

Cloud Virtual Machine

Troubleshooting

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

Troubleshooting

Instance-Related Failures

CVM Login Failures

Windows Instance Login Failures

Windows Instance Login Failures

An authentication error occurred when you tried to log in to a Windows instance remotely

Failed to Reset the CVM Password or the CVM Password Is Invalid

Connection to a Windows CVM through Remote Desktop was denied

Requires network-level identity verification

Problems occurred when you tried to log in to a Windows CVM remotely on Mac

Failed to log in to a Windows CVM due to high CPU and memory usage

Failed to connect to a remote computer through Remote Desktop

Credentials Not Work

Windows instance: no remote Desktop license server can provide license

Remote Login Failure Due To Port Issues

Linux Instance Login Failures

Linux Instance Login Failures

Unable to Log in to a Linux Instance via SSH Key

Failing to log in to a Linux CVM due to high CPU and memory usage

Remote Login Failure due to Port Issues

VNC Login Error (Module is Unknown)

VNC Login Error (Account Locked due to XXX Failed Logins)

VNC Login Error (Login Failed with Correct Password)

VNC or SSH Login Error (Permission Denied)

Login Failure Due to /etc/fstab Configuration Errors

sshd Configuration File Permissions

Infinite Loop Call in /etc/profile

Login Failure Due to Server Isolation

Login Failure Due to High Bandwidth Occupation

Remote Connect Failure Due to Security Group Settings

Troubleshooting Linux Instance Issues via VNC and Rescue Mode

Failed to shut down or restart a CVM

Kernel and IO Issues

Missing System bin or lib Soft Link

Suspected Infection with Virus

"no space left on device" Error During File Creation

Linux CVM Memory Issues

High Memory Utilization

Log Error "fork: Cannot allocate memory"

VNC Login Error "Cannot allocate memory"

Triggering Out of Memory When There is Available Memory

Network Related Failures

Cross-MLC-Boarder Linkage Latency

Website Access Failure

Slow Website Access

Incorrect Multi-Queue ENI Configuration

CVM Network Latency and Packet Loss

Network Packet Loss

Ping Failures

Domain Name Resolution Failure (CentOS 6.X System)

Domain Name Resolution Failure (Linux System)

Troubleshooting

Instance-Related Failures

CVM Login Failures

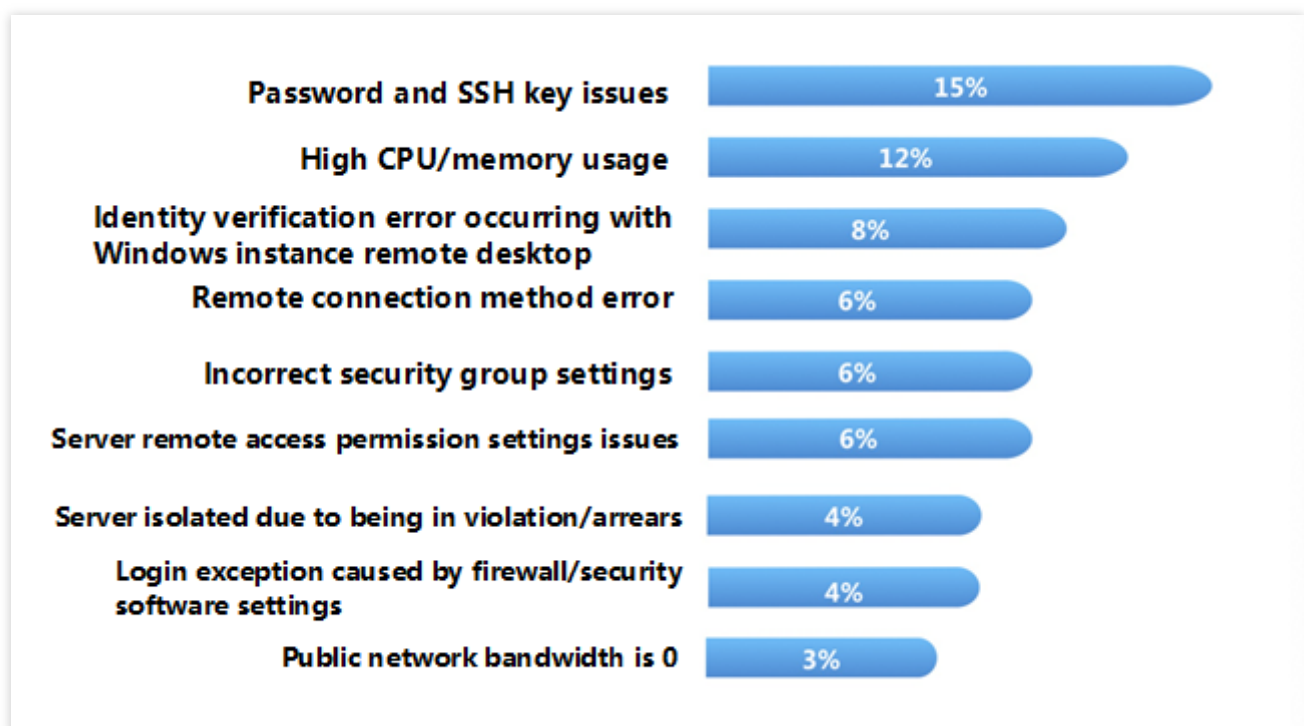
Last updated : 2024-01-06 17:32:18

This document describes how to troubleshoot instance login failures after you purchase Cloud Virtual Machine (CVM) instances, helping you locate and resolve CVM login failures.

This document describes how to determine possible causes of instance login failures after you purchase Cloud Virtual Machine (CVM) instances, helping you locate and resolve CVM login failures.

Possible Causes

The following figure shows the primary causes of CVM instance login failures and their probabilities. If you cannot connect to an instance, we recommended you use the diagnosis tool and perform troubleshooting as instructed below.



Troubleshooting

Confirming the instance type

You must first determine whether your purchased instance is a Windows system instance or Linux system instance. The causes of login failures vary by instance types. According to your purchased instance type, refer to the following documentation to locate and resolve the issue.

[Unable to log into a Windows instance](#)

[Unable to log into a Linux instance](#)

Using the diagnosis tool to locate the causes

Tencent Cloud provides [Port Verification](#) to help you determine possible causes of login failures. More than 70% of login issues can be checked and located by this tool.

Self-Diagnosis Tool

Problems that can be diagnosed include high bandwidth usage rate, zero public network bandwidth, high server workload, improper security group rules, DDoS attack blocking, security isolation, and account in arrears.

Port Verification Tool

This tool can diagnose security group- and port-related problems. If there is a security group configuration issue, you can use **Open All Ports** function of the tool to open all commonly used interfaces of the security group.

If you locate the cause of the issue using the tool, we recommend you follow the corresponding issue guidelines to resolve it.

Restarting Instance

After the diagnosis tool has located and managed the corresponding issue, or it is still not possible to locate the cause of the login failure using the diagnosis tool, you can restart the instance and connect remotely again to see whether the connection succeeds.

For information about how to restart an instance, see [Restart Instance](#).

Other common causes of login failures

If you cannot locate the cause of the issue following the above-mentioned steps, or you receive the following error messages when logging in to the CVM, refer to the following solutions.

Windows Instances

[Windows instance: Unauthorized to log in via remote desktop service](#)

[Windows instance: Mac remote login exception](#)

[Windows instance: Authentication error](#)

[Windows instance: Remote desktop cannot connect to the remote computer](#)

Linux Instances

[Linux instance: Unable to login due to high CPU and memory usage rates](#)

Subsequent Operations

If you still cannot log in remotely following the above-mentioned steps, save the related logs and self-diagnosis results, then [Submit Ticket](#).

Windows Instance Login Failures

Windows Instance Login Failures

Last updated : 2024-01-06 17:32:18

This document describes the possible causes of Windows instance login failures and their troubleshooting methods.

Possible Cause

Common login failure reasons:

- [Incorrect password](#)
- [High bandwidth utilization](#)
- [High server load](#)
- [Improper remote port configuration](#)
- [Improper security group rules](#)
- [Exception caused by the firewall or security software](#)
- [Authentication error in access through remote desktop](#)

Using Self-Diagnosis Tool

Tencent Cloud provides a self-diagnosis tool to help you determine whether the failure is caused by common problems with the bandwidth, firewall, and security group configurations. 70% of faults can be located with this tool. You can locate the faults that may result in the login failure based on the detected causes.

1. Click [Self-diagnose](#) to open the self-diagnosis tool.
2. Select the target CVM instance as prompted and click **Start Detection**.

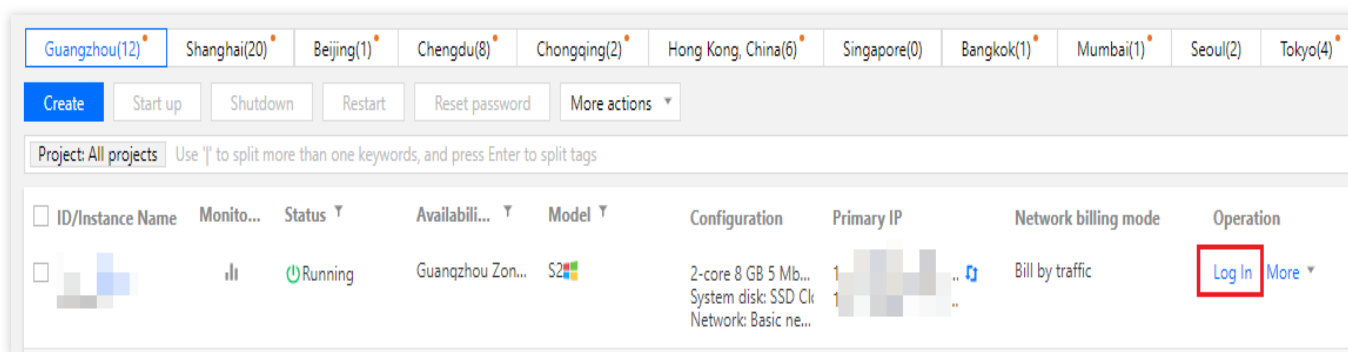
If you cannot troubleshoot with the diagnosis tool, we recommend you [log in to the CVM instance via VNC](#) and follow the instructions.

Troubleshooting

Logging in via VNC

If you cannot log in to a Windows instance through RDP or remote access software, you can log in through VNC for troubleshooting.

1. Log in to the [CVM console](#).
2. On the **Instances** page, select the target instance and click **Log in**.

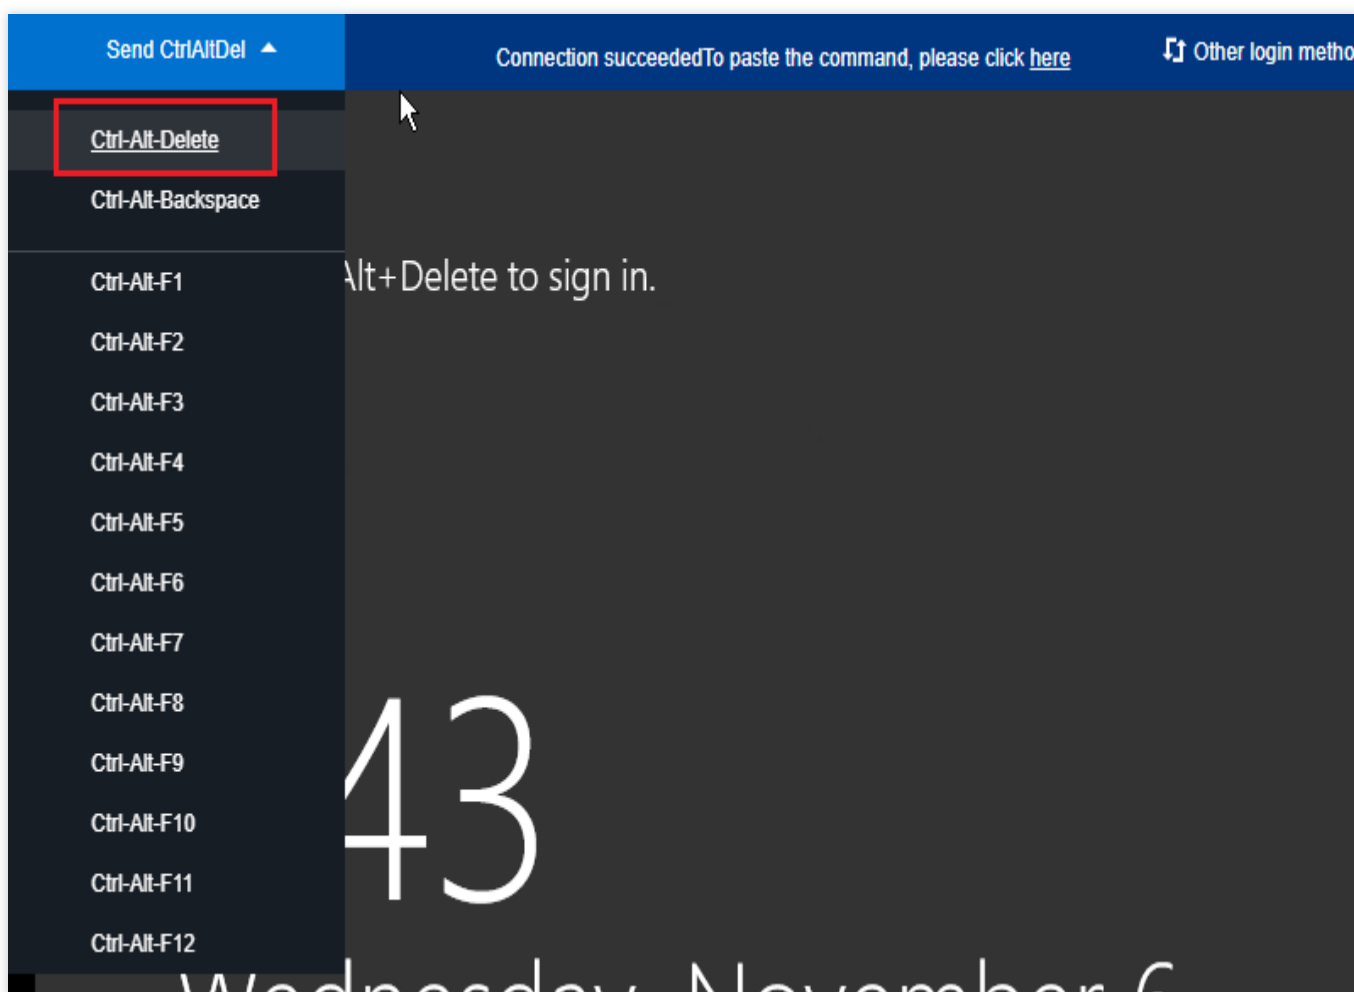


3. In the **Standard Login | Windows Instance** pop-up window, click **Login via VNC**.

Note:

If you forgot the password for the instance, you can reset it in the console. For more information, see [Resetting Instance Password](#).

4. In the login pop-up window, click **Ctrl-Alt-Delete** from the top left list.



Login failure due to password issue

Problem: The login attempt failed because you forgot the password, entered an incorrect password, or failed to reset your password.

Solution: Reset the password for this instance in the [CVM console](#) and restart the instance. For more information, see [Resetting Instance Password](#).

High bandwidth utilization

Problem: The self-diagnosis tool shows that bandwidth utilization is too high.

Procedure:

1. Log in to the instance by using [VNC login](#).
2. Check the bandwidth utilization of the instance and perform troubleshooting accordingly. For details, see [Login Failure Due to High Bandwidth Occupation](#).

High server load

Problem: The self-diagnosis tool or Tencent Cloud Observability Platform shows that server CPU workload is too high, and the system is unable to perform remote connection or access is slow.

Possible cause: Viruses, trojans, third-party antivirus software, application exceptions, driver exceptions, and automatic updates of software on the backend may lead to high CPU utilization.

Procedure:

1. Log in to the instance by using [VNC login](#).
2. In **Task manager**, find the process with high load. For details, see [Failed to log in to a Windows CVM due to high CPU and memory usage](#).

Improper remote port configuration

Problem: Failed to access the instance remotely, the remote access port is not the default port or has been modified, or port 3389 is not open.

Diagnosis: Ping the public IP address of the instance to check network connectivity and run telnet to check whether the port is open.

Procedure: See [Remote Login Failure Due to Port Issues](#) for the detailed procedure.

Improper security group rules

Problems: Security group rule configuration is improper, leading to login failures.

Procedure: Troubleshoot with the [Port Verification](#) feature on the VPC console.

Note:

Open 3389 must be open for remote login.

Testing Details ×

Protocol	Port	Direction	Policy	Effects
TCP	3389	Inbound	Not opened ⓘ	Unable to log into C...
TCP	22	Inbound	Open	None
TCP	443	Inbound	Not opened ⓘ	Unable to use Web ...
TCP	80	Inbound	Not opened ⓘ	Unable to use Web ...
TCP	21	Inbound	Not opened ⓘ	Unable to access FTP
TCP	20	Inbound	Not opened ⓘ	Unable to access FTP
ICMP	0	Inbound	Open	None
ALL	ALL	Outbound	Open	None

Open all ports Cancel

Open all ports Cancel

To define a custom rule for the security group, see [Adding Security Group Rules](#).

Login failure due to firewall or security software

Problem: The login attempt failed due to the CVM firewall or security software.

Diagnosis: Log in to a Windows instance through VNC to check whether the login is blocked by the firewall policies or security software installed on the server.

Note:

This operation involves shutting down the CVM firewall. To perform it, check whether you have the corresponding permission.

Procedure: Shut down the firewall or the installed security software, and then try to access remotely again. For example, you can shut down the firewall of Windows Server 2016 as follows:

1. Log in to the instance by using [VNC login](#).
2. On the desktop, click



and select **Control Panel**.

3. Click **Windows Defender Firewall**.

4. In the **Windows Defender Firewall** window, click **Turn Windows Firewall on or off** on the left to open **Customize Settings**.

5. Set **Private network settings** and **Public network settings** to **Turn off Windows Firewall** and click **OK**.

6. Restart the instance and try to access remotely again.

Identity verification error in access through remote desktop

Problem: When you tried to log in to a Windows instance through the remote desktop, the prompt stating "Authentication error. Invalid flag is provided to the function." or **Authentication error. The required function is not supported.** appears.

Possible cause: Microsoft released a security update in March 2018. This update fixes a remote code execution vulnerability in the Credential Security Supporting Program (CredSSP) by correcting how CredSSP validates requests during the authentication process. Both the client and server need to be updated or the preceding error may occur.

Procedure: Install the security update (recommended). For details, see [An Authentication Error Occurred when You Tried to Log In to a Windows Instance Remotely](#).

Other Solutions

If you still cannot connect to the Windows instance, and [submit a ticket](#) for assistance.

An authentication error occurred when you tried to log in to a Windows instance remotely

Last updated : 2024-01-06 17:32:18

Problem Description

When a Remote Desktop Connection is used to log in to a Windows instance, an error is displayed.

An authentication error has occurred. The token supplied to the function is invalid.

An authentication error has occurred. The function requested is not supported.

Problem Analysis

Microsoft published a security update in March 2018. By correcting how the Credential Security Support Provider protocol (CredSSP) validates requests during authentication, this update fixes the remote code execution vulnerability in the CredSSP. Both the client and server need to install the security update, or the preceding error may occur.

Remote connection fails in the following three scenarios:

Scenario 1: The security update is installed on the server but not on the client, and the "force updated clients" policy is configured.

Scenario 2: The security update is installed on the client but not on the server, and the "force updated clients" policy is configured.

Scenario 3: The security update is installed on the client but not on the server, and the "mitigated" policy is configured.

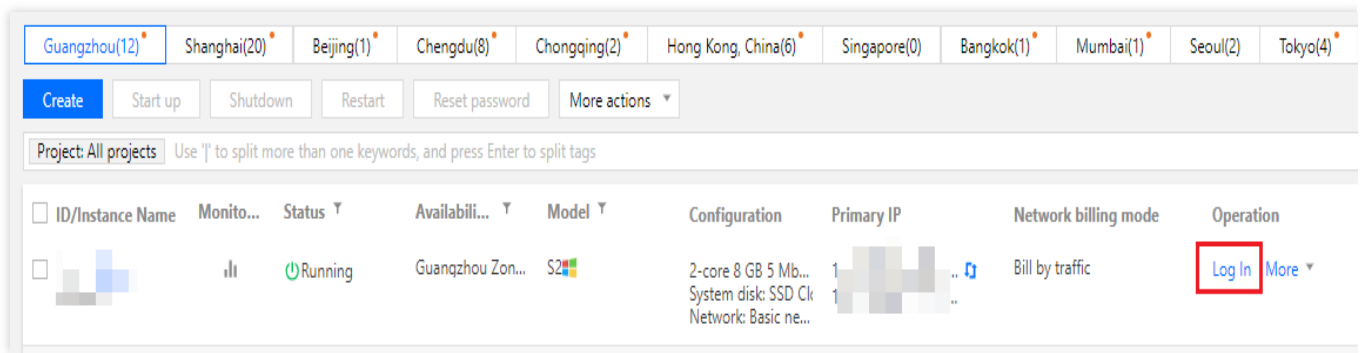
Solution

Note:

If you only update the client locally, use [Solution 1. Install the security update \(recommended\)](#).

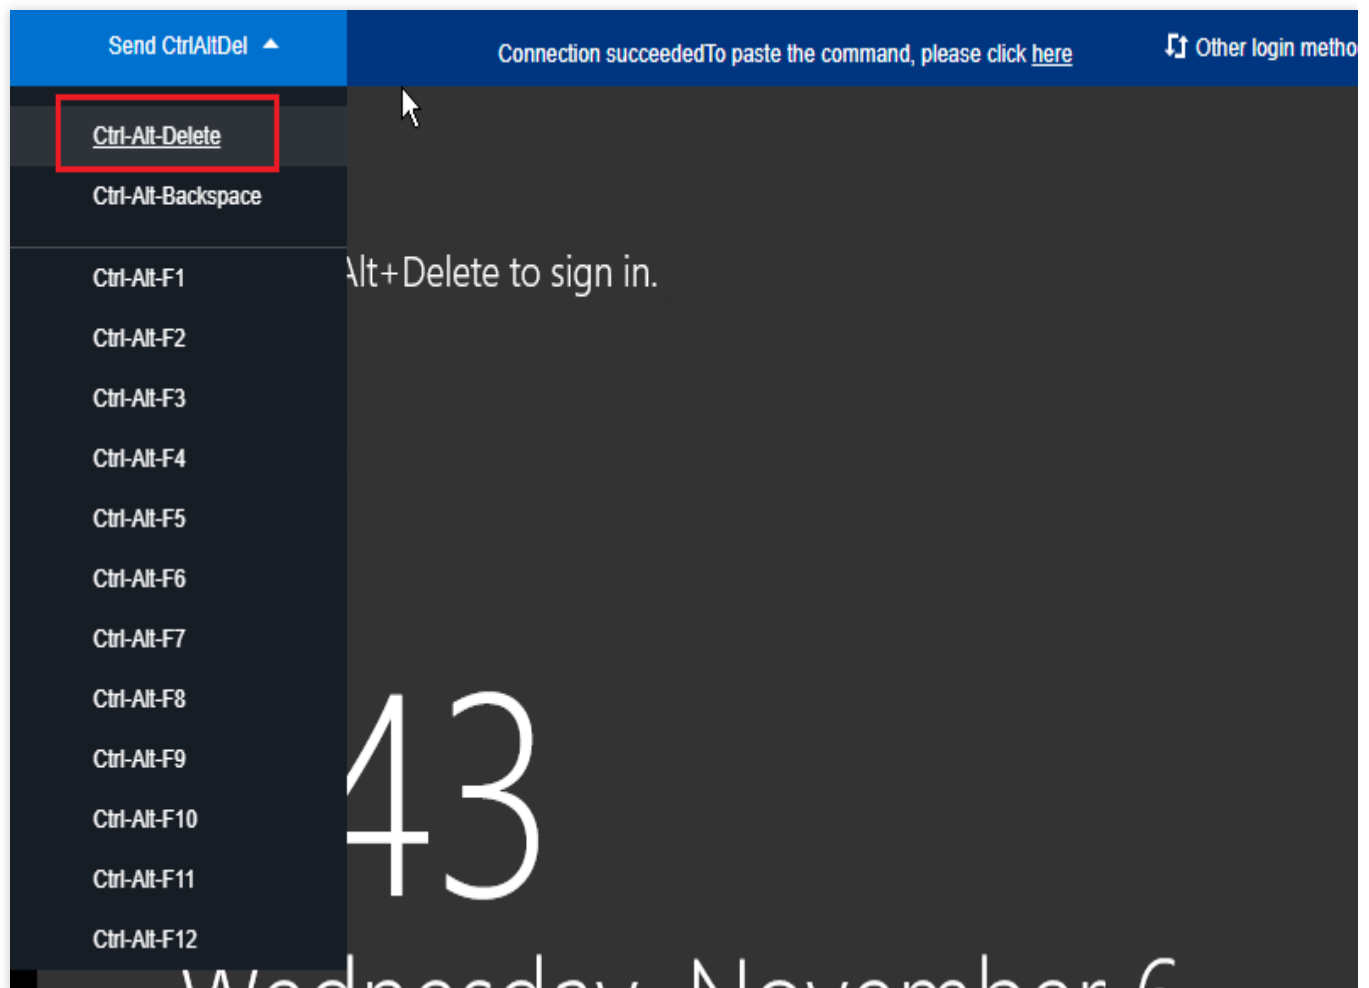
Logging in to CVM via VNC

1. Log in to the [CVM console](#).
2. On the **Instances** page, find the target CVM instance and click **Log in**.



3. In the **Standard Login | Windows Instance** pop-up window, select **Login via VNC**.

4. In the login pop-up window, select **Send remote command** in the top-left corner and press **Ctrl-Alt-Delete** to open the system login window as shown below:



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

Solution 1. Install the security update (recommended)

Install the security update on the unpatched client or server. For updates for different operating systems, see [CVE-2018-0886 | CredSSP remote code execution vulnerability](#). This solution uses Windows Server 2016 as an example.

In other operating systems, you may use the following methods to enter **Windows Update**:

Windows Server 2012:



> **Control Panel** > **System and Security** > **Windows Update**

Windows Server 2008: **Start** > **Control Panel** > **System and Security** > **Windows Update**

Windows 10:



> **Settings** > **Update & Security**

Windows 7:



> **Control Panel** > **System and Security** > **Windows Update**

1. On the desktop, click



and select **Settings**.

2. In the *Settings* pop-up window, select **Update & Security**.

3. In **Update & Security**, select **Windows Update** and click **Check for updates**.

4. Click **Start Installation**.

5. After the installation is complete, restart the instance to finish the update.

Solution 2. Modify the policy

In a CVM instance that has the security update installed, set the **Encryption Oracle Remediation** policy to **Vulnerable**. This solution uses Windows Server 2016 as an example. Follow the steps below:

Note:

If no group policy editor is available in the Windows 10 Home operating system, you can modify the registry to edit the policy as instructed in [Solution 3. Modify the registry](#).

1. On the desktop, click



, enter "gpedit.msc", and press **Enter** to open **Local Group Policy Editor**.

Note:

You can also press **Win+R** to open the **Run** window.

2. On the left sidebar, select **Computer Configuration** > **Administrative Templates** > **System** > **Credentials Delegation** and double-click **Encryption Oracle Remediation**.

3. In the **Encryption Oracle Remediation** pop-up window, select **Enabled** and set **Protection level** to **Vulnerable**.

4. Click **OK**.

Solution 3. Modify the registry

1. On the desktop, click



, enter "regedit", and press **Enter** to open the Registry Editor.

Note:

You can also press **Win+R** to open the **Run** window.

2. On the left sidebar, select **Computer > HKEY_LOCAL_MACHINE > SOFTWARE > Microsoft > Windows > CurrentVersion > Policies > System > CredSSP > Parameters**.

Note:

If the directory path does not exist, create one manually.

3. Right-click **Parameters**, select **New > DWORD (32-bit) value**, and name the file **AllowEncryptionOracle**.

4. Double-click the newly created "AllowEncryptionOracle" file, set **Value data** to "2", and click **OK**.

5. Restart the instance.

References

[CVE-2018-0886 | CredSSP remote code execution vulnerability](#)

[CredSSP updates for CVE-2018-0886](#)

Failed to Reset the CVM Password or the CVM Password Is Invalid

Last updated : 2024-01-06 17:32:18

This document describe how to troubleshoot failed or invalid password resets for a Windows Server 2012 CVM.

Problem

After the CVM password is reset, the system prompts **The system is busy, and your instance password failed to be reset (7617d94c)**.

After the CVM password is reset, the new password does not take effect, and the login password remains the old one.

Cause

Possible causes are:

The `cloudbase-init` component in the CVM is damaged, modified, disabled, or not started.

The `cloudbase-init` component for password reset is blocked by the third-party security program (e.g., 360 Total Security or Huorong Security) installed on the CVM.

Troubleshooting

Try the following according to the failure reason.

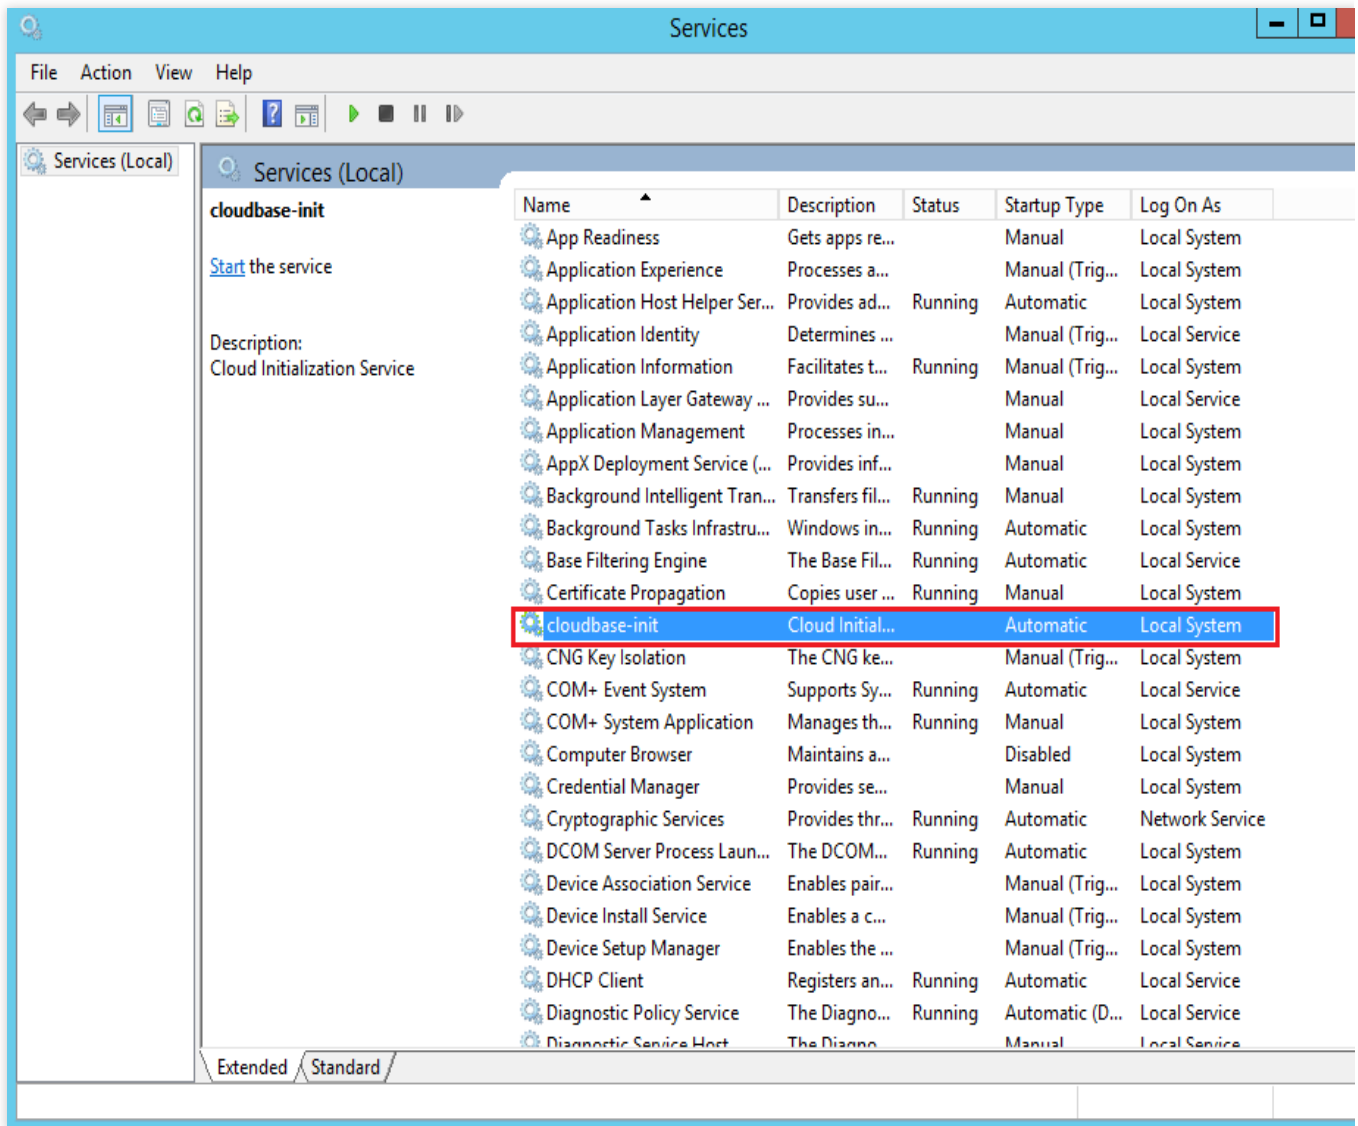
Checking the cloudbase-init service

1. [Logging in to Windows Instance \(WebRDP\)](#).
2. On the desktop, right-click



and choose **Run**. Enter **services.msc** in the **Run** dialog box, and press **Enter** to open the **Services** window.

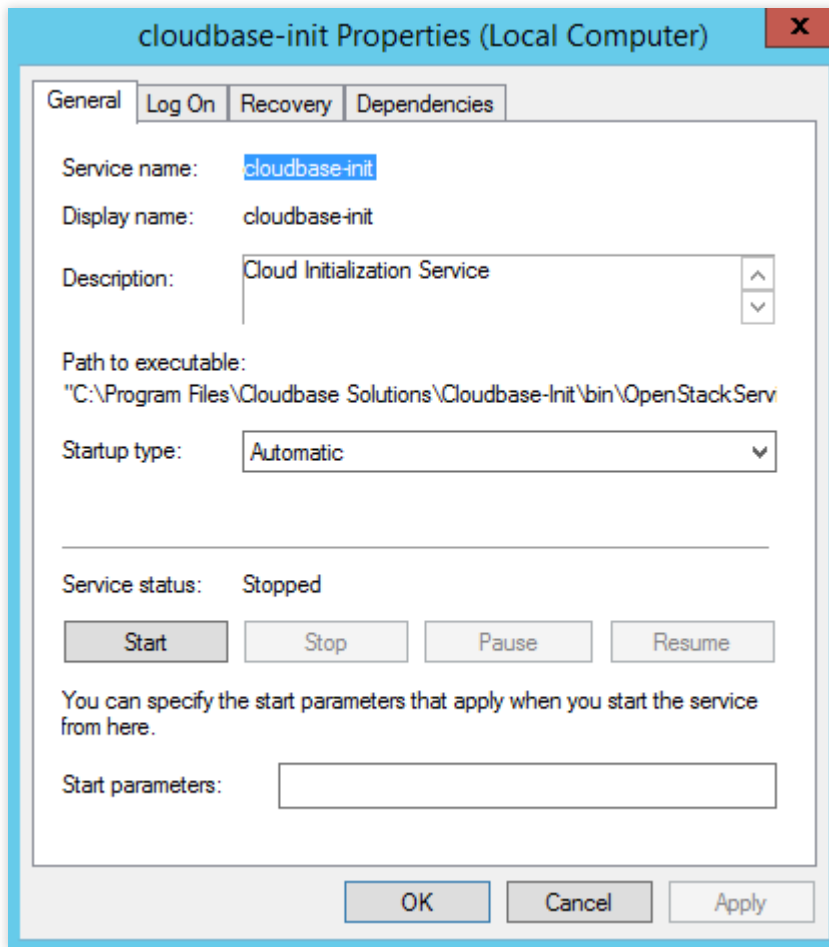
3. Check whether the `cloudbase-init` service exists, as shown in the following figure:



If yes, proceed to the next step.

If no, reinstall the `cloudbase-init` service. For more information, see [Installing Cloudbase-Init on Windows](#).

4. Double-click the `cloudbase-init` service to open the cloudbase-init properties dialog box, as shown in the following figure:



5. Select the **General** tab and check whether the `cloudbase-init` startup type is **Automatic**.

If yes, proceed to the next step.

If no, set the `cloudbase-init` startup type to **Automatic**.

6. Switch to the **Log On** tab and check whether **Local System account** is selected for the `cloudbase-init` service.

If yes, proceed to the next step.

If no, select **Local System account** for the `cloudbase-init` service.

7. Switch to the **General** tab, click **Start** in **Service status** to manually enable the `cloudbase-init` service, and check whether an error occurs.

If yes, [check the security program installed on the CVM](#).

If no, proceed to the next step.

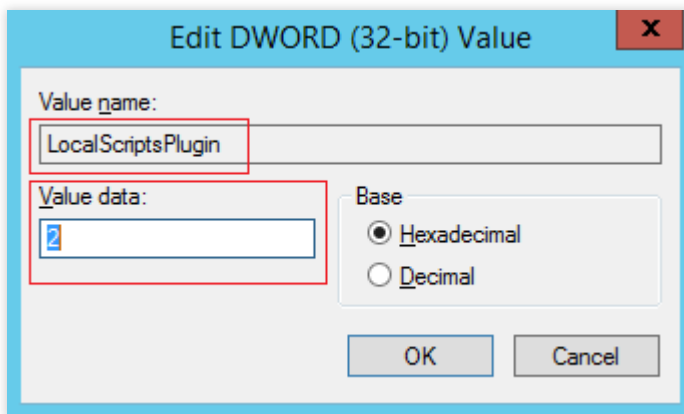
8. On the desktop, right-click



and choose **Run**. Enter **regedit** in the **Run** dialog box, and press **Enter** to open the **Registry Editor** window.

9. In the registry navigation pane on the left, expand the following hierarchies in order: **HKEY_LOCAL_MACHINE > SOFTWARE > Cloudbase Solutions > Cloudbase-Init**.

10. Locate all "LocalScriptsPlugin" registry keys under **ins-xxx** and check whether the LocalScriptsPlugin value is 2.



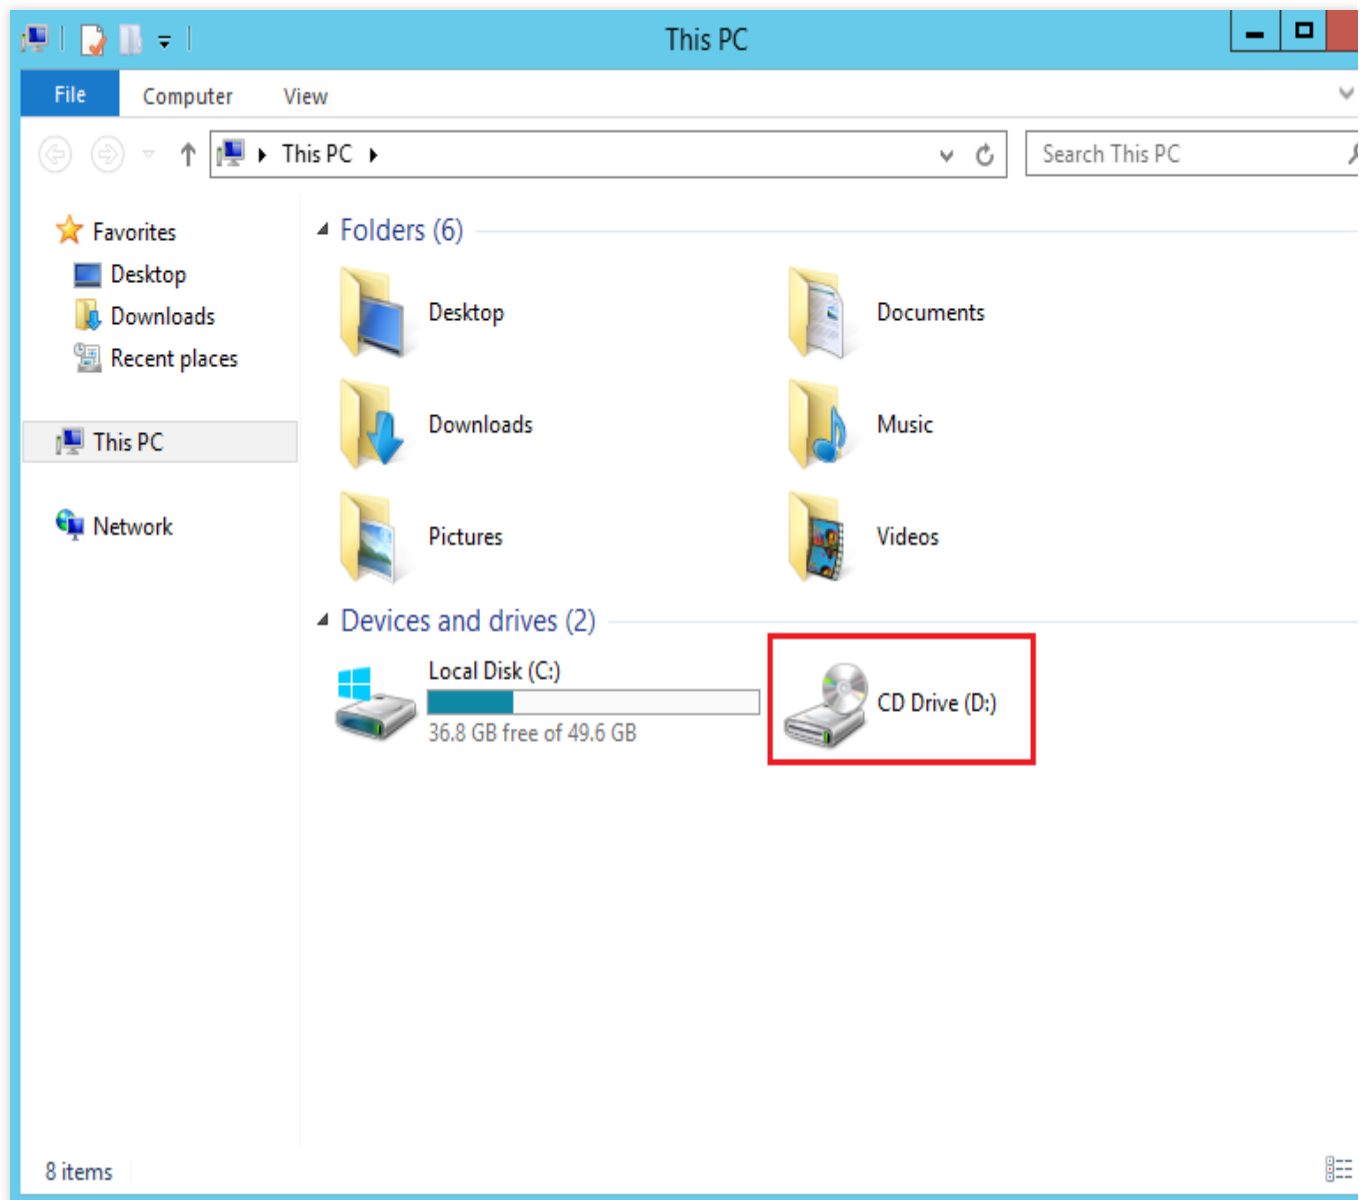
If yes, proceed to the next step.

If no, set the LocalScriptsPlugin value to 2.

11. On the desktop, click



and choose **This PC**. Check whether the CD drive is loaded under **Devices and drives**, as shown in the following figure:



If yes, [check the security program installed on the CVM](#).

If no, start the CD-ROM drive in Device Manager.

Checking the security program installed on the CVM

Scan for CVM vulnerabilities using the installed security program and check whether `cloudbase-init` components are blocked.

If the CVM has vulnerabilities, fix them.

If core components are blocked, unblock them.

Check and configure the `cloudbase-init` components as instructed below.

1. [Logging in to Windows Instance \(WebRDP\)](#).

2. Restore and set the `cloudbase-init` components according to the actually installed third-party security program.

Connection to a Windows CVM through Remote Desktop was denied

Last updated : 2024-01-06 17:32:18

Error Description

Case 1

: When trying to connect to a Windows instance via Remote Desktop from Windows, the user sees an error that says **The connection was denied because the user account is not authorized for remote login.**

Case 2

: When trying to connect to a Windows instance via Windows Remote Desktop, the user sees an error that says **To sign in remotely, you need the right to sign in through Remote Desktop Services. By default, members of the Remote Desktop Users group have this right. If the group you're in doesn't have the right, or if the right has been removed from the Remote Desktop Users group, you need to be granted the right manually.**

Possible Reasons

The user is not allowed to log in to the Windows instance via Remote Desktop connections.

Solution

For [Case 1](#), add the user account to the list of accounts that are permitted by the Windows instance to log in through Remote Desktop Services. For detailed directions, see [Allowing remote login](#).

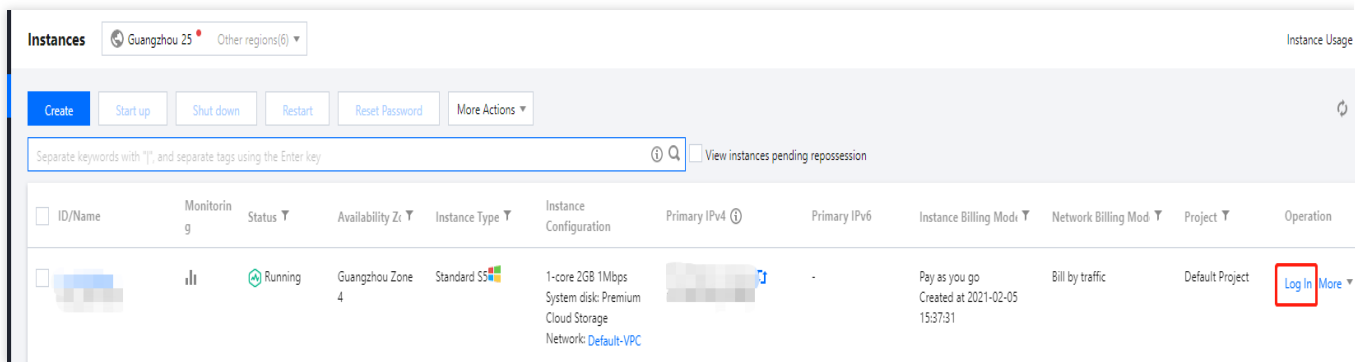
For [Case 2](#), remove the user account from the list of accounts that are denied by the Windows instance to log in through Remote Desktop Services. For detailed directions, see [Denying remote login](#).

Directions

Logging in to the CVM using VNC

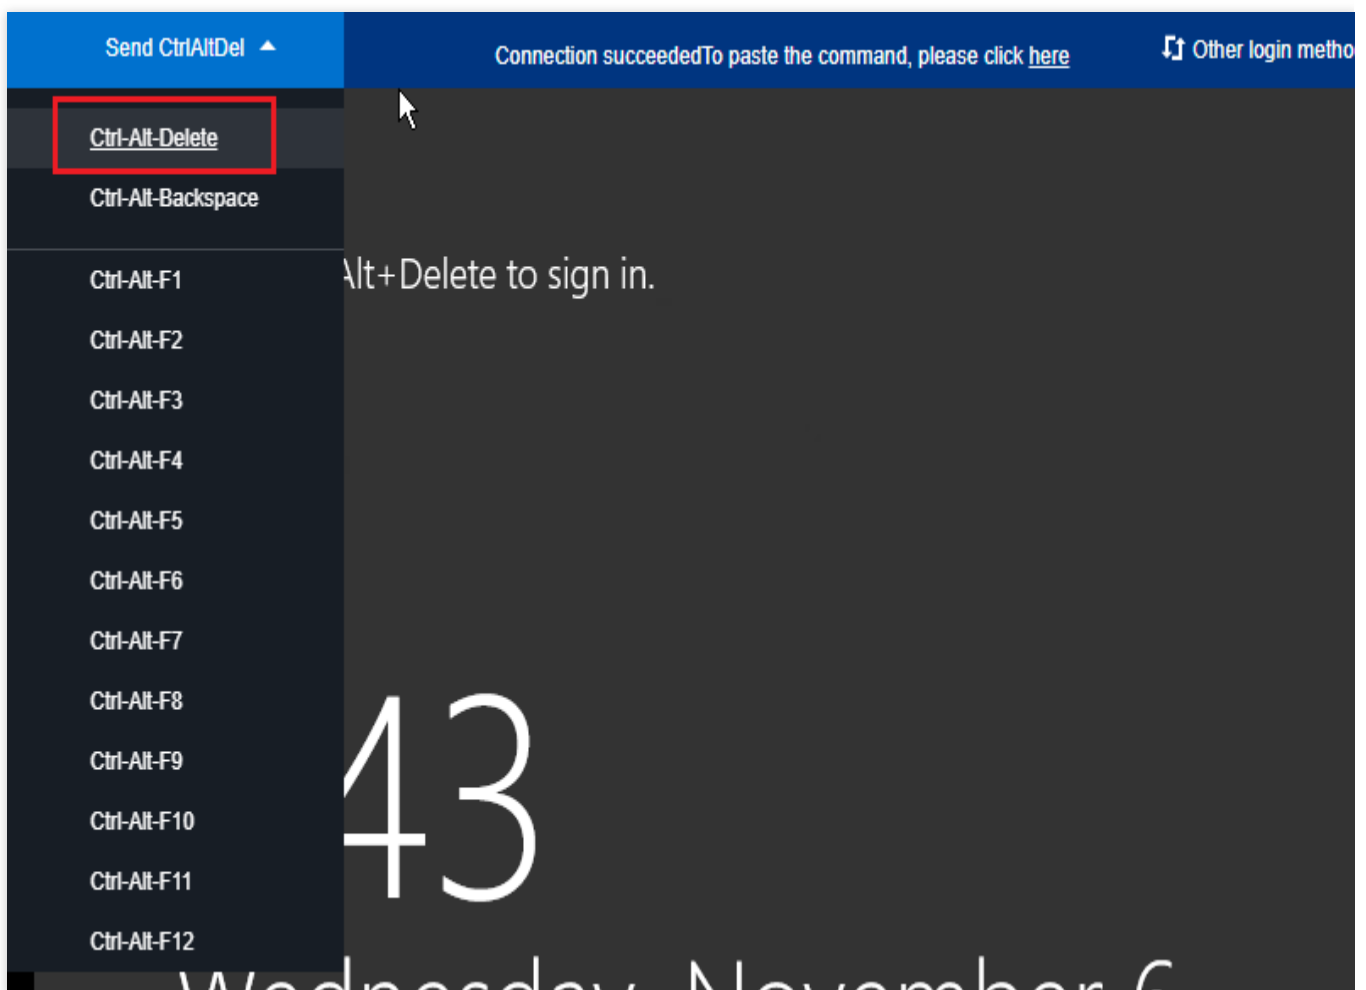
1. Log in to the [CVM console](#).

2. On the instance management page, locate the target CVM instance and click **Log In**, as shown in the following figure:



3. In the **Log in to Windows Instance** window that appears, select “Alternative login methods (VNC)”, click **Log In Now** to log in to the CVM.

