

Elastic Cloud Server

Troubleshooting

Issue 01
Date 2026-03-23



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2026. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 General Issues.....	1
1.1 Why Accessing a Website Outside the Chinese Mainland Is Slow on an ECS?.....	1
1.2 How Do I Troubleshoot a Ping Failure or Packet Loss Using a Link Test?.....	10
1.3 How Do I Troubleshoot Slow Connections to a Website Hosted on My ECS?.....	15
1.4 How Do I Troubleshoot an Unresponsive Website Hosted on My ECS?.....	17
1.5 Why Am I Unable to Connect to a Port on an ECS?.....	23
1.6 How Can I Resolve High Bandwidth Usage on My ECSs?.....	26
1.7 Why Is My Windows ECS Running Slowly?.....	28
1.8 Why Is My Linux ECS Running Slowly?.....	33
1.9 How Can I Handle Slow ECS Startup?.....	38
1.10 How Do I Configure Multiple IP Addresses for an ECS with Multiple NICs Attached?.....	38
2 Windows ECS Issues.....	40
2.1 How Can I Retain a Session on a Windows ECS?.....	40
2.2 How Can I Fix the Difference Between the System Time and the Local Standard Time?.....	42
2.3 How Do I Attach an Extension NIC to a Windows ECS for Accessing the Internet?.....	46
2.4 How Can I Fix Grayed Out Copy and Paste Options for a Windows ECS?.....	47
2.5 How Do I Configure File Sharing and Network Disk Mapping for a Windows ECS?.....	50
2.6 How Do I Troubleshoot an In-Service Port During Tomcat Startup?.....	60
2.7 How Do I Troubleshoot Unavailable IME Issues?.....	61
2.8 How Can I Set the Input Method for a Windows ECS?.....	68
2.9 How Do I Share Files Between Windows ECSs?.....	74
2.10 How Do I Restore Data in the Event of a Startup Failure on a Windows ECS?.....	76
2.11 How Do I View Login Logs of a Windows ECS?.....	77
2.12 What Can I Do If My Windows ECS Can Ping a Website but Cannot Access it?.....	79
2.13 Why Can't I Open the Start Menu and Search Box on a Windows ECS?.....	81
3 Linux ECS Issues.....	82
3.1 Why Is My Linux ECS Not Booting and Going Into Emergency Mode?.....	82
3.2 How Do I Fix a "Read-Only" Error When I Edit the /etc/fstab File?.....	86
3.3 How Do I Change the Time Zone on ECSs Running CentOS or EulerOS?.....	86
3.4 How Do I Fix "nf_conntrack:table full, dropping packet"?.....	87
3.5 How Do I Change the Default Boot Kernel in Ubuntu?.....	89
3.6 How Do I Configure atop and kdump on Linux ECSs?.....	90

3.7 Why Is the OS Version of My ECS Not the One in the Image I Selected During ECS Creation?.....	103
3.8 How Do I Enable My ECS to Boot From the Second Kernel If It Fails to Boot from the First Kernel?.	104
3.9 How Can I Make /etc/rc.local Run at Startup in CentOS 7?.....	105
3.10 What OSs Are Supported If I Want to Install Docker on a Linux ECS?.....	107
3.11 Why Do the Modifications to /etc/security/limits.conf Not Take Effect After the ECS Restarts?.....	107
3.12 How Do I Set vCPU Affinity for Processes Using taskset?.....	108
3.13 What Should I Do If Error "command `gcc` failed with exit status 1" Occurs During PIP-based Software Installation?.....	109
3.14 What Can I Do If Switching from a Non-root User to a Root User Timed Out?.....	110
3.15 What Can I Do If the Permissions on the Root Directory of My CentOS ECS Changed to 777?.....	111
3.16 What Should I Do If the IP Settings of My Linux ECS Are Lost?.....	113
3.17 Why Does My Linux ECS Restart Unexpectedly?.....	115
3.18 What Do I Do If Error "Cannot allocate memory" Is Displayed?.....	117
3.19 What Can I Do If the Fork Process Failed and New Threads Cannot Be Created?.....	118
3.20 What Can I Do If the ECS Startup or Remote Login Fails Due to Incorrect System Configurations?.	122
3.21 Why Do df and du Commands Show Different Disk Usage?.....	127
3.22 What Can I Do If NetworkManager Cannot Be Started? (Error Message: Failed to restart NetworkManager.service: Unit NetworkManager.service is masked).....	129
3.23 Why Is the IP Address Lost After the System Time of an ECS Is Modified?.....	129
4 Remote Login.....	131
4.1 VNC Login.....	131
4.1.1 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC?.....	131
4.1.2 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?.....	132
4.1.3 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?.....	132
4.1.4 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console?.....	132
4.2 SSH Connection Issues.....	133
4.2.1 How Do I Keep an SSH Session Alive?.....	133
4.2.2 How Can I Allow or Deny Login from Specific Users or IP Addresses to an ECS Using SSH?.....	134
4.2.3 Why Can't I Access an ECS Running CentOS 7 Using SSH After I Changed the Default SSH Port?..	135
4.2.4 How Can I Resolve ECS Login Failures Due to Corrupt /etc/passwd?.....	136
4.2.5 Why Does It Take a Long Time to Connect to an ECS Using SSH After UseDNS Is Enabled?.....	138
4.2.6 Why Does sshd Fail to Be Started on a Linux ECS?.....	139
4.2.7 How Do I Disable Login to an ECS Using SSH Password?.....	140
4.2.8 Why Are Connections to a Linux ECS Using SSH or to Applications on the ECS Interrupted Occasionally?.....	141
4.2.9 What Do I Do If I Can Log In to an Ubuntu 16.04 ECS Using SSH But the VNC Login Page Cannot Be Displayed?.....	142
4.3 Remote Login Errors on Linux.....	143
4.3.1 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?	143
4.3.2 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?.....	145

4.3.3 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?.....	147
4.3.4 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?.....	148
4.3.5 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?.....	149
4.3.6 What Do I Do If "Authentication refused: bad ownership or modes for directory /root" Is Displayed and I Can't Log In to an ECS Using SSH Key?.....	150
5 Configuring the Network.....	152
5.1 Why Does My ECS Running CentOS 7 Fail to Obtain an IP Address Using dhclient?.....	152
5.2 Why Does the NIC Names Change After I Start a Linux ECS?.....	153
5.3 Why an Entry Is Automatically Added to /etc/hosts After a Linux ECS Is Restarted?.....	154
5.4 How Do I Fix a Network Startup Failure Due to Multiple NIC Configuration Files?.....	154
5.5 Why Do I Get the Error "Name or service not known" When I Ping a Public Domain Name Configured for a Linux ECS?.....	156
5.6 Why Cannot the EIP Bound to the Extension NIC of My ECS Access the Internet?.....	158
5.7 How Do I Fix Too High Memory Usage by NetworkManager When Multiple Docker Containers Are Running?.....	159
5.8 Why Is the ECS IP Address Lost After the System Time Changes?.....	161
5.9 What Can I Do If resolv.conf Gets Reset?.....	163
5.10 What Can I Do If /etc/resolv.conf Is Restored After an ECS Running Ubuntu Is Restarted?.....	163
6 Disk Space Management Issues.....	166
6.1 Why Can't I Mount a Disk on an Old Mount Point by Modifying fstab in CentOS 7?.....	166
6.2 How Do I Create a Swap Partition or File in Linux?.....	167
6.3 Why Is the Space Not Released After I Delete a Large File on a Linux ECS?.....	168
6.4 What Should I Do If the Error Message "Read-only file system" Is Displayed When I Attempt to Delete a File on a Linux ECS?.....	170
6.5 How Do I Fix File Creation Failures Due to Inode Exhaustion?.....	172
6.6 Why Do I Get the Error "No space left on device" When I Create a File on a Linux ECS?.....	173
6.7 What Can I Do If the Buffer and Cache Occupy Too Much Memory of a Linux ECS?.....	175
6.8 What Can I Do If the Partition Capacity Fails to Be Expanded Using growpart After the EVS Disk Capacity Is Expanded?.....	176
6.9 What Can I Do If Disk Scale-Out Fails When There Is Heavy I/O Workload for SCSI Disks?.....	178
6.10 Handling Insufficient Disk Space on a Windows ECS.....	179
6.11 Handling Insufficient Disk Space on a Linux ECS.....	182
7 Self-diagnosis of Faulty GPU-accelerated ECSs.....	188
7.1 How Do I Troubleshoot a Faulty GPU-accelerated ECS?.....	188
7.2 What Types of GPU-accelerated ECS Faults Are There?.....	189
7.3 Fault Information Collection.....	191
7.3.1 How Do I Collect Fault Information?.....	191
7.3.2 How Do I Obtain the Graphics Card ID?.....	192
7.3.3 How Do I Query Graphics Card Details?.....	193
7.3.4 How Do I Query the Status of a Graphics Card?.....	194
7.3.5 How Do I Query NVIDIA Error Messages?.....	194

7.3.6 How Do I Query XID Errors?.....	195
7.3.7 How Do I Collect NVIDIA Logs?.....	195
7.3.8 How Do I Query Kernel Information?.....	196
7.3.9 How Do I Collect the Driver Installation Information?.....	196
7.4 Self-recovery from Non-hardware Faults.....	197
7.4.1 What Do I Do If the Nouveau Driver Is Not Disabled?.....	197
7.4.2 What Do I Do If There Are Retired Pages?.....	198
7.4.3 What Do I Do If the Driver Is Unavailable After the Kernel Is Upgraded?.....	200
7.4.4 Why Is the Number of Queried Graphics Cards Different from the Actual One?.....	201
7.4.5 What Do I Do If ERR! Is Displayed?.....	202
7.4.6 What Do I Do If an Error Occurs During the Installation of the NVIDIA Driver and CUDA Software?	202
7.4.7 How Do I Handle Driver Compatibility Issues?.....	203
7.4.8 How Do I Handle Recoverable Xid Errors?.....	204
7.4.9 What Do I Do If the Error Message "Windows has stopped this device because it has reported problems" Is Displayed?.....	204
7.4.10 What Do I Do If the Driver and Image I Have Selected Do Not Match My Service Scenario?.....	205
7.4.11 What Do I Do If I Have Installed the GRID Driver but Have Not Purchased or Configured the License?.....	206
7.5 Fault Diagnosis and Handling of Graphics Cards.....	206
7.5.1 How Do I Handle the infoROM Error?.....	206
7.5.2 What Do I Do If ECC Error "double bit ecc error" Occurs and There Are No Retired Pages Shown in the nvidia-smi -q Command Output?.....	207
7.5.3 What Do I Do If nvidia-smi Command Output Shows the SRAM ECC Error (V100 GPUs)?.....	208
7.5.4 What Do I Do If the GPU Is Disconnected or the Graphics Card Can't Be found, or rev ff Is Displayed After lspci grep -i nvidia Is Executed?.....	210
7.5.5 What Do I Do If the nvidia-smi Command Output Shows Overheated GPUs?.....	210
7.5.6 What Do I Do If "Unable to load the kernel module 'nvidia.ko'" Is Displayed During Driver Installation?.....	211
7.5.7 What Can I Do If an Xid Error Is Displayed in the Message Log When a GPU-accelerated ECS Is Faulty?.....	212
8 GPU Driver Issues.....	213
8.1 Why Is the GPU Driver Abnormal?.....	213
8.2 Why Is the GPU Driver Unavailable?.....	214
8.3 Why Is the GPU Display Abnormal?.....	216
8.4 Why Is the T4 GPU Display Abnormal?.....	217
8.5 How Do I Troubleshoot GPU Startup Failures Caused by NULL Pointer Dereference on NVIDIA?.....	218
9 Multi-User Login Issues.....	220
9.1 How Do I Configure Multi-User Logins for an ECS Running Windows Server 2012?.....	220
9.2 Why Does a Browser Launch Error Occur in Multi-User Login?.....	233
9.3 How Do I Apply for a License for Authenticating Multi-User Sessions and Activate an ECS?.....	234
9.4 How Do I Troubleshoot Login Screen Flickering After Configuring Multi-User Login?.....	247
10 Password and Key Pair Issues.....	249

10.1 What Should I Do If the passwd Failed to Be Executed, with Error "Authentication token manipulation error" Displayed?.....	249
10.2 How Do I Change the Key Pair for a Linux ECS?.....	250
10.3 How Do I Change the Login Mode of a Linux ECS from Key Pair to Password?.....	252
11 Firewall Configuration Issues.....	253
11.1 How Do I Enable or Disable a Windows ECS Firewall and Add Port Exceptions to the Firewall?.....	253
11.2 How Do I Enable or Disable a Linux ECS Firewall and Add a Port Exception on a Linux ECS Firewall?	258
11.3 Why Does My Linux ECS Fail to Access the Internet After Port 80 Is Allowed by the Firewall Rules?	260
12 BSOD Issues.....	263
12.1 How Do I Fix a BSOD on a Windows ECS?.....	263
12.2 How Do I Troubleshoot Blue Screen or Black Screen Errors After an ECS Is Started?.....	264
12.3 Why Does BSOD Occur When I Log In to an ECS Using Remote Desktop Connection?.....	266
13 IIS Installation Issues.....	267
13.1 How Do I Install IIS on a Windows ECS?.....	267
13.2 Why Does an Error Occur When I Attempt to Change a Domain Name on IIS Manager?.....	271
13.3 How Do I Redirect Web Pages?.....	272
13.4 What Can I Do If the Error Message "Bad Request - Invalid Hostname" Is Displayed When I Use a Windows ECS to Access an IIS Website?.....	273

1 General Issues

1.1 Why Accessing a Website Outside the Chinese Mainland Is Slow on an ECS?

Symptom

Websites outside the Chinese mainland, including those in Hong Kong (China), Macao (China), Taiwan (China), and other countries and regions, may be slow to access.

Generally, an international line is used for accessing websites outside the Chinese mainland. However, the international line may inevitably pass through network nodes distributed around the world, resulting in high latency.

Solution

- Purchase an ECS in a region (such as **CN-Hong Kong**) outside the Chinese mainland.

Considering the physical distance and network infrastructure, you can purchase an ECS in a region outside the Chinese mainland if you need to access websites outside the Chinese mainland.

For example, select the **CN-Hong Kong** region during the ECS purchase.

- Improve the access speed.

Alternatively, perform the following operations to speed up the access.

- [Modifying the DNS Configuration](#)
- [Modifying the hosts File](#)

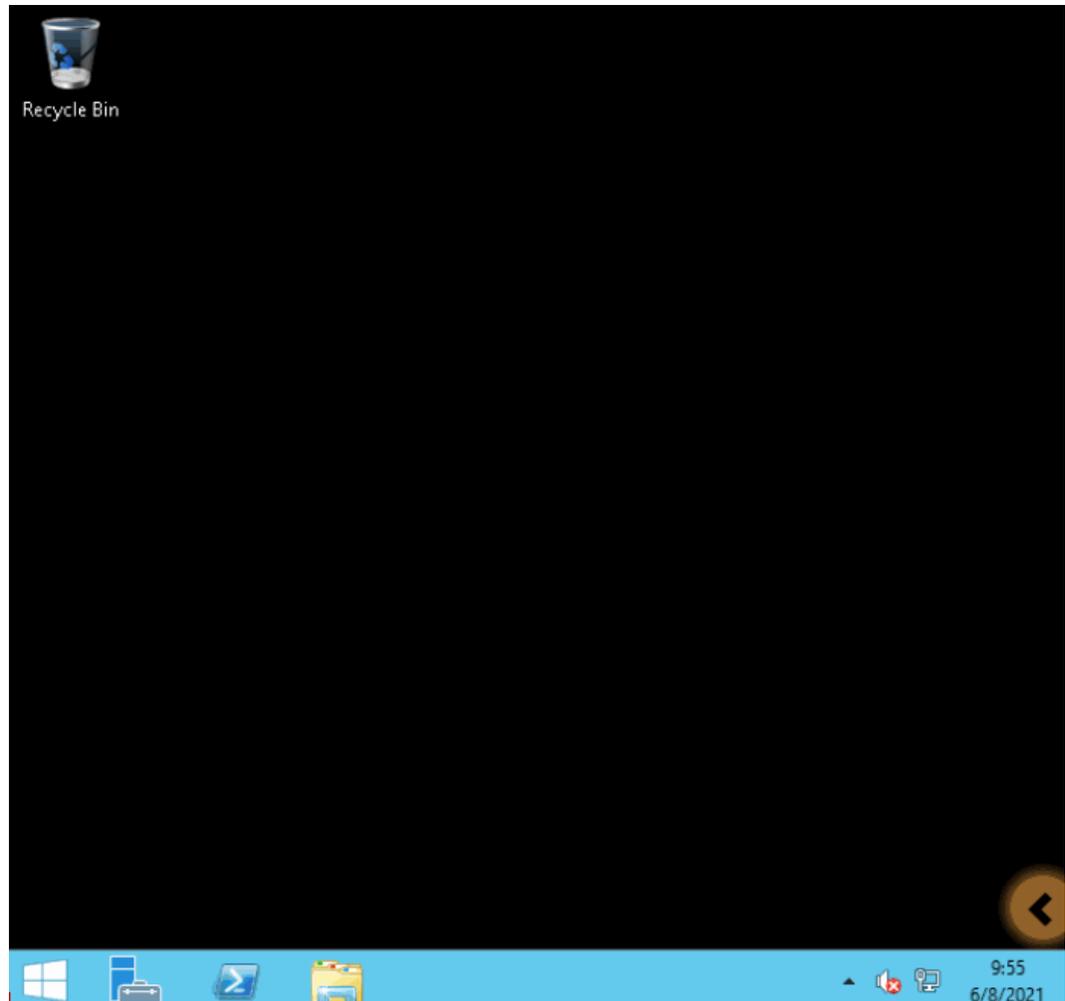
After that, run the `ping -t Website address` command to check the packet loss. For details, see [Checking Whether the Request Is Responded](#).

Modifying the DNS Configuration

Change the DNS server addresses to public DNS server addresses, for example, 101.226.4.6 and 1.1.1.1.

The following figure demonstrates how to modify the DNS configuration on an ECS running Windows Server 2012.

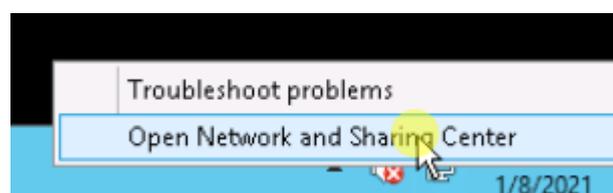
Figure 1-1 Modifying the DNS configuration



The following are detailed operations:

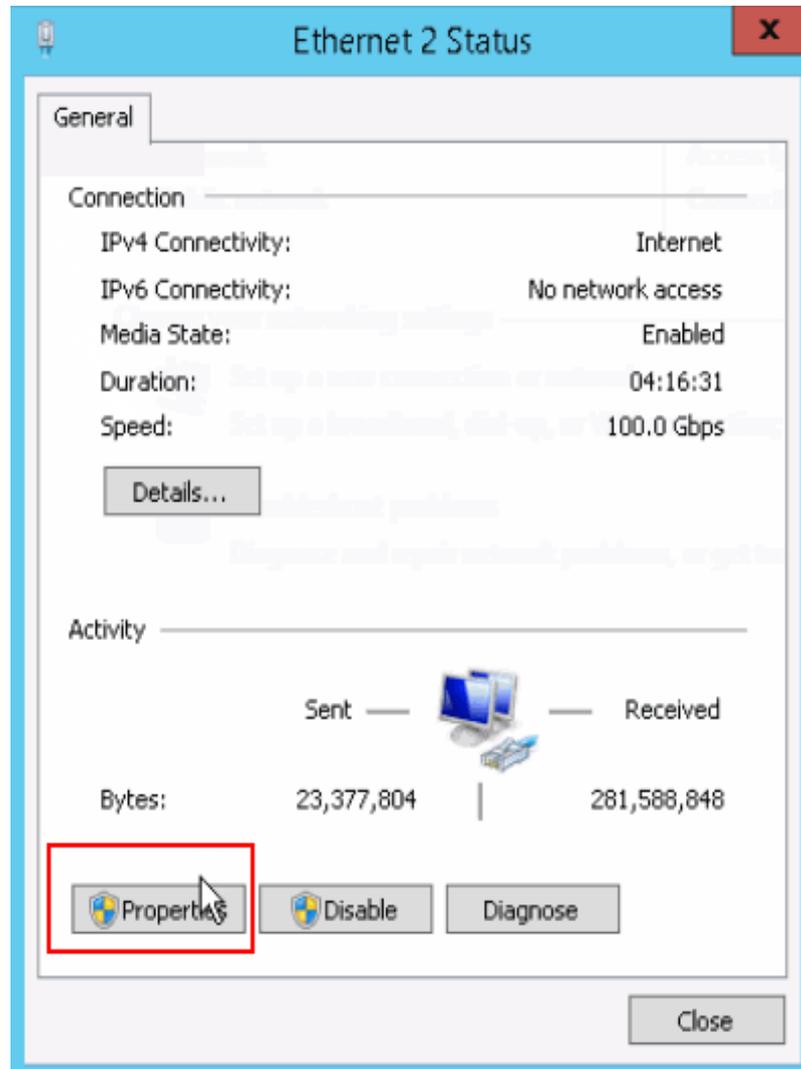
1. Log in to the Windows ECS as user **Administrator**.
2. Enable the local area connection.
 - a. In the lower right corner of the taskbar, right-click the network connection icon.
 - b. Click **Open Network and Sharing Center**.

Figure 1-2 Open Network and Sharing Center



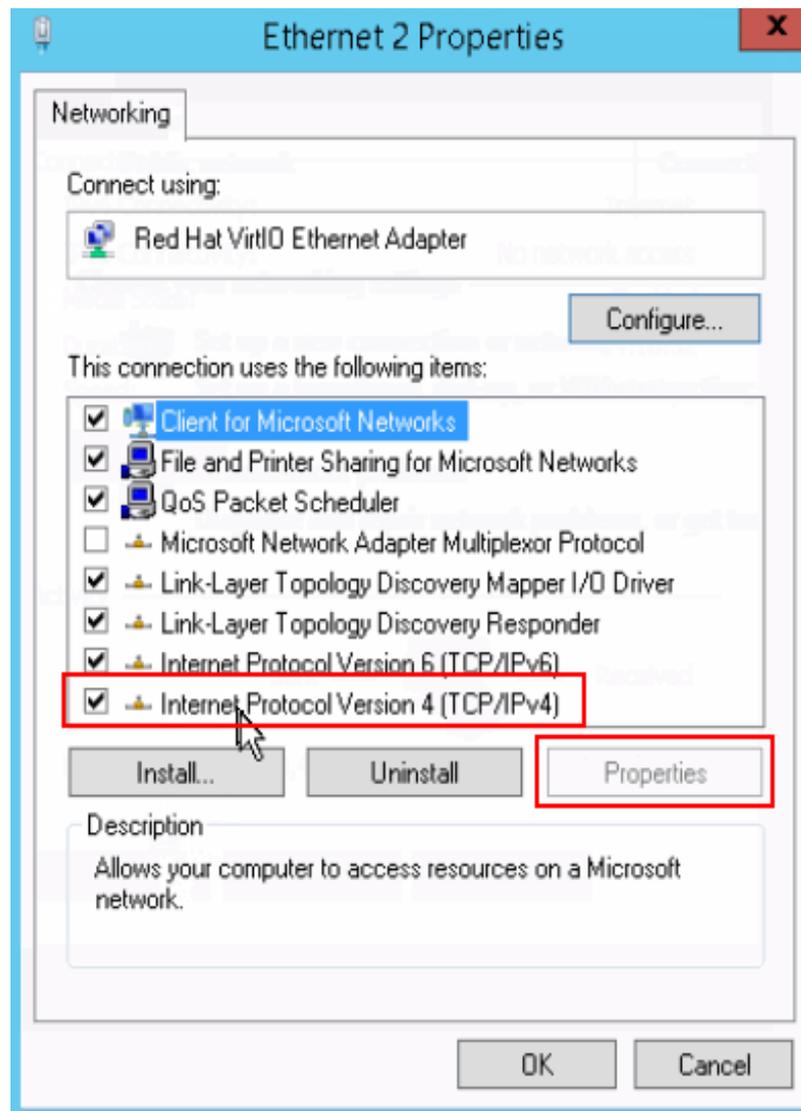
- c. In the navigation pane on the left, click **Change adapter settings**.
- 3. Configure the DNS server for the ECS.
 - a. Double-click network connections.
 - b. Click **Properties** in the lower left corner.

Figure 1-3 Local area connection



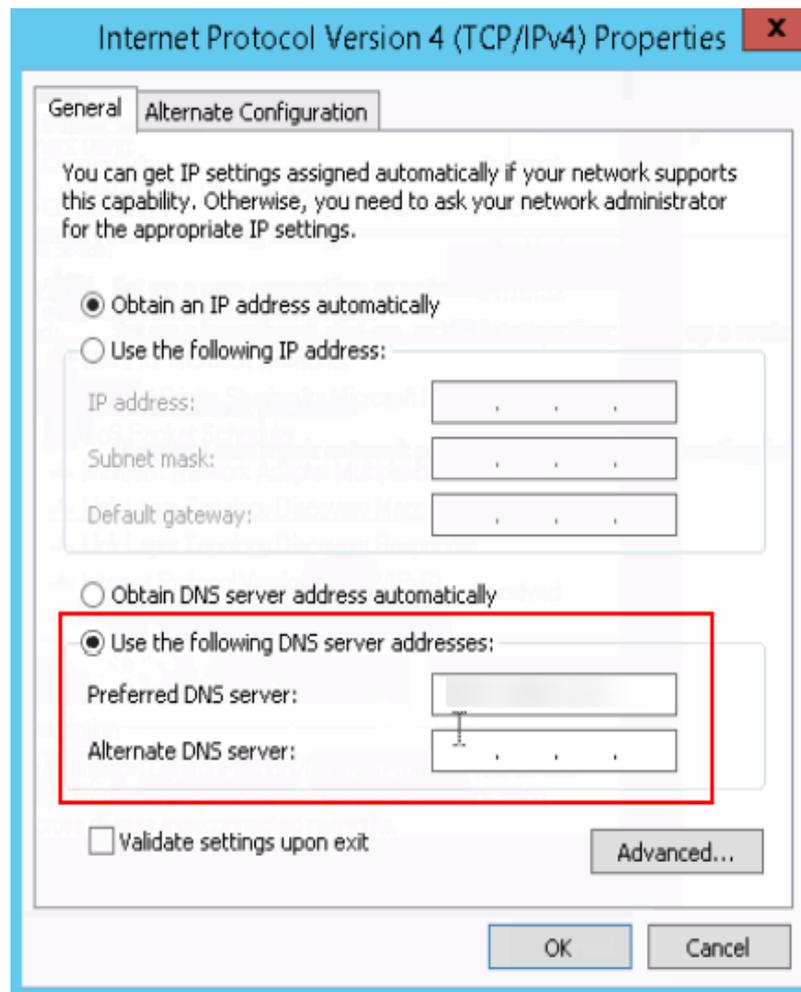
- c. Select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**.

Figure 1-4 Selecting a protocol type



- d. Select **Use the following DNS server addresses** and set the IP addresses of the DNS servers as prompted.

Figure 1-5 Setting the DNS server addresses



Modifying the hosts File

Select a server that allows you to access the website at the fastest speed and add its IP address and the domain name of the website to the **hosts** file.

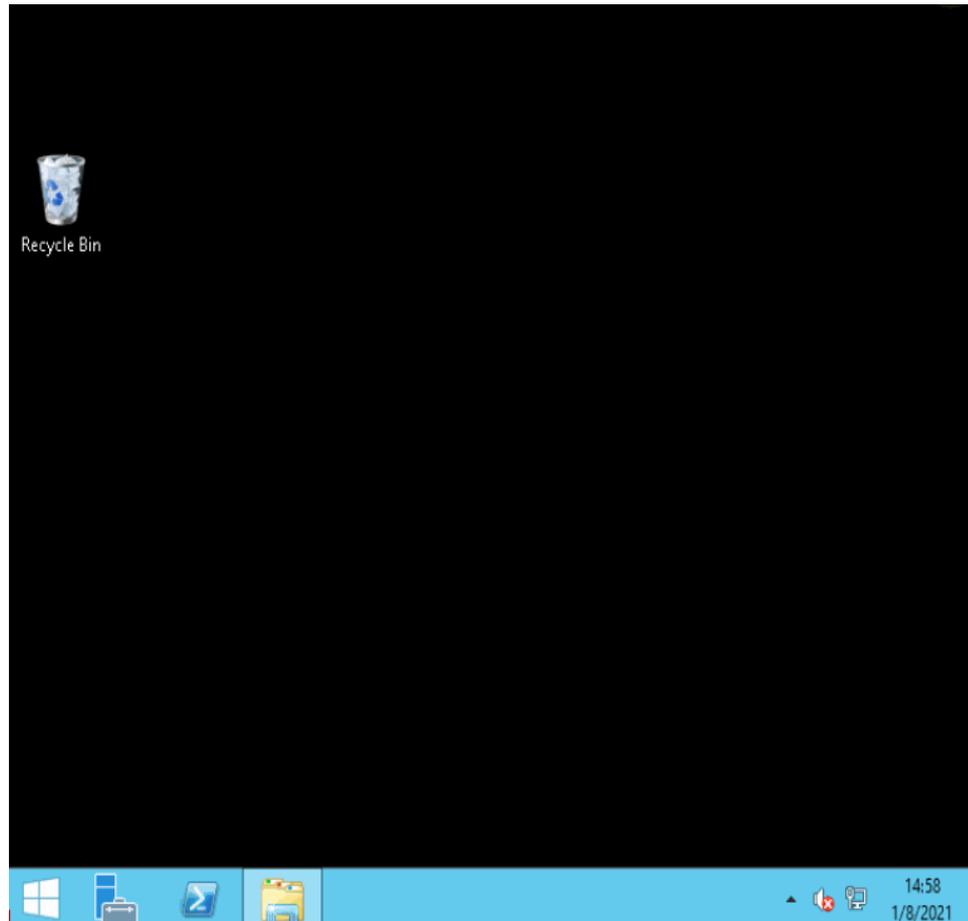
Use either of the following methods to obtain the IP address of the server that allows you to access the website at the fastest speed:

- Ping the domain name.
For details, see [Method 1: Pinging the Domain Name](#).
- Use a ping tool and PingInfoView.
For details, see [Method 2: Using a Ping Tool and PingInfoView](#).

Method 1: Pinging the Domain Name

The following figure demonstrates how you can ping the domain name on Windows Server 2012 to obtain the IP address of the server with the fastest access speed. (www.example.com is used as the example domain name.)

