script.

Here, `ethx` indicates the name of the ENI (required). `XX.XX.XX.XX` indicates the EIP address (optional). You may leave it blank and run `./eip_direct.sh install ethx` directly.

```
./eip_direct.sh install ethx XX.XX.XX.XX
```

The script for Windows is applicable to the following scenarios: Public network traffic goes through the primary ENI, and private network traffic goes through the secondary ENI.

Note:

To use EIP direct connection in Windows, you need one ENI for private IP and one ENI for public IP, and bind the public IP to the primary ENI and bind the private IP to the secondary ENI.

During configuration of EIP direct connection in Windows, your internet connection may be interrupted. Therefore, we recommend that you log in to a Windows instance via VNC.

Step one: download the script for EIP direct connection

During configuration of EIP direct connection, the internet connection will be interrupted. Therefore, you need to download the script for EIP direct connection and upload it to CVM in advance.

1. Log into Windows Instance via VNC to access the CVM that requires EIP direction connection.
2. Open the following link in the browser of the CVM to download the script for EIP direct connection:

```
https://eip-public-read-1255852779.cos.ap-
guangzhou.myqcloud.com/eip_windows_direct.bat
```

Step two: configure the secondary ENI

Given that the Windows script is designed for scenarios where auxiliary network cards handle internal network traffic, it is therefore necessary to configure auxiliary network cards for the CVM.

1. Log in to the CVM Console.
2. On the Instances page, click the configured CVM ID to go to the Basic Information page.
3. Select the **ENI tab** and click **Bind ENI** to create an ENI that is in the same subnet as the primary ENI.

4. In the pop-up window, select **Create** and **Bind** an ENI, enter the information, select A**utomatic Assignment in Assign IP** section and click **OK**.

Subnet: Select the subnet to which the cloud server belongs.

IP assignment: You can select A**utomatic Assignment in Assign IP** or enter an IP manually.

Step three: configure EIP direct connection for the primary ENI

Upon completion of the auxiliary network card configuration, configure the EIP passthrough for the primary network card in the EIP console.

1. Log in to the Public IP Console.

2. Find the EIP that is bound to the primary ENI and choose **More > Direct Connection** in the Operation column on the right.

3. Click on **OK** on the pop-up window.

Step four: configure IP in CVM

After configuring the EIP direct connection for the the primary ENI in the EIP console, you need to log into the CVM to configure the EIP.

1. Log in to the CVM Console. This operation may cause public network interruption. Therefore, you need to log in to Windows Instance via VNC.

2. On the operating system page, select

in the lower-left corner and click

to open the Windows PowerShell window. Enter `firewall.cpl` and press **Enter** to open the Windows Firewall page.

3. Click **Turn Windows Firewall on or off** to go to the **Customize Settings** page.

4. Select **Turn off Windows Firewall** both in the **Private network settings pane** and the **Public network settings pane**.

5. Double-click to run the script downloaded in Step 1. Enter the public IP address and press Enter twice.

6. Enter `ipconfig` in the **Windows PowerShell** window and press **Enter**. You can see that the IPv4 address on the primary ENI changes to the public network address.

Note:

When the direct connection is enabled, you cannot assign a private IP to the primary ENI. Otherwise, the CVM cannot access the public network.

Security
Security Groups
Security Group

Last updated：2024-01-08 09:41:35

A security group is a virtual firewall that features stateful data packet filtering. It is used to configure the network access control of CVM, Cloud Load Balancer, TencentDB, and other instances while controlling their outbound and inbound traffic. It is an important means of network security isolation.
You can configure security group rules to allow or reject inbound and outbound traffic of instances within the security group.

Security Group Features

A security group is a logical group. You can add CVM, ENI, TencentDB, and other instances in the same region with the same network security isolation requirements to the same security group.
If a security group has no rules, it will reject all traffic by default, and you need to add rules to it to allow traffic.
Security groups are stateful. Inbound traffic you have allowed can automatically become outbound and vice versa.
You can modify security group rules at any time, and the new rules will take effect immediately.

Usage Limits

For use limits and quotas of security groups, see security group limits in Use Limits Overview.

Security Group Rules

Components

A security group rule consists of:
Source: IP address of the source data (inbound) or target data (outbound)
Protocol type and protocol port: Protocol type such as TCP and UDP
Policy: Allow or reject the access request.

Rule priorities

The rules in a security group are prioritized from top to bottom. The rule at the top of the list has the highest priority and will take effect first, while the rule at the bottom has the lowest priority and will take effect last.

If there is a rule conflict, the rule with the higher priority will prevail by default.

When the traffic goes in to or out from an instance bound to a security group, the security group rules are calculated from top to bottom. If a rule is matched and executed (allow/reject requests), the subsequent rules will not be matched.

Multiple security groups

An instance can be bound to one or multiple security groups. When it is bound to multiple security groups, the security group rules are calculated from top to bottom. You can adjust the priorities of security groups at any time.

Security Group Templates

Tencent Cloud provides the following two security group templates:

Open all ports: All inbound and outbound traffic are allowed

Open common ports : It opens port TCP 22 (for Linux SSH login), ports 80 and 443 (for Web service), port 3389 (for Windows remote login), the ICMP protocol (for Ping commands), and allows all traffic from the private network.

Note:

If these templates cannot meet your actual needs, you can create custom security groups. For more information, see Creating a Security Group and Security Group Use Cases.

If you need to protect the application layer (HTTP/HTTPS), please activate Tencent Cloud Web Application Firewall (WAF), which provides web security at the application layer to defend against web vulnerabilities, malicious crawlers, and CC attacks, protecting your websites and web applications security.

Directions

The following figure shows you how to use a security group:

Security Group Best Practices

Creating a security group

We recommend that you specify a security group while you're purchasing a CVM via the API. Otherwise, the default security group will be used. The default security group cannot be deleted, and it adopts the default security rule (i.e., allowing all IPv4 addresses). You can modify the security rule after the security group is created.

If you need to change the instance protection policy, we recommend modifying the existing rules rather than creating a new security group.

Managing rules

Export and back up the security group rules before you modify them, so you can import and restore them if an error occurs.

To create multiple security group rules, please use the parameter template.

Associating a security group

You can add instances with the same protection requirements to the same security group, instead of configuring a separate security group for each instance.

It's not recommended to bind one instance to too many security groups, which may cause rule conflicts and result in network disconnection.

Creating a Security Group

Last updated：2024-01-08 09:41:35

Scenario

Security Groups act as virtual firewalls for CVMs. Each CVM instance must associate with at least one security group. By default, each CVM instance has two templates (**Open all ports** and **Open port 22, 80, 443, 3389 and ICMP protocol**) for creating a default security group. For details, refer to Security Group Overview.

If the default security group does not meet your needs, you can create your own security group as instructed below.