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



Allowing remote login

Note:

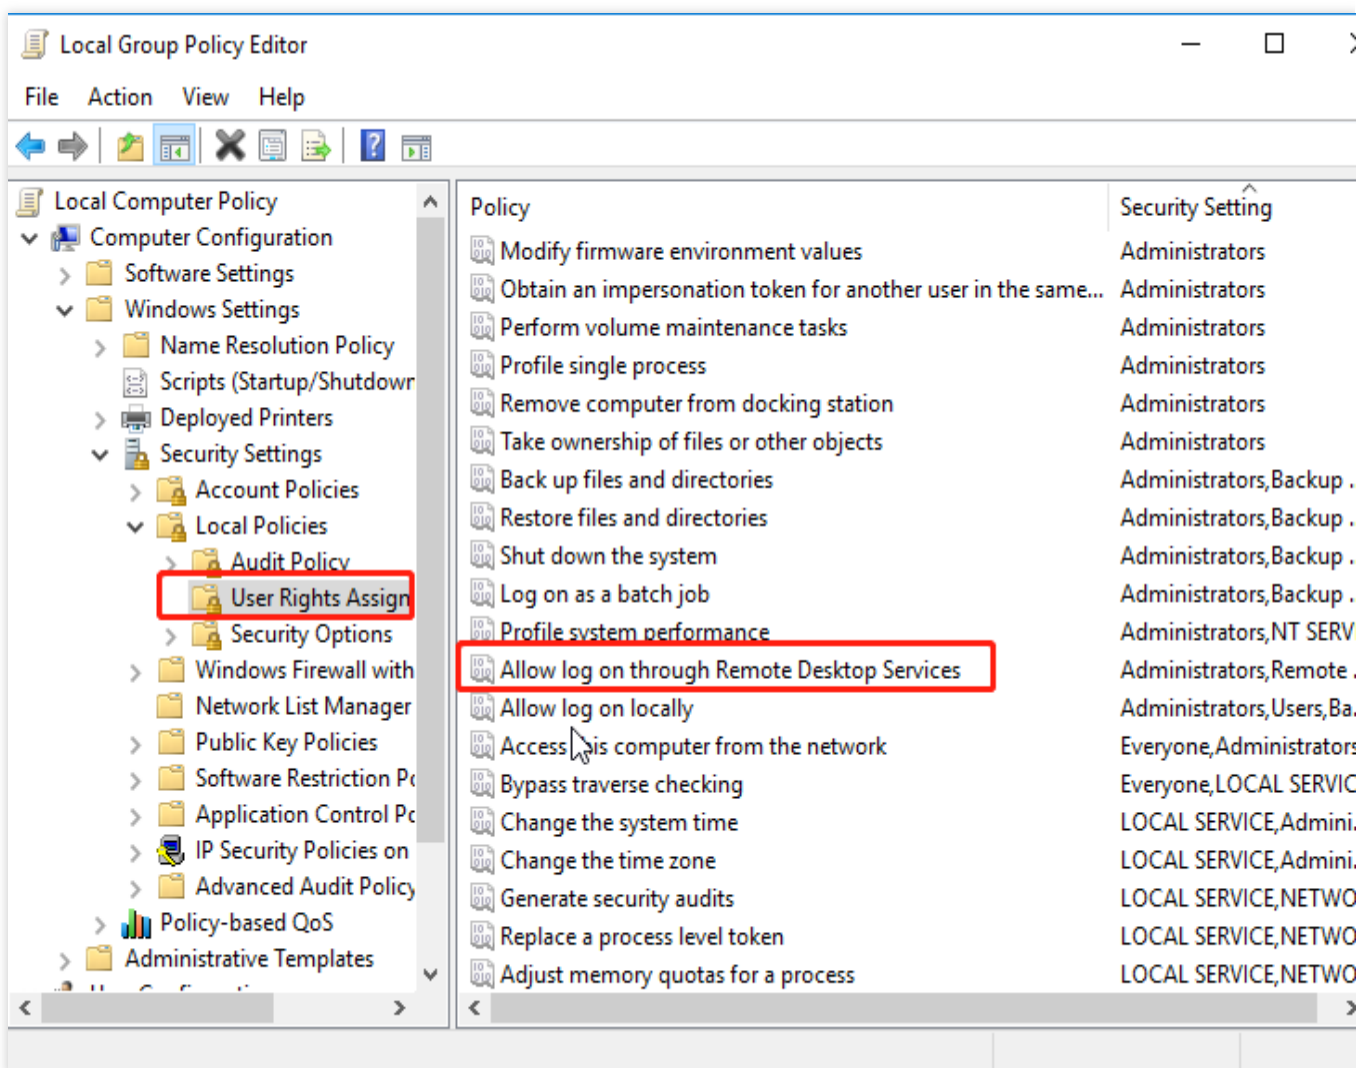
The following operations take Windows Server 2016 as an example.

1. On the desktop, click

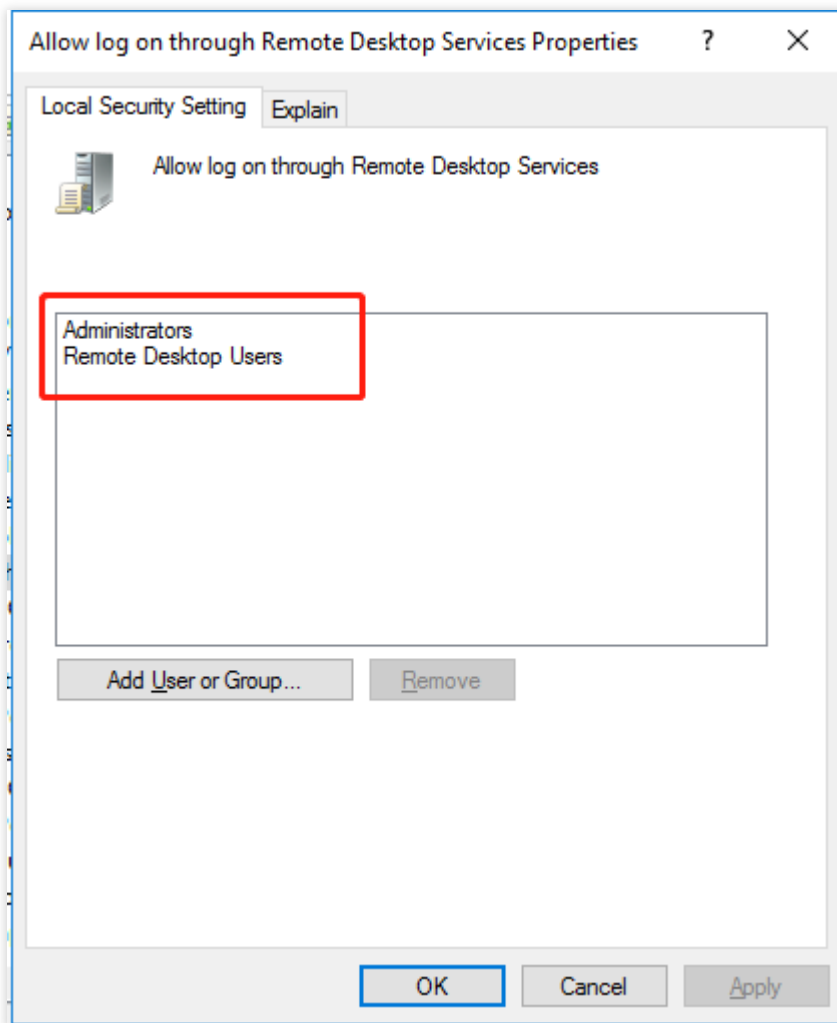


, enter **gpedit.msc**, and press **Enter** to open “Local Group Policy Editor”.

2. In the left navigation tree, choose **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment**, and right-click **Allow log on through Remote Desktop Services**.



3. In the “Allow log on through Remote Desktop Services Properties” window that appears, check whether the user account you want to use for remote login is in the user list of “Allow log on through Remote Desktop Services”.



If the user is not in the list, go to [step 4](#).

If the user is in the list, please [submit a ticket](#).

4.

Click **Add User or Group** to go to the **Select User or Group** window.

5. Enter the account you want to use for remote login and click **OK**.

6. Click **OK** and close “Local Group Policy Editor”.

7. Restart the instance and try to connect to the Windows instance with the account through Remote Desktop again.

Denying remote login

Note:

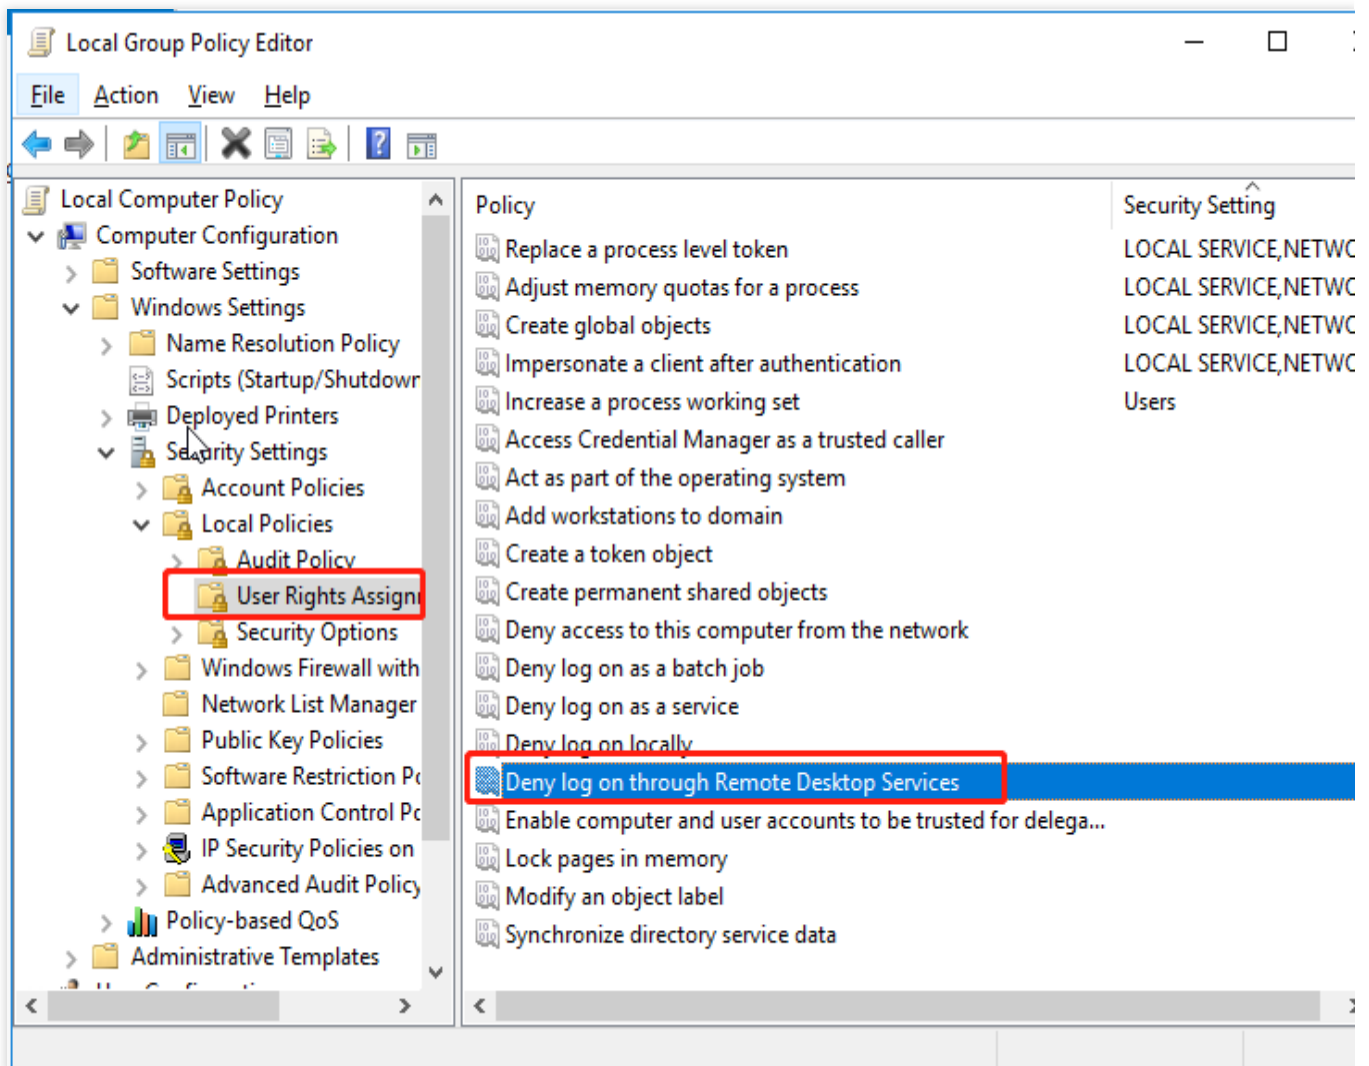
The following operations take Windows Server 2016 as an example.

1. On the desktop, click



, enter **gpedit.msc**, and press **Enter** to open “Local Group Policy Editor”.

2. In the left navigation tree, choose **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment**, and double-click **Deny log on through Remote Desktop Services** as shown below:



3. In the pop-up window, check whether the account you want to use for remote login is in the user list of "Deny log on through Remote Desktop Services".

If the user is in the list, remove the user account from the list and restart the instance.

If the user is not in the list, please [submit a ticket](#).

Requires network-level identity verification

Last updated : 2024-01-06 17:32:18

This document describes how to solve the issue of **network level authentication** when connecting to a Windows instance using Remote Desktop.

Issue

Windows Remote Desktop fails to connect to your Windows instance with the error message **The remote computer requires Network Level Authentication, which your computer does not support. For assistance, contact your system administrator or technical support.**



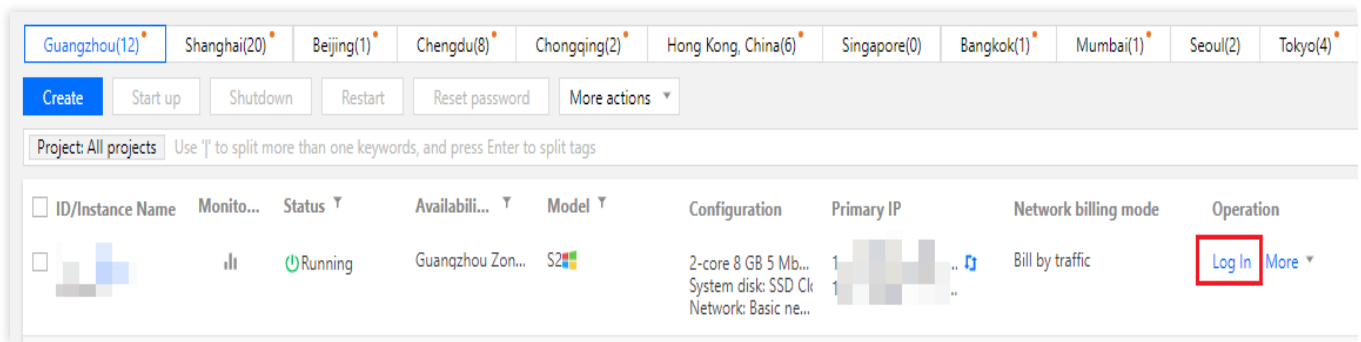
Troubleshooting

Note:

In the following steps, we use Windows Server 2016 as an example.


Logging in to the CVM using VNC

1. Log in to the [CVM Console](#).
2. On the instance management page, find the desired CVM instance. Click **Log In**, as shown in the following figure:



3. In the **Log in to Windows instance** window that appears, select **Alternative login methods (VNC)**. Click **Log In Now** to log in to the CVM.

4. In the login window that appears, select **Send Remote Command** in the top-left corner. Click **Ctrl-Alt-Delete** to enter the system login interface as shown below:




Modifying the Windows registry

1. In the operating system interface, click




. Enter **regedit** and press **Enter** to open the Registry Editor.

2. Using navigation tree on the left-side, navigate to **Computer > HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control > Lsa**. In the right-side pane, select **Security Packages**, as shown in the following figure:

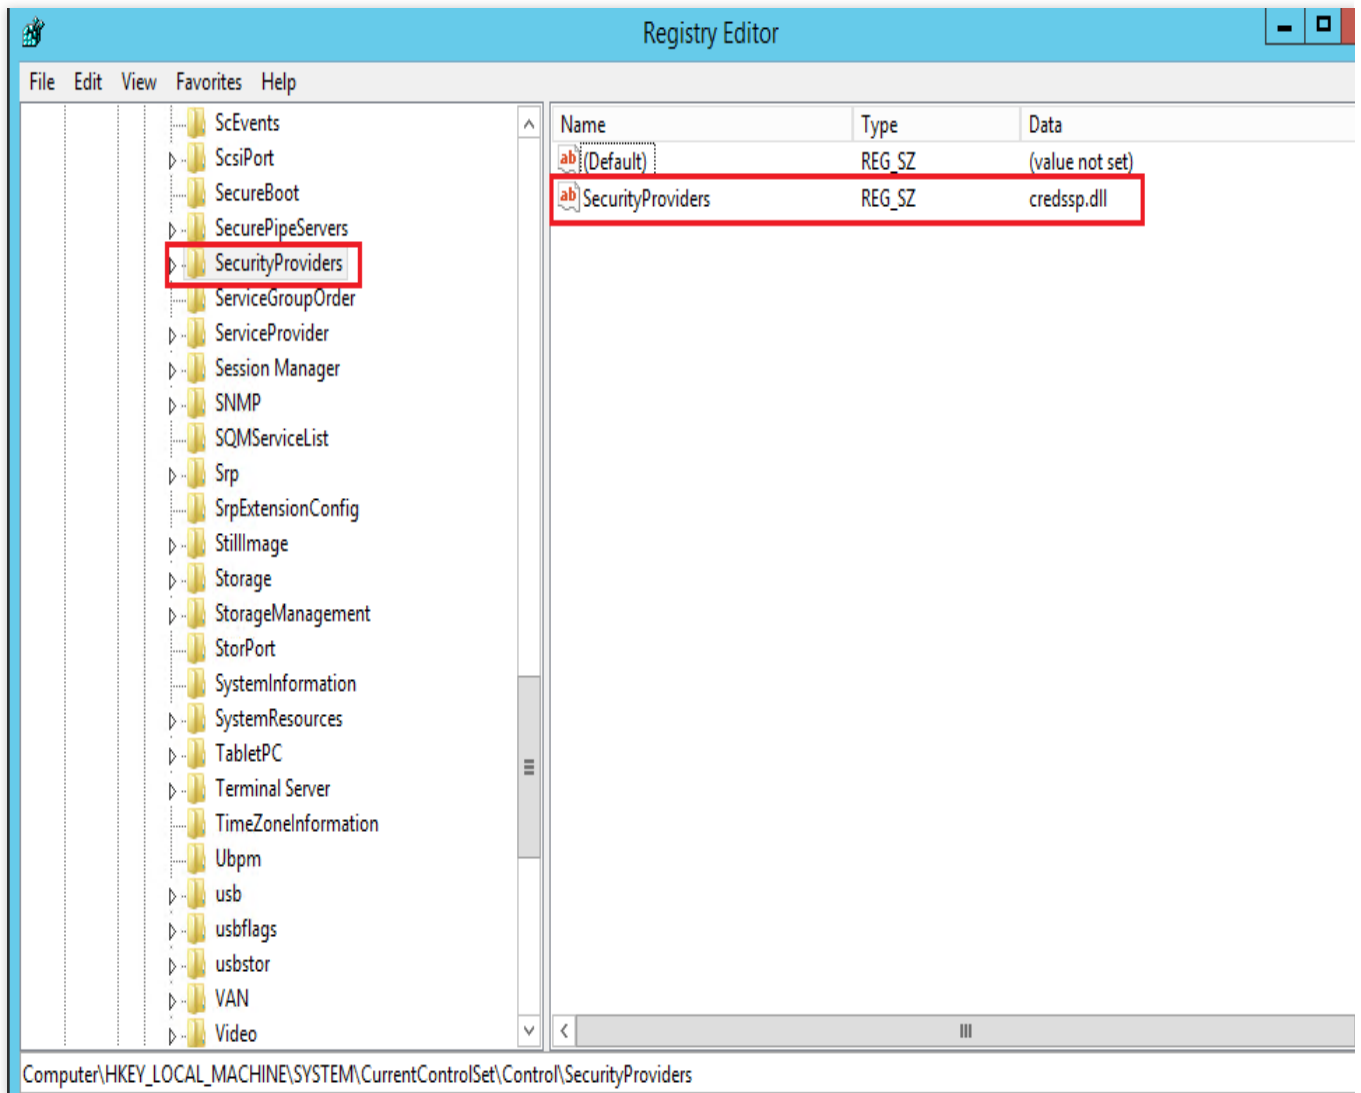


3. Double click **Security Packages** to open the **Edit Multi-String** dialog box.

4. In the **Edit Multi-String** dialog box, add **tspkg** under **Value Data** and click **OK**, as shown in the following figure.

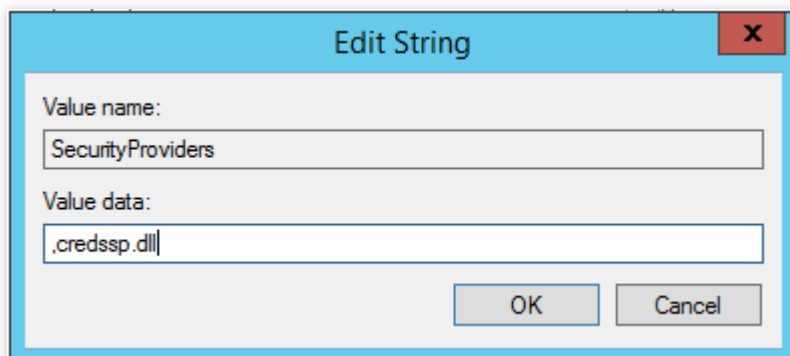


5. Using navigation tree on the left-side, navigate to **Computer > HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Control > SecurityProviders**. In the right-side pane, select **SecurityProviders**, as shown in the following figure:



6. Double-click **SecurityProviders** to open the **Edit Multi-String** dialog box.

7. Append `, credssp.dll` to the end of the **Value Data** field in the **Edit Multi-String** dialog box. Click **OK** as shown in the following figure:



8. Close the registry editor and restart the instance. You can now log in remotely.

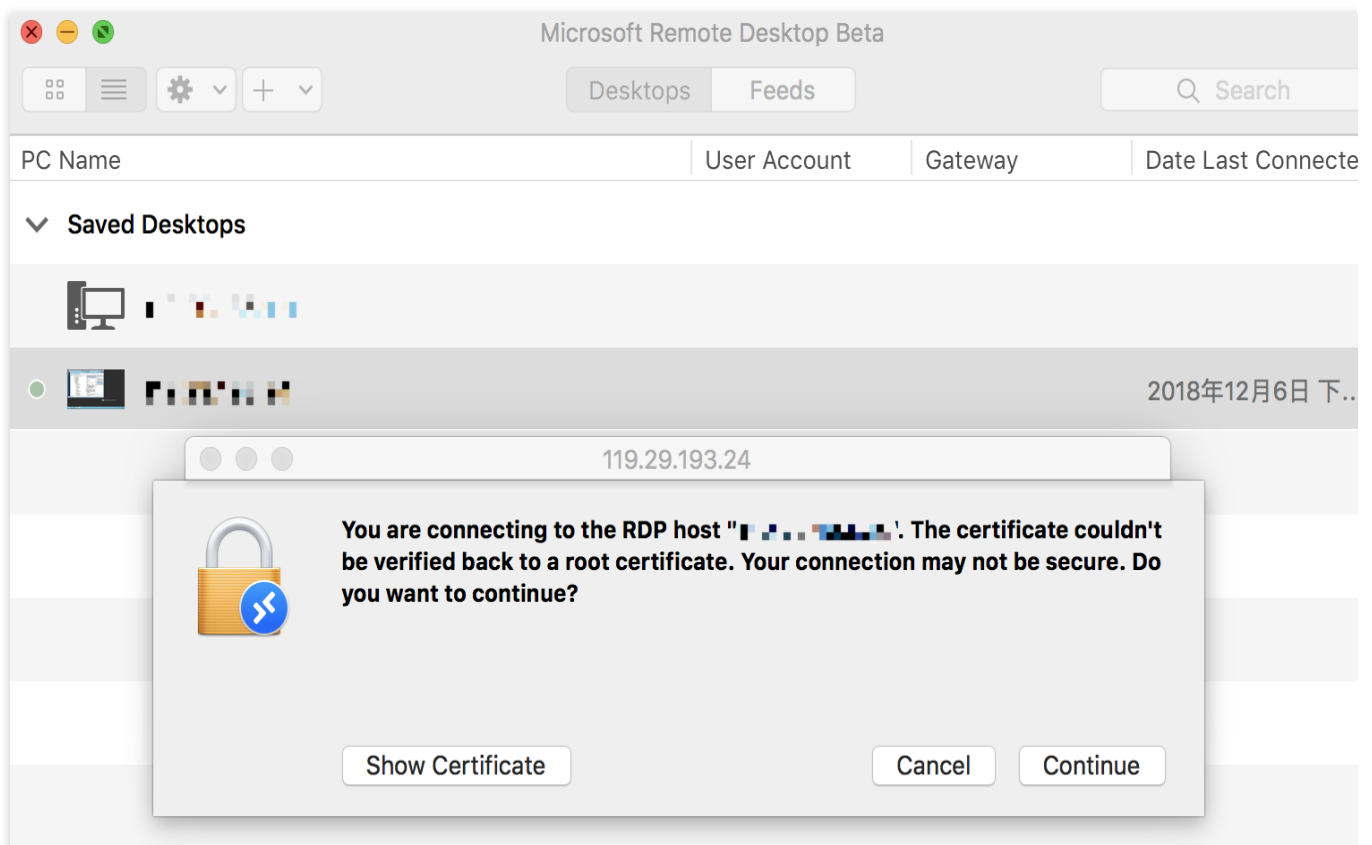
Problems occurred when you tried to log in to a Windows CVM remotely on Mac

Last updated : 2024-01-06 17:32:18

This document describes common problems you may encounter when logging in to Windows CVM on Mac through Microsoft Remote Desktop and how to solve them.

Problems

When logging in to Windows CVM through Microsoft Remote Desktop, you get a **The certificate couldn't be verified back to a root certificate** prompt.



When using Remote Desktop Connection on Mac, you get a **Remote Desktop Connection cannot verify the identity of the computer that you want to connect to** prompt.

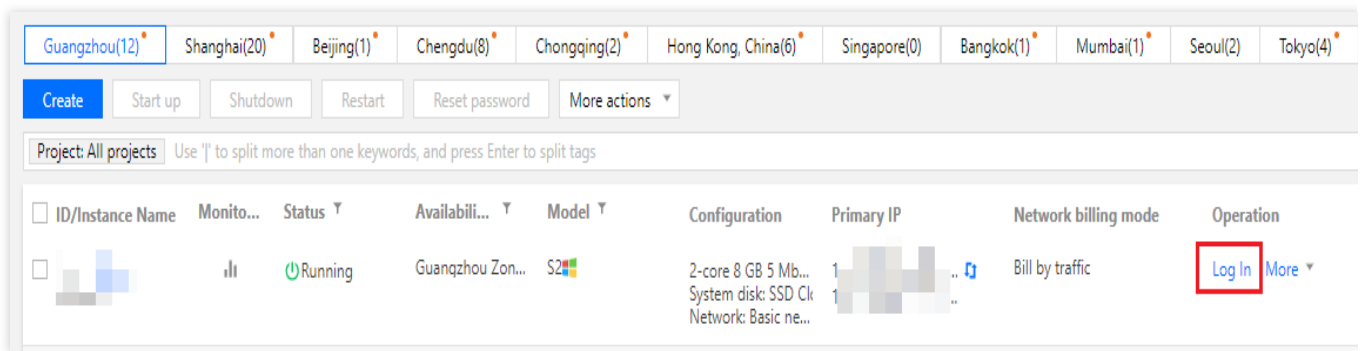
Troubleshooting

Note:

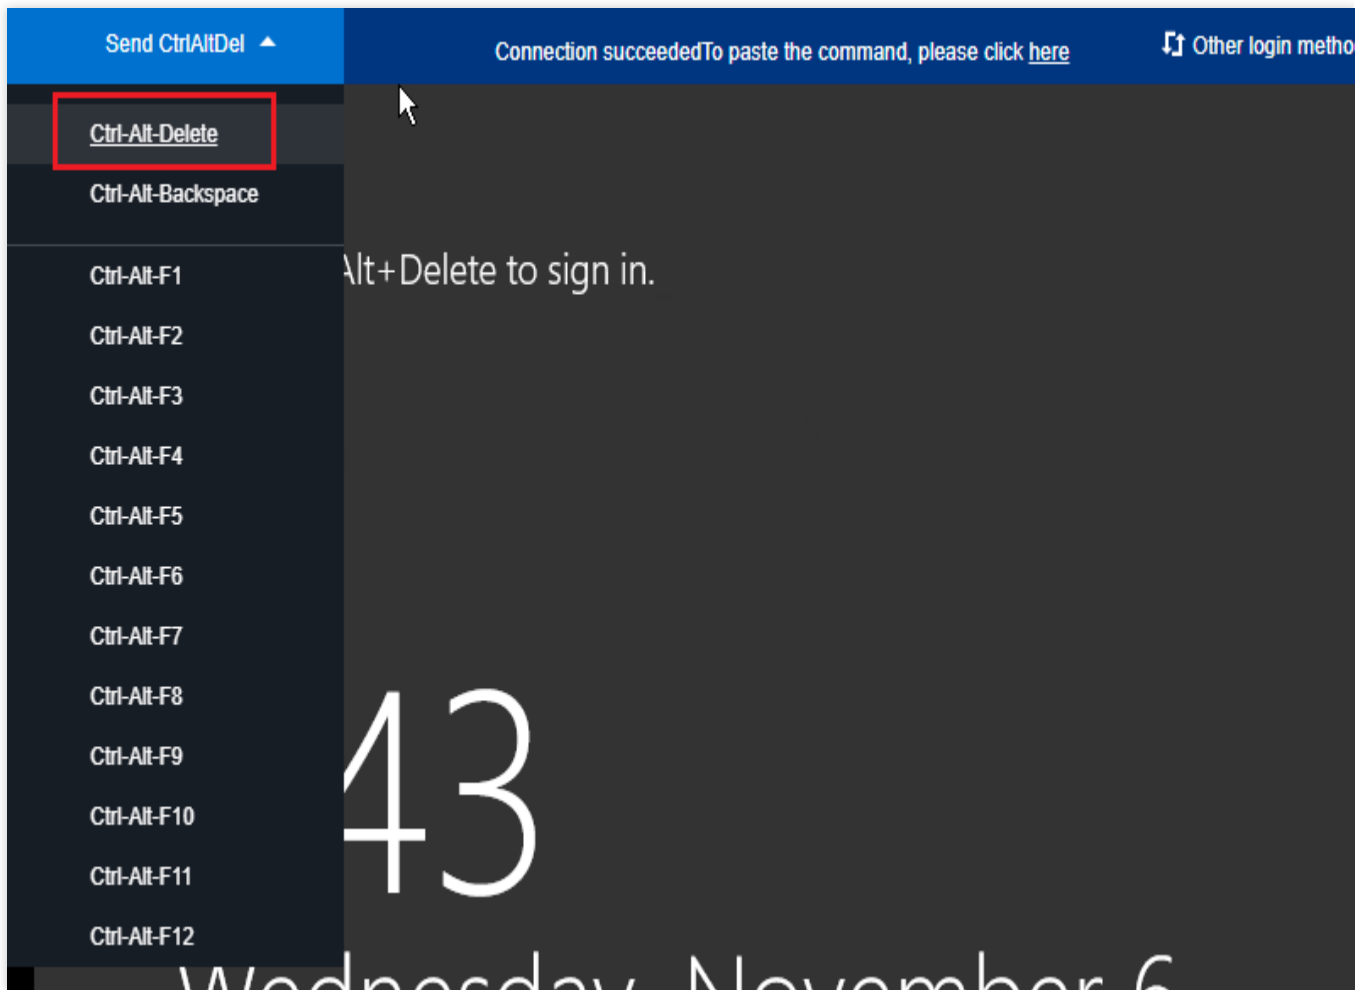
The following operations take Windows Server 2016 as an example.

Logging in to the CVM using VNC

1. Log in to the [CVM Console](#).
2. In the instance management page, locate the CVM you need, and click **Log In**. This is shown in the following figure:



3. In the **Log into Windows instance** window that pops up, select **Alternative login methods (VNC)**, and click **Log In Now** to log in to the CVM.
4. In the login window that pops up, select **Send CtrlAltDel** in the top left corner, and click **Ctrl-Alt-Delete** to enter the system login interface as shown below:



Modifying the local group policy of the instance

1. In the operating system interface, click



, enter **gpedit.msc**, and press **Enter** to open the Local Group Policy Editor.

Note:

You can also use the shortcut **Win+R** to open the Run interface.

2. In the left navigation tree, select **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Security**, double-click **Require use of specific security layer for remote (RDP) connections**.

3. In the “Require use of specific security layer for remote (RDP) connections” window, select **Enabled**, and set the **Security Layer** to **RDP**.

4. Click **OK** to complete the configuration.

5. Restart the instance and try to connect again.

If the connection fails again, please [submit a ticket](#).

Failed to log in to a Windows CVM due to high CPU and memory usage

Last updated : 2024-01-06 17:32:18

This document describes how to troubleshoot Windows CVM login failures due to high CPU or memory utilization.

Note:

The following uses Windows Server 2012 R2 as an example. The steps may vary by operating system (OS) versions.

Possible Causes

Hardware, system processes, service processes, trojans, and viruses may cause high CPU or memory utilization, resulting in slow service response speed or CVM login failure. You can use [Cloud Monitor](#) to create an alarm threshold for CPU or memory usage. You will be notified promptly when the configured threshold is exceeded.

Troubleshooting

1. Identify the process that causes high CPU or memory utilization.
2. Analyze the process.

If it is an unhealthy process, it may be caused by a virus or trojan. Terminate the process or use an antivirus application to scan the system.

If it is a service process, check whether the high CPU or memory utilization is caused by access traffic and whether it can be optimized.

If it is a Tencent Cloud component process, [submit a ticket](#) for assistance.

Tools

Task Manager: an application and process manager included with the Microsoft Windows OS. It provides information on computer performance and running software, such as the names of running processes, CPU load, memory usage, I/O, logged-in users, and Windows services.

Processes: a list of all running processes.

Performance: system performance statistics such as the overall CPU usage and current memory usage.

Users: all users with sessions.

Details: an enhanced version of the processes tab, including detailed information on PID, status, CPU usage, and memory usage.

Services: a list of all services, including those that are not running.

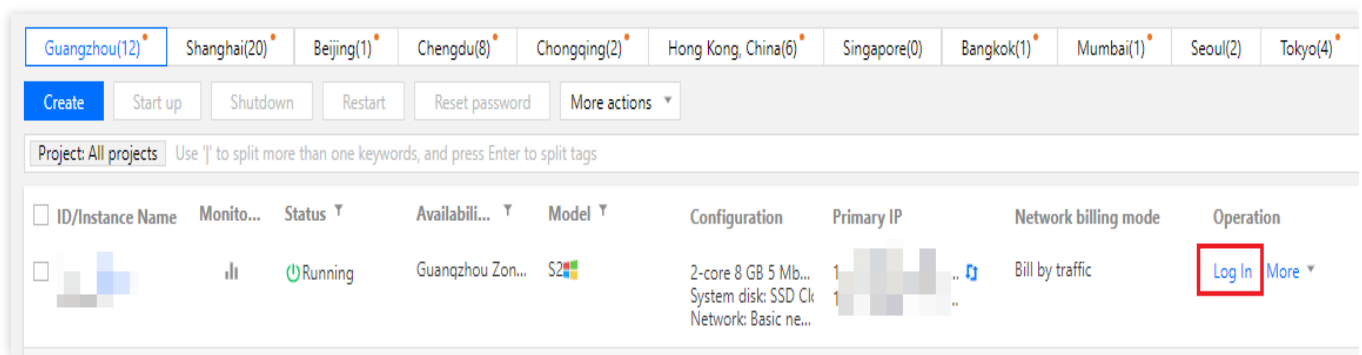
Troubleshooting Method

Logging in to the CVM instance using VNC

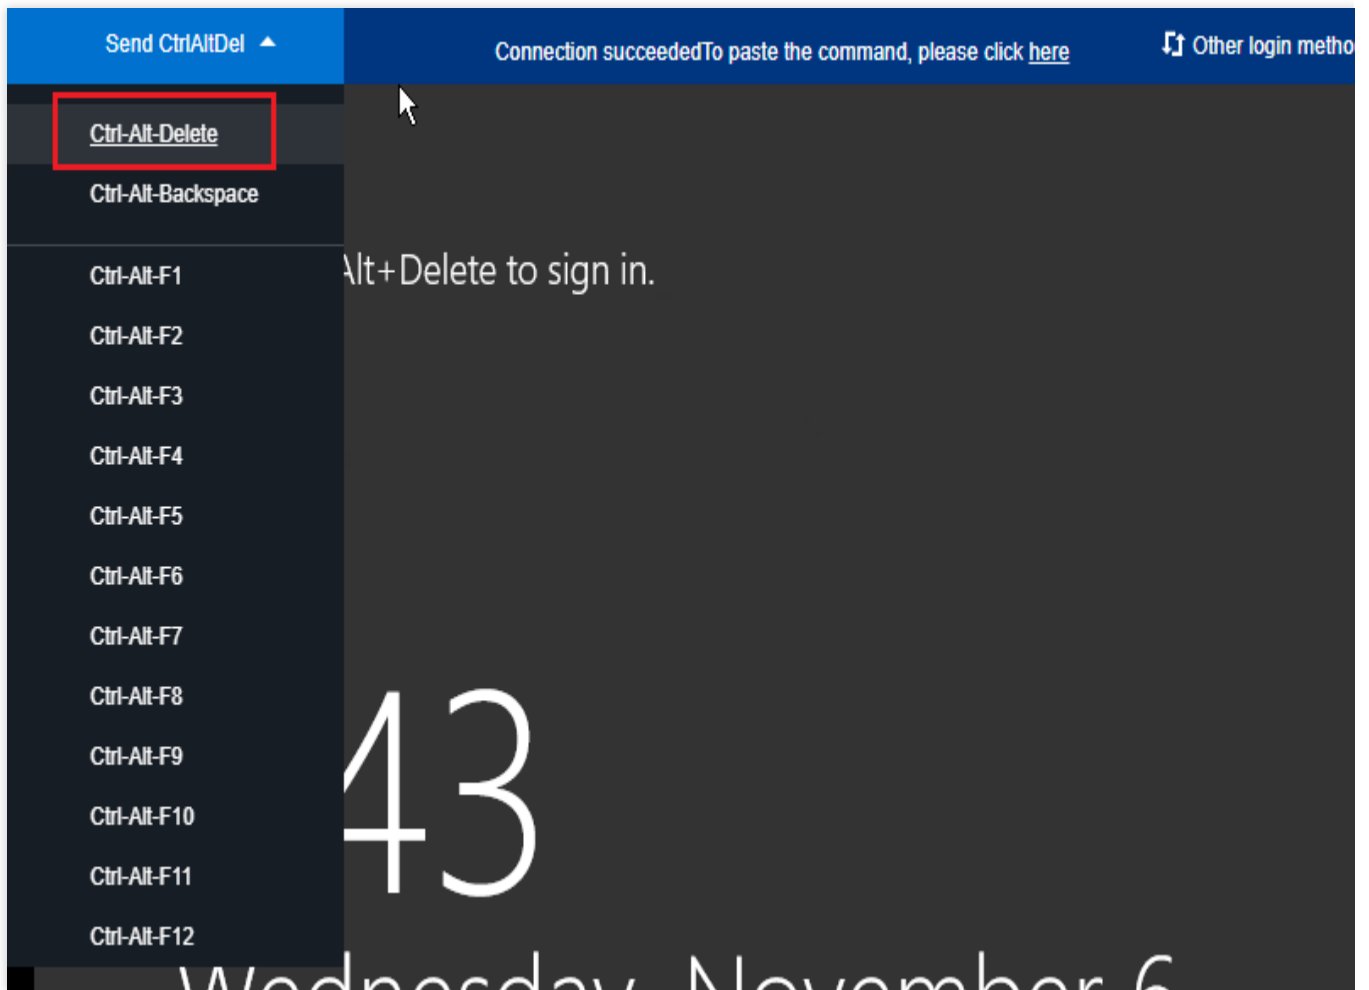
Note:

If you cannot log in to your CVM instance due to high CPU or memory utilization, we recommend [logging into Windows instance via VNC](#).

1. Log in to the [CVM console](#).
2. On the instance management page, locate the CVM instance and click **Log In**, as shown in the following figure:

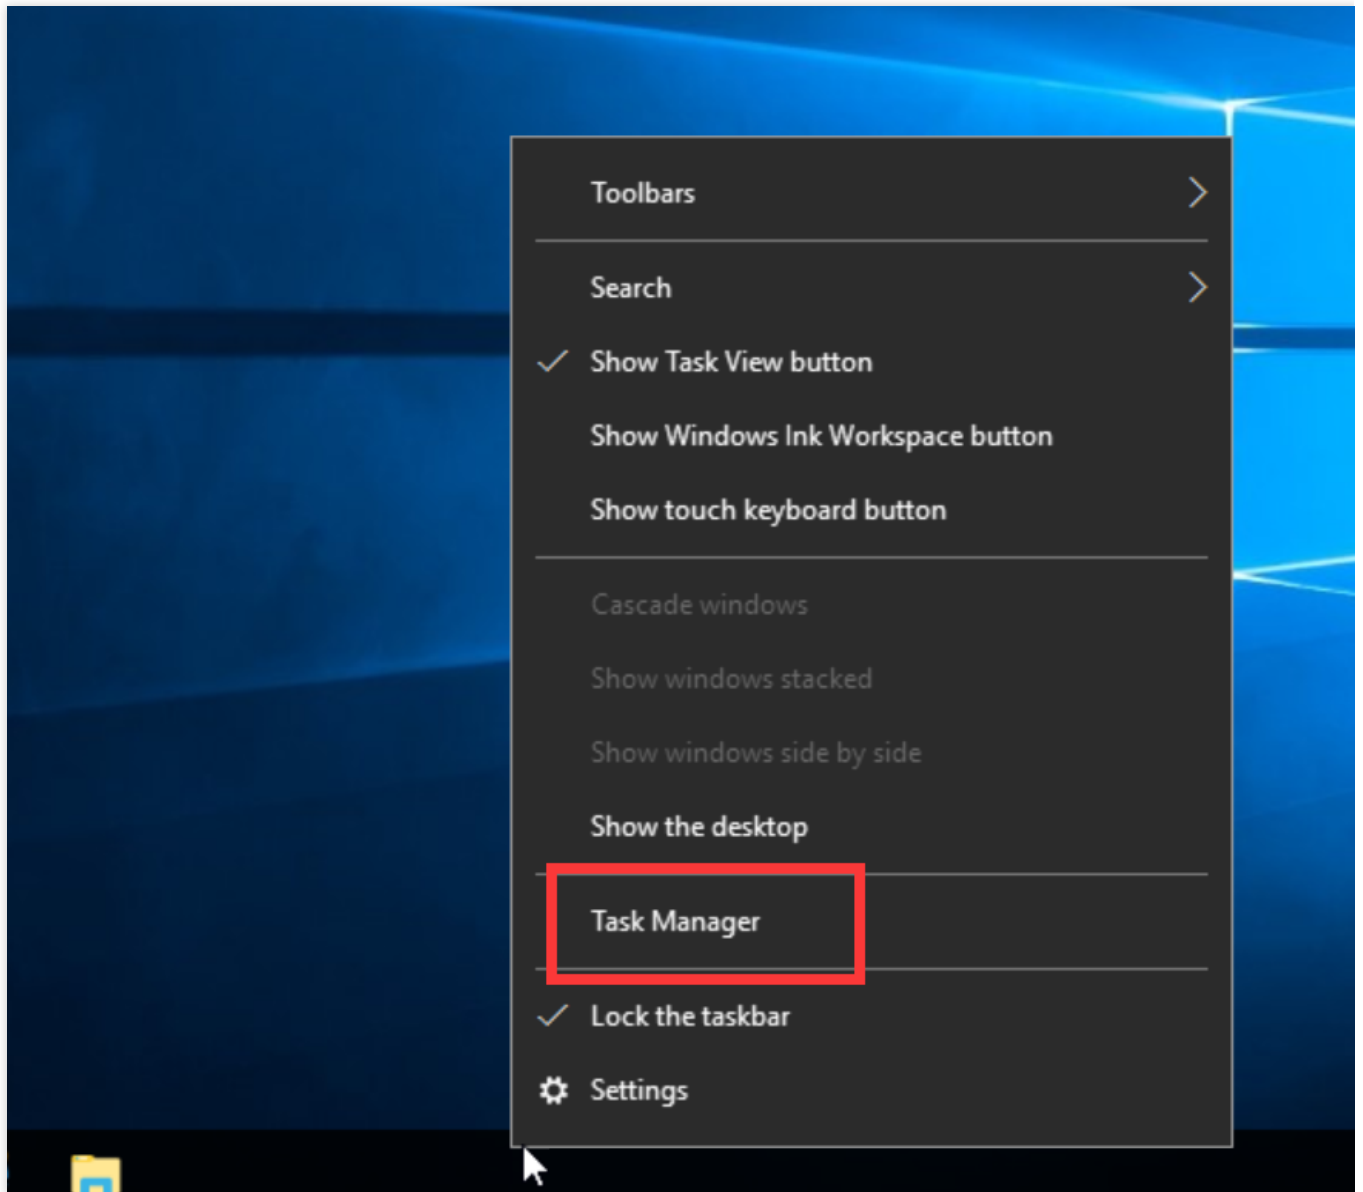


3. In the pop-up "Log into Windows instance" window, select **Alternative login methods (VNC)** and click **Log In Now** to log in to the CVM instance.
4. In the pop-up login window, select "Send CtrlAltDel" in the upper-left corner and click **Ctrl-Alt-Delete** to go to the OS login page, as shown in the following figure:



Viewing the resource usage of processes

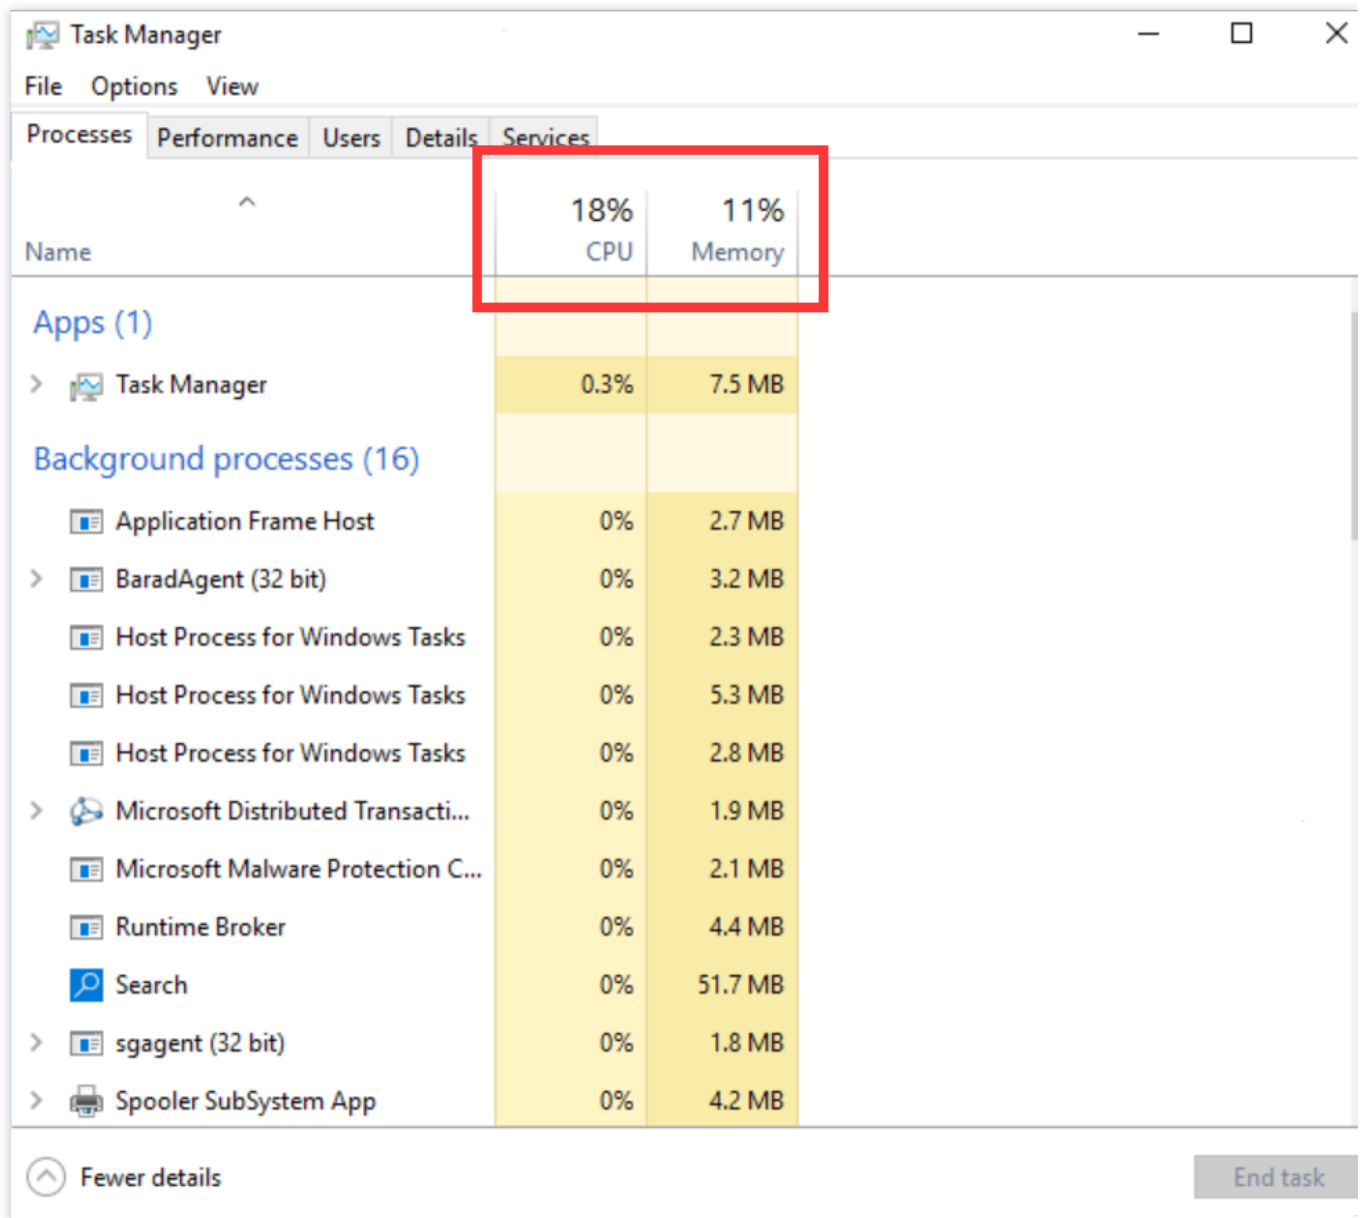
1. In the CVM, right-click the "taskbar" and choose **Task Manager**, as shown in the following figure:



2. View resource usage in the "Task Manager" window, as shown in the following figure:

Note:

You can click the CPU or memory column to sort the processes in ascending or descending order.



Analyzing processes

Analyze processes in Task Manager to identify the causes and troubleshoot accordingly.

A system process causes the issue

If a system process causes high CPU and memory utilization, troubleshoot as follows:

1. Check the name of the process.

Some viruses use names similar to system processes, such as `svch0st.exe`, `explore.exe`, `iexplorer.exe`, etc.

2. Check the location of the executable file of a process.

The executable file of a system process is usually located in `C:\Windows\System32` with valid signatures and descriptions. To locate the executable file of a process, such as `svchost.exe`, right-click the process in Task Manager and choose **Open file location**.

If the executable file is not in `C:\Windows\System32`, your CVM instance may have a virus. Scan for viruses with an antivirus application or manually fix the issue.