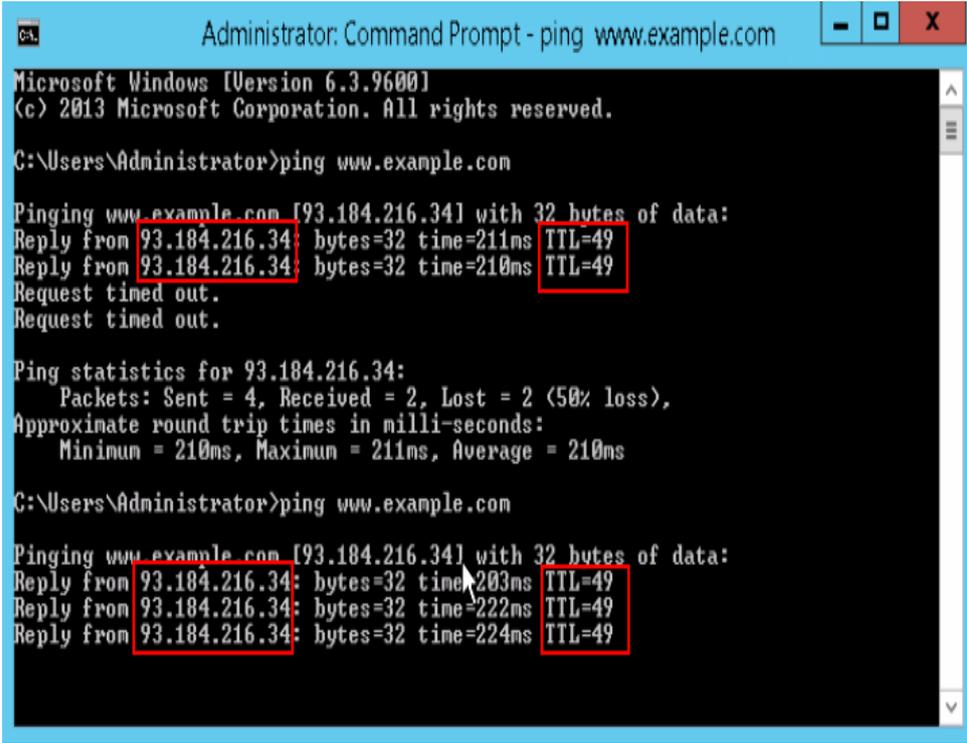
Figure 1-6 Modifying the **hosts** file



The following are detailed operations:

1. Ping `www.example.com` and wait for the result.

Figure 1-7 Command output



```
Administrator: Command Prompt - ping www.example.com
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ping www.example.com

Pinging www.example.com [93.184.216.34] with 32 bytes of data:
Reply from 93.184.216.34: bytes=32 time=211ms TTL=49
Reply from 93.184.216.34: bytes=32 time=210ms TTL=49
Request timed out.
Request timed out.

Ping statistics for 93.184.216.34:
    Packets: Sent = 4, Received = 2, Lost = 2 (50% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 210ms, Maximum = 211ms, Average = 210ms

C:\Users\Administrator>ping www.example.com

Pinging www.example.com [93.184.216.34] with 32 bytes of data:
Reply from 93.184.216.34: bytes=32 time=203ms TTL=49
Reply from 93.184.216.34: bytes=32 time=222ms TTL=49
Reply from 93.184.216.34: bytes=32 time=224ms TTL=49
```

2. Ping the domain name repeatedly and record a stable IP address with the smallest TTL value.

CAUTION

During the ping operation, run the **ipconfig /flushdns** command to refresh the DNS resolution cache. Otherwise, the same IP address will be pinged continuously.

In this example, IP address 93.184.216.34 has the smallest TTL value.

3. Modify the **hosts** file.

Open the **C:\Windows\System32\drivers\etc\hosts** file and add the mapping between the IP address and the domain name to the end of the file.

For example, if the obtained IP address is 93.184.216.34, enter **93.184.216.34 www.example.com** at the end of the **hosts** file, save and exit the file.

CAUTION

- Exercise caution when you modify the **hosts** file.
You are advised to back up the **hosts** file using either of the following methods: Copy and paste the **hosts** file, or copy and paste the content of the **hosts** file.
- Only the IP address you have configured in the **hosts** file will be returned when the domain name is used to access the website.
- If access is still slow and you want to replace the IP address, delete the existing mapping from the **hosts** file and repeat the preceding operations to obtain a new IP address.

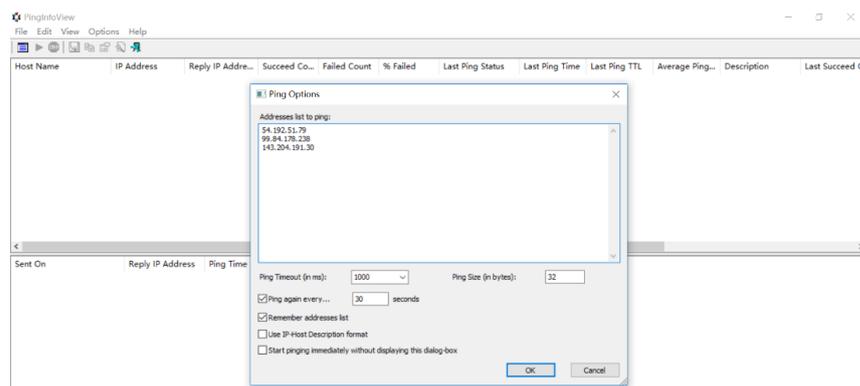
4. Access the website again.

Modifying the **hosts** file can only speed up the website access. If the problem persists, purchase an ECS in a region outside the Chinese Mainland, for example, **CN-Hong Kong**.

Method 2: Using a Ping Tool and PingInfoView

You can also try to speed up website access by modifying the **hosts** file. To do so, perform the following operations:

1. Log in to your ECS as user **Administrator**.
2. Use a browser to access the ping tool.
3. Enter the domain name of the website and record the IP addresses with the lowest response time. (**www.example.com** is used as an example.)
4. Download **PingInfoView**, decompress it, and run **PingInfoView.exe**.
5. Open **PingInfoView**, copy the IP addresses obtained in step 3 to the text box, and click **OK**.



6. Copy one IP address in the search result.

secutive...	% Failed	Last Ping Status	Last Ping Time	Last Ping TTL	Average Ping...	Descript
	0%	Succeeded	248	234	248	
	0%	Succeeded	226	238	226	
	0%	Succeeded	222	234	217	

Sent On	Reply IP Address	Ping Time	Ping TTL	Ping Status
2019/1/29 17:3...	3	216	234	Succeeded
2019/1/29 17:3...	3	217	234	Succeeded
2019/1/29 17:3...	3	217	234	Succeeded
2019/1/29 17:4...	3	216	234	Succeeded
2019/1/29 17:4...	3	216	234	Succeeded
2019/1/29 17:4...	3	222	234	Succeeded

7. Open the **hosts** file in **C:\Windows\System32\drivers\etc** and add the mapping between the IP address and the domain name to the end of the file.

CAUTION

- Exercise caution when you modify the **hosts** file.
You are advised to back up the **hosts** file using either of the following methods: Copy and paste the **hosts** file, or copy and paste the content of the **hosts** file.
- Only the IP address you have configured in the **hosts** file will be returned when the domain name is used to access the website.
- If access is still slow and you want to replace the IP address, delete the existing mapping from the **hosts** file and repeat the preceding operations to obtain a new IP address.

For example, if the selected IP address is 99.84.178.238, enter **99.84.178.238 www.example.com** at the end of the **hosts** file, save and exit the file.

8. Access the website again.
If the fault persists, use an ECS purchased in a region outside the Chinese Mainland to access the target website.

Checking Whether the Request Is Responded

Try to access the target website. If the website can be accessed but the loading is still slow, packet loss may occur. In such a case, run the **ping -t Website address** command to check the packet loss. For details, see [How Do I Troubleshoot a Ping Failure or Packet Loss Using a Link Test?](#)

For example, run **ping -t www.example.com**.

 NOTE

In Windows, you can also [download the curl client](#), decompress it, open the **bin** folder, copy the path, and configure the environment variables.

If a response status code is displayed, the request has been sent and received. Slow website access may be caused by loss of packets sent to the destination server.

Contact customer service to check for packet loss.

1.2 How Do I Troubleshoot a Ping Failure or Packet Loss Using a Link Test?

Symptom

When you accessed other resources from an ECS, network freezing occurred. The ping command output showed that packet loss occurred or the network delay was long.

This section uses Tracert and MTR as an example to describe how to troubleshoot packet loss or long delay.

Possible Cause

Packet loss or long delay may be caused by link congestion, link node faults, high server load, or incorrect system settings.

After verifying that the issue was not caused by the ECS, use Tracert or MTR for further fault locating.

MTR is used to detect network faults.

You can choose to use Tracert or MTR depending on the ECS OS:

- Windows
 - (Recommended) [Using Tracert in Windows](#)
 - [Using WinMTR in Windows](#)
- Linux
 - [Using MTR in Linux](#)

Using Tracert in Windows

Tracert shows the path through which packets reach the destination server and the time when the packets reach each node. Tracert offers similar functions as the ping command but it provides more detailed information, including the entire path the packets take, IP address of each pass-through node, and time when the packets arrive at each node.

1. Log in to the Windows ECS.
2. Open the **cmd** window and run the following command to trace the IP address:

tracert *IP address or website*

For example, **tracert www.example.com**

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

C:\Users\Administrator>tracert www.████████.com

Tracing route to www.████████.com [████████]
over a maximum of 30 hops:

  0      1 ms    1 ms    1 ms  ██████████
  1     30 ms    9 ms    7 ms  ██████████
  2      5 ms    5 ms    5 ms  ██████████

Trace complete.

C:\Users\Administrator>_
```

The command output shows that:

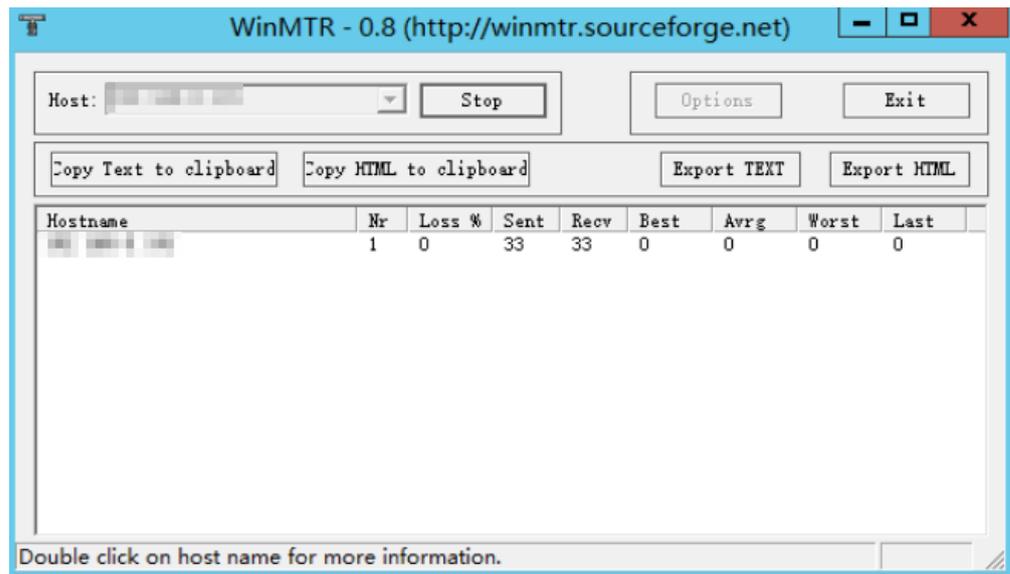
- The maximum number of hops for Tracert is 30 by default. The first column shows the sequence number of each hop.
- Tracert sends three packets each time. The second, third, and fourth columns show the time the three packets take to arrive their destination. The last column shows the IP addresses of the nodes where the packets were redirected.
- If the message ***** request timed out** is reported, troubleshoot the affected link and node.

Using WinMTR in Windows

1. Log in to the Windows ECS.
2. Download the WinMTR installation package from the official website.
3. Decompress the WinMTR installation package.
4. Double-click **WinMTR.exe** to start the tool.
5. In the WinMTR window, enter the IP address or domain name of the destination server in **Host** and click **Start**.



6. Wait for WinMTR to run for a period and click **Stop** to stop the test.



The test results are as follows:

- **Hostname:** IP address or domain name of each node that the packets pass through to the destination server
- **Nr:** number of nodes that the packets pass through
- **Loss%:** packet loss rate of a node
- **Sent:** number of sent packets
- **Recv:** number of received responses
- **Best:** shortest response time
- **Avrg:** average response time
- **Worst:** longest response time
- **Last:** last response time

Using MTR in Linux

Installing MTR

MTR has been installed on all Linux distributions. If MTR is not installed on your Linux ECS, run the following command to install it:

- CentOS
`yum install mtr`
- Ubuntu
`sudo apt-get install mtr`

MTR parameters

- **-h/--help:** help menu
- **-v/--version:** MTR version
- **-r/--report:** results of all traces
- **-p/--split:** results of each trace
- **-c/--report-cycles:** number of packets (**10** by default) sent per second
- **-s/--psize:** size of a packet

- **-n/--no-dns**: no domain name resolution performed for IP addresses
- **-a/--address**: IP address for sending packets, which is set if a single host has multiple IP addresses
- **-4**: IPv4
- **-6**: IPv6

The following uses the link between the local server and the destination server with IP address `119.xx.xx.xx` as an example.

Run the following command to obtain the MTR diagnosis results in a report:

```
mtr 119.xx.xx.xx --report
```

Information similar to the following is displayed:

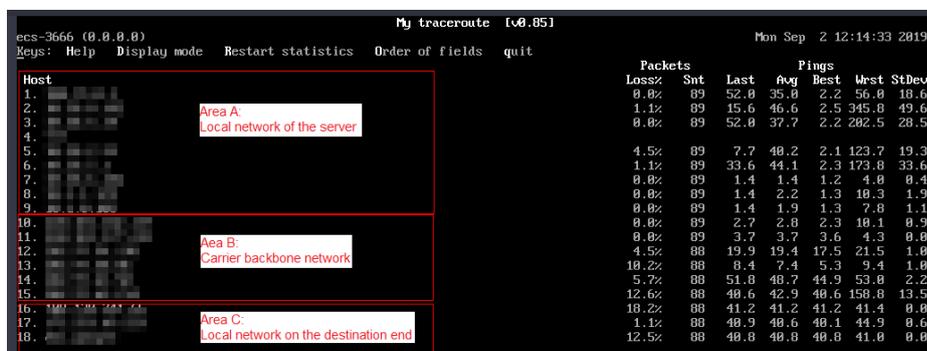
```
[root@ecs-0609 ~]# mtr 119.xx.xx.xx --report
Start: Thu Aug 22 15:41:22 2019
HOST: ecs-652          Loss% Snt  Last  Avg  Best  Wrst  StDev
 1|-- 100.xx.xx.xx      0.0%  10   3.0  3.4  2.8  7.5  1.3
 2|-- 10.xx.xx.xx      0.0%  10  52.4 51.5 34.2 58.9  6.3
 3|-- 10.xx.xx.xx      0.0%  10   3.2  5.0  2.7 20.8  5.5
 4|-- 10.xx.xx.xx      0.0%  10   1.0  1.0  1.0  1.1  0.0
 5|-- 192.xx.xx.xx     0.0%  10   3.5  4.2  2.8 11.6  2.5
 6|-- 10.xx.xx.xx      0.0%  10  35.3 34.5  6.0 56.4 22.6
 7|-- 10.xx.xx.xx      0.0%  10   3.3  4.7  3.1 14.7  3.6
 8|-- ???             100.0  10   0.0  0.0  0.0  0.0  0.0
```

The parameters in the preceding command output are described as follows:

- **HOST**: IP address or domain name of the node
- **Loss%**: packet loss rate
- **Snt**: number of packets sent per second
- **Last**: last response time
- **Avg**: average response time
- **Best**: shortest response time
- **Wrst**: longest response time
- **StDev**: standard deviation, a larger value indicates a larger difference between the response time for each data packet on the node

Handling WinMTR and MTR Reports

The following figure is an example of analyzing the reports of WinMTR and MTR.



- **Local network of the server (area A)**: the local area network and local ISP network

- If a node in the local network malfunctions, check the local network.
- If the local ISP network malfunctions, report the issue to the local carrier.
- Carrier backbone network (area B): If an error occurs in this area, identify the carrier to which the faulty node belongs based on the node IP address and report the issue to the carrier.
- Local network on the destination server (area C): the network of the provider to which the destination server belongs
 - If packet loss occurs on the destination server, the network configuration of the destination server may be incorrect. Check the firewall configuration on the destination server.
 - If packet loss occurs on certain nodes with several hops close to the destination server, the network of the provider to which the destination server belongs may be faulty.

Common Link Faults

- Incorrect destination server configurations

As shown in the following example, if the packet loss rate is 100%, the packets are not received by the destination server. The fault might be caused by incorrect network configuration on the destination server. In such a case, check the firewall configuration on the destination server.

Host	Loss%	Snt	Last	Avg	Best	Wrst	StDev
1. ???							
2. ???							
3. 1XX.X.XX	0.0%	10	521.3	90.1	2.7	521.3	211.3
4. 11X.X.XX	0.0%	10	2.9	4.7	1.6	10.6	3.9
5. 2X.X.XX	80.0%	10	3.0	3.0	3.0	3.0	0.0
6. 2X.XX.XX.XX	0.0%	10	1.7	7.2	1.6	34.9	13.6
7. 1XX.1XX.XX.X	0.0%	10	5.2	5.2	5.1	5.2	0.0
8. 2XX.XX.XX.XX	0.0%	10	5.3	5.2	5.1	5.3	0.1
9. 1XX.1XX.XX.X	100.0%	10	0.0	0.0	0.0	0.0	0.0

- ICMP rate limit

As shown in the following example, packet loss occurs on the fifth hop, but the issue does not persist on subsequent nodes. It is determined that the fault is caused by ICMP rate limit on the fifth node. This issue does not affect data transmission to the destination server, so ignore this issue.

Host	Loss%	Snt	Last	Avg	Best	Wrst	StDev
1. 1XX.XX.XX.XX	0.0%	10	0.3	0.6	0.3	1.2	0.3
2. 1XX.XX.XX.XX	0.0%	10	0.4	1.0	0.4	6.1	1.8
3. 1XX.XX.XX.XX	0.0%	10	0.8	2.7	0.8	19.0	5.7
4. 1XX.XX.XX.XX	0.0%	10	6.7	6.8	6.7	6.9	0.1
5. 1XX.XX.XX.XX	60.0%	0	27.2	25.3	23.1	26.4	2.9
6. 1XX.XX.XX.XX	0.0%	10	39.1	39.4	39.1	39.7	0.2
7. 1XX.XX.XX.XX	0.0%	10	39.6	40.4	39.4	46.9	2.3
8. 1XX.XX.XX.XX	0.0%	10	39.6	40.5	39.5	46.7	2.2

- Loop

As shown in the following example, the data packets are cyclically transferred after the fifth hop, and they cannot reach the destination server. This fault is caused by incorrect routing configuration on the nodes of the carrier. Contact the carrier to rectify the fault.

Host	Loss%	Snt	Last	Avg	Best	Wrst	StDev
1. 1XX.XX.XX.XX	0.0%	10	0.3	0.6	0.3	1.2	0.3
2. 1XX.XX.XX.XX	0.0%	10	0.4	1.0	0.4	6.1	1.8
3. 1XX.XX.XX.XX	0.0%	10	0.8	2.7	0.8	19.0	5.7
4. 1XX.XX.XX.XX	0.0%	10	6.7	6.8	6.7	6.9	0.1
5. 1XX.XX.XX.65	0.0%	10	0.0	0.0	0.0	0.0	0.0

```
6. 1XX.XX.XX.65          0.0%  10  0.0  0.0  0.0  0.0  0.0
7. 1XX.XX.XX.65          0.0%  10  0.0  0.0  0.0  0.0  0.0
8. 1XX.XX.XX.65          0.0%  10  0.0  0.0  0.0  0.0  0.0
9. ???                   0.0%  10  0.0  0.0  0.0  0.0  0.0
```

- Link interruption

As shown in the following example, no response can be received after the data packets are transferred to the fourth hop. This is generally caused by link interruption between the affected nodes. You are advised to perform a further check using a reverse link test. In such a case, contact the carrier to which the affected nodes belong.

```
Host          Loss%  Snt  Last  Avg  Best  Wrst  StDev
1. 1XX.XX.XX.XX  0.0%  10  0.3  0.6  0.3  1.2  0.3
2. 1XX.XX.XX.XX  0.0%  10  0.4  1.0  0.4  6.1  1.8
3. 1XX.XX.XX.XX  0.0%  10  0.8  2.7  0.8  19.0  5.7
4. 1XX.XX.XX.XX  0.0%  10  6.7  6.8  6.7  6.9  0.1
5. 1XX.XX.XX.XX  0.0%  10  0.0  0.0  0.0  0.0  0.0
6. 1XX.XX.XX.XX  0.0%  10  0.0  0.0  0.0  0.0  0.0
7. 1XX.XX.XX.XX  0.0%  10  0.0  0.0  0.0  0.0  0.0
8. 1XX.XX.XX.XX  0.0%  10  0.0  0.0  0.0  0.0  0.0
9 1XX.XX.XX.XX  0.0%  10  0.0  0.0  0.0  0.0  0.0
```

1.3 How Do I Troubleshoot Slow Connections to a Website Hosted on My ECS?

Symptom

A complete HTTP request includes domain name resolution, TCP connection establishment, request initiation, processing of the request and returning a processing result by the server, parsing of the HTML code and requesting other resources by the browser, and rendering and presentation of the page. The HTTP request goes through a local client of the user, network nodes between the client and the access server, and the access server. An error occurred on any of the preceding nodes will lead to network freezing on the ECS.

Checking DNS Configuration

1. Open the **cmd** window and run **ipconfig /all** to check whether a default Huawei Cloud DNS server address is used.

You are advised to use the default Huawei Cloud DNS server addresses.

NOTE

To obtain the addresses, see [What Are Huawei Cloud Private DNS Server Addresses?](#)

2. Run the following command to check whether the network between your ECS and the DNS server is connected:

```
ping IP address of the DNS server
```

3. Run the following command to check whether domain name resolution is functional:

```
nslookup Target website address
```

For example, **nslookup www.example.com**

Visit websites outside Chinese mainland, including those in Hong Kong (China), Macao (China), Taiwan (China), and other countries and regions, to check whether the access issue is resolved.

If the fault persists, perform the following operations to continue the fault locating.

Checking Network Links

1. On the local client, ping the public IP address of the server to check whether packet loss or network delay occurs.
 - If packet loss or long network delay occurs, use MTR to locate the fault. For details, see [How Do I Troubleshoot a Ping Failure or Packet Loss Using a Link Test?](#)
 - If not, go to step 2.
2. Run the **dig/nslookup** command to check whether the DNS resolution is functional. Alternatively, use the public IP address to access the target web page and check whether the slow access is caused by a DNS fault. For details, see:
 - [What Should I Do If a Record Set Does Not Take Effect?](#)
 - [How Do I Check Whether Record Sets Have Taken Effect?](#)

Checking ECS Resource Usage

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, choose **Elastic Cloud Server**.
4. In the upper right corner of the ECS list, enter an ECS name, IP address, or ID.
5. Click the name of the target ECS to go to the ECS details page.
6. Click the **Monitoring** tab to view the monitoring data.

Check whether any applications running on the ECS are using too many network or CPU resources.

- If CPU or memory is overused, see [Why Is My Windows ECS Running Slowly?](#) or [Why Is My Linux ECS Running Slowly?](#) for troubleshooting.
- If bandwidth is overused, see [How Can I Resolve High Bandwidth Usage on My ECSs?](#) for troubleshooting.
 - To upgrade ECS specifications, see [General Operations for Modifying ECS Specifications](#).
 - To increase the bandwidth, see [Modifying an EIP Bandwidth](#).

1.4 How Do I Troubleshoot an Unresponsive Website Hosted on My ECS?

Symptom

Websites running on an ECS might become unreachable for multiple reasons. Check whether the configurations of network, port, firewall, or security group of the ECS are correct.

Fault Locating

If an error is displayed when you access a website, identify possible causes based on the error message.

Identify possible causes based on error code description in [Returned Values for General Requests](#).

NOTE

If the error message cannot help you locate the fault, record the resource details and fault occurred time. Then, choose **Service Tickets** > **Create Service Ticket** in the upper right corner of the management console to submit a ticket.

You can also locate the fault based on the following possible causes which are listed in order of their probability.

If the fault persists after you have ruled out one cause, move on to the next one.

Figure 1-8 Fault locating



Table 1-1 Possible causes and solutions

Possible Causes	Solution
Port communication	Check whether the web port used by the target website is properly listened to on the ECS. For details, see Checking Port Communication .
Security group rules	Check whether the access to the Web port is allowed in the security group of the ECS. For details, see Checking Security Group Rules .
Firewall configuration	Disable the firewall and try again. For details, see Checking the Firewall Configuration .
Route configuration	Check whether the gateway configurations in the ECS route table are correct. For details, see Checking the ECS Route Configuration .
Local network	Check whether you can use another hotspot or network to access the website. For details, see Checking the Local Network .
CPU usage	Identify and optimize the processes leading to high vCPU usage. For details, see Checking the CPU Usage .

Checking Port Communication

Ensure that service processes and ports are in **LISTEN** state. [Table 1-2](#) lists the common TCP statuses.

- Linux

Run the **netstat -antpu** command to check whether the port used by the target website is in **LISTEN** status.

For example, run **netstat -antpu |grep sshd**.

Figure 1-9 Checking port listening status (Linux)

```
[root@elb-mq02 ~]# netstat -antpu | grep sshd
tcp        0      0 0.0.0.0:22          0.0.0.0:*           LISTEN    7178/sshd
```

- If the port status is **LISTEN**, go to [Checking Security Group Rules](#).
- If the port status is not **LISTEN**, check whether the web service process has been started and correctly configured.

- Windows

Perform the following operations to check port communication:

- Run **cmd.exe**.
- Run the **netstat -ano | findstr "Port number"** command to obtain the port number used by the process.
For example, run **netstat -ano | findstr "80"**.

Figure 1-10 Checking port listening status (Windows)

```
C:\Users\Administrator>netstat -ano |findstr "80"
TCP    0.0.0.0:80          0.0.0.0:0        LISTENING        4
TCP    0.0.0.0:49155    0.0.0.0:0        LISTENING        880
TCP    [::]:80         [::]:0           LISTENING        4
TCP    [::]:49155     [::]:0           LISTENING        880
UDP    0.0.0.0:123     *:*              808
UDP    [::]:123       *:*              808
```

- If the port status is **LISTEN**, go to [Checking Security Group Rules](#).
- If the port status is not **LISTEN**, check whether the web service process has been started and correctly configured.

Table 1-2 Common TCP statuses

TCP Status	Description	Application Scenario
LISTEN	Listens for network connection requests from a remote TCP port.	The TCP server is running properly.
ESTABLISHED	Indicates that a connection has been set up.	A TCP connection is properly set up.
TIME-WAIT	Waits until the remote TCP server receives the acknowledgment after sending a disconnection request.	The TCP connection is disconnected, and this state is cleared in 1 minute.
CLOSE-WAIT	Waits for a disconnection request sent by a local user.	An application program fault leads to an open socket. This state is displayed after the network is disconnected, indicating that a process is in an infinite loop or waiting for certain requirements to be met. To resolve this issue, restart the affected process.
FIN-WAIT-2	Waits for the network disconnection request from a remote TCP server.	The network has been disconnected and requires 12 minutes to automatically recover.
SYN-SENT	Waits for the matched network connection request after a network connection request is sent.	The TCP connection request failed, which is generally caused by the delayed handling of high CPU usage on the server or by a DDoS attack.

TCP Status	Description	Application Scenario
FIN-WAIT-1	Waits for the remote TCP disconnection request, or the acknowledgment for previous disconnection request.	If the network has been disconnected, this state may not automatically recover after 15 minutes. If the port has been used for a long period, restart the OS to resolve this issue.

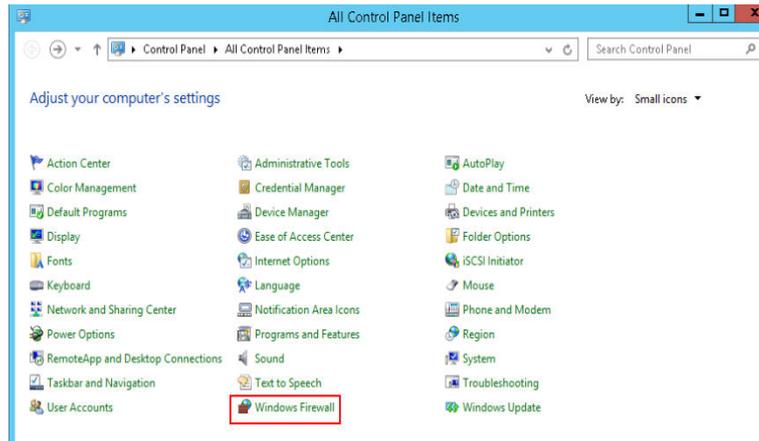
Checking Security Group Rules

If the port used by the target website is denied in the security group, add a rule to the security group to allow the access of the port.

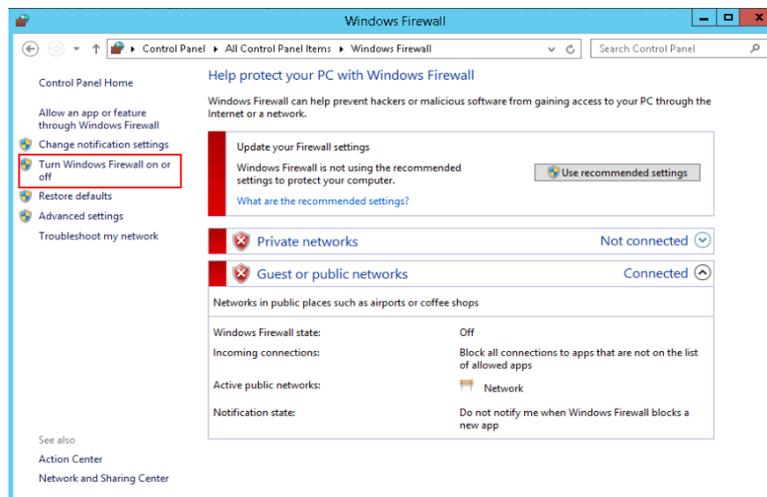
1. Log in to the management console.
2. Under **Compute**, choose **Elastic Cloud Server**.
3. In the ECS list, click the name of the target ECS.
4. Click the **Security Groups** tab.
5. In the **All** area on the left, click **Manage Rule** on the right of the target security group.
6. Configure the rule to allow the access of the port used by the website.
For details, see [Configuring Security Group Rules](#).

Checking the Firewall Configuration

- Linux ECS
The following uses port 80 and CentOS 6.8 as an example.
 - a. Run the **iptables -nvL --line-number** command to obtain firewall policies.
 - b. Run the following commands to allow access to port 80:
iptables -A INPUT -p tcp --dport 80 -j ACCEPT
iptables -A OUTPUT -p tcp --sport 80 -j ACCEPT
 - c. Run the **service iptables save** command to save the added rules.
 - d. Run the **service iptables restart** command to restart iptables.
 - e. Run the **iptables -nvL --line-number** command to check whether the added rules have taken effect.
 - f. Disable the firewall and try again.
- Windows ECS
 - a. Log in to the Windows ECS.
 - b. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.



- c. Click **Turn Windows Firewall on or off**.
View and set the firewall status.



- d. Disable the firewall and try again.

Checking the ECS Route Configuration

- Linux ECS
 - a. Run the **route** command to check the routing policy. Ensure that the default route of 0.0.0.0 is destined for the gateway and that the IP address and the gateway are in the same network segment, as shown in the first and third lines in the following figure.

