

Cloud Virtual Machine

FAQs

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

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

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by the Tencent corporate group, including its parent, subsidiaries and affiliated companies, as the case may be. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

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

Contents

FAQs

- Region and Availability Zone

Billing

- Purchasing and Renewing

Instance

- Issues on Logging in and Connecting to Instances

- Purchasing

- A CVM Instance Cannot Be Found

- Usage

- Use Limits

- Adjusting Configuration

- Spot Instances

- Big Data Instance

- Others

Storage

- System Disks

- Using Cloud Disks

- Mounting and Unmounting Cloud Disks

- Resizing Cloud Disks

- Snapshots

- Issues on Data Backup and Recovery

Images

- Custom Images

- Copying Images

- Issues on Custom Image

- Changing Images (Reinstalling System)

- Cloud-Init & Cloudbase-Init

- FAQs about Red Hat Enterprise Linux Image

- Others

Service Migration

Network

- IP Address

- Elastic Public IP issues

- Elastic Network Interface

- Public Network Bandwidth

Network

Security

 Passwords

 Password Login and SSH Key Login

 Security Groups

 Security Group Rules

 Port

 Firewall

Operating Systems

OPS & Monitoring

 Uploading Files

 OPS

 Common Operations and Commands

Access Control

NTP Service

Use Cases

FAQs

Region and Availability Zone

Last updated : 2024-01-06 17:21:38

How can I view regions and availability zones?

To view regions and availability zones, try the followings:

Refer to [Regions and Availability Zones](#) document.

Use APIs to query regions and availability zones.

To view regions, use the [DescribeRegions](#) API.

To view availability zones, use the [DescribeZones](#) API.

What CVM regions and availability zones are available and how to choose?

For more information on available CVM regions and availability zones, see [Regions and Availability Zones](#).

For more information on how to choose regions and availability zones, see [How to Select the Region and Availability Zone](#).

Can I change the region of a purchased CVM?

No. You cannot change the regions of purchased CVMs. If you need to deploy an existing CVM to another region or availability zone, please try the following:

[Terminate the instance](#) and then purchase a new one.

Create a custom image for the original instance. Then, use the custom image to create an instance, launch it and update its configuration in a new availability zone.

1.1 Create a custom image for the current instance. For more information, see [Create Custom Images](#).

1.2 If the [network environment](#) of the current instance is VPC and you need to retain the private IP address after migration, please delete the subnet in the current availability zone and then create a subnet in the new availability zone with the same IP address range as that of the original subnet.

Note:

If the subnet to be deleted contains available instances, first migrate all instances in the subnet to the new subnet.

1.3 Use the newly created custom image to create a new instance in the new availability zone.

You can choose the same instance type and configuration as that of the original instance, or configure new ones for the new instance. For more information, see [Creating an Instance](#).

1.4 If an elastic public IP address is associated with the original instance, disassociate it from the original instance and associate it with the new instance. For more information, see [Elastic IP Addresses \(EIPs\)](#).

1.5 (Optional) If the [pricing modes](#) of the original instance is pay-as-you-go, you can terminate it. For more information, see [Terminate Instances](#).

Can Tencent Cloud users in China enjoy the same product quality and services for resources purchased in China and other regions or countries?

Yes. Tencent Cloud China Console provides all users with the same product quality and services. The purchase region will not affect your user rights on the console.

Can I use the replication feature of a custom image to migrate a CVM in Mainland China to another country or region?

No. Images can only be replicated within the same country or region. To replicate images between countries, please [submit a ticket](#) to apply.

What are the differences between instances in and outside Mainland China, and how do I determine which region fits me the best?

Instances in other countries or regions are deployed outside Mainland China, providing geographical and market advantages for global users. The fast local network can also meet international customer needs, which is suitable for users who run businesses in other countries or regions.

For more information on supported regions, see [Regions and Availability Zones](#).

Can I change between Linux and Windows operating systems for a CVM instance purchased in a region outside Mainland China?

No. Switching OS between Windows and Linux is only available for CVM instances in Mainland China.

How do I apply for after-sale services for products purchased in regions outside Mainland China?

If you purchased a product on the Tencent Cloud official website, call the Tencent Cloud 24/7 service hotline (4009100100) or [submit a ticket](#).

How do I deploy CVM instances in regions outside China?

CVM instances in regions outside China have the same deployment method if they have the same type of operating system.

Can I migrate instances to regions outside Mainland China?

The region and availability zone of an instance cannot be changed. To use an instance in other countries or regions, you need to purchase another instance.

Why some instance types can only be purchased in China regions?

Some regions may not support certain instance types. For information on supported instance types, go to the [CVM purchase page](#) to view instance purchase information.

If I build a website using a CVM instance from a region outside Mainland China and my users access the website through a domain name, does this domain name need ICP filing?

For websites in regions in Mainland China, ICP filing is required for domain names. For websites in other countries or regions, ICP filing is not required.

Do instances in different regions have the same price?

The pricing of a CVM instance includes its specifications, storage, network bandwidth, etc. The prices of these components vary by regions, so the instance price will be different.

For more information on pricing, see [Pricing](#).

Billing

Purchasing and Renewing

Last updated : 2024-01-06 17:26:45

Purchasing CVMs

What should I do if the CVM is not created successfully?

If it takes a long time to create a CVM, wait to see if the CVM is created successfully. If the creation fails, you can [submit a ticket](#) to report your problems and receive assistance.

How do I terminate a CVM if it fails to be delivered?

You can [submit a ticket](#) to contact Customer Service and provide complete screenshots of the server information and termination failure indicating **delivery failure** to facilitate troubleshooting.

Renewing CVMs

How do I renew an expired CVM?

For more information, see [Renewing Instances](#).

How do I set auto renewal for CVMs?

For more information, see the steps for auto renewal in [Renewing Instances](#).

Do pay-as-you-go instances need to be renewed?

For pay-as-you-go instances, charges are automatically deducted from the account every hour, so renewal is not needed.

Other

Can I set auto release (termination) for CVM instances?

Yes. You can set auto release (termination) for pay-as-you-go instances and specify when to release them automatically. For more information, see [Terminating Instances](#).

If you have additional questions, please [submit a ticket](#).

I have purchased a Linux CVM with cloud disks that are over 20 GB. How will it be charged if I reinstall the operating system as Windows?

The billing for the parts exceeding 20 GB of the previously purchased system disks will stop once the Windows OS is successfully installed.

I have purchased a Windows CVM with cloud disks. How will it be charged if I reinstall the operating system as Linux?

Because the capacity of the system disk cannot be reduced, it needs to be reinstalled at the current capacity. You will not incur additional charges. If you need to expand the system disk capacity after reinstalling the system as Linux, you will incur additional fees. For more information, see [System and Data Disks](#).

For more information on cloud disk prices, see [Pricing List](#).

Note:

Changing OS is not supported for CVMs in regions outside the Chinese mainland.

What do the computing components in a CVM bill mean?

The computing components correspond to instance specifications. For example, the computing components for S5.SMALL4 include the CPU, the memory, and the NVMe local disk.

Instance

Issues on Logging in and Connecting to Instances

Last updated : 2024-01-06 17:26:45

How do I log in to a CVM instance using VNC?

Tencent Cloud allows users to log in to a CVM instance by using the VNC web client. If you do not have a remote login client installed or cannot log in using the client, you can use a VNC web client to log in to a CVM instance. You can use the web client to log in remotely, view the CVM instance status, and perform basic instance management. For detailed instructions, refer to the following documents:

[Logging into Linux Instances via VNC](#)

[Logging into Windows Instances via VNC](#)

How do I enable multi-user remote login on a Windows server?

The Windows server supports multi-user remote login. For detailed configuration directions, refer to [Setting Up Multi-user Remote Login for Windows CVMs](#).

If you fail to enable multi-user remote login on a Windows server after following the directions in the document above, restart the instance and try logging in again.

How can I log in to an instance running Ubuntu as root?

The default user for Ubuntu is `ubuntu`. `root` is not enabled during the installation process by default. If necessary, you can enable `root` by following the directions below:

1. Log in to the CVM instance using `ubuntu`.
- 2.

Run the following command to set a password for `root`:

```
sudo passwd root
```

3. Enter a password for `root` and press **Enter**.
4. Enter the password again and press **Enter**.

If the following message is displayed, the password was set successfully:

```
passwd: password updated successfully
```

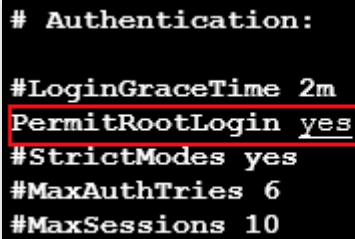
5. Run the following command to open `sshd_config`:

```
sudo vi /etc/ssh/sshd_config
```

6. Press **i** to switch to the editing mode. Find `#Authentication` and change the value of `PermitRootLogin` to `yes`, as shown in the following figure:

Note:

If `PermitRootLogin` is commented, remove the comment marks (`#`).



```
# Authentication:

#LoginGraceTime 2m
PermitRootLogin yes
#StrictModes yes
#MaxAuthTries 6
#MaxSessions 10
```

7. Press **Esc** and enter `:wq` to save the file and exit vi.

8. Run the following command to restart the SSH service.

```
sudo service ssh restart
```

9. Refer to [Logging into Linux Instances Using the Standard Login Method](#) on how to use the following information to log in to your CVM instance running Ubuntu.

Username: root

Password: the password you set in [Step 2](#).

A successful login using root is shown in the following figure:

```
* Socket connection established *
Welcome to Ubuntu 18.04.3 LTS (GNU/Linux 4.15.0-46-generic amd64)

* Documentation:  https://help.ubuntu.com
* Management:    https://landscape.canonical.com
* Support:        https://ubuntu.com/advantage

System information as of Tue Jan 21 09:27:02 CST 2020

System load:  0.08               Processes:            89
Usage of /:   4.6% of 49.15GB    Users logged in:     0
Memory usage: 11%               IP address for eth0: 172.31.0.12
Swap usage:   0%

* Overheard at KubeCon: "microk8s.status just blew my mind".

https://microk8s.io/docs/commands#microk8s.status

Last login: Tue Jan 21 08:46:18 2020 from 172.31.0.12
root@VM-0-12-ubuntu:~#
```

What do I do if I cannot connect (log in) to my CVM instance after restarting it?

This problem may have been caused by the high CPU and memory usage of your instance. For more information, please refer to the following documents:

[Failing to Log in to a Linux CVM Due to High CPU and Memory Usage](#)

[Failing to Log in to a Windows CVM Due to High CPU and Memory Usage](#)

What do I do when I connect to my CVM instance remotely and the connection times out?

Make sure that:

The instance is running.

The instance is associated with security groups with the necessary security group rules. For detailed information, refer to [Adding Security Group Rules](#).

The instance has enabled SSH or RDP.

The instance has opened the corresponding ports. SSH uses port 22 by default and RDP uses port 3389 by default.

Why is my remote connection to my Linux instance refused?

Make sure that:

The instance has enabled SSH or RDP.

The instance has opened the corresponding ports. SSH uses port 22 by default and RDP uses port 3389 by default.

Why do I get the invalid username or password message when I remotely connect to my Linux instance?

Make sure that:

You used the correct username. The default username for most Linux distributions is `root`. However, Ubuntu uses `ubuntu`.

You entered your password correctly. If you forgot your password, reset it. For detailed directions, refer to [Resetting Instance Passwords](#).

Why do I get the invalid username or password message when I remotely connect to my Windows instance?

Make sure that:

You used the correct username. The default username for Windows is `Administrator`.

You entered your password correctly. If you forgot your password, reset it. For detailed directions, refer to [Resetting Instance Passwords](#).

To use a user without administrator privileges to log in, that user needs to belong to the Remote Desktop Users group.

Does the Console support multi-user login via VNC?

No. If one user has already logged in, other users will not be able to log in.

What do I do if I forgot my CVM instance login password?

You can reset your password. For detailed directions, refer to [Resetting Instance Passwords](#).

Why can't I use Internet Explorer 8.0 to log in to my instances?

The VNC web client supports Internet Explorer 10 or later. Please download the latest version of Internet Explorer. In addition, we recommend you use Google Chrome because the Tencent Cloud Console offers a better user experience on Google Chrome than on Internet Explorer.

How do I log in to my Linux instance remotely?

Tencent Cloud recommends [logging in to a Linux instance using the standard login mode](#). You can also:

[Log in to a Linux instance using remote login software](#).

[Log in to a Linux instance using SSH](#).

[Log in to a Linux instance using VNC](#).

What do I do if I cannot connect to my Linux instance?

If you cannot connect to your Linux instance, refer to [Linux Instance Login Failures](#) for troubleshooting instructions.

If this problem persists, please [submit a ticket](#).

What do I do if I cannot connect to my Windows instance?

If you cannot connect to your Windows instance, refer to [CVM Login Failure](#) for troubleshooting instructions.

If this problem persists, please [submit a ticket](#).

What do I do if I find that my CVM instance was logged into from unusual locations?

If you find that your CVM instance was logged into from locations you do not recognize:

1. Double check the login time from the unusual login location and check whether you or other administrators logged in at that time.
2. If no, follow the steps below:
 - 2.1 [Reset the password](#) immediately.
 - 2.2 Check your instance for viruses or trojans.
 - 2.3 Use security groups to [limit logins to specific IP addresses](#).

Purchasing

Last updated : 2024-01-06 17:26:45

How do I purchase a CVM?

You can purchase both Tencent Cloud services and pay-as-you-go CVMs (billing is accurate to the seconds and settled on an hourly basis) on the Tencent Cloud Console. For more information, see [Billing Modes](#).

What instance types does CVM provide?

CVM provides multiple instance specifications. For more information, see [Instance Types](#). You can select the appropriate instance type based on your business needs.

To manage sudden business surges, you can choose the pay-as-you-go billing method, which allows you to activate/terminate instances at any time and only pay for resources that you actually use. Billing is accurate to the seconds to maximize your savings.

Can I purchase a Windows 2003 CVM?

Because Microsoft discontinued support for Windows Server 2003 and Windows Server 2003 R2 since July 14th, 2015, Tencent Cloud no longer provides the Windows 2003 server, which is not available for purchase.

What are some channels through which I can purchase CVMs?

You can purchase CVMs either from the Console or via an API.

When can I use the CVM I purchased?

After the CVM system installation is completed, you can log in to and use the CVM once its status becomes **running**.

How do I check which instances are available for purchase in a specific region or availability zone?

You can go to the [CVM purchase page](#) to view the instance types available in each region or availability zone.

What can I do if the resource of an instance I want to purchase is sold out?

If the resource of a specific instance you want to purchase is sold out, we recommend you:

Select another region

Select another availability zone

Select another resource

If the resource that you want to purchase is still not available after you take all of the measures above, try again later.

Instance resources are dynamic. When resources are insufficient, Tencent Cloud will replenish them as soon as possible.

How do I select a CVM instance that is suitable for my business?

Consider the following factors when you select a CVM instance. For more information about how to select a suitable CVM instance for your business, see [Instance Types](#).

Your business needs.

Your website type.

The average number of page views (PV) per day on your website.

The size of your homepage.

The data capacity of your website.

How do I pay for a CVM instance?

You can pay for a CVM instance [online](#) using your credit card.

How long does it take to launch a CVM?

Generally, it takes 1 to 2 minutes to create an instance. After the instance is created:

You can connect to the Linux instance by [logging in to the Linux instance](#).

You can connect to the Windows instance by [logging in to the Windows instance](#).

Note:

If an error occurs during instance creation, please [submit a ticket](#).

I paid for a CVM instance but the instance was not created. Why did this happen?

If there are insufficient resources to create an instance of your selected type within the availability zone, the instance will fail to be created. Your account will automatically be refunded the cost of the instance. If you do not receive a refund within 30 minutes, please [submit a ticket](#).

You can go to the [CVM purchase page](#) to view the instance types available in each availability zone.

A CVM Instance Cannot Be Found

Last updated : 2024-01-06 17:26:45

Problem Description

A purchased CVM instance cannot be found in the CVM console.

Possible Causes

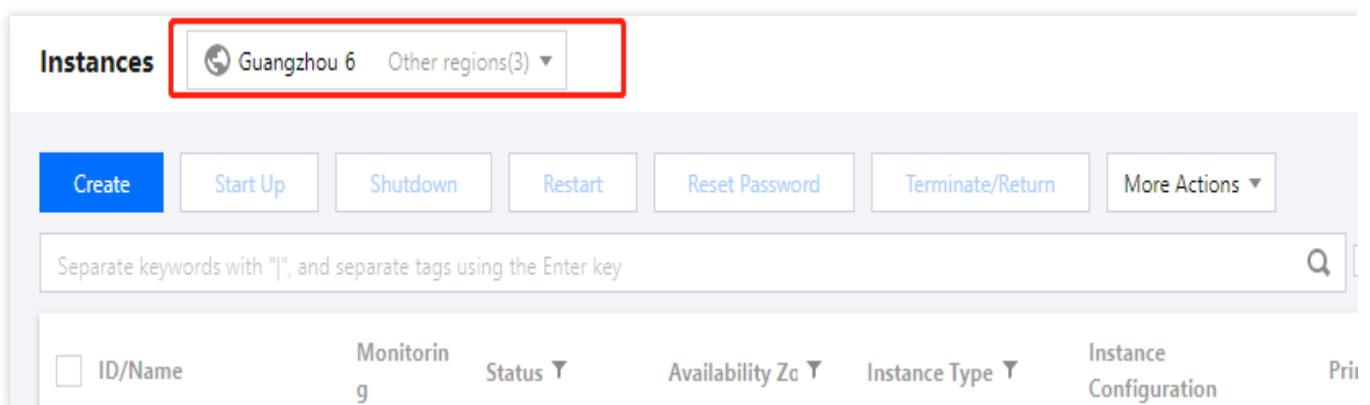
1. [The instance is not located in the current region](#)
2. [You did not log in to the right console](#)
3. [No instance is available under the current account](#)
4. [The instance is released due to overdue payment or expiration](#)
5. [The spot instance is automatically repossessed](#)
6. [Refund is caused by insufficient resources](#)

Solutions

1. The instance is not located in the current region

If you did not switch to the region where the instance is located, you cannot find the instance. You can check this with the following method.