If the executable file is in `C:\Windows\System32`, restart your CVM instance or terminate unnecessary but secure system processes.

The following lists common system processes:

System Idle Process: a process that displays the percentage of time that the processor is idle

system: indicates the memory management process

explorer: indicates the desktop and file management process

ieexplore: indicates the Microsoft Internet Explorer process

csrss: indicates the runtime subsystem on the Microsoft client or server

svchost: indicates the system process for running DLL

Taskmgr: indicates the task manager

lsass: indicates the local security authority service

An unhealthy process causes the issue

If high CPU or memory utilization is caused by a process that has a strange name, such as xmr64.exe (a cryptomining malware), your CVM instance may have a virus or trojan. We recommend using a search engine to verify.

If the process is a virus or trojan, use an antivirus application to delete the virus or trojan. If necessary, back up your data and reinstall the operating system.

If the process is not a virus or trojan, restart your CVM instance or terminate unnecessary but secure processes.

A service process causes the issue

If a service process such as IIS, HTTPD, PHP, or Java causes the issue, we recommend analyzing it further.

For example, check whether your business volume is high.

If yes, we recommend [upgrading configuration](#) for your CVM instance. Alternatively, you can optimize service processes.

If no, use service error logs to further analyze the issue. For example, check whether incorrect parameter configurations lead to resource waste.

A Tencent Cloud process causes the issue

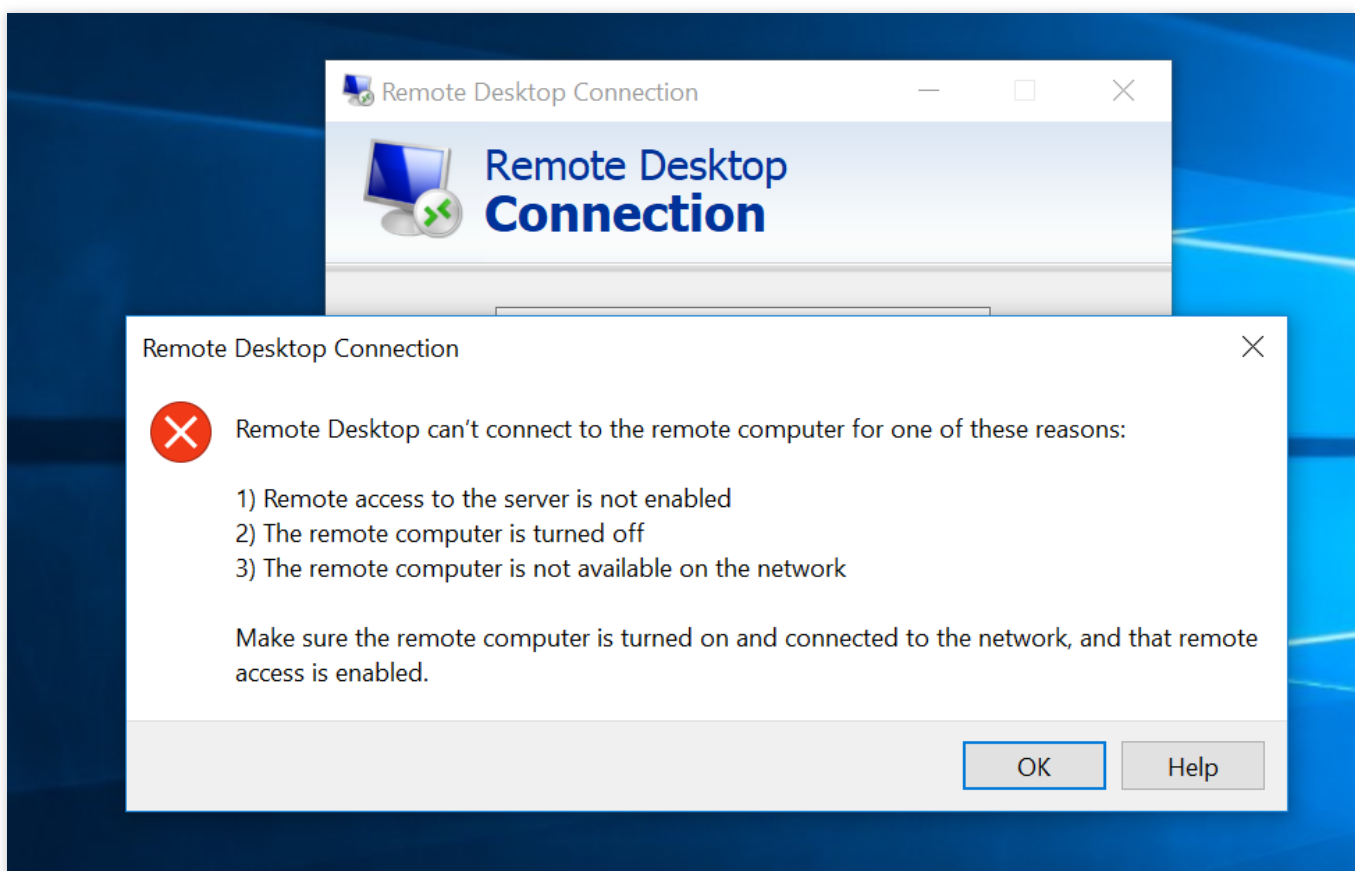
If a Tencent Cloud component process causes the issue, [submit a ticket](#) to contact us for assistance.

Failed to connect to a remote computer through Remote Desktop

Last updated : 2024-01-06 17:32:18

Symptom

When trying to connect to a Windows instance remotely from Windows, you get the following message:



The remote desktop can't connect to the remote computer for one of the following reasons:

1. Remote access to the server is not enabled
2. The remote computer is turned off
3. The remote computer is not available on the network

Make sure the remote computer is turned on and connected to the network, and that remote access is enabled.

Possible Causes

Possible causes include but are not limited to the following:

The instance is in an abnormal status.

The CVM instance does not have a public IP or the public network bandwidth is 0.

The remote login port (3389 by default) is not opened in the security group bound to the instance.

Remote Desktop Services is not started.

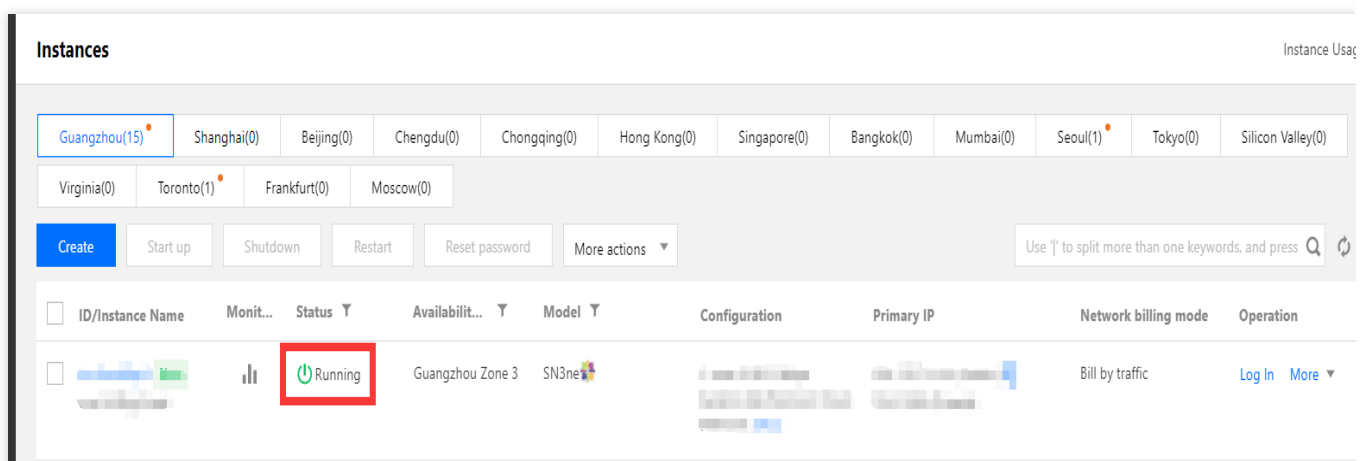
Incorrect Remote Desktop settings

Incorrect Windows Firewall settings

Troubleshooting Directions

Checking instance status

1. Log in to the [CVM console](#).
2. On the **Instances** page, check whether the instance is "Running".



If yes, [check whether the CVM instance has a public IP](#).

If no, start up the Windows instance.

Checking public IP

Check whether the CVM instance has a public IP in the CVM console.

Instances Instance Usage

Guangzhou(15) Shanghai(0) Beijing(0) Chengdu(0) Chongqing(0) Hong Kong(0) Singapore(0) Bangkok(0) Mumbai(0) Seoul(1) Tokyo(0) Silicon Valley(0)

Virginia(0) Toronto(1) Frankfurt(0) Moscow(0)

Create Start up Shutdown Restart Reset password More actions

Use " " to split more than one keywords, and press Q

ID/Instance Name	Monit...	Status	Availabilit...	Model	Configuration	Primary IP	Network billing mode	Operation
		Running	Guangzhou Zone 3	SN3ne	1-core 2 GB 1 Mbps	193.112.71.133 (Public)	Bill by traffic	Log In More

If yes, [check whether you have purchased public network bandwidth](#).

If no, [apply for an EIP and bind it to the CVM instance](#).

Checking public network bandwidth

Check whether the public network bandwidth is 0 Mbps. At least 1 Mbps of public network bandwidth is required.

If yes, increase the bandwidth to 5 Mbps or above by [adjusting network configuration](#).

Instances Instance Usage Gu

Guangzhou(15) Shanghai(0) Beijing(0) Chengdu(0) Chongqing(0) Hong Kong(0) Singapore(0) Bangkok(0) Mumbai(0) Seoul(1) Tokyo(0) Silicon Valley(0)

Virginia(0) Toronto(1) Frankfurt(0) Moscow(0)

Create Start up Shutdown Restart Reset password More actions

Use " " to split more than one keywords, and press Q

ID/Instance Name	Monit...	Status	Availabilit...	Model	Configuration	Primary IP	Network billing mode	Operation
		Running	Guangzhou Zone 3	SN3ne	1-core 2 GB 1 Mbps System disk: Premium Cloud Network: VPC2		Bill by traffic	Log In More

If no, [check whether the remote login port \(3389\) of the instance is opened](#).

Checking remote login port (3389)

1. On the instance management page in the CVM console, click the target instance ID/name to enter its details page.
2. In the **Security Groups** tab, check whether the remote login port (3389 by default) is opened.

The screenshot displays the 'Security Groups' management interface. On the left, under 'Bound to security group', there is a table with two entries:

Prior...	Security Group ID/name	Operation
1	Open all ports-2	Unbind
2	Open all ports	Unbind

On the right, the 'Rule preview' section shows the 'Inbound rule' details for 'Open all ports-2'. A red box highlights the following rule:

Source	Port Protocol	Policy	Notes
0.0.0.0/0	TCP:3389	Allow	-
0.0.0.0/0	ALL	Allow	-

If yes, [check the remote desktop service](#).

If no, edit the corresponding security group rules to open the port as instructed in [Adding Security Group Rules](#).

Checking remote desktop service

1. [Log in to the instance via VNC](#) and check whether the remote desktop service of the Windows instance is enabled.

Note:

The following operations use an instance on Windows Server 2016 as an example.

2. Right-click



and select **System** in the pop-up menu.

3. In the *System* pop-up window, select **Advanced System Settings**.

4. In the *System Properties* pop-up window, select the **Remote** tab and check whether the **Allow remote connections to this computer** is selected:

If yes, proceed to [step 5](#).

If no, select it and click **OK**.

5.

Right-click



and select **Computer Management** in the pop-up menu.

6. On the left sidebar in the **Computer Management** window, select **Services and Applications > Services**.

7. In the service list on the right, check whether **Remote Desktop Services* is started:

If yes, proceed to [step 8](#).

If no, start the service.

8.

Right-click



and select **Run** in the pop-up menu.

9. In the **Run** pop-up window, enter **msconfig** and click **OK**.

10. In the **System Configuration** pop-up window, check whether **Normal startup** is selected:

If yes, [check Windows instance system settings](#).

If no, select it and click **OK**.

Checking Windows instance system settings

1. [Log in to the Windows instance via VNC](#) and check the system settings of the instance.

Note:

The following operations use an instance on Windows Server 2012 as an example.

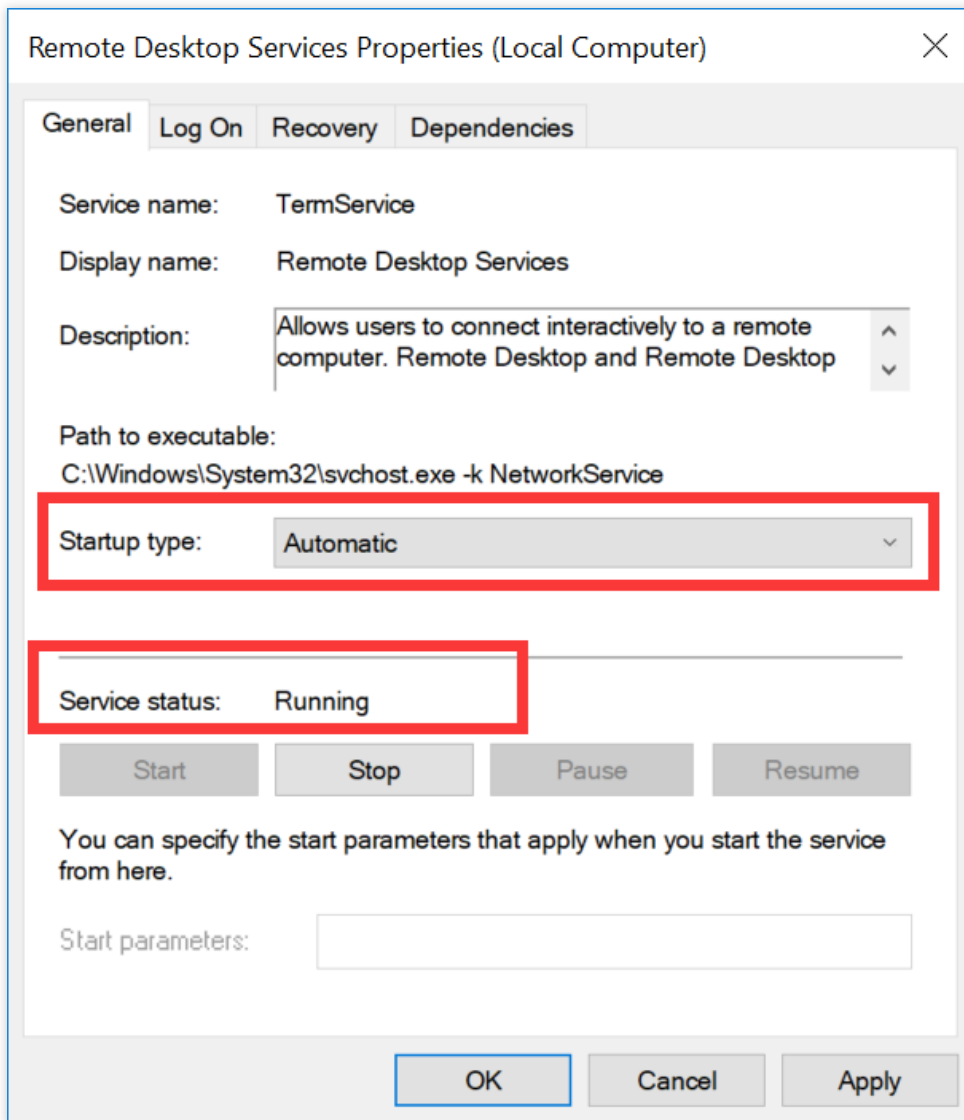
2. Right-click



and select **Run** in the pop-up menu.

3. In the **Run** pop-up window, enter **services.msc** and press **Enter** to open the **Services** window.

4. Double-click to open the **Remote Desktop Services** properties and check whether **Remote Desktop Services** is running as shown below:



If yes, proceed to [step 5](#).

If no, set **Startup Type** to **Automatic** and **Service Status** to **Running** (i.e., clicking **Start**).

5.

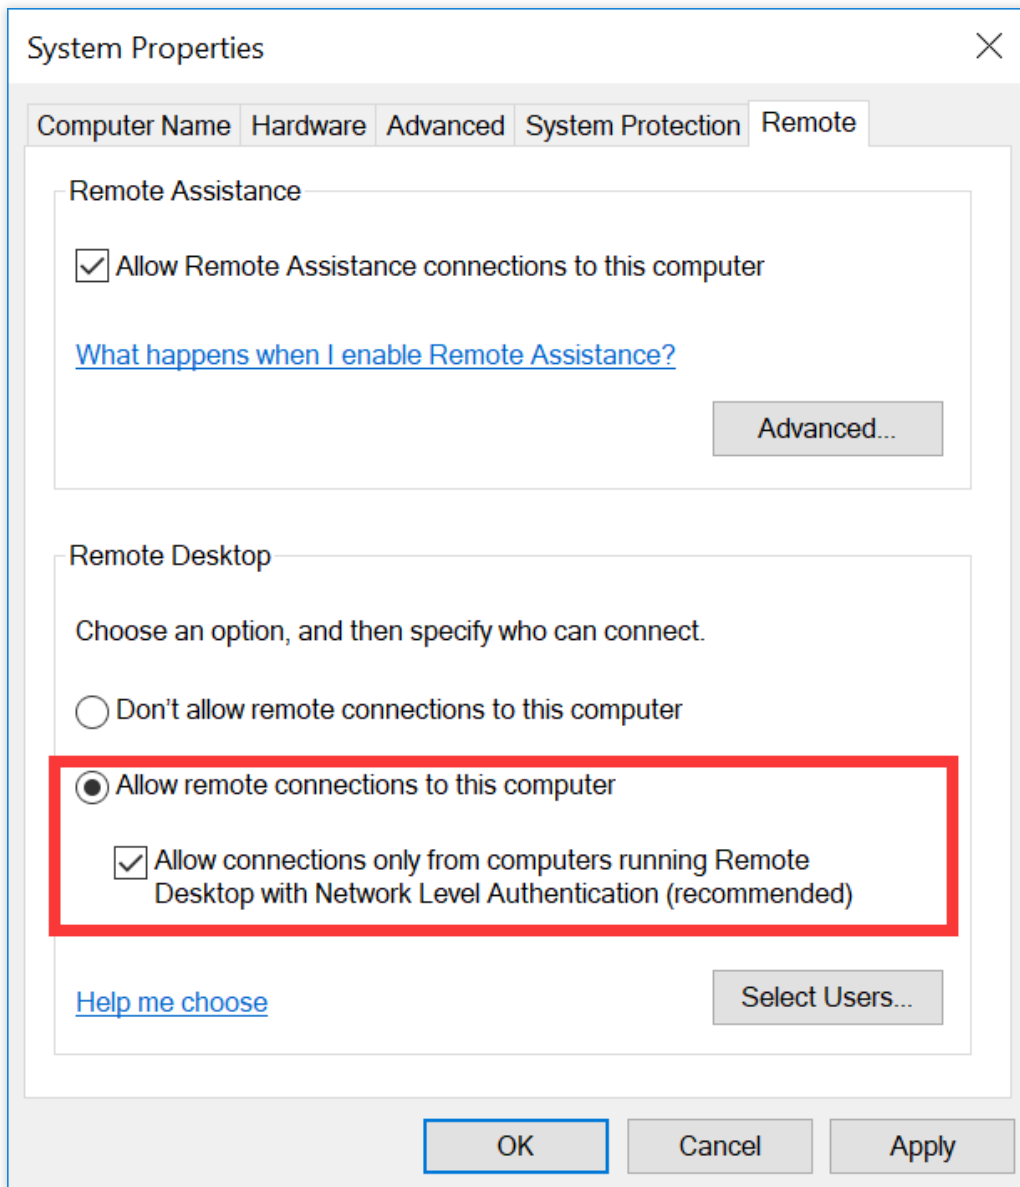
Right-click



and select **Run** in the pop-up menu.

6. In the **Run** pop-up window, enter **sysdm.cpl** and press **Enter** to open the **System Properties** window.

7. On the **Remote** tab, check whether the **Remote Desktop** is set to **Allow remote connections to this computer** as shown below:



If yes, proceed to [Step 8](#).

If no, set **Remote Desktop** to **Allow remote connections to this computer**.

8.

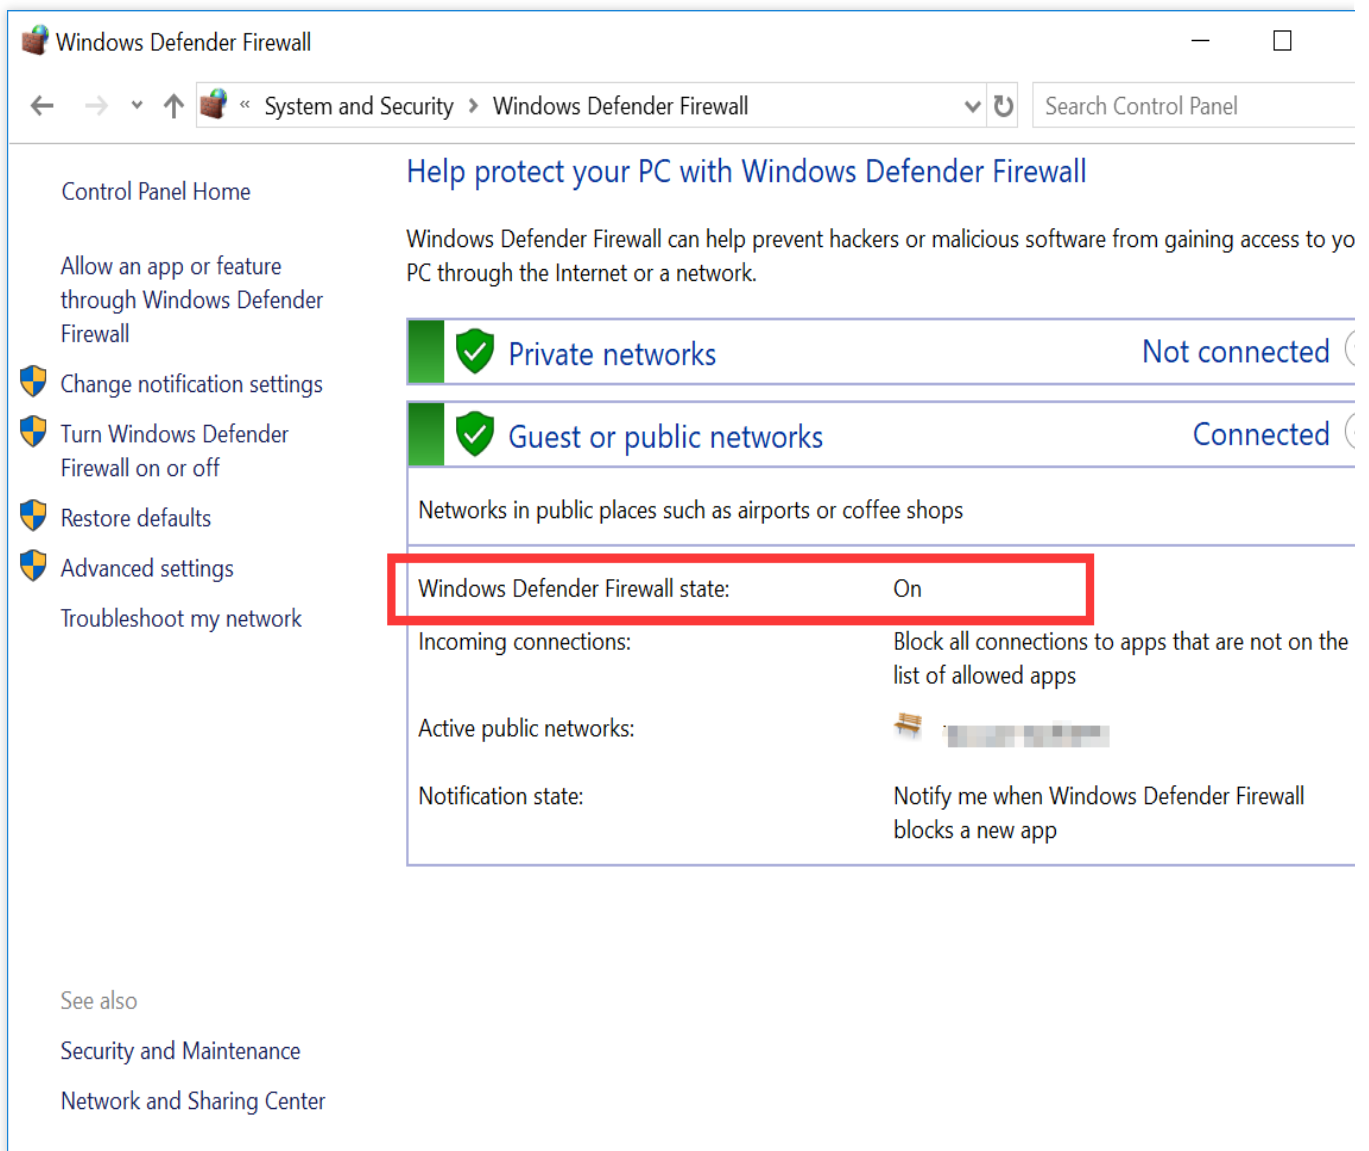
Click



and select **Control Panel** in the pop-up menu.

9. In **Control Panel**, select **System and Security** > **Windows Defender Firewall**.

10. In **Windows Defender Firewall**, check the Windows Defender Firewall status as shown below:



If the status is **On**, proceed to [Step 11](#).

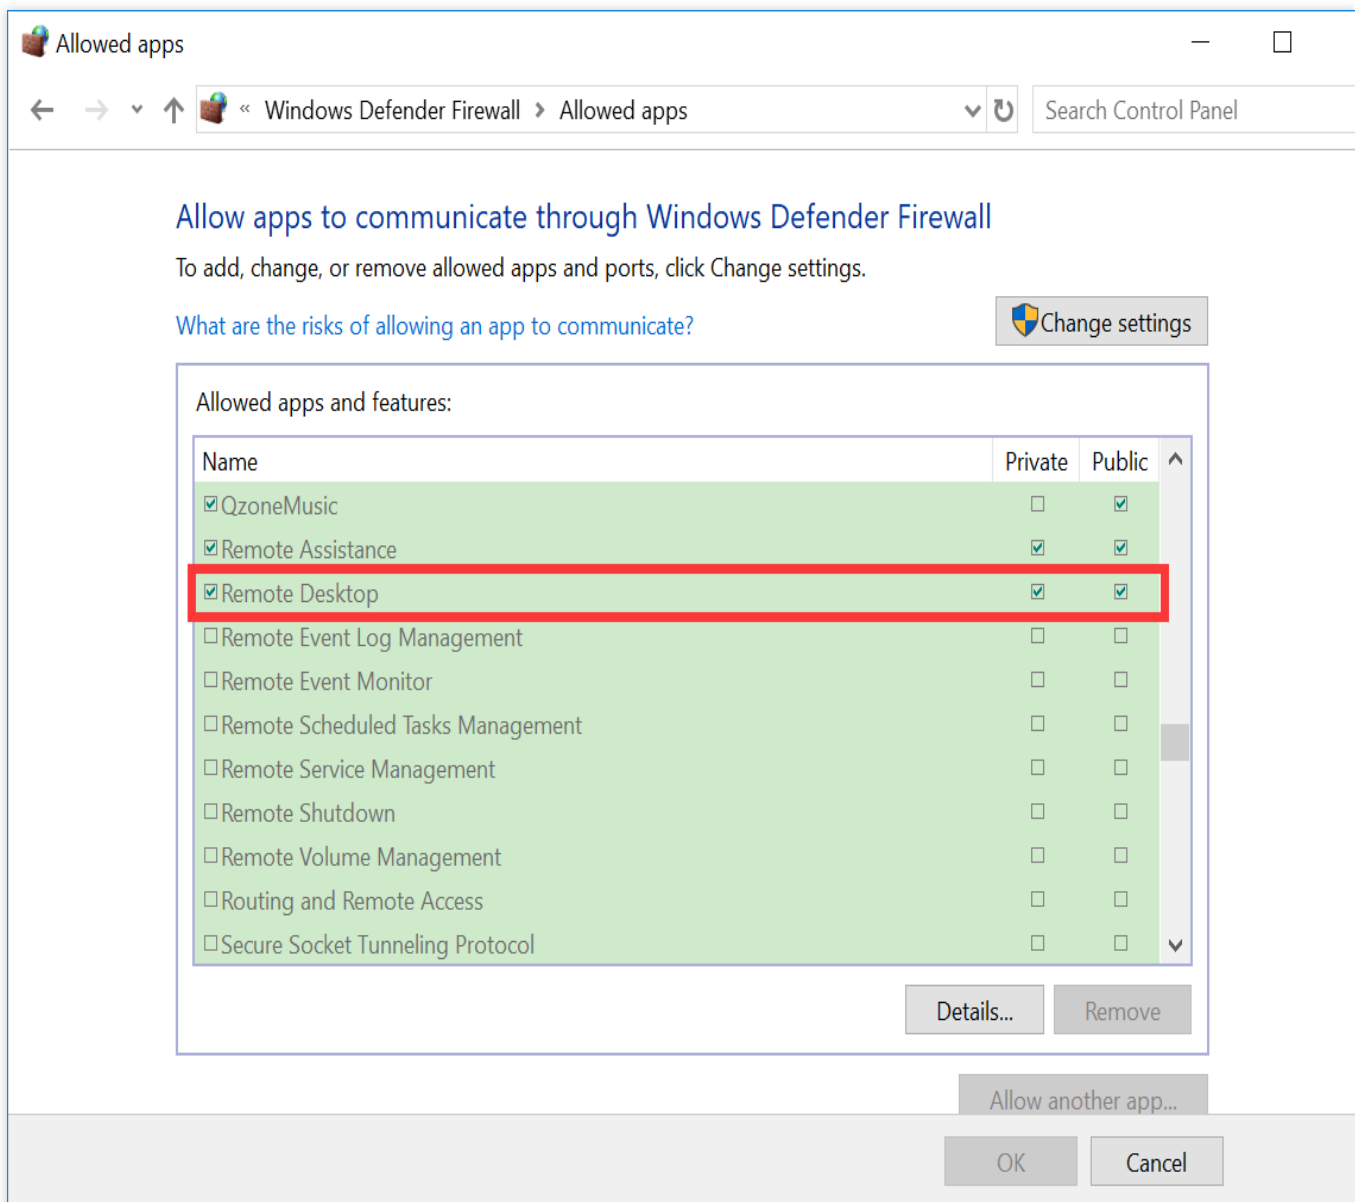
If the status is **Off**, contact [submit a ticket](#) for assistance.

11.

In Windows Defender Firewall

, click **Allow an app or feature through Windows Defender Firewall** to open the **Allowed apps** window.

12. In the **Allowed apps** window, check whether **Remote Desktop** is selected in **Allowed apps and features**.



If yes, proceed to [Step 13](#).

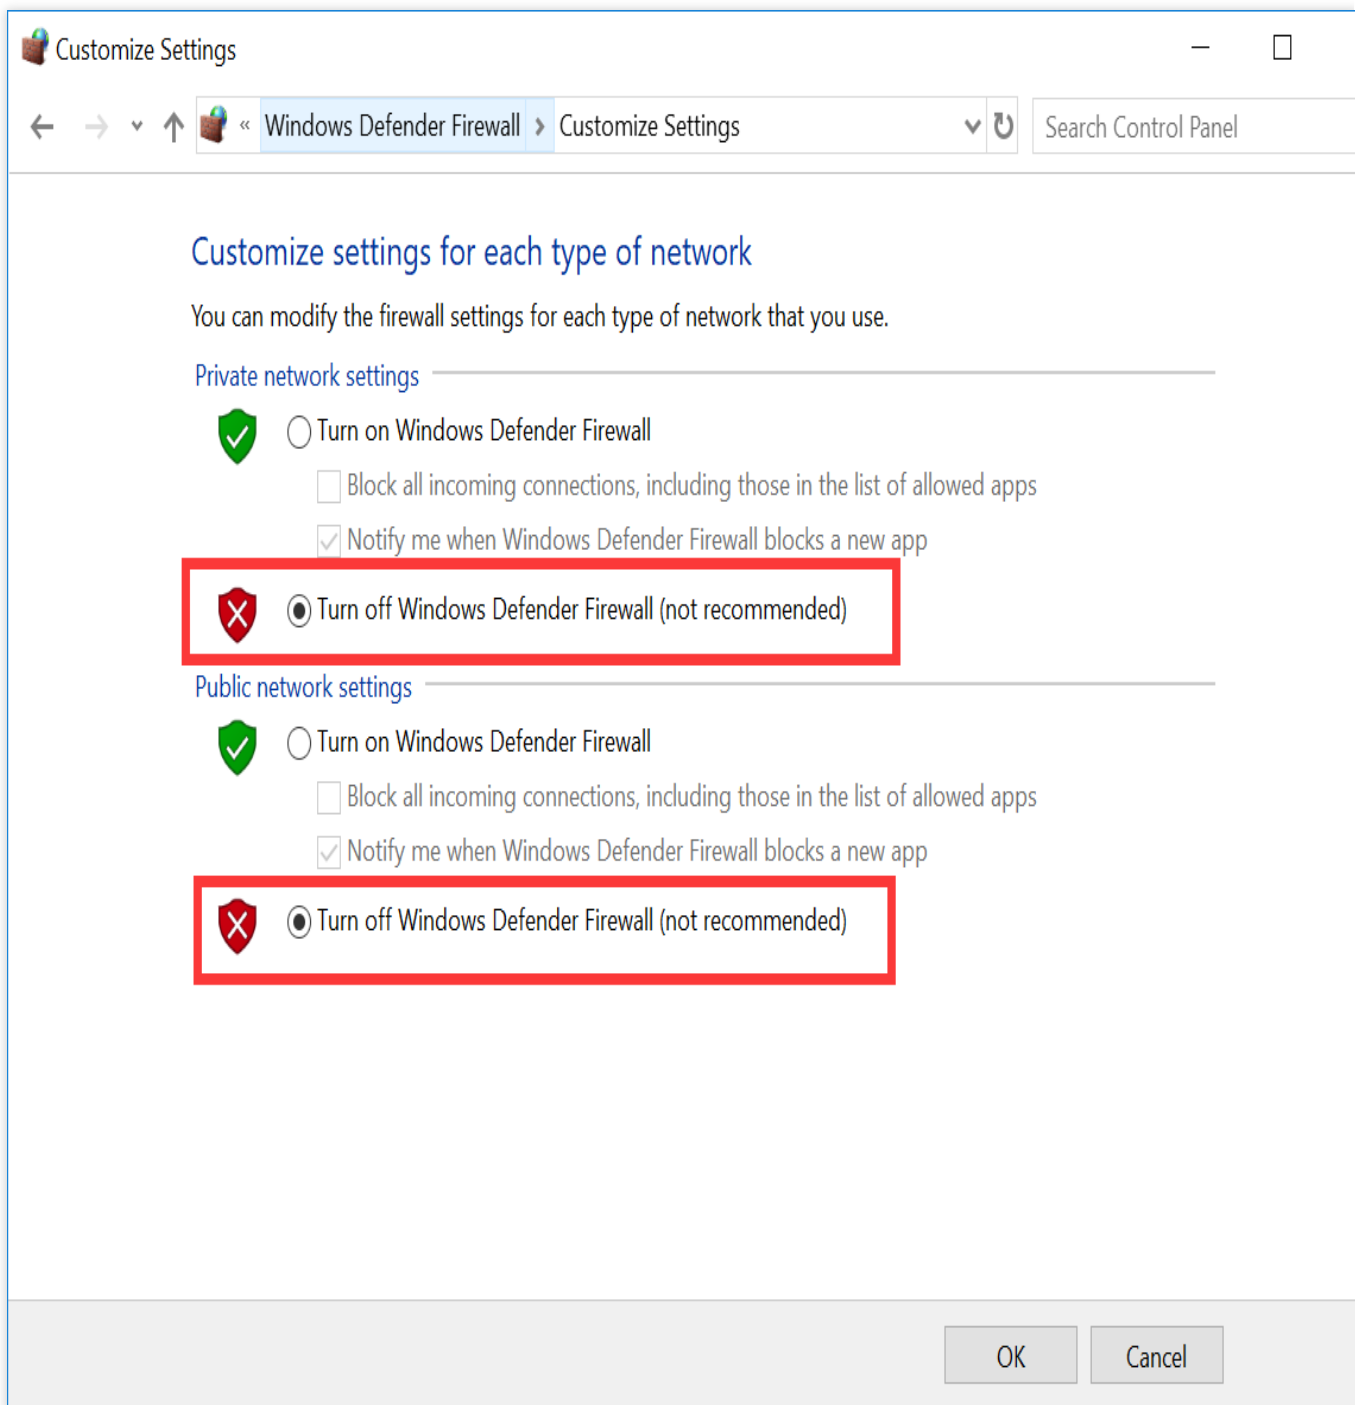
If no, select **Remote Desktop** to allow **Remote Desktop** through Windows Defender Firewall.

13.

In **Windows Defender**

Firewall, click **Turn Windows Firewall on or off** to open the **Customize Settings** window.

14. In the **Customize Settings** window, set **Private network settings** and **Public network settings** to **Turn off Windows Defender Firewall (not recommended)** as shown below:



If you still cannot connect to the Windows instance on a remote desktop, please [submit a ticket](#) for assistance.


Credentials Not Work

Last updated : 2024-01-06 17:32:18

Issue

The following error message appears when trying to log in to a Windows CVM remotely via RDP protocol, such as using MSTSC.

The credentials that were used to connect to `xxx.xxx.xxx.xxx` did not work. Please enter new credentials.



Solutions

Note:

This document uses a Tencent Cloud CVM with the Windows Server 2012 operating system as an example. The procedure may vary slightly according to the operating system version.

Follow the instructions below and try to connect to your Windows CVM after each step. If the issue persists, proceed to the next step.

Step 1: modify the network access policy

1. [Log in to the Windows instance using VNC.](#)

2. Once you log in, click

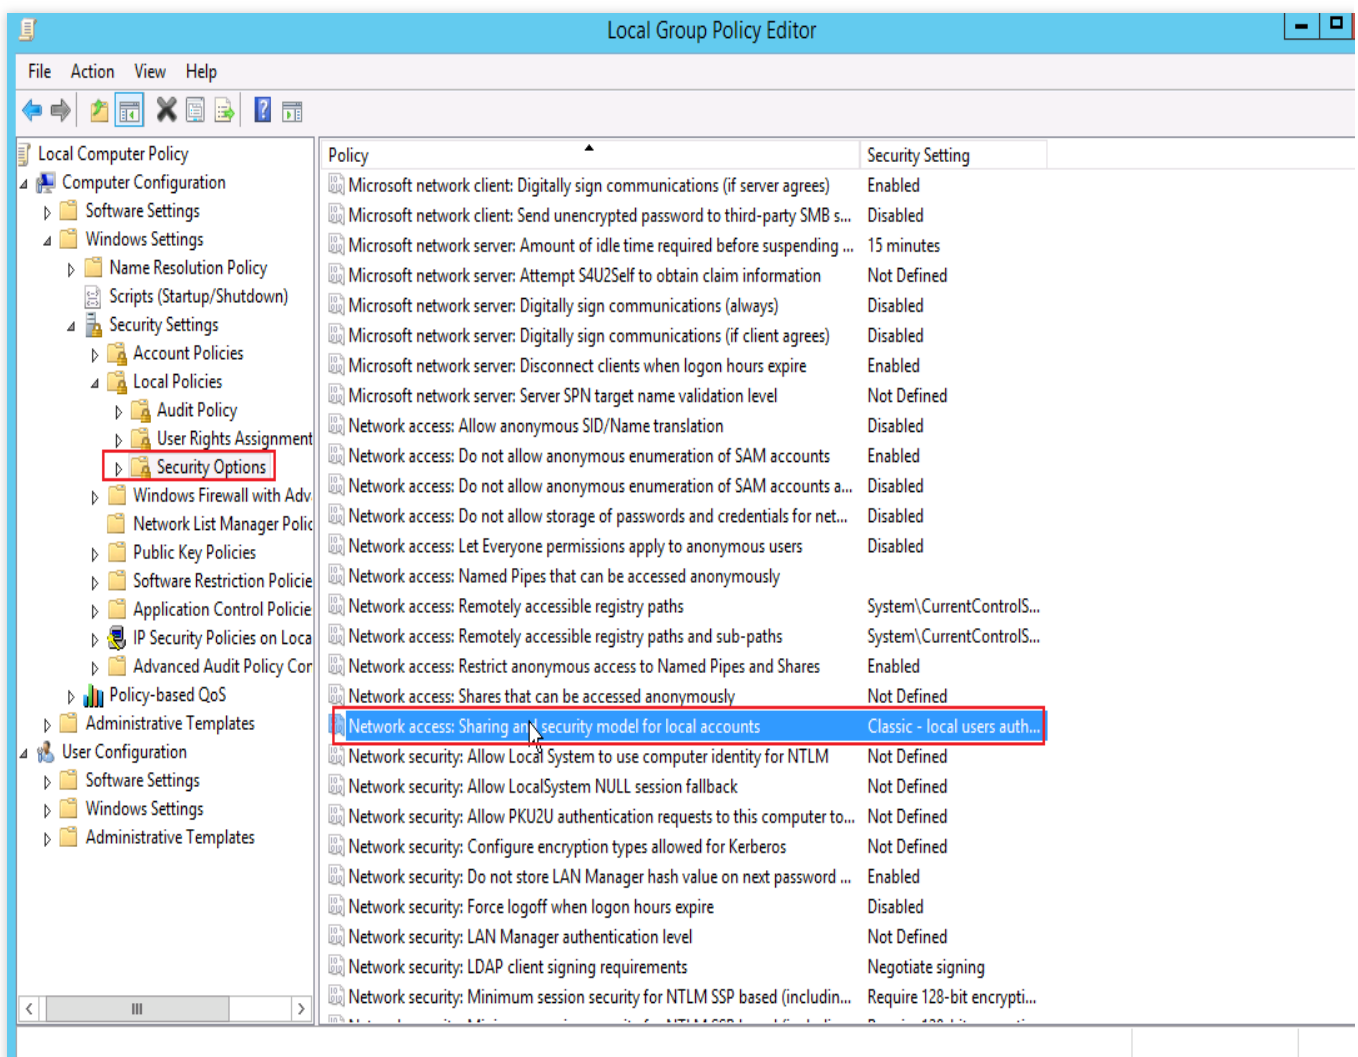


to open the **Windows PowerShell** window.

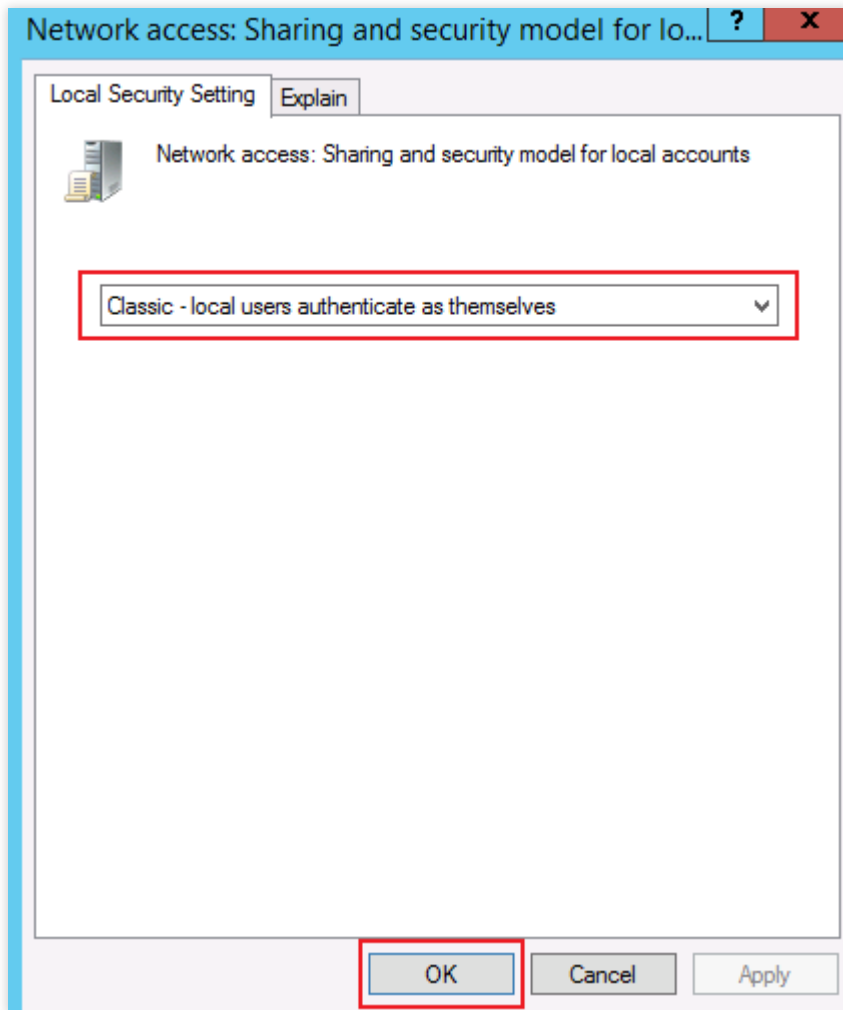
3. In the **Windows PowerShell** window, enter **gpedit.msc** and press **Enter**. The **Local Group Policy Editor** window appears.

4. In the left sidebar, expand the following directories in order: **Computer Configuration > Windows Settings > Security Settings > Local Policies > Security Options**.

5. Locate and open **Network access: Sharing and security model for local accounts** under **Security Options**, as shown below:



6. Select **Classic - local users authenticate as themselves** and click **OK**, as shown below:



7. Check whether you can connect to your Windows CVM now.

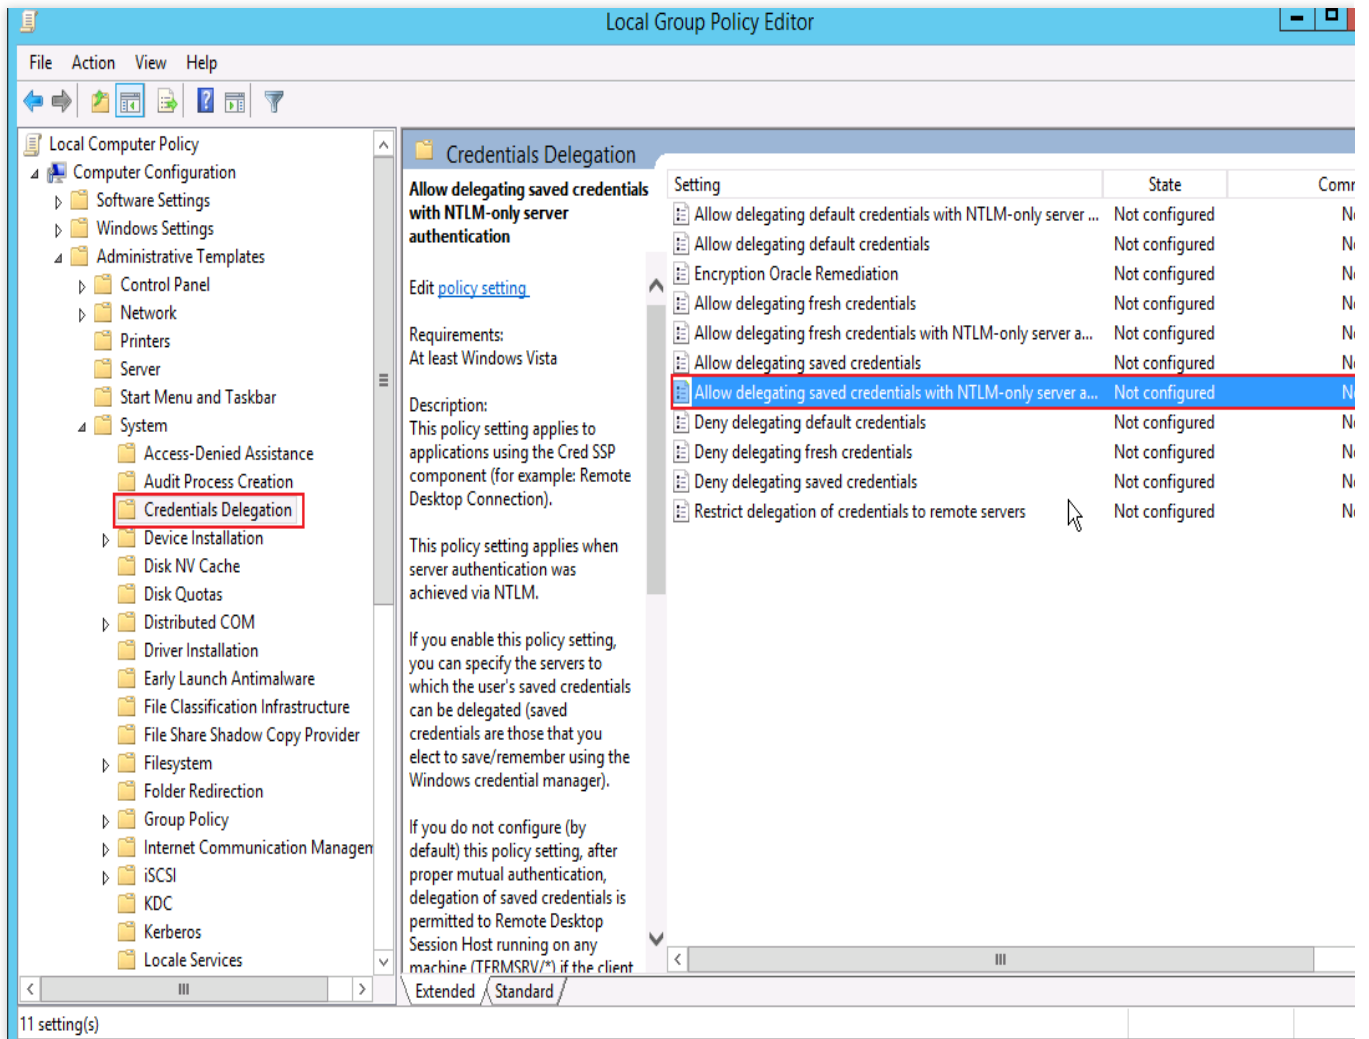
If yes, the problem has been solved.

If no, proceed to "Step 2: modify credentials delegation".