```
[root@ ~]# route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
default gateway 0.0.0.0 UG 100 0 0 eth0
gateway 255.255.255.255 UGH 100 0 0 eth0
0.0.0.0 255.255.255.0 U 100 0 0 eth0
0.0.0.0 255.255.255.0 U 101 0 0 eth1
0.0.0.0 255.255.255.0 U 102 0 0 eth2
[root@ ~]#
```

- b. Run the **ifconfig** or **ip addr** command to obtain the ECS IP address.

Figure 1-11 ifconfig command output

```
[root@... ~]# ifconfig -a
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet [redacted] netmask 255.255.255.0 broadcast 1[redacted]
    inet6 fe80::f816:3eff:fe24:1e7f prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:24:1e:7f txqueuelen 1000 (Ethernet)
    RX packets 227250083 bytes 21176207838 (19.7 GiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 149514101 bytes 276209392634 (257.2 GiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1000 (Local Loopback)
    RX packets 14 bytes 1088 (1.0 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 14 bytes 1088 (1.0 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

Figure 1-12 ip addr command output

```
[root@... ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether fa:16:3e:24:1e:7f brd ff:ff:ff:ff:ff:ff
    inet [redacted]/24 brd 1[redacted] scope global noprefixroute dynamic eth0
        valid_lft 77109sec preferred_lft 77109sec
    inet6 fe80::f816:3eff:fe24:1e7f/64 scope link
        valid_lft forever preferred_lft forever
```

- c. Run the **route -n** command to obtain the gateway in the routing table.

Figure 1-13 shows an example just for reference.

Figure 1-13 route -n command output

```
[root@... ~]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
0.0.0.0 [redacted] 0.0.0.0 UG 100 0 0 eth0
1 [redacted] 255.255.255.255 UGH 100 0 0 eth0
1 [redacted] 0.0.0.0 255.255.255.0 U 100 0 0 eth0
```

- Windows ECS
 - a. Run **cmd.exe**.
 - b. Run the **ipconfig** command to obtain the ECS IP address.

Figure 1-14 ipconfig command output

```
Administrator: Command Prompt
Microsoft Windows [Version 10.0.14393]
(c) 2016 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Ethernet 4:

    Connection-specific DNS Suffix . : openstacklocal
    Link-local IPv6 Address . . . . . : [redacted]
    IPv4 Address. . . . . : [redacted]
    Subnet Mask . . . . . : [redacted]
    Default Gateway . . . . . : [redacted]
```

- c. Run the **route print** command to obtain the gateway in the routing table.

Figure 1-15 route print command output

```
Select Administrator: Command Prompt

C:\Users\Administrator>route print
=====
Interface List
10...fa 16 3e 90 4b b3 .....Red Hat VirtIO Ethernet Adapter
1.....Software Loopback Interface 1
2...00 00 00 00 00 00 e0 Microsoft ISATAP Adapter
9...00 00 00 00 00 00 e0 Microsoft Teredo Tunneling Adapter
=====

IPv4 Route Table
=====
Active Routes:
Network Destination    Netmask          Gateway          Interface Metric
=====
```

Checking the Local Network

Try another hotspot or network for access.

If the access is successful, the fault may occur in the local carrier network. In such a case, rectify the local network fault and try again.

Checking the CPU Usage

If the bandwidth or vCPU usage of an ECS is too high, website access failures may occur. If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

1. Identify the processes leading to a high bandwidth or vCPU usage.
 - Windows
Windows offers multiple tools to locate faults, including Task Manager, Performance Monitor, Resource Monitor, Process Explorer, Xperf (supported by versions later than Windows Server 2008), and full memory dump analysis.
 - Linux
Run the **top** command to check the OS running status.
2. Check whether the processes are normal and handle any issues accordingly.
 - If the processes are normal, optimize them or modify ECS configurations by referring to [Modifying Specifications](#).
 - If the processes are abnormal, stop these processes manually or use a third-party tool to stop them automatically.

1.5 Why Am I Unable to Connect to a Port on an ECS?

Scenarios

A connection to an ECS on a specific port may be prevented for multiple reasons, for example, a security group is blocking traffic to the port.

This section uses port 80 as an example to describe how to troubleshoot an unreachable ECS port.

Fault Locating

If the ECS cannot provide the HTTP service, check whether the port used by the web service (TCP port 80 by default) is working properly.

1. On the ECS management console, ensure that the port is permitted in the security group.
2. Remotely log in to the ECS and ensure that HTTP is enabled on it.
3. Ensure that the port is listened to properly. If it is not, change the listened IP address.
4. Ensure that HTTP is permitted on the ECS firewall.

Windows

The following uses an ECS running Windows Server 2012 with Internet Information Services (IIS) deployed as an example.

Step 1 Ensure that port 80 is permitted in the security group.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, choose **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, click the name of the target ECS. The page providing details about the ECS is displayed.
5. Click the **Security Groups** tab and view security group rules.
6. Make sure that the following security group rules have been added to the security group which the ECS belongs to.

Step 2 Remotely log in to the ECS and verify that IIS is enabled on it.

1. In the **Server Manager** window, choose **Tools > Internet Information Services (IIS) Manager**.
If this option is unavailable, IIS is not successfully deployed. In such a case, deploy IIS again.
2. In the **Internet Information Services (IIS) Manager** window, check the following information:
 - In the **Connections** navigation pane, right-click the ECS ID. If **Connect** is unavailable, IIS has been enabled.
 - Click **Sites**. Then, view the website status on the right side of the page. If the website is stopped, click the website and then **Start** under **Manage Server** on the right side of the page to start the website.

Step 3 Check whether the port is properly listened to on the ECS.

Open the cmd window and run the following command:

```
netstat -ano | findstr: 80
```

If information similar to the following is displayed, port 80 is being properly listened to on the entire network. If it is not, change the listened IP address.

```
C:\Users\Administrator>netstat -ano | findstr :80
TCP 0.0.0.0:80 0.0.0.0:0 LISTENING 4
TCP [::]:80 [::]:0 LISTENING 4
```

Step 4 Ensure that HTTP is permitted on the firewall.

1. Choose **Control Panel > Windows Firewall**.
2. Perform operations based on the firewall status.
 - If the firewall is disabled, no further action is required.
 - If the firewall is enabled, perform the following operations:
 - i. Click **Advanced settings**.
 - ii. In the navigation pane, click **Inbound Rules**.
 - iii. Select **World Wide Web Services (HTTP Traffic-In)**. If it is disabled, enable the rule.

----End

Linux

The following uses an ECS running CentOS 7 with Nginx deployed as an example.

Step 1 Ensure that port 80 is permitted in the security group.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, choose **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, click the name of the target ECS. The page providing details about the ECS is displayed.
5. Click the **Security Groups** tab and view security group rules.
6. Make sure that the following security group rules have been added to the security group which the ECS belongs to.

Step 2 Remotely log in to the ECS and ensure that Nginx is enabled on it.

Run the following command to check whether Nginx has been enabled:

```
systemctl status nginx
```

If the following information is displayed, Nginx has been enabled:

```
[root@i27jvkxk5y1phzZ ~]# systemctl status nginx
● nginx.service - The nginx HTTP and reverse proxy server
   Loaded: loaded (/usr/lib/systemd/system/nginx.service; disabled; vendor prese
   t: disabled)
   Active: active (running) since Tue 2017-09-12 21:14:08 CST; 44s ago
     Process: 9624 ExecStart=/usr/sbin/nginx (code=exited, status=0/SUCCESS)
     Process: 9622 ExecStartPre=/usr/sbin/nginx -t (code=exited, status=0/SUCCESS)
     Process: 9620 ExecStartPre=/usr/bin/rm -f /run/nginx.pid (code=exited, status=
   0/SUCCESS)
   Main PID: 9627 (nginx)
     CGroup: /system.slice/nginx.service
            └─9627 nginx: master process /usr/sbin/nginx
              └─9628 nginx: worker process

Sep 12 21:14:08 i27jvkxk5y1phzZ systemd[1]: Starting The nginx HTTP and reve...
Sep 12 21:14:08 i27jvkxk5y1phzZ nginx[9622]: nginx: the configuration file /...k
Sep 12 21:14:08 i27jvkxk5y1phzZ nginx[9622]: nginx: configuration file /etc/...l
Sep 12 21:14:08 i27jvkxk5y1phzZ systemd[1]: Failed to read PID from file /ru...t
Sep 12 21:14:08 i27jvkxk5y1phzZ systemd[1]: Started The nginx HTTP and rever...
Hint: Some lines were ellipsized, use -l to show in full.
```

If Nginx has not been enabled, run the following command to enable it:

```
systemctl start nginx
```

Step 3 Run the following command to check whether the port is properly listened to on the ECS:

```
netstat -an | grep 80
```

If information similar to the following is displayed, port 80 is being properly listened to on the entire network. If it is not, change the listened IP address.

```
tcp 0 0 0.0.0.0:80 0.0.0.0:* LISTEN
```

Step 4 View the iptables rules of the ECS firewall.

- Run the following command to view the firewall status:

```
systemctl status firewalld
```

or

```
firewall-cmd --state
```

- If the firewall is disabled, run the following command to enable it:

```
systemctl start firewalld
```

If "Failed to start firewalld.service: Unit is masked." is displayed after the command is executed, run the following command and then run the preceding command to enable the firewall again:

```
systemctl unmask firewalld
```

- Run the following command to view the allowed ports:

```
firewall-cmd --zone=public --list-ports
```

- Run the following command to allow TCP port 80:

```
firewall-cmd --zone=public --add-port=80/tcp --permanent
```

- Run the following command to update the firewall rules:

```
firewall-cmd --reload
```

----End

1.6 How Can I Resolve High Bandwidth Usage on My ECSs?

Scenarios

If an ECS responds slowly or even cannot be accessed, this issue may be caused by high bandwidth usage.

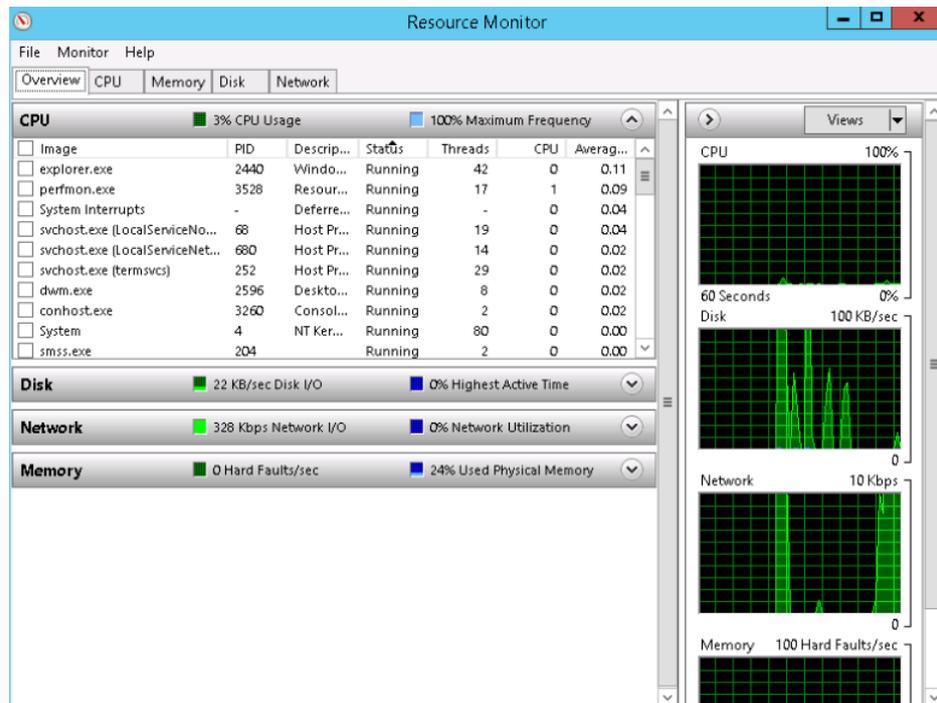
Windows ECS

1. Remotely log in to the ECS on the management console.
The Windows 2012 operating system is used as an example.
2. Start the **Run** dialog box, and then enter **perfmon -res**.

Figure 1-16 Starting the Resource Monitor

```
C:\Users\Administrator>perfmon -res  
C:\Users\Administrator>
```

3. On the **Resource Monitor** page, click the **CPU** or **Network** tab to view the CPU or bandwidth usage.



4. View the processes with high bandwidth usage.
 - If such processes are service processes, you are advised to **modify ECS specifications**.
 - If such processes are abnormal ones, they may be caused by viruses or Trojan horses. You are advised to stop the processes or use security software to scan for and stop them.

Linux ECS

1. Log in to the ECS through the management console.
The following uses an ECS running CentOS 6.8 64bit as an example.
2. Run the following command to install the Linux traffic monitoring tool iftop:
yum install iftop -y

```
[root@ecs-fccf ~]# yum install iftop -y
Loaded plugins: fastestmirror, security
Setting up Install Process
Determining fastest mirrors
epel/metalink
* base: mirrors.163.com
* epel: mirrors.njupt.edu.cn
* extras: ftp.sjtu.edu.cn
* updates: ftp.sjtu.edu.cn
base
base/primary_db 69%
```

3. Run the following command to check the ports that cause heavy traffic and the IP addresses that consume high bandwidth (taking port **eth0** as an example):

```
iftop -i eth0 -P
```

```
12.5Kb 25.0Kb 37.5Kb 50.0Kb 62.5Kb
:38364 => :http 2.23Kb 457b 114b
:38366 <=> :http 6.88Kb 1.36Kb 348b
:38366 <=> :http 2.82Kb 413b 183b
:38366 <=> :http 6.98Kb 1.38Kb 353b
TX: cum: 3.61KB peak: 4.24Kb rates: 4.25Kb 878b 217b
RX: 18.8KB 13.7Kb 13.7Kb 13.7Kb 13.7Kb
TOTAL: 14.5KB 17.9Kb 17.9Kb 17.9Kb 17.9Kb
```

4. Run the following command to check the processes related to the port (taking port 38366 as an example):

```
netstat -tunlp |grep 38366
```

- => indicates transmitted data, and <=> indicates received data.
 - **TX** indicates TX traffic, **RX** indicates RX traffic, and **TOTAL** indicates the total traffic.
 - **cum** indicates the total traffic in the first column.
 - **peak** indicates the peak traffic in the first column.
 - **rates** indicates the average traffic within 2, 10, and 40 seconds in the first column.
5. View the processes with high bandwidth usage.
 - If such processes are service processes, you are advised to **modify ECS specifications**.
 - If such processes are abnormal ones, they may be caused by viruses or Trojan horses. You are advised to stop the processes or use security software to scan for and stop them.

1.7 Why Is My Windows ECS Running Slowly?

If your ECS runs slowly or is disconnected suddenly, the possible causes are as follows:

- Your ECS is a shared ECS.

Multiple ECSs share CPU resources. When resources are insufficient, ECSs may contend for CPU resources, causing slow responses.

- The bandwidth or CPU usage of the ECS may be excessively high.
If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

If your ECS is a shared ECS, perform the following steps:

1. Fault locating: Check the instance type. For details about dedicated and shared ECSs, see [ECS Types](#).
2. Troubleshooting: If you have high requirements on service stability, you are advised to change a shared ECS to a dedicated ECS by referring to [Modifying Individual ECS Specifications](#).

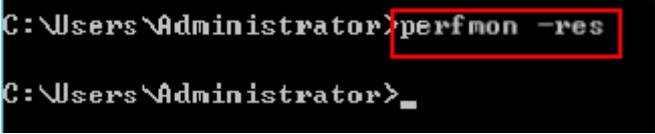
To handle this issue, perform the following operations:

1. Fault locating:
Identify the drivers from unknown sources and processes leading to high bandwidth or CPU usage.
Windows offer multiple tools to locate faults, including Task Manager, Performance Monitor, Resource Monitor, Process Explorer, Xperf (supported by versions later than Windows Server 2008), and full memory dump.
2. Check whether the processes and drivers are malicious and handle the issue accordingly.
 - If the processes are not malicious, optimize their programs or [modify ECS specifications](#).
 - If the processes are abnormal, stop these processes manually or use a third-party tool to stop them automatically.
 - If the drivers are from official sources, there is no need to deal with system built-in drivers. Determine whether to uninstall the third-party software based on your requirements.
 - If the drivers are from unknown sources, you are advised to uninstall them by using commercial antivirus software or third-party security management tools.

Fault Locating

1. Log in to the ECS using VNC available on the management console.
2. Start the **Run** dialog box, and then enter **perfmon -res**.

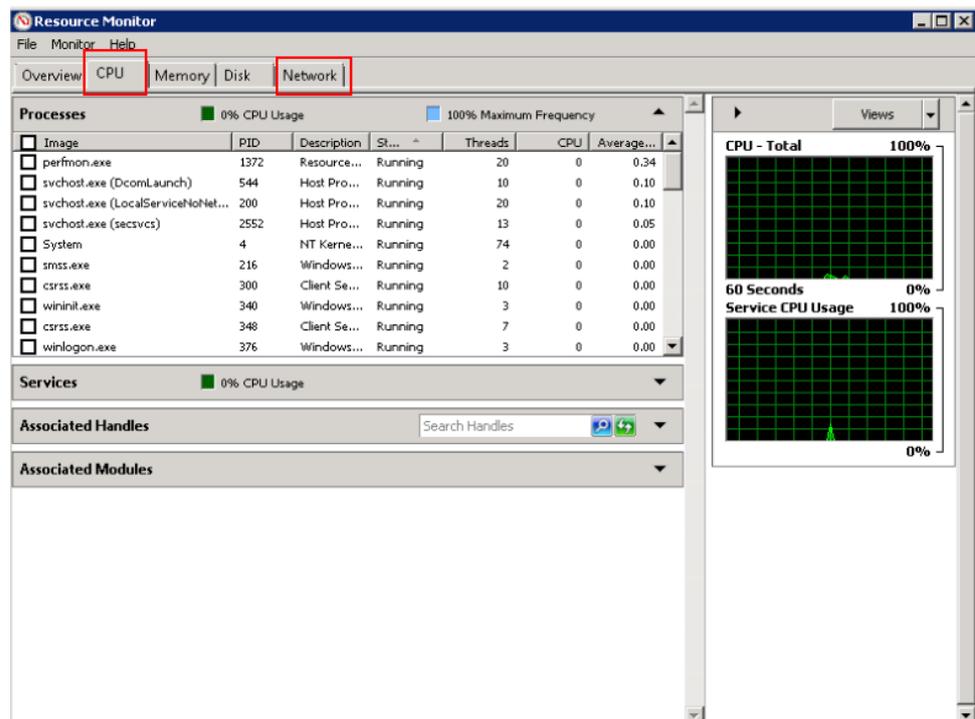
Figure 1-17 Starting the Resource Monitor



```
C:\Users\Administrator>perfmon -res
C:\Users\Administrator>
```

3. On the **Resource Monitor** page, click the **CPU** or **Network** tab to view the CPU or bandwidth usage.

Figure 1-18 Resource Monitor



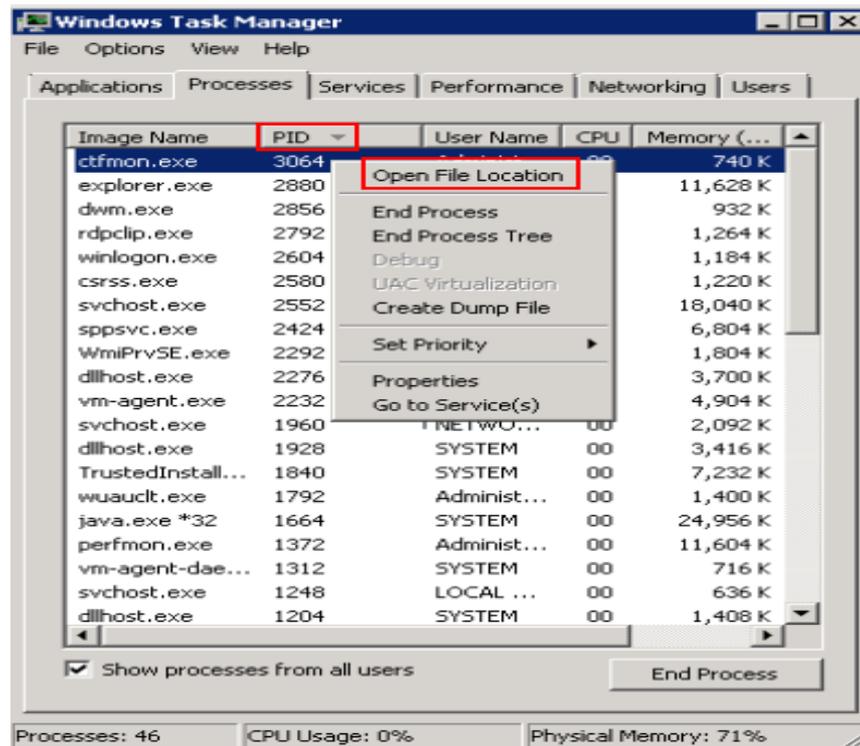
4. Obtain the IDs and names of the processes with high CPU or bandwidth usage.
5. On the VNC remote login page, click **Ctrl+Alt+Del** to start the **Windows Task Manager**.

Alternatively, start the **Run** dialog box and enter **taskmgr** to start the **Windows Task Manager**.

The following describes how to display PIDs in **Windows Task Manager**, locate a process, and check whether it is normal.

- a. Click the **Processes** tab.
- b. Click **PID** to sort the data.
- c. Right-click the process with high CPU or bandwidth usage and choose **Open File Location** from the shortcut menu.
- d. Check whether the process is malicious.

Figure 1-19 Checking the process



6. Open the **Run** dialog box and enter **fltmc** to view the filter drivers of the system.

The following figure uses Windows 10 as an example. Different OSs have different built-in drivers. For details, see their official websites. If a third-party driver is installed, it is also displayed in this figure.

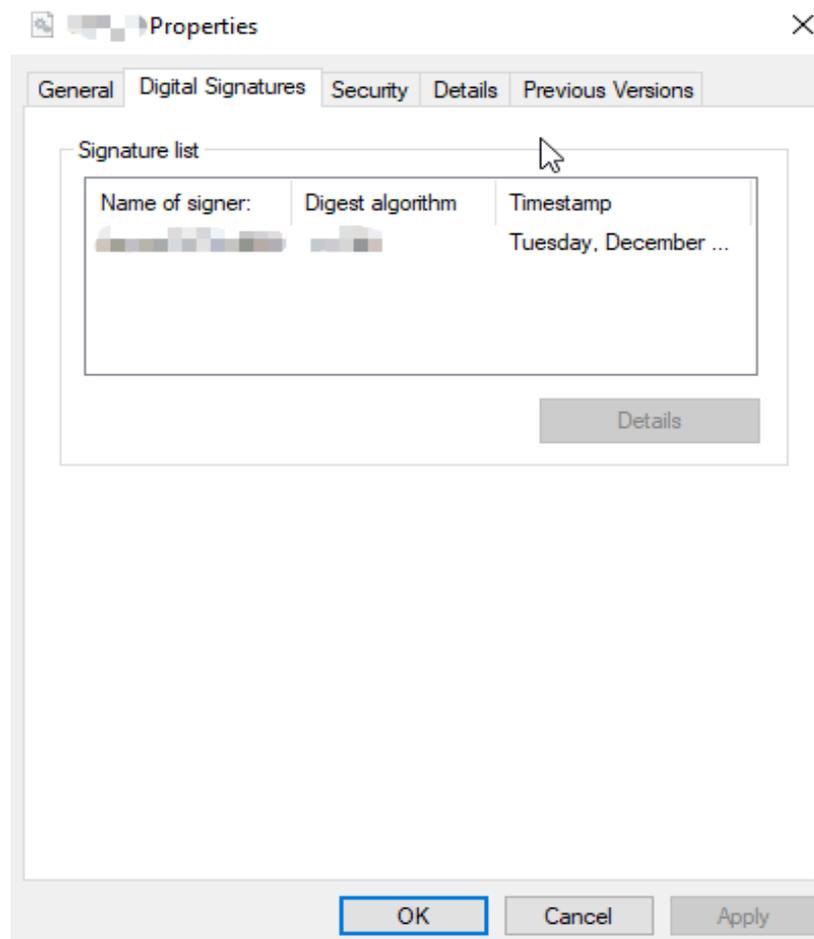
Figure 1-20 Viewing the system drivers

Filter Name	Num Instances	Altitude	Frame
WdFilter	3	328010	0
storqosflt	0	244000	0
wcifs	0	189900	0
CldFlt	0	180451	0
FileCrypt	0	141100	0
luafv	1	135000	0
npsvctrig	1	46000	0
Wof	1	40700	0

The following describes how to view a driver source and check whether the source is unknown.

- a. Go to the **C:\Windows\System32\drivers** directory on the local PC.
- b. Click the name of the unknown driver and choose **Properties** to view its details.
- c. Click the **Digital Signatures** tab to view the driver source.

Figure 1-21 Viewing the driver source



Troubleshooting

Before the troubleshooting, check whether the processes or drivers leading to the high CPU or bandwidth usage are normal, and handle the issue accordingly.

Suggestions for non-malicious processes

1. If your ECS runs Windows Server 2008 or 2012, ensure that the available memory is 2 GiB or larger.
2. Check whether Windows Update is running.
3. Check whether the antivirus software is scanning files and programs on the backend.
4. Check whether any applications requiring high CPU or bandwidth resources are running on the ECS. If yes, [modify ECS specifications](#) or [increase bandwidth](#).
5. If the ECS configuration meets the application requirements, deploy applications separately. For example, deploy the database and applications separately.

Suggestions for malicious processes

If the high CPU or bandwidth usage is caused by viruses or Trojan horses, manually stop the affected processes. You are advised to troubleshoot the issue as follows:

1. Use the commercial-edition antivirus software or install [Microsoft Safety Scanner](#) to scan for viruses in security mode.
2. Install the latest patches for Windows.
3. Run **MSconfig** to disable all drivers that are not delivered with Microsoft and check whether the fault is rectified. For details, see the official Microsoft document *How to perform a clean boot in Windows*.
4. If the ECS or site encounters a DDoS or CC attack, and a large number of access requests are generated within a short period, log in to the management console and perform the following operations:
 - Check whether Anti-DDoS has been enabled and whether the protection rules are proper. To configure a protection rule, see [Configuring an Anti-DDoS Protection Policy](#).
 - Check whether CC attack protection has been enabled and whether the protection rules are appropriate. To configure a protection rule, see [Configuring a CC Attack Protection Rule](#).

Suggestions for drivers from unknown sources

Some viruses and Trojan horses are loaded through the filter drivers of the system. If you find a driver from an unknown source, you are advised to uninstall it. You can also use commercial antivirus software or third-party security management tools to delete it.

If an unknown driver cannot be deleted, or will appear again after being deleted, it is usually a virus or Trojan horse driver. If the driver cannot be completely deleted using commercial antivirus software or third-party security management tools, you are advised to reinstall the OS and back up data before the reinstallation.

1.8 Why Is My Linux ECS Running Slowly?

If your ECS runs slowly or is disconnected suddenly, the possible causes are as follows:

- Your ECS is a shared ECS.
Multiple ECSs share CPU resources. When resources are insufficient, ECSs may contend for CPU resources, causing slow responses.
- The bandwidth or CPU usage of the ECS may be excessively high.
If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

If your ECS is a shared ECS, perform the following steps:

1. Fault locating: Check the instance type. For details about dedicated and shared ECSs, see [ECS Types](#).
2. Troubleshooting: If you have high requirements on service stability, you are advised to change a shared ECS to a dedicated ECS by referring to [Modifying Individual ECS Specifications](#).

To handle this issue, perform the following operations:

1. Fault locating:
Identify the processes leading to high bandwidth or CPU usage.
2. Troubleshooting: Check whether the processes are malicious and handle the issue accordingly.
 - If the processes are normal, optimize them or [modify ECS specifications](#).
 - If the processes are malicious, stop these processes manually or use a third-party tool to stop them automatically.

Common Commands

The following uses the CentOS 7.2 64bit OS as an example to describe common commands. The commands may vary depending on Linux OS editions. For details, see the official documentation for the specific OS edition.

The common commands for checking Linux ECS performance metrics, such as the CPU usage, are as follows:

- `ps -aux`
- `ps -ef`
- `top`

Locating High CPU Usage

1. Log in to the ECS using VNC.
2. Run the following command to check the OS running status:

`top`

Information similar to the following is displayed.

```
top - 20:56:02 up 37 days, 9:09, 1 user, load average: 0.00, 0.01, 0.05
Tasks: 80 total, 1 running, 79 sleeping, 0 stopped, 0 zombie
%Cpu(s): 0.2 us, 0.3 sy, 0.0 ni, 99.5 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 2963304 free, 178384 used, 738336 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3434808 avail Mem

  PID USER      PR  NI   VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 8115 root        20   0 161896  2216 1564 R   0.3  0.1   0:00.01 top
    1 root        20   0 125480  3884  2604 S   0.0  0.1   0:11.32 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   0:00.04 ksoftirqd/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   0:00.18 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   7:32.18 rcu_sched
   10 root        0 -20     0     0     0 S   0.0  0.0   0:00.00 lru-add-drain
```

3. View the command output.
 - The first line in the command output is "20:56:02 up 37 days, 1 user, load average: 0.00, 0.01, 0.05", indicating that:
The current system time is 20:56:02; the ECS has been running for 37 days; there is one login user; the last three values indicate the average CPU load in the last 1 minute, 5 minutes, and 15 minutes, respectively.
 - The third line in the command output shows the overall CPU usage.
 - The fourth line in the command output shows the overall memory usage.
 - The lower part of the command output shows the resource usage of each process.