Directions

1. Log in to the CVM Console.
2. In the left sidebar, select **Security Group**. The Security Group page then appears.
3. Select a region for the security group. Click **+New**.
4. In the **Create a security group** page, complete the following configurations:

Template: select a template that suits your needs, as shown below:

Template	Description	Notes

Open all ports	All ports are open. May present security issues.	-
Open TCP port 22, 80, 443, 3389 and ICMP	TCP port 22, 80, 443 and 3389, and the ICMP are open. All ports are open internally.	Suitable for instances with web services.
Custom	Creates a blank security group in which rules are added afterwards. For details on how to add rules, refer to this article.	-

Name: name of the security group.

Project: by default, **Default project** is selected. Select a project for better management.

Notes: a short description for the security group.

5. Click **OK** to create the security group.

If you select **Custom** as the template for your security group, click **Add rules now** to add security group rules.

Adding Security Group Rules

Last updated：2024-01-08 09:41:35

Overview

Security groups are used to manage traffic to and from public and private networks. For the sake of security, most inbound traffic is denied by default. If you selected **Open all ports** or **Open ports 22, 80, 443, 3389 and ICMP protocol** as the template when creating a security group, rules are automatically created and added to the security group to allow traffic on those ports. For more information, please see Security Groups.
This document describes how to add security group rules to allow or reject traffic to and from public or private networks.

Notes

Security group rules support IPv4 and IPv6 rules.
Open all ports allows both IPv4 and IPv6 traffic.

Prerequisites

You should have an existing security group. If you do not, refer to Creating a Security Group for details.
You should know which traffic is allowed or rejected for your CVM instance. For more information on security group rules and their use cases, please see Security Group Use Cases.

Directions

1. Log in to the CVM console.

2. Select **Security Group** on the left sidebar to access the security group management page.

3. Select a region, and locate the security group for which you want to set rules.

4. Click **Modify Rules** in the **Operation** column.

5.

Click **Inbound rules** and choose

 either of the following methods to add rules.

Note:

The following instructions use **Add a Rule** as an example.

Open all ports: this method is ideal if you do not need custom ICMP rules and all traffic goes through ports 20, 21, 22, 80, 443, and 3389 and the ICMP protocol.

Add a Rule: this method is ideal if you need to use multiple protocols and ports other than those mentioned above.

6. In the pop-up window, set rules.

Configure the following parameters:

Type: **Custom** is selected by default. You can also choose another system rule template including **Login Windows CVMs (3389)**, **Login Linux CVMs (22)**, **Ping**, **HTTP (80)**, **HTTPS (443)**, **MySQL (3306)**, and **SQL Server (1433)**.

Source or **Destination**: traffic source (inbound rules) or destination (outbound rules). You need to specify one of the following options:

Source or Destination	Description
A single IPv4 address or an IPv4 range	In CIDR notation, such as 203.0.113.0, 203.0.113.0/24 or 0.0.0.0/0, where 0.0.0.0/0 indicates all IPv4 addresses will be matched.
A single IPv6 address or an IPv6 range	In CIDR notation, such as FF05::B5, FF05:B5::/60, ::/0 or 0::0/0, where ::/0 or 0::0/0 indicates all IPv6 addresses will be matched.
ID of the referenced security group. You can reference the ID of: Current security group Other security group	To reference the current security group, please enter the ID of security group associated with the CVM. You can also reference another security group in the same region and belongs to the same project by entering the security group ID. **Note:** The referenced security group is available to you as an advanced feature. The rules of the referenced security group are not added to the current security group. If you enter the security group ID in Source/Destination when configuring security group rules, the private IP addresses of the CVM instances and the ENIs that are associated with this security group ID are used as the source/destination. This does not include public IP addresses.
Reference an IP address object or IP address group object in a parameter template.	-

Protocol port: enter the protocol type and port range or reference a protocol/port or protocol/port group in a parameter template. The supported protocol type includes TCP, UDP, ICMP, ICMPv6 and GRE in the following formats.

Single port: such as `TCP:80` .

Multiple ports: such as `TCP:80,443` .

Port range: such as `TCP:3306-20000` .

All ports: such as `TCP:ALL` .

Policy: **Allow** or **Refuse**. **Allow** is selected by default.

Allow: traffic to this port is allowed.

Refuse: data packets will be discarded without any response.

Notes: a short description of the rule for easier management.

7.

Click **Complete** to finish adding the rule.

8. To add an outbound rule, click **Outbound rule** and refer to Step 5 to Step 7.

Associating CVM Instances with Security Groups

Last updated：2024-01-08 09:41:35

Note:
Security groups can be associated with CVMs, ENIs, TencentDB for MySQL, and CLBs. This topic describes how to associate a security group with a CVM.

Overview

Security groups can be associated with one or more CVMs for network access control. They are an important part of CVM network security measures. You can associate your CVM with one or more security groups if necessary. The following are detailed instructions.

Prerequisites

You should already have a CVM instance created before starting.

Directions

1. Log in to the CVM console.
2. On the left sidebar, choose **Security Group** to go to the **Security Group** page.
3. On the **Security Group** page, select a region and locate the security group for which you want to set rules.
4. In the row of the target security group, click **Manage Instances** in the **Operation** column of the target security group. The **Associated Instances** page then appears.
5. On the **Associates Instances** page, click **Add Instance**.
6. In the pop-up **Add Instance** window, select the instances to bind and click **OK**.
Note:
After multiple security groups are bound to an instance, they are executed based on their priorities, which are consistent with their binding sequence. To adjust the priorities of the security groups, see Adjusting Security Group Priority.

Subsequent Operations

You can check all security groups in a specific region.

For operation details, see Viewing Security Groups.

If you want to disassociate a CVM instance from one or more security groups, you can remove it from the security group.

For operation details, see Remove from Security Groups.

If you no longer need a security group, you can delete it. Once a security group is deleted, all rules within it are also deleted.

For operation details, see Deleting a Security Group.

Managing Security Groups
Viewing Security Groups

Last updated：2024-01-08 09:41:35

Scenario

This article describes how to view all security groups of a region.

Directions

View security groups

1. Log in to the CVM Console.
2. In the left sidebar, select **Security Group**. The Security Group page then appears.
3. Select a region to see a list of security groups under that region.

Search for a security group

You can also use the search bar on the Security Group page to quickly find a specific security group.

1. Log in to the CVM Console.
2. In the left sidebar, select **Security Group**. The Security Group page then appears.
3. On the Security Group Management page, select **Regions**.
4. Click the search bar and use one of the following fields to search for a security group.

Security Group ID: input the desired ID and click

to see the corresponding security group.

Security Group Name: input the desired name and click

to see the corresponding security group.

Tag: input a tag and click

to see a list of all security groups with that tag.

Other Operations

To learn more about how to search for a security group, click

.

Remove from Security Groups

Last updated：2024-01-08 09:41:35

Scenario

You can remove a CVM instance from a security group if necessary.