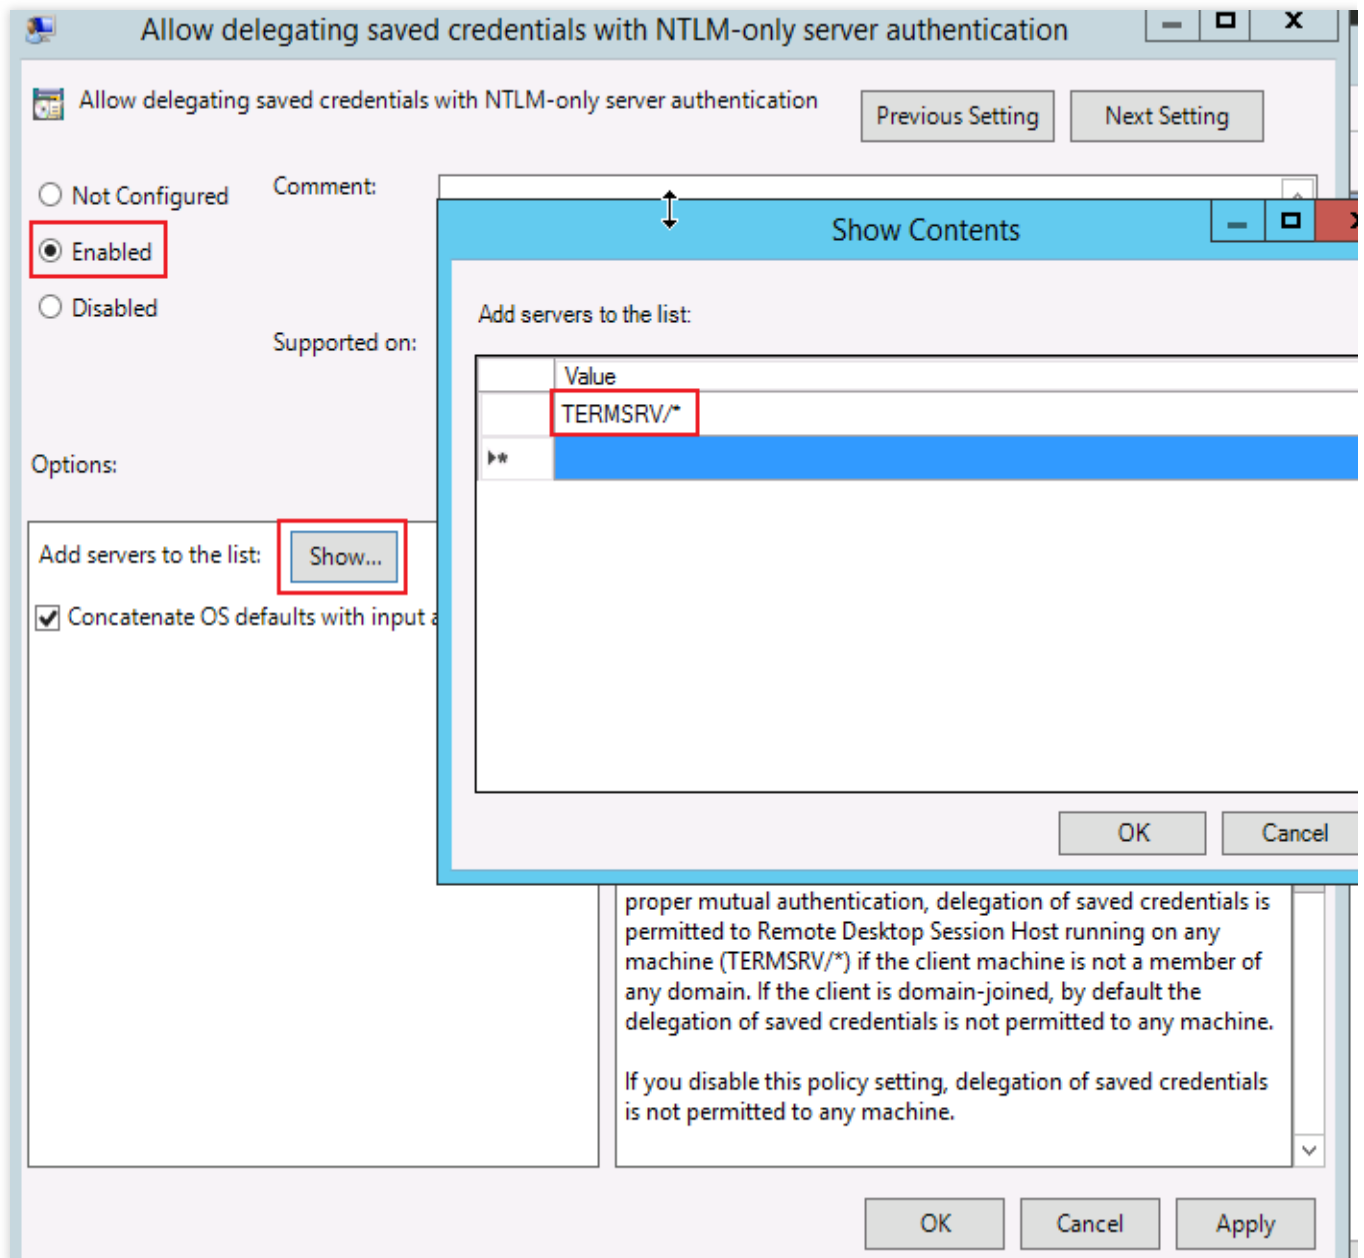
Step 2: modify credentials delegation

1. In the left sidebar of **Local Group Policy Editor**, expand the following directories in order: **Computer Configuration > Administrative Templates > System > Credentials Delegation**.

2. Locate and open **Allow delegating saved credentials with NTLM-only server authentication** under **Credentials Delegation**, as shown below:



3. In the pop-up window, select **Enabled**. Click **Show...** under **Options**, enter `TERMSRV/*` for **Show Contents** and click **OK**, as shown below:



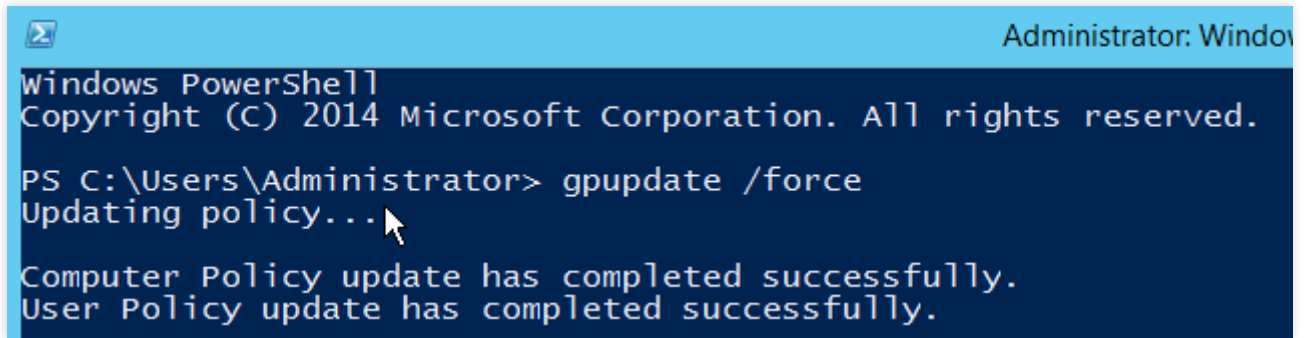
4. Click **OK**.

5. On the desktop, click



to open the **Windows PowerShell** window.

6. In the **Windows PowerShell** window, enter **gpupdate /force** and press **Enter** to update the group policy, as shown below:



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2014 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> gpupdate /force
Updating policy...
Computer Policy update has completed successfully.
User Policy update has completed successfully.
```

7. Check whether you can connect to your Windows CVM now.

If yes, the problem has been solved.

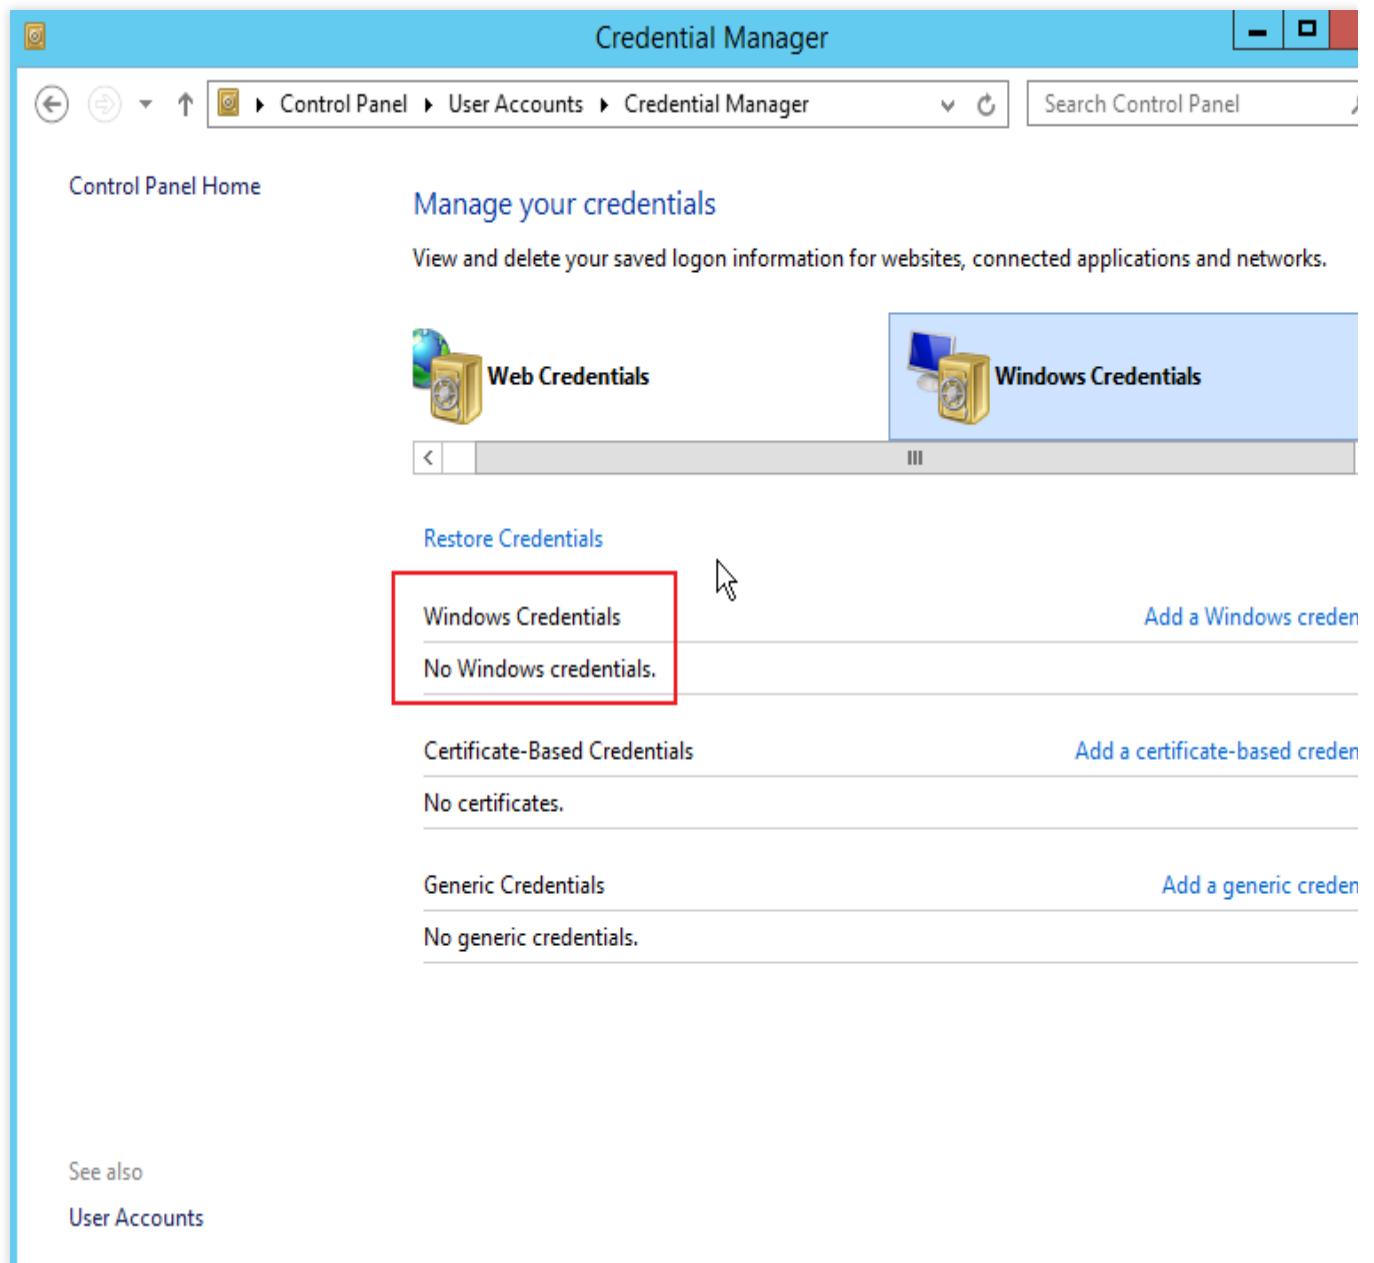
If no, proceed to “Step 3: configure local credentials”.

Step 3: configure local credentials

1. On the desktop, click



and navigate to **Control Panel > Users Accounts > Credential Manager**. Click **Windows Credentials** to display Windows credentials, as shown below:

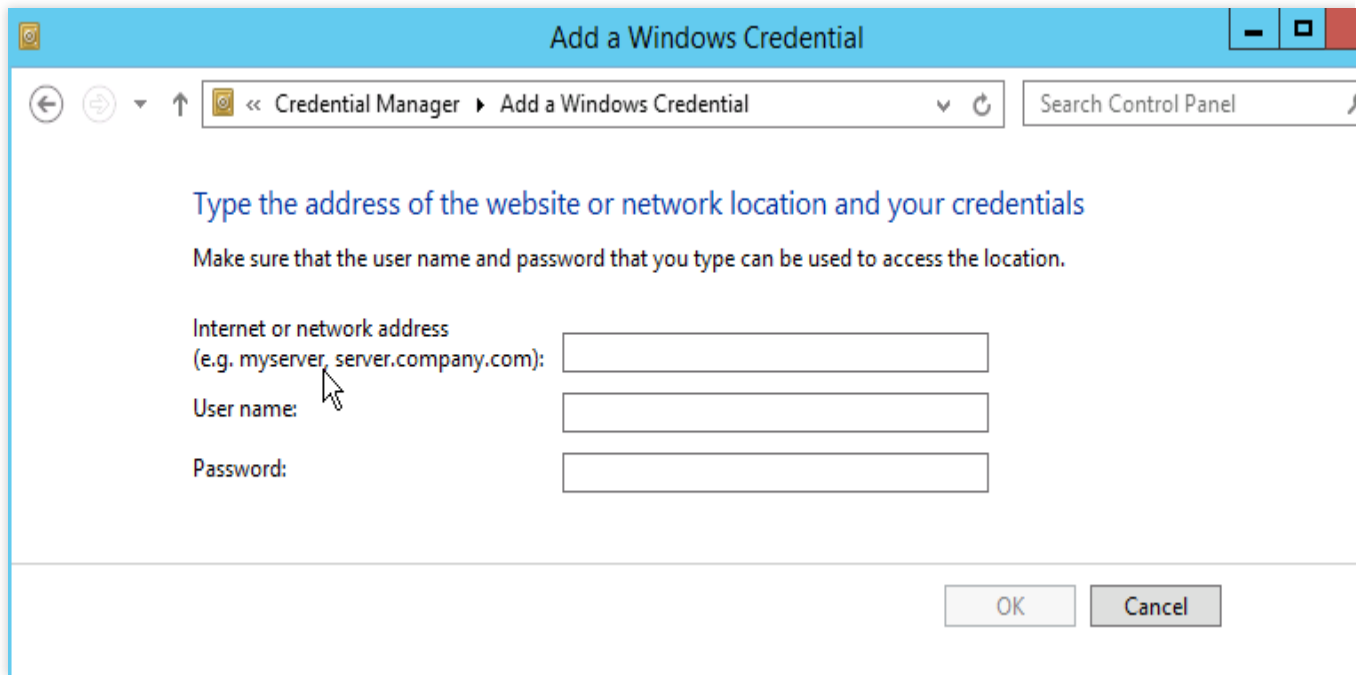


2. Check whether the credential you used to log in to your Windows CVM exists.

If no, proceed to the next step to add Windows credentials.

If yes, proceed to "Step 4: turn off password protected sharing".

3. Click **Add a Windows credential** to go to the **Add a Windows Credential** window, as shown below:



4. Enter the IP address, username, and password of the CVM instance and click **OK**.

Note:

The IP address refers to the public IP address of your CVM instance. For more information, see [Getting Public IP Addresses](#).

The default username for a Windows instance is `Administrator` and the password is set when you create the instance. If you've forgotten your password, you can [reset the instance password](#).

5. Check whether you can connect to your Windows CVM now.

If yes, the problem has been solved.

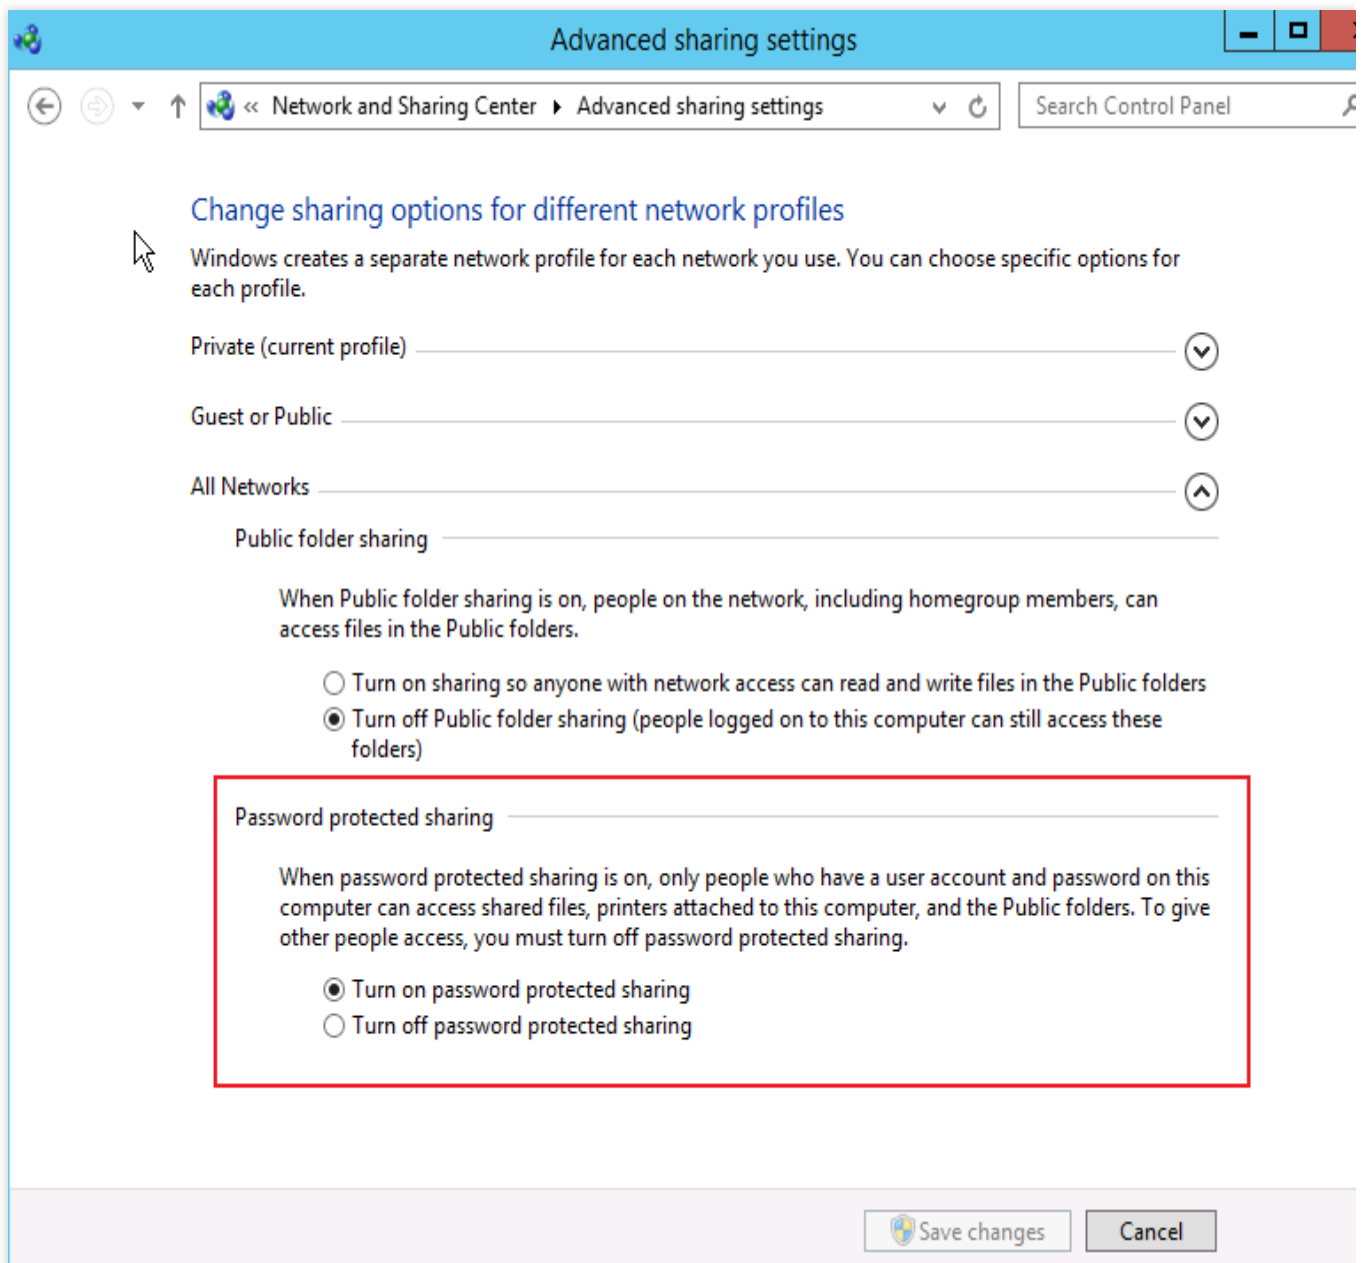
If no, proceed to "Step 4: turn off password protected sharing".

Step 4: turn off password protected sharing

1. On the desktop, click



and navigate to **Control Panel > Network and Sharing Center > Advanced sharing settings**, as shown below:



2. Expand the **All Networks** tab, select **Turn off password protected sharing** under **Password protected sharing**, and click **Save changes**.

3. Check whether you can connect to your Windows CVM now.

If yes, the problem has been solved.

If no, please [submit a ticket](#) for assistance.


Windows instance: no remote Desktop license server can provide license

Last updated : 2024-01-06 17:32:18

This document describes how to manage alarm prompts such as "Remote session has been disconnected as no remote desktop authorization server is available for licensing" when you try to remotely connect to a Windows instance.

Problem

When you try to connect to a Windows instance by using Windows Remote Desktop, a prompt stating "Remote session has been disconnected as no remote desktop authorization server is available for licensing. For assistance, please contact your system administrator" appears, as shown in the following figure:



Problem Analysis

The possible causes to this problem include but are not limited to the following. Therefore, always analyze the problem based on the actual situation.

The RDP-TCP limit is set by the system by default, and it allows only one session for each user. If the account has been logged in, no additional sessions can be established.

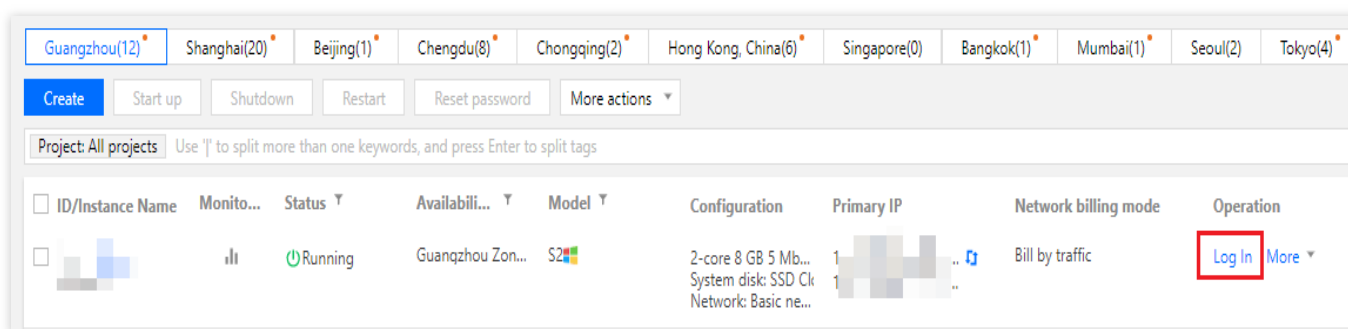
The "Remote Desktop Session Host" role feature was added by the system, but the validity period of the feature has expired.

The "Remote Desktop Session Host" role feature is free for use for 120 days. After the period, you must pay for the feature to continue to use it.

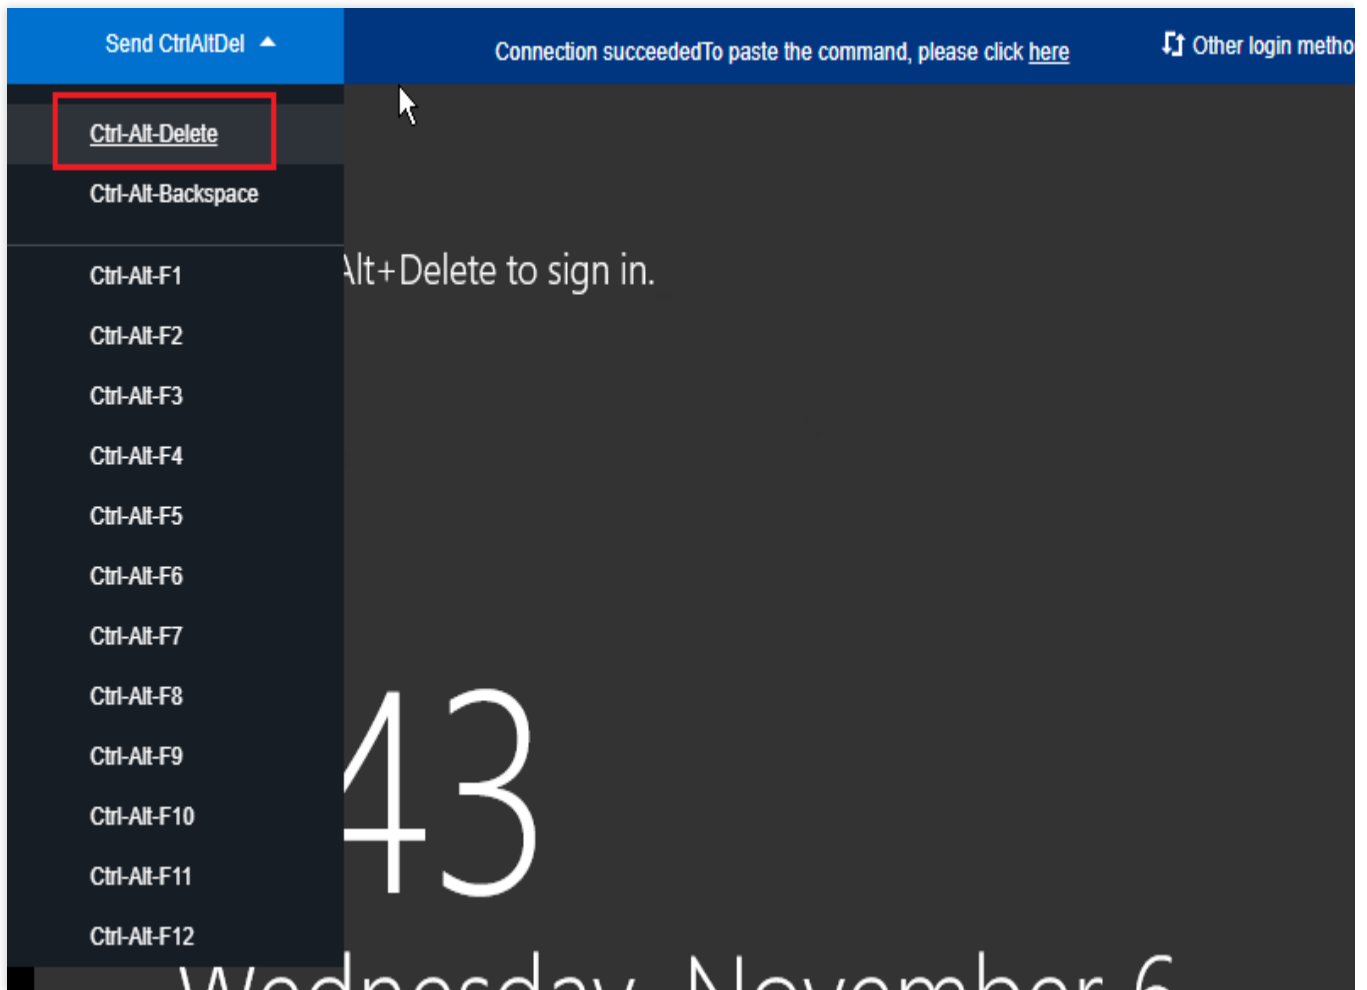
Solution

Logging in to the CVM through VNC

1. Log in to the [CVM console](#).
2. On the instance management page, locate the target CVM instance and click **Log In**, as shown in the following figure:



3. In the **Log in to Windows instance** window that appears, select **Alternative login methods (VNC)**, and click **Log In Now** to log in to the CVM.
4. In the login window that appears, select **Send Remote Command** in the upper-left corner, and press **Ctrl-Alt-Delete** to go to the system login interface, as shown in the following figure:



Solution 1: Modify the policy configuration

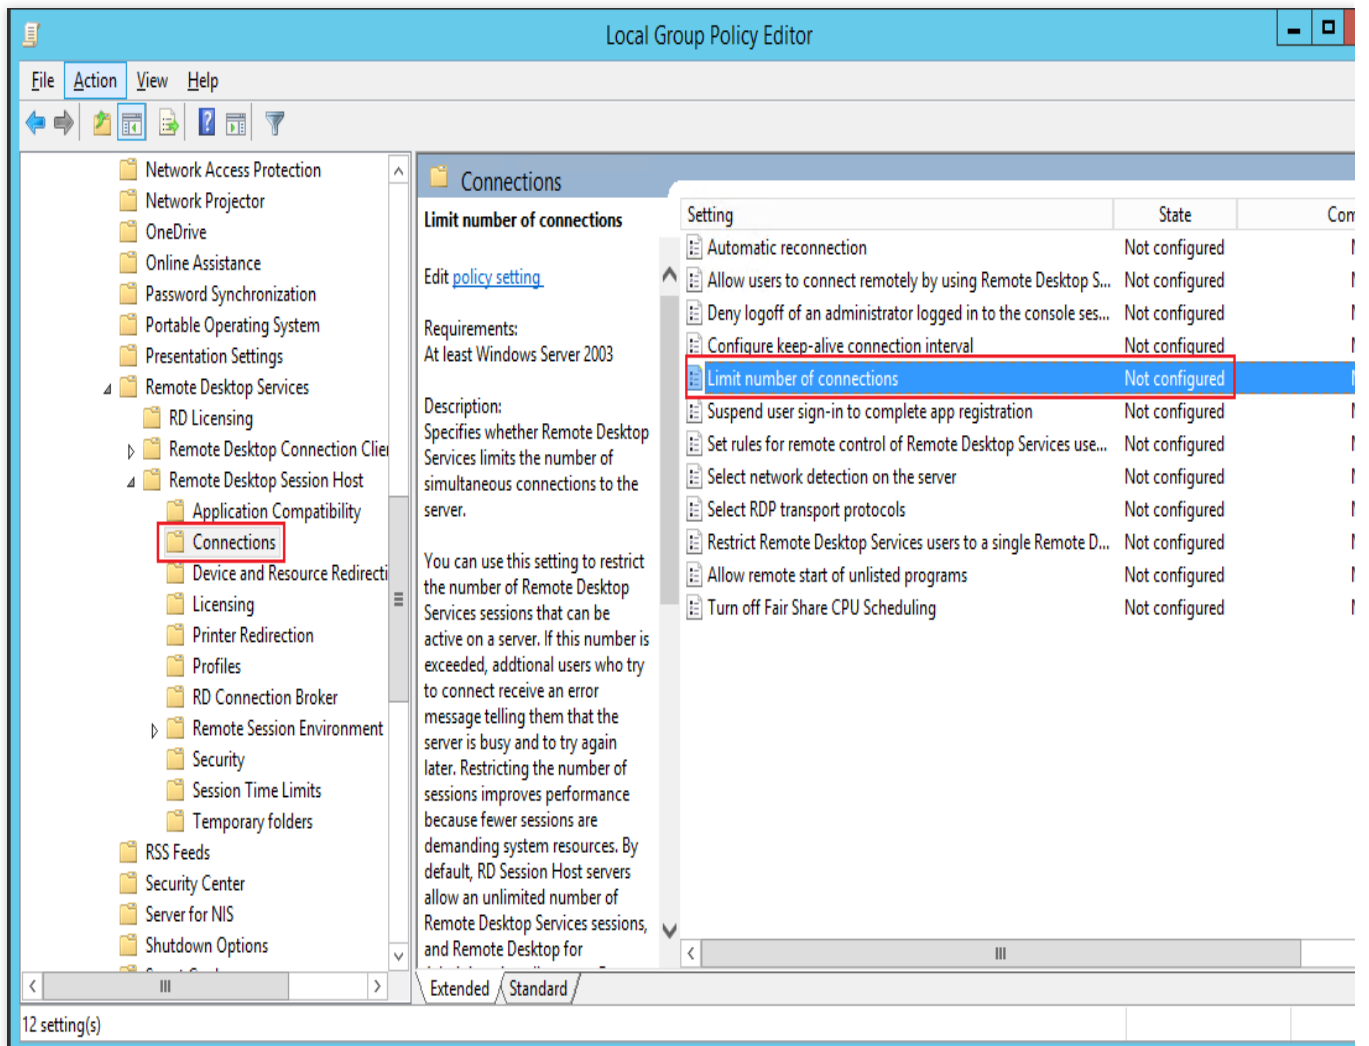
1. On the operating system interface, click




to open a Windows PowerShell window.

2. In the Windows PowerShell window, enter **gpedit.msc** and press **Enter** to open **Local Group Policy Editor**.

3. In the left navigation tree, choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections**, and double-click **Limit number of connections**, as shown in the following figure:



4. In the "Limit number of connections" window that appears, modify **Maximum RD connections supported** and click **OK**, as shown in the following figure:



5. Switch to the Windows PowerShell window.

6. In the Windows PowerShell window, enter **gpupdate** and press **Enter** to update the policy.

Solution 2: Delete the "Remote Desk Session Host" role

Note:

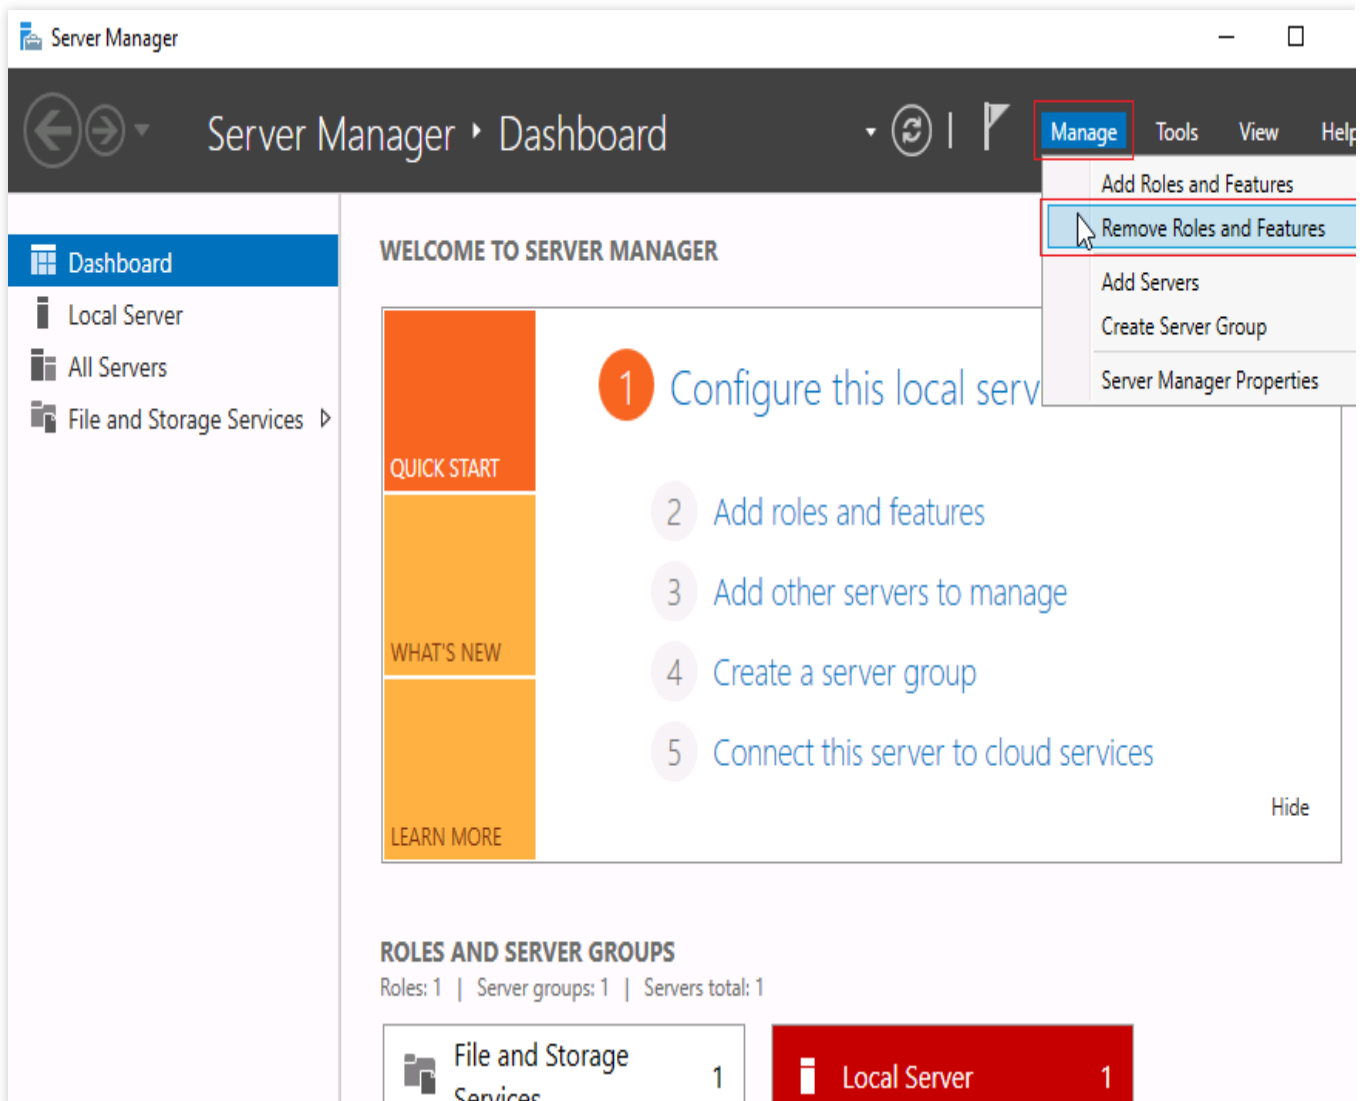
If you want to keep the "Remote Desktop Session Host" role, skip this step and go to the [Microsoft official website](#) to purchase and configure the appropriate certificate.

1. On the operating system interface, click



to open **Server Manager**.

2. Click **Manage** in the upper-right corner of the "Server Manager" window and select **Delete roles and features**, as shown in the following figure:



3. In the "Delete roles and features" wizard, click **Next**.

4. On the "Delete server roles" page, uncheck **Remote Desktop Services**. In the prompt box that appears, select **Remove Feature**.

5. Click **Next** twice.

6. Check **Restart the destination server automatically if required**, and click **Yes** in the prompt box that appears.

7. Click **Delete**.

Wait for the CVM to restart.

Remote Login Failure Due To Port Issues

Last updated : 2024-01-06 17:32:18

This article describes how to troubleshoot remote login failures caused by port problems.

Note:

The following uses a CVM instance running Windows Server 2012 as an example to describe the steps.

Tools

You can use the following tools to check if the login issues are related to ports and security group configurations:

[Self-diagnosis](#)

[Security group \(port\) helper](#)

If the problem is indeed a security group configuration problem, use **Open all ports** in the [Security group \(port\) helper](#) to open related ports and try to log in again. If you still cannot log in after opening the ports, refer to the following for troubleshooting.

Troubleshooting

Checking network connectivity

You can use the Ping command to test network connectivity from your PC. You should run the test from computers in different network environments (such as different IP ranges or ISPs) to check whether it is a local network problem or a server problem.

1. Open the command line tool on your local computer.

Windows: Click **Start** -> **Run** and enter **cmd**. A Command Prompt window appears.

MacOS: Open a Terminal window.

2. Run the following command to test network connection.

```
ping + CVM_Instance_public_IP_address
```

For example, `ping 139.199.xxx.xxx` .

If you see results similar to what is shown in the following figure, your network connection to the CVM instance is normal. In this case, [check the remote desktop configuration](#).

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping 193.112.1.1

Pinging 193.112.1.1 with 32 bytes of data:
Reply from 193.112.1.1: bytes=32 time<1ms TTL=127
Reply from 193.112.1.1: bytes=32 time<1ms TTL=127
Reply from 193.112.1.1: bytes=32 time<1ms TTL=127
Reply from 193.112.1.1: bytes=32 time<1ms TTL=127

Ping statistics for 193.112.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

If **Request Timeout** appears, your network connection to the CVM instance is not working properly. In this case, refer to [Instance IP Address Ping Failure](#) for troubleshooting instructions.

3. Run the following command to check whether the remote port is open.

```
telnet CVM_instance_public_IP_address port_number
```

For example, `telnet 139.199.xxx.xxx 3389`, as shown in the following figure:

```
telnet 139.199.XXX.XXX 3389_
```

If you see a black screen with only the cursor, that indicates the port (3389) is open. For the next step, [check whether remote desktop service is enabled on the instance](#).

If the connection fails, as shown in the following figure, that means a network exception occurred. Check the corresponding part of the network.

```
C:\Users\Administrator>telnet 139.199.XXX.XXX 3389
Connecting To 139.199.XXX.XXX...Could not open connection to the host, on port 3
389: Connect failed
```


Checking remote desktop configuration

Logging in to the CVM instance using VNC

Note:

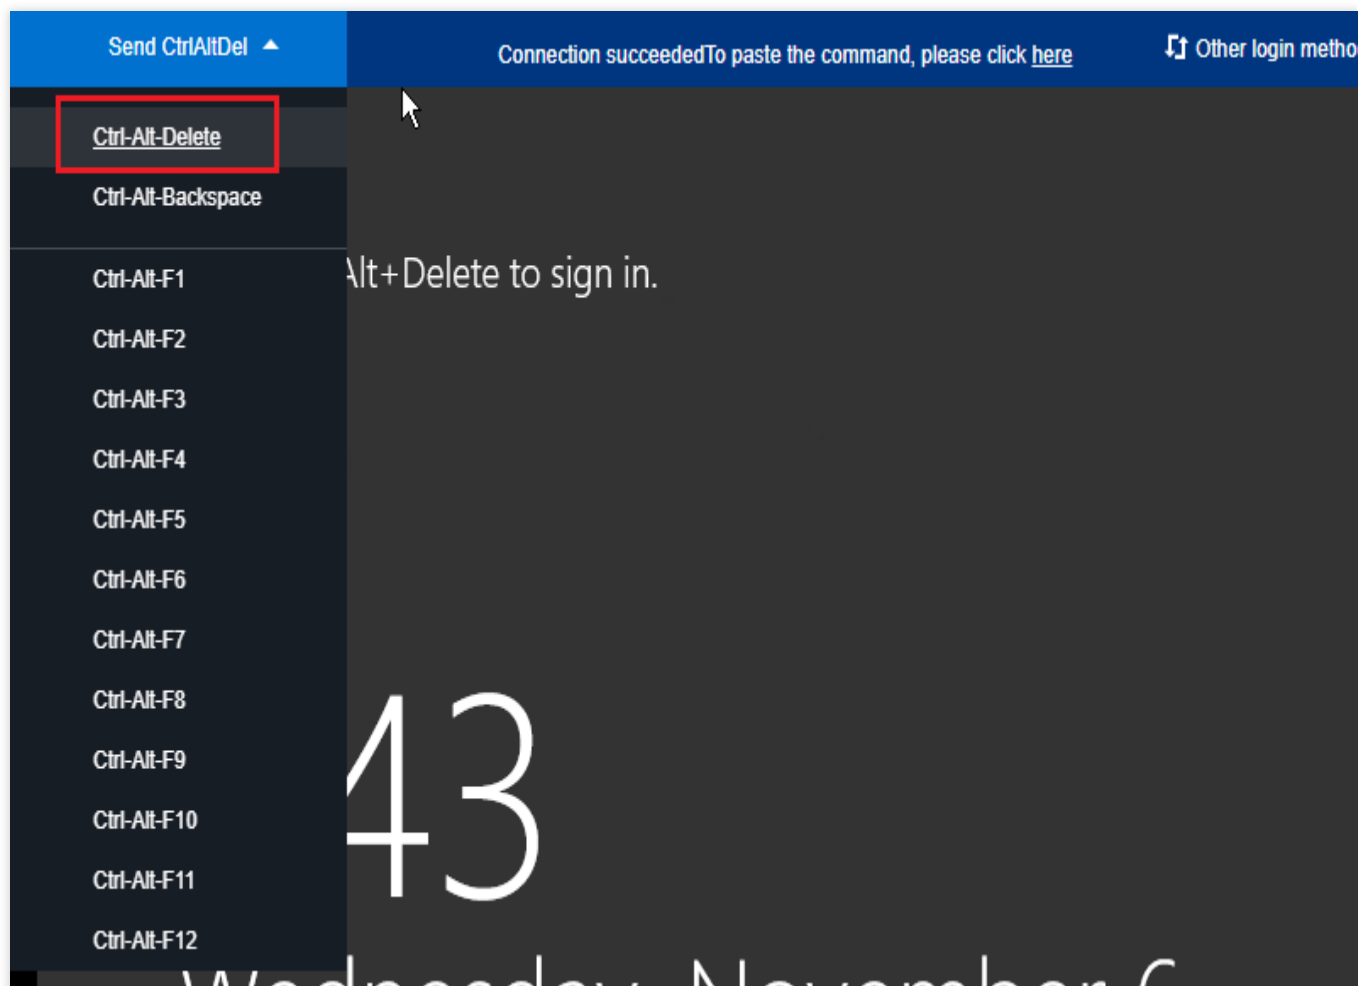
It is recommended that you use VNC to login only if the standard login methods fail.

1. Log in to the [CVM Console](#).
2. Select the desired CVM and click **Log In**, as shown in the following figure:



3. The **Log into Windows instance** page appears. Select **Alternative login methods (VNC)** and click **Log In Now** to log in to the CVM instance.

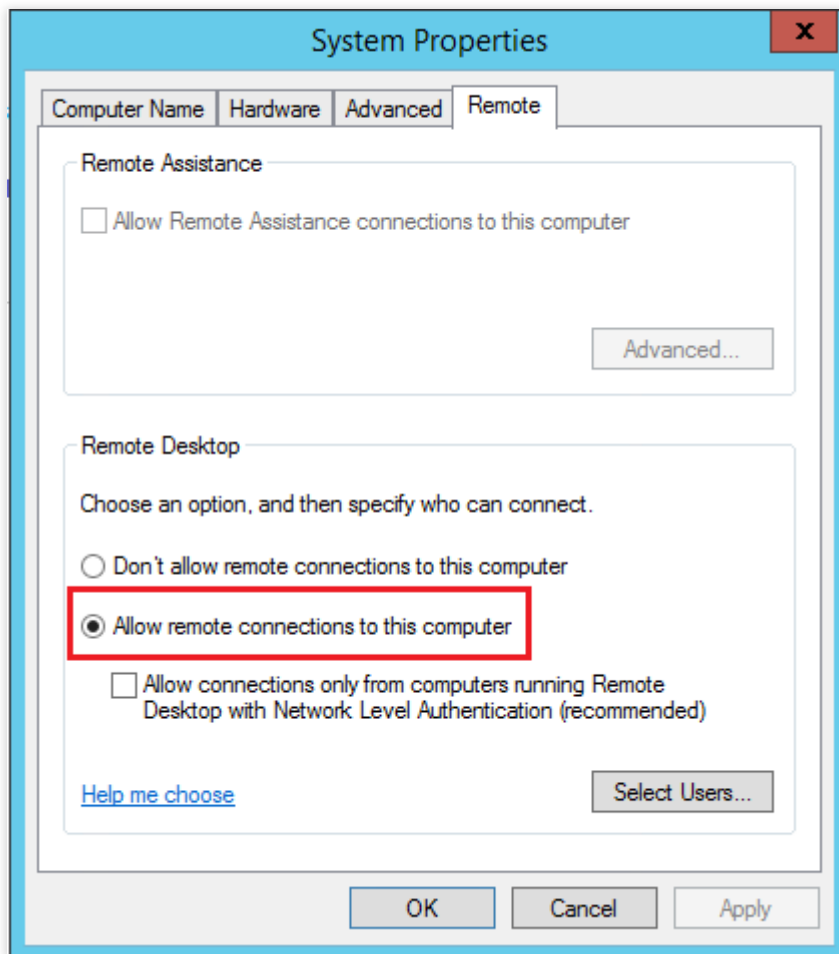
4. The log in page appears. Select **Send Ctrl-Alt-Del** in the top left corner and click **Ctrl-Alt-Delete** to enter the system login interface, as shown in the following figure:



Checking if remote desktop is enabled on the CVM instance

1. Log in to the CVM instance. Right click **This Computer** from the Desktop and select **Properties** to open the **System** window.

2. In the **System** window, select **Advanced System Configurations** to open the **System Properties** window.
3. In the **System Properties** window, select the **Remote** tab. Check whether **Allow remote connections to this computer** under **Remote Desktop** is selected, as shown in the following figure:



If it is selected, remote connection is enabled. Next, you should [check whether remote access ports are open](#). If it is cleared, select **Allow remote connections to this computer** and try to connect to the instance again.

Checking whether remote access ports are open

1. Log in to the CVM instance. Click




to open a **Windows PowerShell** window.

2. In the **Windows PowerShell** window, run the following command to check the status of remote desktop (by default, the remote desktop uses port 3389).

```
netstat -ant | findstr 3389
```

If you see results similar to what is shown in the following figure, the status is normal. You can try to [restart remote desktop](#) and connect to the instance again to check whether the connection is successful.