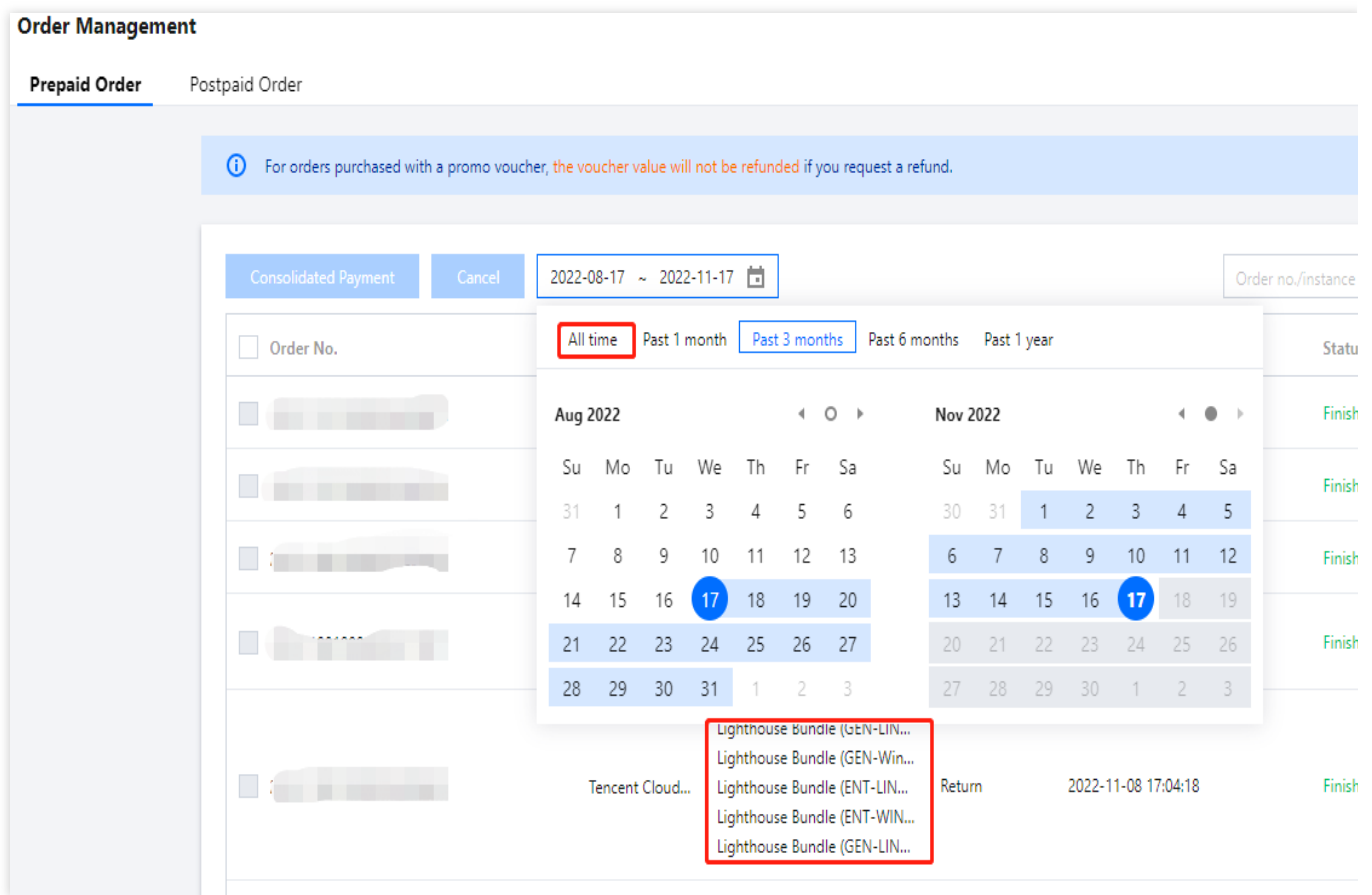
Log in to the CVM console, and click [Instances](#) in the left sidebar. Switch to another region to check the instance at the top of the page.




2. You did not log in to the right console

You maybe purchased a [Lighthouse instance](#) or another instance instead of a CVM instance, and you did not log in to the right console.

Method 1: Go to [Overview](#) > **Order Management**, and check the prepaid and postpaid orders for the last three months. If the instance was purchased long ago, please click [All time](#).



Method 2: View all the resources under your current account by clicking [Overview](#) > **My Resources** in the console.



3. No instance is available under the current account

You can find your account ID in an SMS you received related to CVMs.

The instance may not be purchased with an account under your name.

4. The instance is released due to overdue payment or expiration

If an instance is expired or with overdue payment, it will be released, and so you cannot find it in the console.

Terminated instances cannot be recovered. You can check the Prepaid Service Expiration and Suspension Notification via [Message Center](#). If you did not receive a relevant message in Message Center, maybe you have unsubscribed the notification.

If your account has overdue payment, please see [Payment Overdue](#) for status of the CVM instances and public network bandwidth.

5. The spot instance is automatically repossessed

An allocated [Spot Instance](#) may be repossessed based on the prices or the relationship between supply and demand. You will be notified via Message Center or metadata before the spot instance is terminated. You can check Notification on Spot Instance Repossessing via [Message Center](#). If you did not receive a relevant message in Message Center, maybe you have unsubscribed the notification.

6. Refund is caused by insufficient resources

If you cannot find the newly purchased instance in the console, maybe it is because a refund is caused by insufficient resources.

You can check the refund in [Order Management](#), and then purchase again in another region or availability zone.

Usage

Last updated : 2024-01-06 17:26:45

How do I view CVMs that are currently in use?

You can log in to the [CVM console](#) to view CVMs that are currently in use.

Can VM be installed on a CVM?

No.

How do I shut down an instance?

For more information, see [Shutdown Instances](#).

How do I restart an instance?

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

How do I terminate an instance?

For more information, see [Terminating Instances](#).

How do I query the username and password of a Linux instance?

After you create a CVM instance, its username and password will be delivered to your account through the [Message Center](#). The admin account of a Linux instance is `root` by default.

How do I query, partition, and format the disks of a Linux instance?

You can run the `df -h` command to query the total and used disk capacities and run the `fdisk -l` command to view the disk information. For information about how to partition and format the disks of Linux instances, see [Initializing Cloud Disks \(Smaller than 2TB\)](#) and [Initializing Cloud Disks \(Larger than 2TB\)](#).

How do I upload files to a Linux instance?

You can [use SCP to upload files to a Linux instance](#).

You can [use FTP to upload files to a Linux instance](#).

How do I change the owners and owner groups of directories and files on a Linux instance?

If the file or directory permissions are incorrectly configured on the Web server, a 403 error will occur when you access a website hosted on the instance. Therefore, before you adjust a file or directory, you must confirm the identity under which the file or directory process is running.

You can run the `ps` and `grep` commands to query the identities under which the file or directory process is running.

You can run the `ls -l` command to query the owners and owner groups of files and directories.

You can run the `chown` command to modify the permissions. For example, you can run the `chown -R www.www /tencentcloud/www/user/` command to change the owners and owner groups of all files and directories under the `/tencentcloud/www/user` directory to account "www".

Do Linux instances support GUI?

Yes. If you need to build a GUI in a Linux instance, see [Building a Visual Ubuntu Desktop](#).

Why can't I add sound or video cards to CVM instances?

Tencent Cloud CVMs do not provide multimedia servers and sound and video card components by default. Therefore, sound and video cards cannot be added to CVM instances.

Can I specify a MAC address when purchasing a CVM instance?

No. A MAC address is randomly assigned when a CVM instance is created and cannot be specified.

Can I transfer the unused time of a CVM instance to another CVM instance?

No. If you want higher flexibility and cost-efficiency, we recommend that you purchase pay-as-you-go instances.

How do I query the region where the IP address of a CVM instance is located in?

The IP address of a CVM instance is located in the same region where you purchased the CVM instance.

Do CVM instances provide databases by default?

No. To use database services, you can:

Deploy your own database. For example, you can [install and build MySQL](#).

Purchase [TencentDB for MySQL](#) separately.

Configure the environment database using the image marketplace.

Can I build a database on a CVM instance?

Yes. You can install database software and configure a database environment on a CVM instance as needed. You can also purchase [TencentDB for MySQL](#) separately.

When can I forcibly stop a CVM instance? What are the consequences of forcibly stopping a CVM instance?

You can forcibly stop a CVM instance when normal shutdown fails. Please note that forced shutdown is equivalent to a power outage and can result in the loss of unsaved data.

Use Limits

Last updated : 2024-01-06 17:26:45

What are the limitations of changing public IP addresses of CVM instances?

The limitations of changing public IP addresses of CVM instances are:

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

Each instance can only change its public IP once.

The old public IP will be released after the public IP is changed.

How many websites can be hosted on a CVM instance?

A CVM instance can host up to 5 websites, and each website can bind to multiple domain names.

Can a CVM instance access the Amazon website?

Yes, as long as your CVM instance can connect to the public network properly.

How do I purchase more pay-as-you-go CVMs?

If you have reached the maximum number of pay-as-you-go instances that you can purchase, you cannot purchase more pay-as-you-go instances. For more information, see [Purchase Limits](#).

How can I query the CVM resource quota?

To query the use limits and quotas of CVM resources, see CVM instance use limits in [Use Limits Overview](#).

Adjusting Configuration

Last updated : 2024-01-06 17:26:45

How do I upgrade/downgrade the configurations of CVMs?

You can only adjust the configurations of instances **whose system disks and data disks are both cloud disks**.

For more information about how to upgrade/downgrade instance configurations, see [Changing Instance Configurations](#).

For more information about how to adjust bandwidth/network configurations, see [Adjusting Network Configurations](#).

If your configuration adjustments do not take effect, please [submit a ticket](#) to contact us.

How do I check the records of configuration adjustments?

1. Log in to the [CloudAudit Console](#).
2. On the **Event history** page, set filters such as the username, the resource type, and the resource name as needed to view the record list.

For more information, see [Getting Started](#).

Can I adjust the configurations of CVM instances?

Yes. You can adjust the configurations of instances whose system disks and data disks are both cloud disks. For pay-as-you-go instances, you can upgrade or downgrade their configurations as many times as you want.

What is the maximum number of times I can downgrade the configurations of CVMs?

You can upgrade or downgrade the configurations of pay-as-you-go instances as many times as you want.

Can I upgrade the specifications and the configurations of pay-as-you-go instances?

Yes. You can upgrade the configurations of pay-as-you-go CVM instances by following the directions listed in [Changing Instance Configurations](#) or through the [ResetInstancesType](#) API.

How long does it take to upgrade a CVM instance?

About 1 to 2 minutes.

How is the cost incurred for a CVM instance upgrade calculated?

You can view the cost details at the [Billing Center](#) after you upgrade the specifications or the configurations of the CVM instance.

Will upgrading CVM instances affect my business configurations on the instance?

After you upgrade a CVM instance, you must restart it to validate the new configurations. The upgrade operation will interrupt your business for a short period of time. We recommend that you upgrade instances when business is slow.

After being upgraded, instances will seamlessly resume your business and will not require environment reconfiguration.

Why is the estimated refund of a CVM downgrade 0?

One possible reason is that you have purchased the instance at a discount but the downgraded configuration is calculated according to the original price of the instance. When the price of the instance purchased at a discount is lower than or equal to the original price of the instance, the estimated refund will be displayed as 0.

Why did the instance upgrade not take effect?

After you upgrade an instance, you must restart it on the Console or through an API.

Spot Instances

Last updated : 2025-03-18 10:27:19

Instance Release

Why is a spot instance released automatically?

An important feature of spot instances is that the system will repossess assigned instances based on prices or the supply-demand relationship. If the market price is higher than your bid or if the CVM resource pool corresponding to your spot instances is in short supply, the process will be interrupted by the system.

Is it possible to avoid being repossessed by the system through bidding?

No. Because repossession triggered by insufficient inventory is unavoidable. You need to accept that instance repossession may occur when you deploy businesses on the spot instance.

How do I know that an instance is about to be interrupted?

Two minutes before the interruption, we will notify you in the form of metadata that the instance is about to be interrupted and repossessed.

For more information, see [Querying the Repossession Status of a Spot Instance](#).

How to automatically apply for spot instances after inventory recovery?

You can use cloud products that can automatically maintain the CVM cluster, such as [BatchCompute](#), [Auto Scaling](#). With their cross-model and cross-availability zone capabilities, you can maintain a specified number of CVM clusters more effectively.

Price and Billing

What are the similarities and differences between spot instances and pay-as-you-go instances?

Billing Method	Similarities	Differences
Spot instances	Both of them are pay-as-you-go. There is no need to pay in advance but certain costs must be frozen. You can enable/terminate the CVM at any time and pay according to actual usage. The billing time granularity is accurate to the second, and the account will be settled every hour.	Price: In the vast majority of cases, the price of spot instances is 3%-20% of the pay-as-you-go price under the same specification, and fluctuates in response to market supply and demand dynamics.
Pay as you go		

		<p>Release mechanism: The lifecycle of a pay-as-you-go instance is controlled by the user, while spot instance may be actively repossessed by the system.</p> <p>Feature limitations: Spot instances Configuration adjustments are not allowed.</p>
--	--	---

Which price, the market price and the highest bid specified by the user, will be used for billing?

The market price will be used for billing. You can specify a high bid to prevent instances from being repossessed due to the price. However, the system will only charge you at the current market price (the current market price will be fixed).

How are billing periods calculated for spot instances?

You will be billed for the period from the moment you apply for a spot instance to the moment the spot instance is manually released or interrupted by the system. The billing period is accurate to the second.

Where can I find the current market prices of all the spot instances?

During the beta testing period, we cannot provide a page where you can query the market prices of all instances, but it will be available in the future. Currently, most of the spot instances will be priced at 20% of the regular pay-as-you-go instances of the same model and specification.

Where can I view the consumption details regarding spot instances?

As with pay-as-you-go instances, you can find detailed usage and billing information of spot instances in **Billing Center > Bills** at the top of the console. Spot instances are pay-as-you-go services.

Quotas and Limitations

In which regions are spot instances available? Which instance models and specifications do spot instances support?

Region	Models Supported by Spot Instances	Discounts
Beijing, Shanghai, Chengdu, Chongqing, and Guangzhou Open	All models supported by pay-as-you-go	3%-20% off the published prices of pay-

Guangzhou (excluding Guangzhou Zone 1)	instances	as-you-go instances with the same specifications
Hong Kong (China), Singapore, Bangkok, Seoul, Tokyo, Silicon Valley, Virginia, and Frankfurt		

Are quota limits of spot instances shared by pay-as-you-go instances?

No. Each user can have up to 50 spot instance vCPU cores in each availability zone. To raise the quota limits, submit a ticket.

Can I upgrade or downgrade the specifications of spot instances?

No.

Do spot instances support no charges when shut down?

No.

Do spot instances support system re-installation?

No.

Can I use vouchers for purchasing spot instances?

No.

Big Data Instance

Last updated : 2024-01-06 17:26:45

What is a big data instance?

Big data instances are designed specifically for scenarios like Hadoop distributed computing, massive log processing, distributed file system, large data warehouse, and more. This CVM instance type is mainly used to solve the cloud computing and storage problems of massive business data.

Which industry customers and business scenarios are big data instances applicable to?

Big data instances are applicable to customers in the Internet, game, finance and other industries that require big data computing and storage analysis, as well as business scenarios that require massive data storage and offline computing. They can meet the storage, capacity and private network bandwidth requirements of distributed computing businesses represented by Hadoop.

In addition, with the highly available architectural framework of distributed computing services such as Hadoop, big data instances use local storage, making the total cost close to that of a self-built Hadoop cluster in IDC, while ensuring massive storage capacity and high performance.

Features

A single instance has a throughput capacity up to 2.8 GB/s. HDD local disk is ideal for throughput-intensive storage. With stable and high-performing sequential read/write throughput, big data instances are designed specifically for Hadoop distributed computing, massive log processing, large data warehouse and other business scenarios. Local storage has a unit price as low as 1/10 of S2 instances, making its total cost close to that of a self-built Hadoop clusters in IDC while ensuring massive storage capacity and high performance. Big data instances deliver the optimal cost-efficiency for big data scenarios.

Read/write latency is as low as 2-5ms. Big data instances are high-performing models suitable for enterprise developers.

It supports the pay-as-you-go billing method.

Specifications

Note:

For more information on instance specifications, see the “Big Data Family” section in [Instance Types](#).

Notes on local data storage

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

The table below shows you the local disk data status after you perform different operations on an instance with the local disk.

Operation	Local Disk Data Status	Description
Log in to an instance to restart it, restart an instance on the console, or forcibly restart an instance	Retained	Local disk storage as well as the data is retained.
Log in to an instance to shut it down, shut down an instance on the console, or forcibly shut down an instance	Retained	Local disk storage as well as the data is retained.
Terminate (instance) on the console	Erased	Local disk storage is erased. No data is retained.

Note:

Do not store business data that needs to be retained for a long time on a local disk. Back up data in advance and use a highly available architecture. We recommend you store the data on a CBS disk for long-term retention.

How can I purchase a local disk for big data instances?

Local disk can only be purchased when a big data instance is created. The instance specifications determine the number and capacity of local disks you can purchase.

Does the local storage of big data instances support snapshot?

No.

Does a big data instance support configuration adjustment and failover?

No.

Big data instances feature massive data storage and use local HDD as data disk. To prevent data loss (when the host crashes or local disk is damaged), we recommend you use a redundancy policy, for example, a file system that supports redundancy and fault tolerance (such as HDFS and MapR-FS). We also recommend you regularly back up data to a persistent storage system, such as Tencent COS. For more information, please see [Cloud Object Storage](#). If the local disk is damaged, you will need to shut down the CVM instance for us to replace the local disk. We will notify you and perform fixes if the CVM instance crashes.

What is the difference between the Big Data family and High IO I2 instances?

High IO I2 instances featuring ultra-high IOPS are designed specifically for business scenarios with low latency and high random IO. They are suitable for high-performance database (relational database, NoSQL, etc.). Big data instances are designed specifically for business scenarios that require high sequential read/write and low-cost massive data storage. This type features high storage cost-efficiency and private network bandwidth.

How is the disk throughput of big data instances?

Big data D2 instances boast local disks with the sequential read/write throughput as follows.

For a single disk, the sequential read throughput is 220+ MB/s and sequential write throughput is 220+ MB/s (block size of 128 KB and queue depth of 32).

For the CVM, the throughput can reach up to 2.8 GB/s (block size of 128 KB and queue depth of 32).

What is the difference between the local disk of big data instances and CBS?

[Cloud Block Storage](#) is a highly available, highly reliable, low-cost, and customizable block storage device. It can be used as an independent and scalable disk for CVM, providing efficient and reliable storage devices. It provides data storage at the data block level and employs a three-copy distributed mechanism to ensure data reliability for CVM instance, meeting the requirements of different use cases. The local disk of big data instances is designed specifically for business scenarios that require high sequential read/write for massive local data sets, such as Hadoop distributed computing, large-scale parallel computing, and data warehouse.

Others

Last updated : 2024-01-06 17:26:45

How do I view the operation logs of a CVM?

You can view the operation logs of a CVM on the console as follows:

1. Log in to the [CVM console](#) and select **Tool > Operation Log** in the top-right corner.
2. Select **CVM** for the resource event and configure other filters as needed.
3. Click **Query** to obtain the CVM operation logs.

For more information about CloudAudit operation records, see [Viewing the Event Details of Operation Records](#).

What if I cannot find my CVM in the console?

The reasons why a purchased CVM instance cannot be found in the CVM console include:

1. The instance is not located in the current region.
2. You did not log in to the right console.
3. No resource is available under the current account.
4. The instance is released due to overdue payment or expiration.
5. The spot instance is automatically repossessed.
6. Refund is caused by insufficient resources.

For detailed solutions, see [A CVM Instance Cannot Be Found](#).

How to view the CPU frequency of a CVM?

The frequency viewed via `/proc/cpuinfo` on a Linux instance is only the base frequency, because some privileged registers on the virtual machine cannot be accessed and the underlying real-time running frequency information of the CPU cannot be read. For information on CPU models and frequencies of instances, see [Instance Types](#).

How to enable turbo frequency for a CVM?

Turbo frequency is enabled automatically when a CVM is running. The processor will automatically adjust the running frequency based on the CPU loading requirements of the application. The max turbo frequency can be reached when a CPU intensive application is running.

Storage

System Disks

Last updated : 2024-01-06 17:26:45

What is the default capacity of a CVM system disk?

Currently, a new CVM system disk has a default capacity of 50 GB.

Can I change a CVM system disk from a local disk to a cloud disk?

When purchasing a CVM instance, you can select the desired disk type for the CVM system disk.

After purchasing a CVM instance if there are available cloud disks in the availability zone where the purchased CVM resides, you can use the [change disk media type](#) feature to change the local system disk to a cloud disk.

Which regions and availability zones support expanding the system disk capacity to more than 50 GB?

If the system disk is a cloud disk, the system disk capacity can be adjusted to more than 50 GB in all regions that support snapshots.

Can I expand the system disk capacity of a CVM instance when reinstalling the operating system?

It depends on the system disk type.

If the system disk is a cloud disk,
the system disk capacity cannot be expanded.

If the system disk is a local disk,
it depends on the system disk size.

If the default system disk capacity of the purchased instance is 50 GB, the system disk cannot be expanded.

For instances purchased at an earlier time: if the system disk capacity is 20 GB or less, the system disk can be expanded to 20 GB by default. If the system disk capacity is more than 20 GB, the system disk can be expanded to 50 GB by default.

Can I reinstall the operating system to reduce the system disk capacity after the expansion?