 NOTE

1. On the **top** page, enter **q** or press **Ctrl+C** to exit.
2. Alternatively, click **Input Command** in the upper right corner of the VNC login page, paste or enter commands in the displayed dialog box, and click **Send**.
3. Common parameters in **top** commands are as follows:
 - s**: Change the image update frequency.
 - l**: Show or hide the first line for the top information.
 - t**: Show or hide the second line for tasks and the third line for CPUs.
 - m**: Show or hide the fourth line for Mem and the fifth line for Swap.
 - N**: Sort processes by PID in ascending or descending order.
 - P**: Sort processes by CPU usage in ascending or descending order.
 - M**: Sort processes by memory usage in ascending or descending order.
 - h**: Show help for commands.
 - n**: Set the number of processes displayed in the process list.
4. Run the **ll /proc/PID/exe** command to obtain the program file specified by a PID.

```
lroot@elb-mg01 sysconfig]# ll /proc/4243/exe  
lrwxrwxrwx 1 root root 0 Mar 18 11:46 /proc/4243/exe -> /CloudResetPwdUpdateAgent/depend/jre1.8.0_131/bin/java
```

Troubleshooting High CPU Usage

If the processes leading to high CPU usage are malicious, run the **top** command to stop them. If the **kswapd0** process leads to high CPU usage, optimize the program for the process or upgrade the ECS specifications for a larger memory capacity.

kswapd0 is a virtual memory management process. When the physical memory becomes insufficient, **kswapd0** runs to allocate disk swap capacity for caching. This uses a large number of CPU resources.

- For the detected malicious processes

Quickly stop such processes on the top page. To do so, perform the following operations:

- a. Press the **k** key during the execution of the **top** command.
- b. Enter the PID of the process to be stopped.

The PID of the process is the value in the first column of the top command output. For example, to stop the process with PID 52, enter **52** and press **Enter**.

```
top - 21:07:38 up 37 days, 9:21, 1 user, load average: 0.01, 0.02, 0.05  
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie  
%Cpu(s): 0.0 us, 3.2 sy, 0.0 ni, 96.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st  
KiB Mem : 3880024 total, 2961520 free, 178960 used, 739544 buff/cache  
KiB Swap: 0 total, 0 free, 0 used. 3434216 avail Mem  
PID to signal/kill [default pid = 1] 52  
PID USER PR NI VIRT RES SHR S %CPU %MEM TIME+ COMMAND  
1 root 20 0 125480 3884 2604 S 0.0 0.1 0:11.32 systemd  
2 root 20 0 0 0 0 S 0.0 0.0 0:00.00 kthreadd
```

- c. After the operation is successful, information similar to the following is displayed. Press **Enter**.

```
top - 21:07:38 up 37 days, 9:21, 1 user, load average: 0.01, 0.02, 0.05
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie
%Cpu(s): 0.0 us, 3.2 sy, 0.0 ni, 96.8 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 2961520 free, 178960 used, 739544 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3434216 avail Mem
Send pid 52 signal [15/sigterm]
```

PID	USER	PR	NI	UIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1	root	20	0	125480	3884	2604	S	0.0	0.1	0:11.32	systemd
2	root	20	0	0	0	0	S	0.0	0.0	0:00.00	kthreadd

- For the **kswapd0** process

To check the memory usage of a process, perform the following operations:

- a. Run the **top** command to check the resource usage of the **kswapd0** process.
- b. If the process remains in non-sleeping state for a long period, you can preliminarily determine that the system is consistently paging. In such a case, the high CPU usage is caused by insufficient memory.

```
Tasks: 81 total, 1 running, 79 sleeping, 1 stopped, 0 zombie
%Cpu(s): 0.2 us, 52.2 sy, 0.0 ni, 99.7 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
KiB Mem : 3880024 total, 3014820 free, 179024 used, 686180 buff/cache
KiB Swap: 0 total, 0 free, 0 used. 3433948 avail Mem
```

PID	USER	PR	NI	UIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
36	root	20	0	0	0	0	S	99.0	0.0	964:10.45	kswapd0
4595	nginx	20	0	125392	3576	1040	S	0.3	0.1	60:04.91	nginx
1	root	20	0	125480	3884	2604	S	0.0	0.1	0:11.47	systemd

- c. Run the **vmstat** command to check the virtual memory usage of the system.

If the **si** and **so** values are large, the system is frequently paging and the physical memory of the system is insufficient.

- **si**: Volume of data written from the swap partition to the memory per second, which is transferred from the disk to the memory.
 - **so**: Volume of data written from the memory to the swap partition per second, which is transferred from the memory to the disk.
- d. Further identify the causes of high memory usage by running commands such as **free** and **ps** to check the memory usage of the system and processes in the system.
 - e. Restart the application or release the memory when traffic is light.

To handle this issue, expand the ECS memory. If memory expansion is not allowed, optimize the application and enable hugepage memory.

Handling High Bandwidth Usage

If the high bandwidth usage is caused by normal service access of non-malicious processes, enlarge the bandwidth to handle this issue. If the high bandwidth usage is caused by abnormal service access, for example, malicious access from certain IP addresses, CC attacks on the ECS, or malicious processes, use the traffic monitoring tool **nethogs** to monitor the bandwidth usage of each process in real time and identify faulty processes.

- Using **nethogs** for troubleshooting
 - a. Run the following command to install **nethogs**:
yum install nethogs -y

After the installation, run the **nethogs** command to check bandwidth usage.

Parameters in the **nethogs** command are as follows:

- **-d**: Set the update interval in the unit of second. The default value is 1s.
- **-t**: Enable tracing.
- **-c**: Set the number of updates.
- **device**: Set the NIC to be monitored. The default value is **eth0**.

The following parameters are involved in command execution:

- **q**: Exit **nethogs**.
 - **s**: Sort processes in the process list by TX traffic in ascending or descending order.
 - **r**: Sort processes in the process list by RX traffic in ascending or descending order.
 - **m**: Switch the display unit in the sequence of KB/s, KB, B, and MB.
- b. Run the following command to check the bandwidth usage of each process on the specified NIC:

nethogs eth1

```
Nethogs version 0.8.5
```

PID	USER	PROGRAM	DEV	SENT	RECEIVED
4596	nginx	nginx: worker process	eth1	34.360	3.267 KB/sec
?	root	192.168.0.92:90-100.125.60.19:17873		0.179	0.246 KB/sec
?	root	192.168.0.92:11211-213.32.10.149:44945		0.000	0.000 KB/sec
?	root	192.168.0.92:20101-105.176.26.66:43400		0.000	0.000 KB/sec
?	root	unknown TCP		0.000	0.000 KB/sec
TOTAL				34.540	3.512 KB/sec

The parameters in the command output are as follows:

- **PID**: ID of the process.
 - **USER**: user who runs the process.
 - **PROGRAM**: IP addresses and port numbers of the process and connection, respectively. The former is for the server and the latter is for the client.
 - **DEV**: Network port to which the traffic is destined.
 - **SENT**: Volume of data sent by the process per second.
 - **RECEIVED**: Volume of data received by the process per second.
- c. Stop malicious programs or blacklist malicious IP addresses.
- To stop a malicious process, run the **kill PID** command.
- To blacklist a malicious IP address or limit its rate, use iptables.
- Using Web Application Firewall (WAF) to protect the ECS against CC attacks
If your ECS has encountered a CC attack, enable CC security protection on the WAF console. For instructions about how to use WAF, see [Configuring CC Attack Protection Rules to Defend Against CC Attacks](#).

1.9 How Can I Handle Slow ECS Startup?

If an ECS requires a long period of time to start, you can change the default timeout to speed up the startup.

1. Log in to the ECS.
2. Run the following command to switch to user **root**:
sudo su
3. Run the following command to obtain the grub version:
rpm -qa | grep grub

Figure 1-22 Viewing the grub version

```
[root@ ~]# rpm -qa | grep grub  
grub2-2.02-0.44.el7.centos.x86_64
```

4. Change the timeout in the grub file to 0s.

CAUTION

This operation involves modifying the GRUB configuration file. Misoperations may cause the system to fail to start. Before the operation, back up the GRUB configuration file so that it can be restored in case of misoperations.

- If the grub version is earlier than 2:
Open the **/boot/grub/grub.cfg** or **/boot/grub/menu.lst** file and change the **timeout** value to **0**.
- If the grub version is 2:
Open the **/boot/grub2/grub.cfg** file and change the **timeout** value to **0**.

Figure 1-23 Changing timeout duration

```
#boot=/dev/sda  
default=0  
timeout=0  
splashimage=(hd0,1)/boot/grub/splash.xpm.gz  
hiddenmenu  
title CentOS (2.6.32-696.16.1.el6.x86_64)  
root (hd0,1)  
kernel /boot/vmlinuz-2.6.32-696.16.1.el6.x86_64 ro root=UUID=2bc0f5fd-e0  
19-4ba5-8ce0-0fe12b6efc24 rd_NO_LUKS rd_NO_LVM LANG=en_US.UTF-8 rd_NO_MD SYSFONT  
=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb q
```

1.10 How Do I Configure Multiple IP Addresses for an ECS with Multiple NICs Attached?

Symptom

If an ECS has multiple NICs, perform the following operations to configure policy-based routing for the ECS to enable the network communication of secondary NICs.

Possible Causes

No routing rules are configured for the EIPs bound to the extension NICs.

Operation Guide

Table 1-3 provides the operation guides on how to configure policy-based routes for Linux and Windows ECSs.

Table 1-3 Operation instructions on how to configure policy-based routes

OS Type	IP Address Version	Procedure
Linux	IPv4	Take an ECS running CentOS 8.0 (64-bit) as an example.
	IPv6	For details, see section <i>Configuring Policy-based Routes for a Linux ECS with Multiple Network Interfaces</i> under <i>Configuring Policy-based Routes for an ECS with Multiple Network Interfaces</i> in the <i>Virtual Private Cloud Best Practices</i> .
Windows	IPv4	Take an ECS running Windows Server 2012 (64-bit) as an example.
	IPv6	For details, see section <i>Configuring Policy-based Routes for a Windows ECS with Multiple Network Interfaces</i> under <i>Configuring Policy-based Routes for an ECS with Multiple Network Interfaces</i> in the <i>Virtual Private Cloud Best Practices</i> .

2 Windows ECS Issues

2.1 How Can I Retain a Session on a Windows ECS?

Scenarios

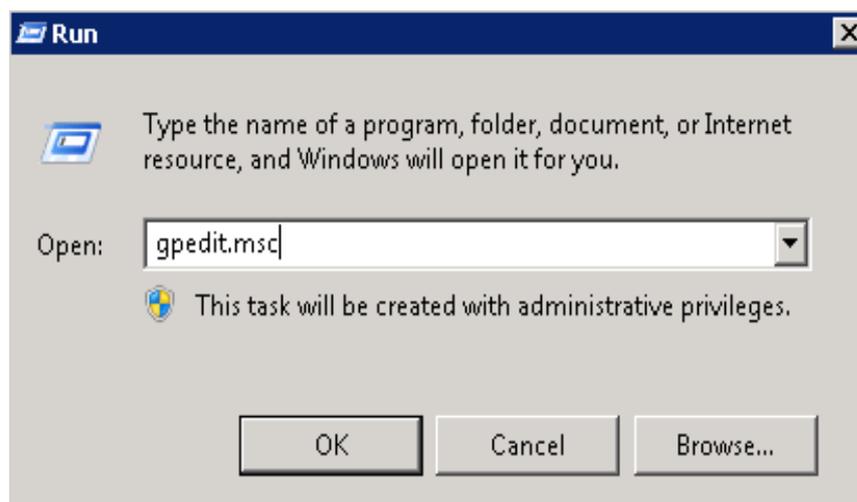
This section describes how to configure a Windows ECS so that its remote desktop connection will not be automatically disconnected.

Procedure

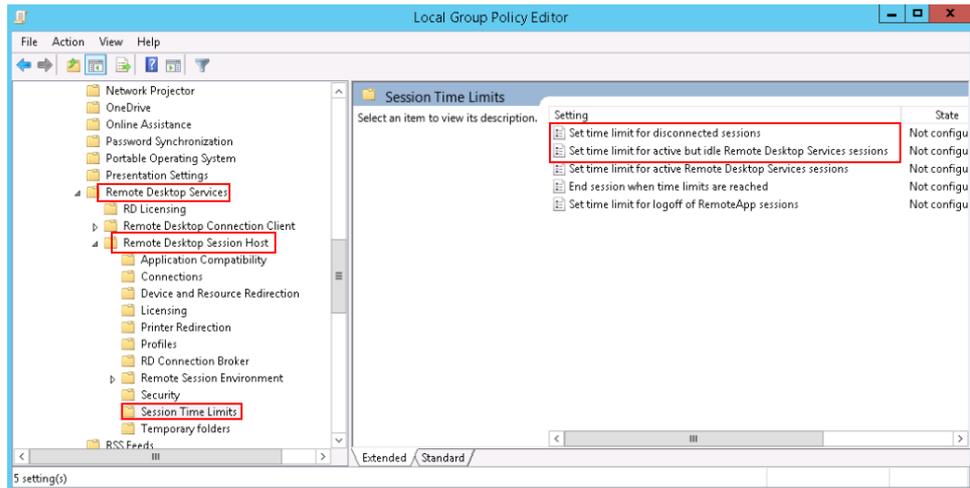
The following uses an ECS running Windows Server 2008 as an example.

1. Choose **Start > Run**. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start Local Group Policy Editor.

Figure 2-1 gpedit.msc

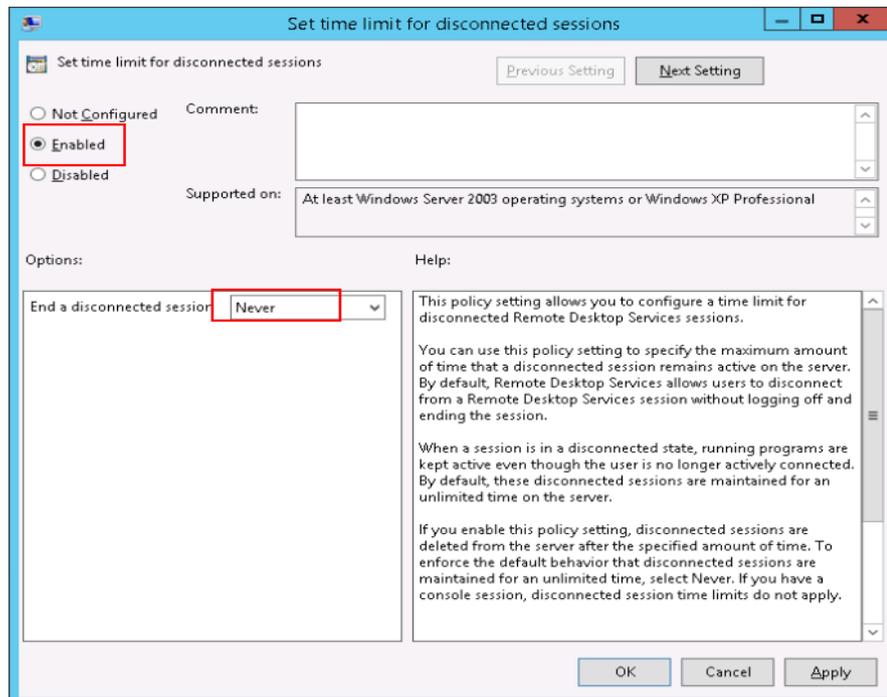


2. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Session Time Limits**.



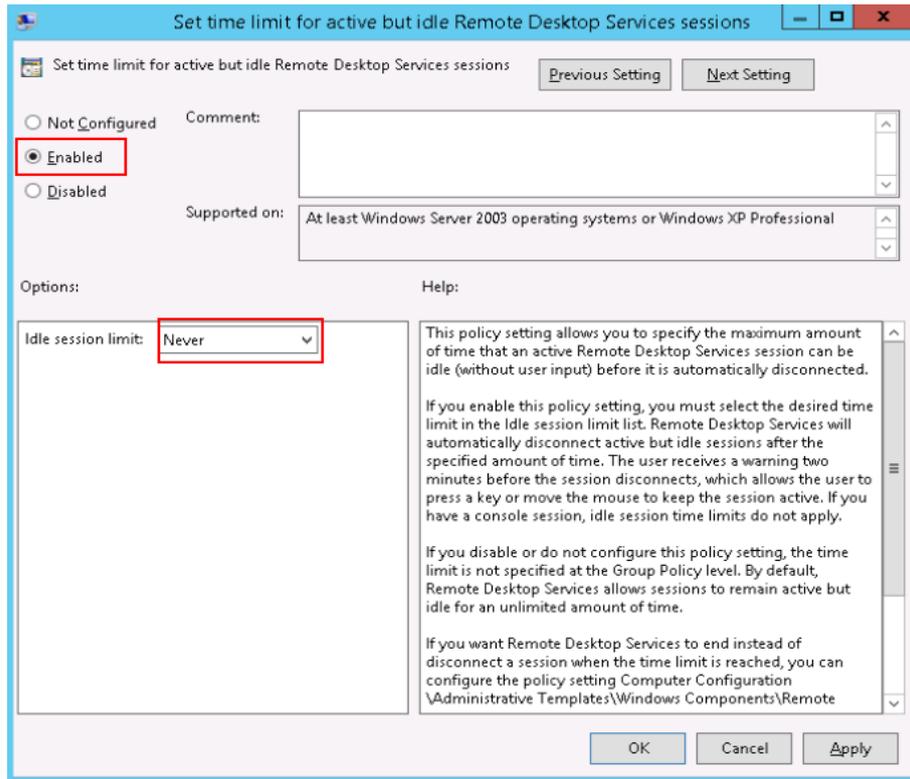
3. Set time limit for disconnected sessions.

- Select **Enabled**.
- Set **End a disconnected session** to **Never**.



4. Set time limit for active but idle Remote Desktop Services sessions.

- Select **Enabled**.
- Set **Idle session limit** to **Never**.

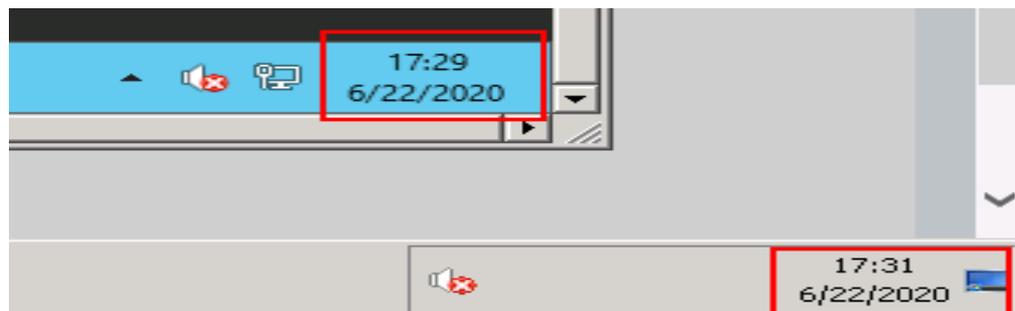


2.2 How Can I Fix the Difference Between the System Time and the Local Standard Time?

Symptom

The system time on your Windows ECS is different from the local standard time.

Figure 2-2 Difference between system time and local standard time



Possible Causes

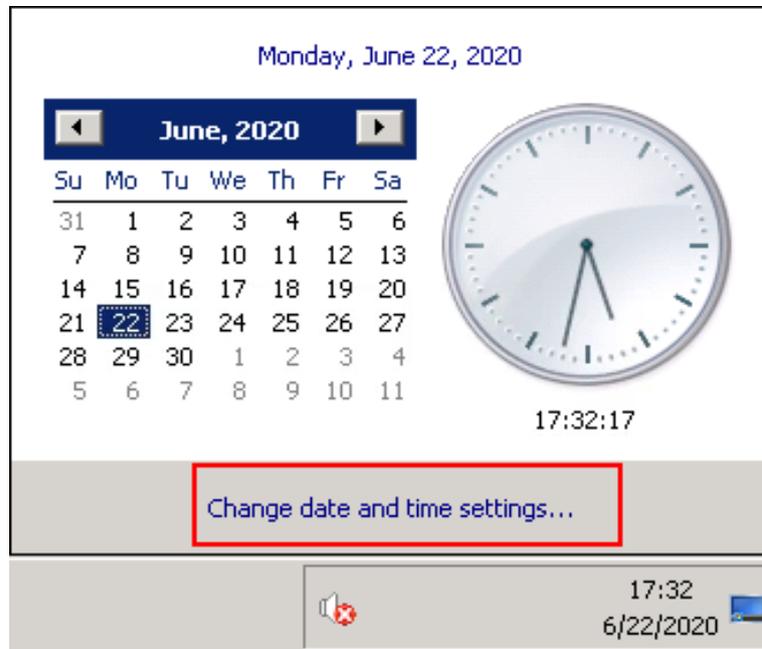
Affected by the network state, some drivers, or processes, the system time may be different from the standard time.

Solution 1

Manually synchronize system time.

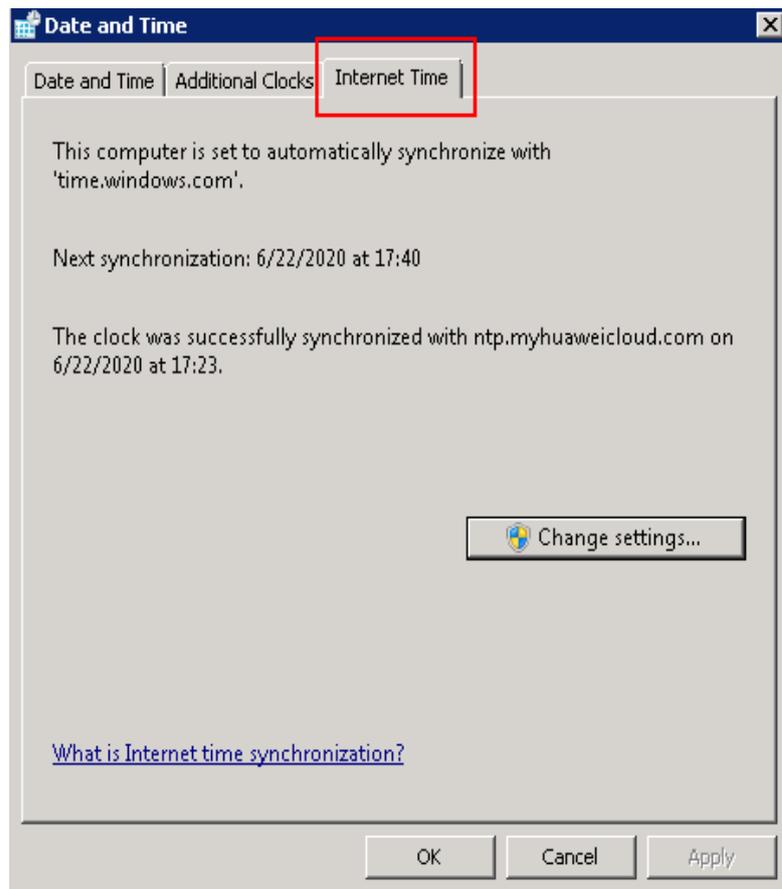
1. Click **Change date and time settings** in the lower right corner of the desktop. The **Date and Time** window is displayed.

Figure 2-3 Date and time



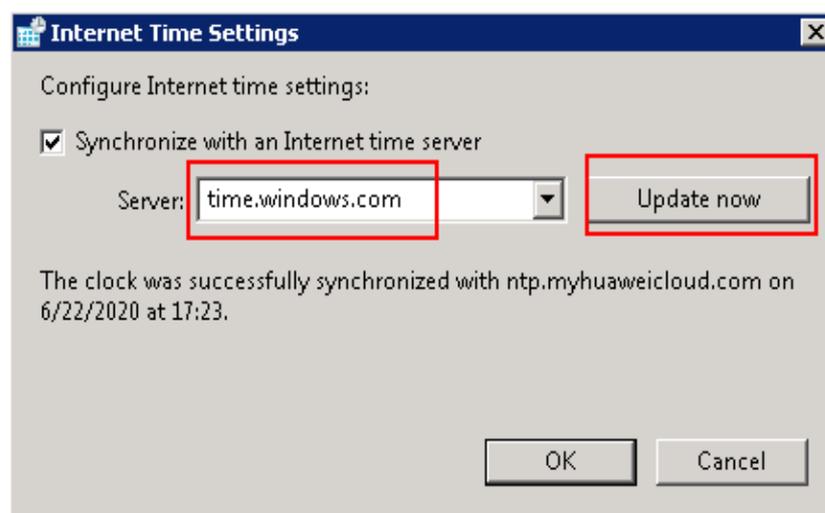
2. Click the **Internet Time** tab.

Figure 2-4 Internet time



3. Click **Change settings** and select a time source. The default time source is **time.windows.com**.
4. Click **Update now** and then click **OK**.

Figure 2-5 Selecting a time source



5. Check whether system time is consistent with the local standard time.

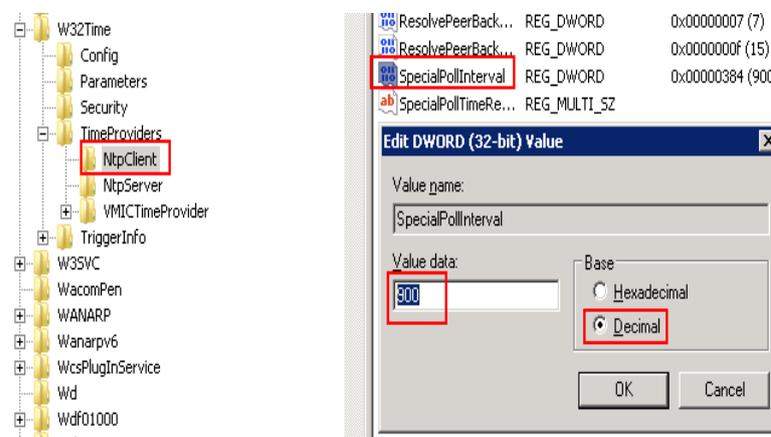
Solution 2

Change the update frequency of system time by editing the Windows registry.

1. In the **Run** dialog box, enter **regedit** to access the registry editor.
2. Choose **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Services > W32Time > TimeProviders > NtpClient** and double-click **SpecialPollInterval**.
3. In the dialog box that is displayed, set **Base** to **Decimal**.
4. Set the time synchronization interval.

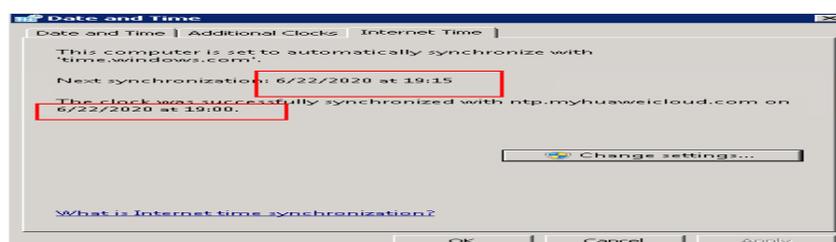
The value displayed in **Value data** is the interval (in seconds) for automatic time synchronization. Set this parameter as required.

Figure 2-6 Setting the time synchronization interval



5. Click **OK**.
6. After the configuration is complete, enter **cmd** to open the CLI and run the following command to update the group policy:
gpupdate
7. Check the Internet time. As shown in [Figure 2-7](#), the time synchronization frequency is changed to once every 15 minutes.

Figure 2-7 Viewing time synchronization frequency



Solution 3

Install the NTP server. For details, see [Does Huawei Cloud Provide the NTP Server and How Can I Configure It?](#)

2.3 How Do I Attach an Extension NIC to a Windows ECS for Accessing the Internet?

Scenarios

A Windows ECS has one primary NIC and one extension NIC attached. Both the NICs have an EIP bound to access the Internet.

Constraints

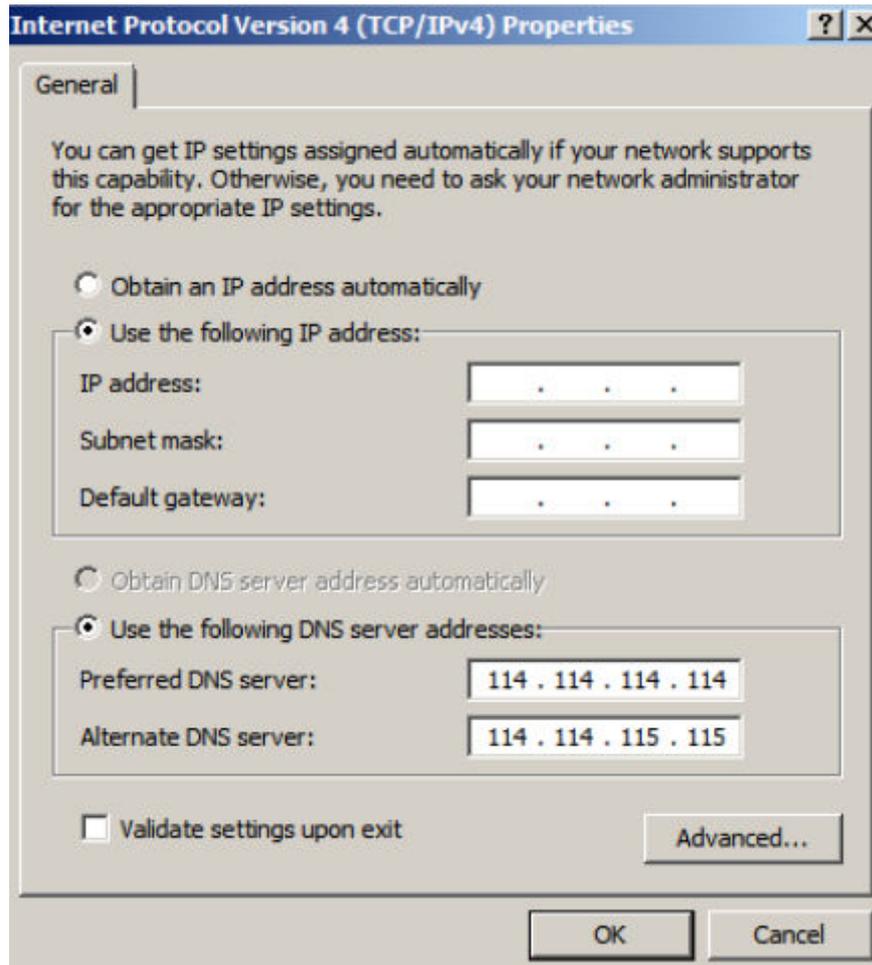
Do not modify the primary NIC settings.

Procedure

1. Log in to the management console and choose **Compute > Elastic Cloud Server**.
2. In the ECS list on the displayed page, select the ECS for which you want to add the NIC.
3. Click the name of the target ECS to go to the ECS details page.
4. Click the **Network Interfaces** tab.
5. Click **Attach Network Interface** and add an extension NIC as prompted.
Ensure that the security group and subnet of the extension NIC are the same as those of the primary NIC.
After the extension NIC is added, the system automatically obtains a private IP address.
6. Remotely log in to the ECS.
7. Click **Open Network and Sharing Center** in the lower right corner. Then, click **Change adapter settings** in the upper left corner and find the newly added NIC.