```
Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2014 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> netstat -ant | findstr 3389
TCP 0.0.0.0:3389 0.0.0.0:0 LISTENING InHost
TCP [REDACTED]:[REDACTED] [REDACTED]:[REDACTED] ESTABLISHED InHost
TCP [REDACTED]:[REDACTED] [REDACTED]:50709 ESTABLISHED InHost
TCP [::]:3389 [::]:0 LISTENING InHost
UDP 0.0.0.0:3389 *:*
UDP [::]:3389 *:*

PS C:\Users\Administrator>
```

If no connection is shown, remote desktop is not functioning properly. You can [check whether the remote desktop ports in registry are consistent](#).


Checking whether the remote desktop ports in registry are consistent

Caution:

This section describes how to check the values of **TCP PortNumber** and **RDP Tcp PortNumber**. They must be the same.

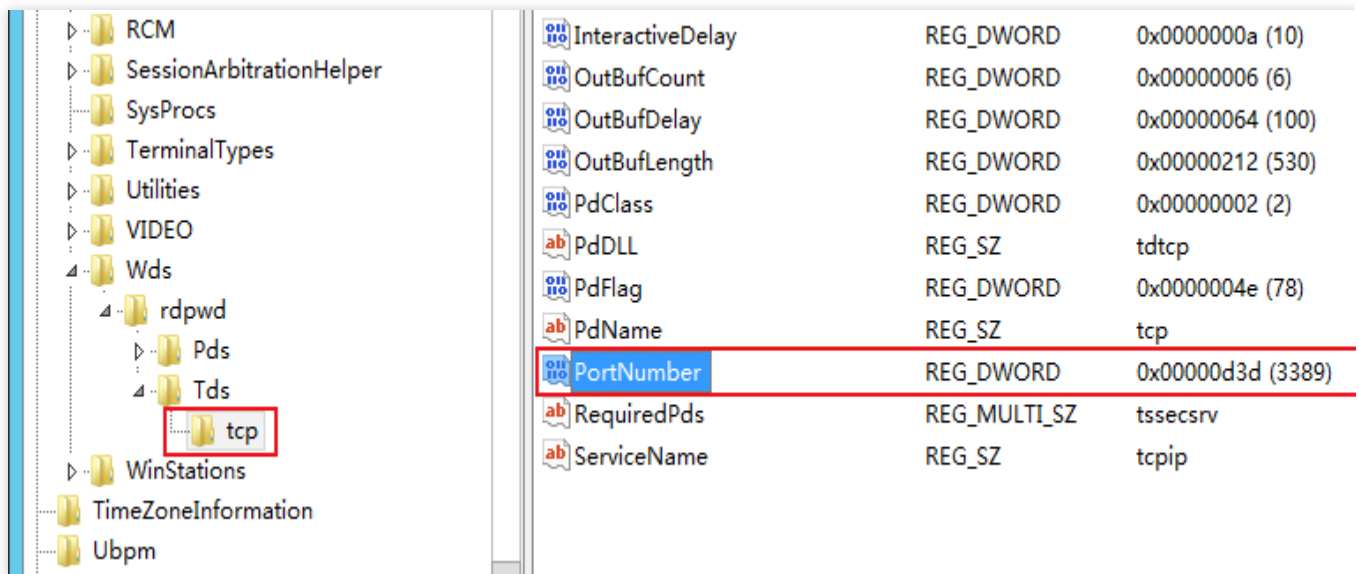
1. Log in to the CVM instance, click



 , and enter **regedit**. Press **Enter** to open the **Registry Editor** window.

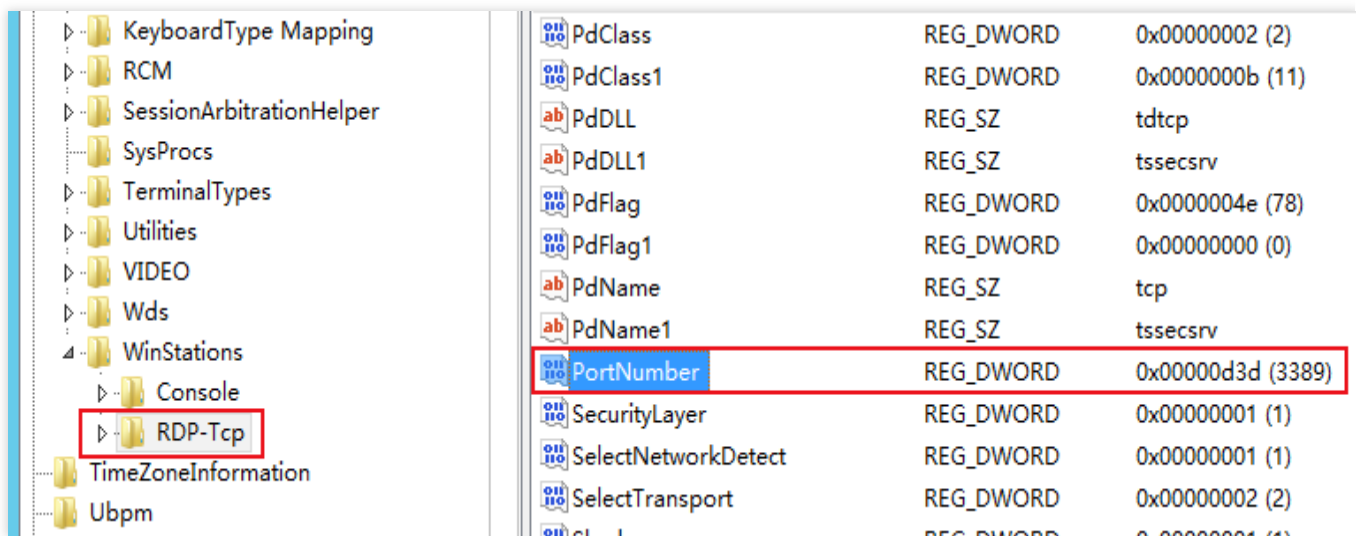
2. In the navigation pane on the left, expand the following directories: **HKEY_LOCAL_MACHINE -> SYSTEM -> CurrentControlSet -> Control -> Terminal Server -> Wds -> rdpwd -> Tds -> tcp**.

3. Locate the PortNumber in **tcp** and record the port number (3389 by default), as shown in the following figure:



4. In the navigation pane on the left, expand the following directories: **HKEY_LOCAL_MACHINE** -> **SYSTEM** -> **CurrentControlSet** -> **Control** -> **Terminal Server** -> **WinStations** -> **RDP-Tcp**.

5. Locate **PortNumber** in **RDP-Tcp** and check whether the **PortNumber** value in **RDP-Tcp** is the same as the one in **tcp**, as shown in the following figure:



If they are not the same, follow [Step 6](#).

If they are the same, [restart remote desktop](#).

6.

Double click **PortNumber** in **RDP-Tcp**.

7. In the dialog box that appears, modify **Value Data** to an unoccupied port number between 0 - 65535. Ensure **TCP PortNumber** and **RDP Tcp PortNumber** are the same, and click **OK**.

8. Restart the instance using the [CVM Console](#), and try to remotely connect to the instance again to check whether the connection is successful.

Restarting remote desktop

1. Log in to the CVM instance. Click

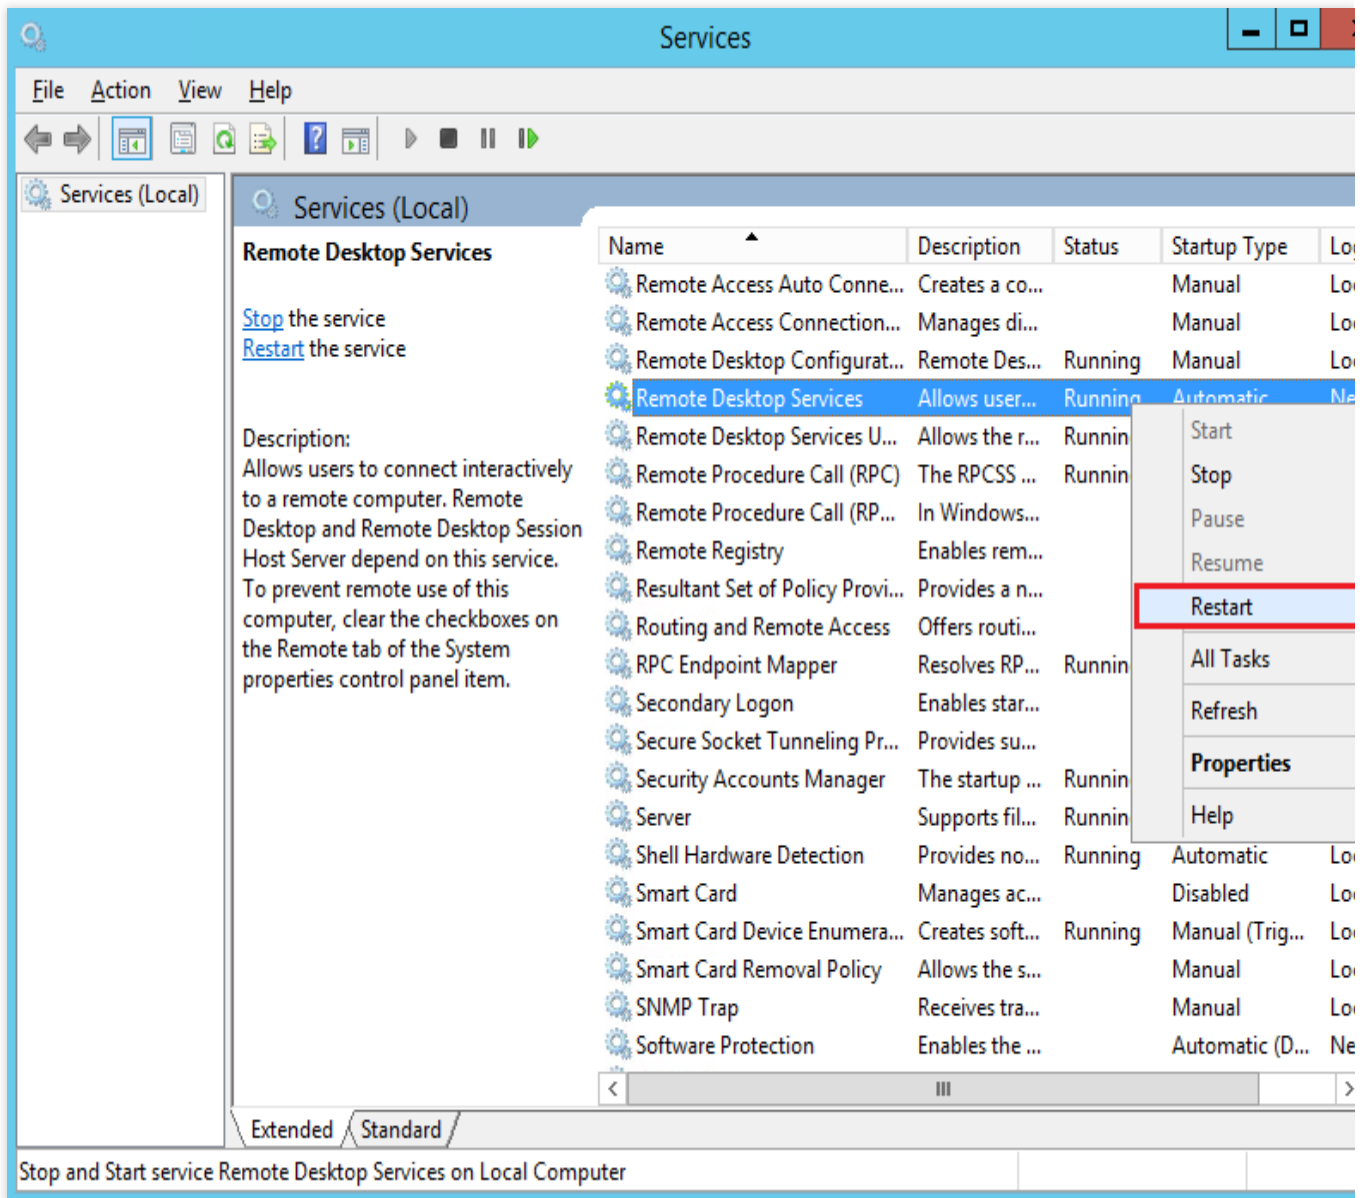


and select



. Enter **services.msc** and press **Enter** to open the **Services** window.

2. In the **Services** window, right-click **Remote Desktop Services** and select **Restart** to restart the remote desktop service, as shown in the following figure:



If All Else Fails

If you are still unable to remotely log in after performing the above-mentioned steps, [submit a ticket](#) for further assistance.

Linux Instance Login Failures

Linux Instance Login Failures

Last updated : 2024-01-06 17:32:18

This document describes possible causes of Linux instance login failures and troubleshooting methods, helping you detect, locate and resolve problems.

Troubleshooting the Issue

Using the Diagnosis Tool

Tencent Cloud provides a diagnosis tool to help you determine whether the problem is due to common bandwidth, firewall, or security group configurations issues. More than 70% of the problems can be located by this tool. You can locate problems that result in login failures based on the detected causes.

1. Click [Self-diagnose](#) to open the self-diagnosis tool.
2. Select the target CVM instance as prompted and click **Start Detection**.

Using TAT to send commands

You can use TAT to send commands to an instance for troubleshooting and problem locating. The directions are as follows:

1. Log in to the [CVM console](#) and click the target instance ID in the instance list.
2. On the instance details page, select the **Run Commands** tab and click **Execute command**.
3. In the **Execute command** pop-up window, select a command as needed. Click **Execute command** and view the result.

For example, enter `df -TH` and click **Execute command** to view the result of an instance without logging in to it.

For more information, see [TencentCloud Automation Tools](#).

Note:

If you cannot troubleshoot with the diagnosis tool, we recommend you [log in to the CVM instance via VNC](#) and follow the instructions.

Possible Cause

Common login failure reasons:

[SSH key problems](#)

[Password issues](#)

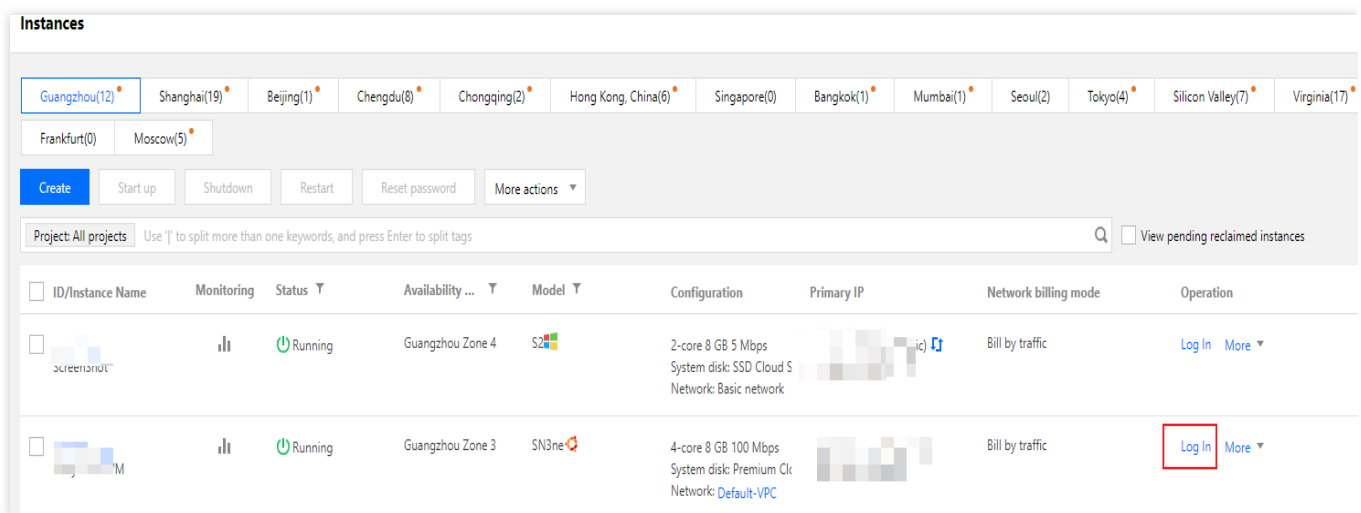
High bandwidth utilization
High server load
Improper security group rules

Troubleshooting

Logging in via VNC

If you cannot log in to the Linux instance by using Orcaterm or remote login software, you can try log in with VNC for troubleshooting.

1. Log in to the [CVM console](#).
2. On the **Instances** page, select the target instance and click **Log in**.



3. In the **Standard Login | Linux Instance** window that pops up, select **Login via VNC**.

Note:

If you forgot the password for the instance, you can reset it in the console. For more information, see [Resetting Instance Password](#).

4. Enter the username and password to login.

Login failure due to SSH issue

Problem: During [login to a Linux instance using SSH](#), a message appears indicating that the connection is unavailable or failed.

Steps: See [Unable to Use the SSH Method to Log in to a Linux Instance](#) to perform troubleshooting.

Login failure due to password issue

Problem: The login attempt failed because you forgot the password, entered an incorrect password, or failed to reset your password.

Solution: Reset the password for this instance in the [CVM console](#) and restart the instance.

Procedure: See [Resetting Instance Password](#) for the detailed procedure.

High bandwidth utilization

Problem: The self-diagnosis tool shows that bandwidth utilization is too high.

Procedure:

1. Log in to the instance by using [VNC login](#).
2. Check the bandwidth utilization of the instance and perform troubleshooting accordingly. For details, see [Login Failure Due to High Bandwidth Occupation](#).

High instance load

Problem: The self-diagnosis tool or Tencent Cloud Observability Platform shows that server CPU workload is too high, and the system is unable to perform remote connection or access is slow.

Possible cause: Viruses, trojans, third-party antivirus software, application exceptions, driver exceptions, and automatic updates of software on the backend may lead to high CPU utilization, causing CVM login failures or slow access.


Procedure:

1. Log in to the instance by using [VNC login](#).
2. In **Task manager**, find the process with a high load. For more information, see [Failing to log in to a Linux CVM due to high CPU and memory usage](#).

Improper security group rules

Problems: The self-diagnosis tool shows that the security group rule configuration is improper, leading to login failures.

Procedure: Troubleshoot with the [security group \(port\) verification tool](#).



If the problem is caused by a port issue of the security group, you can use the **Open all ports** feature to open all ports.

Testing Details ×

Protocol	Port	Direction	Policy	Effects
TCP	3389	Inbound	Open	None
TCP	22	Inbound	Open	None
TCP	443	Inbound	Open	None
TCP	80	Inbound	Open	None
TCP	21	Inbound	Not opened ⓘ	Unable to access FTP
TCP	20	Inbound	Not opened ⓘ	Unable to access FTP
ICMP	0	Inbound	Open	None
ALL	ALL	Outbound	Open	None

Open all ports Cancel

To define a custom rule for the security group, see [Adding Security Group Rules](#).

Other Solutions

If you still cannot connect to the Linux instance, save your self-diagnosis results and [submit a ticket](#) for assistance.

Unable to Log in to a Linux Instance via SSH Key

Last updated : 2024-01-06 17:32:18

Note:

This document is a general instructional guide only for reference.

Perform the file operations with caution. You can create a snapshot or use other methods to back up data when necessary.

Error Description

During [login to a Linux instance using SSH key](#), a message indicating that the connection is unavailable or failed appears.

Troubleshooting

Check below for the solutions for common SSH login errors.

SSH login error "User root not allowed because not listed in AllowUsers"

Problem

Login to a Linux instance via SSH key fails, and a message similar to the following appears in the `secure` log of the client or server:

Permission denied, please try again.

User test from 192.X.X.1 not allowed because not listed in AllowUsers.

User test from 192.X.X.1 not allowed because listed in DenyUsers.

User root from 192.X.X.1 not allowed because a group is listed in DenyGroups.

User test from 192.X.X.1 not allowed because none of user's groups are listed in AllowGroups.

Cause

The login is restricted by the user login control parameters.

AllowUsers: only users listed in this parameter are allowed to log in.

DenyUsers: users listed in this parameter are denied to log in.

AllowGroups: only user groups listed in this parameter are allowed to log in.

DenyGroups: user groups listed in this parameter are denied to log in.

Note:

The **Deny** policy takes priority over the **Allow** policy.

Solutions

1. Open the SSH config file `sshd_config` as instructed in [Steps](#).
2. Delete the user login control parameters and restart the SSH service.

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Press **i** to enter the edit mode. Locate and delete the following configurations, or add a pound sign (`#`) at the beginning of each line to comment them out.

```
AllowUsers root test
DenyUsers test
DenyGroups test
AllowGroups root
```

4. Press **Esc** to exit the edit mode, and enter **:wq** to save the modification.
5. Run the following commands based on the operating system to restart the SSH service.

CentOS

```
systemctl restart sshd.service
```

Ubuntu

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH login error "Disconnected:No supported authentication methods available"

Problem

When logging in via SSH key, the following error message appears:

```
Permission denied (publickey,gssapi-keyex,gssapi-with-mic).
sshd[10826]: Connection closed by xxx.xxx.xxx.xxx.
Disconnected:No supported authentication methods available.
```

Cause

SSH service modifies the `PasswordAuthentication` parameter and disables the password login.

Solutions

1. Open the SSH config file `sshd_config` as instructed in [Steps](#).
2. Modify the `PasswordAuthentication` parameter and restart the SSH service.

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Press **i** to enter the edit mode and change `PasswordAuthentication no` to `PasswordAuthentication yes`.
4. Press **Esc** to exit the edit mode, and enter **:wq** to save the modification.
5. Run the following commands based on the operating system to restart the SSH service.

CentOS

```
systemctl restart sshd.service
```

Ubuntu

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH login error "ssh_exchange_identification: read: Connection reset by peer"

Problem

When logging in via SSH key, the error message "ssh_exchange_identification: read: Connection reset by peer" or as shown below appears:

```
"ssh_exchange_identification: Connection closed by remote host"
```

```
"kex_exchange_identification: read: Connection reset by peer"
```

```
"kex_exchange_identification: Connection closed by remote host"
```

Cause

The common causes are as follows:

The connection is blocked by the local access control policy.

The firewall rules are modified by anti-intrusion software, such as Fail2ban, denyhost, etc.

Reaches the maximum number of connections configured in sshd

Local network problem

Solutions

Check the access policy, firewall rules, sshd configuration, and network environment as instructed in [Steps](#).

Steps

Checking and modifying the access policy settings

In Linux, the allowed and denied access policy are respectively set in the `/etc/hosts.allow` and `/etc/hosts.deny` files. You can set the trusted hosts in the `hosts.allow` file, and deny all other hosts in the `hosts.deny` file. The Deny policy can be set as follows:

```
in.sshd:ALL                                # Deny all SSH connections
in.sshd:218.64.87.0/255.255.255.128        # Deny SSH connections ranging from
218.64.87.0 to -127.
ALL:ALL                                    # Deny all TCP connections
```

[Log in to the Linux instance via VNC](#) and check the `/etc/hosts.deny` and `/etc/hosts.allow` files:

Correct the configurations if needed. The modification takes effect immediately.

If the configuration is correct, proceed to the next step.

Note:

If no access policy is configured, both files will be empty by default and all connections are allowed.

Checking the iptables firewall rules

Check whether the iptables firewall rules are modified, for example, by the anti-intrusion software (such as Fail2ban, denyhost, etc.). Run the following command to check whether the firewall denies SSH connections.

```
sudo iptables -L --line-number
```

If yes, please modify the allowlist settings.

If no, proceed to the next step.

Checking and modifying the sshd configuration

1. Run the following command to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

2. Check the `MaxStartups` value, which specifies the maximum number of connections allowed. If many connections are required to establish in a short period, adjust the value as needed.

Follow the steps below to modify the value:

2.1.1 Press **i** to enter the edit mode. After the modification is complete, press **Esc** to exit the edit mode and enter **:wq** to save the modification.

Note:

This parameter specifies the maximum number of unauthenticated concurrent connections the SSH daemon allows. The default value is `MaxStartups 10:30:100`, which means to allow the first ten unauthenticated connections, refuse connection attempts with a probability of 30%, and reject all new connections when there are 100 connections.

2.1.2 Run the following command to restart the sshd service.

```
service sshd restart
```

If no modification is required, proceed to the next step.

Testing the network environment

1. Check whether a [private IP](#) is used for login.

If yes, use a [public IP](#) and try again.

If not, proceed to the next step.

2. Test the connection by using other network environments.

If the connection is normal, restart the instance and log in to it via VNC.

If there is a connection error, solve it according to the test result.

If the SSH login error persists, it may be caused by kernel exceptions or other unknown reasons. In this case, please [submit a ticket](#).

SSH login error "Permission denied, please try again"

Problem

When the root user logs in to a Linux instance via SSH key, the "Permission denied, please try again" error message appears.

Cause

This is because that the SELinux service is enabled, or the SSH service modifies the `PermitRootLogin` configuration.

Solutions

Check the configuration of SELinux service and the `PermitRootLogin` parameter in `sshd_config` as instructed in [Steps](#).

Steps

Checking and disabling the SELinux service

1. [Log in to the Linux instance via VNC](#).

2. Run the following command to check the current SELinux service status.

```
/usr/sbin/sestatus -v
```

If `enabled` as shown below is returned, the service is enabled. The `disabled` response indicates that the service is disabled.

```
SELinux status:          enabled
```

3. Disable the SELinux service temporarily or permanently as needed.

Disable the SELinux service temporarily

Perform the following command to temporarily disable the SELinux service. The change takes effect immediately without needing to restart the system or instance.

```
setenforce 0
```

Disable the SELinux service permanently

Run the following command to disable the SELinux service.

```
sed -i 's/SELINUX=enforcing/SELINUX=disabled/' /etc/selinux/config
```

Note:

This command is only applicable to the SELinux service in the `enforcing` status.

After running the command, restart the system or instance for the modification to take effect.

Checking and modifying the sshd configuration

1. [Log in to the Linux instance via VNC](#).

2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Press `i` to enter the edit mode and change `PermitRootLogin no` to `PermitRootLogin yes`.

Note:

If this parameter is not configured in `sshd_config`, the root user is allowed to log in by default.

This parameter only affects the SSH key login of the root user, who can log in to the instance normally using other methods.

4. Press **Esc** to exit the edit mode, and enter `:wq` to save the modification.

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

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH login error "Too many authentication failures for root"

Problem

The error message "Too many authentication failures for root" is returned and the connection is interrupted after many failed attempts to enter the password to login via SSH key.

Cause

I was requested to reset the SSH service password after consecutive failed password entry.

Solutions

1. Open the SSH config file `sshd_config` as instructed in [Steps](#).
2. Check and modify the `MaxAuthTries` parameter for the password reset policy, and restart the SSH service.

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Check for the configuration similar to what is shown below.

```
MaxAuthTries 5
```

Note:

This parameter is not enabled by default. It specifies the number of consecutive incorrect password attempts for each SSH key login, and displays the error message. However, the account will not be locked, which can be used for another SSH key login.

Determine whether a modification is required. If so, we recommend you back up the `sshd_config` configuration file.

4. Press **I** to enter the edit mode, and modify the following configuration or add a pound sign (`#`) at the beginning of the line to comment it out.

```
MaxAuthTries <number of incorrect password attempts allowed>
```

5. Press **Esc** to exit the edit mode, and enter **:wq** to save the modification.
6. Run the following command to restart the SSH service.

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH startup error "error while loading shared libraries"

Problem

When the SSH service is started in a Linux instance, an error message similar to the following is displayed in the `secure` log or directly returned:

"error while loading shared libraries: libcrypto.so.10: cannot open shared object file: No such file or directory"

"PAM unable to dlopen(/usr/lib64/security/pam_tally.so): /usr/lib64/security/pam_tally.so: cannot open shared object file: No such file or directory"

Cause

This is because that the relevant system library files the SSH service depends on are lost or have an exception, such as the permission configuration exception.

Solutions

Check and fix the system library files as instructed in [Steps](#).

Steps

Note:

The following example guides you through resolving a `libcrypto.so.1.0` library file exception.

Obtaining the library file

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the `libcrypto.so.1.0` library file information.

```
ll /usr/lib64/libcrypto.so.1.0
```

The information similar to what is shown below is returned, indicating that `/usr/lib64/libcrypto.so.1.0` is the soft link of the `libcrypto.so.1.0.2k` library file.

```
lrwxrwxrwx 1 root root 19 Jan 19 2021 /usr/lib64/libcrypto.so.1.0 ->
libcrypto.so.1.0.2k
```

3. Run the following command to view the `libcrypto.so.1.0.2k` library file information.

```
ll /usr/lib64/libcrypto.so.1.0.2k
```

Information similar to the following is returned:

```
-rwxr-xr-x 1 root root 2520768 Dec 17 2020 /usr/lib64/libcrypto.so.1.0.2k
```

4. Note down the path, permission, group and other information of a normal library file, and try the following.

[Find and replace the library file](#)

[Upload an external file](#)

[Recover with snapshots](#)

Finding and replacing the library file

1. Run the following command to find the `libcrypto.so.1.0.2k` file.

```
find / -name libcrypto.so.1.0.2k
```

2. Run the following command to copy the library file to a normal directory.

```
cp <absolute path of the library file obtained in the step 1>
/usr/lib64/libcrypto.so.1.0.2k
```

3. Run the following commands to modify the file permission, owner, and group.

```
chmod 755 /usr/lib64/libcrypto.so.1.0.2k  
  
chown root:root /usr/lib64/libcrypto.so.1.0.2k
```

4. Run the following commands to create a soft link.

```
ln -s /usr/lib64/libcrypto.so.1.0.2k /usr/lib64/libcrypto.so.1.0
```

5. Run the following command to start the SSH service.

```
service sshd start
```

Uploading an external file

1. Upload the `libcrypto.so.1.0.2k` library file of a normal CVM to the `\\tmp` directory of the target CVM using FTP.

Note:

`\\tmp` is used in this example. You can replace with the actual directory.

2. Run the following command to copy the library file to a normal directory.

```
cp /tmp/libcrypto.so.1.0.2k /usr/lib64/libcrypto.so.1.0.2k
```

3. Run the following commands to modify the file permission, owner, and group.

```
chmod 755 /usr/lib64/libcrypto.so.1.0.2k  
  
chown root:root /usr/lib64/libcrypto.so.1.0.2k
```

4. Run the following commands to create a soft link.

```
ln -s /usr/lib64/libcrypto.so.1.0.2k /usr/lib64/libcrypto.so.1.0
```

5. Run the following command to start the SSH service.

```
service sshd start
```

Recovering with snapshot

You can roll back the system disk snapshot of the instance to recover the library file. For more information, please see [Rolling Back Snapshots](#).

Note:

Note that snapshot rollback will cause loss of data stored after the snapshot creation.

We recommend you roll back to snapshots one by one from the latest to the earliest till the SSH service resumes. If the SSH service still cannot run properly after the rollback, the system at that point in time has been abnormal.

SSH service startup error "fatal: Cannot bind any address"

Problem

When the SSH service is started in a Linux instance, an error message similar to the following is directly returned or appears in the `secure` log:

```
FAILED.  
fatal: Cannot bind any address.  
address family must be specified before ListenAddress.
```

Cause

Incorrect configuration of `AddressFamily`. This parameter specifies the protocol suite used at runtime. If only IPv6 is configured here, but the IPv6 is not enabled or invalidly configured in the system, this problem may occur.

Solutions

1. Open the SSH config file `sshd_config` as instructed in [Steps](#).
2. Modify the `AddressFamily` parameter and restart the SSH service.

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Check for the configuration similar to what is shown below.

```
AddressFamily inet6
```

The common parameters are described as follows:

inet: IPv4 protocol suite, the default value

inet6: IPv6 protocol suite

any: both IPv4 and IPv6 protocol suites

4. Press `i` to enter the edit mode, and modify the configuration as follows or add a pound sign (`#`) at the beginning of the line to comment it out.

```
AddressFamily inet
```

Note:

The `AddressFamily` parameter takes effect only after being configured before `ListenAddress` .

5. Press **Esc** to exit the edit mode, and enter **:wq** to save the modification.

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

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH service startup error "Bad configuration options"

Problem

When the SSH service is started in a Linux instance, an error message similar to the following is directly returned or appears in the `secure` log:

```
/etc/ssh/sshd_config: line 2: Bad configuration options:\\\\\\  
/etc/ssh/sshd_config: terminating, 1 bad configuration options
```

Cause

This is caused by the incorrect encoding or configuration of the configuration file.

Solutions

Fix the `sshd_config` configuration file as instructed below.

[Modifying the configuration file according to the error message](#)

[Uploading an external file](#)

[Reinstalling the SSH service](#)

[Using the snapshot rollback for recovery](#)

Steps

Modifying the configuration file according to the error message

If the error message tells the incorrect configuration, modify `/etc/ssh/sshd_config` with VIM editor by referring to the correct configuration file of another instance.

Uploading an external file

1. Upload the `/etc/ssh/sshd_config` library file of a normal CVM to the `\\tmp` directory of the target CVM using FTP.

Note:

`\\tmp` is used in this example. You can replace with the actual directory.

2. Run the following command to copy the library file to a normal directory.

```
cp /tmp/sshd_config /etc/ssh/sshd_config
```

3. Run the following commands to modify the file permission, owner, and group.

```
chmod 600 /etc/ssh/sshd_config  
  
chown root:root /etc/ssh/sshd_config
```

4. Run the following command to start the SSH service.

```
service sshd start
```

Reinstalling the SSH service

1. [Log in to the Linux instance via VNC](#).

2. Run the following command to uninstall the SSH service.

```
rpm -e openssh-server
```

3. Run the following command to install the SSH service.

```
yum install openssh-server
```

4. Run the following command to start the SSH service.

```
service sshd start
```

Recovering with snapshot

You can roll back the system disk snapshot of the instance to recover the library file. For more information, please see [Rolling Back Snapshots](#).

Caution:

Note that snapshot rollback will cause loss of data stored after the snapshot creation.

We recommend you roll back to snapshots one by one from the latest to the earliest till the SSH service resumes. If the SSH service still cannot run properly after the rollback, the system at that point in time has been abnormal.

Enabling UseDNS for SSH led to slow login or data transfer via SSH key

Problem

Login or data transfer via SSH key over the private network in a Linux instance is slow, and after switch to the private network, the problem persists.

Cause

This may be because that the UseDNS feature of the SSH service is enabled. This feature is a security enhancement, which is not enabled by default. After it is enabled, the server will first perform a reverse DNS PTR record query based on the client IP to get the client's host name, then perform a forward DNS A record query based on the client's host

name, and finally compare whether the obtained IP is the same as the original IP so as to prevent client frauds. In general, the client uses a dynamic IP, and there is no corresponding PTR record; therefore, after this feature is enabled, no comparison can be made, and the relevant query operations increase the delay, eventually leading to slow client connection.

Solutions

1. Open the SSH config file `sshd_config` as instructed in [Steps](#).
2. Check and modify the UseDNS configuration and restart the SSH service.

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to open the `sshd_config` configuration file with VIM editor.

```
vim /etc/ssh/sshd_config
```

3. Check for the following configuration:

```
UseDNS yes
```

4. Press `i` to enter the edit mode, and delete the configuration or add a pound sign (`#`) at the beginning of the line to comment it out.
5. Run the following command to restart the SSH service.

```
service sshd restart
```

Then you can [log in to the Linux CVM instance via SSH key](#) normally.

SSH login error "No supported key exchange algorithms"

Problem

Login to a Linux instance via SSH key fails, and an error message similar to the following may appear in the

`secure` log of the client or server:

Read from socket failed: Connection reset by peer.

Connection closed by 192.X.X.1.

sshd error: could not load host key.

fatal: No supported key exchange algorithms [preauth].

DSA host key for 192.X.X.1 has changed and you have requested strict checking.

Host key verification failed.

ssh_exchange_identification: read: Connection reset by peer.

Cause

Usually, this is because that the relevant key file of the SSH service is exceptional, so the sshd daemon cannot load the correct SSH host key. Common causes include the following:

The relevant key file is exceptional; for example, the file is corrupted, deleted, or tampered with.

The permission configuration of the relevant key file is exceptional, so it cannot be read correctly.

Solutions

Check and modify the configuration as instructed below.

[Checking and modifying file permission](#)

[Checking and modifying file validity](#)

Steps

Checking and modifying file permission

The SSH service will check the permission of the relevant key file. For example, the default permission of a private key file is `600`, and if other permissions such as `777` are configured, then other users also have permissions to read or modify the file. In this case, the SSH service will deem that the configuration involves security risks, which causes client connection failures. The troubleshooting process is as follows:

1. [Log in to the Linux instance via VNC](#).
2. Run the following commands to restore the default permission of the relevant file.

```
cd /etc/ssh/

chmod 600 ssh_host_*

chmod 644 *.pub
```

3. Run the `ll` command to view the file permission. If the following result is returned, the file permission is normal.

```
total 156
-rw----- . 1 root root 125811 Nov 23 2013 moduli
-rw-r--r-- . 1 root root 2047 Nov 23 2013 ssh_config
-rw----- 1 root root 3639 May 16 11:43 sshd_config
-rw----- 1 root root 668 May 20 23:31 ssh_host_dsa_key
-rw-r--r-- 1 root root 590 May 20 23:31 ssh_host_dsa_key.pub
-rw----- 1 root root 963 May 20 23:31 ssh_host_key
-rw-r--r-- 1 root root 627 May 20 23:31 ssh_host_key.pub
-rw----- 1 root root 1675 May 20 23:31 ssh_host_rsa_key
-rw-r--r-- 1 root root 382 May 20 23:31 ssh_host_rsa_key.pub
```

Checking and modifying file validity

1. The SSH service will automatically rebuild a lost key file upon start. Run the following commands to check for the `ssh_host_*` file.

```
cd /etc/ssh/
```

```
ll
```

If the following result is returned, the `ssh_host_*` file exists.

```
total 156
-rw-----. 1 root root 125811 Nov 23 2013 moduli
-rw-r--r--. 1 root root 2047 Nov 23 2013 ssh_config
-rw----- 1 root root 3639 May 16 11:43 sshd_config
-rw----- 1 root root 672 May 20 23:08 ssh_host_dsa_key
-rw-r--r-- 1 root root 590 May 20 23:08 ssh_host_dsa_key.pub
-rw----- 1 root root 963 May 20 23:08 ssh_host_key
-rw-r--r-- 1 root root 627 May 20 23:08 ssh_host_key.pub
-rw----- 1 root root 1675 May 20 23:08 ssh_host_rsa_key
-rw-r--r-- 1 root root 382 May 20 23:08 ssh_host_rsa_key.pub
```

2. Run the following command to delete the relevant file.

```
rm -rf ssh_host_*
```

On Ubuntu or Debian, run the following command to delete the relevant file.

```
sudo rm -r /etc/ssh/ssh*key
```

3. Run the `ll` command to check whether the file has been deleted successfully. If the following result is returned, it has been deleted successfully.

```
total 132
-rw-----. 1 root root 125811 Nov 23 2013 moduli
-rw-r--r--. 1 root root 2047 Nov 23 2013 ssh_config
-rw----- 1 root root 3639 May 16 11:43 sshd_config
```

4. Run the following command to restart the SSH service, and the relevant file will be generated automatically.

```
service sshd restart
```

On Ubuntu or Debian, run the following command to restart the SSH service.

```
sudo dpkg-reconfigure openssh-server
```

5. Run the `ll` command to check whether the `ssh_host_*` file has been generated successfully. If the following result is returned, it has been generated successfully.


```
total 156
-rw-----. 1 root root 125811 Nov 23 2013 moduli
-rw-r--r--. 1 root root 2047 Nov 23 2013 ssh_config
-rw----- 1 root root 3639 May 16 11:43 sshd_config
-rw----- 1 root root 668 May 20 23:16 ssh_host_dsa_key
-rw-r--r-- 1 root root 590 May 20 23:16 ssh_host_dsa_key.pub
-rw----- 1 root root 963 May 20 23:16 ssh_host_key
-rw-r--r-- 1 root root 627 May 20 23:16 ssh_host_key.pub
-rw----- 1 root root 1671 May 20 23:16 ssh_host_rsa_key
-rw-r--r-- 1 root root 382 May 20 23:16 ssh_host_rsa_key.pub
```

Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH service startup error "must be owned by root and not group or word-writable"

Problem

When the SSH service is started on a Linux instance, the "must be owned by root and not group or word-writable" error message is returned.

Cause

Usually, this is because that the relevant permission or grouping of the SSH service is exceptional. For security reasons, the service has certain requirements for the permission configuration and grouping of directories or files.

Solutions

Check and modify the incorrect configuration as instructed below.

[Checking and repairing the configuration of `/var/empty/sshd` directory](#)

[Checking and repairing the configuration of `/etc/securetty` file](#)

Steps

Note:

This document uses CentOS 7.6 as an example. Please proceed based on the actual business conditions.

Checking and repairing the configuration of `/var/empty/sshd` directory

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the permission configuration of the `/var/empty/sshd` directory.

```
ll -d /var/empty/sshd/
```

The following is the default permission configuration.

```
drwx--x--x. 2 root root 4096 Aug 9 2019 /var/empty/sshd/
```

3. Compare the returned result with the default permission configuration, and if they are different, run the following command to restore the default configuration.

Note:

The `/var/empty/sshd` directory has the permission `711` and is a root user in the root group by default.

```
chown -R root:root /var/empty/sshd
```

```
chmod -R 711 /var/empty/sshd
```

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

```
systemctl restart sshd.service
```

Checking and repairing the configuration of `/etc/securetty` file

1. Run the following command to view the permission configuration of the `/etc/securetty` file.

```
ll /etc/securetty
```

The following is the default permission configuration.

```
-rw-----. 1 root root 255 Aug  5 2020 /etc/securetty
```

2. Compare the returned result with the default permission configuration, and if they are different, run the following command to restore the default configuration.

Note:

The `/etc/securetty` file has the permission `600` and is a root user in the root group by default.

```
chown root:root /etc/securetty
```

```
chmod 600 /etc/securetty
```

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

```
systemctl restart sshd.service
```

Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH login error "Host key verification failed"**Problem**

Login to a Linux instance via SSH key fails, and the following error message appears:

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@ WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED! @
```

```

@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!
Someone could be eavesdropping on you right now (man-in-the-middle attack)!
It is also possible that the RSA host key has just been changed.
The fingerprint for the RSA key sent by the remote host is
ae:6e:68:4c:97:a6:91:81:11:38:8d:64:ff:92:13:50.
Please contact your system administrator.
Add correct host key in /root/.ssh/known_hosts to get rid of this message.
Offending key in /root/.ssh/known_hosts:70
RSA host key for x.x.x.x has changed and you have requested strict checking.
Host key verification failed.
```

If the client runs on Windows, the following error message will usually appear during an SSH client connection:

```
The host key of `X.X.X.X` (port: XX) is not the same as the one saved in the
host key database. The host key has been changed or someone is attempting to
eavesdrop this connection. If you are not sure, we recommend you cancel this
connection.
```

Cause

After the Linux instance is reinstalled, changes in the account information lead to the change in the SSH public key, so the public key fingerprint saved on the client is not the same as that saved on the server, resulting in SSH authentication failure and denied login.

Solutions

Fix the problem as instructed below based on the operating system of the client.

[Windows client](#)

[Linux client](#)

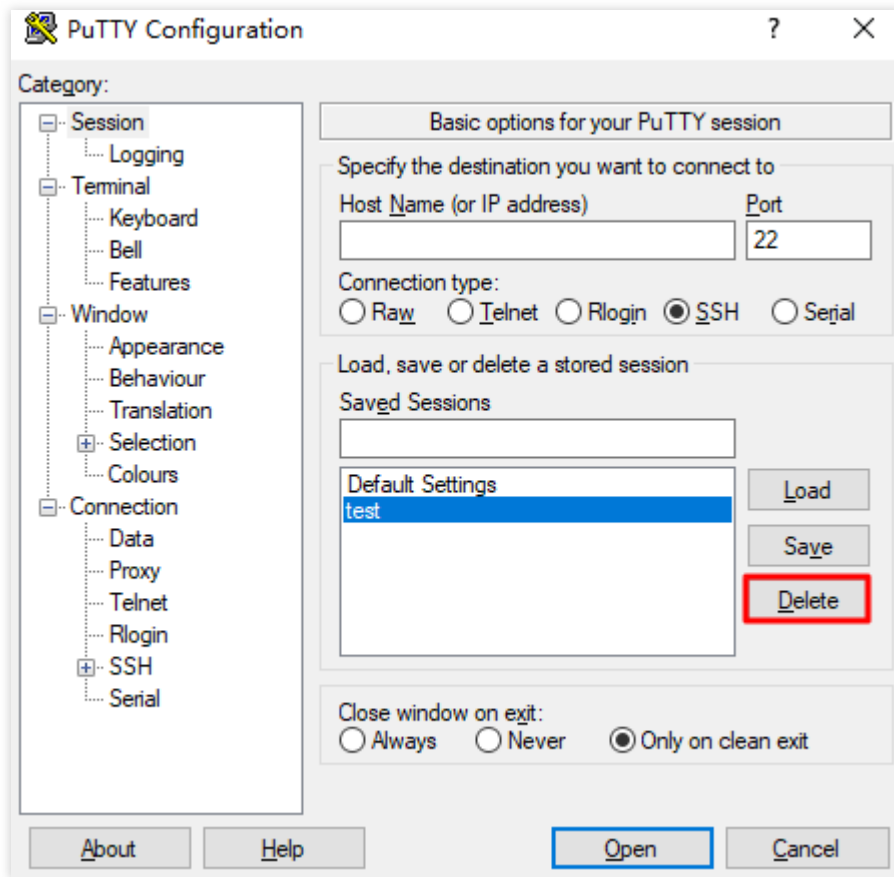
Steps

Windows client

Note:

This document uses PuTTY as an SSH client example. Please proceed based on the actual conditions.

1. Start PuTTY.
2. On the login page, select a session and click **Delete** to delete it as shown below:



3. Log in to the instance with the username and password again as instructed in [Logging in to Linux Instances via Remote Login Tools](#). After confirming that the new public key fingerprint is saved, you can log in successfully.

Linux client

Note:

This document uses CentOS 6.5 as an operating system example of the Linux instance. As the steps may vary by operating system, please proceed based on the actual conditions.

1. [Log in to the Linux instance via VNC](#).

2. Run the following command to open the `known_hosts` file of the corresponding account.

```
vi ~/.ssh/known_hosts
```

3. Press **i** to enter the edit mode and delete the entry of the Linux instance IP similar to the following:

```
1.14.xxx.xx  
skowcenw96a/pxka32sa....  
dsaprgpck2wa22mvi332ueddw...
```

4. Press **Esc** and enter **:wq** to save and close the file.

5. Reconnect to the Linux instance as instructed in [Logging into Linux Instance via SSH Key](#). After confirming that the new public key fingerprint is saved, you can log in successfully.

SSH login error "pam_listfile(sshd:auth): Refused user root for service sshd"

Problem

Login to a Linux instance via SSH key fails with the correct password entered. When this problem occurs, login in the console and via SSH key may both fail, or only one of them can succeed. An error message similar to the following appears in the `secure` log:

```
sshd[1199]: pam_listfile(sshd:auth): Refused user root for service sshd
sshd[1199]: Failed password for root from 192.X.X.1 port 22 ssh2
sshd[1204]: Connection closed by 192.X.X.2
```

Cause

The relevant access control policy of the PAM module (`pam_listfile.so`) causes the user login to fail.

Overview of PAM module

Pluggable Authentication Module (PAM) is an authentication mechanism proposed by Sun. It provides some dynamic link libraries and a set of unified APIs to separate a system service from its authentication method. This allows the system admin to flexibly configure different authentication methods for different services as needed without modifying service programs and add new authentication methods to the system easily.

Each application with the PAM module enabled has a configuration file named after it in the `/etc/pam.d` directory; for example, the configuration file of the `login` command is `/etc/pam.d/login`, where you can configure specific policies.

Solutions

Check and repair the PAM module as instructed in [Steps](#).

Steps

Note:

The steps in this document are for CentOS 6.5 as an example. As the steps may vary by operating system, please proceed based on the actual conditions.

1. [Log in to the Linux instance via VNC](#).
2. Use the `cat` command to view the corresponding PAM configuration file as described below:

File	Feature Description
<code>/etc/pam.d/login</code>	Configuration file of the console (VNC)
<code>/etc/pam.d/sshd</code>	Configuration file of SSH login
<code>/etc/pam.d/system-auth</code>	Global configuration file of the system

3. Check for the configuration similar to what is shown below.

```
auth required pam_listfile.so item=user sense=allow file=/etc/ssh/whitelist
onerr=fail
```

Description:

item: set the object type for access control. Valid values: tty, user, rhost, ruser, group, shell.

sense: find the control method that meets the conditions in the configuration file. Valid values: allow (allowlist), deny (blocklist).

file: specify the full path name of the configuration file.

onerr: define the default value to be returned when an error occurs; for example, the configuration file cannot be opened.

4. Use Vim to delete the policy configuration or add a pound sign (#) at the beginning of the line to comment it out.

Note:

The relevant policy configuration can improve the server security. We recommend you back up the configuration before modifying it as needed.

```
# auth required pam_listfile.so item=user sense=allow file=/etc/ssh/whitelist
onerr=fail
```

5. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH login error "requirement "uid >= 1000" not met by user "root""

Problem

Login to a Linux instance via SSH key fails with the correct username and password entered. When this problem occurs, login in the console and via SSH key may both fail, or only one of them can succeed. An error message similar to the following appears in the `secure` log:

```
pam_succeed_if(sshd:auth): requirement "uid >= 1000" not met by user "root".
```

Cause

The policy configuration of the PAM module bans users with a UID of below `1000` from logging in.

Solutions

Check and repair the PAM module as instructed in [Steps](#).

Steps

Note:

The steps in this document are for CentOS 6.5 as an example. As the steps may vary by operating system, please proceed based on the actual conditions.

1. [Log in to the Linux instance via VNC.](#)

2. Use the `cat` command to view the corresponding PAM configuration file as described below:

File	Feature Description
<code>/etc/pam.d/login</code>	Configuration file of the console (VNC)
<code>/etc/pam.d/sshd</code>	Configuration file of SSH login
<code>/etc/pam.d/system-auth</code>	Global configuration file of the system

3. Check for the configuration similar to what is shown below.

```
auth required pam_succeed_if.so uid >= 1000
```

4. Use Vim to modify or delete the policy configuration or add a pound sign (`#`) at the beginning of the line to comment it out. We recommend you back up the configuration before modifying it as needed.

```
auth          required      pam_succeed_if.so uid <= 1000      # Modify the policy
# auth        required      pam_succeed_if.so uid >= 1000    # Comment out the
relevant configuration
```

5. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH login error "Maximum amount of failed attempts was reached"

Problem

During login to a Linux instance via SSH key, the "Maximum amount of failed attempts was reached" error message appears.

Cause

Entering incorrect passwords several consecutive times triggered the policy restriction of the PAM module, causing the user account to be locked.

Solutions

Please proceed based on the actual conditions as instructed below:

[Root user not locked](#)

[Root user locked](#)

Steps

Note:

The steps in this document are for CentOS 7.6 and 6.5 as an example. As the steps may vary by operating system, please proceed based on the actual conditions.

Root user not locked

1. Log in to the instance as the root user. For more information, please see [Logging into Linux Instances via VNC](#).
2. Run the following command to view the global PAM configuration file of the system.

```
cat /etc/pam.d/system-auth
```

3. Run the following command to view the PAM configuration file of the local terminal.

```
cat /etc/pam.d/login
```

4. Run the following command to view the PAM configuration file of the SSH service.

```
cat /etc/pam.d/sshd
```

5. Use Vim to modify or delete the configuration in the above configuration files or add a pound sign (#) at the beginning of the line to comment it out. This document uses commenting the configuration out as an example. After the modification, the relevant configuration is as shown below:

```
#auth required pam_tally2.so deny=3 unlock_time=5
#auth required pam_tally.so onerr=fail no_magic_root
#auth requieired pam_tally2.so deny=5 lock_time=30 unlock_time=10 even_deny_root
root_unlock_time=10
```

Description:

The `pam_tally2` module is used here; if it is not supported, use the `pam_tally` module. The settings may vary by PAM version. For more information on how to use a specific module, please see the corresponding rules. Both the `pam_tally2` and `pam_tally` modules can be used for account lockout policy control. They differ in that the former has the automatic unlock time feature.

`even_deny_root` indicates to restrict the root user.

`deny` indicates to set the maximum number of consecutive incorrect login attempts for general users and root users. After it is exceeded, the user will be locked.

`unlock_time` indicates to unlock general users after they are locked for a specified period of time in seconds.

`root_unlock_time` indicates to unlock root users after they are locked for a specified period of time in seconds.

6. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

Root user locked

1. Enter the single user mode and log in to the instance. For detailed directions, please see [Booting into Linux Single User Mode](#).
2. Run the following commands to manually unlock the root user.

```
pam_tally2 -u root # View the number of consecutive incorrect password attempts
made by the root user
```



```
pam_tally2 -u root -r # Clear the number of consecutive incorrect password attempts made by the root user
```

```
authconfig --disableldap --update # Update the PAM authentication record
```

3. Restart the instance.

4. Comment out, modify, or update the corresponding PAM configuration file as instructed in [Root user not locked](#).

5. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH login error "login: Module is unknown"

Problem

Login to a Linux instance via SSH key fails, and a message similar to the following is displayed in the `secure` log:

```
login: Module is unknown.  
login: PAM unable to dlopen(/lib/security/pam_limits.so):  
/lib/security/pam_limits.so: cannot open shared object file: No such file or directory.
```

Cause

Each application with the PAM module enabled has a configuration file named after it in the `/etc/pam.d` directory; for example, the configuration file of the `login` command is `/etc/pam.d/login`, where you can configure specific policies as shown below:

File	Feature Description
<code>/etc/pam.d/login</code>	Configuration file of the console (VNC)
<code>/etc/pam.d/sshd</code>	Configuration file of SSH login
<code>/etc/pam.d/system-auth</code>	Global configuration file of the system

During remote login, certain applications with PAM enabled may fail to load the module, causing the login method with the corresponding policy configured to fail.

Solutions

Check and repair the configuration file as instructed in [Steps](#).

Steps

Note:

This document discusses the `/etc/pam.d/sshd` and `/etc/pam.d/system-auth` files. If `/etc/pam.d/login` is exceptional, please [submit a ticket](#) for assistance.