The capacity of a system disk cannot be reduced.

How do I save the data on the CVM instance and then expand the system disk capacity?

For detailed directions, see [Expanding Cloud Disks](#).

What is the system disk capacity if I use an image of less than 50 GB to create or reinstall the CVM?

The image capacity does not affect the system disk capacity. The minimum system disk capacity is 50 GB.

Does the CVM system disk support partitioning?

Yes, but we recommend you not partition a system disk.

If you partition the system disk with a third-party tool, system crash or data loss may occur. You can perform partitioning after expanding the system disk capacity.

Using Cloud Disks

Last updated : 2024-01-06 17:26:45

How can I view the data disk?

1. Log in to the [CVM console](#).
2. Click **Cloud Block Storage** in the left sidebar to access the **Cloud Block Storage** management page.
3. Click the filter icon next to the **Attribute** column title, select **Data disk** and click **OK** to view all data disks in the region.

How do I read and write the original NTFS data disk after the Windows operating system is reinstalled to Linux?

Windows employs two major file systems, NTFS or FAT32, while EXT is the file system for Linux. When an operating system is changed from Windows to Linux after reinstallation, the data disk remains in its original format. Therefore, the system might be unable to access data disk's file system. In these cases, you will need to use a format converter to read the data disk. For details, see [Reading/Writing NTFS Data Disks after Reinstalling a Windows CVM to Linux CVM](#).

How do I read the data disk in EXT format after the Linux operating system is reinstalled to Windows?

Windows employs two major file systems, NTFS or FAT32, while EXT is the file system for Linux. When an operating system is changed from Linux to Windows after reinstallation, the data disk remains in its original format. Therefore, the system might be unable to access data disk's file system. In these cases, you will need to use a format converter to read the data disk. For details, see [Reading/Writing EXT Data Disks after Reinstalling a Linux CVM to Windows CVM](#).

What are the differences among Premium Cloud Storage, SDD and Enhanced SSD?

Premium Cloud Storage: Tencent Cloud Premium Cloud Storage is a hybrid storage type. It adopts the Cache mechanism to provide a high-performance SSD-like storage, and employs a three-copy distributed mechanism to ensure data reliability. Premium Cloud Storage is suitable for small and medium applications with high requirements for data reliability and standard requirements for performance.

SSD: SSD uses NVMe SSD as the storage media, and employs a three-copy distributed mechanism. It provides high-performance storage with low latency, high random IOPS, high throughput I/O, and data security up to 99.9999999%, making it suitable for applications with high requirements for I/O performance.

Enhanced SSD: Enhanced SSD is based on Tencent Cloud's latest storage engine, NVMe SSD storage media and the latest network infrastructure. It employs a three-copy distributed mechanism to provide high-performance storage with low latency, high random IOPS, high throughput I/O, and data security up to 99.9999999%, making it suitable for I/O-intensive applications with high requirements for latency, such as large databases and NoSQL.

Pricing of the three types of cloud disks varies with region. You can select the cloud disk that best suits your application requirements and budget. For pricing details, see [Price Overview](#).
For more information on disk types and performance, see [Cloud Disk Types](#).

How do I test the disk performance?

We recommend using FIO to perform stress test and verification on cloud disks. For instructions, see [Measuring Cloud Disk Performance](#).

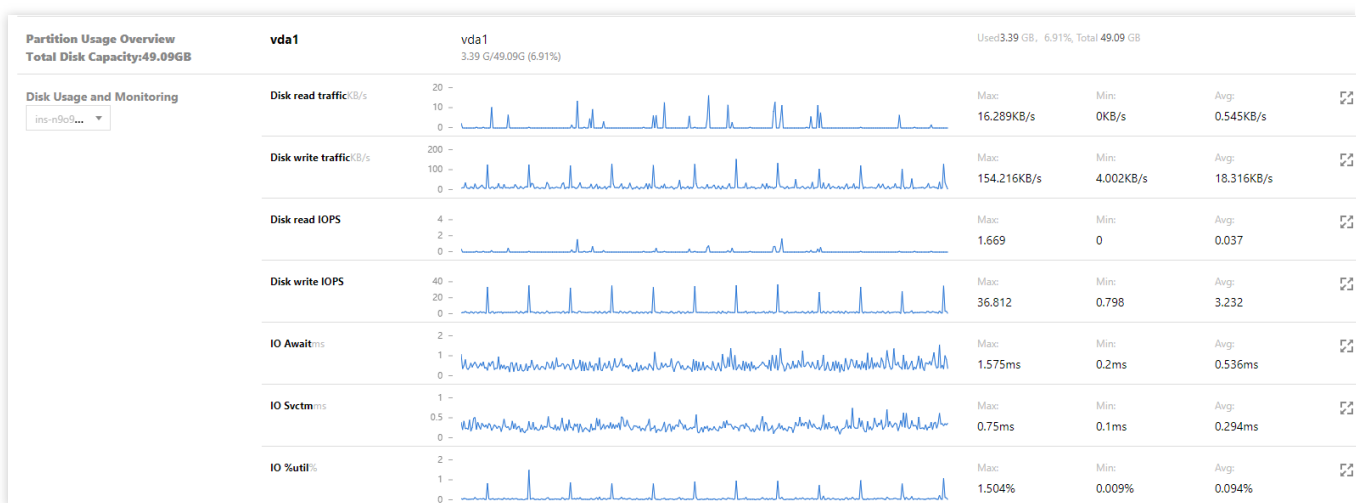
What are the most common cloud disk operations?

For common cloud disk operations, see [Operation Overview](#).

How do I check the available and used space on the cloud disks?

You can log in to the CVM instance to check the available and used space on the cloud disks. This information can also be seen on the CVM console as follows.

1. Log in to the [CVM console](#) and access the **Instances** page.
2. Select the ID/Name of the target instance to access the details page.
3. Click on the **Monitoring** tab to view the instance disk usage:



Why was the cloud disk I separately created released together with my instance?

When mounting a cloud disk, you can decide if it should be released with the instance automatically. This can be configured via the [CBS console](#) or [ModifyDiskAttributes](#) API.

How do I partition and format the a mounted cloud disk?

For more information, see [Initializing Cloud Disks \(Smaller than 2TB\)](#) or [Initializing Cloud Disks \(Larger than 2TB\)](#).

Mounting and Unmounting Cloud Disks

Last updated : 2024-01-06 17:26:45

What is a device name (mount point)?

A device name (mount point) is the location of a cloud disk mounted to the CVM instance on the disk controller bus. The selected device name matches the disk device number in Linux and matches the disk sequence number of the disk manager in Windows.

Can I mount one cloud disk to multiple CVM instances?

No. You can mount up to 20 cloud disks to the same CVM, but you cannot mount the same cloud disk to multiple CVMs. You can only share data by [unmounting the cloud disk](#) from CVM A and then [mounting](#) it to CVM B.

Do I need to partition the cloud disk after purchasing and mounting it to a CVM instance?

Yes. After you purchase a cloud disk, you must mount it to a CVM instance in the same availability zone, and then initialize it by formatting, partitioning, and creating the file system before using it as a data disk. For more information, see [Mounting Cloud Disks](#) and [Initializing Cloud Disks](#).

Why am I unable to locate the data disk that I purchased for a Linux instance?

If you purchase a data disk separately, you must partition, format, and mount it to an instance to view and use its storage space. For more information, see [Mounting Cloud Disks](#) and [Initializing Cloud Disks](#).

How many cloud disks can be mounted to one CVM instance?

A maximum of 20 data disks can be mounted to one CVM instance.

Why can't I locate the CVM to which I want to mount a cloud disk?

Make sure that the CVM instance has not been released and ensure that it is in the same region and availability zone as the cloud disk is in.

Can I mount a cloud disk to a CVM instance in another availability zone?

No. A cloud disk can only be mounted to or unmounted from a CVM instance in the same availability zone as the cloud disk is in.

When I unmount a cloud disk (data disk), will its data be lost?

Data in cloud disks will not be modified during mounting or unmounting. To ensure data consistency, we recommend that you follow the steps below:

In Windows, we recommend that you stop all read and write operations on all file systems of the cloud disk to ensure data integrity. Otherwise, the data that has not finished being read or written will be lost.

In Linux, log in to the CVM instance and run the `umount` command on the cloud disk. After the command is executed, log in to the CVM Console to unmount the cloud disk.

How do I unmount a cloud disk?

For more information, see [Unmounting Cloud Disks](#).

Can cloud disks be mounted and unmounted?

Cloud disks can be mounted and unmounted.

System disks cannot be mounted or unmounted.

Can cloud disks be mounted and unmounted in batches?

Cloud disks can be mounted and unmounted in batches.

System disks cannot be mounted or unmounted.

Can system disks be unmounted?

No.

Resizing Cloud Disks

Last updated : 2024-01-06 17:26:45

How do I expand the capacity of a cloud disk?

If your CVM uses a cloud disk, you can expand the disk capacity. For more information on how to expand the disk capacity, see [Expanding Cloud Disks](#).

Can I reduce the capacity of a cloud disk?

No. If you want to reduce the capacity of a cloud disk you purchased, we recommend that you create a cloud disk of your desired capacity and mount it to the same instance that the original cloud disk is mounted to. Then, copy the data stored in the original disk to the new disk and release the original disk.

How do I expand a system disk?

For data security reasons, a CVM system disk cannot be expanded directly on the Console. You must [reinstall the operating system of the instance](#) to expand the capacity of its system disk.

Can the system disks of all types of cloud disks be expanded?

Yes. The system disks of all types of cloud disks, including SSD, Premium Cloud Storage, and HDD, can be expanded.

Can the system disks of pay-as-you-go CVM instances be expanded?

Yes. The system disks of pay-as-you-go CVM instances can be expanded.

What is the capacity range of the system disk? What is its maximum capacity?

The expanded system disk's capacity must be greater than its existing capacity but less than or equal to 500 GB.

Snapshots

Last updated : 2024-01-06 17:26:45

Is snapshot available in all availability zones?

Yes.

Will creating a snapshot affect the disk performance?

Creating a snapshot will occupy a small amount of the cloud disk I/O. We recommend creating snapshots when business is slow.

How long does it take to create a snapshot?

The time it takes to create a snapshot is influenced by factors such as the number of disk writes and the underlying read-write operations. Creating a snapshot will not affect your disk use.

Do I need to shut down the CVM to roll back snapshots?

For a cloud disk that has been mounted to a CVM, you have to shut down the CVM to roll back snapshots.

For a cloud disk that has not been mounted, you can directly roll back snapshots.

How is the size of the first full snapshot of a cloud disk calculated?

The first snapshot created on a cloud disk is a full snapshot that copies all the data of the cloud disk at a point in time. The snapshot size equals the used capacity of the cloud disk. For example, if the total capacity of a cloud disk is 200 GB and 122 GB have been used, the size of the first full snapshot is 122 GB.

Can I download or export CVM instance snapshots to a local device?

No. Snapshots cannot be downloaded or exported to local devices. You need to create a custom image from a snapshot and then export the image.

Do automatic snapshots differ from or conflict with manual snapshots?

No. You can use both of them at the same time. However, you cannot create manual snapshots while automatic snapshots are being created.

You cannot create a custom disk snapshot until the automatic disk snapshot has finished being created, and vice versa.

If one snapshot created from a large disk lasts longer than the interval between two automatic snapshots, the next automatic snapshot will be skipped. For example, if you configure the automatic snapshots to be created at 9:00, 10:00, and 11:00 and the automatic snapshot created at 9:00 lasts for 70 minutes and ends at 10:10, the next automatic snapshot will be created at 11:00 and not at 10:00.

Can I create snapshots for local disks?

No. We recommend that you use data redundancy at the application layer or create deployment sets for clusters to improve the availability of applications.

After I release the cloud disk, will local snapshots be deleted?

No. You need to delete snapshots via the Console or APIs. For more information, see [Deleting Snapshots](#).

Why does the used disk capacity displayed in the file system differ from the snapshot size?

A cloud disk snapshot is a block-level clone or backup. In general, the snapshot capacity will be larger than the data capacity displayed in the file system because:

The underlying data block stores the metadata of the file system.

Some data are deleted. Deleting data modifies the written-in data block, which will be backed up to snapshots.

How can I prevent snapshots from being deleted by Tencent Cloud?

Manual snapshots: Tencent Cloud will never delete them, regardless of whether their corresponding disk or instance has been released.

Scheduled snapshots: you can set the retention period of scheduled snapshots to long-term retention to prevent them from being deleted. This way, only the snapshot created earliest will be deleted when the snapshot quota is reached.

For detailed directions, see [Scheduled Snapshot](#) and [Use Limits](#).

How can I delete snapshots to reduce backup costs?

For cloud disk snapshots, you can delete them directly via the Console or APIs. For more information, see [Deleting Snapshots](#).

For snapshots associated with custom images, you must first delete the custom images and then [delete the snapshots](#).

Will automatic snapshots be deleted after the instance expires or the cloud disk is released?

No. Automatic snapshots will not be deleted after the instance expires or the cloud disk is released because they are retained based on the retention period policies of scheduled snapshots. To modify the policies of scheduled snapshots, see [Scheduled Snapshot](#).

How can I delete snapshots from which an image or cloud disk was created?

Snapshots from which a cloud disk was created: you can delete the snapshots separately. After you delete a snapshot, you will not be able to operate a business that depends on the original snapshot data.

Snapshots from which a custom image was created: you have to first delete the image and then delete the snapshots.

Snapshots from which an instance was created: you can delete the snapshots separately. After you delete a snapshot, you will not be able to operate a business that depends on the original snapshot data.

Will the snapshot policy fail to be executed if I create a custom image or cloud disk based on the scheduled snapshot?

No.

Can a cloud disk have multiple automatic snapshot policies?

No.

How can I prevent data loss caused by incorrect operations?

You can create snapshots to back up data before you perform risky operations. For example, you can create a snapshot before you modify critical system files, migrate instances from a classic network to a VPC, back up data, restore an instance that was released accidentally, prevent network attacks, change operating systems, or provide data support for a production environment. For more information, see [Creating Snapshots](#). If an error occurs, you can [roll back snapshots](#) to reduce risks.

I created an instance in the Guangzhou region and created snapshots for the data disks of the instance. After the instance expired and was released, I purchased another instance in the Guangzhou region. Can I roll back the snapshots on the new instance to restore the original instance?

No. Snapshots can only be rolled back to the cloud disk in which they are created. You can use the snapshots of the previous data disk to create a cloud disk and mount the cloud disk to the new instance. For more information, see [Creating Cloud Disks Using Snapshots](#) and [Mounting Cloud Disks](#).

What are the differences between snapshots and images?

If no data disk is mounted to an instance and all data is written on the system disk, the data on the system disk cannot be protected by creating an image. Images cannot be scheduled for continuous backup. Once the system disk data is damaged, you can only recover the data to the state when the image was initially created. Therefore, images are not suitable for data protection. Specific differences are as follows:

Name	Snapshots	Images
Characteristics	Backup data of a cloud disk at a certain point in time	CVM software configuration template, which contains information about the operating system and pre-installed programs
Use Cases	Regularly back up important business data Back up data before major operations Produce multiple copies of data	Back up systems that will remain unchanged in the short term Deploy applications in batches Migrate the system

How can I migrate snapshot data from account A to account B?

Snapshots cannot be migrated. Instead, you can create an image from the snapshot that you want to migrate and share the image with another account.

Can I create a custom image from data disk snapshots?

No. You can only create a custom image from system disk snapshots.

Issues on Data Backup and Recovery

Last updated : 2024-01-06 17:26:45

How do I back up data for CVM?

If your CVM uses a cloud disk, you can back up your business data by creating a system disk custom image and a data disk snapshot.

For more information on how to create a custom image, see [Create Custom Images](#).

For more information on how to create a snapshot, please see [Creating Snapshots](#).

If your CVM uses a local disk, you can back up data on system disk by creating a custom image. For business data in your data disk, you still need to customize the backup policy.

You can use FTP to back up data in the server to other places. For specific FTP deployment methods, see:

For Windows: [Build the FTP Service \(Windows\)](#)

For Linux: [Build the FTP Service \(Linux\)](#)

What are common data backup and recovery solutions?

The data backup and recovery solutions vary by application scenarios and businesses. The following recommendations can be used based on your actual needs:

Back up the instance regularly using the [CBS Snapshot](#) feature.

Deploy key components of an application across multiple availability zones, and copy the data as needed.

Use [Elastic Public IP](#) for domain name mapping to ensure that the service IP can be quickly redirected to another CVM instance when the server is unavailable.

View the monitoring data regularly and configure corresponding alarms. For more information, please see [Cloud Monitor](#).

Process emergency requests with auto scaling. For more information, please see [Auto Scaling](#).

Images

Custom Images

Last updated : 2024-01-06 17:26:45

What should I do if the Windows system fails to create a custom image?

If the Windows system fails to create a custom image, you can troubleshoot the issue as follows:

1. The creation of a custom image relies on the Windows Modules Installer provided by Microsoft. Make sure this service is working properly.
2. Some antivirus tools or Safedog may block custom image creation scripts from being executed. To avoid creation failure, we recommend you disable these tools before creating a custom image.
3. If the image creation tool is interrupted by system pop-ups, remotely log in to the CVM and check and adjust the CVM configuration to block pop-ups.

Can a custom image be created from the data disk snapshot?

No. A custom image can be created from the system disk snapshot, but not from the data disk snapshot.

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, see [Creating Cloud Disks using Snapshots](#).

Do local disks support custom images?

If the system disk of an instance is a local disk, you can create custom images.

If the data disk of an instance is a local disk, you can create only images of the system disk of the instance.

How do I confirm that a data disk has been unmounted and a custom image can be created?

1. Confirm that the partition statement row of the automatically mounted data disk has been deleted in the `/etc/fstab` file.
2. Execute the `mount` command to query the mounting information of all devices. Ensure that the execution result does not include the corresponding data disk partition information.

Will the custom image still exist after the instance is released?

Yes.

Can I change the operating system of an instance created from a custom image? After I change the operating system, will I still be able to use the original custom image?

Yes. You can continue to use the original custom image after changing the operating system.

Can I upgrade the CPU, memory, bandwidth, disk, and other configurations of a CVM instance enabled by a custom image?

Yes. You can upgrade all of them. For more information, see [Change Instance Configuration](#) and [Adjust Network Configuration](#).

Can a custom image be used across regions?

No. A custom image can only be used in the same region. For example, you cannot directly launch a CVM instance in East China (Nanjing) using the custom image created from a CVM instance in East China (Shanghai).

To use a custom image across regions, copy the image to the target region first. For more information, see [Copying Images](#).

Where can I view the image creation progress? How long does it take to create an image?

You can view the image creation progress on the **Images** page of the CVM Console. The time it takes to create an image depends on the size of the instance's data.

Copying Images

Last updated : 2024-01-06 17:26:45

Under what circumstances do I need to copy an image?

Custom images can only be used in the same region. You can perform the following operations by [copying images](#):

Deploy applications on the CVM instance in multiple regions.

Migrate the CVM instance to other regions.

Use custom images across regions.

Which images can I copy?

You can only copy custom images and cannot copy public images, service marketplace images, and images that others have shared with you.

Which regions currently support copying images?

All Tencent Cloud regions support copying images.

How long does it take to copy an image?

It depends on the network transmission speed and the task waiting queue length. To copy an image, you need to transfer the image files from a region's availability zone to the target region's availability zone via the network. Please wait patiently.

How much does it cost to copy an image ?

Copying images across regions is free for now, but the copied image will occupy the snapshot capacity and incur a cost. For more information, see [Pricing List](#).

How do I copy the image resources under my account to other regions under another Tencent Cloud account?

You need to first copy your image to the target region, and then share the image under the target region with another Tencent Cloud account. After the image is shared, it will appear in the target account's shared image list.

Is there a capacity limit for copying images?

No.

Issues on Custom Image

Last updated : 2024-05-16 10:19:57

What is the maximum number of users an image can be shared with?

500 users.