8. Right-click the target network connection and choose **Properties** from the shortcut menu.
9. In the **Local Area Connection 3 Properties** dialog box, click the **Networking** tab, select **Internet Protocol Version 4 (TCP/IPv4)**, and click **Properties**.
10. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box, click the **General** tab and configure the private IP address automatically obtained in step 5 for the newly added NIC.



Example configurations:

- Private IP address: 192.168.1.11
- IP address: 192.168.1.11
- Subnet mask: 255.255.255.0
- Default gateway: 192.168.1.1
- Preferred DNS server: 114.114.114.114
- Alternate DNS server: 114.114.115.115

11. After the configuration is complete, select **Validate settings upon exit**.
12. Restart the NIC and bind an EIP to it.

2.4 How Can I Fix Grayed Out Copy and Paste Options for a Windows ECS?

Symptom

After remotely logging in to a Windows ECS, I cannot copy or paste content and find that the **Paste** option grayed out.

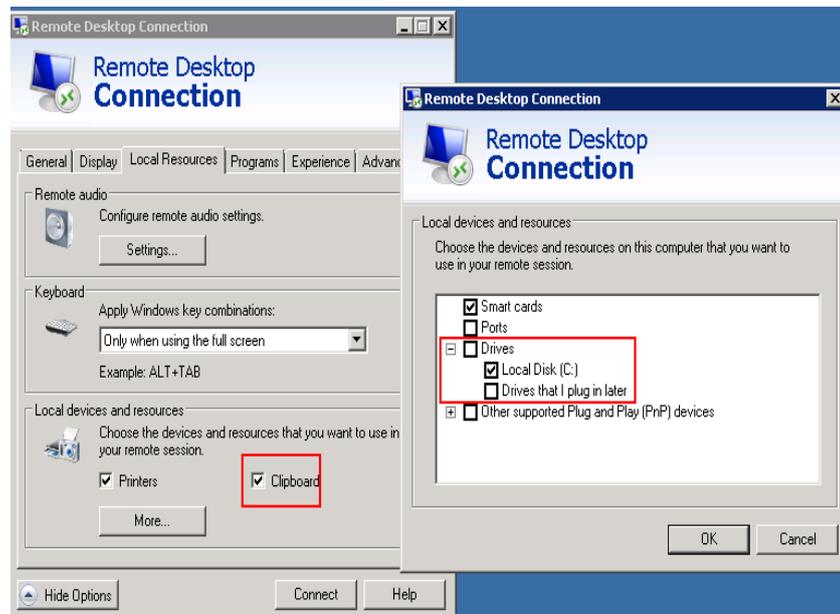
Possible Causes

- Local drive is not mapped.
- An error has occurred in the `rdpclip.exe` process.
- Data cannot be copied and pasted between the ECS and the local server.

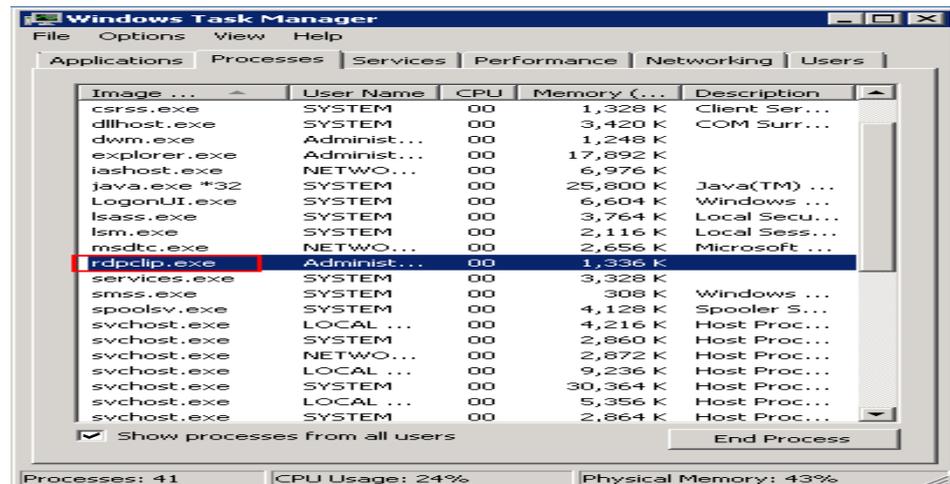
Solution

- The local drive is not mapped. **Figure 2-8** uses local drive C and drive D as an example.
 - a. Open the **Run** dialog box, enter `mstsc`, and click **OK** to start **Remote Desktop Connection**.
 - b. Click the **Local Resources** tab and select **Clipboard** in the **Local devices and resources** pane. Click **More** and select the drives you want to use in the remote session.
 - c. Click **OK** and check whether the copy and paste options work.

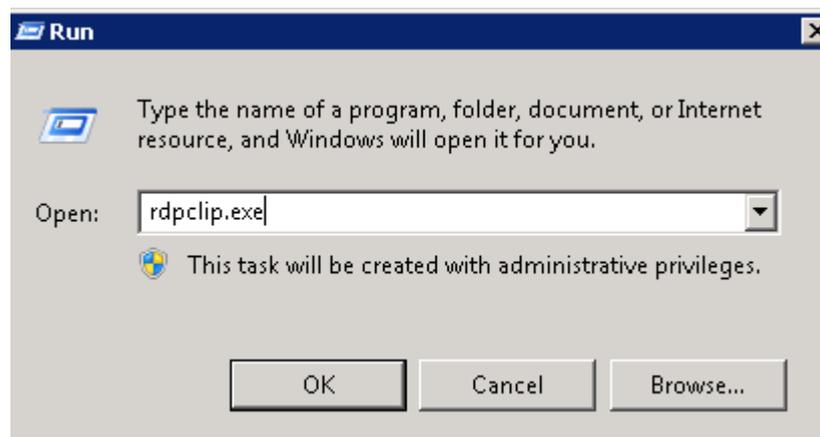
Figure 2-8 Local resources



- An error has occurred in the `rdpclip.exe` process.
 - a. Enable remote desktop connection, start **Windows Task Manager**, and stop the `rdpclip.exe` process on the **Processes** tab.

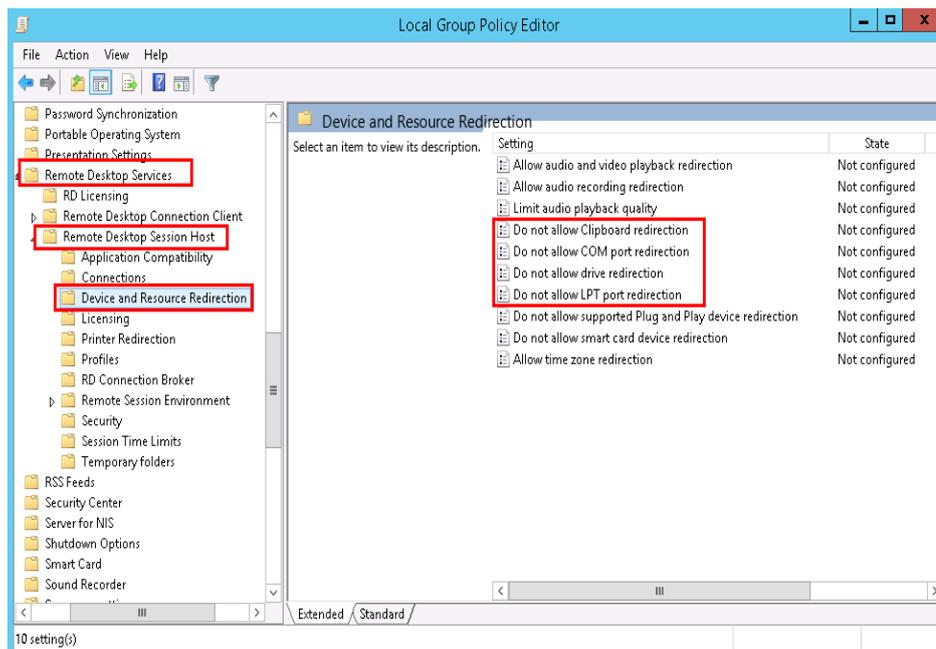
Figure 2-9 Ending the `rdpclip.exe` process

- b. Open the **Run** dialog box, enter `rdpclip.exe`, and click **OK** to restart it.

Figure 2-10 Starting the `rdpclip.exe` process

- c. Check whether the copy and paste options work.
- Data cannot be copied and pasted between the ECS and the local server.
 - a. Choose **Start > Run**. In the **Run** dialog box, enter `gpedit.msc` and click **OK** to start **Local Group Policy Editor**.
 - b. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Device and Resource Redirection**.
 - c. Set **Do not allow Clipboard redirection**, **Do not allow COM port redirection**, **Do not allow drive redirection**, and **Do not allow LPT port redirection** to **Disabled**.

Figure 2-11 Device and resource redirection



- d. Restart the ECS and check whether the copy and paste options work.

2.5 How Do I Configure File Sharing and Network Disk Mapping for a Windows ECS?

Scenarios

This section describes how to share a folder between Windows ECSs over an intranet.

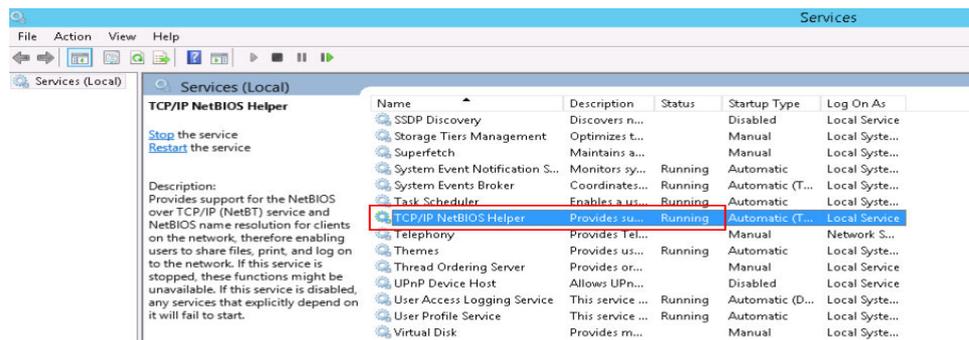
Constraints

Some carriers may block ports 139 and 445. As a result, you cannot access the shared folders over the Internet. Therefore, you are advised to share a folder between Windows ECSs only over an intranet.

Procedure

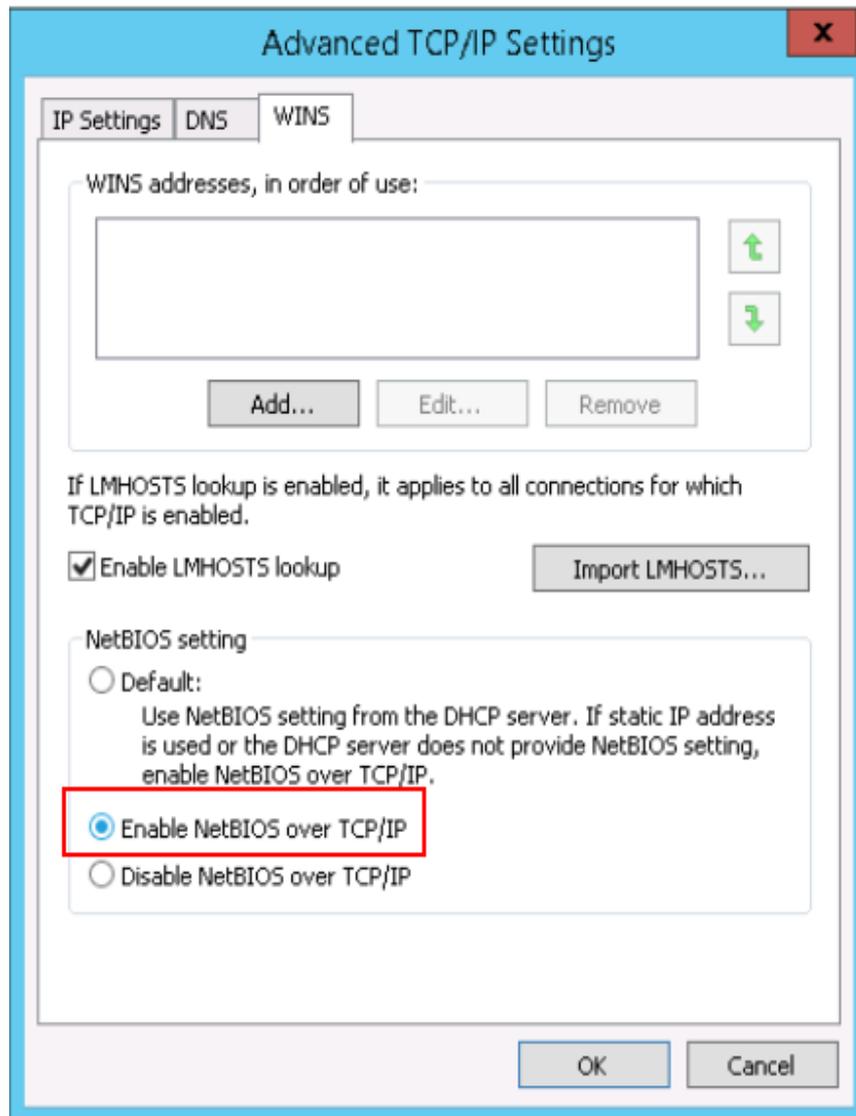
- Step 1** Check whether the two Windows ECSs that are to share a folder are correctly configured.
- Ensure that **TCP/IP NetBIOS Helper** is started.
Open the CLI, run the **services.msc** command, and locate **TCP/IP NetBIOS Helper** to check its status.

Figure 2-12 TCP/IP NetBIOS Helper



- Ensure that **Enable NetBIOS over TCP/IP** is selected on the NIC.
 Right-click the **Network** icon in the lower right corner and choose **Open Network and Sharing Center** from the shortcut menu. Click **Change adapter settings**. Right-click **Ethernet** and choose **Properties** from the shortcut menu. Double-click **Internet Protocol Version 4 (TCP/IPv4)**. Choose **Advanced TCP/IP Settings > WINS** and select **Enable NetBIOS over TCP/IP**.

Figure 2-13 Enabling NetBIOS over TCP/IP

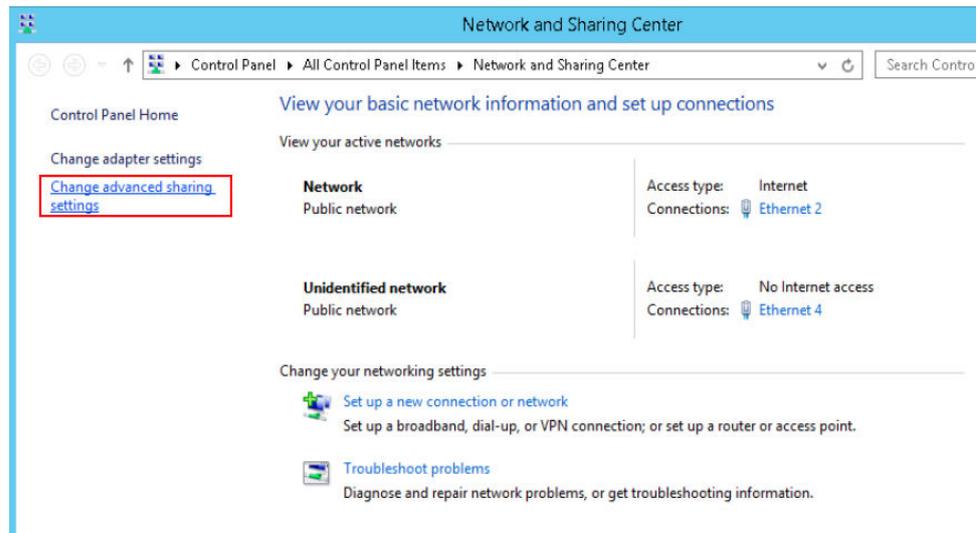


- Ensure that the inbound rules are added for ports 139 and 445 on the Windows ECS firewall.

Step 2 Configure **Network and Sharing Center**

1. Open **Open Network and Sharing Center**.
2. Click **Change advanced sharing settings**.

Figure 2-14 Change advanced sharing settings



3. Ensure the network types of the two Windows ECSs are same. For example, both are public or private. Select **Turn on network discovery** and **Turn on sharing so anyone with network access can read and write files in the Public folders**.

Figure 2-15 Enabling network discovery

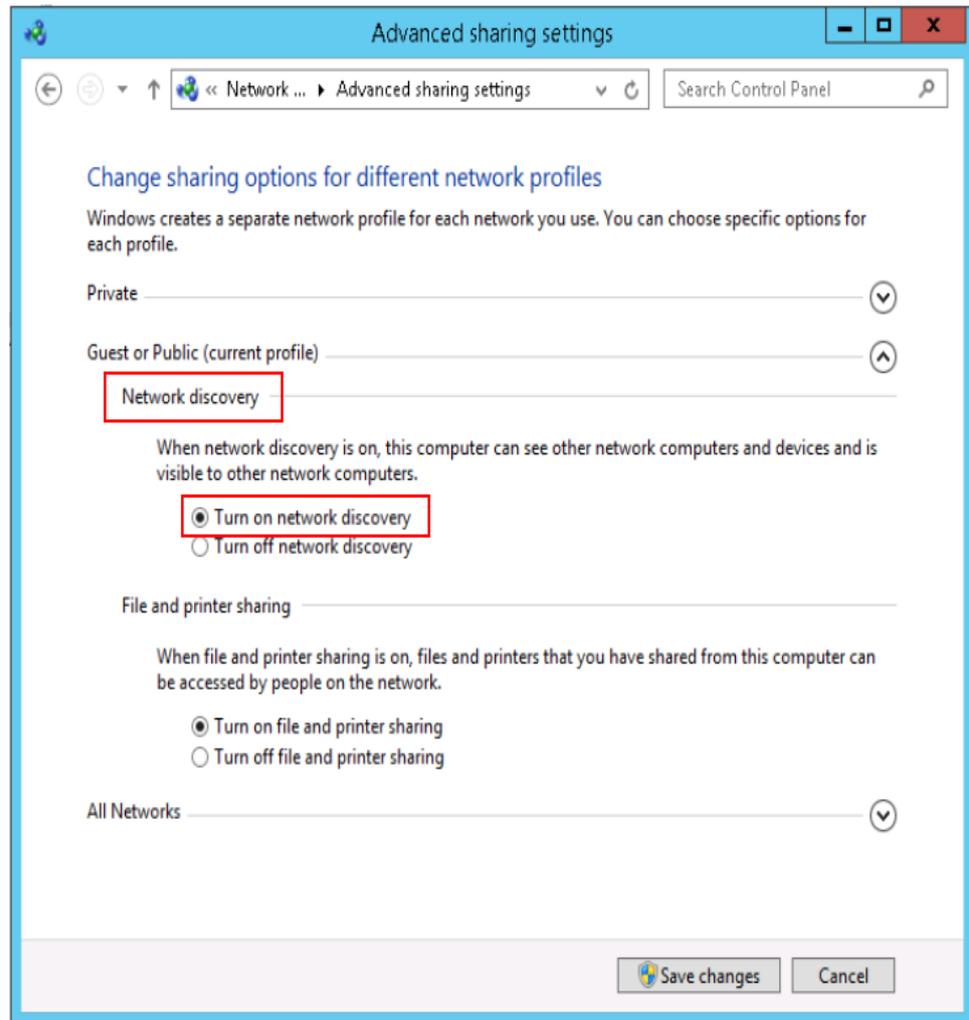
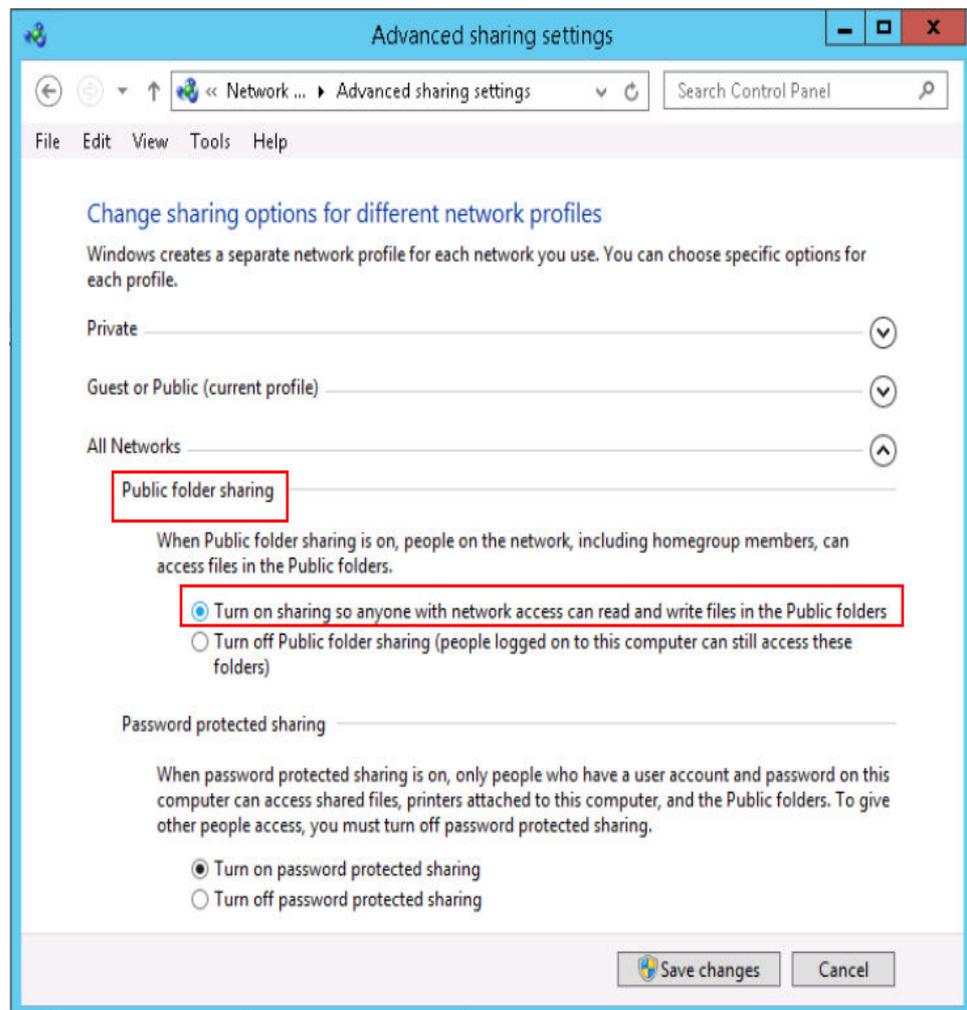


Figure 2-16 Enabling public folder sharing



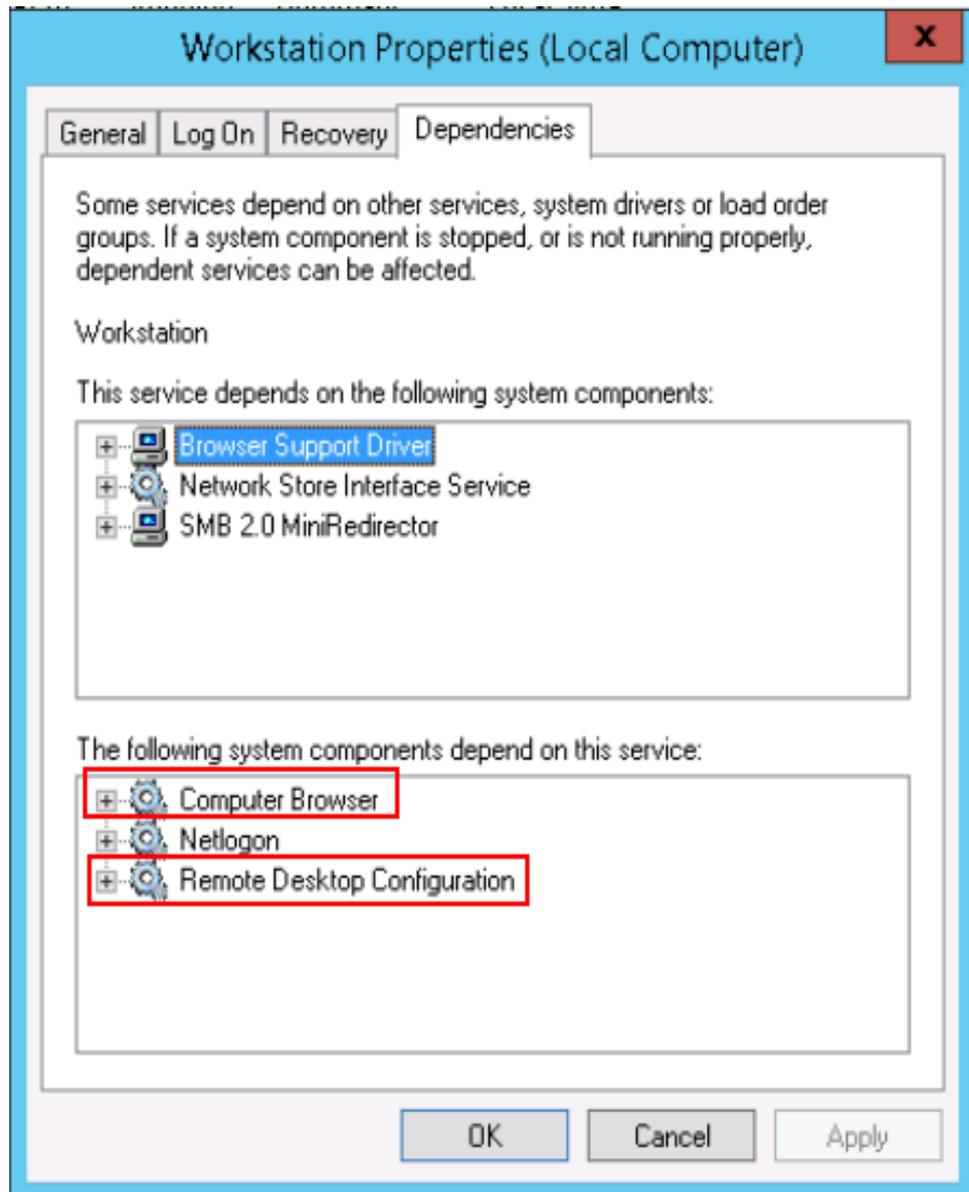
If the network discovery function cannot be enabled, you can run the **services.msc** command to start the Services manager. Check whether the following services on which the network discovery function depends are enabled:

- Function Discovery Resource Publication
- SSDP Discovery
- UPnP Device Host

4. In the Services manager, start the **Workstation** service.

This service depends on the **Computer Browser** and **Remote Desktop Configuration** components. Before starting the **Workstation** service, start the two components.

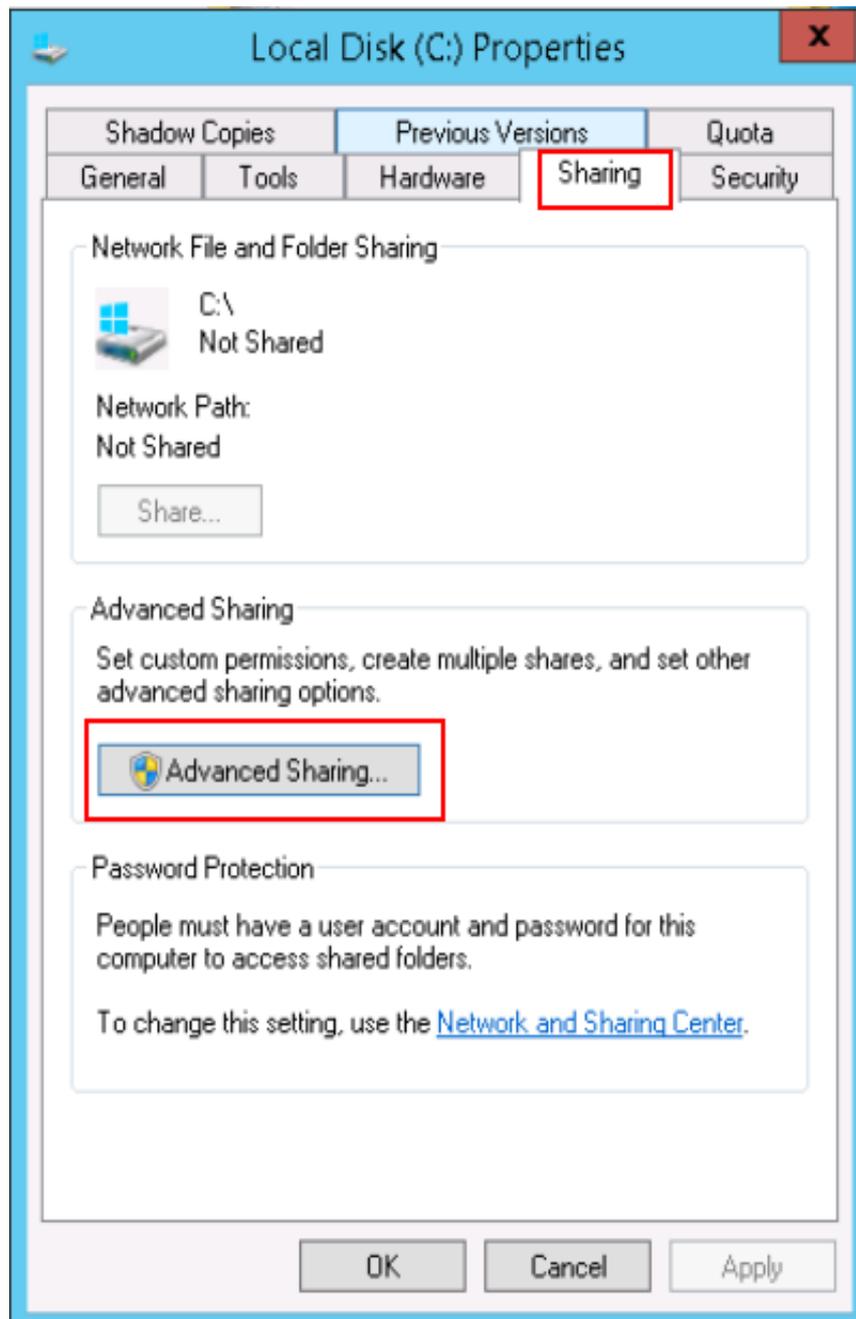
Figure 2-17 Enabling the **Workstation** service



Step 3 Configure disk sharing on the ECS that needs to access the shared disk.