1. [Log in to the Linux instance via VNC.](#)
2. Run the following command to view the PAM configuration file.

```
cat [absolute path of the corresponding PAM configuration file]
```

Check for the configuration similar to what is shown below. The file path is `/lib/security/pam_limits.so`.

```
session      required      pam_limits.so
```

3. Run the following command to check whether the `/lib/security/pam_limits.so` path is incorrect.

```
ll /lib/security/pam_limits.so
```

If yes, use Vim to edit the PAM configuration file and repair the path of the `pam_limits.so` module. The correct path should be `/lib64/security` on a 64-bit Linux instance. The modified configuration information should be as shown below:

```
session      required      /lib64/security/pam_limits.so
```

If no, please [submit a ticket](#) for assistance.

4. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH service exception error due to virus "fatal: mm_request_send: write: Broken pipe"

Problem

The "fatal: mm_request_send: write: Broken pipe" error message is displayed due to an SSH service exception caused by a virus.

Cause

This may be because that viruses such as udev-fall affect the normal execution of the SSH service.

Solutions

Fix the virus issue based on actual conditions as instructed below.

[Temporary solution](#)

[Reliable solution](#)

Steps

Temporary solution

This document uses the udev-fall virus as an example. You can temporarily make the SSH service work normally as instructed below.

1. [Log in to the Linux instance via VNC.](#)
2. Run the following command to view the process information of the udev-fall virus and record the process ID.

```
ps aux | grep udev-fall
```

3. Run the following command to end the process.

```
kill -9 [virus process ID]
```

4. Run the following command to ban udev-fall from automatically starting.

```
chkconfig udev-fall off
```

5. Run the following command to delete all the instructions and startup configuration related to udev-fall.

```
for i in `find / -name "udev-fall"`;  
do echo '' > $i && rm -rf $i;  
done
```

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

```
systemctl restart sshd.service
```

Reliable solution

As it is unclear whether the virus or intruder has made other changes to the system or hides other virus files, to ensure the long-term stable operations of the server, we recommend you restore the server to its normal state by rolling back the system disk snapshot of the instance. For more information, please see [Rolling Back Snapshots](#).

Note:

Note that snapshot rollback will cause loss of data stored after the snapshot creation.

We recommend you roll back to snapshots one by one from the latest to the earliest till the SSH service resumes. If the SSH service still cannot run properly after the rollback, the system at that point in time has been abnormal.

SSH service start error "main process exited, code=exited"

Problem

When the SSH service is started by the `service` or `systemctl` command on a Linux instance, the command line does not return any error message, but the service cannot run properly, and an error message similar to the following is displayed in the `secure` log:

```
sshd.service: main process exited, code=exited, status=203/EXEC.  
init: ssh main process (1843) terminated with status 255.
```

Cause

Usually, this is because that configuration of the `PATH` environment variable is exceptional or the relevant files of the SSH software package are removed.

Solutions

Check and repair the `PATH` environment variable or reinstall the SSH software package as instructed in [Steps](#).

Steps

Note:

The steps in this document are for CentOS 6.5 as an example. As the steps may vary by operating system, please proceed based on the actual conditions.

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to check the environment variable configuration.

```
echo $PATH
```

3. Compare the returned value of the `PATH` environment variable with its default value as shown below:

```
/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:/root/bin
```

If the returned value of the `PATH` environment variable is not the same as the default value, you need to run the following command to reset it.

```
export  
PATH=/usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/sbin:/usr/bin:/root/bin
```

4. Run the following command to find and confirm the path of the `sshd` program.

```
find / -name sshd
```

If the following result is returned, the `sshd` program files already exist.

```
/usr/sbin/sshd
```

If the corresponding files do not exist, please reinstall the SSH software package.

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

```
service sshd restart
```

Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).

SSH login error "pam_limits(sshd:session): could not sent limit for 'nofile'"

Problem

During login to a Linux instance via SSH key, the following error message is returned:

```
-bash: fork: retry: Resource temporarily unavailable.  
pam_limits(sshd:session): could not sent limit for 'nofile':operation not  
permitted.  
Permission denied.
```

Cause

Usually, this is because that the number of currently opened shell processes or files exceeds the ulimit system environment limit of the server.

Solutions

Modify the `limits.conf` file to permanently change the ulimit system environment limit based on the operating system version as instructed in [Steps](#).

Steps

Note:

Since CentOS 6, the `x-nproc.conf` file is used to manage the ulimit system environment limits. The steps for versions below and above CentOS 6 are differentiated here. The prefix number of the `x-nproc.conf` file varies by system version; for example, it is `90-nproc.conf` on CentOS 6 and `20-nproc.conf` on CentOS 7. Please proceed based on the actual environment.

This document uses CentOS 7.6 and CentOS 5 as an example. Please proceed based on the actual business conditions.

Below CentOS 6

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the current ulimit system resource limits.

```
cat /etc/security/limits.conf
```

Description:

<domain>: target system user. You can use `*` to indicate all users.

<type>: `soft`, `hard`, and `-`.

`soft` is the <value> of the current system that has taken effect.

`hard` is the maximum <value> set in the system.

The limit of `soft` cannot be greater than that of `hard`. `-` indicates to set the values of `soft` and `hard` at the same time.

<item>: target resource type.

`core` limits the kernel file size.

`rss` is the maximum resident set size.

`nofile` is the maximum number of opened files.

`nproc` is the maximum number of processes.

3. No system resource limits are set by default. Please judge based on the actual conditions. If system resource limits are enabled and configured, you need to edit the `limits.conf` file to comment out, modify, or delete the resource type code limited by the `nproc` or `nofile` parameter.

We recommend you run the following command to back up the `limits.conf` file before modifying it.

```
cp -af /etc/security/limits.conf /root/limits.conf_bak
```

4. After completing the modification, restart the instance.

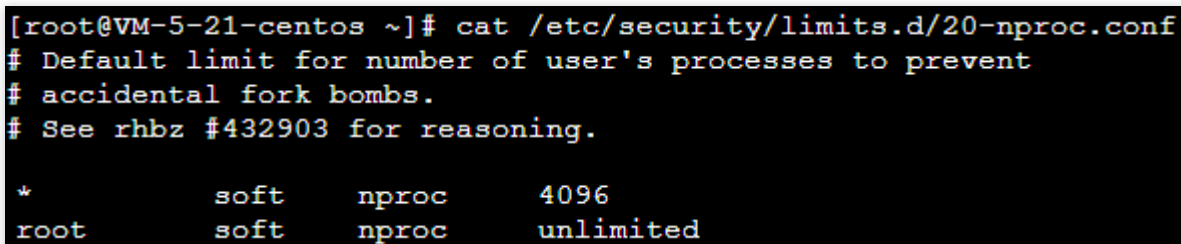
CentOS 6 and above

1. [Log in to the Linux instance via VNC](#).

2. Run the following command to view the current ulimit system resource limits.

```
cat /etc/security/limits.d/20-nproc.conf
```

If the returned result is as shown below, system resource limits have been enabled and the maximum number of connection processes allowed for all users except root users is 4,096.



```
[root@VM-5-21-centos ~]# cat /etc/security/limits.d/20-nproc.conf
# Default limit for number of user's processes to prevent
# accidental fork bombs.
# See rhbz #432903 for reasoning.

*                soft    nproc           4096
root             soft    nproc           unlimited
```

3. Modify the `/etc/security/limits.d/20-nproc.conf` file as instructed in [Below CentOS 6](#). We recommend you back up the file before doing so.

4. After completing the modification, restart the instance.

SSH login error "pam_unix(sshd:session) session closed for user"

Problem

Login to a Linux instance via SSH key fails with the correct username and password entered. An error message similar to the following is directly returned or appears in the `secure` log:

This account is currently not available.

Connection to 127.0.0.1 closed.

Received disconnect from 127.0.0.1: 11: disconnected by user.

pam_unix(sshd:session): session closed for user test.

Cause

Usually, this is because that the default shell of the corresponding user is modified.

Solutions

Check and repair the default shell configuration of the corresponding user as instructed in [Steps](#).

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the default shell of the test user.

```
cat /etc/passwd | grep test
```

If a message similar to the following is returned by the system, the shell of the test user has been changed to `nologin`.

```
test:x:1000:1000::/home/test:/sbin/nologin
```

3. Run the following command and use Vim to edit the `/etc/passwd` file. We recommend you back up the file before doing so.

```
vim /etc/passwd
```

4. Press **i** to enter the edit mode and change `/sbin/nologin` to `/bin/bash`.
 5. Press **Esc** and enter **:wq** to save and close the file.
 6. Log in to the instance via SSH key. For more information, please see [Logging into Linux Instance via SSH Key](#).
- If the problem persists, please [submit a ticket](#) for assistance.

Failing to log in to a Linux CVM due to high CPU and memory usage

Last updated : 2024-01-06 17:32:18

This document describes how to troubleshoot and handle Linux CVM login failures due to high CPU or memory usage.

Possible Cause

Hardware, system processes, business processes, and trojans may cause high CPU or memory usage, resulting in slow service response or CVM login failure. You can use [Cloud Monitor](#) to create an alarm threshold for CPU or memory usage. Then, you will be notified promptly when the configured threshold is exceeded.

Tools

Top: A monitoring tool on Linux commonly used to obtain the CPU or memory usage by process. The output information of the `top` command is as shown below:


```
top - 22:16:25 up 6:18, 1 user, load average: 0.00, 0.01, 0.05
Tasks: 68 total, 1 running, 67 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.0 us, 0.3 sy, 0.0 ni, 99.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 1016516 total, 605016 free, 77224 used, 334276 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 778708 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
257	root	20	0	0	0	0	S	0.3	0.0	0:00.73	jbd2/vda1-8
984	root	20	0	569592	5068	2568	S	0.3	0.5	0:16.51	YDService
1253	root	20	0	534620	12288	2104	S	0.3	1.2	0:34.21	barad_agent
1	root	20	0	43104	3512	2404	S	0.0	0.3	0:01.87	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd
3	root	20	0	0	0	0	S	0.0	0.0	0:00.33	ksoftirqd/0
4	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kworker/0:0
5	root	0	-20	0	0	0	S	0.0	0.0	0:00.00	kworker/0:0H
7	root	rt	0	0	0	0	S	0.0	0.0	0:00.00	migration/0
8	root	20	0	0	0	0	S	0.0	0.0	0:00.00	rcu_bh
9	root	20	0	0	0	0	S	0.0	0.0	0:01.20	rcu_sched
10	root	rt	0	0	0	0	S	0.0	0.0	0:00.05	watchdog/0

The `top` command output consists of two parts. The upper part displays the general usage of CPU and memory resources:

Line 1: The current system time, current number of logged-in users, and system load.

Line 2: The total number of system processes and the numbers of running, hibernating, sleeping, and zombie processes.

Line 3: The current CPU usage.

Line 4: The current memory usage.

Line 5: The current swap space usage.

The lower part displays the resource usage by process:

PID: Process ID.

USER: Process owner.

PR: Process priority. NI is the NICE value. The smaller the NICE value, the higher the priority.

VIRT: Used virtual memory size in KB.

RES: Currently used memory size in KB.

SHR: Used shared memory size in KB.

S: Process state.

%CPU: Percentage of CPU time used by the process within the update time interval.

%MEM: Percentage of memory used by the process within the update time interval.

TIME+: CPU time used by the process, accurate down to 0.01s.

COMMAND: Process name.

Troubleshooting

Logging in to CVM

Select a CVM login method based on your actual needs.

Log in to the Linux CVM instance remotely via third party software.

Note:

If the Linux CVM instance has a high CPU load, the login may fail.

[Logging into Linux Instances via VNC.](#)

Note:

If the Linux CVM instance has a high CPU load, you can log in to it in the console.

Viewing the resource usage of processes

Run the following command to view the system load. View the **%CPU** and **%MEM** columns and identify which processes consume more resources.

```
top
```

Analyzing process

Analyze the processes on the **Task Manager** page to troubleshoot and solve the problem.

If the problem is caused by a service process, analyze whether the service process can be optimized and accordingly optimize the process or [upgrade the CVM configuration](#).

If the problem is caused by a process with an exception, the instance may have a virus. In this case, you can terminate the process or use an antivirus application to kill the virus. When necessary, back up the data and reinstall the operating system.

If the problem is caused by a Tencent Cloud component process, [submit a ticket](#) for assistance.

Common Tencent Cloud components include:

sap00x: A security component

Barad_agent: A monitoring component

secu-tcs-agent: A security component

Terminating process

1. Compare the resource consumption of different processes and record the PID of the process that needs to be terminated.
2. Enter `k` .

3. Enter the PID of the process that needs to be terminated and press the **Enter** key to terminate it, as shown in the following figure:

Suppose you need to terminate a process whose PID is 23.

```
top - 09:58:45 up 51 min, 1 user, load average: 0.00, 0.01, 0.05
Tasks: 351 total, 1 running, 350 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.0 us, 0.1 sy, 0.0 ni, 99.9 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 1870516 total, 1441292 free, 127068 used, 302156 buff/cache
KiB Swap: 2097148 total, 2097148 free, 0 used, 1537932 avail Mem
PID to signal/kill [default pid = 293] 23
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND
293 root 20 0 0 0 0 S 0.2 0.0 0:03.24 kworker/2:1
524 root 20 0 0 0 0 S 0.1 0.0 0:03.53 kworker/0:2
137 root 20 0 0 0 0 S 0.1 0.0 0:02.70 rcu_sched
141 root 20 0 0 0 0 S 0.0 0.0 0:00.73 rcuos/3
15672 root 20 0 130156 2028 1260 R 0.0 0.1 0:04.61 top
1 root 20 0 57592 7436 2612 S 0.0 0.4 0:03.44 systemd
310 root 20 0 0 0 0 S 0.0 0.0 0:00.64 kworker/u256:1
333 root 20 0 0 0 0 S 0.0 0.0 0:00.26 kworker/3:1
540 root 20 0 0 0 0 S 0.0 0.0 0:00.11 jbd2/sda2-8
619 root 20 0 43016 2076 2564 S 0.0 0.2 0:00.33 systemd-journal
730 root 20 0 329592 23192 6252 S 0.0 1.2 0:01.02 firewallld
745 root 20 0 19204 1236 944 S 0.0 0.1 0:00.67 irqbalance
754 dbus 20 0 34000 1904 1420 S 0.0 0.1 0:00.27 dbus-daemon
853 root 20 0 509040 9620 5956 S 0.0 0.5 0:00.30 NetworkManager
901 polkitd 20 0 514364 12260 4568 S 0.0 0.7 0:00.17 polkitd
1016 root 20 0 91064 2064 1064 S 0.0 0.1 0:00.09 master
15601 root 20 0 0 0 0 S 0.0 0.0 0:00.06 kworker/1:1
15699 root 20 0 0 0 0 S 0.0 0.0 0:00.01 kworker/1:0
2 root 20 0 0 0 0 S 0.0 0.0 0:00.09 kthreadd
```

Caution:

If `kill PID 23 with signal [15]:` appears after you press **Enter**, press **Enter** again to keep the default settings.

4. If the operation is successful, the message `Send PID 23 signal [15/sigterm]` will appear. Press **Enter** to confirm the termination.

Other Related Problems

Low CPU usage but high load average

Cause

The load average is an indicator of CPU load. The higher the load average, the longer the queue of pending processes is.

After the `top` command is executed, information similar to the following is returned, indicating that the CPU usage is low but the load average is very high.

```
top - 19:46:57 up 27 days, 5:33, 1 user, load average: 23, 22, 23
Tasks: 94 total, 1 running, 93 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.3 us, 0.0 sy, 0.0 ni, 99.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem: 1016656 total, 950428 used, 66228 free, 170148 buffers
KiB Swap: 0 total, 0 used, 0 free. 452740 cached Mem
```

Solution

Run the following command to check whether any process is in the D state as shown below:

```
ps -axjf
```

```
1 516 516 516 ? -1 Ss 0 0:00 /sbin/iprinit --daemon
1 569 569 569 ? -1 Ss 0 0:00 /sbin/iprdump --daemon
1 863 863 863 ? -1 D+ 38 0:16 /usr/sbin/ntpd -u ntp:ntp -g
1 874 874 874 ? -1 Ss 0 0:01 /usr/sbin/sshd -D
874 8823 8823 8823 ? -1 Ss 0 0:03 \_ sshd: root@pts/0
8823 8825 8825 8825 pts/0 9006 Ss 0 0:00 \_ -bash
8825 9006 9006 8825 pts/0 9006 D+ 0 0:00 \_ ps -axjf
```

Note: The D state refers to the uninterrupted sleep state. A process in this state cannot be terminated nor can it be exited by itself.

If there are many processes in the D state, restore the resources on which the processes depend or restart the operating system.

High CPU usage by kswapd0 process

Cause

Linux manages memory by using the pagination mechanism and sets aside a portion of the disk as virtual memory. kswapd0 is the process responsible for page replacement in the virtual memory management of the Linux system. When system memory becomes insufficient, kswapd0 will frequently replace pages, which will result in high CPU usage.

Solution

1. Run the following command and find the kswapd0 process.

```
top
```

2. Check the state of the kswapd0 process.

If the process is not in the D state and has been running for a long time and consuming too many CPU resources, perform [step 3](#) to check the memory usage.

3. Run commands such as `vmstat` , `free` , and `ps` to check how much memory is used by processes in the system.

Based on the memory usage, restart the system or terminate safe but unnecessary processes. If the `si` and `so` values are also high, pages are frequently replaced in the system. If the physical memory of the current system can no longer meet your requirements, consider upgrading your system memory.

Remote Login Failure due to Port Issues

Last updated : 2024-01-06 17:32:18

This document describes how to diagnose and troubleshoot remote login failures caused by port problems.

Note:

The following operations use a CVM instance with the CentOS 7.8 operating system as an example.

Tools

You can use the following tools to check whether the login issues are related to port or security group configuration:

[Self-diagnose](#)

[Port Verification](#)

If the problem is indeed caused by security group configurations, click **Open all ports** in [Port Verification](#) and try to log in again. If you still cannot log in after opening the ports, refer to the following for troubleshooting.

Troubleshooting

Checking network connectivity

You can use the `Ping` command to test network connectivity from your PC. You should run the test from computers in different network environments (such as different IP ranges or ISPs) to check whether it is a local network problem or a server problem.

1. Open the command line tool on your local computer.

Windows: click **Start > Run** and enter `cmd` . A Command Prompt window appears.

MacOS: open a Terminal window.

2. Run the following command to test network connection.

```
ping [CVM instance's public IP address]
```

You should first [obtain the public IP address](#) of the CVM instance. For example, `ping 81.71.XXX.XXX` .

If the result similar to the following is returned, your network connection to the CVM instance is normal.

```
ping 81.71.
Pinging 81.71. with 32 bytes of data:
Reply from 81.71. : bytes=32 time=13ms TTL=44
Reply from 81.71. : bytes=32 time=12ms TTL=44
Reply from 81.71. : bytes=32 time=12ms TTL=44
Reply from 81.71. : bytes=32 time=12ms TTL=44

Ping statistics for 81.71. :
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 12ms, Maximum = 13ms, Average = 12ms
```

If **Request Timeout** appears, your network connection to the CVM instance is not working properly. In this case, refer to [Instance IP Address Ping Failure](#) for troubleshooting instructions.

Checking port connectivity

1. [Log in to a Linux instance via VNC](#).
2. Run the following command and press **Enter** to check whether the remote port is open and accessible.

```
telnet [CVM instance's public IP address] [Port number]
```

For example, run the `telnet 119.XX.XXX.67 22` command to check whether the port 22 is open.

If information similar to what is shown below is returned, the port 22 is accessible.

```
[root@VM-8-25-centos ~]# telnet 119.29.115.67 22
Trying 119.29.115.67...
Connected to 119.29.115.67.
Escape character is '^]'.
SSH-2.0-OpenSSH_7.4
```

If information similar to what is shown below is returned, the port 22 is unaccessible. Check the corresponding network configuration. For example, check whether the port 22 is open in the firewall or security group of the instance.

```
[root@VM-8-25-centos ~]# telnet 119.29.115.67 22
Trying 119.29.115.67...
telnet: connect to address 119.29.115.67: Connection timed out
```

Checking the SSHD service

If you are unable to [log in to a Linux instance via SSH key](#) due to connection failures, it may be because the SSHD port is not being listened on or the SSHD service is not started. In this case, refer to [Unable to Log into a Linux Instance via SSH](#) for troubleshooting instructions.

VNC Login Error (Module is Unknown)

Last updated : 2024-01-06 17:32:18

Error Description

I entered the correct password, but still could not log in to the CVM using VNC. The message "Module is unknown" appears.

```
login: root
Password:
Last failed login: Mon Oct 26 10:24:25 CST 2020 from 31.154.9.174 on ssh:notty
There were 46 failed login attempts since the last successful login.

Module is unknown
```

Possible Reasons

This issue may be caused by the `/etc/pam.d/system-auth` configuration in `/etc/pam.d/login` file.

```
%PAM-1.0
auth [user_unknown=ignore success=ok ignore=ignore default=bad] pam_securetty.so
auth      substack      system-auth
auth      include       postlogin
account   required      pam_nologin.so
account   include       system-auth
password  include       system-auth
# pam_selinux.so close should be the first session rule
session   required      pam_selinux.so close
session   required      pam_loginuid.so
session   optional      pam_console.so
# pam_selinux.so open should only be followed by sessions to be executed in the user cont
session   required      pam_selinux.so open
session   required      pam_namespace.so
session   optional      pam_keyinit.so force revoke
session   include       system-auth
session   include       postlogin
-session  optional      pam_ck_connector.so
~
```

If the path of the `pam_limits.so` module is not configured correctly in the `system-auth` configuration file, the login fails.

```

##PAM-1.0
# This file is auto-generated.
# User changes will be destroyed the next time authconfig is run.
auth      required      pam_env.so
auth      required      pam_faildelay.so delay=2000000
auth      sufficient    pam_unix.so nullok try_first_pass
auth      requisite     pam_succeed_if.so uid >= 1000 quiet_success
auth      required      pam_deny.so

account    required      pam_unix.so
account    sufficient    pam_localuser.so
account    sufficient    pam_succeed_if.so uid < 1000 quiet
account    required      pam_permit.so

password   requisite     pam_pwquality.so try_first_pass local_users_only retry=3 authtok_type=
password   sufficient    pam_unix.so md5 shadow nullok try_first_pass use_authtok
password   required      pam_deny.so

[
session    optional     pam_keyinit.so revoke
session    required     /lib/security/pam_limits.so
-session   optional     pam_systemd.so
session    [success=1 default=ignore] pam_succeed_if.so service in crond quiet use_uid
session    required      pam_unix.so

```

Note:

The `pam_limits.so` module limits the system resource usage of a user during the session. If the module path is not configured correctly according to the actual operating system, the login authentication fails.

Solutions

1. Perform the [troubleshooting procedure](#) to locate the `pam_limits.so` path configuration in the `system-auth` file.
2. Correct the `pam_limits.so` module path.

Troubleshooting Procedure

1. Try to [log in to Linux CVM via SSH key](#).

If the login succeeded, proceed to the next step.

If login failed, use single user mode. For more information, see [Booting into Linux Single User Mode](#).

2. Run the following command to view logs.

```
vim /var/log/secure
```

This file records the security information, mostly CVM login logs. You can check the error logs of

`/lib/security/pam_limits.so` as shown below.

```
Oct 28 15:51:01 VM-96-4-centos crond[3803]: PAM unable to dlopen(/lib/security/pam_limits.so): /lib/security/pam_limits.so: cannot open shared object file: No such file or directory
Oct 28 15:51:01 VM-96-4-centos crond[3803]: PAM adding faulty module: /lib/security/pam_limits.so
Oct 28 15:51:06 VM-96-4-centos login: PAM unable to dlopen(/lib/security/pam_limits.so): /lib/security/pam_limits.so: cannot open shared object file: No such file or directory
```

3. Run the following commands in sequence to enter `/etc/pam.d` directory and search for

`/lib/security/pam_limits.so` .

```
cd /etc/pam.d
```

```
find . | xargs grep -ri "/lib/security/pam_limits.so" -l
```

If the result similar to the following figure is returned, `/lib/security/pam_limits.so` is configured in the `system-auth` file.

```
bash-4.2# find . | xargs grep -ri "/lib/security/pam_limits.so" -l
./system-auth-ac
./system-auth
./system-auth-ac
```

4. Access the `system-auth` file to correct the `pam_limits.so` module path.

For example, you can use the absolute path `/lib64/security/pam_limits.so` or a relative path


`pam_limits.so` in a 64-bit operating system.

VNC Login Error (Account Locked due to XXX Failed Logins)

Last updated : 2024-01-06 17:32:18

Error Description

The error “Account locked due to XXX failed logins” appears before I enter the login password.



```
CentOS Linux 7 (Core)  
Kernel 3.10.0-1062.18.1.el7.x86_64 on an x86_64  
  
login: root  
Account locked due to 10 failed logins  
Password:
```

Possible Reasons

This issue may be caused by the `pam_tally2.so` configuration in `/etc/pam.d/login` file. For VNC login, `/etc/pam.d/login` is called for authentication, while `pam_tally2.so` indicates to automatically lock the user account temporarily or permanently after the specified number of consecutive failed logins. When an account is permanently locked, you need to unlock it manually.

The login account will be locked when the number of consecutive failed logins exceeds the configured value. Note that the account may also be locked for brute force attacks.

```

#%PAM-1.0
auth      required      pam_tally2.so deny=6 un_lock_time=300 even_deny_root root_unlock_time=
auth [user_unknown=ignore success=ok ignore=ignore default=bad] pam_securetty.so
auth      substack      system-auth
auth      include       postlogin
account   required      pam_nologin.so
account   include       system-auth
password  include       system-auth
# pam_selinux.so close should be the first session rule
session   required      pam_selinux.so close
session   required      pam_loginuid.so
session   optional      pam_console.so
# pam_selinux.so open should only be followed by sessions to be executed in the user context
session   required      pam_selinux.so open
session   required      pam_namespace.so
session   optional      pam_keyinit.so force revoke
session   include       system-auth
session   include       postlogin
-session  optional      pam_ck_connector.so
~

```

See below for the parameters of `pam_tally2` module.

Parameter	Description
deny=n	Lock the account if the number of consecutive failed logins exceeds n.
lock_time=n	Lock the account for n seconds when the number of consecutive failed logins exceeds the limit
un lock_time=n	Unlock the account automatically n seconds later
no_lock_time	Do not use .fail_locktime field in /var/log/faillog
magic_root	If the module is invoked by a root user (uid=0), the counter is not incremented.
even_deny_root	The root user will be locked after deny=n consecutive failed logins.
root_unlock_time=n	This parameter is required if even_deny_root is configured. It indicates how long the the root user is locked when the number of consecutive failed logins exceeds the limit.

Solutions

1. Refer to [troubleshooting procedure](#) to access the login configuration file and temporarily comment out the configuration of the `pam_limits.so` module.
2. Find the reason why your account is locked, and improve your security policy.

Troubleshooting Procedure

1. Try to [log in to Linux CVM via SSH key](#).

If the login succeeded, proceed to the next step.

If the login failed, try the single user mode.

2. Run the following command to view logs.

```
vim /var/log/secure
```

This file records the security information, mostly CVM login logs. You can check the error logs of `pam_tally2` as shown below.

```
Oct 28 17:14:45 VM-96-4-centos sshd[167041]: Failed password for invalid user dell from 202.153.37.205 port 13069 ssh2
Oct 28 17:14:45 VM-96-4-centos sshd[167041]: Received disconnect from 202.153.37.205 port 13069:11: Bye Bye [preauth]
Oct 28 17:14:45 VM-96-4-centos sshd[167041]: Disconnected from 202.153.37.205 port 13069 [preauth]
Oct 28 17:14:59 VM-96-4-centos login: pam_tally2(login:auth): user root (0) tally 12, deny 2
Oct 28 17:14:59 VM-96-4-centos login: pam_succeed_if(login:auth): requirement "uid >= 1000" not met by user "root"
Oct 28 17:15:01 VM-96-4-centos login: FAILED LOGIN 2 FROM tty1 FOR root, Authentication failure
Oct 28 17:15:01 VM-96-4-centos login: pam_tally2(login:auth): unknown option: un_lock_time=300
Oct 28 17:15:03 VM-96-4-centos login: pam_tally2(login:auth): user root (0) tally 13, deny 2
Oct 28 17:15:04 VM-96-4-centos sshd[167041]: pam_unix(sshd:auth): authentication failure; logname= uid=0 euid=0 tty=ssh ruser=
ost=203.213.66.170 user=root
```

3. Run the following commands in sequence to open `/etc/pam.d` directory and search for `pam_tally2`.

```
cd /etc/pam.d
```

```
find . | xargs grep -ri "pam_tally2" -l
```

If the result similar to the following figure is returned, `pam_tally2` is included in `login` file.

```
bash-4.2# find . | xargs grep -ri "pam_tally2" -l
./login
./login
bash-4.2# _
```

4. Run the following command to temporarily comment out the `pam_tally2.so` configurations. Then you can log in normally.

```
sed -i "s/.*pam_tally.*/#&/" /etc/pam.d/login
```

5. Check whether the account is locked due to misoperations or brute force attacks. In the later case, it is recommended to strengthen the security policy as follows:

Change the CVM password to a stronger password containing 12-16 characters, including uppercase letters, lowercase letters, special characters, and numbers. For more information, see [Resetting Instance Password](#).

Delete unused CVM login accounts.

Change the default sshd port 22 to a less common port between 1024-65525. For more information, see [Modifying the Default Remote Port of CVM](#).

Manage the associated security group rules to open only ports and protocols required by your business. For more information, see [Adding Security Group Rules](#).

Close the port for internet access for core applications such as MySQL and Redis databases.

Install security software (such as CWPP agent), and configure real-time alarms to get notices about suspicious logins instantly.

VNC Login Error (Login Failed with Correct Password)

Last updated : 2024-01-06 17:32:18

Problem

When you try to log in to the CVM via VNC, the following message appears even you enter the correct password. Later, you are required to enter the account name again.

```
VM-55-10-centos login: root
Password:
-
```

And when you try to log in remotely using the SSH key, the message **Permission denied, please try again** appears.

```
[root@VM-96-14-centos ~]# ssh root@
root@4 's password:
Permission denied, please try again.
```

Common Cause

The `/var/log/btmp` log file is oversized due to brute force attacks. This file keeps logs of failed logins. If it is too large, logs can not be written into it, which may cause login error.

```

bash-4.2# ll -h
bash: ll: command not found
bash-4.2# ls -alh
total 9.8G
drwxr-xr-x 10 root  root      4.0K Oct 28 17:53 .
drwxr-xr-x 19 root  root      4.0K Apr 22 2020 ..
drwxr-xr-x  2 root  root      4.0K Mar  7 2019 anaconda
drwx----- 2 root  root      4.0K Aug  8 2019 audit
-rw-----  1 root  root      24K Oct 28 17:30 boot.log
-rw-----  1 root  root         1 Oct 28 15:43 boot.log-20191106
-rw-----  1 root  root         1 Oct 28 15:43 boot.log-20200807
-rw-----  1 root  utmp      9.8G Oct 28 17:41 btmp
-rw-----  1 root  utmp         1 Oct 28 15:43 btmp-20200807
drwxr-xr-x  2 chrony chrony    4.0K Aug  8 2019 chrony
-rw-r--r--  1 syslog adm     181K Oct 28 17:30 cloud-init.log
-rw-r--r--  1 root  root      7.8K Oct 28 17:30 cloud-init-output.log
-rw-----  1 root  root      14K Oct 28 17:42 cron
-rw-r--r--  1 root  root     36K Oct 28 17:30 dmesg
-rw-r--r--  1 root  root     36K Oct 28 16:26 dmesg.old

```

Solutions

1. Check whether the `/var/log/btmp` log file is oversized as instructed in [Troubleshooting Procedure](#).
2. Confirm whether it is caused by brute force attacks and improve security policy.

Troubleshooting Procedure

1. Try to [log in to a Linux CVM instance via SSH key](#).

If the login succeeds, proceed to the next step.

If the login fails, try the single user mode. For detailed directions, see [Booting into Linux Single User Mode](#).

2. Access `/var/log` and check the size of the `/var/log/btmp` log file.
3. Run the following command to clear the oversized `/var/log/btmp` log file. Then you can log in normally.

```
cat /dev/null > /var/log/btmp
```

4. Check whether the account lock is caused by misoperations or brute force attacks. In the later case, it is recommended to strengthen the security policy as follows:

Change the CVM password to a stronger password containing 12-16 characters, including uppercase letters, lowercase letters, special characters, and numbers. For more information, see [Resetting Instance Password](#).

Delete unused CVM login accounts.

Change the default sshd port 22 to a less common port between 1024-65525. For more information, see [Modifying the Default Remote Port of CVM](#).

Manage the associated security group rules to open only ports and protocols required by your business. For more information, see [Adding Security Group Rules](#).

Close the port for internet access for core applications such as MySQL and Redis databases.

Install security software (such as CWPP agent), and configure real-time alarms to get noticed about suspicious logins instantly.

VNC or SSH Login Error (Permission Denied)

Last updated : 2025-06-09 18:30:12

Error Description

The error message "Permission denied" is reported when I log in using VNC or SSH key.

The VNC login error is shown below:

The SSH login error is shown below:

Possible Reasons

Using the VNC or SSH login will call `system-auth` for authentication if this module is configured in the `/etc/pam.d/login` configuration file. By default, the `system-auth` module introduces the `pam_limits.so` module. The default `system-auth` configuration is as shown below:

The `pam_limits.so` module is mainly used to limit the use of system resources during the user session. Its default configuration file `/etc/security/limits.conf` specifies the maximum number of files, the maximum number of threads, the maximum memory and other resources that a user can use. See the table below for details.

Parameter	Description
<code>soft nofile</code>	The maximum number of open file descriptors (soft limit)
<code>hard nofile</code>	The maximum number of open file descriptors (hard limit), which cannot be exceeded.
<code>fs.file-max</code>	The maximum number of open file handles (struct file in the kernel) at the system level.
<code>fs.nr_open</code>	The maximum number of file descriptors (fd) assigned to a process

The login failure may be caused by incorrect configurations of the maximum number of open file descriptors for the root account in the `/etc/security/limits.conf` configuration file. The set value of `soft nofile` should be no more than `hard nofile`, and `hard nofile` should be no more than `fs.nr_open`.

Solutions

Perform the [troubleshooting procedure](#) to correct the relationship configurations of `soft nofile` , `hard nofile` and `fs.nr_open` .

Troubleshooting Procedure

1. Try to [log in to Linux CVM via SSH key](#).

If login succeeded, proceed to the next step.

If login failed, use single user mode. For more information, see [Booting into Linux Single User Mode](#).

2. Check whether the set values meet the relationship `soft nofile ≤ hard nofile ≤ fs.nr_open` .

Run the following command to obtain the values of `soft nofile` and `hard nofile` .

```
cat /etc/security/limits.conf
```

In this example, their values are 3000001 and 3000002 respectively, as shown below.

Run the following command to check the `fs.nr_open` value.

```
sysctl -a 2>/dev/null | grep -Ei "file-max|nr_open"
```

In this example, its value is 1048576, as shown below.

3. Edit the `/etc/security/limits.conf` file to add or modify the following configurations at the end of the file.

```
root soft nofile : 100001
```

```
root hard nofile : 100002
```

4. Edit the `/etc/sysctl.conf` file to add or modify the following configurations at the end of the file.

Note:

This step is optional when the relationship `soft nofile ≤ hard nofile ≤ fs.nr_open` is met. Perform this step to increase the system limit.

```
fs.file-max = 2000000
```

```
fs.nr_open = 2000000
```

5. Run the following command for the configuration to take effect immediately. Then you can log in normally.

```
sysctl -p
```

Login Failure Due to /etc/fstab Configuration Errors

Last updated : 2024-01-06 17:32:18

Error Description

Error: Unable to remotely log in to the Linux CVM via SSH. After logging in to the CVM via VNC, you can check the system start-up failure and view the prompt message "Welcome to emergency mode".

```
[ OK ] Reached target Remote File Systems (Pre).
[ OK ] Reached target Remote File Systems.
       Starting Crash recovery kernel arming...
[ OK ] Started Security Auditing Service.
       Starting Update UTMP about System Boot/Shutdown...
[ OK ] Started Update UTMP about System Boot/Shutdown.
       Starting Update UTMP about System Runlevel Changes...
[ OK ] Started Update UTMP about System Runlevel Changes.
Welcome to emergency mode! After logging in, type "journalctl -xb" to view c   i
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or press Control-D to continue):
```

Possible Causes

The `/etc/fstab` is not properly configured.

For example, you've configured the auto-attaching of disk based on device name in the `/etc/fstab` file. If the device name is changed when the CVM restarts, this configuration will cause the system to fail to start up normally.

Solutions

See [Steps](#) to repair the `/etc/fstab` configuration file. Then, restart the CVM to verify the repaired file.

Steps

You can access the instance in the following 2 methods for troubleshooting:

Method 1: Login via VNC (recommended)

Method 2: Use rescue mode

1. [Log in to Linux Instances \(VNC\)](#).

2. After entering the VNC interface, you see the interface shown in [Error Description](#). Enter the root account password and press **Enter** to log in to the server.

The entered password is not displayed by default.

If you do not have or forgot the root account password, see Method 2 for troubleshooting.

3.

Run the following command to back up the `/etc/fstab` file to the `/home` directory, for example:

```
cp /etc/fstab /home
```

4. Run the following command to use VI editor to open the `/etc/fstab` file.

```
vi /etc/fstab
```

5. Press **i** to enter the edit mode. Move the cursor to the beginning of the error line and enter `#` to comment out this configuration.

Note:

If you cannot determine the error, it is recommended that you comment out the configurations of all attached disks except the system disks, and then configure the file as instructed in [Step 8](#) after the server recovers.

```
# /etc/fstab
# Created by anaconda on Tue Nov 26 02:11:36 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
UUID=659e6f89-71fa-463d-842e-ccdf2c06e0fe / ext4 defaults
# /dev/vdc1_data auto rw,relatime,data=ordered 0 2
```

6.

Press **Esc**, enter `:wq`, and press **Enter** to save the configuration and exit the editor.

7. Restart the instance in the CVM console to see if it can be started and logged in properly.

Note:

Restart the instance through the console. For more information, see [Restarting Instances](#).

8. After successful login, if you need to configure the auto-attaching of the disk, see [Cloud Disk Automount Failed upon Linux CVM Restart](#).

1. Enter the instance rescue mode, see [Using Rescue Mode](#).

Note:

Run the `mount` and `chroot` commands as instructed in [Using rescue mode to repair system](#), and make sure that you have entered the target system.

2. Follow [Step 3 - Step 6](#) in Method 1 to repair the `/etc/fstab` file.

3. Exit the rescue mode as instructed in [Exiting rescue mode](#).

4. After the instance exits the rescue mode, it is still shut down. Start it up as instructed in [Starting Up Instances](#). Then, verify whether the system can be started up and logged in normally.

5. After successful login, if you need to configure the auto-attaching of the disk, see [Cloud Disk Auto-mount Failed upon Linux CVM Restart](#).

sshd Configuration File Permissions

Last updated : 2024-01-06 17:32:18

Issue Description

During login to a Linux instance via SSH key, "ssh_exchange_identification: Connection closed by remote host" or "no hostkey alg" is displayed.

Common Causes

sshd configuration file permissions, such as the permissions of the `/var/empty/sshd` or `/etc/ssh/ssh_host_rsa_key` configuration file, are modified, which may cause a failure in login via SSH key.

Solution

Perform the steps based on the actual error message to modify the configuration file permissions:

If the error message is "ssh_exchange_identification: Connection closed by remote host", see [Modifying permissions of `/var/empty/sshd` file](#).

If the error message is "no hostkey alg", see [Modifying permissions of `/etc/ssh/ssh_host_rsa_key` file](#).

Troubleshooting Procedure

Modifying permissions of `/var/empty/sshd` file

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the error cause:

```
ssh -t
```

Information similar to the following is returned:

```
"/var/empty/sshd must be owned by root and not group or world-writable."
```

3. Run the following command to modify the permissions of the `/var/empty/sshd/` file:

```
chmod 711 /var/empty/sshd/
```

Modifying permissions of `/etc/ssh/ssh_host_rsa_key` file

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the error cause:

```
sshd -t
```

The returned information contains the following field:

```
"/etc/ssh/ssh_host_rsa_key are too open"
```

3. Run the following command to modify the permissions of the `/etc/ssh/ssh_host_rsa_key` file:

```
chmod 600 /etc/ssh/ssh_host_rsa_key
```

Infinite Loop Call in /etc/profile

Last updated : 2024-01-06 17:32:18

Issue Description

When you use SSH to log in to a Linux instance, the SSH command gets stuck after outputting the "Last login: " information.

Common Causes

This problem is probably because the `/etc/profile` file was modified, so `/etc/profile` was called in `/etc/profile`, resulting in an infinite loop call and inability to log in successfully.

Solution

Check and repair the `/etc/profile` file as instructed in [Steps](#).

Steps

1. [Log in to the Linux instance via VNC](#).
2. Run the following command to view the `/etc/profile` file.

```
vim /etc/profile
```

3. Check whether the `/etc/profile` file contains commands related to `/etc/profile`.

If yes, go to the next step.

If no, [submit a ticket](#) for assistance.

4. Press **i** to enter the edit mode and add `#` before the relevant commands in `/etc/profile` to comment them.
5. Press **Esc** to exit the edit mode, and enter `:wq` to save the modification.
6. Log in to the Linux instance again by [using SSH](#).

Login Failure Due to Server Isolation

Last updated : 2024-01-06 17:32:18

CVM may be blocked and isolated due to security violations (content or behavior violations) or DDoS attacks. This document describes how to solve the problem of login failure due to CVM isolation caused by security violations.

Problems

The CVM isolation may be brought by violations of applicable laws and regulations. You can check whether the CVM is isolated in the following ways.

When a CVM is isolated from the public network, you will be notified of the isolation via an [internal message in the console](#) or a text message.

The **Monitoring/Status** tab in the [CVM Console](#) displays the status of the CVM: Isolated.

Causes

When a CVM is involved in a violation event or risk event, it will be partially isolated (except for the private network login interfaces 22, 36000, and 3389, all other network access is isolated, and the developers can log in to the CVM through a jump server).

Solutions

1. Delete the violating content as instructed by the internal message or text message. Resolve security risks and reinstall the system if necessary.
2. If the violation is not caused by your own action, your CVM may have encountered malicious intrusion. To resolve this, see [Host Security](#).
3. After eliminating potential safety hazards or stopping violations, please [submit a ticket](#) to contact customer service for removing the isolation.

Login Failure Due to High Bandwidth Occupation

Last updated : 2024-01-06 17:32:18

This document describes how to diagnose and troubleshoot Linux or Windows CVM login issues caused by high bandwidth utilization.

Problems

In the [CVM console](#), the bandwidth monitoring data of the CVM shows that bandwidth usage is too high, and connection to CVM fails.

A [Health check](#) shows that the bandwidth utilization is too high.

Fault Locating and Troubleshooting

1. Log in to the target CVM instance using VNC:

For Windows instances, see [Logging into Windows Instance via VNC](#).

For Linux instances, see [Logging In to Linux Instances \(VNC\)](#).

2. Troubleshoot:

Windows CVM instances

Linux CVM instances

After logging in to the Windows CVM instance via VNC, perform the following operations:

Note:

The following operations use a CVM instance running the Windows Server 2012 operating system as an example.

1. In the CVM instance, click



. Select **Task Manager** to open the **Task Manager** page.

2. Select **Performance** tab, and click **Open resource monitor**.

3. On the **Resource Monitor** page, identify the process that consumes a lot of bandwidth. Based on your actual business, determine whether the process is normal.

If this process is a service process, check whether the high bandwidth utilization is caused by changes in access traffic and whether you need to optimize the capacity or upgrade the CVM configuration as instructed in [Changing Instance Configuration](#).

If this process has an exception, the high bandwidth utilization may be caused by a virus or a trojan. If so, you can manually terminate the process or use security software to kill the virus. You can also back up data and then reinstall the operating system.

Note:

In Windows, many virus processes can disguise themselves as system processes. You can select **Task Manager > Processes** to check the process information and preliminarily identify the virus.

Normal system processes have complete signatures and descriptions, and most of them are located in the `C:\Windows\System32` directory. While virus programs may have the same names as system processes, they lack signatures and descriptions. In addition, their locations are often abnormal.

If this process is a Tencent Cloud component process, please [submit a ticket](#), and we will help you locate and troubleshoot the problem.

After logging in to the Linux CVM instance via VNC, perform the following operations:

Note:

The following operations use a CVM instance running the CentOS 7.6 operating system as an example.

1. Run the following command to install the iftop tool. This tool monitors traffic for Linux CVM instances.

```
yum install iftop -y
```

Note:


For Ubuntu system, run the `apt-get install iftop -y` command.

2. Run the following command to install lsof.

```
yum install lsof -y
```

3. Run the following command to run iftop.

```
iftop
```



"<=" and ">=" indicate the direction of the traffic.

"TX" indicates the traffic is outbound.

"RX" indicates the traffic is inbound.

"TOTAL" indicates the total traffic.

"Cum" indicates the total traffic from the moment iftop started to run until now.

"peak" indicates the traffic peak.

"rates" indicates the average traffic over the last 2, 10, and 40 seconds.

4. Based on the IP address of the consumed traffic in iftop, run the following command to check the process connected to this IP address.

```
lsof -i | grep IP
```

For example, if the IP address of the consumed traffic is 201.205.141.123, run the following command:

```
lsof -i | grep 201.205.141.123
```

If the following result is returned, the majority of the CVM bandwidth is consumed by the SSH process.

```
sshd      12145    root      3u  IPV4  3294018      0t0    TCP  10.144.90.86:ssh->203.205.141.123:58614 (ESTABLISHED)
sshd      12179    ubuntu    3u  IPV4  3294018      0t0    TCP  10.144.90.86:ssh->203.205.141.123:58614 (ESTABLISHED)
```

5. View the process that consumes a lot of bandwidth and check whether the process is normal.

If this process is a service process, check whether the high bandwidth utilization is caused by changes in access traffic and whether you need to optimize the capacity or upgrade the CVM configuration as instructed in [Changing Instance Configuration](#).

If this process has an exception, the high bandwidth utilization may be caused by a virus or a trojan. If so, you can manually terminate the process or use security software to kill the virus. You can also back up data and then reinstall the operating system.

If this process is a Tencent Cloud component process, please [submit a ticket](#), and we will help you locate and troubleshoot the problem.

We strongly recommend that you query the location of the destination IP on the [IP138 query website](#). Note that the security risk is greater if the destination IP is found to be located in a region outside the Chinese mainland.

Remote Connect Failure Due to Security Group Settings


Last updated : 2024-01-06 17:32:18

This document describes how to troubleshoot the problem where you are unable to connect remotely to a CVM due to security group configuration issues.

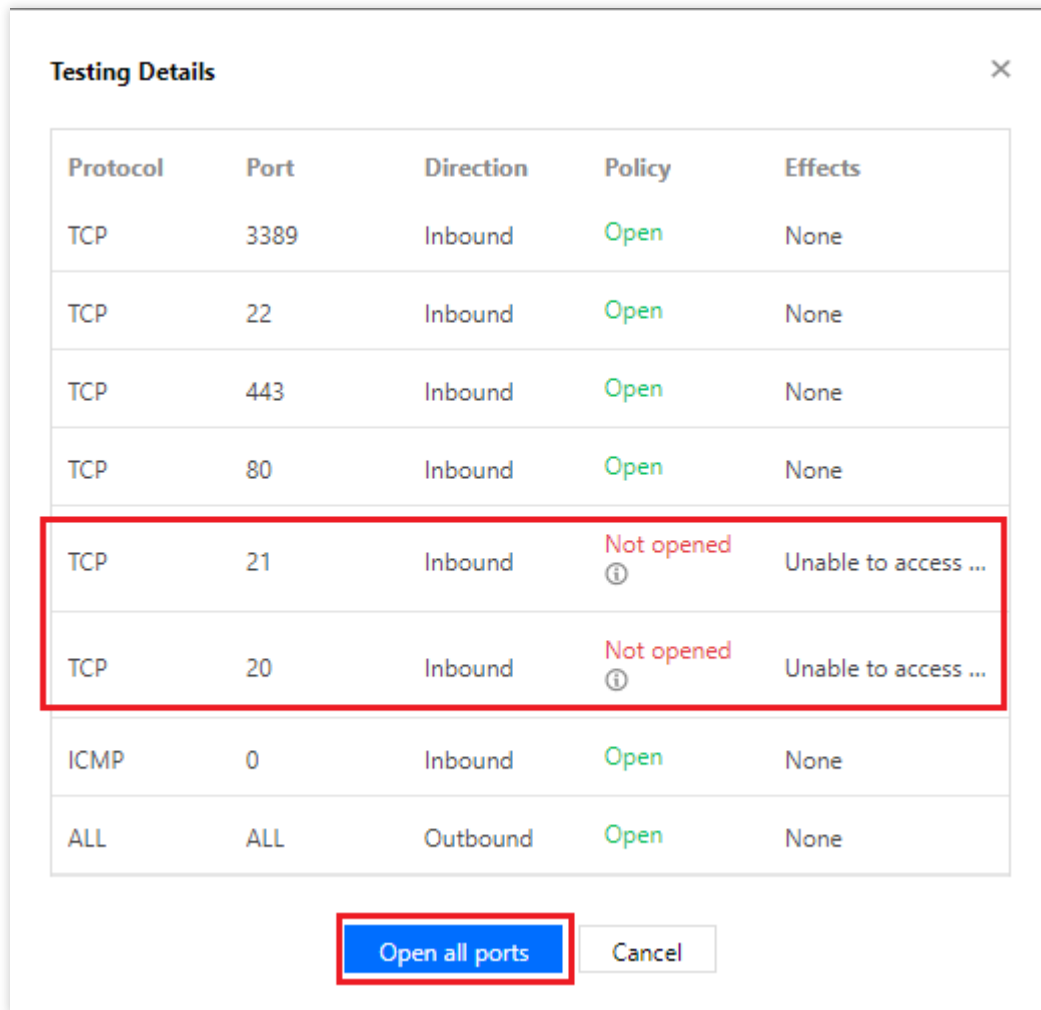
Diagnostic Tool

You can use [Port Verification](#) to check whether the problem is caused by security group configurations.

1. Log in to the [Port Verification](#).
2. In the **Port Verification** page, select the instance to check and click **Quick Check**. This is shown in the following figure:



If the check shows that the instance has ports that are not open, you can select **Open all ports** to open all commonly used ports of the CVM to the Internet, and try to log in remotely again.



Modifying security group configurations

If you do not want to use **Open all ports** to open all commonly used ports of the CVM to the Internet, or you need to customize the remote login port, you can use custom configuration of the inbound and outbound rules of the security group to resolve remote login failures. For more information, see [Modifying Security Group Rules](#).

Troubleshooting Linux Instance Issues via VNC and Rescue Mode

Last updated : 2024-01-06 17:32:18

Generally, you can troubleshoot most Linux system issues via VNC and rescue mode. This document describes how to troubleshoot the issues such as a failure to log in to a Linux Instance via SSH Key and system failures.

Troubleshooting Tool

Login via VNC is a method of remotely connecting to a CVM through a Web browser. This enables you to directly observe the CVM status, or modify the configuration file in the system. Usually when you cannot remotely log in to the instance via SSH key, you can use this login method.

Rescue mode is generally used when the Linux system cannot be started up normally, or the Linux instance cannot be logged in via VNC. Common usage cases: abnormal fstab configuration, missing key system files, and damaged/missing .lib and .dll files, etc.

Identifying and Fixing Issues

Using VNC to troubleshoot the SSH key login failure

Error description

When I tried to log in to a Linux instance via SSH key, the error message **ssh_exchange_identification: Connection closed by remote host** is displayed.