Can the name and description of a shared image be changed?

No.

Does a shared image count toward my image quota?

No.

Are there any region limitations for a shared image on creating and reinstalling a CVM instance?

Yes. The shared image should be in the same region as the source image. The CVM instance can only be created and reinstalled in the same region.

Can a shared image be copied to other regions?

No.

Can a custom image that has been shared with other users be deleted?

Yes, but you need to cancel all sharing for this custom image first.

Can an image shared by other users be deleted?

No.

What are the risks of using a custom image shared by other users?

Tencent Cloud does not guarantee the integrity and security of images shared by other users. Please select images shared by a trusted account.

Can I share with others an image another user shared with me?

No.

What are the risks of sharing a custom image with another account?

Data and software leakage may occur. Before sharing a custom image with another account, make sure the image has no sensitive data or software. Be aware that the account you share the image with can create CVM instances from the shared image to create more custom images. The constant transfer of data will increase the risk of leakage.

Can I create an instance using an image that I shared with other users?

Yes. You can create a CVM instance using the image you shared with another account. You can also create a custom image based on the CVM instance.

Can an image created on Server A in North China be shared with Server B in East China?

If both servers are under the same account, you can directly copy the image to Server B in East China. For detailed directions, see [Copying Images](#).

If both servers are under different accounts, you need to copy the image to the East China region and then share it with the account of Server B. For detailed directions, see [Copying Images](#), [Sharing Custom Images](#), and [Cancelling Image Sharing](#).

Changing Images (Reinstalling System)

Last updated : 2024-01-06 17:26:45

Does CVM support reinstalling the operating system?

Reinstalling the operation system is an important method to restore instances to their initial status in the event of a system failure. For more information, see [Reinstalling System](#).

How long does it take to reinstall the operating system for an instance?

Reinstallation generally takes 10-30 minutes.

What should I do if re-installation is very slow or fails?

Reinstallation generally takes 10-30 minutes.

If it has been less than 30 minutes, please continue to wait.

If reinstallation does not complete within 30 minutes or fails, contact us by [submitting a ticket](#).

Will reinstalling the operating system cause data loss?

After reinstallation, all data on the server's system disk will be cleared and the system disk will be restored to its initial status. All data on the server's data disk will be retained but cannot be used until the data disk is mounted manually.

How can I replace the operating system of a CVM with an existing image?

See [Reinstalling System](#).

Cloud-Init & Cloudbase-Init

Last updated : 2024-03-26 10:28:49

Cloud-init

What is cloud-init?

Cloud-init is an open source tool that runs inside a CVM instance as a non-resident service. It is executed at startup and exits immediately after execution. It does not listen to any ports.

All the Linux public images of Tencent Cloud are pre-installed with the cloud-init service. You need to run the service as the root user because the service is mainly used for the initialization of CVM instances such as configuring DNS, hostname, and IP, and the execution of some custom scripts that users specify to be executed during the first boot when creating the CVM instances.

How do I check whether the cloud-init service inside a Linux instance is working properly?

Troubleshooting cloud-init service errors

[Log in to a Linux instance using the standard login method \(recommended\)](#). Then, run the following commands in sequence and observe whether an error is reported. If the execution result is displayed, the service is running normally. Otherwise, the cause of the error will be reported. You can troubleshoot the error accordingly.

Note:

This step only applies to CVM instances created with public Linux images. If you installed cloud-init by yourself, adjust the commands according to the actual situation.

1. Delete the cloud-init cache directory.

```
rm -rf /var/lib/cloud
```

2. Perform complete cloud-init initialization.

```
/usr/bin/cloud-init init --local
```

3. Pull data from the configured data source.

```
/usr/bin/cloud-init init
```

4. Cloud-init initialization involves multiple stages. To ensure sufficient dependencies between the stages, specify the config stage for the cloud-init modules.

```
/usr/bin/cloud-init modules --mode=config
```

5. Specify the final stage for the cloud-init modules.

```
/usr/bin/cloud-init modules --mode=final
```

What initialization operations does cloud-init perform on instances?

Tencent Cloud implements all instance initialization operations through cloud-init, ensuring the transparency of the operations inside an instance. The following briefly covers some initialization operations. For more details, see [cloud-init documentation](#).

Initialization Type	Default Behavior	Customization	Notes
hostname initialization	During the first launch of an instance, cloud-init sets the hostname of the instance according to the hostname information in <code>vendor_data.json</code> .	If you create or reinstall an instance with a custom image and you want to keep the custom `hostname` settings of the image, you can set <code>preserve_hostname</code> in <code>/etc/cloud/cloud.cfg</code> to true and delete <code>- scripts-user</code> before creating the custom image.	If <code>preserve_hostname</code> is disabled, the <code>scripts-user</code> inside the instance will not be executed. This configuration will also affect the initialization of sub-items such as the installation of the Observability Platform and cloud software source settings. Also, the scripts will not be executed when you create the instance.
/etc/hosts initialization	During the first launch of an instance, cloud-init initializes <code>/etc/hosts</code> to <code>127.0.0.1 \$hostname</code> by default.	If you create or reinstall an instance with a custom image and you want to keep the custom <code>/etc/hosts</code> settings of the image, of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- scripts-user</code> and <code>- ['update_etc_hosts', 'once-per-instance']</code> .	After you disable <code>- scripts-user</code> initialization script, the <code>/var/lib/cloud/instance</code> inside the instance will not be executed. This configuration will also affect the initialization of sub-items such as the installation of the Observability Platform and cloud software source settings. Also, the scripts will not be executed when you create the instance. During instance restart, the settings of some existing instance will be restored. For information about the solution, see etc/hosts Configuration of Linux I .
DNS initialization (non-DHCP scenario)	During the first launch of an instance, cloud-init sets the DNS of the instance according to the nameservers information in <code>vendor_data.json</code> .	If you create or reinstall an instance with a custom image and you want to keep the custom DNS settings of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- resolv_conf</code> .	None.

		and <code>unverified_modules:</code> <code>['resolv_conf']</code> .	
Software source initialization	During the first launch of an instance, cloud-init sets the software source of the instance according to the <code>write_files</code> information in vendor_data.json .	If you create or reinstall an instance with a custom image and you want to keep the custom software source settings of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- write-files</code> .	None.
NTP initialization	During the first launch of an instance, cloud-init sets the NTP server configuration of the instance according to the NTP Server information in <code>vendor_data.json</code> and starts the NTP service.	If you create or reinstall an instance with a custom image and you want to keep the custom NTP settings of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- ntp</code> .	None.
Password initialization	During <code>the first launch</code> of an instance, cloud-init sets the default account password of the instance according to the <code>chpasswd</code> information in <code>vendor_data.json</code> .	If you create or reinstall an instance with a custom image and you want to keep the custom default account password of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- set-passwords</code> .	None.
Key binding	During the first launch of an instance, cloud-init sets the default account key of the instance according to the <code>`ssh_authorized_keys`</code> information in <code>vendor_data.json</code> .	If you create or reinstall an instance with a custom image and you want to keep the custom key of the image, you can delete the following configuration from <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>- users-groups</code> .	If you manually bind the instance instance, the previous key will be perform the key binding operatio

Network initialization (non-DHCP scenario)	During the first launch of an instance, cloud-init sets the IP, Gateway, and Mask of the instance according to the information in <code>network_data.json</code> .	If you create or reinstall an instance with a custom image and you want to keep the custom network information of the image, you can add the following configuration to <code>/etc/cloud/cloud.cfg</code> before creating the custom image: <code>network: {config: disabled}</code> .	None.
--	---	--	-------

How can I fix common issues related to Cloud-Init?

1. Error due to the unmounting of Cloud-Init dependencies

Problem description:

When commands are used to check whether the Cloud-Init service is working properly, the following error is returned:

```
Traceback (most recent call last):
File "/usr/bin/cloud-init", line 5, in
*****
raise DistributionNotFound(req)
pkg_resources.DistributionNotFound: pyyaml
```

Problem analysis:

"pkg_resources.DistributionNotFound: xxxxx" indicates that the cloud-init dependencies have been uninstalled.

Solution:

1.1 Reinstall the dependencies.

1.2 Follow [Troubleshooting cloud-init service errors](#) to see if the error is returned again.

2. Error due to the modification of the default Python interpreter

Problem description:

An error is returned when cloud-init is run on startup.

Problem analysis:

When cloud-init is installed, Python 2 is used as the default Python interpreter, which means that symbolic links,

`/usr/bin/python` and `/bin/python`, are linked to Python 2. Users may change the default Python

interpreter to Python 3 inside the instance by directing the symbolic links, `/usr/bin/python` and

`/bin/python`, to Python 3. Due to compatibility issues, an error will be returned when cloud-init is run on startup.

Solution:

1.1 Taking Python 2.7 for example, modify the Python interpreter specified in the `/usr/bin/cloud-init` file by changing `#!/usr/bin/python` or `#!/bin/python` to `#!/usr/bin/python2.7`.

Note:

Do not use symbolic links. Point directly to a specific interpreter.

1.2 Follow [Troubleshooting cloud-init service errors](#) to see if the error is returned again.

Cloudbase-Init

What is Cloudbase-Init?

Like cloud-init, Cloudbase-Init is a bridge by which you can communicate with Windows CVM instances. The Cloudbase-Init service is run when an instance boots up for the first time. The service will read the configuration information of the instance and initialize it. Following operations such as resetting password and modifying IP addresses are also done via Cloudbase-Init.

How do I check whether the Cloudbase-Init service inside a Windows instance is working properly?

Troubleshooting Cloudbase-Init service errors

1. Log in to the instance.

Note:

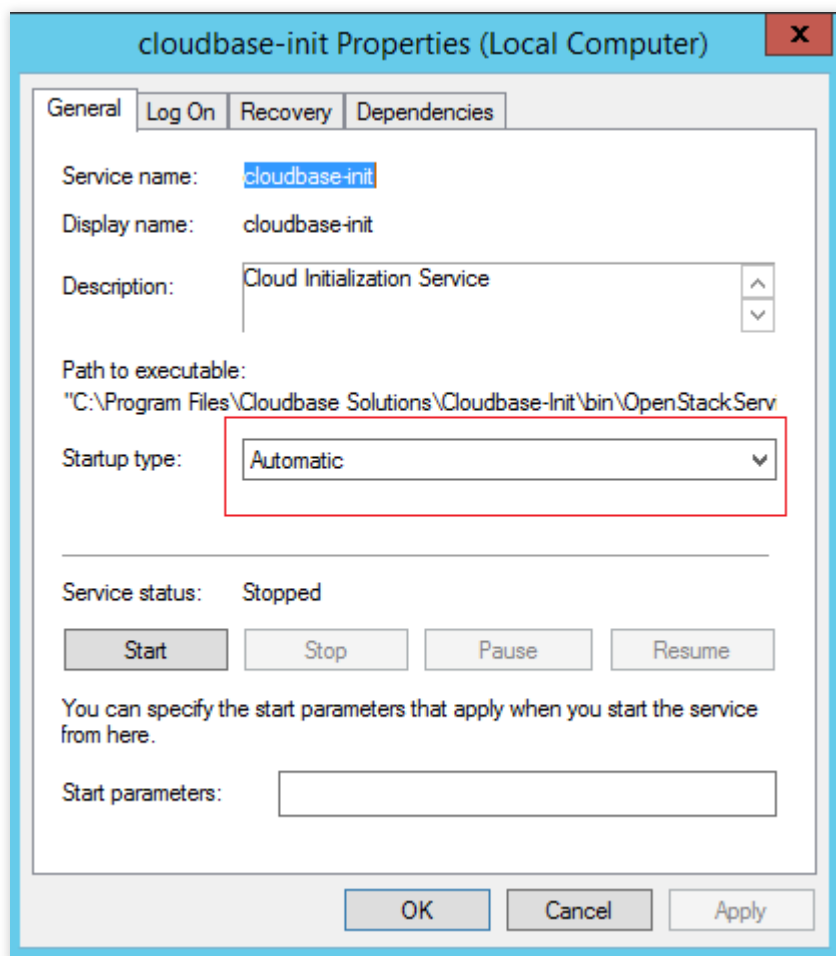
If you forget your password or fail to reset your password because of Cloudbase-Init service exceptions, you can reset your password by following [Step 2](#).

2.

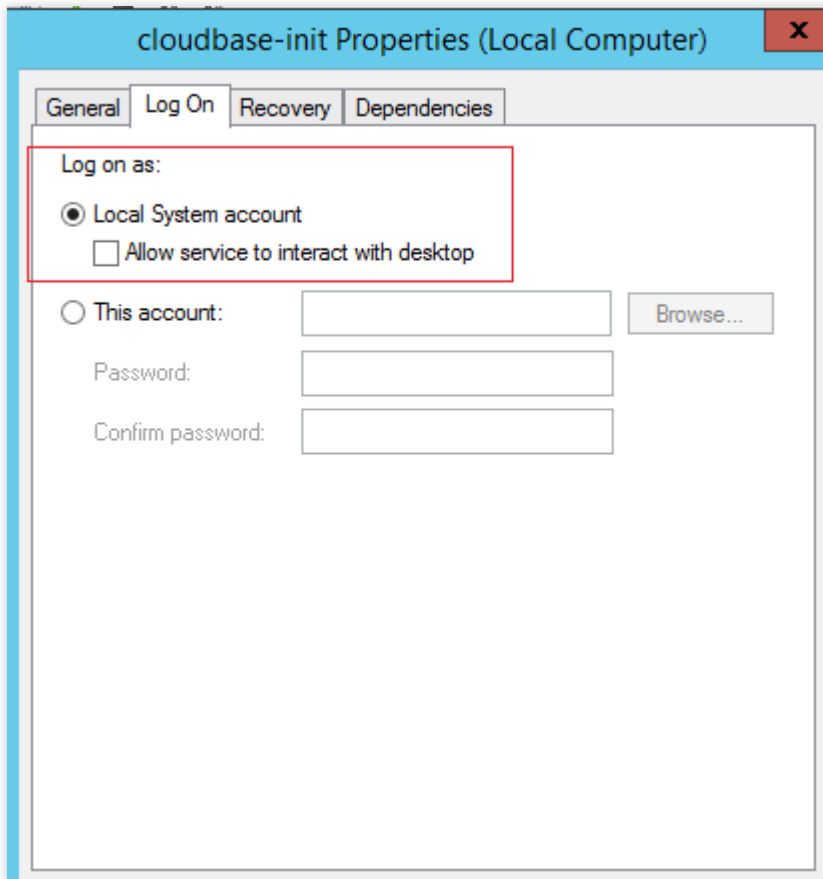
Choose **Control panel > Administrative tools > Services**.

3. Find the Cloudbase-Init service, right-click it, and go to **Properties**.

View **Startup type** and make sure it is set to **Automatic**, as shown in the figure below:

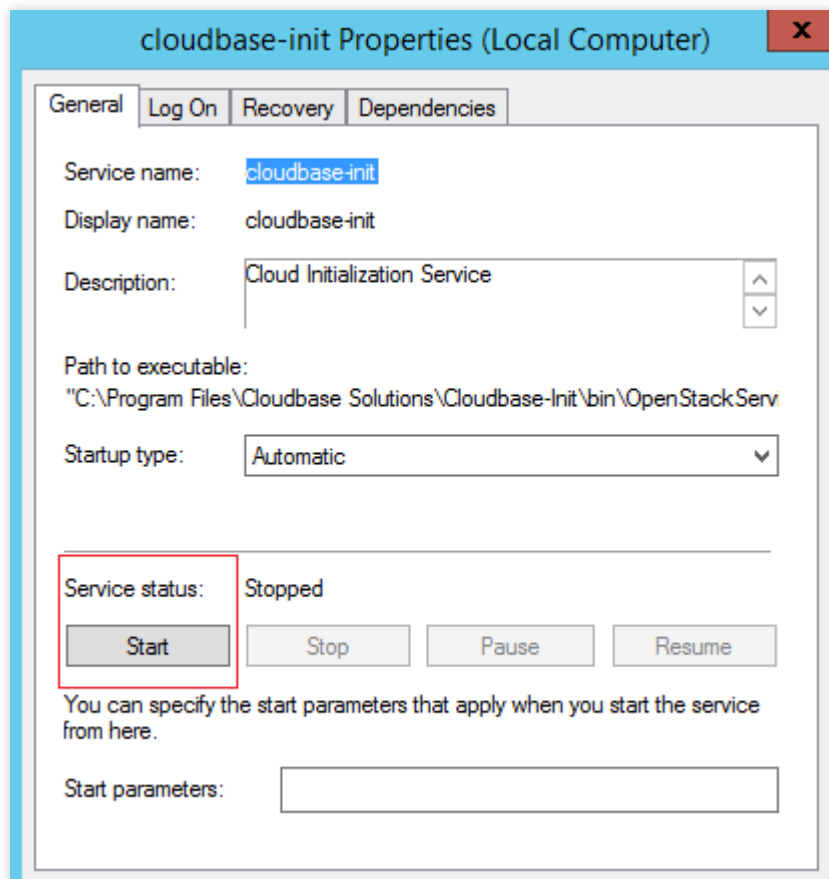


View **Log on as** and make sure that **Local System account** is selected, as shown in the figure below:

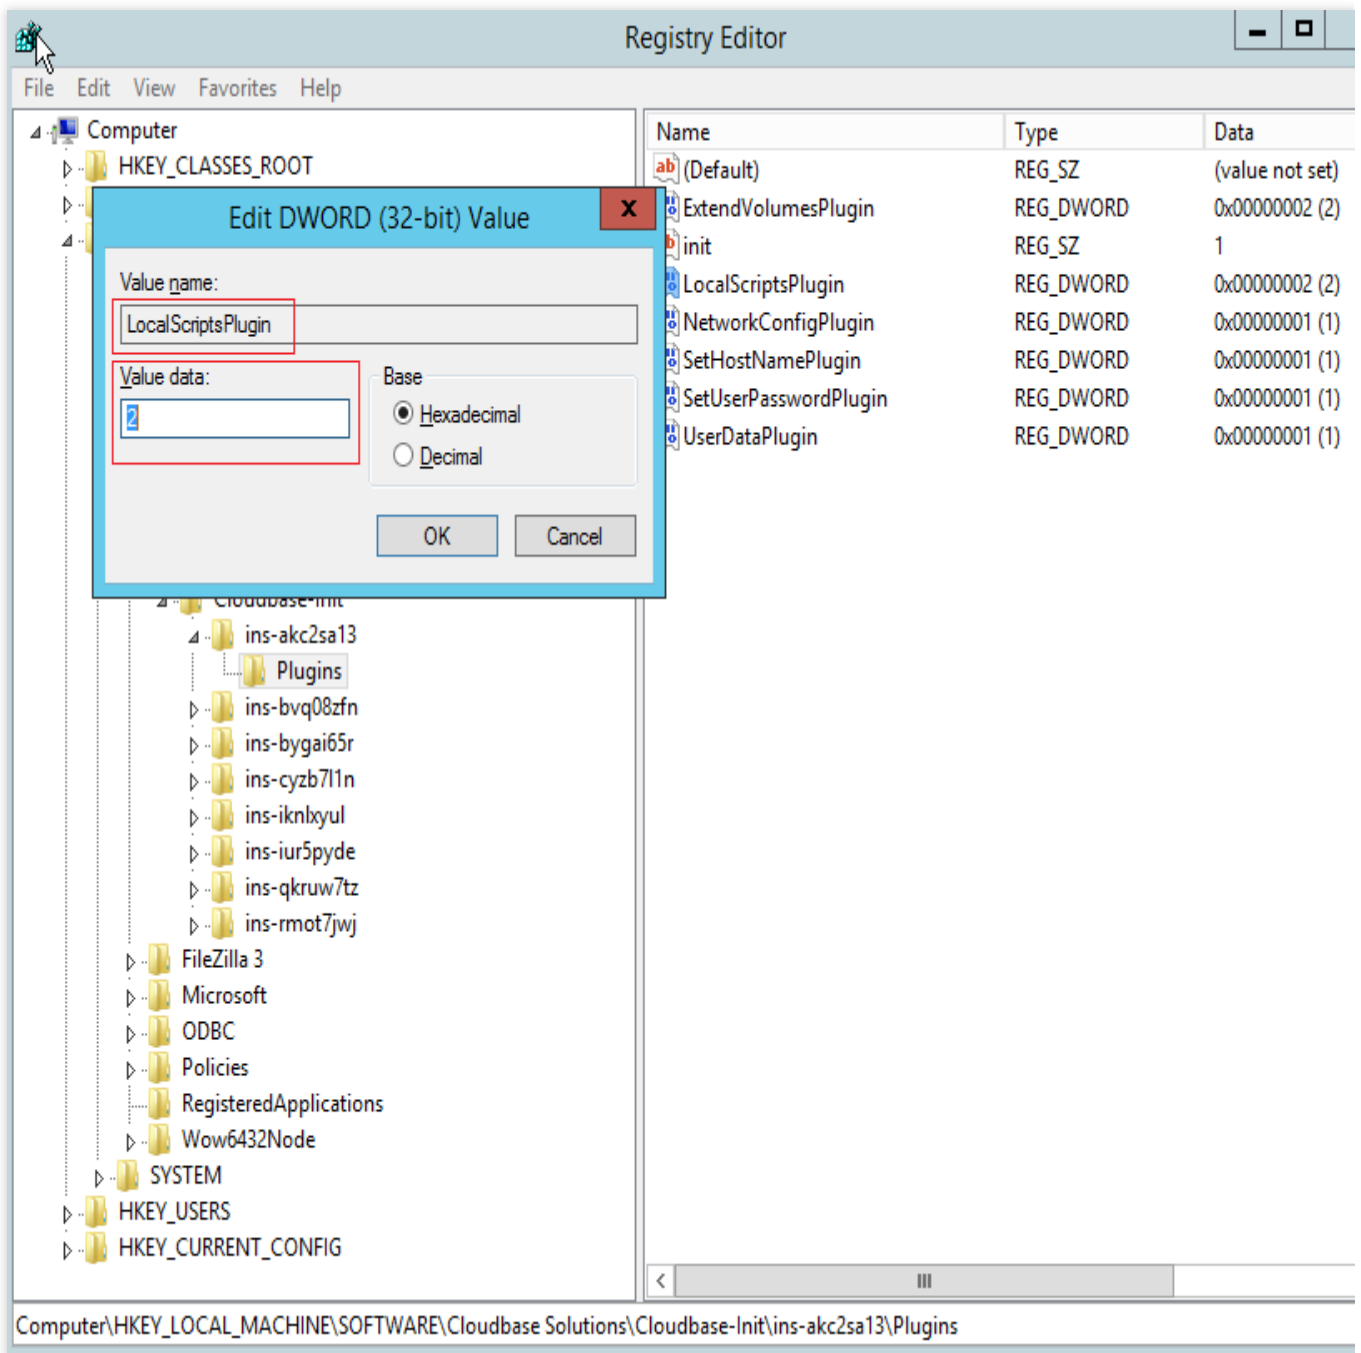


Manually start the Cloudbase-Init service and see if any error is returned.

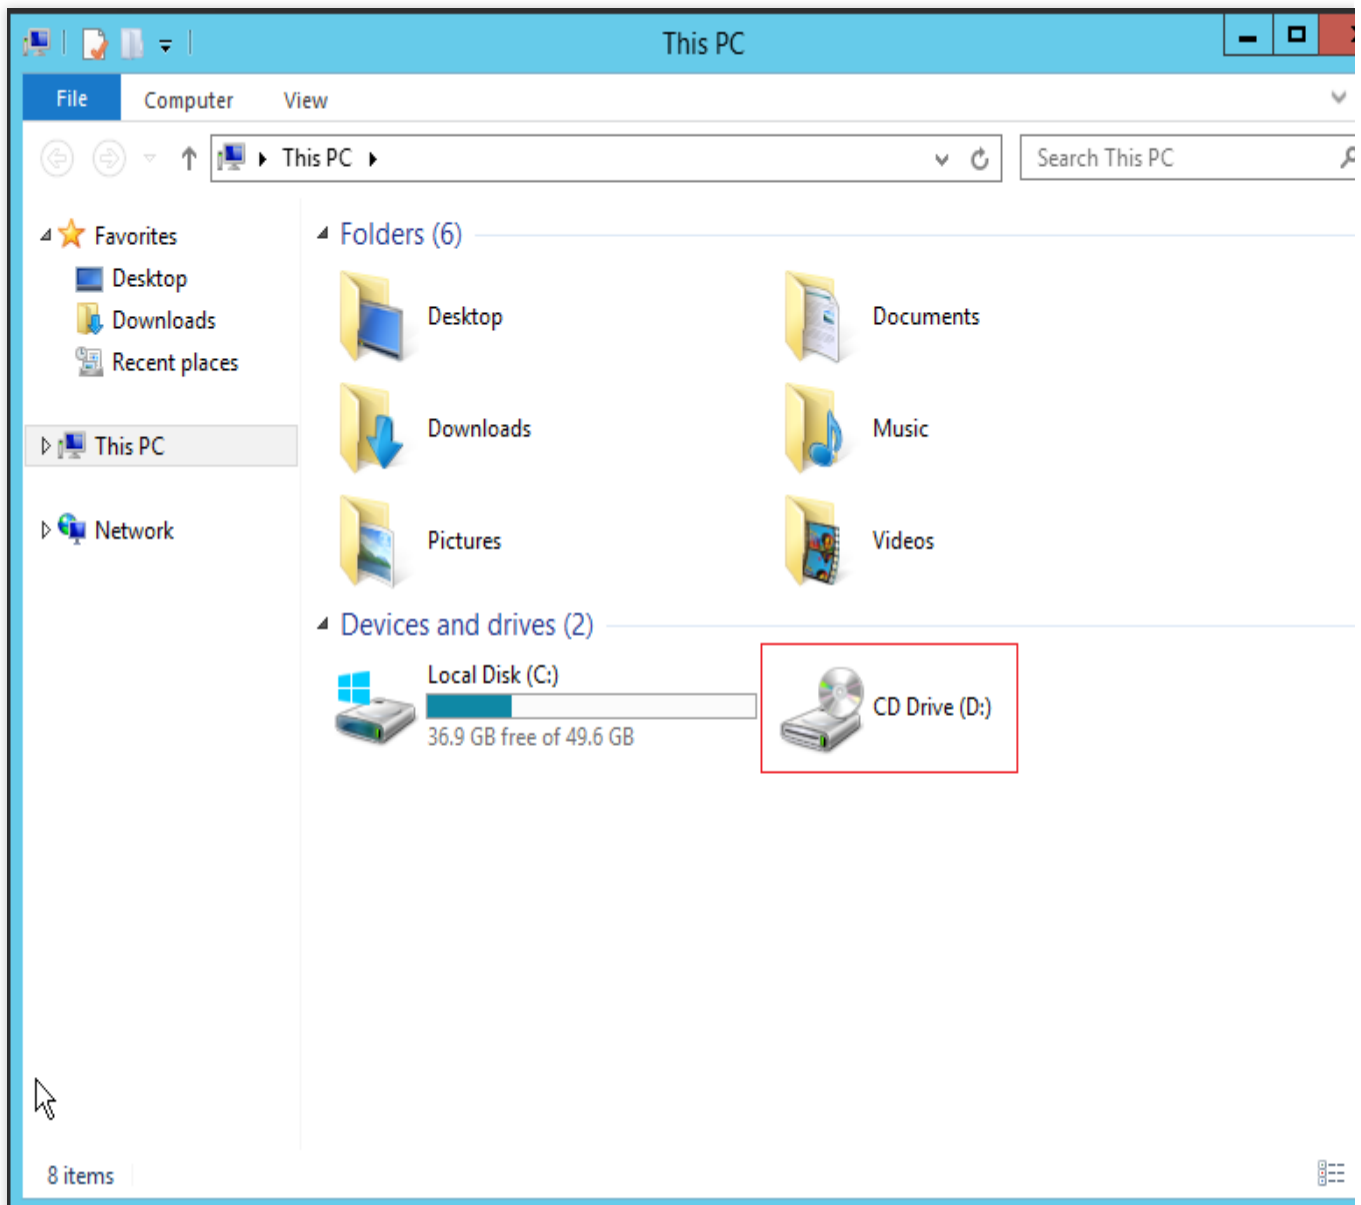
If any error is returned, you need to fix the issue first and check whether you have installed any security software which may stop Cloudbase-Init from performing related operations.



Open the registry, find each **LocalScriptsPlugin**, and make sure its value is 2, as shown in the figure below:



Check whether CD-ROM loading is disabled. If there is an optical disc drive as shown in the figure below, the loading has not been disabled. Otherwise, you need to enable the loading.



How do I view Cloudbase-Init execution logs?

You can view the log files for different operating systems:

Linux: `/var/log/cloud-init-output.log`

Windows: `C:\Program Files\Cloudbase Solutions\Cloudbase-Init\log\cloudbase-init.log`

How do I fix common issues related to Cloudbase-Init?

Failed to reset password during initialization

Possible reasons:

The Cloudbase-Init account password is manually changed, which results in the failure to start the Cloudbase-Init service, which further led to the failure of operations such as resetting password during initialization.

The Cloudbase-Init service is disabled, which led to the failure of operations such as resetting password during initialization.

The security software installed stops the Cloudbase-Init service from resetting password so that the operation returns a successful result but actually failed.

Solution:

Follow the corresponding solution to each possible reason to fix the issue.

1.1 Change the Cloudbase-Init service to LocalSystem service. For details, see [step 2](#) in [Checking the operation of the Cloudbase-Init service](#).

1.2 Change the startup type of the Cloudbase-Init service to automatic. For details, see [step 2](#) in [Checking the operation of the Cloudbase-Init service](#).

1.3 Uninstall the security software involved or add the relevant operations of the Cloudbase-Init service to the white list of the security software.

FAQs about Red Hat Enterprise Linux Image

Last updated : 2024-03-28 10:28:31

Which Red Hat Enterprise Linux image tags are available? Which models support Which Red Hat Enterprise Linux?

Note:

To use Red Hat Enterprise Linux images, select an instance type that is verified by Red Hat Enterprise Linux when you purchase the CVM.

Instance specifications	RHEL 7 Series	RHEL 8 Series	RHEL 9 Series
Standard SA5	Not Supported	Supported	Supported
Standard SA3, S6, SA2, S5	Supported	Supported	Supported
MEM optimized M6, M5	Supported	Supported	Supported
Computing C4, C3	Supported	Supported	Supported

Why can't I see the Red Hat Enterprise Linux image when purchasing certain types of instances?

Please check if your selected instance model is the one listed in the table above.

What should I do if my existing image tag or model doesn't meet the requirements?

You can submit a ticket to [contact us](#).

How is the Red Hat Enterprise Linux image billed?

For billing details, see [Billing Description](#).

After-sales support

Free source updates (including version updates and vulnerability patches) are available to the officially licensed Red Hat Enterprise Linux OS sold on Tencent Cloud official website. If you have any problems during use, please [submit a ticket](#).

Are refunds supported after I purchased the Red Hat Enterprise Linux image?

The Red Hat Enterprise Linux images are licensed through official authorization from Red Hat. Once a license has been purchased and activated on Tencent Cloud, it is not eligible for a refund.

For instances with a **pay-as-you-go** billing mode, those using Tencent Cloud authorized Red Hat Enterprise Linux images can be terminated or returned.

For instances with a **monthly subscription** billing mode, those using Tencent Cloud authorized Red Hat Enterprise Linux images are not eligible for termination or refund. If you wish to receive a refund of the instance fee, you can reinstall the system with another open-source Linux operating system, such as CentOS, and then proceed with the regular refund process for the instance.

Note :

After the system is reinstalled, the fee of the image will not be refunded.

How is the Red Hat Enterprise Linux image billed under the RI mode?

For details, see [Billing Description - Red Hat Enterprise Linux Image](#).

Does the same discount apply to both the CVM instance and the Red Hat Enterprise Linux image?

The discount available to the CVM instance is not applied to the Red Hat Enterprise Linux image. No other discounts are available.

Can I use a voucher to purchase Red Hat Enterprise Linux images?

No. Vouchers cannot be used to purchase Hat Enterprise Linux images.

Does Tencent Cloud support BYOL for Red Hat Enterprise Linux images?

Yes. You can import a custom image to create an instance. No commercial image license fee occurs. For more information, see [Importing Images](#).

Can I change the configuration and the billing mode for an instance with Red Hat Enterprise Linux OS?

No.

Others

Last updated : 2024-01-06 17:26:45

What is an image?

An image is the template for CVM software configuration (operating systems, pre-installed programs, etc.). Tencent Cloud requires users to use images to launch instances. An image can launch multiple instances, and users can use it repeatedly. For more information on images, please see [Overview](#).

What do I need to do before importing the image?

Before importing an image, you need to complete two major steps: applying for permissions and preparing image files. For more information, please see [Overview](#).

How do I export an image for local testing?

Tencent Cloud Service Migration supports images in qcow2, vhd, raw, and vmdk formats.

You can use the image export tools of virtualization platforms such as VMWare vCenter Convert or Citrix XenConvert. For more information, please reference each platform's export tool documentation. You can also export images [using Disk2vhd \(Windows\)](#) or [by running commands \(Linux\)](#).

Can I delete a custom image that has been used to create a CVM instance?

Yes. After a custom image is deleted, it can no longer be used to launch a new CVM instance. It will not affect instances that have already been launched. If you want to delete all instances launched from this image, see [Reclaiming Instances](#) or [Terminating Instances](#).

Can I delete a custom image that has been shared with another account?

No. You cannot delete a custom image that has been shared with another account. To delete it, you need to cancel image sharing first. For more information, see [Cancelling Image Sharing](#).

Service Migration

Last updated : 2024-03-26 10:29:23

Offline Migration

Why do COS upload and migration take so long?

The time it takes to upload images to COS is related to the image file size and bandwidth. We recommend you use compressed image formats (QCOW2 or VHD) to shorten the transmission and migration time.

Why did the migration task fail?

Tencent Cloud's service migration supports images in QCOW2, VHD, VMDK and RAW formats. Make sure that your image is in one of these formats.

Make sure that your image file has been uploaded to COS, check COS permissions, and check that the temporary URL is valid and that the file is not corrupted.

Make sure that the target CVM or cloud disk are not expired.

How do I troubleshoot the error occurred in a migration task?

Failed to obtain the metadata of the image file: The image file is damaged or the image file format is not supported. Supported image file format: QCOW2, VHD, VMDK, and RAW.

Image decompression failed: Image creation error. Check the image package and export the image again. Make sure that the destination disk has a larger capacity than the source disk.

The storage space of the target device is too small: The capacity of the system disk or data disk you want to migrate to is smaller than the capacity of the source disk or the size of the image file.

Failed to access the COS image file: Not authorized to access COS. Make sure the COS file is in the current region, COS permission is assigned and the temporary URL is valid.

Failed to migrate to the destination disk:

The destination disk does not have enough capacity.

An error occurred during the image creation.

Task timeout, System error, and Other reasons: Try again. If the problem persists, [contact us](#).

What should I do if the network is disconnected after the offline migration of the Windows server?

You can reset the Windows network configuration as instructed below. For more information, see [Ping Failures](#).

1. For VPCs created before June 2018, check if the server supports DHCP. If not, check the static IP.
2. If the server supports dhcp, check if the assigned private IP is correct. If it is incorrect, run the command as the admin without restarting the server: `ipconfig /release; ipconfig /renew`.

3. If the assigned IP is correct, but the network is disconnected, run `ncpa.cpl` to open the local connection, and disable and enable the ENI.

4. Run the command as the admin and restart the server: `reg delete`

```
"HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Windows  
NT\\CurrentVersion\\NetworkList\\Profiles" /f .
```

Online Migration

Which types of operating systems and disks are supported by online migration?

Mainstream Linux and Windows operating systems are supported.

Migration is irrelevant to disk type and usage.

What is the difference between online migration and image import?

Both can migrate source systems to Tencent Cloud. However, online migration allows you to transfer the system and data from the source server to Tencent Cloud without manually preparing, exporting and importing an image.

Will the public IP of the source server be migrated?

No. The online migration only syncs the source system and data, rather than the public IP.

Does the online migration tool support checkpoint restart?

Yes. You can run the tool again to resume the transfer after an interruption.

Do I need to retain the tool after migration is completed?

No. Once migration is completed, you can directly delete the tool on the source server.

What about migration speed and cost?

Speed: The speed is subject to the destination CVM instance bandwidth. In the test where a 1CIG pay-as-you-go instance with a bandwidth of 12 Mbps is used, the actual migration speed is about 9 Mbps. For the specific calculation method, see [Migration Time Estimation](#).

Cost: The migration tool is free of charge. You will be charged for bandwidth and other resources used during migration. They will be billed by traffic and duration.

Can I migrate multiple CVMs at the same time?

Yes. You can migrate multiple CVMs to different destination CVMs simultaneously.

How do I check and install Virtio?

Check and install Virtio by referring to [Checking Virtio Drivers in Linux](#).

How do I install Rsync?

Select one of the following commands based on the operating system of the source server to install Rsync:

CentOS: `yum -y install rsync`

Ubuntu : `apt-get -y install rsync`

SUSE: `zypper install rsync`

For the specific commands of other Linux distributions, see the installation document at the corresponding official website.

How do I disable SELinux?

Edit the `/etc/selinux/config` file and set `SELINUX=disabled` .

FAQs about Console-Based Online Migration

Where can I download the online migration tool?

[Click here](#) to download the online migration tool package and use it as instructed in [Migration in Console](#).

How do I import a migration source?

1. Download and decompress [online migration tool package](#).
2. Run the executable file for the source server architecture to import the migration source to the online migration page in the CVM console.

If the source server is 32-bit, run `go2tencentcloud_x32` .

If the source server is 64-bit, run `go2tencentcloud_x64` .

How do I update a migration source or import a migration source again?

If information such as migration source disk changes, you need to update the migration source information or import the source again. Delete the existing migration source first and then run the client to import the source again.

How do I delete a migration source?

Log in to the [CVM console](#), go to the online migration page, and click **Delete** on the right of the target migration source. If your migration source is associated with a migration task that has not been completed, pause and delete the task first and then delete the source.

How do I select the destination CVM instance before migration?

We recommend you use the same operating system for the source server and destination CVM instance. For example, to migrate a CentOS 7 source server, select a CentOS 7 CVM as the destination.

Which migration task status indicates that the task has been completed?

A migration task has been completed only when it is in the **Successful** status; otherwise, it has not been completed.

How do I delete a migration task?

Perform the following operations based on the actual migration task status:

When a migration task is in the **Not started** or **Successful** status, you can click **Delete** on the right to directly delete it.

When a migration task is in the **Failed** status, to delete it, you need to keep the migration source online and click **Delete** on the right. Then the task will enter the **Deleting** status, and the migration tool will automatically clear the resources generated by the migration task and delete the task.

How do I cancel a migration task if it takes too long?

In this case, log in to the [CVM console](#), go to the online migration page, pause the migration task, and click **Delete** on the right to delete it.

What should I do if the relay instance is terminated?

During migration to the destination Tencent Cloud image, if the created relay instance is terminated by mistake, you can delete the current migration task and create and start a migration task for the migration source again.

What should I do if an error or failure occurred during migration in the console?

Check the common errors below:

During migration to the destination CVM instance, ports 80 and 443 are not open in the security group of the destination CVM instance.

Solution: modify the security group rules of the destination CVM instance to open ports 80 and 443.

During migration to the destination CVM instance, the disk capacity of the destination CVM instance is insufficient.

Solution: Select a CVM instance with a sufficient capacity, mount a disk with a sufficient capacity to it, and create a migration task again to start migration.

An error occurred when you create a relay instance during migration to the destination Tencent Cloud image; for example, a log error message **Failed create transit instance, maybe no available source in target region** is displayed.