Prerequisites

The instance is associated with two or more security groups.

Directions

1. Log in to the CVM Console.

2. In the left sidebar, select **Security Group**. The Security Group page then appears.

3. Select the desired region and find the desired security group.

4. Click the corresponding **Manage Instances** button to go to the **Bind with Instance** page.

5. Select the instances to be removed and click **Remove Selected**.

6. In the pop-up window, click **OK**.

Cloning Security Groups

Last updated：2024-01-08 09:41:35

Scenario

You might need to clone a security group if you:

Have created a security group sg-A in region A and you want to apply the same rules to an instance in region B. You can clone sg-A to region B, instead of creating a new security group from scratch.

Need a new security group for your service but want to clone the old security group as a backup.

Notes

By default, when you clone a security group, only the rules are cloned, not the association with instances.

You can clone a security group across projects and regions.

Directions

1. Log in to the CVM Console.

2. In the left sidebar, select **Security Group**. The Security Group page then appears.

3. Select desired region. A list of security groups under the region then appears.

4. Locate the desired security group and click **More**. Then click **Clone**. The **Clone security group** page then appears.

5. Select a **Target region** and **Target project** and input a **New name** for the new security group. Click **OK**.

Deleting a Security Group

Last updated：2024-01-08 09:41:35

Scenario

If you no longer need a security group, you can delete it. Once a security group is deleted, all rules within it are also deleted.

Prerequisites

Before deleting a security group, you must remove all associated CVM instances. Otherwise, the operation will fail. For details, refer to Removing From Security Group.

Directions

1. Log in to the CVM Console.
2. In the left sidebar, select **Security Group**. The Security Group page then appears.
3. Select the desired region and find the security group to be deleted.
4. Locate the desired security group and click **Delete**.
5. In the pop-up window, click **OK**.

Adjusting Security Group Priority

Last updated：2024-01-08 09:41:35

Overview

You can bind one or more security groups to a CVM. If you have bound multiple security groups, these security groups are executed based on their priorities. You can adjust the priorities as follows.

Prerequisite

The CVM instance is bound to two or more security groups.

Directions

1. Log in to the CVM console.
2. On the instance management page, click the ID of the CVM instance to go to the details page.
3. Click the **Security Groups** tab to enter the security group management page.
4. In the **Bound Security Groups** module, click **Sort**.

5. Click the following icon and drag it up/down to adjust the priority of the security group. The higher the position is, the higher the priority of the security group becomes.

6. After completing the adjustment, click **Save**.

Managing Security Group Rules
Viewing Security Group Rules

Last updated：2024-01-08 09:41:35

Scenario

After adding a security group rule, you can view its details in the console.

Prerequisites

You have created a security group and added at least one rule.

For information on how to create a security group and a security group rule, refer to Creating a Security Group and Adding Security Group Rules.

Directions

1. Log in to the CVM Console.

2. In the left sidebar, select **Security Group**. The Security Group page then appears.

3. On the **Security Group** page, select a region, and find the security group for which you want to view rules.

4. Click the ID or the desired security group to go to the details page.

5. Select **Inbound rule** or **Outbound rule** to view all inbound or outbound security group rules.

Modifying Security Group Rules

Last updated：2024-01-08 09:41:35

Scenario

This article describes how to modify a security group rule. Rules are important because they protect you CVM instance from malicious attacks. For example, they can protect certain ports from being abused.

Prerequisites

Make sure you have created a security group with rules.

Refer to Creating Security Groups and Adding Security Group Rules.

Directions

1. Log in to the CVM Console.

2. In the left sidebar, select **Security Group**. The Security Group page then appears.

3. Select the desired region and find the security group.

4. Locate desired security group and click **Modify Rules**. The Security Group Rule page then appears.

5. Use **Inbound rule** and **Outbound rule** to switch between inbound and outbound security group rules.

6. Locate the desired rule and click **Edit** to modify it.

Note:

You don't need to reboot the CVM for the rule changes to take effect.

Deleting Security Group Rules

Last updated：2024-01-08 09:41:35

Scenario

If you no longer need a security group rule, you can delete it.

Prerequisites

You have created a security group and added at least one rule to it.

For information on how to create a security group and add security group rules to it, see Creating a Security Group

and Adding Security Group Rules.

You have confirmed that your CVM instance does not need to permit or forbid Internet access or private network

access.

Directions

1. Log in to the CVM console.

2. In the left sidebar, click **Security Group**. The "Security Group" page then appears.

3. On the security group management page, select **Region** and locate the security group whose rules you want to

delete.

4. In the action column, click **Modify Rule** to go to the security group rule page.

5. Select inbound or outbound rules by clicking **Inbound Rules** or **Outbound Rules**.

6. Locate the security group rule to delete and click **Delete** in the action column.

7. In the window that appears, click **OK**.

Exporting Security Group Rules

Last updated：2024-01-08 09:41:35

Scenarios

You can export security group rules and save them locally for backup.

Directions

1. Log into the CVM Console.
2. On the left sidebar, choose **Security Group** to go to the **Security Group** page.

3. Select a region and locate the target security group.

4. Click the name or ID of the desired security group. The details page of the selected security group appears.

5. Use **Inbound Rule** and **Outbound Rule** to switch between inbound and outbound security group rules.

6. Click

to export security group rules to a file and save it to your local device.

Inbound rules	Outbound rules					

Add rule | Import rule | Sort | Edit all | Delete | Open all common ports | How to Set

☐ Source ⓘ ▼	Protocol+Port ⓘ	Policy	Remark	Modification time	Operation
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete
☐		Allow		2022-10-18 11:39:20	Edit Insert ▼ Delete

Total items: 6

10 ▼ / page |◄ ◄   1   / 1 page   ►

Importing Security Group Rules

Last updated：2024-01-08 09:41:35

Scenario

Security group rules can be imported from a file. You can use this feature to quickly restore or create security group rules.

Directions

1. Log in to the CVM Console.
2. In the left sidebar, select **Security Group**. The Security Group page then appears.
3. Select desired region to see a list of security groups.
4. Locate desired security group and click its name. Security Group Rule page appears.
5. Select inbound or outbound rules by clicking **Inbound rule** or **Outbound rule**.
6. Click **Import rules**. The **Batch import - Inbound/Outbound Rules** page appears.
7. Click **Browse** and select a rule template file. Click **Import**.
   Note:
   If there are existing rules in the security group, export them before importing new rules. Existing rules are overwritten after importing.
   If there is no existing rules in the security group, download the template first. Use it as a start to modify rules to your liking. Import them once you are finished.

Security Group Use Cases

Last updated：2024-01-08 09:41:35

Security groups can manage the access to CVMs. You can configure inbound and outbound rules for security groups to specify whether your server can be accessed by or can access other network resources.

The default inbound and outbound rules for security groups are as follows:

To ensure data security, the inbound rule for a security group is a rejection policy that forbids remote access from external networks. To enable public access to your CVM, you need to open the corresponding port to the Internet in the inbound rule.