```
[root@~]# ssh root@  
kex_exchange_identification: read: Connection reset by peer
```

Possible cause

The connection reset error in the kex_exchange_identification stage generally means that the ssh-related process has been started, but the configuration may be abnormal, for example, the sshd configuration file permissions have been modified.

Solutions

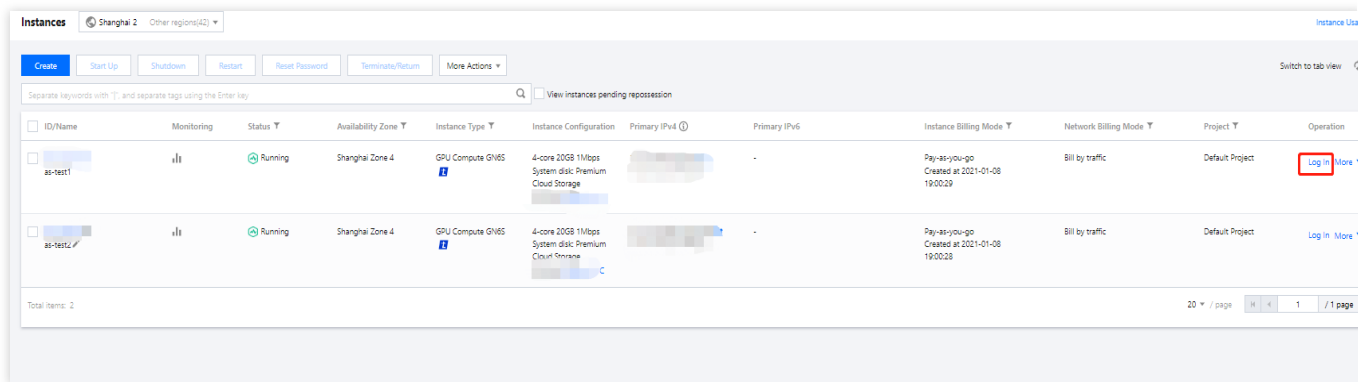
See [Steps](#) to check the sshd process, locate and fix the problem.

Steps

Follow the steps to log in to


o the Linux instance via VNC:

1.1 Log in to the [CVM console](#), find the target Linux CVM, and click **Login** under the "Operation" column.



1.2 In the **Standard login | Linux instance** window, click **Login via VNC**.

1.3 Enter the username after "login" and press **Enter**, and enter the password after "Password" and press **Enter**. The login is successful when the following information is displayed.



2. Run the following command to check whether the sshd process is running normally.

```
ps -ef | grep sshd
```

If the result shown below is returned, the sshd process is normal.

```
[root@UM-0-11-centos ~]# ps -ef | grep sshd
root      1173      1  0 22:08 ?        00:00:00 /usr/sbin/sshd -D -oCiphers=aes256-gcm@openssh.com,chacha20-poly1305@openssh.com,aes256-ctr,aes256-cbc,aes128-gcm@openssh.com,aes128-ctr,aes128-cbc -oMACs= hmac-sha2-256-etm@openssh.com,hmac-sha1-etm@openssh.com,umac-128-etm@openssh.com,hmac-sha2-512-etm@openssh.com,hmac-sha2-256,hmac-sha1,umac-128@openssh.com,hmac-sha2-512 -oGKexAlgorithms=gss-curve25519-sha256-,gss-nistp256-sha256-,gss-group14-sha256-,gss-group16-sha512-,gss-gex-sha1-,gss-group141- -oKexAlgorithms=curve25519-sha256,curve25519-sha256@libssh.org,ecdh-sha2-nistp256,ecdh-sha2-nistp384,ecdh-sha2-nistp521,diffie-hellman-group-exchange-sha256,diffie-hellman-group14-sha256,diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,diffie-hellman-group-exchange-sha1,diffie-hellman-group14-sha1 -oHostKeyAlgorithms=ecdsa-sha2-nistp256,ecdsa-sha2-nistp256-cert-v01@openssh.com,ecdsa-sha2-nistp384,ecdsa-sha2-nistp384-cert-v01@openssh.com,ecdsa-sha2-nistp521,ecdsa-sha2-nistp521-cert-v01@openssh.com,ssh-ed25519,ssh-ed25519-cert-v01@openssh.com,rsa-sha2-256,rsa-sha2-256-cert-v01@openssh.com,rsa-sha2-512,rsa-sha2-512-cert-v01@openssh.com,ssh-rsa,ssh-rsa-cert-v01@openssh.com -oPubkeyAcceptedKeyTypes=ecdsa-sha2-nistp256,ecdsa-sha2-nistp256-cert-v01@openssh.com,ecdsa-sha2-nistp384,ecdsa-sha2-nistp384-cert-v01@openssh.com,ecdsa-sha2-nistp521,ecdsa-sha2-nistp521-cert-v01@openssh.com,ssh-ed25519,ssh-ed25519-cert-v01@openssh.com,rsa-sha2-256,rsa-sha2-256-cert-v01@openssh.com,rsa-sha2-512,rsa-sha2-512-cert-v01@openssh.com,ssh-rsa,ssh-rsa-cert-v01@openssh.com -oSignatureAlgorithms=ecdsa-sha2-nistp256,ecdsa-sha2-nistp384,ecdsa-sha2-nistp521,ssh-ed25519,rsa-sha2-256,rsa-sha2-512,ssh-rsa
root      2473    1722  0 22:13 tty1      00:00:00 grep --color=auto sshd
```

3. Run the following command to view the error cause:

```
sshd -t
```

If a message similar to `"/var/empty/sshd must be owned by root and not group or world-writable."` shown below is returned,

the error can be caused by a `/var/empty/sshd/` permission issue.

```
[root@UM-0-11-centos ~]# sshd -t
/var/empty/sshd must be owned by root and not group or world-writable.
[root@UM-0-11-centos ~]#
```

You can also check the error messages in the `/var/log/secure` logs to facilitate troubleshooting.

```
systemd[1]: Started Session 174 of user root.
systemd[1]: session-174.scope: Succeeded.
systemd[1]: Started Session 175 of user root.
systemd[1]: session-175.scope: Succeeded.
systemd[1]: Started Session 176 of user root.
systemd[1]: session-176.scope: Succeeded.
sshd[123651]: fatal: /var/empty/sshd must be owned by root and not group or world-writable
```

4. Run the following command to view the permission of the `/var/empty/sshd` directory.

```
ll -d /var/empty/sshd/
```

The returned result is shown below. The permission configuration is `777`.

```
[root@UM-0-11-centos ~]# ll -d /var/empty/sshd/
drwxrwxrwx. 2 root root 4096 Jul 13  2021 /var/empty/sshd/
```

5. Run the following command to modify the permissions of the `/var/empty/sshd/` file.

```
chmod 711 /var/empty/sshd/
```

The instance can be logged in remotely after a test of [Logging into Linux Instance \(SSH Key\)](#).

Using VNC to troubleshoot Linux system startup failures

Error description

Error: Unable to remotely log in to the Linux CVM via SSH. After logging in to the CVM via VNC, you can check the system start-up failure and view the prompt message "Welcome to emergency mode".

```
[ OK ] Reached target Remote File Systems (Pre).
[ OK ] Reached target Remote File Systems.
      Starting Crash recovery kernel arming...
[ OK ] Started Security Auditing Service.
      Starting Update UTMP about System Boot/Shutdown...
[ OK ] Started Update UTMP about System Boot/Shutdown.
      Starting Update UTMP about System Runlevel Changes...
[ OK ] Started Update UTMP about System Runlevel Changes.
Welcome to emergency mode! After logging in, type "journalctl -xb" to view c   i
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or press Control-D to continue):
```

Possible cause

The `/etc/fstab` is not properly configured.

For example, you've configured the auto-attaching of disk based on device name in the `/etc/fstab` file. If the device name is changed when the CVM restarts, this configuration will cause the system to fail to start up normally.

Solutions

See [Steps](#) to repair the `/etc/fstab` configuration file. Then, restart the CVM to verify the repaired file.

Steps

1. Follow the [step](#) to log in to the Linux instance via VNC.
2. After entering the VNC interface, you see the interface shown in [Error Description](#). Enter the root account password (which is not displayed by default) and press **Enter** to log in to the server.

```
Give root password for maintenance
(or press Control-D to continue):
[root@i3-8-11-centos ~]#
```

3. After entering the system, run the following command to check whether the drive letter information in the `fstab` file is correct.

```
lsblk
```

If the result shown below is returned, the drive letter in the file is incorrect:

```
[root@localhost ~]# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sr0         11:0    1 184.1M  0 rom
vda         253:0    0   50G  0 disk
└─vda1      253:1    0   50G  0 part /
[root@localhost ~]# cat /etc/fstab

#
# /etc/fstab
# Created by anaconda on Tue Nov 26 02:11:36 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
UUID=659e6f89-71fa-463d-842e-ccdf2c06e0fe /          ext4 defaults
/dev/vdb1    /data     ext3     defaults 0        0
[root@localhost ~]#
```

4. Run the following command to back up the `fstab` file.

```
cp /etc/fstab /home
```

5. Run the following command to use VI editor to open the `/etc/fstab` file.

```
vi /etc/fstab
```

6. Press `i` to enter the edit mode. Move the cursor to the beginning of the error line and enter `#` to comment out this configuration.

```
#
# /etc/fstab
# Created by anaconda on Tue Nov 26 02:11:36 2019
#
# Accessible filesystems, by reference, are maintained under '/dev/disk/'.
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info.
#
# After editing this file, run 'systemctl daemon-reload' to update systemd
# units generated from this file.
#
UUID=659e6f89-71fa-463d-842e-ccdf2c06e0fe / ext4 defaults 1
# /dev/vdb1 /data ext3 defaults 0 0
```

7. Press **ESC**, enter **:wq**, and press **Enter** to save the configuration and exit the editor.

8. Restart the instance in the CVM console. For more information, see [Restarting Instances](#).

9. Check if it can be started and logged in properly.

Using the rescue mode to troubleshoot Linux system startup failures

Error description

The instance cannot be started up normally after the Linux system restarts. There are many **FAILED** startup failure items in the prompt message.

```
[ OK ] Reached target Local File Systems (Pre).
[ OK ] Reached target Local File Systems.
        Starting Restore /run/initramfs on shutdown...
        Starting Tell Plymouth To Write Out Runtime Data...
        Starting Create Volatile Files and Directories...
[FAILED] Failed to start Restore /run/initramfs on shutdown.
See 'systemctl status dracut-shutdown.service' for details.
[ OK ] Started Tell Plymouth To Write Out Runtime Data.
[FAILED] Failed to start Create Volatile Files and Directories.
See 'systemctl status systemd-tmpfiles-setup.service' for details.
        Starting Security Auditing Service...
[FAILED] Failed to start Security Auditing Service.
See 'systemctl status auditd.service' for details.
        Starting Update UTMP about System Boot/Shutdown...
[FAILED] Failed to start Update UTMP about System Boot/Shutdown.
See 'systemctl status systemd-update-utmp.service' for details.
[DEPEND] Dependency failed for Update UTMP about System Runlevel Changes.
[ OK ] Reached target System Initialization.
[ OK ] Listening on D-Bus System Message Bus Socket.
[ OK ] Listening on Open-iSCSI iscsid Socket.
[ OK ] Started daily update of the root trust anchor for DNSSEC.
[ OK ] Started Daily Cleanup of Temporary Directories.
```



```
[ OK ] Started dnf makecache --timer.
[ OK ] Reached target Timers.
[ OK ] Listening on ACPID Listen Socket.
[ OK ] Listening on SSSD Kerberos Cache Manager responder socket.
[ OK ] Listening on Open-iSCSI iscsiuiio Socket.
[ OK ] Reached target Sockets.
[ OK ] Reached target Basic System.
      Starting Authorization Manager...
[ OK ] Started libstoragemgmt plug-in server daemon.
[ OK ] Started Machine Check Exception Logging Daemon.
      Starting System Security Services Daemon...
[ OK ] Started ACPI Event Daemon.
      Starting Hardware RNG Entropy Gatherer Wake threshold service...
[FAILED] Failed to start NTP client/server.
See 'systemctl status chronyd.service' for details.
      Starting VDO volume services...
[ OK ] Started D-Bus System Message Bus.
      Starting Network Manager...
[ OK ] Reached target sshd-keygen.target.
[FAILED] Failed to start Hardware RNG Entropy Gatherer Wake threshold service.
See 'systemctl status rngd-wake-threshold.service' for details.
[DEPEND] Dependency failed for Hardware RNG Entropy Gatherer Daemon.
[FAILED] Failed to start VDO volume services.
See 'systemctl status vdo.service' for details.
[ OK ] Started D-Bus System Message Bus.
```

Possible cause

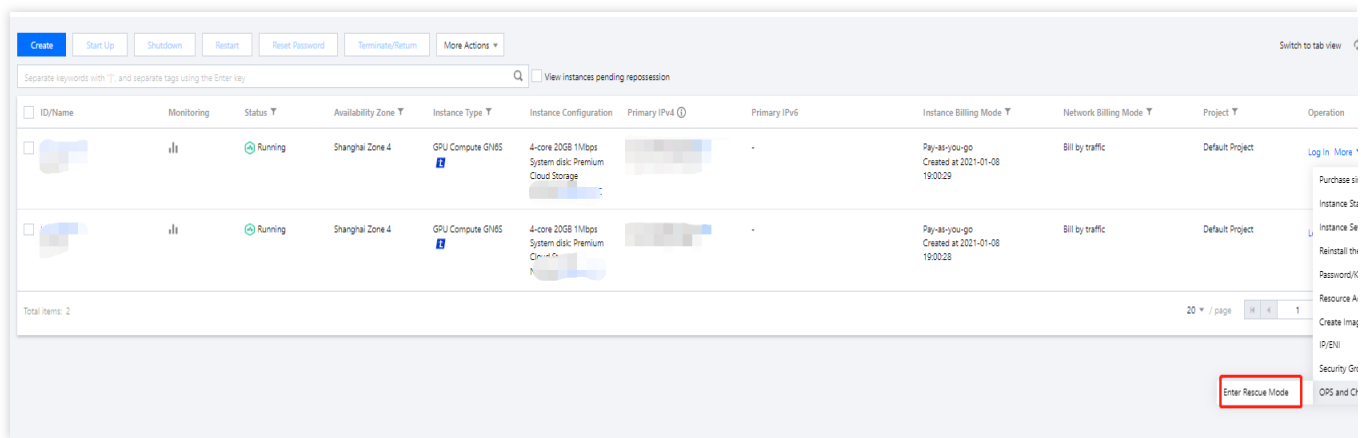
The key system files such as `.bin` and `.lib` files are missing.

Solutions

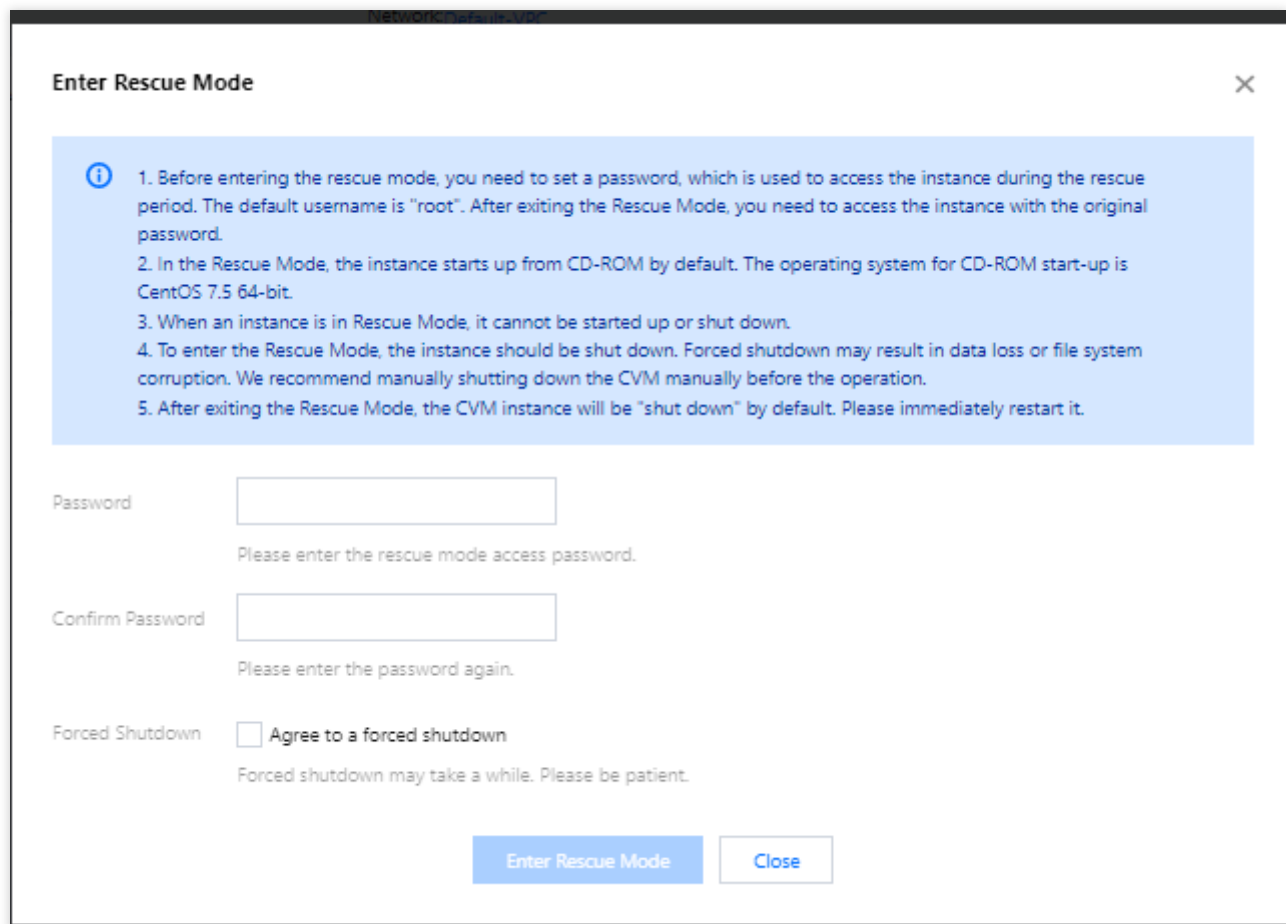
See [Steps](#) to enter the instance rescue mode through the console to troubleshoot.

Steps

1. Before you enter the rescue mode, we strongly recommend you back up the instance to avoid the impact of maloperations. You can [create snapshots](#) to back up cloud disks and [create a custom image](#) to back up local system disks.
2. Log in to the [CVM console](#). On the "Instances" page, find the target instance, select **More > Ops and check > Enter rescue mode**.



3. In the pop-up window, set the instance login password for the rescue mode.



4. Click **Enter rescue mode**, and the instance status will change to "Entering rescue mode", which typically completes within a few minutes:

ID/Name	Monitoring	Status	Availability Zone	Instance Type	Instance Configuration	Primary IPv4	Primary IPv6	Instance Billing Mode	Network Billing Mode	Project	Operation
		Entering Rescue Mode...	Shanghai Zone 4	GPU Compute G4ES	4-core 20GB 11Mbps System disk: Premium Cloud			Pay-as-you-go Created at 2021-01-08 19:00:29	Bill by traffic	Default Project	
		Stop	Shanghai Zone 4	GPU Compute G4ES	4-core 20GB 11Mbps System disk: Premium Cloud	171.171.171.171		Pay-as-you-go	Bill by traffic	Default Project	

The status of instance entered the rescue mode changes to "Rescue mode" with a red exclamation mark.

<div> <div>Create Start Up Shutdown Restart Reset Password Terminate/Return More Actions</div> <div>Switch to tab view</div> </div> <div> <div>Separate keywords with ", and separate tags using the Enter key</div> <div>View Instances pending reposition</div> </div>											
ID/Name	Monitoring	Status	Availability Zone	Instance Type	Instance Configuration	Primary IPv4	Primary IPv6	Instance Billing Mode	Network Billing Mode	Project	Operation
		Rescue Mode	Shanghai Zone 4	GPU Compute G4ES				Pay-as-you-go Created at 2021-01-08 19:00:29	Bill by traffic	Default Project	Log In More

5. Use the `root` account and the password set in [step 3](#) to log in to the instance as follows:

If the instance has a public IP, log in to it as instructed in [Logging in to Linux Instance \(SSH Key\)](#).

If the instance has no public IPs, log in to it as instructed in [Logging in to Linux Instances \(VNC\)](#).

6. This document takes login via VNC as an example. After successful login, run the following commands in sequence to mount the root partition of the system disk.

Note:

In rescue mode, the device name of the instance system disk is `vda`, and its root partition is `vda1`, which is unmounted by default.

```
mkdir -p /mnt/vm1

mount /dev/vda1 /mnt/vm1
```

The returned result is shown below:

```
[root@ ~]# mkdir -p /mnt/vm1
[root@ ~]# mount /dev/vda1 /mnt/vm1
[root@ ~]# _
```

7. After successful mounting, you can manipulate the data in the root partition of the original system.

You can also use the `mount -o bind` command to mount some sub-directories in the original file system and use the `chroot` command to run commands in the specified root directory. Below are the specific commands:

```
mount -o bind /dev /mnt/vm1/dev
mount -o bind /dev/pts /mnt/vm1/dev/pts
mount -o bind /proc /mnt/vm1/proc
mount -o bind /run /mnt/vm1/run
mount -o bind /sys /mnt/vm1/sys
chroot /mnt/vm1 /bin/bash
```

When running the `chroot` command:

If there is no error message, you can continue to run the `cd /` command.

If the error message as shown below appears, the root directory cannot be switched normally. In this case, you can run `cd /mnt/vm1` to view the root partition data.

```
[root@UM-0-11-centos ~]# mkdir -p /mnt/vm1
[root@UM-0-11-centos ~]# mount /dev/vda1 /mnt/vm1
[root@UM-0-11-centos ~]# mount -o bind /dev /mnt/vm1/dev
[root@UM-0-11-centos ~]# mount -o bind /dev/pts /mnt/vm1/dev/pts
[root@UM-0-11-centos ~]# mount -o bind /proc /mnt/vm1/proc
[root@UM-0-11-centos ~]# mount -o bind /run /mnt/vm1/run
[root@UM-0-11-centos ~]# mount -o bind /sys /mnt/vm1/sys
[root@UM-0-11-centos ~]# chroot /mnt/vm1 /bin/bash
chroot: failed to run command '/bin/bash': No such file or directory
[root@UM-0-11-centos ~]#
```

8. Through the command, you can check that all files in the `/usr/bin` directory in the original system root partition have been deleted.

```
bin boot data dev etc home lib lib64 lost+found media mnt opt proc root run sbin srv sys tmp usr
[root@UM-0-11-centos vm1]# ll
total 72
lrwxrwxrwx. 1 root root 7 Nov 3 2020 bin -> usr/bin
dr-xr-xr-x. 5 root root 4096 Apr 14 17:53 boot
drwxr-xr-x. 2 root root 4096 Dec 10 2019 data
drwxr-xr-x. 19 root root 3268 Apr 14 18:09 dev
drwxr-xr-x. 100 root root 12288 Apr 14 17:53 etc
drwxr-xr-x. 2 root root 4096 Jun 28 2021 home
lrwxrwxrwx. 1 root root 7 Nov 3 2020 lib -> usr/lib
lrwxrwxrwx. 1 root root 9 Nov 3 2020 lib64 -> usr/lib64
drwx----- 2 root root 16384 Nov 26 2019 lost+found
drwxr-xr-x. 2 root root 4096 Nov 3 2020 media
drwxr-xr-x. 2 root root 4096 Nov 3 2020 mnt
drwxr-xr-x. 2 root root 4096 Nov 3 2020 opt
dr-xr-xr-x. 125 root root 8 Apr 14 18:08 proc
dr-xr-xr-x. 5 root root 4096 Mar 10 19:24 root
drwxr-xr-x. 37 root root 1140 Apr 14 18:10 run
lrwxrwxrwx. 1 root root 8 Nov 3 2020 sbin -> usr/sbin
drwxr-xr-x. 2 root root 4096 Nov 3 2020 srv
dr-xr-xr-x. 13 root root 8 Apr 14 18:12 sys
drwxrwxrwt. 8 root root 4096 Apr 14 17:56 tmp
drwxr-xr-x. 12 root root 4096 Jun 10 2021 usr
drwxr-xr-x. 20 root root 4096 Jun 10 2021 var
[root@UM-0-11-centos vm1]# cd /usr/bin/
[root@UM-0-11-centos bin]# pwd
/mnt/vm1/usr/bin
[root@UM-0-11-centos bin]# ls
-
[root@UM-0-11-centos bin]#
```

9. In this case, you can create a normal instance using the same operating system, and run the following commands to compress and remotely copy the files in the `/usr/bin` directory of the normal system to the abnormal instance. For the normal instance, run the following commands in sequence:

```
cd /usr/bin/ && tar -zcvf bin.tar.gz *

scp bin.tar.gz root@abnormal instance ip:/mnt/vm1/usr/bin/
```

Note:

If the instances have public network IPs, the copy can be performed through the public network; otherwise, the copy is performed through the private network.

The execution result is as shown below:

```
[root@M-10-12-centos bin]# scp bin.tar.gz root@81.70.163.225:/mnt/vm1/usr/bin/
The authenticity of host '81.70.163.225 (81.70.163.225)' can't be established.
ECDSA key fingerprint is SHA256:ey4JYiXm44, GLmHaVo8ihDvNtbeA.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '81.70.163.225' (ECDSA) to the list of known hosts.
root@81.70.163.225's password:
bin.tar.gz                                100% 33MB 121.9KB/s
[root@M-10-12-centos bin]#
```

For the abnormal instance, run the following commands in sequence in the rescue mode:

```
cd /mnt/vm1/usr/bin/

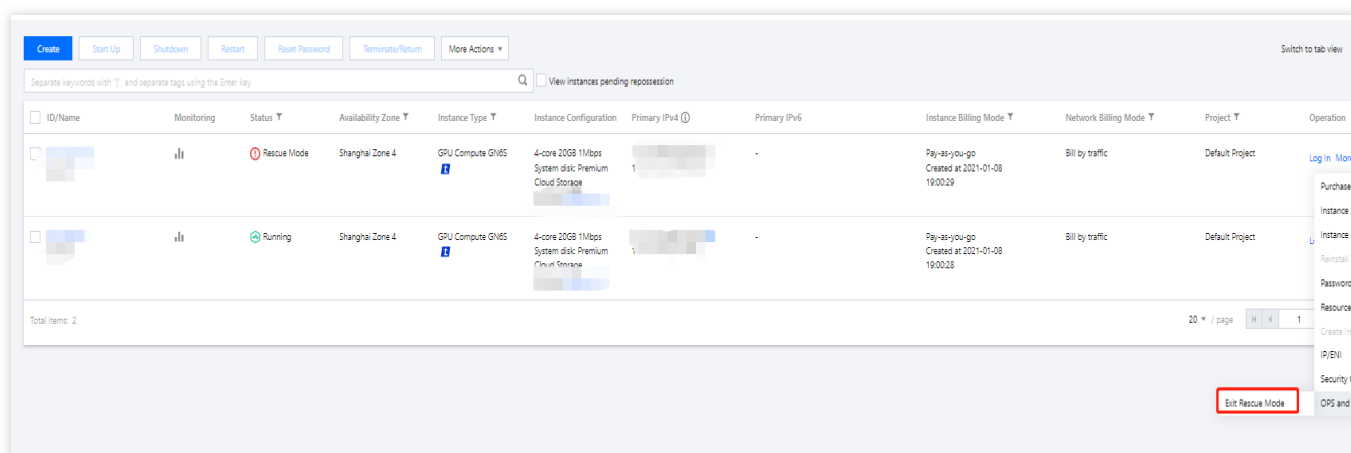
tar -zxvf bin.tar.gz

chroot /mnt/vm1 /bin/bash
```

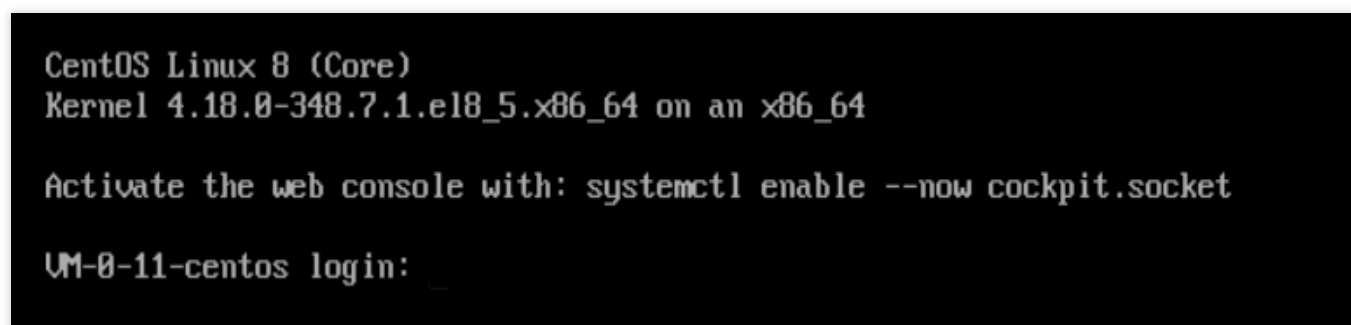
The execution result is as shown below:

```
[root@M-8-11-centos /]# chroot /mnt/vm1 /bin/ba
baobab      base64      basename   bash        bashbug     bashbug-64  batch
[root@M-8-11-centos /]# chroot /mnt/vm1 /bin/bash
[root@M-8-11-centos /]#
```

10. After repairing the instance, select **More > Ops and check > Exit rescue mode** under the **Operation** column of the target instance.



11. After exiting the rescue mode, the instance is in a shutdown status. Start up the instance to verify the system. As shown below, the system has been restored.



Failed to shut down or restart a CVM

Last updated : 2024-01-06 17:32:18

When you shut down or restart the CVM, a failure may occur. While it is a rare event, you can troubleshoot as follows:

Possible Causes

High CPU or memory usage.

ACPI has not been installed on the Linux CVM.

System update of the Windows CVM takes too long.

Windows CVM has not completed initialization yet when you purchase it for the first time.

The operating system is damaged due to installed software or viruses such as Trojan.

Troubleshooting

Check CPU/memory usage

1. Check CPU/memory usage based on the operating system of the CVM.

For Windows CVM: Right-click the "Taskbar" and select **Task Manager** on the CVM.

For Linux CVM: Execute the `top` command to view information in `%CPU` and `%MEM` columns.

2. Terminate processes with high CPU or memory usage.

If you still cannot shut down or restart the CVM, please execute [forced shutdown or restart](#).

Check whether ACPI has been installed

Note:

This operation is for Linux CVM.

Execute the following command to see if an ACPI process exists.

```
ps -ef | grep -w "acpid" | grep -v "grep"
```

If an ACPI process exists, please execute [forced shutdown or restart](#).

If no ACPI process exists, please install ACPI. For specific operations, see [Linux Power Management Configuration](#).

Check whether WindowsUpdate is running

Note:

This operation is for Windows CVM.

On the operating system interface of Windows CVM, click **Start > Control Panel > Windows Update** to see if any patches or programs are being updated.

Windows may perform patching when the system is shutting down. The update may take a long time, causing CVM shutdown/restart to fail. We recommend you wait for the Windows update to complete and then try to shut down or restart the CVM.

If no patches or programs are being updated, please execute [forced shutdown or restart](#).

Check whether the CVM has completed initialization

Note:

This operation is for Windows CVM.

When you purchase Windows CVM for the first time, initialization may take longer because Sysprep is used to distribute images. Before the initialization is complete, Windows will ignore shutdown and restart operations.

If the Windows CVM you purchased is initializing, we recommend you wait for the initialization to complete before shutting down or restarting the CVM again.

If the CVM has completed initialization, please execute [forced shutdown or restart](#).

Check whether the software installed is normal

Use a check tool or antivirus software to see if the software installed on the CVM is normal or attacked by viruses such as Trojan.

If an exception is found, the system may be damaged, causing shutdown and restart to fail. We recommend you uninstall the software, back up data or scan with security software, and then reinstall the system.

If no exception is found, please execute [forced shutdown or restart](#).

Forced shutdown/restart

Note:

Forced shutdown/restart provided by Tencent Cloud can be used if you fail to shut down or restart the CVM after multiple attempts. This feature allows you to force a shutdown or restart on the CVM, which may cause data loss or damage the file system.

1. Log in to the [CVM Console](#).
2. On the instance management page, select the CVM you want to shut down or restart.

Shut down CVM: Click **More > Instance Status > Shutdown**.

Restart CVM: Click **More > Instance Status > Restart**.

3. In the **Shutdown** or **Restart Instance** window that pops up, check **Forced Shutdown** or **Forced Restart**, and Click **Ok**.

Check **Forced Shutdown**, as shown below:

Shutdown ×

You have selected **1 Instance**, [Learn More](#) ▼

No.	Instance Name	Instance ID	Operation
1	Unnamed	<div></div>	Can be shut down

Are you sure you want to shut down the selected instances?

☐ CVM No Charge when Shut down

No Charge When Shut Down is available when the following conditions are met:

- Pay-as-you-go Instances
- The instance's system disk and the data disk are both cloud disks.
- Non-GPU-and FPGA-based instances

☒ **Forced shutdown**

Forced shutdown may lead to data loss or file system damage. This is only allowed when the instance cannot be shut down normally.


OK

Cancel

Check **Forced Restart**, as shown below:

Restart Instance

You have selected **1 Instance** , [Learn More](#) ▼

No.	Instance Name	Instance ID	Current Bandwidth C...
1	Unnamed		1 Mbps

Are you sure you want to restart the selected instances?

During restarting, this instance cannot work and your service may be affected.

☒ Forced restart

Just like turning off the computer and then powering it on, forced restart may lead to data loss or damage to file system. This is allowed only when the instance cannot be restarted normally.

OK

Cancel

Kernel and IO Issues

Last updated : 2024-01-06 17:32:18

The status check provides a report of instance exceptions. This document mainly describes the symptoms, causes and solutions of kernel and IO problems shown in the status check report.

Troubleshooting Kernel Failures

Problems

The kernel failure may cause login failure or abnormal restart.

Common causes

Kernel hung_task

The kernel hung task is based on a single kernel thread named as `khungtaskd`, which monitors processes in the `TASK_UNINTERRUPTIBLE` status. If a process stuck in D state during the period specified by `kernel.hung_task_timeout_secs` (defaults to 120 seconds), the stack information of this hung task process will be printed.

If `kernel.hung_task_panic=1` is configured, the hung task will trigger kernel panic and system restart.

Kernel soft lockup

A soft lockup refers to a kernel thread using and not releasing a CPU, without giving other tasks a chance to run. Each CPU is assigned with a timed kernel thread `watchdog/x`. If this thread is not executed during the specified period (the default period is two times the `kernel.watchdog_thresh` value. For example, the default `kernel.watchdog_thresh` value is 10 seconds for a 3.10 kernel), soft lockup occurs.

If `kernel.softlockup_panic=1` is configured, the soft lockup will trigger kernel panic and system restart.

Kernel panic

A kernel panic refers to a kernel crash that causes the abnormal restart. The kernel panic will be generally caused by:

Kernel hung_task, with `kernel.hung_task_panic=1` configured.

Kernel soft lockup, with `kernel.softlockup_panic=1` configured.

Kernel bug

Solution

Due to the difficulty, we recommend you [submit a ticket](#) for the troubleshooting.

Troubleshooting Disk Failures

The inode is full

Problem: the error message “No space left on device” is prompted when you create a file. After you run the `df -i` command, you will see inode is 100% used.

Common causes: the file system exhausted all inodes.

Procedure: delete useless files or expand the disk.

The disk space is full

Problem: the error message “No space left on device” is prompted when you create a file. After you run the `df -h` command, you will see the disk space is 100% used.

Common causes: the disk space runs out.

Procedure: delete useless files or expand the disk.

The disk is read-only

Problem: the file system can read files only without creating one.

Common cause: the file system is damaged.

Procedure:

1. Create a snapshot to back up the disk data. For detailed directions, see [Creating Snapshots](#).
2. Perform the troubleshooting procedure according to the disk type.

System disk

Data disk

We recommend directly restarting the instance, please see [Restart Instances](#).

1. Run the following command to check the type of the read-only disk file system.

```
lsblk -f
```

2. Run the following command to detach the data disk.

```
umount <mount point of the data disk>
```

3. Run the file system-specific command to fix the file system.

Run the following command on the **ext3/ext4** file system:

```
fsck -y /dev/[data disk]
```

Run the following command on the **xfs** file system:

```
xfs_repair /dev/[data disk]
```

The disk %util is high

Problem: the instance lags, and responds slowly or stop responding to the SSH or VNC login.

Common cause: high IO causes the disk %util to reach 100%.

Procedure: check the high IO status, and assess whether to reduce IO reads/writes, or use a disk with higher performance.

Missing System bin or lib Soft Link

Last updated : 2024-01-06 17:32:18

Issue Description

During command execution or system startup, an error message that the command or lib was not found is reported.

Common Causes

The bin, sbin, lib, and lib64 of CentOS 7, CentOS 8, and Ubuntu 20 are soft links as follows:

```
lrwxrwxrwx 1 root root 7 Jun 19 2018 bin -> usr/bin
lrwxrwxrwx 1 root root 7 Jun 19 2018 lib -> usr/lib
lrwxrwxrwx 1 root root 9 Jun 19 2018 lib64 -> usr/lib64
lrwxrwxrwx 1 root root 8 Jun 19 2018 sbin -> usr/sbin
```

If a soft link is deleted, an error will be reported during command execution or system startup.

Solution

Check and create the required soft links as instructed in [Steps](#).

Steps

1. Enter the rescue mode.
2. Run commands such as `mount` and `chroot`. When running the `chroot` command:
If there is an error, run `cd /mnt/vml`.
Otherwise, run `cd /`.
3. Run the following command to check whether the corresponding soft link exists.

```
ls -al / | grep -E "lib|bin"
```

If yes, [submit a ticket](#) for assistance.

If no, run the following commands as needed to create corresponding soft links.

```
ln -s usr/lib64 lib64
ln -s usr/sbin sbin
```

```
ln -s usr/bin bin  
ln -s usr/lib lib
```

4. Run the following command to check the soft links.

```
chroot /mnt/vml /bin/bash
```

If no error is reported, the soft links have been successfully repaired.

5. Exit rescue mode and start the system.

Suspected Infection with Virus

Last updated : 2024-01-06 17:32:18

CVMs may be intruded by hackers due to weak passwords and vulnerabilities of open-source components. This document describes how to determine whether a CVM has been infected with a virus and how to fix it.

Troubleshooting the Issue

Use [SSH](#) or [VNC](#) to log in to the instance and check whether it has been infected with a virus in the following ways:

Malicious commands were added to rc.local

Run the following command to view the `rc.local` file.

```
cat /etc/rc.local
```

If the output information contains a command not added by the business team or public image, such as `wget xx` and `/tmp/xx`, the CVM has probably been infected with a virus.

Malicious tasks were added to crontab

Run the following command to list the current schedule.

```
crontab -l
```

If the output information contains a command not added by the business team or public image, such as `wget xx` and `/tmp/xx`, the CVM has probably been infected with a virus.

Dynamic library hijacking was added to ld.so.preload

Run the following command to view the `/etc/ld.so.preload` file.

```
cat /etc/ld.so.preload
```

If the output information contains a dynamic library not added by the business team, the CVM has probably been infected with a virus.

Huge page memory was configured in sysctl.conf

Run the following command to check the usage of huge page memory.

```
sysctl -a | grep "nr_hugepages "
```

If the output is not 0, and the business program does not use huge page memory, the CVM has probably been infected with a virus.

Troubleshooting Procedure

1. Back up the system data as instructed in [Creating Snapshots](#).
2. Reinstall the instance system as instructed in [Reinstalling System](#) and take the following security hardening measures:

Change the CVM password to a stronger password containing 12-16 characters, including uppercase letters, lowercase letters, special characters, and numbers. For more information, see [Resetting Instance Password](#).

Delete unused CVM login accounts.

Change the default sshd port 22 to a less common port between 1024-65525. For more information, see [Modifying the Default Remote Port of CVM](#).

Manage the associated security group rules to open only ports and protocols required by your business. For more information, see [Adding Security Group Rules](#).

Close the port for internet access for core applications such as MySQL and Redis databases.

Install security software (such as CWPP agent), and configure real-time alarms to get noticed about suspicious logins instantly.

"no space left on device" Error During File Creation

Last updated : 2024-01-06 17:32:18

Issue Description

When you create a file in the Linux CVM, the error "no space left on device" is reported.

Common Causes

The disk space is full

The file system's `inode` is full

`df` and `du` are different

Files have been deleted, but there are still processes holding corresponding file handles, resulting in the disk space not being released.

Mounts are nested. For example, if the `/data` directory of the system disk uses a lot of space, and `/data` is used as a mount point of other data disks, then the `df` and `du` of the system disk will be different.

Solution

Troubleshoot the problems as instructed in [Steps](#).

Steps

Solving the problem of full disk space

1. Log in to the Linux instance [in the standard login method](#).
2. Run the following command to check the disk utilization.

```
df -h
```

3. Locate the mount point with a high disk utilization and run the following command to enter the mount point.

```
cd mount point
```

For example, to enter the system disk mount point, run `cd /`.

4. Run the following command to find directories that occupy a large space.

```
du -x --max-depth=1 | sort -n
```

Perform the following steps for the located directory with the largest occupied space:

If the directory capacity is much lower than the total disk space, proceed to the [Solving the problem of `df` and `du` inconsistency](#) step.

If the directory capacity is large, perform [step 2](#) to locate files that occupy a large space and evaluate whether they can be deleted based on the business conditions. If not, expand the disk capacity as instructed in [Expanding Cloud Disks](#).

Solving the problem of full file system inode

1. Log in to the Linux instance [in the standard login method](#).
2. Run the following command to check the disk utilization.

```
df -h
```

3. Locate the mount point with a high disk utilization and run the following command to enter the mount point.

```
cd mount point
```

For example, to enter the system disk mount point, run `cd /`.

4. Run the following command to find the directory with the largest number of files. This command is time-consuming, so wait patiently.

```
find / -type f | awk -F / -v OFS=/ '{ $NF="" }; dir[$0]++ } END { for (i in dir) print dir[i] " " i }' | sort -k1 -nr | head
```

Solving the problem of `df` and `du` inconsistency

Solving the problem of processes occupying file handles

Run the following command to view the processes occupying files.

```
lsof | grep delete
```

Perform the following steps according to the returned result:

Kill the corresponding processes.

Restart the service.

If many processes occupy file handles, restart the server.

Solving the problem of nested mounts

1. Run the `mount` command to mount a highly utilized disk to `/mnt` ; for example:

```
mount /dev/vda1 /mnt
```

2. Run the following command to enter `/mnt` .

```
cd /mnt
```

3. Run the following command to find directories that occupy a large space.

```
du -x --max-depth=1 | sort -n
```

According to the returned result, evaluate whether the directories or files can be deleted based on the business conditions.

4. Run the `umount` command to unmount the disk; for example:

```
umount /mnt
```

Linux CVM Memory Issues

High Memory Utilization

Last updated : 2024-01-06 17:32:18

Error Description

The Linux CVM encounters memory issues, such as slow service response speed, CVM login failure, or Out of Memory (OOM).

Possible Reasons

These issues may be caused by high memory utilization of the instance, i.e., memory utilization generally stays above 90%.

Troubleshooting Approaches

1. Perform the [troubleshooting procedure](#) to check whether the memory utilization is too high.
2. See the [memory issue analysis](#) to find the causes of problems.

Troubleshooting Procedure

1. Follow the [directions](#) to check whether the memory utilization is too high.

If yes, proceed to the next step.

If not, see the [memory issue analysis](#) to find the causes of problems.

2. Log in to the CVM, run the `top` command, and press **M** to check whether there are processes in the “RES” and “SHR” columns using much memory.

If not, proceed to the next step.

If yes, perform the operations as instructed in [process analysis](#) according to the process type.

3. Run the following command to check the shared memory utilization.

```
cat /proc/meminfo | grep -i shmem
```

The following information will appear:

```
[root@ ~]# cat /proc/meminfo | grep -i shmem
Shmem:          556 kB
```

4. Run the following command to check the non-reclaimable slab memory utilization.

```
cat /proc/meminfo | grep -i SUnreclaim
```

The following information will appear:

```
[root@ ~]# cat /proc/meminfo | grep -i SUnreclaim
SUnreclaim:      13780 kB
```

5. Run the following command to check if there are huge pages.

```
cat /proc/meminfo | grep -iE "HugePages_Total|Hugepagesize"
```

The following information will appear:

```
[root@ ~]# cat /proc/meminfo | grep -iE "HugePages_Total|Hugepagesize"
HugePages_Total:      0
Hugepagesize:         2048 kB
```

If the `HugePages_Total` output is `0`, see the [memory issue analysis](#) to find the causes of problems.


If the `HugePages_Total` output is not `0`, there are huge pages. The huge page size equals to `HugePages_Total * Hugepagesize`. Check whether huge pages are configured by a malicious program, or if they are unnecessary, you can comment out the `vm.nr_hugepage` configuration item in the `/etc/sysctl.conf` file, and then run the `sysctl -p` command to abandon huge pages.

Directions

Viewing memory utilization

The `free` command output may vary with the Linux distributions, which is unreliable for calculating the memory utilization. Perform the following steps to view the memory utilization on the **Monitoring** page of the CVM console.

1. Log in to the [CVM console](#) and access the **Instances** page.
2. Click the **ID/Name** of the instance to enter its details page. Select the **Monitoring** tab.
3. View memory utilization in the **Memory Monitor** section.



Calculating memory utilization

The memory utilization is the ratio of memory used to total memory, excluding the buffer and system cache. The calculation formula is as follows:

$$= \frac{(\text{Total} - \text{available})}{\text{Total}} \times 100\%$$

$$= \frac{(\text{Total} - (\text{Free} + \text{Buffers} + \text{Cached} + \text{SReclaimable} - \text{Shmem}))}{\text{Total}} \times 100\%$$

$$= \frac{(\text{Total} - \text{Free} - \text{Buffers} - \text{Cached} - \text{SReclaimable} + \text{Shmem})}{\text{Total}} \times 100\%$$

The required parameters `Total`, `Free`, `Buffer`, `Cached`, `SReclaimable`, and `Shmem` can be obtained in `/proc/meminfo`. Below is an example of `/proc/meminfo`.

```
1. [root@VM_0_113_centos test]# cat /proc/meminfo
2. MemTotal: 16265592 kB
3. MemFree: 1880084 kB
4. ....
5. Buffers: 194384 kB
6. Cached: 13647556 kB
7. ....
8. Shmem: 7727752 kB
9. Slab: 328864 kB
10. SReclaimable: 306500 kB
11. SUnreclaim: 22364 kB
12. ....
13. HugePages_Total: 0
14. Hugepagesize: 2048 kB
```

The parameters are described as follows:

Parameter	Description
MemTotal	Total system memory
MemFree	Free memory
Buffers	Cached page used by block devices for reads/writes and file system metadata (such as SuperBlock)
Cached	Page cache, including POSIX/SysV shared memory and shared anonymous mmap of tmpfs

Shmem	Including shared memory, tmpfs, etc.
Slab	Memory allocated by the kernel slab memory allocator, which can be viewed using the slabtop command
SReclaimable	Reclaimable slabs
SUnreclaim	Non-reclaimable slabs
HugePages_Total	Total number of huge pages
Hugepagesize	Size of a huge page

Memory issue analysis

If the problem persists, or an error shown below appears during your use of CVM, refer to the corresponding solutions:

[Log Error “fork: Cannot allocate memory”](#)

[VNC Login Error “Cannot allocate memory”](#)

[Triggering Out of Memory When There is Available Memory](#)

Log Error “fork: Cannot allocate memory”

Last updated : 2024-01-06 17:32:18

Error Description

Log contains the error message "fork: Cannot allocate memory".

```
Jan 30 18:26:45 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:26:48 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:27:03 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:27:11 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:27:15 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:33:24 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:35:24 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:14 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:15 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:16 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:17 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:20 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:41:21 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:42:18 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
Jan 30 18:42:22 VM_130_173_centos sshd[2110]: error: fork: Cannot allocate memo:
```

Common Causes

There are too many processes. If a new process is created after the `pid_max` value is reached, the error message "fork: Cannot allocate memory" will appear.

Solution

1. Check the memory utilization as instructed in [Steps](#).
2. Check the number of processes and modify the `pid_max` configuration.

Steps

1. Check the memory utilization as instructed in [High Memory Utilization](#). If the memory utilization is normal, proceed to the next step.

2. Run the following command to obtain the value of `pid_max` .

```
sysctl -a | grep pid_max
```

Perform corresponding operations according to the returned result:

If the returned result is as shown below, where the default value of `pid_max` is 32768, go to the next step.

```
[root@VM-55-2-centos ~]# sysctl -a | grep pid_max
kernel.pid_max = 32768
```

If the error message "fork: Cannot allocate memory" is returned, run the following command to temporarily increase `pid_max` .

```
echo 42768 > /proc/sys/kernel/pid_max
```

Run the following command again to get the value of `pid_max` .

3. Run the following command to view the total number of processes.

```
pstree -p | wc -l
```

When the total number of processes has reached `pid_max` , a new process will cause the "fork: Cannot allocate memory" error.

Note:

You can use the `ps -efL` command to locate the programs for which many processes are running.

4. Change the `kernel.pid_max` value in the `/etc/sysctl.conf` configuration file to `65535` to increase the number of processes. The result should be as follows:

```
kernel.sysrq = 1
net.ipv6.conf.all.disable_ipv6=0
net.ipv6.conf.default.disable_ipv6=0
net.ipv6.conf.lo.disable_ipv6=0
kernel.numa_balancing = 0
kernel.shmmax = 68719476736
kernel.printk = 5
kernel.pid_max=65535
```

5. Run the following command for the configuration to take effect immediately.

```
sysctl -p
```

VNC Login Error “Cannot allocate memory”

Last updated : 2024-01-06 17:32:18

Error Description

I cannot log in to the CVM via VNC, and the error message “Cannot allocate memory” appears.

```
[ OK ] Started LVM2 metadata daemon.
Starting udev Coldplug all Devices...
Starting Configure read-only root support...
Starting Create Static Device Nodes in /dev...
Starting Flush Journal to Persistent Storage... onfig/network.
[ OK ] Started Apply Kernel Variables. for the current kernel
[ OK ] Started udev Coldplug all Devices.
[ OK ] Started Configure read-only root support.
[ OK ] Started Create Static Device Nodes in /dev.
Starting udev Kernel Device Manager...
Starting Load/Save Random Seed...
[ OK ] Started Load/Save Random Seed. e...
[ OK ] Started udev Kernel Device Manager.
[ 15.439583] systemd-udevd[431]: fork of child failed: Cannot allocate memory
[ 25.460271] systemd-udevd[431]: fork of child failed: Cannot allocate memory
[ 35.473367] systemd-udevd[431]: fork of child failed: Cannot allocate memory
[ 45.491094] systemd-udevd[431]: fork of child failed: Cannot allocate memory
[ 55.505765] systemd-udevd[431]: fork of child failed: Cannot allocate memory
[ OK ] Started Flush Journal to Persistent Storage.
```

Possible Reasons

This issue may be caused by too many huge pages. The default huge page size is 2048 KB. The number of huge pages is stated in `/etc/sysctl.conf`. If there are 1280 huge pages (as shown below), 2.5 GB of memory are taken. In case of low instance specification, there may not be enough memory for proper system running, and you cannot enter the system after restarting it.

[illegible]

Solutions

1. Perform the [troubleshooting procedure](#) to check whether the total number of threads exceeds the limit.
2. Modify the configurations of huge pages as needed.

Troubleshooting Procedure

1. Check whether the total number of threads exceeds the limit as instructed in [Log Error “fork: Cannot allocate memory”](#). If not, proceed to the next step.
2. Enter the single user mode and log in to the CVM. For detailed directions, see [Booting into Linux Single User Mode](#).
3. Run the following command to check the configurations of huge pages.

```
cat /etc/sysctl.conf | grep hugepages
```

If there are many huge pages, perform the following steps to modify the configurations.

4. Run the following command to open the `/etc/sysctl.conf` configuration file with VIM editor.