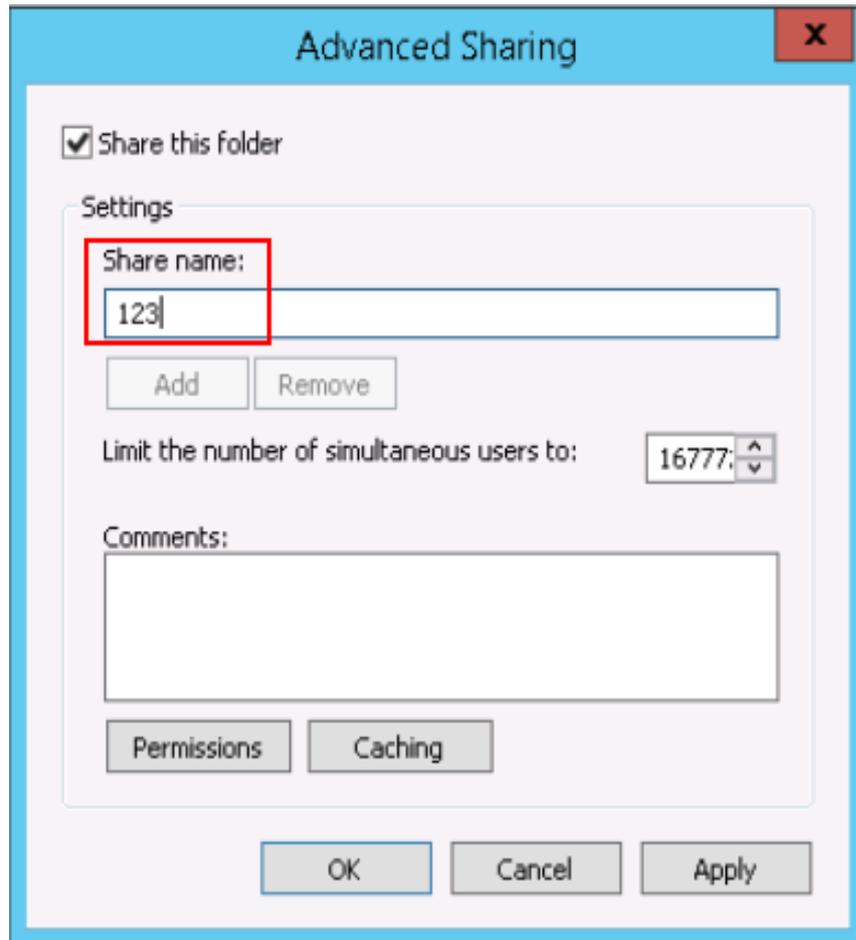
1. Right-click the disk to be shared and select **Properties**.
2. Click the **Sharing** tab and then click **Advanced Sharing**.

Figure 2-18 Configuring disk sharing



3. Set **Share name** and click **OK**.
Customize a name for the folder to be shared.

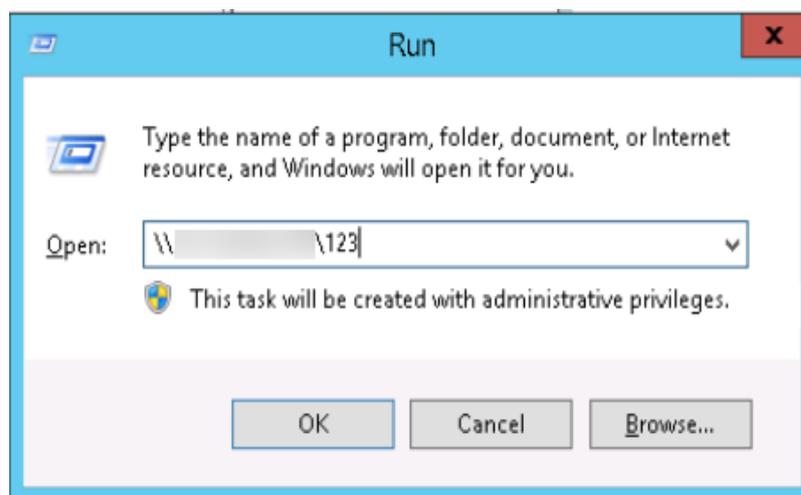
Figure 2-19 Entering a name for the folder to be shared



Step 4 Start another ECS in the same region and access the shared folder over the intranet.

1. Start the **Run** dialog box.
2. Enter `\\Intranet IP address\123` and click **OK** to open the shared folder.

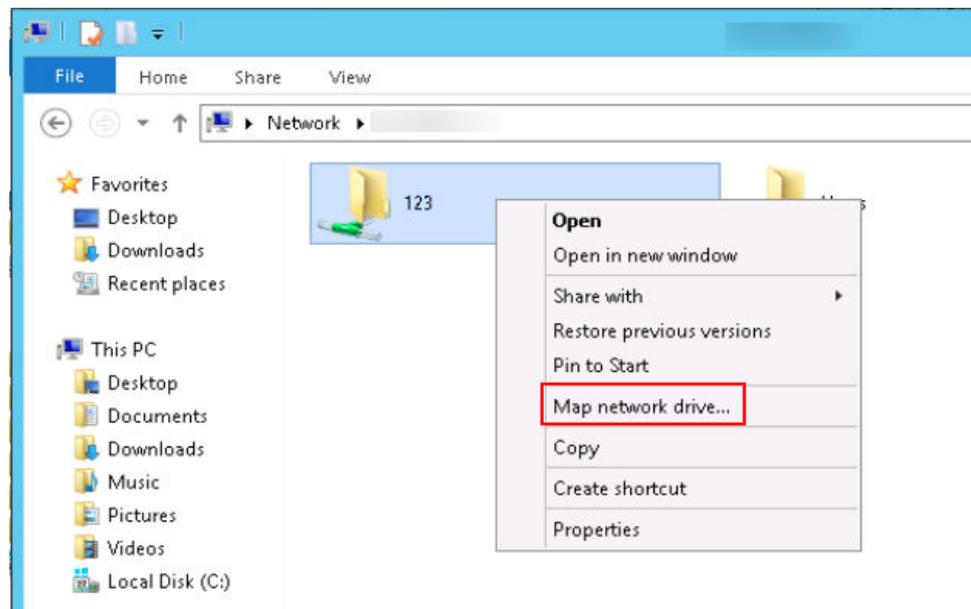
Figure 2-20 Assessing the shared folder



Step 5 (Optional) Create a network drive mapping for the shared path. Perform this operation to access the shared folder more conveniently.

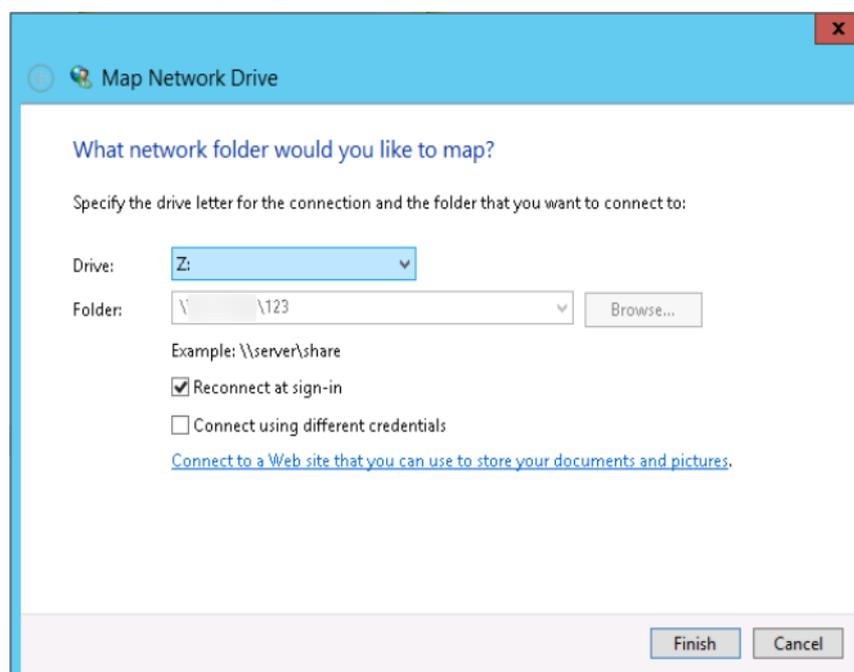
1. Right-click the shared folder and choose **Map network drive** from the shortcut menu.

Figure 2-21 Creating a mapping



2. Map the network drive.

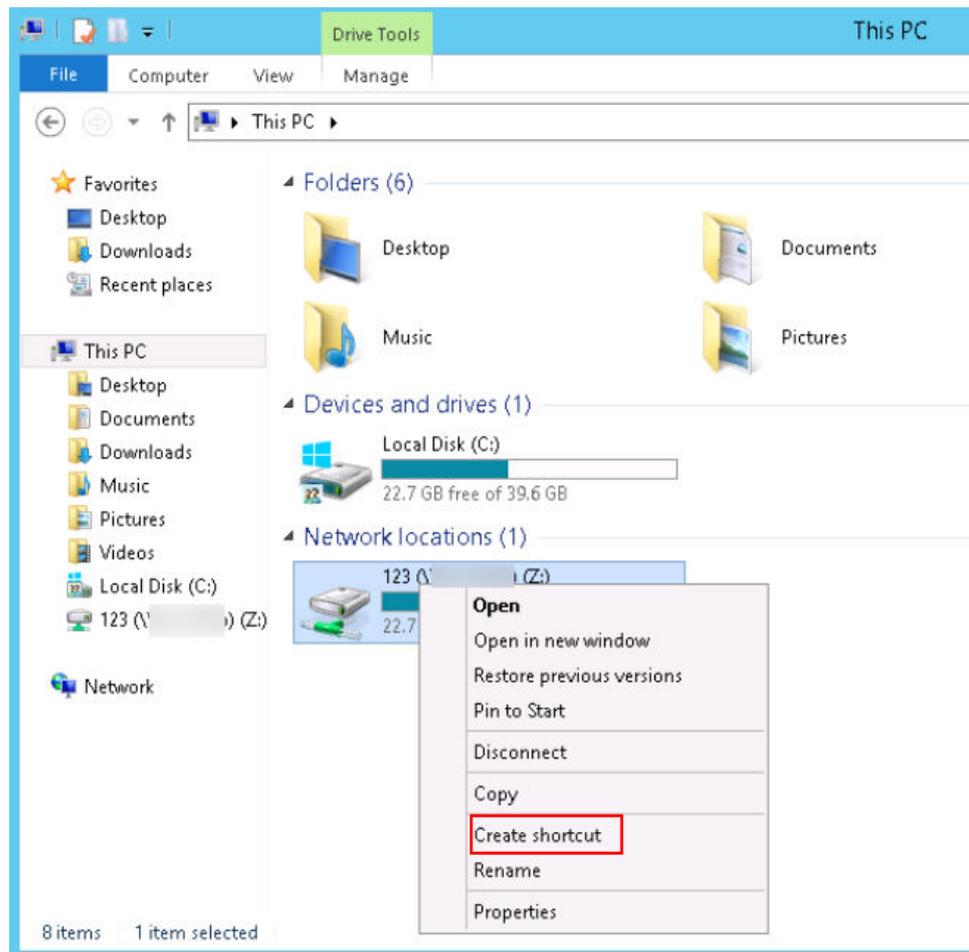
Figure 2-22 Mapping the network drive



3. Click the name of the mapped network drive, choose **Create shortcut**, and send the shortcut to the desktop.

Double-click the shortcut to quickly access the shared folder.

Figure 2-23 Creating a shortcut



----End

2.6 How Do I Troubleshoot an In-Service Port During Tomcat Startup?

Symptom

The system prompts that the required port is being used when Tomcat is started on a Windows ECS.

This section uses Windows Server 2008 R2 and port 80 as an example to describe how to resolve this issue.

Possible Causes

The port 80 required by Tomcat is being used by other programs, viruses, or Trojan horses.

1. Run the **netstat -ano | find "80"** command and find that the ID of the process using port 80 is 4.

NOTE

Change the port number as required.

Figure 2-24 Checking the process using port 80

```
C:\Users>netstat -ano | find ":80"
TCP    0.0.0.0:80          0.0.0.0:0          LISTENING        4
```

2. Run the **tasklist /svc | find "4"** command and find that the process using port 80 is the System process.

NOTE

Change the process ID as required.

Figure 2-25 System process

```
C:\Users>tasklist /svc | find "4"
System 4
```

Port 80 is used by the System process.

Solution

NOTE

Stopping the process that is using port 80 may stop the applications that are running or restart the ECS.

Solution 1:

1. Run the **cmd** command as the administrator and enter **net stop http**.
2. To stop the process that is using port 80, enter **Y**.
3. Run the **sc config http start= disabled** command.

Solution 2:

1. In the **cmd** window, run **regedit** to open the registry editor.
2. In **Registry Editor**, choose **HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Services > HTTP** and change the value of **Start** to **0**.
3. Restart the ECS.

2.7 How Do I Troubleshoot Unavailable IME Issues?

Symptom

- On my Windows ECS, you cannot switch input methods using **Ctrl+Shift**. The input method editor (IME) toolbar is not shown.
- The IME does not work.

Possible Causes

- The **ctfmon.exe** process is not started.

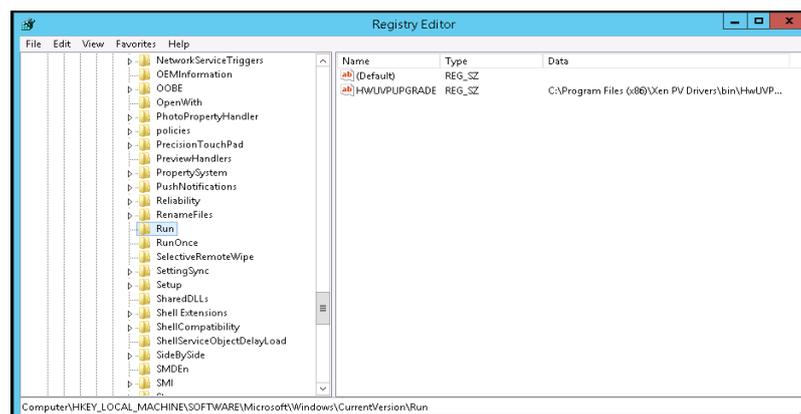
- The input indication icon is turned off.

Solutions to Unavailable IME

Solution 1: Check whether **ctfmon.exe** is started.

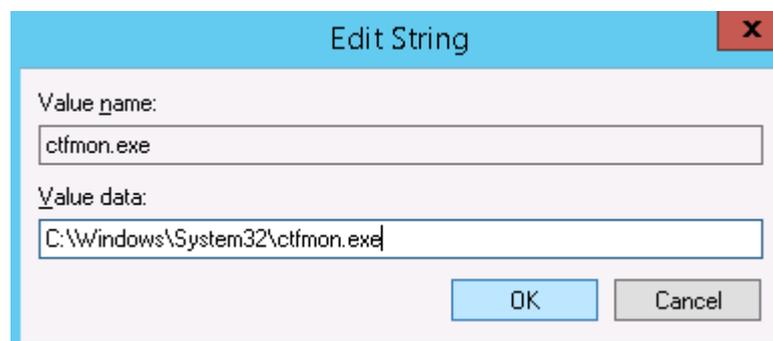
1. Open **C:\Windows\System32**.
2. Double-click **ctfmon.exe** and check whether the IME toolbar is shown in the lower right corner and whether you can switch the input method.
3. Enable the IME to automatically start upon ECS startup.
 - a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
 - b. In the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run** directory, create a file named **ctfmon.exe** in string type.

Figure 2-26 Run directory



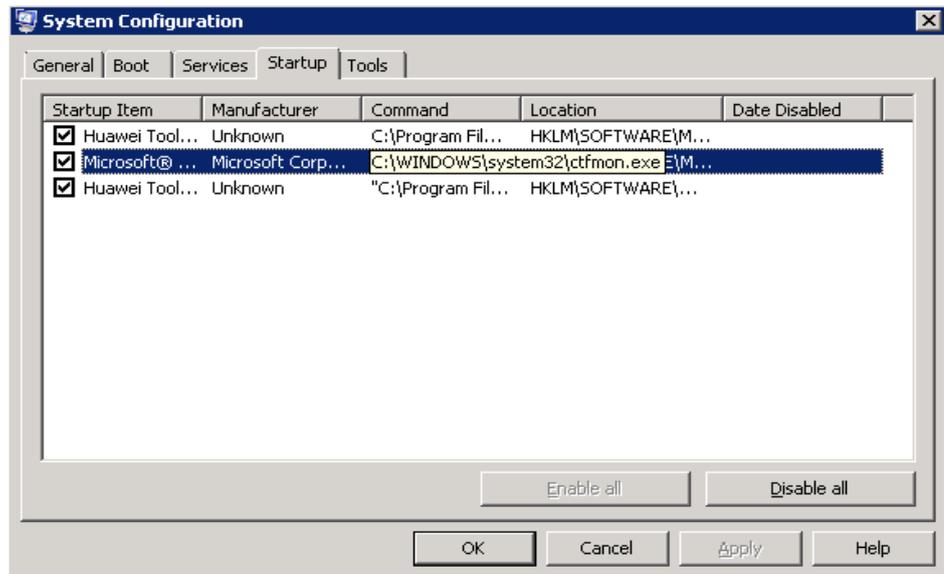
- c. Right-click the file you create and click **Modify**. Change the value data to **C:\Windows\System32\ctfmon.exe** in which **ctfmon.exe** is stored.

Figure 2-27 Editing the value data



- d. In the **Run** dialog box, enter **msconfig** and click **OK** to open the **System Configuration**.
- e. Click the **Startup** tab, find and select **ctfmon.exe**, and click **Apply** and then **OK**. Save the configurations and log out. Restart the ECS.

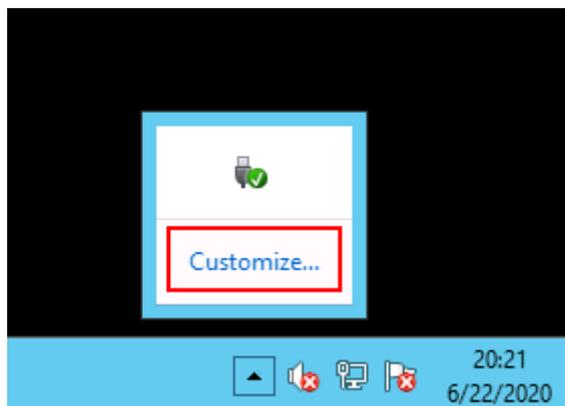
Figure 2-28 Selecting **ctfmon.exe**



Enable the input indication icon.

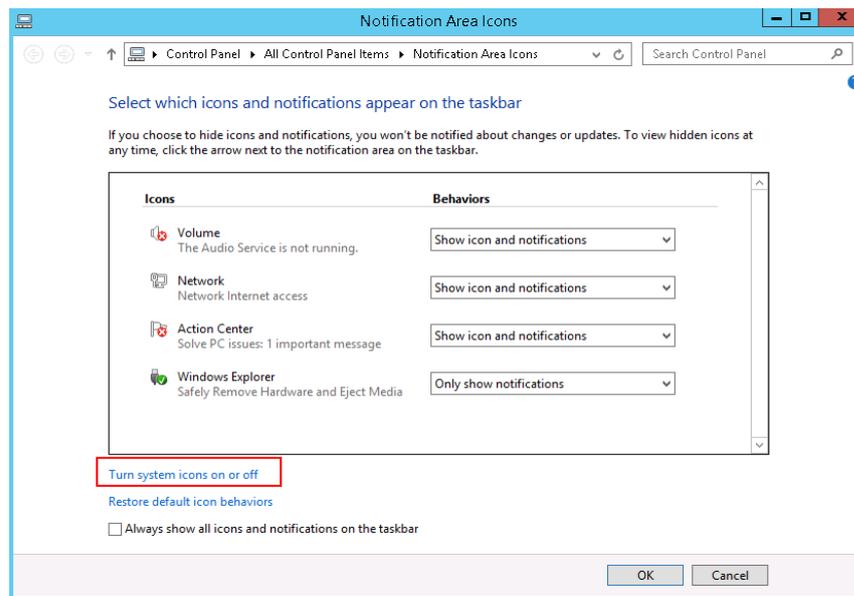
1. Log in to the ECS and click **Customize** in the lower right corner.

Figure 2-29 Clicking **Customize**



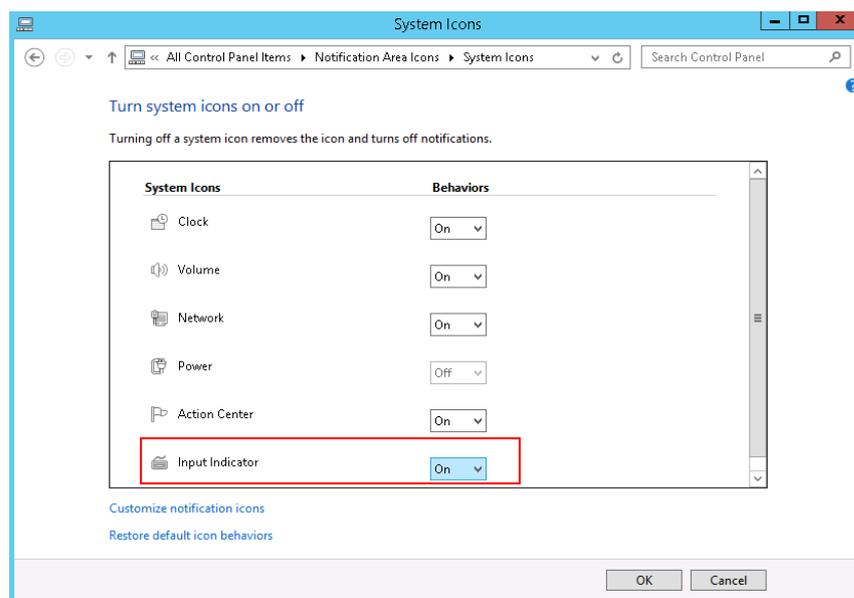
2. In the **Notification Area Icons** window, click **Turn system icons on or off**.

Figure 2-30 Turning system icons on or off



3. Set **Input Indicator** to **On** and click **OK**.

Figure 2-31 Turning input indicator on

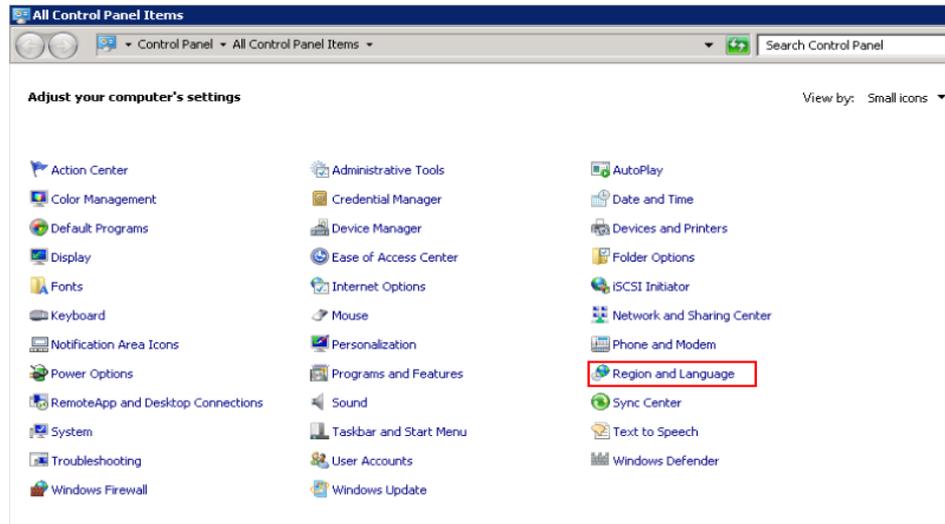


Adding Another Input Language

The following uses an ECS running Windows Server 2008 as an example to describe how to add the Japanese to IME.

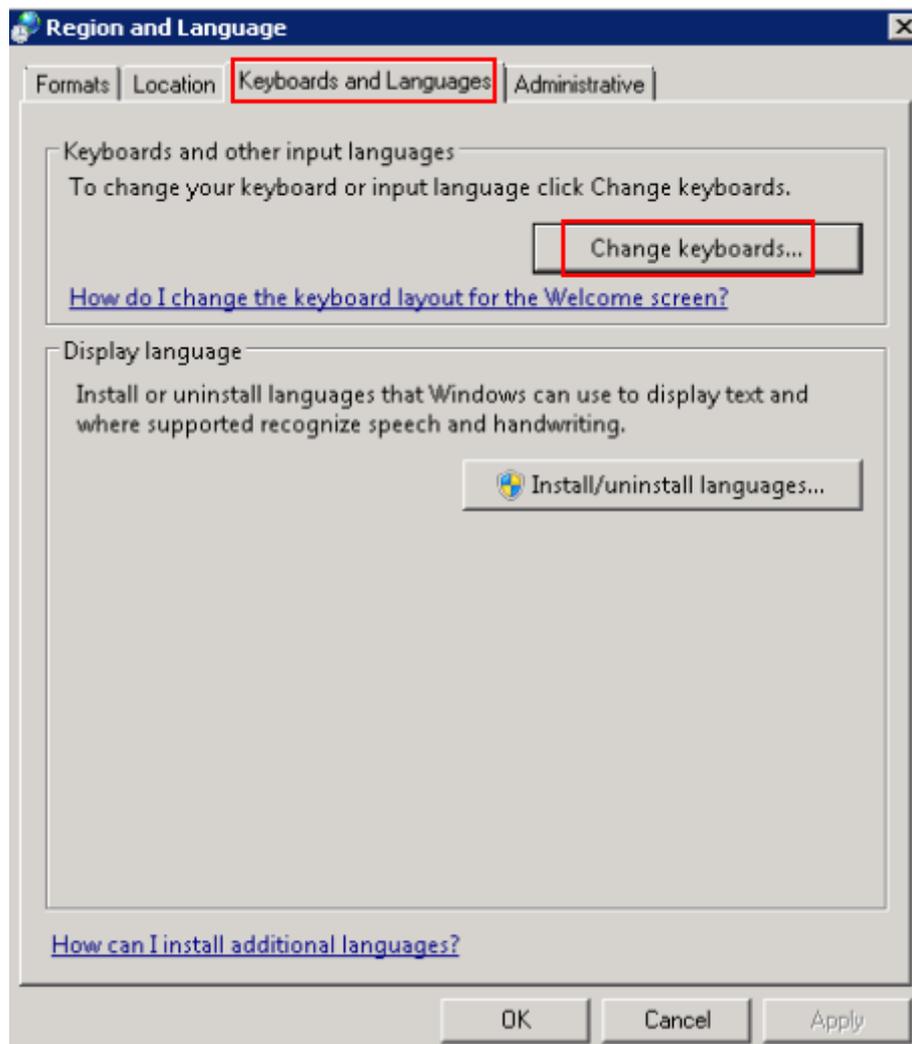
1. Log in to the ECS and open the control panel.
2. Click **Region and Language**.

Figure 2-32 Region and Language



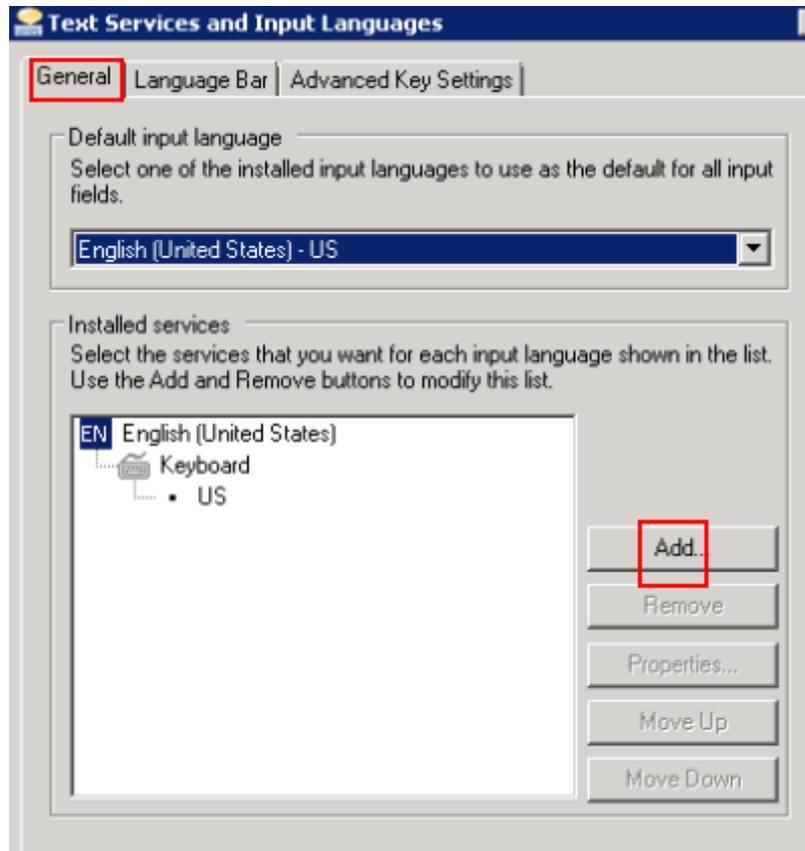
3. On the **Keyboards and Languages** tab, click **Change keyboards**.

Figure 2-33 Keyboards and Languages



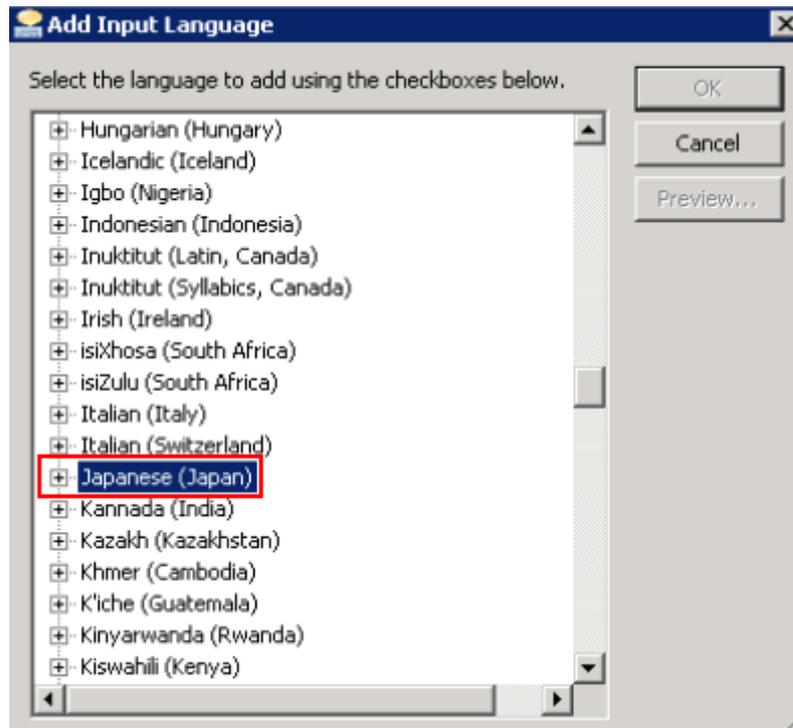
4. Right-click the language bar and click **Settings** to open the **Text Services and Input Languages** dialog box. On the **General** tab, click **Add**.