The outbound rule for a security group specifies whether your CVM can access external network resources. If you select **Open all ports** or **Open ports 22, 80, 443, and 3389 and the ICMP protocol**, the outbound rule for the security group opens all ports to the Internet. If you select a custom security group rule, the outbound rule blocks all ports by default, and you need to configure the outbound rule to open the corresponding port to the Internet.

Common Use Cases

This document provides several common use cases of security groups. You can directly use its recommended security group configurations if a use case meets your requirements.

Scenario 1: remotely connecting to a Linux CVM via SSH

Case: you have created a Linux CVM and want to remotely connect to it via SSH.

Solution: when adding a security group rule, set **Type** to **Login Linux CVMs(22)**, enter WebShell proxy IP address for **Source**, and open TCP port 22 to the Internet to enable Linux login via SSH.

You can open all IP addresses or a specified IP address (or IP range) to the Internet as required. This allows you to configure the source IP addresses of the CVMs that can be remotely connected to through SSH.

Direction	Type	Source	Protocol Port	Policy
Inbound	Linux login	All IP addresses: 0.0.0.0/0 WebShell proxy IP addresses: as detailed in Orcaterm Proxy IP Addresses Updates Specified IP address: enter your specified IP address or IP range	TCP:22	Allow

Scenario 2: remotely connecting to a Windows CVM through RDP

Case: you have created a Windows CVM and want to remotely connect to it by using Remote Desktop (RDP).

Solution: when adding a security group rule, set **Type** to **Login Windows CVMs(3389)**, enter the WebRDP proxy IP

addresses for **Source**, and open TCP port 3389 to the Internet to enable remote login to Windows.

You can open all IP addresses or a specified IP address (or IP range) to the Internet as required. This enables you to configure the source IP addresses of the CVMs that can be remotely connected to via RDP.

Direction	Type	Source	Protocol Port	Policy
Inbound	Windows login	All IP addresses: 0.0.0.0/0 WebRDP proxy IP addresses: 81.69.102.0/24 106.55.203.0/24 101.33.121.0/24 101.32.250.0/24 Specified IP address: enter your specified IP address or IP range	TCP:3389	Allow

Scenario 3: pinging a CVM on the Internet

Case: you have created a CVM and want to test whether its communication with other CVMs is normal.

Solution: test the connection by using the `ping` command. Specifically, when adding a security group rule, set **Type** to **Ping** and open Internet Control Message Protocol (ICMP) ports to the Internet to enable other CVMs to access this CVM through ICMP.

You can open all IP addresses or a specified IP address (or IP range) to the Internet as required. This allows you to configure the source IP addresses of the CVMs that can access this CVM through ICMP.

Direction	Type	Source	Protocol Port	Policy
Inbound	Ping	All IP addresses: 0.0.0.0/0 Specified IP address: enter your specified IP address or IP range	ICMP	Allow

Scenario 4: remotely logging in to a CVM through Telnet

Case: you want to remotely log in to a CVM by using Telnet.

Solution: when adding a security group rule, configure the following security group rule:

Direction	Type	Source	Protocol Port	Policy
Inbound	Custom	All IP addresses: 0.0.0.0/0 Specified IP address: enter your specified IP address or IP range	TCP: 23	Allow

Scenario 5: allowing access to a web service through HTTP or HTTPS

Case: you have built a website and want to allow access to your website through HTTP or HTTPS.

Solution: when adding a security group rule, configure the following security group rules as required:

Allow all public IP addresses to access this website

Direction	Type	Source	Protocol Port	Policy
Inbound	HTTP (80)	0.0.0.0/0	TCP: 80	Allow
Inbound	HTTPS (443)	0.0.0.0/0	TCP: 443	Allow

Allow some public IP addresses to visit this website.

Direction	Type	Source	Protocol Port	Policy
Inbound	HTTP (80)	IP address or IP range that is allowed to access your website	TCP: 80	Allow
Inbound	HTTPS (443)	IP address or IP range that is allowed to access your website	TCP: 443	Allow

Scenario 6: allowing an external IP address to access a specified port

Case: you have deployed a service and want the specified service port (such as port 1101) to be externally accessible.

Solution: when adding a security group rule, set **Type** to **Custom** and open TCP port 1101 to the Internet to allow external access to the specified service port.

You can open all IP addresses or a specified IP address (or IP range) to the Internet as required. This allows the source IP address to access the specified service port.

Direction	Type	Source	Protocol Port	Policy
Inbound	Custom	All IP addresses: 0.0.0.0/0 Specified IP address: enter your specified IP address or IP range	TCP: 1101	Allow

Scenario 7: rejecting an external IP address to access a specified port

Case: you have deployed a service and want to prevent external access to a specified service port (such as port 1102).

Solution: when adding a security group rule, set **Type** to **Custom**, configure the TCP port 1102, and set **Policy** to **Reject**, so that external services cannot access the specified service port.

Direction	Type	Source	Protocol Port	Policy
Inbound	Custom	All IP addresses: 0.0.0.0/0	TCP: 1102	Reject

		Specified IP address: enter your specified IP address or IP range		

Scenario 8: allowing a CVM to access only a specified external IP address

Case: you want your CVM to access only a specified external IP address.

Solution: add two outbound security group rules as follows.

Allow the CVM instance to access a specified external IP address.

Forbid the CVM instance from accessing any public IP addresses via any protocol.

Note:

The first rule takes priority over the second.

Direction	Type	Source	Protocol Port	Policy
Outbound	Custom	Specified public IP address that the CVM can access	Required protocol and port number	Allow
Outbound	Custom	0.0.0.0/0	All	Reject

Scenario 9: prohibiting a CVM from accessing a specified external IP address

Case: you do not want your CVM to access a specified external IP address.

Solution: add a security group rule as follows.

Direction	Type	Source	Protocol Port	Policy
Outbound	Custom	Specified public IP address that your CVM instance cannot access	All	Reject

Scenario 10: uploading or downloading a file from a CVM through FTP

Case: you want to allow uploads and downloads over FTP.

Solution: add a security group rule as follows.

Direction	Type	Source	Protocol Port	Policy
Inbound	Custom	0.0.0.0/0	TCP: 20 to 21	Allow

Multi-scenario Configurations

You can configure multiple security group rules to meet your business requirements. For example, both inbound and outbound runes can be simultaneously configured. A CVM instance can be bound to one or multiple security groups. When it is bound to multiple security groups, the security group rules will be matched sequentially from top to bottom.

You can adjust the priorities of security groups at any time. For more information about the priorities, see Rule Priorities.

Server Common Port

Last updated：2024-01-08 09:41:35

This document describes common server ports. For more information on service application ports for Windows, see Service Overview and Network Port Requirements for Windows.