Solution: there are no available resources in the destination region of the migration task. In this case, you can select another region or migrate the server to the specified CVM instance.

If the problem persists or the reason is not identified, [contact us](#) and provide the migration log file (in the `log` directory of the migration tool by default).

What can I get after migration?

After a migration task is completed, a Tencent Cloud resource is generated according to your specified destination type.

Migrate to a **CVM image**: A destination Tencent Cloud image will be generated for the migration source after migration. You can click the **CVM image ID** in the row of the migration task to view the image details on the [CVM image page](#), and use the image to quickly create CVM instances. After a CVM image is created successfully, a **CVM snapshot** associated with it will be created at the same time.

Migrate to a **CVM instance**: The migration source will be migrated to the destination CVM instance after migration.

How do I check the system after migrating a Linux server?

Check whether the target CVM instance starts up normally, whether data on the target CVM instance is consistent with that on the source server, and whether the network and other system services are normal.

How do I migrate a source again after migration is completed?

Log in to the [CVM console](#), and create and start a migration task for the migration source again.

What should I do if the network is disconnected after the online migration of the Windows server?

You can reset the Windows network configuration as instructed below. For more information, see [Ping Failures](#).

1. Check if the server supports dhcp (all VPCs created after June 2018 support dhcp). If not, check if the static IP is correct.
2. If the server supports dhcp, check if the assigned private IP is correct. If it is incorrect, run the command as the admin without restarting the server: `ipconfig /release; ipconfig /renew`.
3. If the assigned IP is correct, but the network is disconnected, run `ncpa.cpl` to open the local connection, and disable and enable the ENI.

4. Run the command as the admin and restart the server: `reg delete`

```
"HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Windows  
NT\\CurrentVersion\\NetworkList\\Profiles" /f .
```

Online Migration Tool

Where can I download the online migration tool?

[Click here](#) to download the compressed migration tool package and use it as instructed in [Online Migration Tool](#).

How do I use the online migration tool?

You need to download the tool and upload it to the source server and modify the configuration file based on the actual server conditions before running it. To perform batch migration, you can write a script to implement batch processing.

What should I do if an error or failure occurred during migration with the tool?

Check the common errors below:

Ports 80 and 443 are not open in the security group of the destination CVM instance.

The `DataDisks` field in the `user.json` file is not configured for migrating the data disk of the source server.

The system disk capacity of the destination CVM instance is insufficient to store all the data.

In the [private network migration mode](#), the migration stage 3 is not implemented, which causes the destination CVM to exit the migration mode.

If the problem persists or the reason is not identified, [contact us](#) and provide the migration log file (in the `log` directory of the migration tool by default).

Network

IP Address

Last updated : 2024-01-06 17:26:45

Public IP

How can a CVM without public IP access public network?

If you did not purchase the public IP when purchasing the CVM or have returned the public IP, you can apply for an EIP on the [EIP Console](#) and bind it to your CVM to allow public network access.

Can I change my public IP?

You can change the public IP of your CVM. For more information on specific operations, please see [Changing Instance Public IP](#).

How do I keep a public IP unchanged?

If you need to retain a specific public IP under your account, you can convert it to an EIP, which can be bound to the device and used to access the public network. This EIP will be retained under your account until you **release** it. For more information on the related operations, see [EIP](#).

What is a public IP address?

A public IP address is an unreserved IP address on the Internet. A CVM with a public IP address can be accessible to and by other computers on the Internet. For more information, see [Internet Access](#).

How do I obtain the public IP address of an instance?

For details, see [Getting Public IP Addresses](#).

How do I change the public IP address of an instance?

For details, see [Changing Instance Public IP](#).

What are the differences between public gateways and CVMs with public IP addresses?

Public gateways enable the public network traffic forwarding feature in the image. However, CVMs with public IP addresses do not have this feature by default. Therefore, CVMs created on Windows public image cannot function as public gateways because the Windows image has no traffic forwarding feature.

Why can't I change the public IP of my CVM?

Possible reasons include:

The CVM instance is shut down and incurs no charges when shut down.

The public IP of the CVM has been changed.

How can I obtain the public IP address of an instance that has no public IP (IPv4) assigned during the instance creation?

You can obtain the public IP of an instance by applying and binding an EIP to it. For the step-by-step operations, see [EIP](#).

Private IP

What is a private IP address?

A private IP address cannot be accessed through the Internet. That is how Tencent Cloud provides private network services.

For more information, please see [Private Network Access](#).

How do I obtain the private IP address of an instance?

For details, see [Getting Private IP Addresses and Setting DNS](#).

Can I change the private IP address of an instance in addition to its public IP address?

Yes, you can change the private IP of an instance. For operations, see [Modifying Private IP Addresses](#).

Elastic Public IP issues

Last updated : 2024-01-06 17:26:45

What are EIPs used for?

EIPs are applicable to the following scenarios:

Disaster recovery. We strongly recommend you use EIPs for disaster recovery. For example, when one of your servers fails, you can unbind the EIP from this server and then bind it to a healthy server to quickly resume services.

Retaining a specific public IP. If you need to retain a specific public IP under your account, you can convert it to an EIP, which can be bound with or unbound from the device and used to access the public network. This EIP will be retained under your account until you "release" it.

Other special cases. When you need to change an IP in some cases, you can convert the public IP to an EIP and then bind the EIP to or unbind from your device. But the EIP resource under an account is limited in each region, please plan and use EIPs accordingly.

How is an EIP billed?

1. The fee displayed on the console applies to EIPs that have been idle for more than 1 hour. EIPs can be billed accurately to seconds. EIPs that have been bound/unbound multiple times are billed based on the total duration (in sec) for which they are unbound.
2. EIPs that have been idle for less than 1 hour are billed for resource occupation fee on a pro rata basis.

When is an EIP billed?

You can apply for, bind, unbind or release EIPs. With limited EIP resources available, an EIP is billed for a small resource occupation fee only when it is unbound.

How do I stop EIP billing?

When you no longer need an EIP, you can release it to stop the billing.

For more information on the specific operations, see [Releasing EIPs](#).

If you need to retain an EIP but want to stop the billing, bind it to a device (such as CVM or NAT). A bound EIP will not be billed.

Can an EIP be converted back to a public IP?

An EIP cannot be converted back to a public IP.

Can an EIP be retrieved?

You can retrieve public IPs that have not been assigned to other users. For details, see [Retrieving the Public IP Address](#).

Elastic Network Interface

Last updated : 2024-01-06 17:26:45

What is ENI?

[Elastic Network Interface](#) (ENI) is an elastic network interface bound to CVMs in a VPC, which can be migrated among multiple CVMs. ENI is very useful for configuring management networks and establishing highly reliable network solutions.

ENI has VPC, availability zone and subnet attributes. You can only bind it to CVMs under the same availability zone. A CVM can be bound with multiple ENIs, and the maximum number allowed varies by CVM specifications.

What are the restrictions to use ENIs on CVMs?

For details, please see use limits section in [Use Limits Overview](#).

What is the basic information of an ENI?

Please see **Concepts** in [Elastic Network Interface \(ENI\)](#).

How do I create an ENI?

Please see [Creating an ENI](#).

How do I view the ENI information?

Please see [Viewing ENI Information](#).

How do I bind an ENI to a CVM instance?

Please see [Binding and Configuring CVMs](#).

How do I configure an ENI in the CVM instance?

Please see [Binding and Configuring CVMs](#).

How do I modify or customize the private IP of an ENI?

CVMs in VPC support modifying and customizing the private IP of an ENI. Follow the steps below:

1. Log in to [VPC Console](#).
2. In the left side bar, click **IP and Interface** > **ENI** to enter the ENI list page.
3. Click the **ID/Name** of an ENI to enter its details page to view its information.
4. Select the **IPv4 address management** tab and click **Assign Private IP**.
5. In the pop-up window, select the IP assigning method as **Enter manually** to enter the IP address you want to modify.
6. Click **OK** to complete the operation.

After the modification is made on the console, you also need to modify the configuration file of the ENI. For more information, please see [Binding and Configuring CVMs](#).

Public Network Bandwidth

Last updated : 2024-01-06 17:26:45

What do the inbound bandwidth and the outbound bandwidth for CVM instances mean?

Inbound bandwidth: a bandwidth that flows into CVM instances. For example, the bandwidth resulting from downloading external network resources on the CVM instance.

Outbound bandwidth: a bandwidth that flows out of CVM instances. For example, the bandwidth resulting from external access from the CVM instance.

How is the public network bandwidth billed for CVM instances?

For more information, see [Public Network Billing](#).

The bandwidth cap for pay-as-you-go instances using the bill-by-traffic billing mode is 200 Mbit/s. How do I increase the bandwidth cap?

The public network bandwidth cap depends on the billing modes and the configurations of CVM instances. For more information, see [Public Network Bandwidth Cap](#).

Network

Last updated : 2024-01-06 17:26:45

There is no network connection after I log in to the CVM. How do I solve this problem?

If there is no network access (for example, the webpage is inaccessible) after you log in to the CVM, you need to check the DNS configuration. Set the CVM to obtain the DNS address automatically by following the instructions below.

Note:

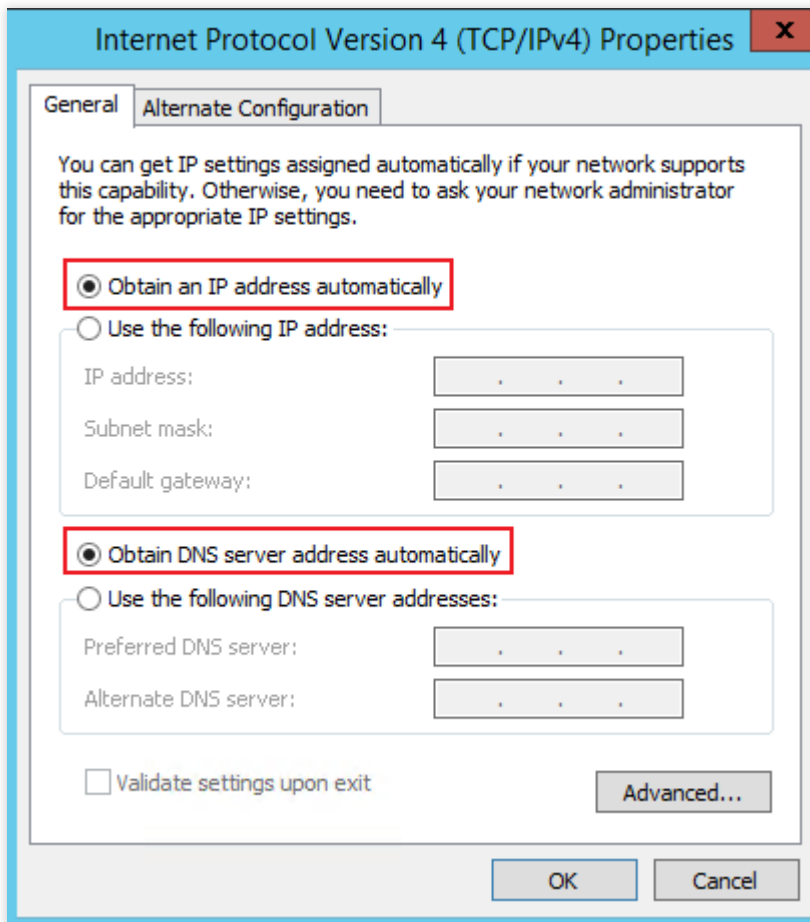
The following operations take Windows Server 2012 as an example.

1. Log in to the Windows CVM.
2. On the desktop, click



and select **Control panel -> Network and Internet -> View network status and tasks -> Change adapter settings**.

3. Right-click **Ethernet** and select **Properties** to open the "Ethernet Properties" window.
4. In the "Ethernet Properties" window, double-click and open the **Internet Protocol Version 4 (TCP/IPv4) Properties** window.
5. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** window, select **Obtain an IP address automatically** and **Obtain DNS server address automatically**, and click **OK**, as shown below:



Can a VPC instance interconnect with the classic network instance?

Yes, but it has the following restriction:

The VPC IP address range (CIDR) must be 10.0.0.0/16 - 10.0.47.0/16 (including subsets). Otherwise, there will be conflicts.

Directions

Log in to the [VPC Console](#), click the VPC ID/name to access its details page, and then click the **Classiclink** tab to associate the classic network CVM to be interconnected.

How do I view classic network CVMs that are interconnected with VPC CVMs?

Log in to the [VPC Console](#), click the VPC ID/name to access its details page, and view classic network CVMs interconnected with VPC CVMs in **Classiclink**.

Can the CVM be switched to an overseas network?

After being purchased, the network cannot be changed for CVM. If you need an overseas network, we recommend that you return the CVM and purchase an overseas CVM.

How do I configure a private network DNS?

Please see [Private Network Access](#).

Within the same IP range, the VPN can obtain the IP address of an IP range but cannot access the internet. How do I solve this problem?

Please check that the following configurations are correct:

1. Are the manually added IP and the automatically obtained IP in the same IP subnet? Are the subnet masks the same? Is the default gateway configured? Is the default gateway address correct?
2. Is DNS configured and is the DNS address correct?
3. If all of the configurations above are correct, check if the statically configured IP address has an IP conflict.

If none of the above works, please [submit a ticket](#) to contact us.

How do I add CVMs under accounts A and B to the same subnet?

By default, accounts are not interconnected by network. To interconnect accounts, see [Creating Cross-account Peering Connection](#) or [Network Instance Interconnection Crossing Account](#).

Security

Passwords

Last updated : 2024-01-06 17:26:45

How can I obtain the initial CVM password?

If you purchase the CVM through [Custom Configuration](#), you can obtain the initial password based on different login methods as follows:

Login Method	Description
Random Password	The initial password will be sent to you via email and Message Center on the console.
SSH Key Pair	Disabled by default you can log in to the CVM using your username and password, but the initial password will still be sent to you via email and Message Center on the console.
Custom Password	The password you set is the initial password.

For more information, see [Managing Login Password](#).

What should I do if I cannot receive the initial password via Message Center?

When you purchase a CVM, the initial password will be sent to you via email and [Message Center](#) on the console. If you cannot find the password in Message Center, please confirm that you log in to the console with a root account. If not, switch to the root account and check again.

What should I do if I forget my password?

You can reset the password. For more information, see [Resetting Instance Password](#).

How do I change or reset the CVM password?

To reset the password of your CVM, see [Resetting Instance Password](#).

What should I do if I fail to reset the password?

For Windows CVM, see [Failed to Reset the CVM Password or the CVM Password Is Invalid](#) to troubleshoot the issue. For Linux CVM, please [submit a ticket](#).

How do I reset CVM passwords in batches?

To reset CVM passwords in batches, see [Resetting Instance Password](#).

What are the default CVM login username and password?

The default username varies by operating systems:

For Windows instance, the default username is `Administrator` .

For Linux instance, the default username is `root` (`ubuntu` for Ubuntu system).

The CVM password for remote login is set by you when you create the instance. For more information, see [Creating Instances](#). If you forgot the login password, you can reset it. For more information, see [Resetting Instance Password](#).

Password Login and SSH Key Login

Last updated : 2024-03-26 10:30:53

What is the difference between SSH key login and password login?

An SSH key allows you to log in to a Linux server remotely. It uses the key generator to create a key pair (public and private). The public key is added to the server, and the client can use the private key to complete the authentication and login. Compared to password login, SSH key login is more secure and efficient.

Currently, Linux instance supports both password and SSH key login, while Windows instance only supports password login. For related documentation, see:

[Logging in to Linux Instance](#)

[Logging in to Windows Instance](#)

I cannot use the password to log in to a Linux instance that is associated with an SSH key.

After the CVM is associated with an SSH key, password login is **disabled by default**. See [Logging in to Linux Instance \(SSH Key\)](#).

Can I use SSH key together with the password login?

When you [log in to Linux instance via SSH key](#), password login is disabled by default to improve security.

How do I create an SSH key and what if I lose it?

For more information on key creation, see [Managing SSH keys](#).

If you lost the key, troubleshoot by the following methods:

Create a new key via [SSH Key](#) on the console and bind it with the original instance.

1.1 [Create an SSH key](#).

1.2 After the key is successfully created, log in to the [CVM Console](#).

1.3 Select the original instance with which you want to bind the key, click **More** -> **Password/key** -> **Load a key**.

You can then use the new key to log in to the instance.

Reset your password through the CVM Console and log in to the instance with your new password. For more information, see [Resetting Instance Password](#).

How do I bind/unbind an SSH key to or from the server?

For more information, see the **Binding/Unbinding a key to or from a CVM** section in [Managing SSH keys](#).

How do I modify the SSH key name/description?

For more information, see the **Modifying the SSH key name and description** section in [Managing SSH keys](#).

How do I delete an SSH key?

For more information, see the **Deleting SSH keys** section in [Managing SSH keys](#).

What are the use limits on SSH keys?

For more information, see the **Use Limits** section in [SSH Keys](#).

How to troubleshoot if I fail to log in to a Linux instance using an SSH key?

For more information, see [Unable to Log in to a Linux Instance via SSH](#).

I can not download my key.

The key can only be downloaded once. If you lost the key, we recommend that you create a new key, download and save it.

How can I query which key is used by the CVM instance?

You can log in to the CVM Console and go to the instance details page to view keys used by the CVM.

Security Groups

Last updated : 2024-01-06 17:26:45

Why does a security group have a reject rule by default?

Security group rules take effect in sequence from top to bottom. After a preceding allow rule is applied, other rules are ignored by default. If all ports are open to the internet, the last deny rule does not take effect. We provide this default configuration for security reasons.

How can I adjust the priority of a security group?

For more information, see [Adjusting Security Group Priority](#).

If I bind an incorrect security group to an instance, what impact will this have on the instance? How can I fix this problem?

Potential problems

You may fail to remotely connect to a Linux instance over SSH or a Windows instance via remote desktop.

You may fail to remotely ping the public IP and private IP addresses of the CVM instance in this security group.

You may fail to access over HTTP the web services exposed by the CVM instance in this security group.

The CVM instance in this security group may fail to access Internet services.

Solutions

If any of the aforementioned problems occur, you can go to **Security Groups** in the console and modify the security group rule. For example, you can change the rule to "bind only all-ports-open security groups by default".

For more information on how to set a security group rule, see [Security Groups](#).

What do the security group direction and policy mean?

The security group policies involve inbound and outbound directions. Outbound policies are used to filter outbound traffic from the CVM, and inbound policies are used to filter inbound traffic to the CVM.

Security group policies are divided into those that **allow** and **refuse** traffic.

In which order do the security group policies take effect?

Security group policies take effect from top to bottom. When traffic passes through the security group, policy matching is performed from top to bottom. A policy takes effect once matching is successful.

Why can an IP address that is rejected by a security group still access the CVM instance?

Possible reasons:

The CVM is bound to multiple security groups, and this IP address is allowed by another security group among them.

This IP address belongs to an approved Tencent Cloud public service.

Can iptables be used along with security groups?

Yes, iptables and security groups can be used at the same time. Your traffic will be filtered twice as follows:

Outbound: processes in your instance > iptables > security groups.

Inbound: security groups > iptables > processes in your instance.

Why can't security groups be deleted even though all CVM instances have been returned?

Check whether there are still bound CVMs in the recycle bin. The security group cannot be deleted even if it is bound to a CVM in the recycle bin.

Can the name of a cloned security group be the same as that of a security group in the target region?

No. The name must be different from names of existing security groups in the target region.

Can a security group be cloned across different users?

No.

Can I clone a security group to another project or region using an API?

To facilitate the use of the console, we provide MC support but not cloud API support. You can use the cloud APIs for importing and exporting security group rules in batches to clone security groups across projects and regions.

When I clone a security group across projects and regions, will CVM instances managed by the security group also be copied?

No. Only the inbound and outbound rules of the security group are cloned. You need to associate CVMs again after cloning.

What is a security group?

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.