Figure 2-34 Text Services and Input Languages



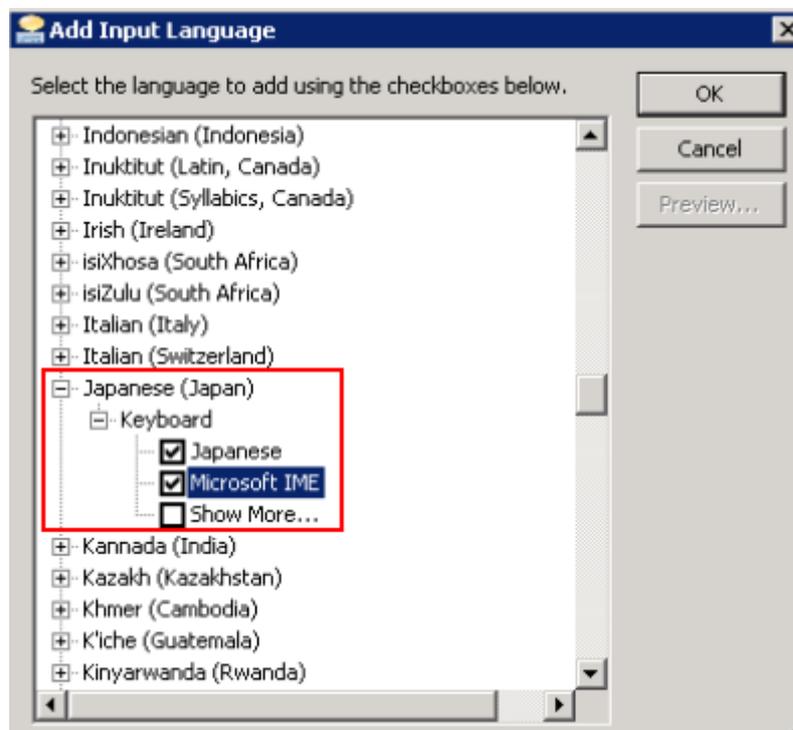
5. In the **Add Input Language** dialog box, find and double-click **Japanese (Japan)** and then **Keyboard**.

Figure 2-35 Adding input language



6. Select **Japanese** and **Microsoft IME**. Click **OK** and then **Apply** to save the configurations.

Figure 2-36 Selecting Japanese



2.8 How Can I Set the Input Method for a Windows ECS?

Symptom

You cannot switch the input method for a Windows ECS.

Possible Causes

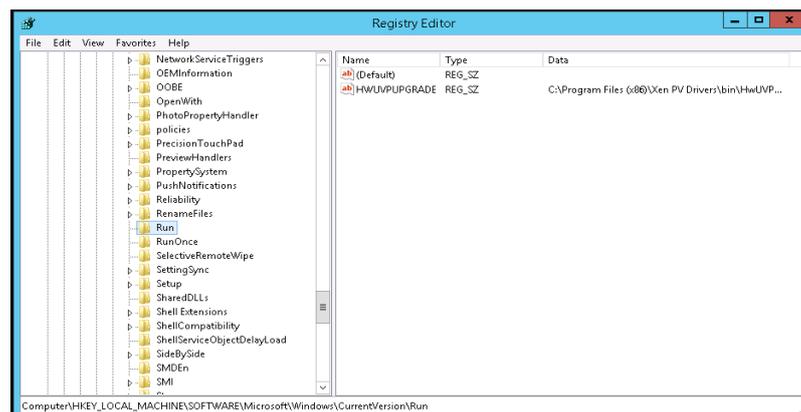
- The **ctfmon.exe** process is not started.
- The input indication icon is turned off.

Solutions to a Missing Language Bar

Solution 1: Check whether **ctfmon.exe** is started.

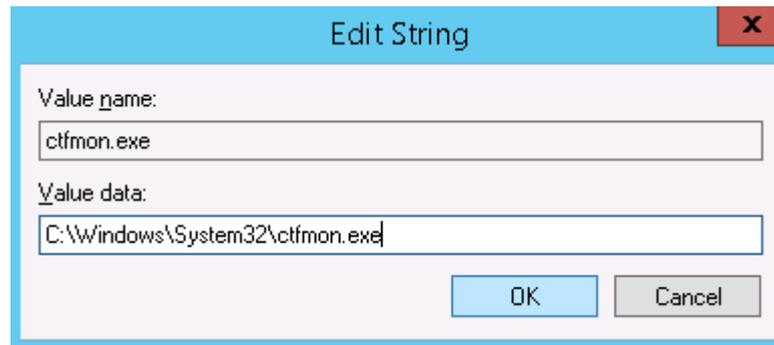
1. Open **C:\Windows\System32**.
2. Double-click **ctfmon.exe** and check whether the input method editor (IME) toolbar is shown in the lower right corner and whether you can switch the input method.
3. Enable the IME to automatically start upon ECS startup.
 - a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
 - b. In the **HKEY_LOCAL_MACHINE > SOFTWARE > Microsoft > Windows > CurrentVersion > Run** directory, create a file named **ctfmon.exe**.

Figure 2-37 Run directory



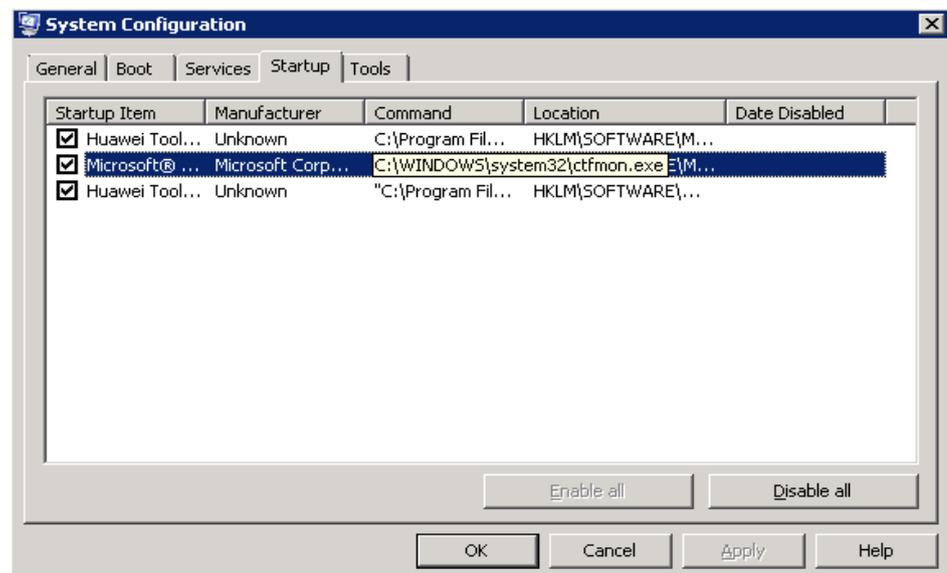
- c. Right-click the file you created and click **Modify**. Change the **Value data** to **C:\Windows\System32\ctfmon.exe**.

Figure 2-38 Editing the value data



- d. In the **Run** dialog box, enter **msconfig** and click **OK** to open **System Configuration**.
- e. Click the **Startup** tab, find and select **C:\WINDOWS\system32\ctfmon.exe**, and click **Apply** and then **OK**. Save the configurations and log out. Restart the ECS.

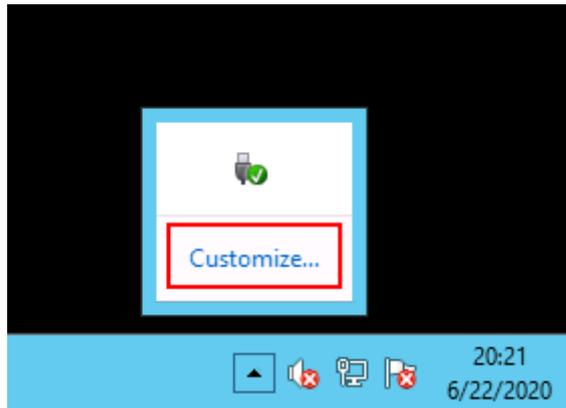
Figure 2-39 Selecting C:\WINDOWS\system32\ctfmon.exe



Solution 2: Enable the input indication icon.

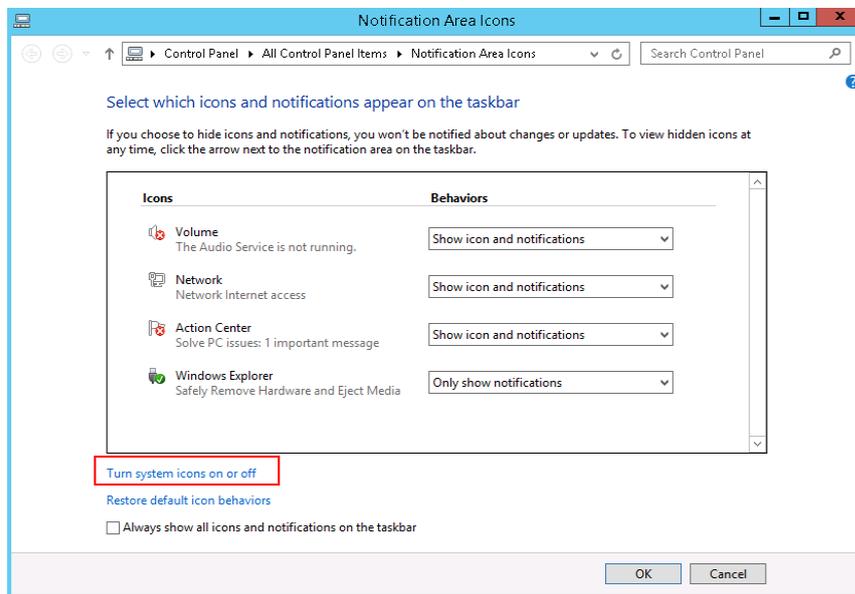
- 1. Log in to the ECS and click **Customize** in the lower right corner.

Figure 2-40 Clicking Customize



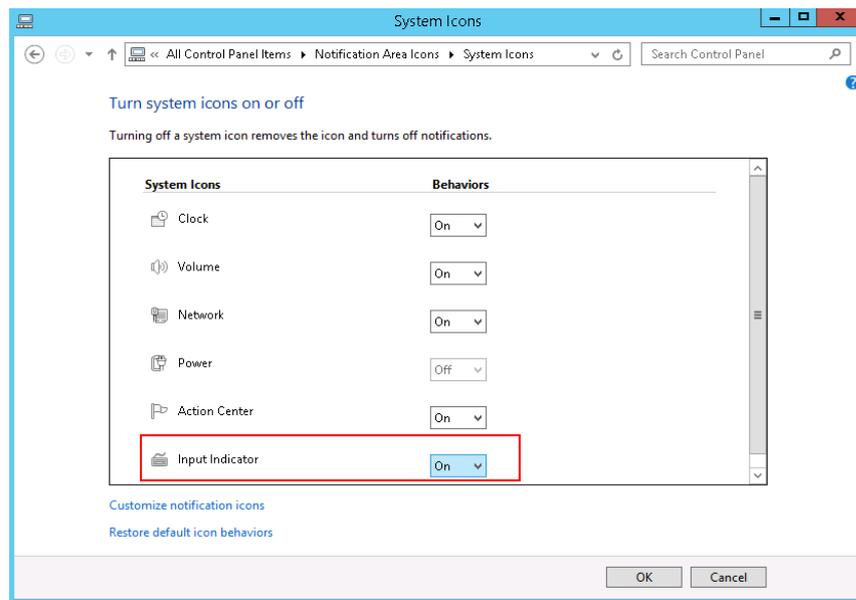
2. In the **Notification Area Icons** window, click **Turn system icons on or off**.

Figure 2-41 Turning system icons on or off



3. Set **Input Indicator** to **On** and click **OK**.

Figure 2-42 Turning input indicator on

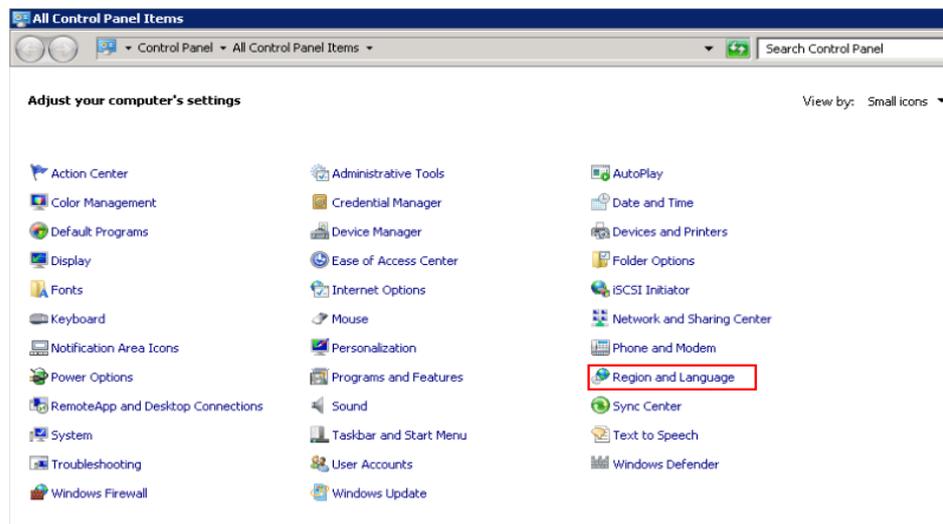


Adding Another Input Language

The following uses an ECS running Windows Server 2008 as an example to describe how to add the Japanese to IME.

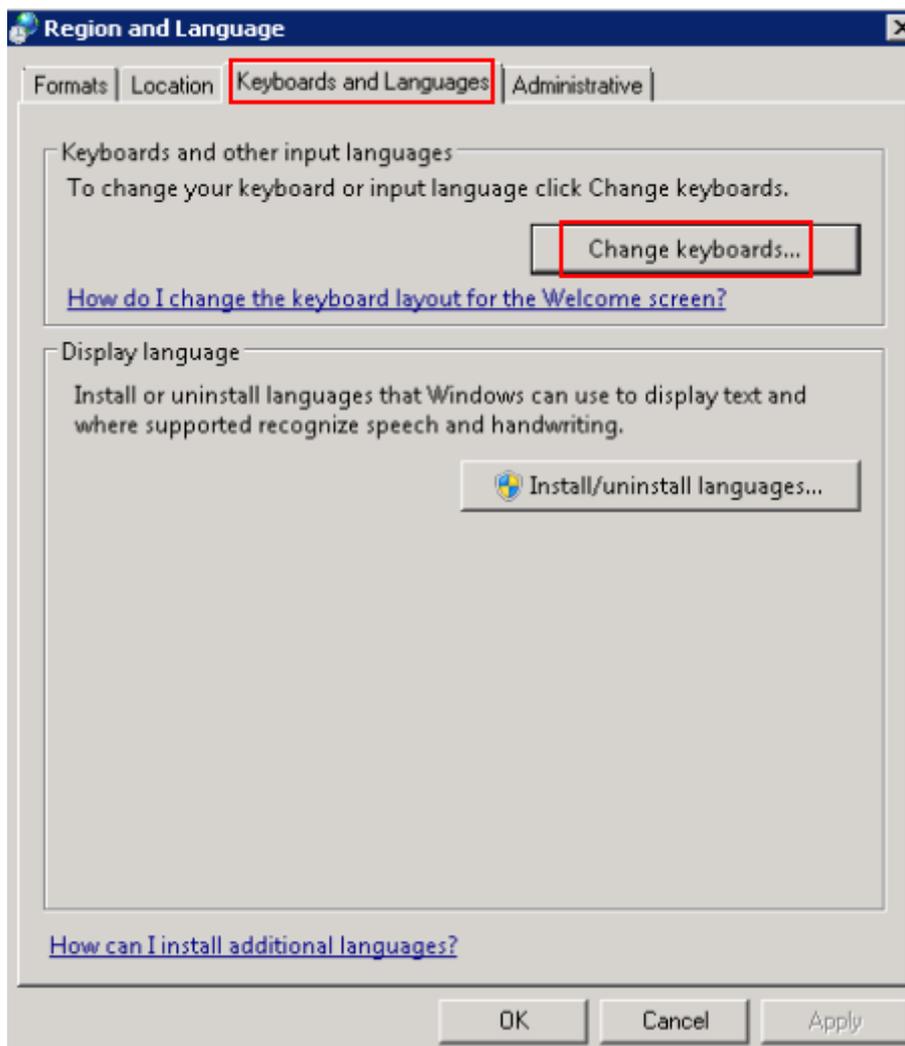
1. Log in to the ECS and open the control panel.
2. Click **Region and Language**.

Figure 2-43 Region and Language



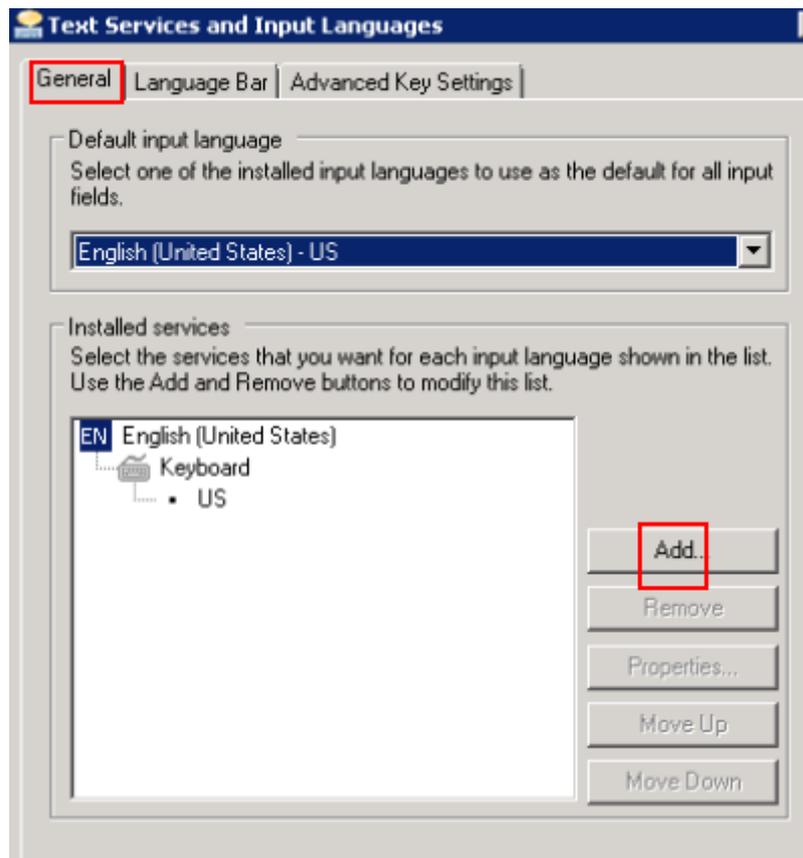
3. On the **Keyboards and Languages** tab, click **Change keyboards**.

Figure 2-44 Keyboards and Languages



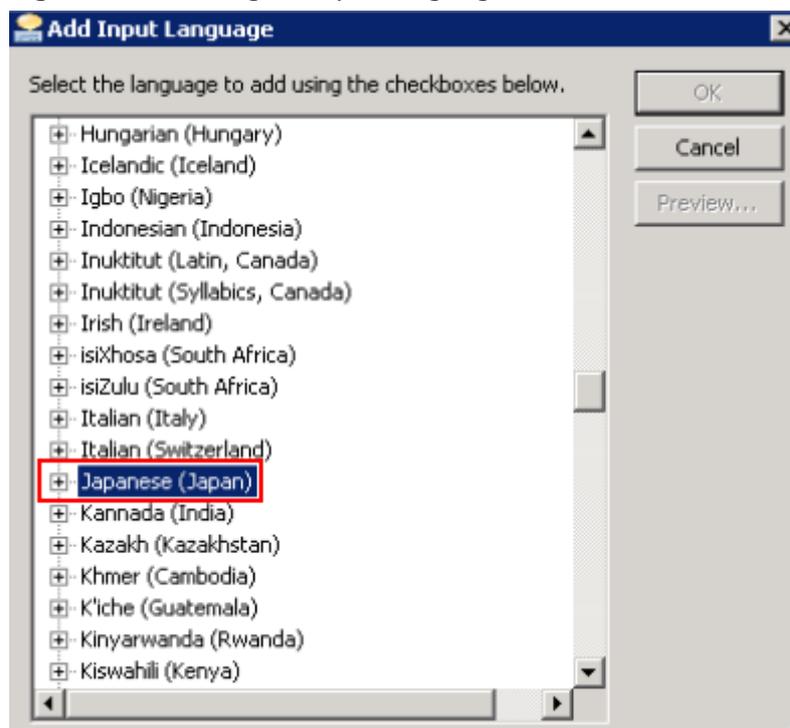
4. Right-click the language bar and click **Settings** to open the **Text Services and Input Languages** dialog box. On the **General** tab, click **Add**.

Figure 2-45 Text Services and Input Languages



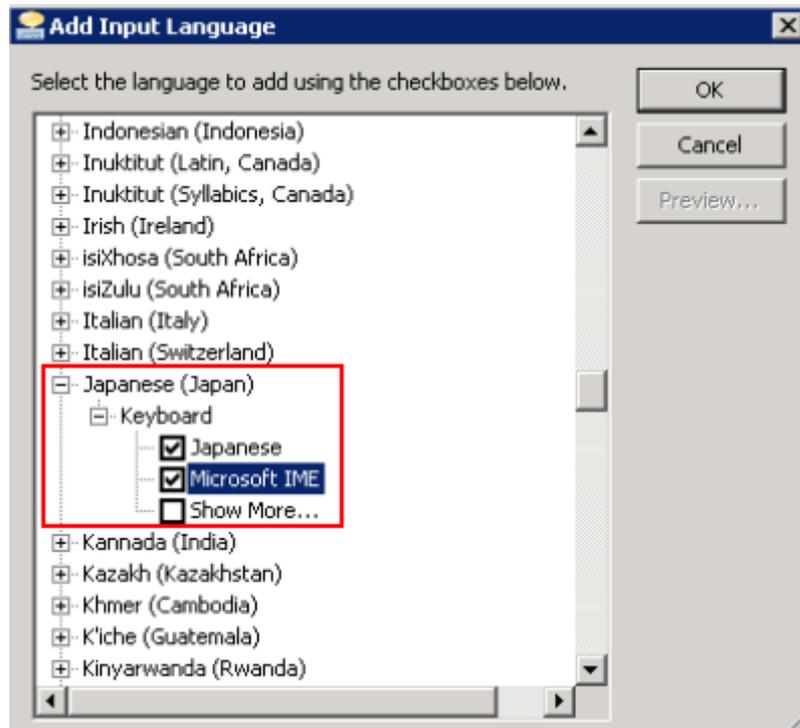
5. In the **Add Input Language** dialog box, find and double-click **Japanese (Japan)** and then **Keyboard**.

Figure 2-46 Adding an input language



6. Select **Japanese** and **Microsoft IME**. Click **OK** and then **Apply** to save the configurations.

Figure 2-47 Selecting Japanese



2.9 How Do I Share Files Between Windows ECSs?

Scenarios

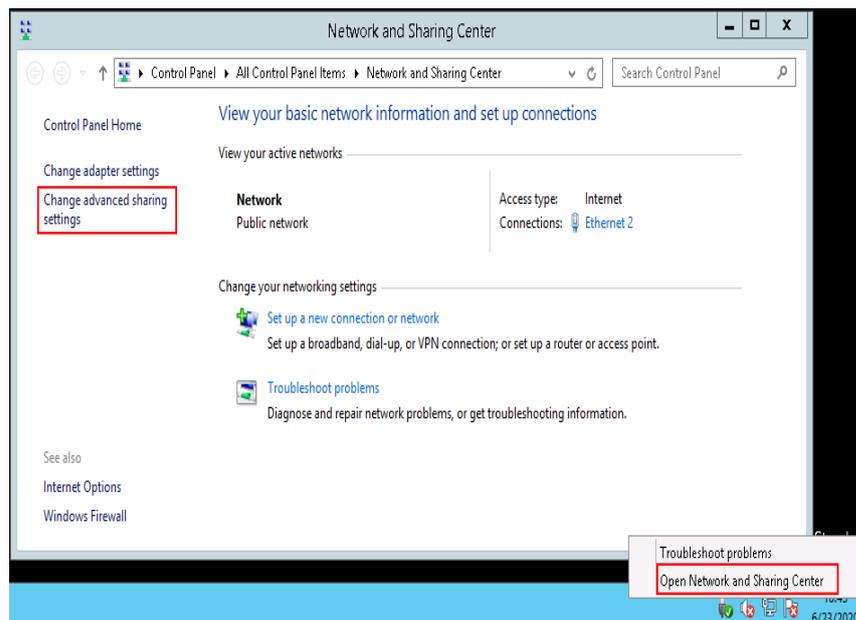
This section describes how to share files between Windows ECSs in the same subnet.

Constraints

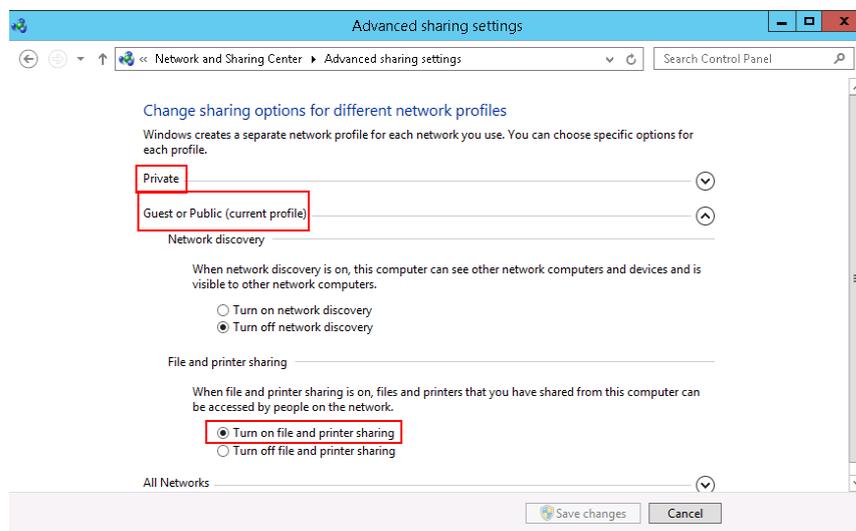
ECSs that share files are in the same subnet and can communicate with each other.

Procedure

1. Log in to the ECS. Right-click the **Network** icon in the lower right corner and choose **Open Network and Sharing Center** from the shortcut menu.

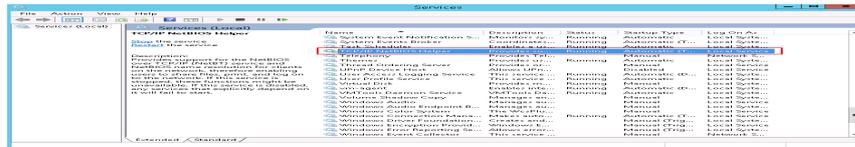
Figure 2-48 Open Network and Sharing Center

2. Click **Change advanced sharing settings**. On the displayed page, select **Turn on file and printer sharing** for the **Private** and **Guest or Public** areas and click **Save changes**.

Figure 2-49 Turning system icons on or off

3. In the **cmd** window, run the **services.msc** command, find and enable the **TCP/IP NetBIOS Helper** service.
If the service is disabled, right-click it and choose **Properties** from the shortcut menu. Set the **Startup Type** to **Automatic** or **Manual** and then enable the service.

Figure 2-50 Enabling TCP/IP NetBIOS Helper



2.10 How Do I Restore Data in the Event of a Startup Failure on a Windows ECS?

Symptom

If logging in to a Windows ECS failed, the system automatically starts the recovery option. However, an error is reported after the recovery option is selected, and the system recovery cannot continue.

Figure 2-51 Auto-starting system recovery mode



Possible Causes

Windows files have been damaged.

Solution

1. Log in to the management console and choose **Compute > Elastic Cloud Server**.
2. Detach the attached data disk.
 Click the name of the target ECS. On the ECS details page, click the **Disks** tab. Locate the target disk and click **Detach**.
3. Reinstall the OS of the ECS
 Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Image > Reinstall OS**

NOTICE

- Only stopped ECSs support OS reinstallation. If the ECS is not stopped, stop it before reinstalling the OS.
 - Reinstalling the OS will clear the system disk. Back up data before proceeding with reinstallation.
4. Attach the data disk to the ECS again and check whether you can log in to the ECS.

2.11 How Do I View Login Logs of a Windows ECS?

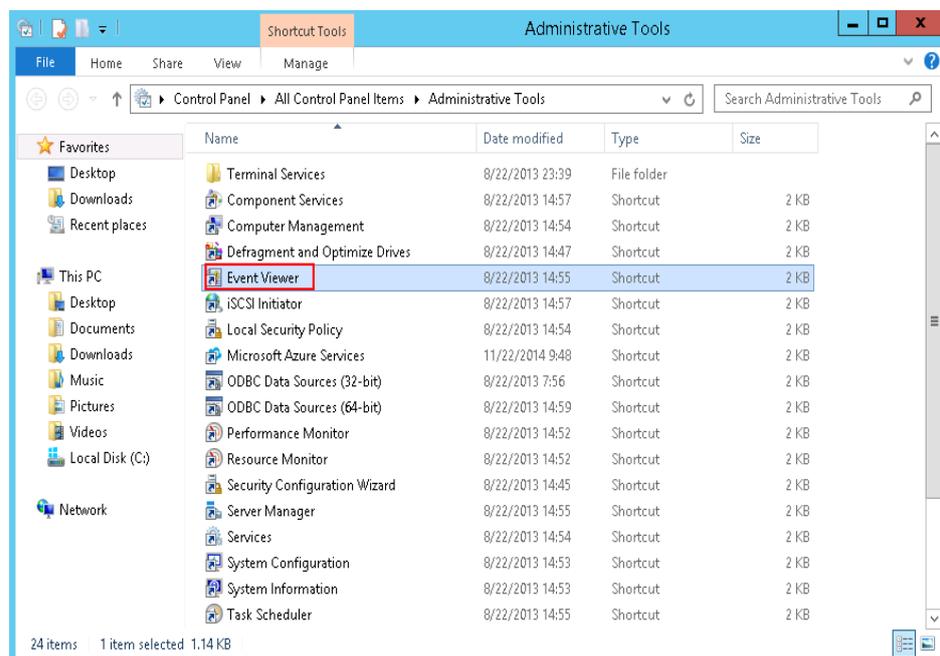
Scenarios

This section describes how to view the login logs of a Windows ECS.