Port	Service	Description
21	FTP	An open FTP server port for uploading and downloading.
22	SSH	An SSH port for remotely connecting to Linux servers in CLI mode.
25	SMTP	An open SMTP server port for sending emails.
80	HTTP	A port for web services, such as IIS, Apache, and Nginx, to provide external access.
110	POP3	A port for the POP3 (email protocol 3) service.
137, 138, 139	NetBIOS protocol	Ports 137 and 138 are UDP ports for transferring files through My Network Places. Port 139: connections established through port 139 attempt to access the NetBIOS/SMB service. This protocol is used for file and printer sharing on Windows and SAMBA.
143	IMAP	A port for Internet Message Access Protocol (IMAP) v2, which is a protocol for receiving emails like POP3.
443	HTTPS	A port for web browsing. HTTPS is a variant of HTTP that provides encryption and transmission over secure ports.
1433	SQL Server	Default port for SQL Server. The SQL Server service uses two ports: TCP-1433 and UDP-1434. Port 1433 is used to provide external services, and port 1434 is used to return a response to the requester to indicate the TCP/IP port used by SQL Server.
3306	MySQL	Default port for MySQL databases, which is used by MySQL to provide external services.
3389	Windows Server Remote Desktop Services	Service port for the Windows Server remote desktop, through which you can connect to a remote server by using the "Remote Desktop" connection tool.
8080	Proxy port	Similar to port 80, port 8080 is used in the WWW proxy service for web

| | | browsing. The port number ":8080" is often appended to the URL when you visit a website or use a proxy. In addition, after the Apache Tomcat web server is installed, its default service port is port 8080. |

Security Group API Overview

Last updated：2024-01-08 09:41:35

API Name	Description
CreateSecurityGroup	Create security groups
CreateSecurityGroupPolicies	Create security group rules
DeleteSecurityGroup	Delete security groups
DeleteSecurityGroupPolicies	Delete security group rules
DescribeSecurityGroupAssociationStatistics	Query the statistics of the instances associated with a security group
DescribeSecurityGroupPolicies	Query security group rules
DescribeSecurityGroups	Query security groups
ModifySecurityGroupAttribute	Modify security group attributes
ModifySecurityGroupPolicies	Modify the inbound and outbound rules of a security group
ReplaceSecurityGroupPolicy	Replace a single security group rule

Protection of Sensitive Operations

Last updated：2024-01-08 09:41:35

Overview

The sensitive operation protection feature is currently available in CVM. Once the feature is enabled, identity verification needs to be completed before performing sensitive operations.

This feature can effectively protect the security of account resources, including shutdown, restart, VNC login, password reset, instance termination, system reinstallation, configuration adjustment, key load and VPC switch.

Enabling Operation Protection

You can enable the operation protection feature in Security Settings console. For more information, see Operation Protection.

Verifying Operation Protection

Once operation protection is enabled, you need to complete identity verification before you can perform a sensitive operation:

If you have enabled **MFA verification** for operation protection, you need to enter the 6-digit dynamic verification code displayed on the MFA device.

If you have enabled **SMS code verification** for operation protection, you need to enter the verification code received on your phone.

Managing Login Password

Last updated：2024-01-08 09:41:35

Overview

CVM accounts and passwords are the login credentials for CVMs. This document describes how to use and manage passwords when logging in to a CVM.

Limits

The password must comply with the following limits:

Linux instance: The password must consist of 8 to 30 characters. We recommend that you use a password of more than 12 characters. The password cannot start with "/" and must contain at least three of the following character types (`a-z` , `A-Z` , `0-9` and special characters `()`~!@#$%^&*-+=_|{}[]:;'<>,.?/`).

Windows instance: The password must consist of 12 to 30 characters. The password cannot start with "/" and must contain at least three of the following character types (`a-z` , `A-Z` , `0-9` and special characters `()`~!@#$%^&*-+=_|{}[]:;'<>,.?/`), excluding user names.

Directions

Setting an initial password

For different configuration methods selected during CVM purchase, the initial password settings will also be different. Instance creation through **Custom configuration**: During creation, the initial password setting methods for different login modes are different.

Login Method	Description
Random Password	The initial password will be sent to you via email and Message Center on the console.
Password Associated with Key	The login with username and password is disabled by default. To use the password, you can log in to the CVM console to reset it, see Resetting Instance Password.
Custom Password	The password you set is the initial password.

Viewing the password

The random password auto-generated by the system will be sent to you through email and the console Message Center. The following operations take Message Center as an example.

1. Log in to the CVM console.

2. Click

in the upper right corner and select the target product message.

Enter the product message page, and you can view the password.

Resetting the password

See Resetting Instance Password.

Managing SSH Keys

Last updated：2024-01-08 09:41:35

Scenarios

This document describes common operations related to using SSH key pair to log in to an instance. For example, you can create, bind, unbind, modify, or delete an SSH key pair.

Caution:

An SSH key can only be bound to or unbound from a CVM instance that is shut down. For directions on how to shut down a CVM instance, see Shutting Down Instances.

Directions

Creating an SSH key

1. Log in to the CVM console and select **SSH Key** on the left sidebar.

2. On the **SSH Key** page, click **Create secret**.

3. In the **Create an SSH key** pop-up window, configure the key.

Creation Method:

Create a new key pair: Enter a key name

Import existing public keys: Enter the key name and existing public key information.

Note:

You need to use a public key without a password; otherwise, you cannot log in to the instance in the console.

Key Name: Customize a name.

Tag (Optional): You can add tags for a key as needed, which can be used to categorize, search for, and aggregate resources. For more information, see Overview.

4. Click **OK**.

Note:

After clicking **OK**, the private key is automatically downloaded. Tencent Cloud will not retain your private key. If you lost the private key, create a new one and bind it with the instance again.

Binding a key to an instance

1. Log in to the CVM console and select SSH Key on the left sidebar.

2. On the **SSH Key** page, click **Bind instance** in the row of the target key.

3. In the pop-up window, select the **Region** and target CVM instance, and click **Bind**.

Unbinding a key from an instance

1. Log in to the CVM console and select SSH Key on the left sidebar.

2. On the **SSH Key** page, click **Unbind instance** in the row of the target key.

3. In the pop-up window, select the **Region** and target CVM instance, and click **Unbind**.

Modifying the SSH key name or description

1. Log in to the CVM console and select SSH Key on the left sidebar.

2. On the **SSH Key** page, select

on the right of the key name.

3. In the pop-up window, enter the new key name or description, and click **OK**.

Deleting an SSH key

Note:

An SSH key that is bound to a CVM instance or custom image cannot be deleted.

1. Log in to the CVM console and select SSH Key on the left sidebar.

2. On the **SSH Key** page, delete one key or batch delete keys as needed.

Deleting one key

Batch deleting keys

1. Click **Delete** in the row of the target SSH key.

2. In the key deletion pop-up window, click **OK**.

1. Select the target keys and click **Delete** at the top of the page.

2. In the key deletion pop-up window, click **OK**.

Only deletable ones of the selected key pairs will be deleted.