Each CVM instance is bound to at least one security group, which must be specified when the instance is created. CVM instances are interconnected in the same security group but cannot communicate with CVM instances across security groups. However, you can authorize the interconnection between two security groups. For more information, see [Security Groups](#).

Why do I need to choose a security group when creating CVM instances?

Before creating CVM instances, you must choose a security group to partition the security domains of the application environment and authorize the security group rules to properly implement network security isolation.

What should I do when I create CVM instances if I have not created a security group?

In this case, you can create a security group.

A security group supports the following rules. You can open IPs and ports as needed.

ICMP: opens to the ICMP protocol and allows the pinging of the server over the public network.

TCP:80: opens port 80 and allows access to Web services over HTTP.

TCP:22: opens port 22 and allows a remote connection to Linux CVMs over SSH.

TCP:443: opens port 443 and allows access to Web services over HTTPS.

TCP:3389: opens port 3389 and allows a remote connection to Windows CVMs over RDP.

Open the private network: opens to the private network and allows private network access among different cloud resources (IPv4).

In what situations will the default security group rule be used to a security group?

The default security group rule will be used in the following situations:

When creating a CVM instance, if you select **Quick Configuration** to purchase the instance, the system will automatically create a default security group for you, which will adopt the default security rule (i.e., allowing all IPv4 addresses).

However, for security reasons, we recommend you associate your CVM with a new security group that only opens ports when absolutely necessary to avoid unnecessary security risks.

You can choose a template when creating a security group in the CVM console. Currently, supported templates include Windows login, Linux login, Ping, HTTP(80) and HTTPS(443).

Security Group Rules

Last updated : 2024-01-06 17:26:45

In what scenarios do I need to add security group rules?

To ensure the accessibility of CVM instances, you need to add security group rules in the following scenarios:

There is no security group rule (added or default) for CVM instances. If your CVM instances need to access the public network or communicate with CVM instances associated to other security groups in the same region, you need to add security rules.

The application you built uses a custom port or port range, instead of the default port. In this case, you need to open the custom port or port range before testing the application connectivity. For example, if you set up an Nginx service on your CVM instances that listens on the TCP 1800 communication port but your security group only opens port 80, you need to add security group rules to ensure the accessibility of the Nginx service.

For other scenarios, see [Security Group Use Cases](#).

What impact will incorrectly configuring security group rules have?

If any security group rule is configured incorrectly, CVM instances will fail to access other devices over the private network or public network. For example:

You may fail to remotely connect to a Linux instance over SSH or a Windows instance via remote desktop.

You may fail to remotely ping the public IP of a CVM instance.

You may fail to access Web services provided by CVM instances over the HTTP or HTTPS protocol.

You may fail to access other CVM instances over the private network.

Are the inbound rules and outbound rules of a security group counted separately?

A maximum of 100 inbound and 100 outbound rules can be set for a security group.

Can I adjust the maximum number of security group rules allowed?

Each security group can configure up to 200 security group rules, including 100 inbound rules and 100 outbound rules. One CVM instance can be associated with up to 5 security groups and can therefore adopt up to 1,000 security group rules, which is enough to meet the needs of most scenarios.

If your usage surpasses this upper limit, check whether redundant rules exist.

If so, remove them.

If not, create more security groups.

In addition, you can [submit a ticket](#) to raise the upper limit of security group rules.

Port

Last updated : 2024-01-06 17:26:45

Which ports should be opened to the Internet before I log in to an instance?

Generally, you need to open port 22 for a Linux instance, or port 3389 for a Windows instance. For more information about ports applicable to other instance types, see [Security Group Use Cases](#).

Which ports are often used in CVM?

For more information, see [Server Common Port](#).

Why should a port be opened to the Internet? How can I open a specific port?

You can use the service only after you open the port to the Internet in the security group. For example:

If you want to access web pages using port 8080, the port must be enabled and opened to the Internet in the security group.

To open a port to the Internet, follow the steps below:

1. Go to the [security group page](#), and click the ID/name of the security group bound with this instance to go to its details page.
2. Select **Inbound/Outbound rule** and click **Add a Rule**.
3. Enter your IP address (range) and port to be opened, and then select **Allow** to open the port.

For more information, see [Adding Security Group Rules](#).

How can I change the default remote port of a CVM instance?

For more information, see [Modifying the Default Remote Port of CVM](#).

Why cannot the service be used after I change the port?

After modifying the service port, you also need to open the corresponding port to the Internet in the security group.

Which ports are not supported by Tencent Cloud?

By default, TCP port 25 is blocked by Tencent Cloud. If you need to unblock this port, see [Unblocking Port 25](#).

Some ports have security risks. Although these ports are not blocked by Tencent Cloud, they will be blocked by ISPs and cannot be accessed. To avoid this, we recommend that you change ports and do not use the following ports for listening:

Protocol	Blocked Port
TCP	42, 135, 137, 138, 139, 445, 593, 1025, 1434, 1068, 3127, 3128, 3129, 3130, 4444, 5554, 5800, 5900, 9996
UDP	1026, 1027, 1434, 1068, 5554, 9996, 1028, 1433, 135 - 139

Why cannot I use TCP 25 port?

TCP port 25 is the default email service port. For security reasons, port 25 of CVM instances is blocked by default. If you need to use it, see [Unblocking Port 25](#).

Firewall

Last updated : 2024-05-29 15:50:55

How should I configure the firewall software iptables on a Linux OS?

Note:

iptables varies greatly between versions earlier and later than CentOS 7.

In versions earlier than CentOS 7, the iptables service is used as the firewall by default. After running the `service iptables stop` command, the iptables service will first clear rules and then unload the iptables module. When the iptables service restarts, it will load rules from the configuration file. You can stop the iptables service to test whether the firewall restrictions were applied.

In versions later than CentOS 7, the firewall service is used as the firewall by default, and the iptables_filter module is loaded to ensure compatibility. You can use the iptables command to add rules. However, the iptables service is disabled by default. After you confirm that the iptable_filter module is loaded, the rules will take effect.

To determine the firewall, run `iptables -nvL` to view the rules.

The following two scenarios describe how to configure the iptables firewall software program.

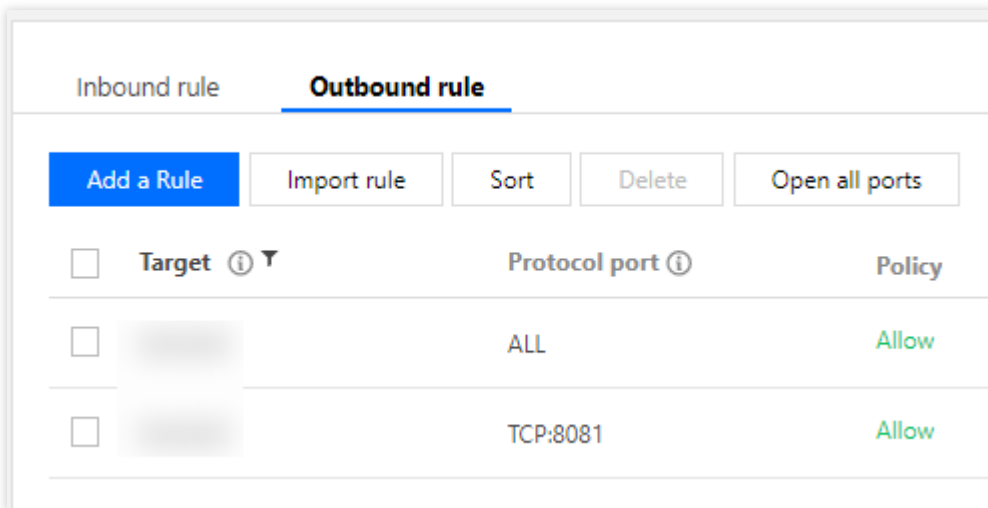
Scenario 1

In an Ubuntu 14 OS, the security group and listening port are enabled, but the Telnet connection fails.

Security group inbound rules:

Security Group Rule			Associate with Instance	
Inbound rule			Outbound rule	
<div>Add a Rule</div>			<div>Import rule</div>	
			<div>Sort</div>	
			<div>Delete</div>	
			<div>Open all ports</div>	
<input type="checkbox"/>	Source ⓘ ▾	Protocol port ⓘ	Policy	
<input type="checkbox"/>		ALL	Allow	
<input type="checkbox"/>		TCP:8081	Allow	

Security group outbound rules:



Telnet connection failure:

```
target      prot opt source                destination              limit: avg 3/se
RETURN     tcp  -- anywhere              anywhere
urst 6
REJECT     all  -- anywhere              anywhere                  reject-with icmp
ort-unreachable
root@VM-166-120-ubuntu:~# service iptables status
iptables: unrecognized service
root@VM-166-120-ubuntu:~# ufw status
Status: inactive
root@VM-166-120-ubuntu:~# netstat -tlnp | grep 8081
tcp        0      0 0.0.0.0:8081          0.0.0.0:*               LISTEN
1687/nginx.conf
root@VM-166-120-ubuntu:~# _
```

Solution

1. Capture the packets on the CVM to check whether the packets were sent to the CVM.

If no, the packets may be blocked by the security group, the upper TGW, or the ISP.

If yes but the CVM does not respond, this issue may be caused by the iptables policy of the CVM. As shown in the following figure, the CVM does not return TCP packets to 64.11 after the Telnet connection is established.

```
0 packets dropped by kernel
root@VM-166-120-ubuntu:/home/ubuntu# tcpdump -i any host 183.60.64.11
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on any, link-type LINUX_SLL (Linux cooked), capture size 65535 bytes
19:02:21.194801 IP 183.60.64.11.58839 > 10.104.166.120.tproxy: Flags [S], seq 668485174, win 5840, options [mss 1424,sackOK,TS val 496201777 e
op,wscale 7], length 0
19:02:21.194824 IP 10.104.166.120 > 183.60.64.11: ICMP host 10.104.166.120 unreachable - admin prohibited, length 68
```

2. After confirming that the issue is caused by the iptables policy, run `iptables -nvL` to check whether the policy opened port 8081. In this example, port 8081 is closed.

```

root@VM-166-120-ubuntu:~# iptables -L -n
Chain INPUT (policy ACCEPT)
target    prot opt source                destination
ACCEPT    all  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    all  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    tcp  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    tcp  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    tcp  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    tcp  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    tcp  --  0.0.0.0/0              0.0.0.0/0
ACCEPT    icmp --  0.0.0.0/0              0.0.0.0/0
ACCEPT    icmp --  0.0.0.0/0              0.0.0.0/0
syn-flood tcp  --  0.0.0.0/0              0.0.0.0/0
REJECT    all  --  0.0.0.0/0              0.0.0.0/0

```

Inbound direction

Inbound packets are allowed when a connection has been established. This is intended for receiving packets for active access.

state RELATED,ESTABLISHED

state NEW tcp dpt:22

state NEW tcp dpt:80

state NEW tcp dpt:21

state NEW tcp dpts:20000:30000

state NEW tcp dpt:443

limit: avg 100/sec burst 100

limit: avg 1/sec burst 10

tcp flags:0x17/0x02

reject-with icmp-host-prohibited

Open the port for access by external network

3. Run the following command to open port 8081:

```
iptables -I INPUT 5 -p tcp --dport 8081 -j ACCEPT
```

4. Check whether port 8081 is open. If yes, the issue has been solved.

Scenario 2

Based on the iptables configuration, the policy has been enabled, but the destination server still cannot be pinged through.

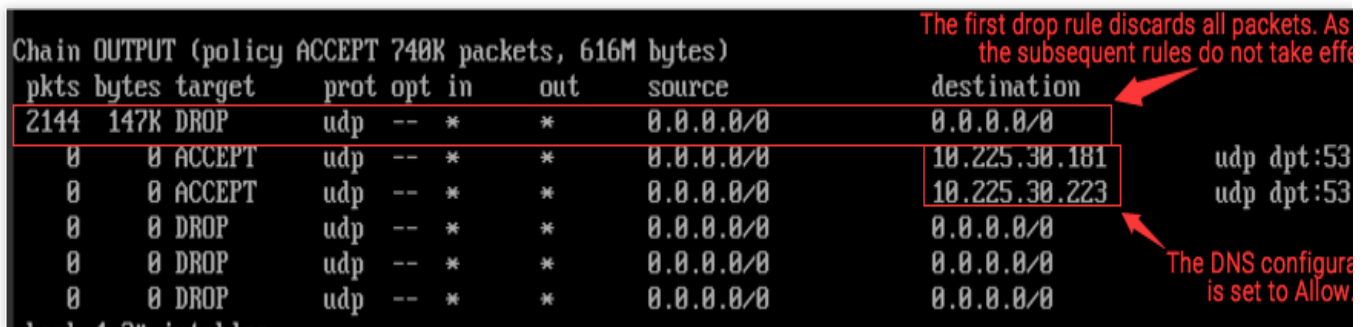
```

-bash-4.2# cat /etc/resolv.conf
#search localdomain
#nameserver 202.98.
#nameserver 61.139.
nameserver 10.225.30.181
nameserver 10.225.30.223
options timeout:1 rotate
-bash-4.2# vi /etc/resolv.conf
-bash-4.2# grep host /etc/nsswitch.conf
#hosts:      db files nisplus nis dns
hosts:      files dns
-bash-4.2# grep GATEWAY /etc/sysconfig/network-scripts/ifcfg*
/etc/sysconfig/network-scripts/ifcfg-eth0:GATEWAY='10.104.61.1'
-bash-4.2#

```

Solution

If the information shown below appears,



The first drop rule discards all packets. As the subsequent rules do not take effect

Chain OUTPUT (policy ACCEPT 740K packets, 616M bytes)									
pkts	bytes	target	prot	opt	in	out	source	destination	
2144	147K	DROP	udp	--	*	*	0.0.0.0/0	0.0.0.0/0	
0	0	ACCEPT	udp	--	*	*	0.0.0.0/0	10.225.30.181	udp dpt:53
0	0	ACCEPT	udp	--	*	*	0.0.0.0/0	10.225.30.223	udp dpt:53
0	0	DROP	udp	--	*	*	0.0.0.0/0	0.0.0.0/0	
0	0	DROP	udp	--	*	*	0.0.0.0/0	0.0.0.0/0	
0	0	DROP	udp	--	*	*	0.0.0.0/0	0.0.0.0/0	

The DNS configuration is set to Allow

run the following command to delete the first output rule:

```
iptables -D OUTPUT 1
```

Test to verify that the issue has been solved.

How can I remove a firewall?

Windows instances:

1. After logging in to the instance, choose **Start > Control Panel > Firewall Settings** to access the "Firewall Settings" page.
2. Check whether the firewall and other security software, such as safedog, have been enabled. If yes, disable them.

Linux instances:

1. Run the following command to check whether the firewall policy has been enabled. If no, skip step 2 and go to step 3.

```
iptables -vnL
```

2. If the firewall policy has been enabled, run the following command to back up the firewall policy:

```
iptables-save
```

3. Run the following command to clear the firewall policy:

```
iptables -F
```

Will the firewall intercept a non-Tencent Cloud CDN CVM?

No. If you are worried that the firewall may disrupt your business, you can disable the firewall.

Operating Systems

Last updated : 2024-01-06 17:26:45

Why an error occurs when I execute an update on Ubuntu OS?

Possible reasons:

Tencent Cloud Ubuntu system source is synced with the official source at 00:00, 08:00 and 16:00 every day (UTC+8). If you run the `apt-get update` and `apt-get install` commands just before and after the time points, the apt verification will be failed and an error will occur.

Solution:

Run the `apt-get clean all` command to purge cache before you run the `apt-get update` command.

How do I change the owners and owner groups of directories and files on a Linux instance?

If the file or directory permissions are incorrectly configured on the Web server, a 403 error will occur when you access a website hosted on the instance. Therefore, before you adjust a file or directory, you must confirm the identity under which the file or directory process is running.

You can run the `ps` and `grep` commands to query the identities under which the file or directory process is running.

You can run the `ls -l` command to query the owners and owner groups of files and directories.

You can run the `chown` command to modify the permissions. For example, you can run the `chown -R www.www /tencentcloud/www/user/` command to change the owners and owner groups of all files and directories under the `/tencentcloud/www/user` directory to account "www".

Do Linux instances support GUI?

Yes. For directions on building GUI on a Linux instance, see [Building a Visual Ubuntu Desktop](#).

How can I activate the operating system of a Windows CVM instance?

You can refer to [Activating Windows with slmgr Command](#) or [Windows Server System Activation](#) for directions.

Why do I need to boot into the single-user mode on a Linux instance, and how to do this?

Linux's users sometimes need to boot into the single-user mode to perform special operations, such as password management, sshd repair, or maintenance before disk attaching. For directions, see [Booting into Linux Single User Mode](#).

How can I check the login logs of a CVM instance?

For more information, see [Getting Instance Login Logs](#).

OPS & Monitoring

Uploading Files

Last updated : 2024-03-26 10:31:21

Does CVM come with FTP?

Yes. You can install and configure FTP as needed.

To build FTP sites using Windows' built-in FTP service, see [Building the FTP Service \(Windows\)](#).

To build FTP sites using vsftpd software, see [Building the FTP Service \(Linux\)](#).

How do I upload files to a Windows CVM?

For a local Windows, see [Uploading Files from Windows to a Windows CVM using MSTSC](#).

For a local Linux, see [Uploading Files from Linux to Windows CVM using RDP](#).

For a local MacOS, see [Uploading Files from MacOS to Windows CVM using MRD](#).

How does a local host transfer data to and from a Windows CVM?

To transfer data:

If you use the RDP file to log in to a Windows CVM from a local Windows computer, you can directly drag and upload local files to the CVM.

[Upload Files from Windows to a Windows CVM using MSTSC](#)

[Build the FTP Service \(Linux\)](#) on the local host.

How do I upload files to a Linux CVM using WinSCP?

For more information, see [Uploading Files via WinSCP to a Linux CVM from Windows](#).

How do I upload files to or download files from a Linux CVM?

For more information, see [Uploading Files via SCP to a Linux CVM from Linux](#).

What should I do if the client-server connection times out during FTP file upload?

The server firewall or security group may have blocked the connection. Follow the steps below to troubleshoot this issue:

1. Check the server's firewall settings.
2. Disable the firewall or add rules.

OPS

Last updated : 2024-01-06 17:26:45

Do you have any advice regarding OPS for small websites hosted on CVM?

To maintain the website applications, you can:

Back up data to the cloud disk daily. For more information, see [Creating Snapshots](#).

Use [SSL Certificates Service](#) for identity verification and encrypted connections.

Install anti-virus plugins or anti-DDoS services or purchase Cloud Workload Protection.

Monitor the traffic to and from the website and identify exceptions in the traffic range. Add the security group rules of denied access to temporarily control the exception request of a single point. For more information, see [Getting Monitoring Statistics](#) and [Adding Security Group Rules](#).

Monitor the performance of CVM instances and the cloud disk, and mark the traffic/access peak period. Familiarize yourself in advance with upgrade/degrade, auto scaling, or cloud disk expansion to respond to request surges. For more information, see [Change Instance Configuration](#), [What is Auto Scaling \(AS\)?](#), or [Cloud Disk Expansion Scenarios](#).

Update the admin password regularly for scenarios where you log in to CVM instances with the root/Administrator username and password. For more information, see [Reset Instance Password](#).

Update the software patches regularly.

Common Operations and Commands

Last updated : 2024-01-06 17:26:45

What is the Linux server load average?

Load is used to measure the workload of a server; that is, the length of the queue of tasks to be executed by the CPU. The greater the value, the more processes that are currently running or waiting to be executed.

How can I check the Linux server load?

You can run the `w` , `top` , `uptime` , and `procinfo` commands or access the `/proc/loadavg` file to view the load.

Please refer to "Installing Software in the Linux Environment" for instructions on how to install the `procinfo` tool.

What should I do when the server load is too high?

The server load/load average is displayed based on the length of the process queue. A high server load (we recommend referencing the 15-minute average) may be caused by insufficient CPU resources, I/O read/write bottlenecks, insufficient memory resources, or intensive computing tasks. We recommend you run the `vmstat` , `iostat` , and `top` commands to identify the reason for the high load and optimize the processes.