```
vim /etc/sysctl.conf
```

5. Press **i** to enter the edit mode and reduce the value of the `vm.nr_hugepages` as needed.
6. Press **Esc**, enter `:wq`, and press **Enter** to save the configurations and exit the VIM editor.
7. Run the following command for the configuration to take effect immediately.

```
sysctl -p
```

8. Then restart the CVM, and you can log in normally.

Triggering Out of Memory When There is Available Memory

Last updated : 2024-01-06 17:32:18

Error Description

The Linux CVM does not run out of memory and triggers OOM (Out of Memory) as shown below:

```
kernel: Out of memory: Kill process 802931 (java) score 620 or sacrifice ch
kernel: Killed process 802931 (java) total-vm:9125940kB, anon-rss:5114236kB
```

Possible Reasons

This issue may be caused by the `min_free_kbytes` configuration. It specifies the minimum idle memory of the Linux system (in kilobytes). When the system's available memory goes below the set value of `min_free_kbytes`, the system will invoke oom-killer or forcibly restart depending on the `vm.panic_on_oom` kernel parameter:

If `vm.panic_on_oom=0` is set, the system prompts OOM and invokes oom-killer to kill the process using the most memory.

If `vm.panic_on_oom =1` is set, the system will restart automatically.

Solutions


1. Perform the [troubleshooting procedure] to check the memory utilization and total number of threads.
2. Correct the `min_free_kbytes` configuration.

Troubleshooting Procedure

1. Check the memory utilization as instructed in [High Memory Utilization](#). If the memory utilization is normal, proceed to the next step.
2. Check whether the number of threads exceeds the limit as instructed in [Log Error "fork: Cannot allocate memory"](#). If the number of threads is within the limit, proceed to the next step.
3. Log in to the CVM and run the following command to view the `min_free_kbytes` configuration.

```
sysctl -a | grep min_free
```

The `min_free_kbytes` value is in kbytes. For example, the `min_free_kbytes = 1024000` shown below is 1 GB.



```
[root@ ~]# sysctl -a | grep min_free
vm.min_free_kbytes = 1024000
```

4. Run the following command to open the `/etc/sysctl.conf` configuration file with VIM editor.

```
vim /etc/sysctl.conf
```

5. Press **i** to enter the edit mode and modify the `vm.min_free_kbytes` configuration item.

Note:

We recommend changing the `vm.min_free_kbytes` value to no more than 1% of the total memory.

6. Press **Esc**, enter **:wq**, and press **Enter** to save the configurations and exit the VIM editor.

7. Run the following command for the configuration to take effect.

```
sysctl -p
```

Network Related Failures

Cross-MLC-Boarder Linkage Latency

Last updated : 2024-01-06 17:32:18

Problem Description

The user experiences high latency when logging in to CVMs located in North America from the Chinese mainland.

Problem Analysis

Due to the limited number of international egress routers within the Chinese mainland, high concurrency may cause linkage congestion and unstable access.

If you are in the Chinese mainland, and need to manage CVMs located in North America, you can purchase a CVM located in Hong Kong (China) and use it as a transfer point to log in to the CVM located in North America.

Solution

1. Purchase a Windows CVM located in Hong Kong (China) as a **jump server**.

Note:

In the “1. Select a model” of the “Custom Configuration” page, choose **Hong Kong, China**.

[Click here to purchase >>](#)

Windows operating system supports login to both Windows and Linux CVMs located in North America, which is recommended to purchase.

When purchasing the Windows CVM located in Hong Kong (China), you need to buy at least 1 Mbps bandwidth. Otherwise, you cannot log in to the jump server.

2. After the purchase is completed, log in to the Windows CVM located in Hong Kong (China) based on your needs:

[Logging in to Windows instance using the RDP file](#)

[Logging in to Windows instance via remote desktop](#)

[Logging in to Windows instance via VNC](#)

3. Log in to your CVM located in North America from the Windows CVM located in Hong Kong (China) based on your needs:

For Linux CVMs

[Logging in to Linux instance using standard login method](#)

[Logging in to Linux instance via remote login tools](#)

[Logging in to Linux instance via SSH key.](#)

For Windows CVMs

[Logging in to Windows instance using the RDP file](#)

[Logging in to Windows instance via remote desktop](#)

[Logging in to Windows instance via VNC.](#)

Website Access Failure

Last updated : 2024-01-06 17:32:18

This document describes how to locate and troubleshoot the problems that cause website access failure.

Possible Causes

Website access failure may be caused by network problems, firewall configurations or CVM overload.

Troubleshooting

Troubleshooting CVM problems

CVM shutdown, hardware failure, and high CPU/memory/bandwidth usage may all cause website access failure.

Thus, we recommend that you check CVM running status and CPU/memory/bandwidth usage.

1. Log in to the [CVM Console](#) and verify whether the running status of the CVM instance is normal on the instance management page, as shown below:

Create

Start up

Shutdown

Restart

Reset Password

More Actions

Separate keywords with ";", press Enter to separate filter tags

View instances pending reposition

ID/Name	Monitoring	Status	Availability Zone	Instance Type	Instance Configuration	Primary IPv4	Primary IPv6	Instance Billing Mode	Network billing mode	Project	Operation
<div><div></div><div></div></div>	<div></div>	<div>Running</div>	Guangzhou Zone 4	Standard S4	1-core 2GB 1Mbps System disk: Premium Cloud Storage Network: Lab1-VPC01	<div></div>	-	<div></div>	Bandwidth Package	Default Project	<div>Log In</div>

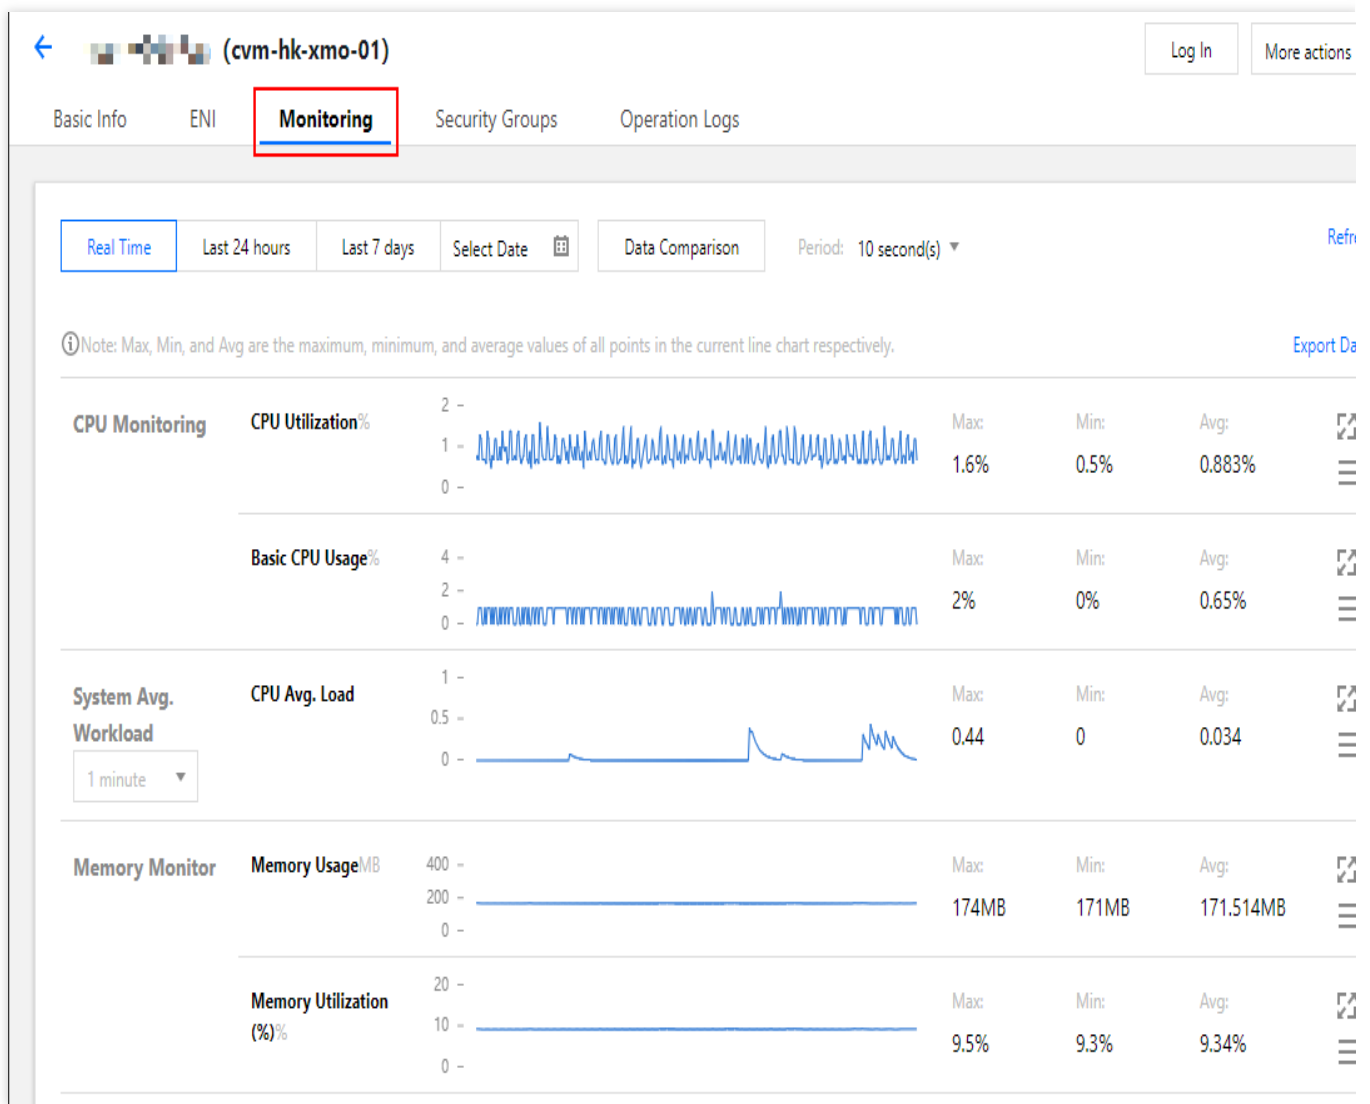
If yes, please execute [step 2](#).

If no, please restart the CVM instance.

2.

Click the ID/name of the instance to enter its details page.

3. Select the **Monitoring** tab to view the instance resource usage, as shown below:



If the CPU/memory usage is too high, please refer to [Failed to log in to a Windows CVM due to high CPU and memory usage](#) and [Failed to log in to a Linux CVM due to high CPU and memory usage](#) for troubleshooting.

If the bandwidth usage is too high, please refer to [Login Failure Due to High Bandwidth Occupation](#) for troubleshooting.

If CPU/memory/bandwidth usage is normal, please execute [step 4](#).

4.
Execute the following command
to check whether the corresponding Web service port is being monitored normally.

Note:

The following operations take port 80, which is commonly used in HTTP service, as an example.

For a Linux instance: execute the `netstat -ntulp |grep 80` command, as shown below:

```
[root@VM_2_184_centos ~]# netstat -ntulp |grep 80
tcp        0      0 0.0.0.0:80          0.0.0.0:*          LISTEN      1309/httpd
```

For a Windows instance: open the CMD command line tool to execute the `netstat -ano|findstr :80` command, as shown below:

```
C:\Users\Administrator>netstat -ano|findstr :80
TCP        0.0.0.0:80          0.0.0.0:0          LISTENING      4
TCP        10.135.182.70:53406 10.225.30.181:80    TIME_WAIT      0
TCP        10.135.182.70:53419 10.225.30.181:80    TIME_WAIT      0
TCP        10.135.182.70:53423 10.225.30.181:80    TIME_WAIT      0
TCP        [::]:80            [::]:0             LISTENING      4
```

If the port is being monitored normally, please execute [step 5](#).

If the port is not being monitored normally, please check whether the Web service process is launched or correctly configured.

5.

Check whether the corresponding We

b service port is opened in the firewall configuration.

For a Linux instance: execute the `iptables -vnL` command to check whether iptables opens port 80.

If port 80 is open, please [troubleshoot network-related problems](#).

If port 80 is not open, please execute the `iptables -I INPUT 5 -p tcp --dport 80 -j ACCEPT` command to open it.

For a Windows instance: click **Start > Control Panel > Windows Firewall** on the OS interface to check whether Windows firewall configuration is off.

- If yes, please [troubleshoot network-related problems](#).
- If no, please turn off the Windows firewall configuration.

Troubleshoot network-related problems

Network problems can also cause network access failure. You can execute the following command to check whether the network has packet loss or high latency.

```
ping the public IP of the server
```

If a result similar to the one below is returned, there is packet loss or high latency. Please use MTR for troubleshooting. For more information, please see [CVM Network Latency and Packet Loss](#).

```

MB0:~ chenhuiping$ ping 193.112.12.138
Pinging 193.112.12.138 (193.112.12.138): 56 data bytes
64 bytes from 193.112.12.138: icmp_seq=0 ttl=43 time=161.240 ms
64 bytes from 193.112.12.138: icmp_seq=1 ttl=43 time=161.996 ms
64 bytes from 193.112.12.138: icmp_seq=2 ttl=43 time=164.837 ms
64 bytes from 193.112.12.138: icmp_seq=3 ttl=43 time=215.650 ms
64 bytes from 193.112.12.138: icmp_seq=4 ttl=43 time=166.375 ms
64 bytes from 193.112.12.138: icmp_seq=5 ttl=43 time=160.576 ms
64 bytes from 193.112.12.138: icmp_seq=6 ttl=43 time=161.016 ms
64 bytes from 193.112.12.138: icmp_seq=7 ttl=43 time=164.129 ms
64 bytes from 193.112.12.138: icmp_seq=8 ttl=43 time=192.682 ms
64 bytes from 193.112.12.138: icmp_seq=9 ttl=43 time=163.376 ms
64 bytes from 193.112.12.138: icmp_seq=10 ttl=43 time=161.859 ms
^C
--- 193.112.12.138 ping statistics ---
11 packets transmitted, 11 packets received, 0.0% packet loss
round-trip min/avg/max/stddev = 160.576/170.340/215.650/16.765 ms

```

If there is no packet loss or high latency, please [troubleshoot security group problems](#).

Troubleshoot security group problems

Security group is a virtual firewall that allows you to control the inbound and outbound traffic of the associated instance. You can specify protocols, ports and policies for security group rules. If you did not open the ports related to the Web processes, website access failure may occur.

1. Log in to the [CVM Console](#) and click the ID/name of the instance to enter its details page.
2. Click the **Security Group** tab to view the bound security groups and their outbound and inbound rules. Confirm that the ports related to the Web processes are open, as shown below:

The screenshot shows the 'Security Groups' tab in the CVM Console. On the left, there is a list of security groups bound to the instance. On the right, the 'Rule preview' section is visible, showing an 'Inbound rule' table. The table has columns for Source, Port Protocol, Policy, and Notes. A single rule is listed with Source 'ALL', Port Protocol 'ALL', and Policy 'Allow'. A note at the bottom states: 'If there is no rule, all traffic is rejected'.

Source	Port Protocol	Policy	Notes
ALL	ALL	Allow	

Slow Website Access

Last updated : 2024-01-06 17:32:18

Problem Description

Website access is slow.

Problem Analysis

A complete HTTP request includes resolving the domain name, establishing the TCP connection, initiating the request, CVM receiving and processing the request, returning the result, the browser parsing the HTML code, requesting other resources, and rendering the page. These processes involve the local client, network nodes between the client and the server, and the server. A problem with any of them may cause network access latency.

Solutions

Check the local client

1. Access the [network testing website](#) to test the access speed to different domain names from the local client.
2. Based on the test result, check whether the local network has an exception.

For example, the test result is as shown below:

The following are the test results of Tencent's domain name.	
inews.qq.com	Normal network , 194 milliseconds delay
www.qq.com	Normal network , 128 milliseconds delay
3g.qq.com	Normal network , 140 milliseconds delay
mail.qq.com	Normal network , 99 milliseconds delay
user.qzone.qq.com	Normal network , 98 milliseconds delay

r.qzone.qq.com	Normal network , 203 milliseconds delay
w.qzone.qq.com	Normal network , 188 milliseconds delay
ptlogin2.qq.com	Normal network , 96 milliseconds delay
check.ptlogin2.qq.com	Normal network , 189 milliseconds delay
ui.ptlogin2.qq.com	Normal network , 91 milliseconds delay
i.mail.qq.com	Normal network , 129 milliseconds delay
v.qq.com	Normal network , 129 milliseconds delay
The following are the test results of other's domain name.	
c.3g.163.com	Normal network , 143 milliseconds delay
weibo.com	Normal network , 211 milliseconds delay
www.baidu.com	Normal network , 94 milliseconds delay
www.sina.com.cn	Normal network , 138 milliseconds delay
www.taobao.com	Normal network , 136 milliseconds delay

The test result shows the access latency for each domain name and whether the network is normal.

If the network has an exception, contact your ISP to locate and solve the problem.

If the network is normal, please [check the network linkage](#).

Check the network linkage

1. Ping the server's public IP from the local client to check if there is packet loss or high latency.

If any of the problems occurs, use MTR for troubleshooting. For more information, please see [CVM Network Latency and Packet Loss](#).

If the ping test shows no packet loss or high latency, please execute [step 2](#).

2.

Use the `dig/nslookup` command

to check whether the problem is caused by DNS resolution.

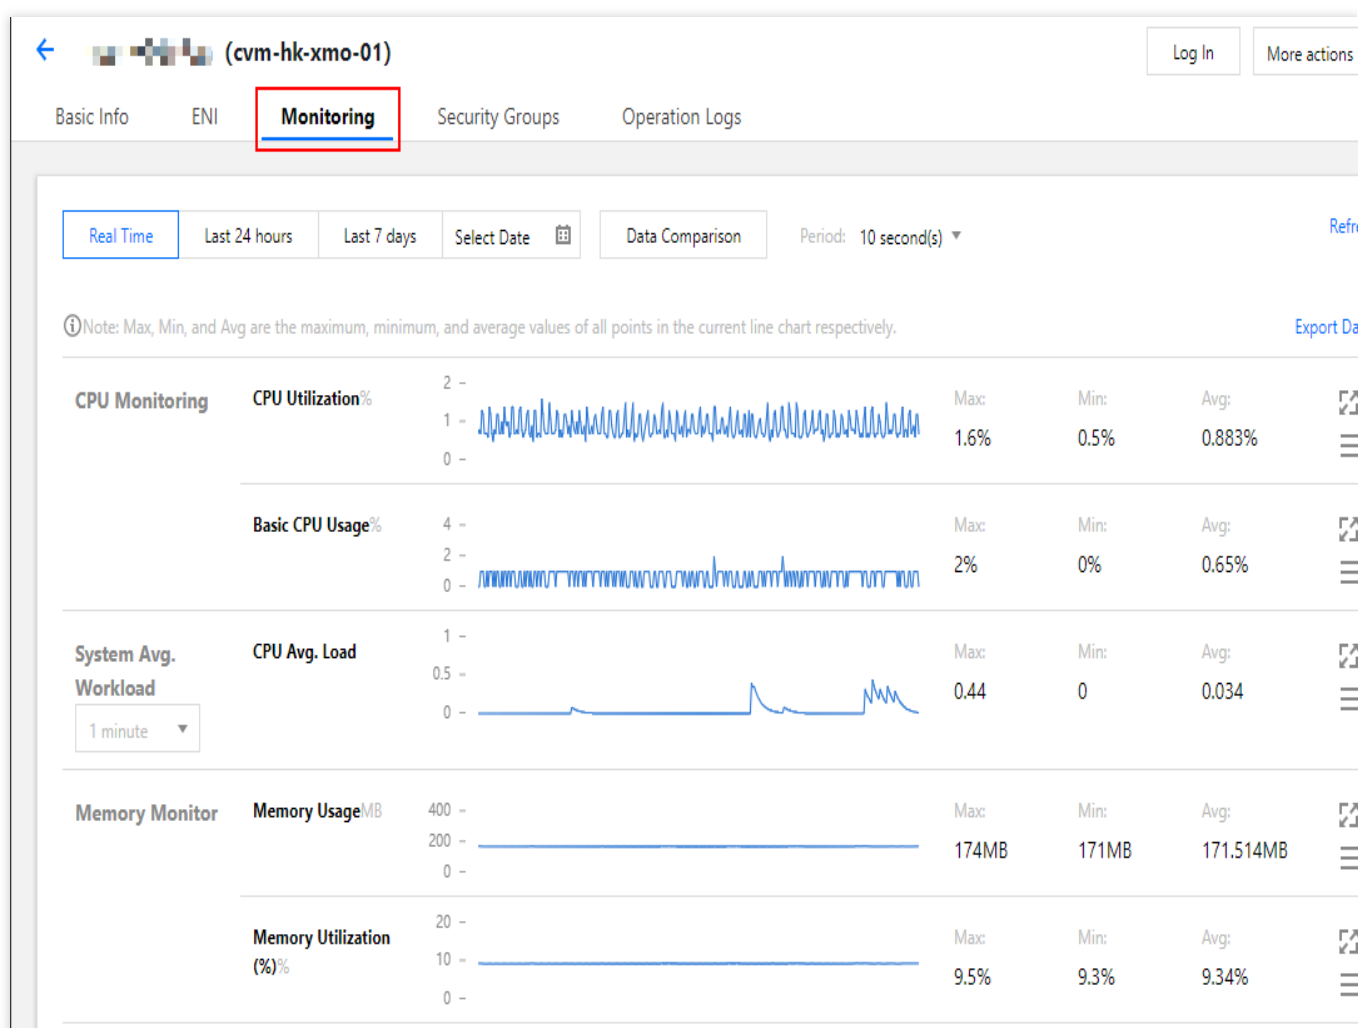
You can also access the page directly with the public network IP to check whether DNS has caused access latency.

If DNS has an exception, check the DNS resolution.

If DNS is normal, please [check the server](#).

Check the server

1. Log in to the [CVM Console](#).
2. Click the ID/name of the instance you want to check to enter its details page.
3. Select the **Monitoring** tab on the details page to view the instance resource usage, as shown below:



If the CPU/memory usage is too high, please see [Failed to log in to a Windows CVM due to high CPU and memory usage](#) and [Failed to log in to a Linux CVM due to high CPU and memory usage](#) for troubleshooting.

If the bandwidth usage is too high, please refer to [Login Failure Due to High Bandwidth Occupation](#) for troubleshooting.

If the instance resource usage is normal, please [check other problems](#).

Check other problems

Based on instance resource usage, check whether the increase in resource consumption is caused by server load.

If yes, we recommend that you optimize the business processes, [change instance configuration](#), or purchase new servers to reduce the pressure on existing servers.

If no, we recommend that you check log files to locate the problem and carry out targeted optimization.

Incorrect Multi-Queue ENI Configuration

Last updated : 2024-01-06 17:32:18

Issue Description

The multi-queue configuration of the CVM ENI is incorrect.

Common Causes

By default, the CVM is configured with multiple queues for its ENI. This method distributes ENI terminals to different CPUs and improves the network processing performance. There may be manual modifications that result in the incorrect multi-queue ENI configuration.

Solution

Correct the number of ENI queues as instructed in [Steps](#).

Steps

In the following steps, the default main ENI of the CVM is `eth0`, and the number of ENI queues is 2.

1. Run the following command to check the current number of ENI queues.

```
ethtool -l eth0
```

If the following result is returned, the currently set number of queues is less than the maximum number of ENI queues, which is unreasonable and needs to be fixed.

```
Channel parameters for eth0:
Pre-set maximums:
RX:                0
TX:                0
Other:             0
Combined:          2      ### Maximum number of ENI queues supported by the
server
Current hardware settings:
RX:                0
TX:                0
```

```
Other:          0
Combined:       1      ### Currently set number of ENI queues
```

2. Run the following command to adjust the current number of ENI queues.

```
ethtool -L eth0 combined 2
```

The number of queues in the command is set to 2, which can be adjusted as needed, up to the maximum number of ENI queues supported by the server.

3. Run the following command to check the current configuration of the number of ENI queues.

```
ethtool -l eth
```

If the maximum number of ENI queues supported by the server is equal to the currently set number of ENI queues, the configuration is successful.

CVM Network Latency and Packet Loss

Last updated : 2024-01-06 17:32:18

Problem Description

When you access the CVM from a local machine or access other network resources from the CVM, network stutters. Packet loss or high latency is found when you execute the `ping` command.

Problem Analysis

Packet loss or high latency may be caused by backbone network congestion, network node failure, high load or system configuration. You can use MTR for further diagnosis after ruling out CVM problems. MTR is a network diagnostic tool and provides reports that help you locate networking problems.

Solution

This document uses Linux and Windows CVM instances as an example to describe how to use MTR and analyze the report.

Note:

If ping is disabled on the local server or in the CVM instance, MTR will not generate any result.

Please see the MTR introduction and instructions corresponding to the host operating system.

[WinMTR Overview and Instructions \(for Windows\)](#)

[MTR Overview and Instructions \(for Linux\)](#)


WinMTR is a free network diagnostic tool for Windows integrated with Ping and traceroute features. Its graphical interface allows you to intuitively see the response time and packet loss of each node.

Installing WinMTR

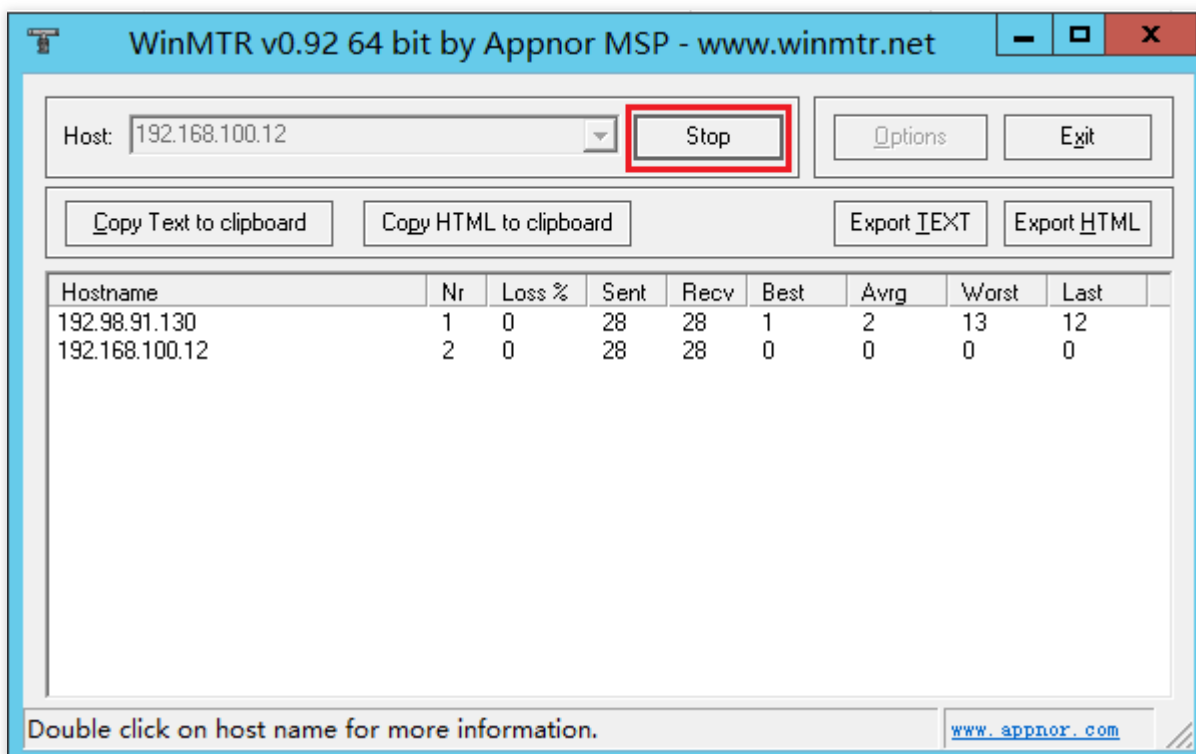
1. Log in to the Windows CVM.
2. On the operating system interface, visit the official website (or other valid channels) through the browser to download the WinMTR installer package corresponding to your operating system.
3. Unzip the WinMTR installer package.

Using WinMTR

1. Double-click WinMTR.exe to open WinMTR tool.
2. Enter the IP or domain name of the host in the **Host** field. Then click **Start** as shown below:



3. Wait for WinMTR to run for a while and click **Stop** to stop the test as shown below:



Key information of the test result is as shown below:

Hostname: IP or name of each host passed through on the path to the destination server.

Nr: Number of nodes that have been passed through.

Loss%: Packet loss of each node.

Sent: Number of data packets sent.

Recv: Number of responses received.

Best: Shortest response time.

Avrg: Average response time.

Worst: Longest response time.

Last: Last response time.

MTR is a network diagnostic tool for Linux integrated with Ping, traceroute and nslookup features. ICMP packets are used by default to test the network connection between two nodes.

Installing MTR Installation

Currently, all released versions of Linux have MTR preinstalled. If not, you can install MTR using the following command:

CentOS:

```
yum install mtr
```

Ubuntu:

```
sudo apt-get install mtr
```

MTR Parameters

-h/--help: Displays help menu.

-v/--version: Displays MTR version information.

-r/--report: Outputs the result in a report.

-p/--split: Different from **--report**, **-p/--split** lists the result of each trace separately.

-c/--report-cycles: Sets the number of data packets sent per second. Default is 10.

-s/--psize: Sets the size of each data packet.

-n/--no-dns: Disables domain name resolution for IP address.

-a/--address: Sets the IP address from which data packets are sent. It is mainly used for scenarios with a single host and multiple IP addresses.

-4: IPv4

-6: IPv6

Sample code

Take a local machine to server (IP: 119.28.98.39) as an example.

Execute the following command to output the diagnostic result of MTR in a report.

```
mtr 119.28.98.39 --report
```

Information similar to the following is returned:

```
[root@VM_103_80_centos ~]# mtr 119.28.98.39 --report
Start: Mon Feb  5 11:33:34 2019
HOST:VM_103_80_centos          Loss%   Snt    Last    Avg     Best    Wrst
StDev
1. |-- 100.119.162.130          0.0%    10     6.5    8.4     4.6    13.7
2.9
2. |-- 100.119.170.58          0.0%    10     0.8    8.4     0.6     1.1
0.0
3. |-- 10.200.135.213          0.0%    10     0.4    8.4     0.4     2.5
0.6
4. |-- 10.200.16.173           0.0%    10     1.6    8.4     1.4     1.6
0.0
5. |-- 14.18.199.58            0.0%    10     1.0    8.4     1.0     4.1
0.9
6. |-- 14.18.199.25            0.0%    10     4.1    8.4     3.3    10.2
1.9
7. |-- 113.96.7.214            0.0%    10     5.8    8.4     3.1    10.1
2.1
8. |-- 113.96.0.106            0.0%    10     3.9    8.4     3.9    11.0
2.5
9. |-- 202.97.90.206           30.0%    10     2.4    8.4     2.4     2.5
0.0
10. |-- 202.97.94.77           0.0%    10     3.5    4.6     3.5     7.0
1.2
11. |-- 202.97.51.142          0.0%    10    164.7    8.4    161.3    165.3
1.2
12. |-- 202.97.49.106          0.0%    10    162.3    8.4    161.7    167.8
2.0
13. |-- ix-xe-10-2-6-0.tcore2.LVW 10.0%    10    168.4    8.4    161.5    168.9
2.3
14. |-- 180.87.15.25            10.0%    10    348.1    8.4    347.7    350.2
0.7
15. |-- 180.87.96.21            0.0%    10    345.0    8.4    343.4    345.0
0.3
16. |-- 180.87.96.142          0.0%    10    187.4    8.4    187.3    187.6
0.0
17. |-- ???                     100.0%    10     0.0    8.4     0.0     0.0
0.0
18. |-- 100.78.119.231          0.0%    10    187.7    8.4    187.3    194.0
2.5
19. |-- 119.28.98.39           0.0%    10    186.5    8.4    186.4    186.5
0.0
```

The main output information is as follows

Host: IP address or domain name of a node.

Loss%: Packet loss.

Snt: Number of data packets sent per second.

Last: Last response time.

Avg: Average response time.

Best: Shortest response time.

Wrst: Longest response time.

StDev: Standard deviation. A higher standard deviation indicates a larger difference in the response time of data packets at this node.

Report analysis and troubleshooting

Note:

Due to network asymmetry, we recommend you collect two-way MTR data (from the local server to the destination server and from the destination server to the local server) if any network error occurs.

1. According to the report, check whether there is packet loss on the destination IP.

If there is no packet loss on the destination IP, network conditions are normal.

If there is packet loss on the destination IP, perform [Step 2](#).

2. Check the result to locate the node where the first packet loss occurred.

If packet loss occurred at the destination server, it may be caused by incorrect network configuration of the destination server. Please check its firewall configuration.

If packet loss occurred at the first three hops, it may be caused by network problems of the local machine's ISP. If the problem also happens when you access other addresses, report it to your ISP.

If packet losses frequently occur and the network is considered unstable, [submit a ticket](#) for assistance and attach the test screenshots to help the engineer locate the problem.

Network Packet Loss

Last updated : 2024-01-06 17:32:18

This document describes the common causes and troubleshooting procedures of CVM network packet loss.

Common Causes

The common causes of CVM network packet loss are as follows:

[TCP packet loss due to the limit setting](#)

[UDP packet loss due to the limit setting](#)

[Packet loss due to soft interrupt](#)

[Full UDP send buffer](#)

[Full UDP receive buffer](#)

[Full TCP accept queue](#)

[TCP request overflow](#)

[Connections exceeding the upper limit](#)

[iptables policy rules](#)

Prerequisites

To troubleshoot a problem, you need to first log in to your CVM instance. For detailed directions, see [Logging into Linux Instance](#) and [Logging into Windows Instance](#).

Troubleshooting

TCP packet loss due to the limit setting

Tencent Cloud provides various types of CVM instances, each of which has different network performance. When the maximum bandwidth or packet size of an instance is reached, packet loss may occur. The troubleshooting procedure is as follows:

1. Check the bandwidth and packet volume of the instance.

For a Linux instance, run the `sar -n DEV 2` command to check its bandwidth and packets. The `rxpck/s` and `txpck/s` metrics indicate the packets received and sent, respectively. The `rxkB/s` and `txkB/s` metrics indicate the inbound and outbound bandwidth, respectively.

2. Compare the result with the performance indicator shown in the [instance type](#) and check if the upper limit is reached.

If yes, upgrade the instance or adjust your business volume.

If no, please [submit a ticket](#) for assistance.

UDP packet loss due to the limit setting

Refer to the [troubleshooting procedure for TCP packet loss due to the limit setting](#), and check whether the upper limit is reached.

If yes, upgrade the instance or adjust your business volume.

If no, the cause may be the frequency limit on DNS requests. After the overall bandwidth or packets hit the performance bottleneck of the instance, the DNS request speed may be limited, which causes packet loss. In this case, please [submit a ticket](#) for assistance.

Packet loss due to soft interrupt

When the operating system detects that the second value of the `/proc/net/softnet_stat` statistics is increasing, a soft interrupt causes the packet loss. The troubleshooting procedure is as follows:

Check whether the RPS (Receive Packet Steering) is enabled:

If yes, increase the value of the kernel parameter `net.core.netdev_max_backlog`. For more information on how to modify a kernel parameter, see [Introduction to Linux Kernel Parameters](#).

If no, check whether the CPU high single-core soft interrupt causes the delayed data receiving and sending. In this case, you can:

Choose to enable RPS to make soft interrupt distribution more balanced.

Check whether the business program will cause uneven distribution of soft interrupts.

Full UDP send buffer

If your instance lost packets due to insufficient UDP buffer, the troubleshooting procedure is as follows:

1. Run the `ss -nump` command to check whether the UDP send buffer is full.
2. If the buffer is full, increase the values of the kernel parameters `net.core.wmem_max` and `net.core.wmem_default`, and restart the UDP program for the configuration to take effect. For more information about kernel parameters, see [Introduction to Linux Kernel Parameters](#).
3. If the packet loss problem persists, run the `ss -nump` command, and you will find that the send buffer size does not increase as expected. In this case, check whether `SO_SNDBUF` is configured through the `setsockopt` function in the code. If so, modify the code to increase the value of `SO_SNDBUF`.

Full UDP receive buffer

If your instance lost packets due to insufficient UDP buffer, the troubleshooting procedure is as follows:

1. Run the `ss -nump` command to check whether the UDP receive buffer is full.

2. If the buffer is full, increase the values of the kernel parameters `net.core.rmem_max` and `net.core.rmem_default`, and restart the UDP program for the configuration to take effect. For more information about kernel parameters, see [Introduction to Linux Kernel Parameters](#).

3. If the packet loss problem persists, run the `ss -nump` command, and you will find that the receive buffer size does not increase as expected. In this case, check whether `SO_RCVBUF` is configured through the `setsockopt` function in the code. If so, modify the code to increase the value of `SO_RCVBUF`.

Full TCP accept queue

The TCP accept queue length is the `net.core.somaxconn` value or the passed-in `backlog` value when a business process calls the listen system, whichever is smaller. If your instance lost packets due to full TCP accept queue, the troubleshooting procedure is as follows:

1. Increase the value of the kernel parameter `net.core.somaxconn`. For more information about kernel parameters, see [Introduction to Linux Kernel Parameters](#).
2. Check whether the business process passes in the `backlog` parameter, and increase its value accordingly.

TCP request overflow

If you lock the socket when TCP receives data, the data will be sent to the backlog queue. If the process fails, packet loss occurs due to the TCP request overflow. Assume the business program performs well, troubleshoot the packet loss problem at the system level.

Check whether the business program sets the buffer size through the `setsockopt` function.

If yes, modify the business program to specify a larger value or abandon the setting.

Note:

The `setsockopt` value is restricted by the kernel parameters `net.core.rmem_max` and `net.core.wmem_max`. You can also adjust the values of the two kernel parameters, and then restart the business program for the configuration to take effect.

If not, increase the respective values of the kernel parameters `net.ipv4.tcp_mem`, `net.ipv4.tcp_rmem` and `net.ipv4.tcp_wmem` to heighten the socket level.

For kernel parameter modifications, see [Introduction to Linux Kernel Parameters](#).

Connections exceeding the upper limit

Tencent Cloud provides various types of CVM instances. Each type has unique connection performance. When instance connections exceed the specified threshold, no connection is allowed, resulting in packet loss. The troubleshooting procedure is as follows:

Note:

The connection refers to the number of CVM instance sessions (including TCP, UDP, and ICMP sessions) saved on a host. If the value is greater than the network connections obtained by using the `ss` or `netstat` command on the instance, the threshold is exceeded.

Compare the network connections on your instance with the number of connections shown in the [instance type](#) and check if the upper limit is reached.

If yes, upgrade the instance or adjust your business volume.

If no, [submit a ticket](#) for assistance.

iptables policy rules

If no relevant rules are set in the iptables of the CVM, the problem may be due to the settings of the iptables policy rules that drop all packets arriving at the CVM. The troubleshooting procedure is as follows:

1. Run the following command to view the iptables policy rules.

```
iptables -L | grep policy
```

The iptables policy rule defaults to `ACCEPT`. If the INPUT chain policy is not `ACCEPT`, all packets to the CVM will be dropped. For example, if the following result is returned, all packets arriving at the CVM will be dropped.

```
Chain INPUT (policy DROP)
Chain FORWARD (policy ACCEPT)
Chain OUTPUT (policy ACCEPT)
```

2. Run the following command to adjust the value after `-P` as needed.

```
iptables -P INPUT ACCEPT
```

After adjustment, run the command in [step 1](#) again to check, and the following result should be returned:

```
Chain INPUT (policy ACCEPT)
Chain FORWARD (policy ACCEPT)
Chain OUTPUT (policy ACCEPT)
```

Ping Failures

Last updated : 2024-12-10 10:47:18

Problems

The CVM instance on the local server may be unreachable when pinged for the following reasons:

Incorrect destination server configuration

The domain name resolution fails

The link is abnormal

If the local network is normal (that is, other websites can be pinged through), troubleshoot the problem as follows:

[Check whether the instance is configured with a public IP address](#)

[Check the security group settings](#)

[Check the OS settings](#)

[Perform other operations](#)

Troubleshooting

Check whether the instance is configured with a public IP address

Note:

Only CVM instances configured with public IP addresses can communicate with other computers on the Internet. For CVM instances that are not configured with public IP addresses, the attempt to ping through the private IP address of an instance through the Internet will fail.

1. Log in to the [CVM console](#).
2. On the **Instances** page, select the ID or the name of the target instance to access its details page.

The screenshot displays the Tencent Cloud CVM console interface. On the left is a navigation sidebar with options like Instances, Dedicated Host, Placement Group, Image, Auto Scaling, Cloud Block Storage, Snapshots, SSH Key, Security Groups, and EIP. The main content area shows the details for an instance named 'ins-llf99epy (Unnamed)'. The 'Basic Info' tab is selected, showing a table of instance properties. The 'Network Information' section is expanded, and the 'Public IP' field is highlighted with a red rectangle. To the right, an 'Architecture diagram' illustrates the instance's configuration, including its connection to a security group, an ENI, and a system disk labeled 'disk-nf3zczes'.

Instance Info	
Name	Unnamed
Instance ID	[Redacted]
Instance specification	[Redacted]
Project	Default Project
Region	South China (Guangzhou)
Availability Zone	Guangzhou Zone 4
Key	None
Tag	None

Network Information	
Network	[Redacted]
Subnet	[Redacted] Default-Subnet
Public IP	[Redacted]
Private IP	[Redacted]
Act as internet gateway	No

Architecture diagram

South China (Guangzhou) / Guangzhou Zone 4 / subnet-31u47d2w

- 1 security group(s)
- 1 ENI(s)
- Tencent Linux Release 1.2 (Final) Running
- System Disk disk-nf3zczes
Premium Cloud Storage, 50 GB in total
Pay as you go Creation time: 2019-08-13 18:15:26

3. Check whether the instance is configured with a public IP address under **Network Information**.

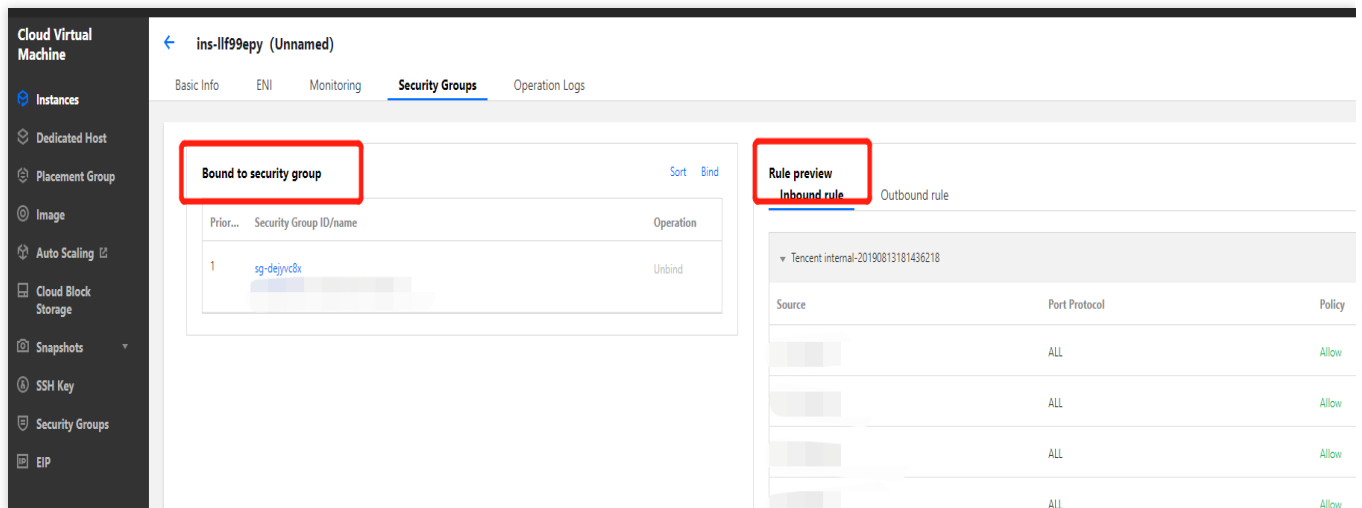
If yes, [check the security group settings](#).

If no, [bind an EIP to the instance](#).

Check the security group settings

A security group is a virtual firewall that allows you to control the inbound and outbound traffic of an associated instance. You can specify the protocol, port, and policy in a security group rule. The ICMP protocol is used in the ping test. Therefore, you need to check whether ICMP is allowed in the security group associated with the instance. To view the security group associated with the instance and its inbound and outbound rules, perform the following steps:

1. Log in to the [CVM console](#).
2. On the **Instances** page, select the ID or the name of the target instance to access the instance details page.
3. Click the **Security groups** tab to access the security group management page of the instance.



4. Check the security group associated with the instance and the detailed inbound and outbound rules to determine whether this security group allows ICMP.

If yes, [check the OS settings](#).

If no, enable the ICMP protocol policy in the security group.

Check the OS settings

Based on the operating system (OS) of the instance, select one of the following methods to check the OS settings:

For the Linux OS, [check the Linux kernel parameters and the firewall settings](#).

For the Windows OS, [check the Windows firewall settings](#). If the firewall settings are correct, try to [reset the Windows network settings](#).

Checking Linux kernel parameters and firewall configurations

Note:

In the Linux system, the kernel and the firewall settings determine whether a ping test is allowed. If the ping test is prohibited in either settings, "Request timeout" will be returned in a ping test.

Checking the kernel parameter `icmp_echo_ignore_all`

1. Log in to the instance via VNC. For details, see:

[Logging into Linux Instances via VNC](#).

[Logging into Windows Instances via VNC](#).

2. Run the following command to view the `icmp_echo_ignore_all` settings of the system:

```
cat /proc/sys/net/ipv4/icmp_echo_ignore_all
```

If 0 is returned, the OS allows all ICMP Echo requests. In this case, [check the firewall settings](#).

If 1 is returned, the OS denies all ICMP Echo requests. In this case, perform [step 3](#).

3. Run the following command to change the settings of the kernel parameter `icmp_echo_ignore_all`.

```
echo "0" >/proc/sys/net/ipv4/icmp_echo_ignore_all
```

Checking the firewall settings

Run the following command to check whether the firewall rule and the corresponding ICMP rule of the current server are disabled:

```
iptables -L
```

If the following result is returned, the corresponding ICMP rules are active.

```
Chain INPUT (policy ACCEPT)
target     prot opt source                destination           icmp echo-request
ACCEPT     icmp -- anywhere             anywhere              icmp echo-request
Chain FORWARD (policy ACCEPT)
target     prot opt source                destination
Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination           icmp echo-request
ACCEPT     icmp -- anywhere             anywhere              icmp echo-request
```

If the return result indicates that the corresponding ICMP rules are inactive, run the following commands to activate them.

```
#Chain INPUT
iptables -A INPUT -p icmp --icmp-type echo-request -j ACCEPT
#Chain OUTPUT
iptables -A OUTPUT -p icmp --icmp-type echo-reply -j ACCEPT
```

Checking the Windows firewall settings

1. Log in to the instance.
2. Open the **Control panel**, and select **Windows firewall settings**.
3. On the **Windows firewall** page, select **Advanced settings**.
4. In the pop-up "Windows firewall with advanced security" window, check whether ICMP inbound and outbound rules are disabled.

If ICMP inbound and outbound rules are disabled, please enable them.

Reset the Windows network settings

1. Check whether your VPC network supports DHCP (DHCP is supported by a VPC network created after June 2018). If it does not support DHCP, check whether the static IP in the network settings is correct.
2. If it supports DHCP, check whether the private IP of DHCP is correct. If it is incorrect, log in via VNC on the official website, and run `PowerShell` as the admin. Implement `ipconfig /release` and `ipconfig/renew` (without the need to restart the instance) to re-obtain the IP.

3. If the IP of DHCP is correct, but the network still cannot be connected, click **Start > Run**, enter `ncpa.cpl` and click **OK**. Open the local connection, try to disable and enable the ENI.

4. If the problem persists, please run the following command in CMD as the admin, and restart the instance.

```
reg delete "HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Windows NT\\CurrentVersion\\Ne
```

Other operations

If you cannot solve the problem by performing the operations above, perform the following operations:

If the domain name cannot be pinged through, check your website configurations.

If the public IP address cannot be pinged through, [submit a ticket](#) and attach relevant information about the instance and the two-way (from the local server to the CVM and from the CVM to the local server) MTR data to receive assistance.

For more information on how to use MTR, see [CVM Network Latency and Packet Loss](#).

Domain Name Resolution Failure (CentOS 6.X System)

Last updated : 2024-01-06 17:32:18

Problem Description

After a CVM with CentOS 6.x operating system is restarted or executes on the `service network restart` command, its domain names cannot be resolved. In addition, DNS information in the configuration file `/etc/resolv.conf` is found to be cleared.

Possible Reasons

In CentOS 6.x operating system, initscripts with versions earlier than 9.03.49-1 has a defect due to different grep versions.

Solution

Upgrade initscripts to the latest version and generate DNS information again.

Directions

1. Log in to the CVM.
2. Execute the following command to check the initscripts version, and verify whether a defect exists because the initscripts version is earlier than 9.03.49-1.

```
rpm -q initscripts
```

A message similar to the one below is returned:

```
initscripts-9.03.40-2.e16.centos.x86_64
```

As shown above, the initscripts version of initscripts-9.03.40-2 is earlier than the defective version of initscripts-9.03.49-1. There is a risk of DNS information being cleared.

3. Execute the following command to upgrade initscripts to the latest version and generate DNS information again.

```
yum makecache  
yum -y update initscripts  
service network restart
```

4. Execute the following command after the upgrade is completed to check the version information of initscripts, and verify whether the upgrade is successful.

```
rpm -q initscripts
```

A message similar to the one below is returned:

```
initscripts-9.03.58-1.el6.centos.2.x86_64
```

As shown above, the version displayed is different from that before the upgrade and is newer than initscripts-9.03.49-

1. This indicates that initscripts has been upgraded successfully.

Domain Name Resolution Failure (Linux System)

Last updated : 2024-12-23 16:07:43

Issue Description

The Linux instances cannot resolve domain name such as mirrors.tencentyun.com and mirror.ccs.tencentyun.com, and so on.

Possible Causes

DNS address is not configured or is configured incorrectly.

Firewall is configured with rules for port 53.

NSCD caching service for DNS is enabled.

The /lib64/libnss_dns.so.2 library file is missing, causing domain name resolution failure.

Solution

Follow the methods below to troubleshoot and locate the issue, and then choose the appropriate solution based on the actual situation:

1. Run the following command to check whether the DNS server is configured correctly.

```
cat /etc/resolv.conf # (tlinux/redhat/centos/rockylinux)
cat /run/systemd/resolve/resolv.conf # (ubuntu)
```

If the DNS server is not configured correctly, modify the DNS server address. It is recommended to refer to the document [customizing DNS configuration in Linux instances](#).

Tencent Cloud's DNS server:

```
nameserver 183.60.82.98
nameserver 183.60.83.19
```

2. Run the following command to check if the firewall has added the rules for port 53.

```
iptables -L
```

If the firewall settings include the rules for port 53, disable the firewall or delete the rules. For more information on using the firewall, see [Firewall](#).

3. Run the following command to check if the NSCD caching service for DNS is enabled.

```
systemctl status nscd
```

If the NSCD caching service is enabled, run the command `systemctl stop nscd` to disable the NSCD caching service.

4. Check if key library files such as `/lib64/libnss_dns.so.2` are missing or modified.

```
ls -l /lib64/libnss_dns.so.2
# The /lib64/libnss_dns.so.2 library file is generated by the glibc package,
and you can check if the package has been modified by verifying glibc.
rpm -V glibc
```

Check if the source file for the symbolic link exists, such as `/usr/lib64/libnss_dns-2.17.so` for CentOS 7. If it does not exist, you can download and replace it with the same version of OS downloaded from a normal server.

If `/usr/lib64/libnss_dns-2.17.so` still exists, you can fix it with a symbolic link.

```
ln -s /usr/lib64/libnss_dns-2.17.so /usr/lib64/libnss_dns.so.2
```