Relevant operations

Using an SSH key to log in to a Linux CVM

1. Create an SSH key.

2. Bind an SSH key to a CVM instance.

3. Log in to a Linux instance using SSH.

Editing a key tag

You can add, modify, or delete a tag for an SSH key in the following steps. For more information on tags, see Overview.

1. On the **SSH Key** page, click **Edit tags** on the right of the key.

2. In the **Edit tags** pop-up window, perform the desired operation.

3. Click **OK**.

Spread Placement Group

Last updated：2024-01-08 09:41:35

Scenario

This document describes how to manage spread placement groups.For more information about the placement group, see Placement Group.

Directions

Creating a placement group

1. Log in to the CVM placement group console.

2. Click **Create**.

3. In the window that appears, enter a name for the placement group, and select the layer of the placement group.

4. Click **OK** to finish the creation.

Starting up an instance in the placement group

1. Go to the CVM purchase page.

2. Complete the purchase as prompted on the page.

During the purchase process, be sure to perform the following operations:

When setting the CVM, click **Advanced Configuration**, select **Add Instance to Placement Group**, and select an existing placement group.

 If no existing placement groups meet your requirement, create one in the console.

When confirming the configuration information, enter the total number of instances to be added to the placement group, which must be less than the quantity limit set for the placement group.

Modifying an instance's placement group

Note:

Currently, you can change only the name of a placement group. To do this, complete the following steps.

1. Log in to the CVM placement group console.

2. Hover the cursor over the ID or name of the target placement group and click

.

3. In the window that appears, enter the new name.

4. Click **OK** to finish the modification.

Deleting a placement group

Note:

You can delete a placement group that needs to be replaced or is no longer needed. You must terminate all instances running in the placement group before you can delete it. To do this, complete the following steps.

1. Log in to the CVM placement group console.

2. Click **Number of Instances** for the placement group to be deleted to go to the instance management page, and terminate all instances in the placement group.

3. Return to the placement group console, select the placement group to be deleted, and click **Delete**.

4. In the window that appears, click **OK** to finish the deletion.

You can delete a single placement group or multiple placement groups in batches.

Unblocking Port 25

Last updated：2024-01-08 09:41:35

Operation Scenarios

In the case of a unique scenario, you must use TCP port 25 for outbound connections on the CVM. This document guides you how to request unblock of port 25.

Notes

It only supports unblocking of prepaid annual or monthly CVMs and currently does not support CVMs operating on a pay-as-you-go basis.

You can only unblock port 25 for five instances for each Tencent Cloud account.

Make sure that you only use port 25 to connect to a third-party SMTP server for sending email. If you use your CVMs to send email directly, we reserve the right to permanently ban you from opening port 25.

It is recommended to prioritize the use of other ports for sending emails. For configuration guidance, see Sending Emails via Port 465.

Directions

1. Log in to the Tencent Cloud console.

2. Click your account name in the upper-right corner. Select **Security Management**.

3. In the left sidebar, click **Unblock port 25** to go to the **Unblock port 25** page.

4. Click **Apply for unblocking port 25** to bring up the **Apply for unblocking port 25** window.

5. In the **Application for Unblocking TCP Port 25** pop-up window, you can select the region and the CVM instance for which port 25 needs to be unblocked, input details of the intended use, and set Reverse DNS (rDNS) records. Then, select **I have read and accepted "Port 25 Protocol"**. For details, see the following figure.

Note:

Make sure you have not used up your unblocking quota. You can check the remaining quota in the lower left of the **Application for Unblocking TCP Port 25** window.

Application for Unblocking TCP Port 25

Note: In order to improve the performance for sending emails from Tencent Cloud IP addresses, your CVMs are restricted from accessing the external TCP Port 25 by default. You can apply for unblocking your CVMs. A maximum of 5 unblocking operations are allowed for each account.

Select Region *   South China (Guangzhou)

CVM *   Search CVM

Purpose Description *   Describe the purpose in a clear and detailed manner and the solution to avoid sending spam emails. This helps reviewers review and confirm your unblocking applications.

Reverse DNS *   Enter reverse DNS (rDNS)

IP address   Optional. IP address for sending ou

ⓘ Remaining quota: 5 times

☐ I have read and accepted "Port 25 Protocol"

OK   Cancel

6. Click **OK** to complete the application. Please patiently await the platform administrator to process. You can check the approval results and causes via **Moderation Status** on the page.

Tags

Managing Instances via Tags

Last updated：2024-01-08 09:41:35

Overview

Tags are key-value pairs provided by Tencent Cloud for easy resource identification. You can use tags to categorize and manage your CVM resources.

Tencent Cloud will not use your tags, they are solely used by you to manage your CVM resources.

Usage Limits

Note the following limits when using tags:

Quantity limits: each Tencent Cloud resource allows up to 50 tags.

Tag key limits:

Tag keys cannot start with `qcloud` , `tencent` , or `project` .

A tag key can contain up to 255 characters, including numbers, letters, and `+=.@-` .

Tag value limits: a tag value can contain up to 127 characters, including numbers, letters and `+=.@-` . It can be left empty if necessary.

Directions and Use Cases

Use case

A company has purchased six CVM instances, of which the business group, scope and owners are as follows:

Instance ID	Business Group	Business Scope	Owner
ins-abcdef1	E-commerce	Marketing campaigns	John Smith
ins-abcdef2	E-commerce	Marketing campaigns	Chris
ins-abcdef3	Games	Game A	Jane Smith
ins-abcdef4	Games	Game B	Chris
ins-abcdef5	Entertainment	Post-production	Chris

| ins-abcdef6 | Entertainment | Post-production | John Smith |

Taking ins-abcdef1 as an example, we can add the following 3 sets of tags to the instance:

Tag Key	Tag Value
dept	ecommerce
business	mkt
owner	John Smith

Similarly, you can add tag key-value pairs to other instances based on the business group, scope and owners.

Setting tags in the CVM console

Take the preceding case as an example. After designing the tag key-value pairs, you can log in to the CVM console to specify the tags.

1. Log in to the CVM console.
2. On the instance management page, proceed according to the actually used view mode:

List view

Tab view

Select the instance for which to edit tags and click **More** > **Instance Settings** > **Edit Tags** as shown below:

Select the instance for which to edit tags and click **More Actions** > **Instance Settings** > **Edit Tags** in the top-right corner as shown below:

3. In the **You have selected 1 resource** window that appears, specify the tags as required. As shown in the following figure.

For example, you can add three tag key-value pairs to the ins-abcdef1 instance.

4. Click **OK**. A message indicating the edit was successful will be prompted.

Filtering instances by tags

To filter instances by tag, follow the steps below:

1. Click the search box and select **Tag** from the drop-down list.

2. Enter the tag, and click

 to search.