How do I check the server memory usage?

You can check the server memory usage by running the `free` , `top` (after running, you can press `shift+m` to sort the memory), `vmstat` , and `procinfo` commands or by accessing the `/proc/meminfo` file.

How do I check the memory usage of a single process?

You can check the memory usage of a single process by running the `top -p PID` , `pmap -x PID` , and `ps aux|grep PID` commands or by accessing the `/proc/$process_id (process PID) /status` file (for example, the `/proc/7159/status` file).

How do I check the services and ports that are currently in use?

You can check the services and ports that are currently in use by running the `netstat -tunlp` , `netstat -antup` , and `lsof -i:PORT` commands.

How do I check the server process information?

You can check the server process information by running the `ps auxww|grep PID` , `ps -ef` , `lsof -p PID` , and `top -p PID` commands.

How do I stop a process?

You can run `kill -9 PID` (PID indicates the process ID) and `killall program name` (for example, `killall cron`) to stop a process.

To stop a zombie process, you need to kill the parent process by running `kill -9 ppid` (ppid indicates the parent process ID, which can be queried by running `ps -o ppid PID` (for example, `ps -o ppid 32535`)).

How do I locate a zombie process?

You can run `top` to view the total of zombie processes, and run `ps -ef | grep defunct | grep -v grep` to locate a specific zombie process.

Why can't I enable the server port?

You need to check the operating system and the application to ensure a port is enabled.

Only the `root` user can enable ports below 1024 on the Linux operating system. Run `sudo su -` to obtain root permissions before enabling the server port.

For application issues such as port conflicts or configuration problems, use the application startup logs to troubleshoot them. The Tencent server system uses port 36000.

What are the commands commonly used to check the performance of a Linux server?

Command Name	Description
top	top is a task manager program that monitors the overall performance of the system. This command can be used to display information such as the system load, the process, the CPU, the memory, and paging. Use <code>shift+m</code> and <code>shift+p</code> to sort the processes by memory usage and CPU usage.
vmstat	vmstat is a computer system monitoring tool mainly used for virtual memory that collects and displays summary information about CPU, processes, memory paging, and IO. For example, <code>vmstat 3 10</code> outputs results every 3 seconds and is executed 10 times.
iostat	iostat is a computer system monitoring tool that collects and displays statistics about CPU and IO. For example, <code>iostat -dxmt 10</code> outputs detailed information about IO in MB every 10 seconds.
df	df is a command used to display the amount of available disk space. For example: <code>df -m</code> displays the disk space usage in MB.
lsof	<p>lsof reports a list of all open files, which is very useful for Linux system management.</p> <p>For example:</p> <p><code>lsof -i: 36000</code> displays the processes using port 36000.</p> <p><code>lsof -u root</code> displays the programs run by root.</p> <p><code>lsof -c php-fpm</code> displays the files opened by the php-fpm process.</p> <p><code>lsof php.ini</code> displays the processes for which php.ini is opened.</p>

ps	<p>ps is a process query command that displays information related to the processes. Commonly used command parameter combinations are <code>ps -ef</code> and <code>ps aux</code>. Use <code>ps -A -o</code> to output custom fields.</p> <p>For example:</p> <p><code>ps -A -o pid,stat,uname,%cpu,%mem,rss,args,lstart,etime sort -k6,6 -rn</code> outputs results according to the listed fields and sorts them using the 6th field.</p> <p><code>ps -A -o comm sort -k1 uniq -c sort -k1 -rn head</code> lists the process with the largest number of running instances.</p>
----	---

Other commonly used commands and files: `free -m` , `du` , `uptime` , `w` , `/proc/stat` , `/proc/cpuinfo` , and `/proc/meminfo` .

What do I do when Cron doesn't work?

Follow the steps below to troubleshoot this problem:

1. Verify whether crontab is running normally.
2. Run `crontab -e` to add the following test item.

```
\\*/1 \\* \\* \\* \\* /bin/date >> /tmp/crontest 2>&1 &
```

3. Check the `/tmp/crontest` file.

If there are any problems, run `ps aux|grep cron` to look for the pid of cron, run `kill -9 PID` to terminate the cron process, and then run `/etc/init.d/cron start` to restart cron.

4. Verify whether the script path in the cron entry is an absolute path.
5. Check whether you used the right account to execute cron. Also, check whether the account is included in `/etc/cron.deny` .

6. Check the execution permission of the script, the script directory, and the log file permission.

7. We recommend that you run the script in the background by adding an "&" at the end of the script entry. For example, `*/1 * * * * /bin/date >> /tmp/crontest 2>&1 &` .

How do I set up a startup task for my CVM instance?

The Linux kernel startup sequence is as follows:

1. Start the `/sbin/init` process.
2. Run the init initial scripts in sequence.
3. Run the level script `/etc/rc.d/rc*.d` , where the value of * means the running mode, which can be queried in `/etc/inittab` .
4. Run `/etc/rc.d/rc.local` .

Note:

The startup task can be configured in `/etc/rc.d/rc.local` or in the `S**rclocal` file in `/etc/rc.d/rc*.d`.

Why is the server disk read-only?

Common reasons for a read-only server disk are as follows:

The disk space is full

You can run `df -m` to check the disk usage and then delete unnecessary files to free disk space. (We do not recommend you delete non-third party files. Please check the files before you delete them).

Disk inode resources are all occupied

You can run `df -i` to view and confirm relevant processes.

There is a hardware failure

If none of the above works, please [submit a ticket](#).

How do I view Linux system logs?

The storage path for system-level log files is `/var/log`.

The commonly used system log is `/var/log/messages`.

How do I find large files in the file system?

You can find them by following the steps below:

1. Run `df` to query the disk partition usage (for example, `df -m`).
2. Run `du` to query the size of a specific folder (for example, `du -sh ./*`, `du -h --max-depth=1|head -10`).
3. Run `ls` to list the files and file sizes (for example, `ls -lSh`).

You can also directly check the size of files under a specific directory by using the `find` command (for example, `find / -type f -size +10M -exec ls -lrt {} \\`).

How do I check the version of a server's operating system?

You can run the following commands to check the version of a server's operating system:

```
uname -a
```

```
cat /proc/version
```

```
cat /etc/issue
```

Why are Chinese characters displayed as unintelligible code in the Linux terminal?

The server itself does not impose restrictions on the display language. Therefore, this is most likely a client problem. If you are using secureCRT, try to change the settings in **Options** -> **Session Options** -> **Appearance**. If you are using other terminal software, please search on Google for solutions.

If this problem appears in a pure Linux shell, run `export` to check the user environment variables such as LANG and LC_CTYPE.

How do I configure a connection timeout using SecureCRT?

You can configure the following settings so that the connection to your CVM instance remains stable when connecting to CVM via SecureCRT:

1. Open SecureCRT and select **Options**.
2. Select **Session Options** and click **Terminal**.
3. Select **Send protocol NO-OP** in the **Anti-idle** box on the right and set the time to "every 120 seconds".

Why isn't disk space on a Linux server freed after a file is deleted?

Cause:

After logging in to your Linux CVM instance and deleting a file using `rm`, you may find that the available disk space does not increase when you use `df` to check the disk space. This is because when the file is deleted, if another process happens to be accessing the file, the space occupied by the deleted file will not be immediately freed at the time you check the disk space.

Solution:

1. Run `ls -l | grep deleted` using the root permission and find the PID of the process that is using the deleted file.
2. Kill the process using `kill -9 PID`.

How do I delete files on a Linux server?

You can run `rm` to delete files. Files deleted with this command cannot be recovered. Therefore, please use this command with caution.

Format of `rm`: `rm (option) (parameter)`.

Options:

- d**: directly deletes all hardwired data contained in the directory to be deleted, and then deletes the directory itself.
- f**: forcibly deletes a file or a directory.
- i**: asks the user before deleting an existing file or directory.
- r** or -**R**: deletes all files and sub-directories under a specific directory together using recursive processing.
- preserve-root**: does not implement recursive operations on the root directory.
- v**: displays the detailed execution processes of the commands.

Parameter: specifies the file or file list to be deleted. If the parameter contains a directory, add the `-r` or `-R` option.

Example:

Run `rm test.txt` to delete the `test.txt` file.

Run `rm -r test` to delete the `test` directory.

Run `rm -r *` to delete all files and sub-directories under the current directory.

Access Control

Last updated : 2024-01-06 17:26:45

How can I create a custom policy?

If preset policies do not meet your requirements, you can create a custom policy.

The syntax of a custom policy is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "Action"
      ],
      "resource": "Resource",
      "effect": "Effect"
    }
  ]
}
```

Action: replace this with the operation to be allowed or denied.

Resource: replace this with the resources that you want to authorize.

Effect: replace this with "Allow" or "Deny".

How should I configure the read-only policy for a CVM?

To allow a user to query CVM instances but not to create, delete, start, or shut down the instances, enable the QcloudCVMInnerReadOnlyAccess policy.

To do this, log in to the CAM Console. On the [Policies](#) page, search for **CVM** to find the policy.

The policy syntax is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "name/cvm:Describe*",
        "name/cvm:Inquiry*"
      ],
      "resource": "*",
      "effect": "allow"
    }
  ]
}
```

The preceding policy **grants users the permissions to perform the following operations**:

All operations starting with "Describe" in CVM.

All operations starting with "Inquiry" in CVM.

How can I configure the read-only policy for CVM-related resources?

To allow a user to query CVM instances and relevant resources (VPC and CLB instances) but not to create, delete, start, or shut down the instances and relevant resources, enable the QcloudCVMReadOnlyAccess policy.

To do this, log in to the CAM Console. On the [Policies](#) page, search for **CVM** to find the policy.

The policy syntax is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "name/cvm:Describe*",
        "name/cvm:Inquiry*"
      ],
      "resource": "*",
      "effect": "allow"
    },
    {
      "action": [
        "name/vpc:Describe*",
        "name/vpc:Inquiry*",
        "name/vpc:Get*"
      ],
      "resource": "*",
      "effect": "allow"
    },
    {
      "action": [
        "name/clb:Describe*"
      ],
      "resource": "*",
      "effect": "allow"
    },
    {
      "effect": "allow",
      "action": "name/monitor:*",
      "resource": "*"
    }
  ]
}
```

The preceding policy **grants 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.

NTP Service

Last updated : 2024-01-06 17:26:45

How can I adjust the sync interval of NTP after configuring the NTP service?

After configuring the NTP service as instructed in [Setting Up NTP Service \(Linux\)](#), you can restart `ntpd` to reset the sync interval. To manually set the `ntpd` synch interval, perform the following steps:

1. Run the following command to modify the NTP configuration file.

```
vi /etc/ntp.conf
```

2. Press `i` to enter the edit mode and configure as follows:

2.1 Add the pound sign (`#`) at the beginning of `server time1.tencentyun.com iburst` , if any, to comment it out.

2.2 ii. Add the following configurations. `minpoll 4` means the minimum interval is 2^4 , and `maxpoll 5` means the maximum interval is 2^5 .

```
server time1.tencentyun.com minpoll 4 maxpoll 5
```

The result should be as follows. Enter `:wq` to save changes and close the file.

```
disable monitor
#server time1.tencentyun.com iburst
server time1.tencentyun.com minpoll 4 maxpoll 5
interface ignore wildcard
interface listen eth0
```

3. Restart `ntpd` and run `ntpq -p` . The `poll` value you will see is 16 (namely, 2^4).

```
[root@ ~]# ntpq -p
  remote           refid      st t when poll reach  delay  offset  jitter
=====
*169.254.0.2       169.254.0.196    2 u   4   16  377   4.199   -0.184   0.062
```

What is the time source of Tencent Cloud's ntpd servers?

NTP servers usually use the BeiDou satellite clock.

What should I do when the NTP configuration reports the `localhost.localdomain timeout` error?

The error is shown below.

```

[root@VM_0_17_centos ~]# netstat -nupl
Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State       PID/Program name
udp      0      0 0.0.0.0:39083           0.0.0.0:*               *          31900/ntpdate
udp      0      0 0.0.0.0:35385           0.0.0.0:*               *          1067/dhclient
udp      0      0 0.0.0.0:68              0.0.0.0:*               *          1067/dhclient
udp      0      0 172.17.0.17:123         0.0.0.0:*               *          31876/ntpd
udp      0      0 127.0.0.1:123           0.0.0.0:*               *          31876/ntpd
udp6     0      0 :::62167                :::*                    *          1067/dhclient

[root@VM_0_17_centos ~]# service ntpd status
Redirecting to /bin/systemctl status ntpd.service
● ntpd.service - Network Time Service
   Loaded: loaded (/usr/lib/systemd/system/ntpd.service; enabled; vendor preset: disabled)
   Active: active (running) since Wed 2020-08-26 19:29:56 CST; 15s ago
     Process: 31874 ExecStart=/usr/sbin/ntpd -u ntp:ntp $OPTIONS (code=exited, status=0/SUCCESS)
    Main PID: 31876 (ntpd)
      CGroup: /system.slice/ntpd.service
              └─31876 /usr/sbin/ntpd -u ntp:ntp -g

Aug 26 19:29:56 VM_0_17_centos systemd[1]: Started Network Time Service.
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: proto: precision = 0.078 usec
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: 0.0.0.0 c01d 0d kern kernel time sync enabled
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: ntp_io: estimated max descriptors: 1000000, initial socket boundary:
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: Listen normally on 0 lo 127.0.0.1 UDP 123
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: Listen normally on 1 eth0 172.17.0.17 UDP 123
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: Listening on routing socket on fd #18 for interface updates
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: 0.0.0.0 c016 06 restart
Aug 26 19:29:56 VM_0_17_centos ntpd[31876]: 0.0.0.0 c012 02 freq_set kernel 0.065 PPM
Aug 26 19:30:03 VM_0_17_centos ntpd[31876]: 0.0.0.0 c615 05 clock_sync
[root@VM_0_17_centos ~]# ntpq -p
localhost.localdomain: timed out, nothing received
***Request timed out

```

Check and confirm whether you ran `POSTROUTING`. If yes, change the source IP in the `ntp.conf` configuration file to eth0's IP address.

Can off-cloud servers share a NTP with Tencent Cloud CVMs? What is the NTP synchronization address?

Tencent Cloud provides private NTP servers for Tencent Cloud resources. For other devices, you can use the public NTP servers provided by Tencent Cloud for synchronization.

Private NTP Server

```

time1.tencentyun.com
time2.tencentyun.com
time3.tencentyun.com
time4.tencentyun.com
time5.tencentyun.com

```

Public NTP Server

```

ntp.tencent.com
ntp1.tencent.com
ntp2.tencent.com
ntp3.tencent.com
ntp4.tencent.com
ntp5.tencent.com

```

The following are old public NTP server addresses. These old addresses can still be used, but you are advised to use the new ones.

```
time.cloud.tencent.com
time1.cloud.tencent.com
time2.cloud.tencent.com
time3.cloud.tencent.com
time4.cloud.tencent.com
time5.cloud.tencent.com
```

Why is the time of a CVM instance created from a custom image incorrect?

Please check whether the NTP service is enabled. To use the NTP synchronization feature later, see the following documents to set NTP. Then re-create the custom image.

[Setting Up NTP Service \(Linux\)](#)

[Setting Up NTP Service \(Windows\)](#)

Why is the ntp.conf content of a CVM created from a custom image restored?

It may be caused by the Cloud-Init initialization. Please delete the NTP-related configurations in

`/etc/cloud/cloud.cfg` before creating a custom image. For more information, see [Cloud-Init & Cloudbase-Init](#).

What impact will a modified private network DNS have?

All services involving Tencent Cloud private domain name resolution will be affected. For example, The YUM repository uses the Tencent Cloud private domain name. You need to modify the YUM repository after changing DNS.

The reporting of monitoring data uses the private domain name and will be affected.

The NTP feature depends on the private domain name to sync server time, and will be affected.

Why does the local time of a Windows instance reset from the configured EST to UTC+8 after restart?

Check that the windowstime service is enabled on the instance. You may need to manually enable it to auto sync the instance's system time. We recommend enabling the service autostart.

Why cannot I use the `ntpq -np` command to view the synchronization time?

The error is shown below.

```
[root@ ~]# ntpq -np
remote               refid               st t when poll reach  delay  offset  jitter
=====
100.115.8.147 .INIT.             16 -   -   64    0   0.000   0.000   0.000
```

Please check whether no IP or incorrect IP is configured for `listen` in the `/etc/ntp.conf` configuration file. Change to the primary private IP of the instance and restart `ntpd`.

How do I fix the error occurred when using the public NTP servers to sync time?

When using the public NTP servers to sync time, the `no server suitable for synchronization found` error occurs.

```
[root@VM-0-192-centos ~]# ntpdate -dv time1.cloud.tencent.com
29 Sep 10:20:45 ntpdate[10750]: ntpdate 4.2.6p5@1.2349-o wed Dec 19 20:22:35 UTC 2018 (1)
Looking for host time1.cloud.tencent.com and service ntp
host found : 139.199.215.251
transmit(139.199.215.251)
transmit(139.199.215.251)
transmit(139.199.215.251)
transmit(139.199.215.251)
transmit(139.199.215.251)
139.199.215.251: Server dropped: no data
server 139.199.215.251, port 123
stratum 0, precision 0, leap 00, trust 000
refid [139.199.215.251], delay 0.00000, dispersion 64.00000
transmitted 4, in filter 4
reference time: 00000000.00000000 Mon, Jan 1 1900 8:05:43.000
originate timestamp: 00000000.00000000 Mon, Jan 1 1900 8:05:43.000
transmit timestamp: e31d1882.6c402d75 Tue, Sep 29 2020 10:20:50.422
filter delay: 0.00000 0.00000 0.00000 0.00000 0.00000
0.00000 0.00000 0.00000 0.00000
filter offset: 0.000000 0.000000 0.000000 0.000000 0.000000
0.000000 0.000000 0.000000 0.000000
delay 0.00000, dispersion 64.00000
offset 0.000000
29 Sep 10:20:51 ntpdate[10750]: no server suitable for synchronization found
```

This may be caused by the NTP protection policy in response to the DDoS attacks to the public IP address of the instance. This policy blocks all public inbound traffic at the Tencent Cloud's source point 123 and causes synchronization exception. We recommend trying to use the private NTP servers to sync time.

Use Cases

Last updated : 2024-01-06 17:26:45

About Building a Website

[Setting up a Website](#)

[OPS](#)

[CVM Login Failures](#)

[Slow Website Access](#)

[Website Access Failure](#)

[Instance IP Address Ping Failure](#)

[How Do I Choose a Certificate?](#)

[Passwords](#)

[Use Limits](#)

About Building an Environment

[Accelerating Software Package Download and Update with Tencent Cloud Software Source](#)

[CVM Login Failures](#)

[Passwords](#)

[Failed to Reset the CVM Password or the CVM Password Is Invalid](#)

[Common Operations and Commands](#)

[Introduction to Linux Kernel Parameters](#)

[Booting into Linux Single User Mode](#)

[System and Data Disks](#)

[Changing Images \(Reinstalling System\)](#)

[Changing Instance Configuration](#)

About Data Backup

[Formatting and Mounting Data Disks](#)

[Initializing Data Disks \(Windows CVM\)](#)

[Expanding Cloud Disks](#)

[Copying Local Files to CVMs](#)

[Backup and Restore](#)

[Creating Snapshots](#)

[Creating Custom Images](#)[Reinstalling System](#)[Snapshots](#)**About Other Use Cases**[Security Group Use Cases](#)[Adjusting Security Group Priority](#)[Modifying VNC Resolution](#)[Modifying the Default Remote Port of CVM](#)[CVM Login Failures](#)[Instance IP Address Ping Failure](#)[Failed to Shut Down or Restart a CVM](#)[Failed to Reset the CVM Password or the CVM Password Is Invalid](#)[System Activation](#)[Setup Windows CVM For Multi-user Remote Login](#)