You can filter instances using tags. For example, you can search instances that are bound with tags `key1` or `key2` by entering `Tag: key1|key2` in the search box.

Editing Tags

Last updated：2024-01-08 09:41:35

Overview

This document describes how to edit the tags of resources.

Usage Limits

There are several limits on editing tags:

Quantity: Each Tencent Cloud resource allows up to 50 tags.

Tag key limits:

A tag key cannot start with `qcloud` , `tencent` , or `project` .

A tag key can contain up to 255 characters, including numbers, letters, and `+=.@-` .

Tag value limits: A tag value can contain up to 127 characters, including numbers, letters and `+=.@-` . It can be left empty if necessary.

Prerequisites

Log in to the CVM console.

Directions

Single instance

Multiple instances

1. Open the tag editing page.

List view: Select the target instance and click **More** > **Instance Settings** > **Edit Tags**.

Tab view: Select the target instance and click **More Actions** > **Instance Settings** > **Edit Tags**.

2. Add, modify or delete the tags in the pop-up window.

Note:

You can batch edit tags for up to 20 resources at a time.

1. On the instance management page, select the target instances and click **More Actions** > **Instance Settings** > **Edit Tags**.

2. Add, modify, and delete tags in the pop-up window.

References

For information on how to use tags, please see User Guide on Tags.

Monitoring and Alarms
Getting Monitoring Statistics

Last updated：2024-01-08 09:41:35

Overview

Tencent Cloud provides the Cloud Monitor feature for all users by default. This feature helps you monitor and collect data from the Tencent Cloud products you are using. This document describes how to obtain the monitoring data.

Directions

CVM console

Cloud Monitor console

Cloud Monitor dashboard

API

Note:

CVM console provides a monitoring page, on which you can view the monitoring data of CPU, memory, network bandwidth and disks in the specified period.

1. Log in to the CVM console.

2. In the instance management page, click the ID/Name of the CVM to enter its details page and view the monitoring data.

3. Click the **Monitoring** tab to get the instance monitoring data.

Note:

Cloud Monitor console provides the monitoring data of all Tencent Cloud products. On the console, you can view the monitoring data of CPU, memory, network bandwidth and disks in the specified period.

1. Log in to the Cloud Monitor console.

2. Select **Cloud Product Monitoring** > **Cloud Virtual Machine** on the left sidebar.

3. Click the ID/Name of the CVM instance to enter its details page and view the monitoring data.

Specify required CVM metrics and create a dashboard, on which you can view monitoring data in intuitive charts, helping you analyze metrics through trends and exceptional values.

1. Log in to the Cloud Monitor console and select **Dashboard** > Default Dashboard.

2. Create a dashboard as instructed in Create Dashboard and get the monitoring data.

You can use the `GetMonitorData` API to get the monitoring data for all Tencent Cloud products. For more information, see GetMonitorData.

Creating Alarm Policies

Last updated：2024-01-08 09:41:35

Overview

You can set threshold-triggered alarm policies to monitor CVM performance, and also event-triggered alarm polices to watch the status of CVM instances and the underlying platform infrastructure. When an exception occurs, you will receive notifications via the specified methods (email, SMS or phone call). A proper alarm policy will help improve the robustness and reliability of your applications. You can also see Creating Alarm Policy.

Directions

1. Log in to the Cloud Monitor console and click **Alarm Configuration** > Alarm Policy on the left sidebar.
2. On the **Alarm Policy** page, click **Create**.
3. In the pop-up window, configure basic information, alarm policy, and notification template as instructed below.

Configuration Type	Configuration Item		Description
Basic Info	Policy Name		A custom policy name
	Remarks		Remarks for the policy
	Monitor Type		Choose Cloud Product Monitoring
	Policy Type		Select the desired policy type for monitoring Tencent Cloud services.
	Project		Choose a project as needed. You can later find this policy quickly by filtering by the project.
Configure Alarm Policies	Alarm Object		Instance ID: associate the policy with the specified CVM instance Tag: associate the policy with CVM instances bound with the specified tag Instance Group: associate the policy with the selected instance group All Objects: associate the policy with all instances under the current account (permission required)
	Trigger Condition	Manual Configuration	Trigger condition: specify the metric, comparison, threshold, statistical period, and the number of

		(Metric Alarm)	consecutive periods. You can expand the trigger condition to view the metric trend, and based on which, set a proper threshold.
		Manual Configuration (Event Alarm)	Create an event alarm policy to get notifications in case of service resources or underlying infrastructure exceptions
		Select template	Choose a configured template as needed. For more information about the configurations, please see Configuring Trigger Condition Template.
Configure Alarm Notification (optional)	Notification Template		It defaults to the preset notification template (sending the notification to the root account admin via SMS and email). Up to 3 notification templates can be bound to each alarm policy. For more information about the configurations of notification templates, see Creating Notification Template.

4. Click **Complete**.

Sample Console Configuration

Last updated：2024-01-06 18:00:25

Introduction

You can use Cloud Access Management (CAM) policies to manage user access to resources using the Cloud Virtual Machine (CVM) console. This document provides examples to help you understand how to use the pre-defined CAM policies using the CVM console.

Examples

Read and write (CVM)

If you want to allow a user to create and manage CVM instances, associate the user with the policy named QcloudCVMFullAccess. This policy is designed to grant users the permissions to access all the resources in CVM, Virtual Private Cloud (VPC), Cloud Load Balancer (CLB), and Cloud Monitor.

The detailed steps are as follows:

Refer to Authorization Management for instructions on how to grant the preset policy QcloudCVMFullAccess to a user.

Read-only (CVM)

If you want to allow a user to only query, but not create, delete or start/shutdown CVM instances, associate the user with the policy named QcloudCVMInnerReadOnlyAccess. This policy is designed to grant users the permissions to perform all operations starting with "Describe" and "Inquiry" in CVM. The detailed steps are as follows:

Refer to Authorization Management for instructions on how to grant the preset policy QcloudCVMInnerReadOnlyAccess to a user.

Read-only (CVM and associated resources)

If you want to to allow a user to only query, but not create, delete or start/shut down CVM instances and associated resources (VPC and CLB), associate the user with the policy named QcloudCVMReadOnlyAccess. This policy is designed to grant users the permissions to perform the following operations:

All operations starting with "Describe" and "Inquiry" in CVM.

All operations starting with "Describe", "Inquiry", and "Get" in VPC.

All operations starting with "Describe" in CLB.

All operations in the Monitor.

The detailed steps are as follows:

Refer to Authorization Management for instructions on how to grant the preset policy QcloudCVMReadOnlyAccess to

a user.

CBS policies

If you want to allow a user to view, create, and use cloud disks on the CVM console, add the following operations to your policy and associate the policy with the user.

CreateCbsStorages: create a cloud disk.

AttachCbsStorages: mount the specified cloud disk to the specified CVM.

DetachCbsStorages: unmount the specified cloud disk.

ModifyCbsStorageAttributes: modify the name or the project ID of the specified cloud disk.

DescribeCbsStorages: query the details of a cloud disk.

DescribeInstancesCbsNum: query the number of mounted cloud disks of a CVM and the maximum number of cloud disks that are allowed to be mounted to the CVM.

RenewCbsStorage: renew the specified cloud disk.

ResizeCbsStorage: resize the specified cloud disk.

The detailed steps are as follows:

1. Refer to Policies for information and create a custom policy that grants the permissions to view cloud disk information on the CVM console and to create and use cloud disks.

Use the following as a syntax reference:

```
{
 "version": "2.0",
 "statement": [
     {
         "effect": "allow",
         "action": [
             "name/cvm:CreateCbsStorages",
             "name/cvm:AttachCbsStorages",
             "name/cvm:DetachCbsStorages",
             "name/cvm:ModifyCbsStorageAttributes",
             "name/cvm:DescribeCbsStorages"
         ],
         "resource": [
             "qcs::cvm::uin/1410643447:*"
         ]
     }
  ]
 }
```

2. Find the created policy, and in the "Action" column of the row, click **Associate User/Group**.

3. In the "Associate User/Group" window, select the user/group you want to associate, and click **OK**.

Security group policies

To allow a user to view and use security groups on the CVM console, add the following operations to your policy, and associate the policy with the user.

DeleteSecurityGroup: delete a security group.

ModifySecurityGroupPolicys: replace all the policies of a security group.

ModifySingleSecurityGroupPolicy: modify a single policy of a security group.

CreateSecurityGroupPolicy: create a security group policy.

DeleteSecurityGroupPolicy: delete a security group policy.

ModifySecurityGroupAttributes: modify the attributes of a security group.

The detailed steps are as follows:

1. Refer to Policies for information and create a custom policy that grants the permissions to create, delete, and modify security groups on the CVM console.

Use the following as a syntax reference:

```
{
 "version": "2.0",
 "statement": [
     {
         "action": [
             "name/cvm:ModifySecurityGroupPolicys",
             "name/cvm:ModifySingleSecurityGroupPolicy",
             "name/cvm:CreateSecurityGroupPolicy",
             "name/cvm:DeleteSecurityGroupPolicy"
         ],
         "resource": "*",
         "effect": "allow"
     }
  ]
}
```

2. Find the created policy, and in the "Action" column of the row, click **Associate User**/**Group**.

3. In the "Associate User/Group" window, select the user/group you want to authorize, and click **OK**.

Policy for EIPs

If you want to allow a user to view and use EIPs on the CVM console, add the following operations to your policy, and associate the policy with the user.

AllocateAddresses: assign an EIP to a VPC or CVM instance.

AssociateAddress: associate an EIP with an instance or a network interface.

DescribeAddresses: view EIPs on the CVM console.

DisassociateAddress: disassociate an EIP from an instance or a network interface.

ModifyAddressAttribute: modify the attributes of an EIP.

ReleaseAddresses: release an EIP.

The detailed steps are as follows:

1. Refer to Policies for information and create a custom policy.

This policy allows users to view an EIP and assign it to and associate it with an instance on the CVM console. Users cannot modify the attributes of the EIP, disassociate it from an instance, or release the EIP. Use the following as a syntax reference:

```
{
 "version": "2.0",
 "statement": [
     {
         "action": [
             "name/cvm:DescribeAddresses",
             "name/cvm:AllocateAddresses",
             "name/cvm:AssociateAddress"
         ],
         "resource": "*",
         "effect": "allow"
     }
 ]
}
```

2. Find the created policy, and in the "Action" column of the row, click **Associate User/Group**.

3. In the "Associate User/Group" window, select the user/group you want to authorize, and click **OK**.

Policy for authorizing users to perform operations on specific CVMs

If you want to authorize a user to perform operations on a specific CVM, associate the following policy with the user. The detailed steps are as follows:

1. Refer to Policies for information and create a custom policy.

This policy authorizes the user to operate a CVM instance with the ID of ins-1 in the Guangzhou region. Use the following as a syntax reference:

```
{
 "version": "2.0",
 "statement": [
     {
         "action": "cvm:*",
         "resource": "qcs::cvm:ap-guangzhou::instance/ins-1",
         "effect": "allow"
     }
 ]
}
```

2. Find the created policy, and in the "Action" column of the row, click **Associate User/Group**.

3. In the "Associate User/Group" window, select the user/group you want to authorize, and click **OK**.

Policy for authorizing users to perform operations on the CVMs in a specific region

If you want to authorize a user to perform operations on the CVMs in a specific region, associate the following policy with the user. The detailed steps are as follows:

1. Refer to on Policies for information and create a custom policy.

This policy authorizes the user to operate CVM instances in the Guangzhou region. Use the following as a syntax reference:

```
{
 "version": "2.0",
 "statement": [
     {
         "action": "cvm:*",
         "resource": "qcs::cvm:ap-guangzhou::*",
         "effect": "allow"
     }
 ]
}
```

2. Find the created policy, and in the "Action" column of the row, click **Associate User**/**Group**.

3. In the "Associate User/Group" window, select the user/group you want to authorize, and click **OK**.

Granting a sub-account all permissions to CVM instances except payment

Assume that the account CompanyExample, whose ownerUin is 12345678, has a sub-account called Developer. Developer requires full management permissions (including all operations such as creation and management) for the CVM instance, except payment, which means Developer can make orders but cannot pay for them.

You can do this by using one of the following two solutions:

Solution A

The account owner of CompanyExample associate the preset policy QcloudCVMFullAccess with Developer. For more information, refer to Authorization Management.

Solution B

1. Use the following as a syntax reference and create a custom policy.

```
  {
"version": "2.0",
"statement":[
     {
         "effect": "allow",
         "action": "cvm:*",
         "resource": "*"
     }
 ]
}
```

2. Associate the policy to the sub-account. For more information, see Authorization Management.

Granting a sub-account the permission to manage projects

Assume that the enterprise account, CompanyExample, with ownerUin of 12345678, has a sub-account called Developer. The owner of CompanyExample wants to allow Developer to manage projects, including assigning and removing resources, on the console.

The detailed steps are as follows:

1. Create a custom policy for project management.

For more information, refer to Policies.

2. Refer to Authorization Management for information on how to associate the custom policy with the sub-account.

If you run into permission issues when attempting to view snapshots, images and EIPs, associate preset policies QcloudCVMAccessForNullProject, QcloudCVMOrderAccess, and QcloudCVMLaunchToVPC with the sub-account.

For more information on authorization, refer to Authorization Management.

Custom policy

If preset policies cannot meet your requirements, you can create custom policies.

For detailed instructions, refer to Policies.

For more information on CVM policy syntax, refer to Authorization Policy Syntax.