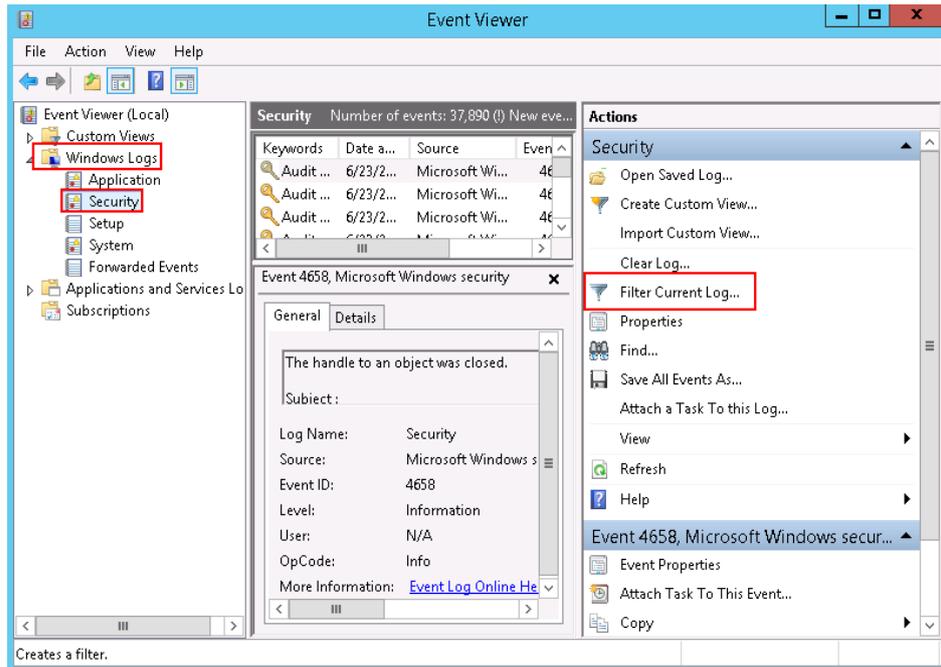
Procedure

The following operations are performed on an ECS running Windows Server 2012.

1. Log in to the Elastic Cloud Service ECS.
2. Choose **Start > Administrative Tools > Event Viewer**.



3. On the **Event Viewer** window, choose **Windows Logs > Security > Filter Current Logs**.

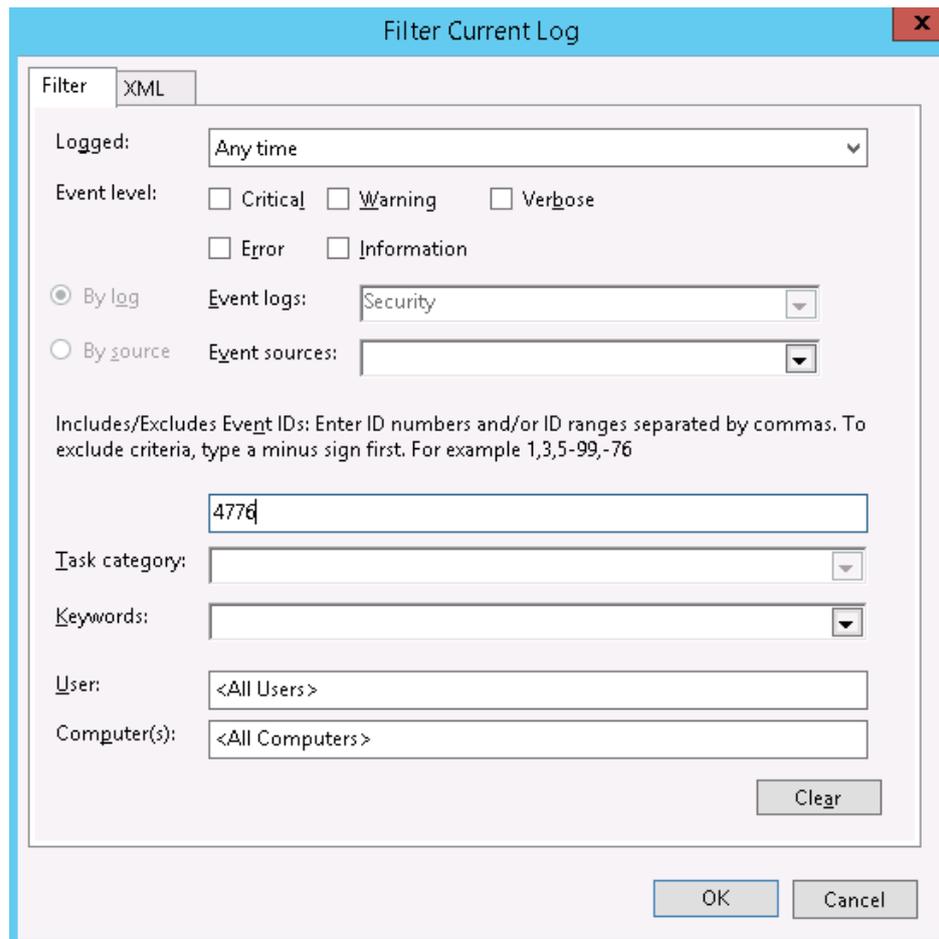


4. Filter events with ID 4776 to obtain remote login logs.

NOTE

You can also view login attempts by filtering events with ID 4624 or 4625.

- 4624: ID of successful login events
- 4625: ID of failed login events



5. Right-click an event and choose **Event Properties** from the shortcut menu. The logon account of the event is displayed.



2.12 What Can I Do If My Windows ECS Can Ping a Website but Cannot Access it?

Symptom

Your Windows ECS cannot access websites or applications after running for a long time.

You can remotely log in to the ECS and ping websites using the ECS, but it failed to access websites or applications.

Possible Causes

The dynamic ports of the Windows ECS are used up.

By default, a TCP connection will stay in the **TIME_WAIT** state for 4 minutes before the port used by this connection is released and available for other connections. If a Windows ECS runs for a long time, there may be too many connections established or in the **TIME_WAIT** state using up all available ports. In this case, any new attempt to establish a connection will fail.

You can run the following command in the CLI of the Windows ECS to check which connections are in the **TIME_WAIT** state:

```
netstat -an |find "TIME_WAIT" /c
```

The command output is as follows:

Figure 2-52 Checking the connections in the **TIME_WAIT** state

```
C:\Users\Administrator>netstat -an |find "TIME_WAIT" /c
15002

C:\Users\Administrator>_
```

Solution

1. Log in to the Windows ECS.
2. Run **cmd** as an administrator.
3. Run the following command to check dynamic ports:

```
netsh int ipv4 show dynamicport tcp
```

Figure 2-53 Checking dynamic ports

```
Protocol tcp Dynamic Port Range
-----
Start Port      : 49152
Number of Ports : 16384
```

4. Run the following commands to increase and then check dynamic ports:

```
netsh int ipv4 set dynamicport tcp start=1025 num=60000
netsh int ipv4 show dynamicport tcp
```

Figure 2-54 Setting dynamic ports

```
C:\Users\Administrator>netsh int ipv4 set dynamicport tcp start=1025 num=60000
Ok.

C:\Users\Administrator>netsh int ipv4 show dynamicport tcp

Protocol tcp Dynamic Port Range
-----
Start Port      : 1025
Number of Ports : 60000
```

5. Access external websites or applications again.

2.13 Why Can't I Open the Start Menu and Search Box on a Windows ECS?

Symptom

You can log in to the newly created Windows ECS, but the system does not respond when you click **Start menu** button and the search box.

Possible Causes

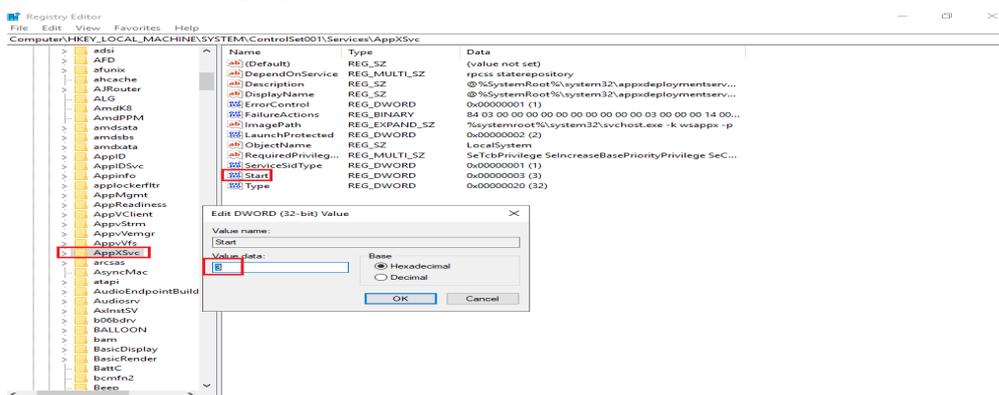
The AppX Deployment Service (AppXSVC) uses too many system resources, such as disks, and cannot be stopped. AppXSVC is a Windows application deployment service. It makes applications ready-to-use when you first log in to a computer and add an application.

Solution

To stop AppXSVC, perform the following steps:

1. Press **Win+R**. In the displayed dialog box, enter **regedit** and press **Enter** to open the registry editor.
2. Open the registry key value: **HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services\AppXSvc**.
3. Double-click **Start** in the right subitem and change **3** in **Value data** to **4**.

Figure 2-55 Changing value data



NOTE

To start AppXSVC, change the value data from **4** to **3**.

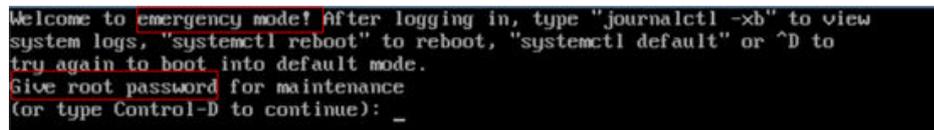
3 Linux ECS Issues

3.1 Why Is My Linux ECS Not Booting and Going Into Emergency Mode?

Symptom

Your Linux ECS enters the emergency mode during startup, and displays the message "Welcome to emergency mode", asking you to enter the password of user **root** for maintenance.

Figure 3-1 Emergency mode



```
Welcome to emergency mode! After logging in, type "journalctl -xb" to view
system logs, "systemctl reboot" to reboot, "systemctl default" or ^D to
try again to boot into default mode.
Give root password for maintenance
(or type Control-D to continue): _
```

Possible Causes

The emergency mode allows you to recover the system even if the system fails to enter the rescue mode. In emergency mode, the system installs only the root file system for data reading. It does not attempt to install any other local file systems or activate network interfaces.

The system enters the emergency mode when:

- An error occurred in the **/etc/fstab** file, leading to the failure in mounting the file system.
- An error occurred in the file system.

Constraints

The operations in this section are applicable to Linux. The operations involve recovering the file system, which may lead to data loss. You need to back up data before recovering the file system.

Solution

1. Enter the password of user **root** and press **Enter** to enter the recovery mode.
2. Run the following command to mount the root partition in read-write mode to modify the files in the root directory:
mount -o rw,remount /
3. Run the following command to try to mount all unmounted file systems:
mount -a
 - If the message "mount point does not exist" is displayed, the mount point is unavailable. In such a case, create the mount point.
 - If the message "no such device" is displayed, the file system device is unavailable. In such a case, comment out or delete the mount line.
 - If the message "an incorrect mount option was specified" is displayed, the mount parameters have been incorrectly set. In such a case, correct the parameter setting.
 - If no error occurs and the message "UNEXPECTED INCONSISTENCY;RUN fsck MANUALLY" is displayed, the file system is faulty. In such a case, go to [7](#).
4. Run the following command to open the **/etc/fstab** file and correct the error:
vi /etc/fstab

The **/etc/fstab** file contains the following parameters separated by space:
[file system] [dir] [type] [options] [dump] [fsck]

Table 3-1 /etc/fstab parameters

Parameter	Description
[file system]	<p>Specifies the partition or storage device to be mounted.</p> <p>You are advised to set file system in UUID format. To obtain the UUID of a device file system, run the blkid command.</p> <p>Format for reference:</p> <pre># <device> <dir> <type> <options> <dump> <fsck> UUID=b411dc99-f0a0-4c87-9e05-184977be8539 /home ext4 defaults 0 2</pre> <p>UUIDs are independent from the disk order. If the sequence of storage devices is changed manually or undergoes some random changes by some BIOSs, or the storage devices are removed and installed again, UUIDs are more effective in identifying the storage devices.</p>
[dir]	Specifies the mount point of a file system.

Parameter	Description
[type]	<p>Specifies the type of the file system to which a device or partition is mounted. The following file systems are supported: ext2, ext3, ext4, reiserfs, xfs, jfs, smbfs, iso9660, vfat, ntfs, swap, and auto.</p> <p>If type is set to auto, the mount command will speculate on the type of the file system that is used, which is useful for mobile devices, such as CD-ROM and DVD.</p>
[options]	<p>Specifies the parameters used for mounting. Some parameters are available only for specific file systems. For example, defaults indicates the default mounting parameters of a file system will be used. The default parameters of the ext4 file system are rw, suid, dev, exec, auto, nouser, and async.</p> <p>For more parameters, run the man mount command to view the man manual.</p>
[dump]	<p>Specifies whether file system data will be backed up.</p> <p>The value can be 0 or 1. 0 indicates that data will not be backed up, and 1 indicates that data will be backed up. If you have not installed dump, set the parameter to 0.</p>
[fsck]	<p>Specifies the sequence of checking file systems.</p> <p>The parameter value can be 0, 1, or 2. 0 indicates that the file systems will not be checked by fsck. 1 indicates the highest priority of the root directory to be checked by fsck, and 2 indicates the lower priority of other systems to be checked.</p>

5. After the modification is complete, run the following command to check the **fstab** file:

```
mount -a
```

6. Run the following command to restart the ECS:

```
reboot
```

7. Run the following command to check for file system errors:

```
dmesg |egrep "ext[2..4]|xfs" |grep -i error
```

NOTE

- If the error message "I/O error... inode" is displayed, the fault is caused by a file system error.
- If no error is found in the logs, the fault is generally caused by the corrupt superblock. The superblock is the header of the file system. It records the status, size, and idle disk blocks of the file system.
- If the superblock of a file system is corrupt, for example, data is written to the superblock of the file system by mistake, the system may fail to identify the file system. As a result, the system enters the emergency mode during startup. The ext2fs file system backs up the superblock and stores the backup at the block group boundary of the driver.

- Run the following command to unmount the directory where the file system error occurred:
umount *Mount point*
- Recover the corrupt file system.

NOTICE

Recovering the file system may lead to data loss. Back up data before the recovery.

- For the ext file system, run the following command to check whether the file system is faulty:

```
fsck -n /dev/vdb1
```

NOTE

If the message "The superblock could not be read or does not describe a correct ext2 filesystem" is displayed, go to step 10.

To recover the file system, run the following command:

```
fsck /dev/vdb1
```

- For the xfs file system, run the following command to check whether the file system is faulty:

```
xfs_repair -n /dev/vdb1
```

To recover the file system, run the following command:

```
xfs_repair /dev/vdb1
```

- (Optional) If the message "The superblock could not be read or does not describe a correct ext2 filesystem" is displayed, as shown in [Figure 3-2](#), the superblock is corrupt. In such a case, use an alternative superblock.

Figure 3-2 Corrupt superblock

```
[root@ecs ~]# fsck -n /dev/vda2
fsck from util-linux-ng 2.17.2
e2fsck 1.41.12 (17-May-2010)
fsck.ext2: No such file or directory while trying to open /dev/vda2

The superblock could not be read or does not describe a correct ext2
filesystem. If the device is valid and it really contains an ext2
filesystem (and not swap or ufs or something else), then the superblock
is corrupt, and you might try running e2fsck with an alternate superblock:
e2fsck -b 8193 <device>
```

Run the following command to replace the corrupt superblock with the superblock backup:

```
e2fsck -b 8193 Device name
```

As shown in [Figure 3-3](#), the corrupt superblock has been replaced.

Figure 3-3 Replacing the corrupt superblock

```
[root@ecs ~]# e2fsck -b 8193 /dev/xvda
e2fsck 1.41.12 (17-May-2010)
/dev/xvda is in use.
e2fsck: Cannot continue, aborting.
```

NOTE

-b 8193 indicates that the backup of superblock 8193 in the file system is used.

The location of the superblock backup varies depending on the block size of the file system. For a file system with a 1 KB block size, locate the backup at superblock 8193; for a 2 KB block size, locate the backup at superblock 16384; for a 4 KB block size, locate the backup at superblock 32768.

Device name is the disk name but not the partition name.

11. Run the following command to restart the ECS:

```
reboot
```

3.2 How Do I Fix a "Read-Only" Error When I Edit the /etc/fstab File?

Symptom

When you try to edit the `/etc/fstab` file, the message "readonly" is displayed.

Figure 3-4 Read-only `/etc/fstab`

```
UUID=1a1ce4de-e56a-4e1f-864d-31b7d9dfb547 / ext4 default
ts 1 1
UUID=25ec3bdb-ba24-4561-bcdc-802edf42b85f swap swap default
ts 0 0
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
UUID=25ec3bdb-ba24-4561-bcdc-801edf42b65f /mnt/sdc ext4 defaults
0 2

"/etc/fstab" [readonly] 15L, 800C 15,1 All
```

Solution

In the read-only mode, the file cannot be modified. Run the following command to obtain the read and write permissions on the file:

```
mount -o remount,rw /
```

After the preceding command is executed, edit the file again.

3.3 How Do I Change the Time Zone on ECSs Running CentOS or EulerOS?

Scenarios

This section describes how to change the time zone in ECSs running CentOS or EulerOS.

Constraints

Operations in this section involve modifying kernel parameters at runtime, which may cause kernel instability. Evaluate risks before modification and restart the system at a proper time after modification.

Possible Causes

The connection-tracking module of iptables uses some system memory to track connections in the table. The message "table full, dropping packet" indicates that the connection tracking table is full and no new entry can be created for new connections. As a result, the "dropping packet" problem occurs.

To resolve the problem, increase the number of connection tracking entries or filter out connections that do not need to be tracked.

Solutions

Solution 1: Adjust parameter values in the `nf_conntrack` module.

1. Increase the value of **`net.netfilter.nf_conntrack_max`**, which indicates the maximum number of connections supported by the `nf_conntrack` module.

For example:

If the memory is 64 GB and you want to change the **`net.netfilter.nf_conntrack_max`** value to **2097152**, run the following command to apply the setting immediately:

```
sysctl -w net.netfilter.nf_conntrack_max=2097152
```

Run the following command to ensure that the setting is still valid after the ECS is restarted:

```
echo "net.netfilter.nf_conntrack_max = 2097152" >> /etc/sysctl.conf
```

NOTE

- A larger **`net.netfilter.nf_conntrack_max`** value is not necessarily better. Set **`net.netfilter.nf_conntrack_max`** based on the memory size.
 - Calculate an appropriate value for **`nf_conntrack_max`** using the following formula (64-bit):
$$\text{CONNTRACK_MAX} = \text{RAMSIZE (inbytes)} / 16384 / 2$$

For example, if your ECS is running a 64-bit OS with 64 GB of memory, calculate an appropriate **`nf_conntrack_max`** value as follows:

$$\text{CONNTRACK_MAX} = 64 \times 1024 \times 1024 \times 1024 / 16384 / 2 = 2097152$$
2. Adjust the value of **`nf_conntrack_buckets`**, which indicates the size of the hash table in the `nf_conntrack` module.

If the number of entries in the `conntrack` table increases significantly (for example, by 4-fold), you should increase the size of the hash table for storing `conntrack` entries.

For CentOS 6 and later versions, calculate a new hash value using formula **`hashsize = nf_conntrack_max/4`**.

Run the following command to set a new hash value:

```
sysctl -w net.netfilter.nf_conntrack_buckets=524288
```

Run the following command to ensure that the setting is still valid after the ECS is restarted:

```
echo "net.netfilter.nf_conntrack_buckets=524288" >> /etc/sysctl.conf
```

NOTE

In earlier versions, you may not be able to adjust the **nf_conntrack_buckets** value. In this case, you can adjust the **hashsize** value.

Method:

In **/etc/modprobe.d/nf_conntrack.conf**, run the following command:

```
echo "options nf_conntrack hashsize=524288" >> /etc/modprobe.d/nf_conntrack.conf
```

Restart the ECS to apply the setting.

Solution 2: Use iptables to filter out connections that do not need to be tracked.

You can add the **-j NOTRACK** parameter to the iptables rule to filter out connections that do not need to be tracked. This method removes the records of the connections that do not need to be tracked from the hash table and prevents excessive connections from causing "nf_conntrack:table full, dropping packet" errors.

```
iptables -t raw -A PREROUTING -p udp -j NOTRACK  
iptables -t raw -A PREROUTING -p tcp --dport 22 -j NOTRACK
```

NOTE

The commands indicate that the UDP connections and TCP connections over port 22 are not tracked. They are provided for reference only.

3.5 How Do I Change the Default Boot Kernel in Ubuntu?

Scenarios

The operations described in this section only apply to ECSs running Ubuntu 16.04.

Constraints

This operation involves modifying the GRUB configuration file. Misoperations may cause the system to fail to start. Before the operation, back up the GRUB configuration file so that it can be restored in case of misoperations.

Procedure

1. Assume that the desired default kernel to boot from is the third one. Open the **/etc/default/grub** file and change the value of **GRUB_DEFAULT** to "1>2", as shown in [Figure 3-6](#).

Figure 3-6 Modifying GRUB_DEFAULT

```
GRUB_DEFAULT="1>2"  
GRUB_HIDDEN_TIMEOUT=  
GRUB_HIDDEN_TIMEOUT_QUIET=true  
GRUB_TIMEOUT=10  
GRUB_DISTRIBUTOR=`lsb_release -i -s 2> /dev/null || echo Debian`  
GRUB_CMDLINE_LINUX_DEFAULT=""  
GRUB_CMDLINE_LINUX="net.ifnames=0 nospectre_v2 nopti noibrs noibpb"
```

- **1** in **1>2** indicates the second entry of the main menu.

```
GNU GRUB version 2.02~beta2-36ubuntu3.20
Ubuntu
#Advanced options for Ubuntu
```

- **2** in **1>2** indicates the third entry of the submenu.

```
GNU GRUB version 2.02~beta2-36ubuntu3.22
Ubuntu, with Linux 4.4.0-151-generic
Ubuntu, with Linux 4.4.0-151-generic (recovery mode)
*Ubuntu, with Linux 4.4.0-142-generic
Ubuntu, with Linux 4.4.0-142-generic (recovery mode)
Ubuntu, with Linux 4.4.0-131-generic
Ubuntu, with Linux 4.4.0-131-generic (recovery mode)
```

- There is no space before and after the greater-than sign (>) in **1>2**.
- Use a set of quotation marks to enclose **1>2**.

2. Run the following command to rebuild a grub configuration file:
update-grub

3.6 How Do I Configure atop and kdump on Linux ECSs?

Scenarios

This section describes how to configure atop and kdump on Linux ECSs. The method for configuring atop varies with the OS version.

atop

- [Introduction to atop](#)
- [Preparing for atop Installation](#)
- [Configuring atop for CentOS 6](#)
- [Configuring atop for CentOS 7/8, AlmaLinux, and Rocky Linux](#)
- [Configuring atop for Ubuntu 16 and Debian 8](#)
- [Configuring atop for Ubuntu 18 and Debian 9](#)
- [Configuring atop for Ubuntu 20 and Debian 10](#)
- [Configuring atop for Ubuntu 22 and Ubuntu 24](#)
- [Configuring atop for Debian 11 and Debian 12](#)
- [Configuring atop for SUSE 12 and SUSE 15](#)
- [Configuring atop for Huawei Cloud EulerOS \(HCE\) 2.0](#)
- [Installing atop by Compiling the Source Code \(for CentOS Stream 8/9, openEuler, or EulerOS\)](#)
- [Analyzing atop Logs](#)
- [\(Optional\) Loading the netatop Kernel Module](#)

kdump

- [Precautions for Configuring kdump](#)
- [Introduction to kdump](#)
- [Configuring kdump](#)
- [Checking Whether kdump Configurations Have Taken Effect](#)

Introduction to atop

atop is a monitor for Linux that can report the activity of all processes and resource consumption by all processes at regular intervals. It shows system-level activity related to the CPU, memory, disks, and network layers for every process. It also logs system and process activities daily and saves the logs in disks for long-term analysis.

Preparing for atop Installation

An EIP is required and can be used to access the yum software source.

Constraints

The atop tool used for node resource monitoring consumes a small number of CPU cores and memory to collect system resources and uses some disk space to record logs. The log path is `/var/log/atop`.

Configuring atop for CentOS 6

1. Run the following command to install atop:

```
yum install -y atop
```

2. Run the following command to modify the configuration file of atop:

```
vi /etc/sysconfig/atop
```

Modify the following parameters, save the modification, and exit:

Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.

```
LOGINTERVAL=15
```

```
vi /etc/logrotate.d/atop
```

Modify the following parameters, save the modification, and exit:

You can change the value of **-mtime** to, for example, **3**. The default retention period of atop logs is **40** days.

```
postrotate
    /usr/bin/find /var/log/atop/ -maxdepth 1 -mount -name atop_\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]\[[0-9\]\]* -mtime +3 -exec /bin/rm {} \;
endscript
```

3. Run the following command to start atop:

```
service atop start
```

4. Run the following command to check the status of atop. **is running** indicates that atop is running properly.

```
service atop status
```

```
atop (pid 3170) is running
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:

service atop stop**Configuring atop for CentOS 7/8, AlmaLinux, and Rocky Linux**

1. Run the following command to install atop:
yum install -y atop
2. Run the following command to modify the configuration file of atop:
vi /etc/sysconfig/atop
Modify the following parameters, save the modification, and exit:
 - Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
 - Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15  
LOGGENERATIONS=3
```

3. Run the following command to start atop:
systemctl enable --now atop atopacct atop-rotate.timer
4. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the service is running properly.

systemctl status atop atopacct atop-rotate.timer

```
atop.service - Atop advanced performance monitor  
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atopacct.service - Atop process accounting daemon  
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atop-rotate.timer - Daily atop restart  
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)  
Active: active (waiting)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:
systemctl disable atop atopacct atop-rotate.timer
systemctl stop atop atopacct atop-rotate.timer

Configuring atop for Ubuntu 16 and Debian 8

1. Run the following command to install atop:
apt-get install -y atop
2. Run the following command to modify the configuration file of atop:
vi /etc/default/atop
Modify the following parameters, save the modification, and exit:
 - Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
 - The default retention period of atop logs is **28** days and cannot be modified.

```
LOGINTERVAL=15
```

3. Run the following command to start atop:

systemctl start atop

4. Run the following command to check the status of atop. **active (running)** indicates that atop is running properly.

systemctl status atop

```
atop.service - Atop advanced performance monitor
Loaded: loaded (/etc/init.d/atop; bad; vendor preset: disabled)
Active: active (running)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:

systemctl stop atop

Configuring atop for Ubuntu 18 and Debian 9

1. Run the following command to install atop:

apt-get install -y atop

2. Run the following command to modify the configuration file of atop:

vi /usr/share/atop/atop.daily

Modify the following parameters, save the modification, and exit:

- Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
- You can change the value of **-mtime** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15
.....
( (sleep 3; find $LOGPATH -name 'atop_*' -mtime +3 -exec rm {} \;)& )
```

3. Restart the atop service (started by default) to apply the configuration:

systemctl restart atop atopacct

4. Run the following command to check the status of atop. **active (running)** indicates that atop is running properly.

systemctl status atop atopacct

```
atop.service - Atop advanced performance monitor
Loaded: loaded (/etc/init.d/atop; enable; vendor preset: enabled)
Active: active (running)
```

```
atopacct.service - Atop process accounting daemon
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)
Active: active (running)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:

systemctl disable atop atopacct**systemctl stop atop atopacct**

Configuring atop for Ubuntu 20 and Debian 10

1. Run the following command to install atop:

apt-get install -y atop

2. Run the following command to modify the configuration file of atop:

vi /etc/default/atop

Modify the following parameters, save the modification, and exit:

- Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
- Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15
LOGGENERATIONS=3
```

3. Restart the atop service (started by default) to apply the configuration:
systemctl restart atop atopacct
4. Run the following command to check the status of atop. **active (running)** indicates that atop is running properly.

systemctl status atop atopacct

```
atop.service - Atop advanced performance monitor
Loaded: loaded (/etc/init.d/atop; enable; vendor preset: enabled)
Active: active (running)
```

```
atopacct.service - Atop process accounting daemon
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)
Active: active (running)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:

systemctl disable atop atopacct
systemctl stop atop atopacct

Configuring atop for Ubuntu 22 and Ubuntu 24

1. Run the following command to install atop:
apt-get install -y atop
2. Run the following command to modify the configuration file of atop:
vi /etc/default/atop
Modify the following parameters, save the modification, and exit:
 - Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
 - Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15
LOGGENERATIONS=3
```

3. Restart the atop service (started by default) to apply the configuration:
systemctl restart atop atopacct atop-rotate.timer
4. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the service is running properly.

systemctl status atop atopacct atop-rotate.timer

```
atop.service - Atop advanced performance monitor
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)
Active: active (running)
```

```
atopacct.service - Atop process accounting daemon
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)
Active: active (running)
```

```
atop-rotate.timer - Daily atop restart
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)
Active: active (waiting)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:
systemctl disable atop atopacct atop-rotate.timer
systemctl stop atop atopacct atop-rotate.timer

Configuring atop for Debian 11 and Debian 12

1. Run the following command to install atop:
apt-get install -y atop
2. Run the following command to modify the configuration file of atop:
vi /etc/default/atop
Modify the following parameters, save the modification, and exit:
 - Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
 - Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15  
LOGGENERATIONS=3
```

3. Restart the atop service (started by default) to apply the configuration:
systemctl restart atop atopacct atop-rotate.timer
4. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the service is running properly.

```
systemctl status atop atopacct atop-rotate.timer
```

```
atop.service - Atop advanced performance monitor  
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atopacct.service - Atop process accounting daemon  
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atop-rotate.timer - Daily atop restart  
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)  
Active: active (waiting)
```

5. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:
systemctl disable atop atopacct atop-rotate.timer
systemctl stop atop atopacct atop-rotate.timer

Configuring atop for SUSE 12 and SUSE 15

1. Run the following command to download the atop source package:
wget https://www.atoptool.nl/download/atop-2.6.0-1.src.rpm
2. Run the following command to install the atop source package:
rpm -ivh atop-2.6.0-1.src.rpm
3. Run the following command to install atop dependencies.
zypper -n install rpm-build ncurses-devel zlib-devel
4. Run the following commands to compile atop:
cd /usr/src/packages/SPECS

rpmbuild -bb atop-2.6.0.spec

5. Run the following commands to install atop:

```
cd /usr/src/packages/RPMS/x86_64
```

```
rpm -ivh atop-2.6.0-1.x86_64.rpm
```

6. Run the following command to modify the configuration file of atop:

```
vi /etc/default/atop
```

Modify the following parameters, save the modification, and exit:

- Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.
- Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGINTERVAL=15  
LOGGENERATIONS=3
```

7. Restart the atop service (started by default) to apply the configuration:

```
systemctl restart atop atopacct atop-rotate.timer
```

8. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the service is running properly.

```
systemctl status atop atopacct atop-rotate.timer
```

```
atop.service - Atop advanced performance monitor  
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)  
Active: active (running)
```

```
atopacct.service - Atop process accounting daemon  
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)  
Active: active (running)
```

```
atop-rotate.timer - Daily atop restart  
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)  
Active: active (waiting)
```

9. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:

```
systemctl disable atop atopacct atop-rotate.timer
```

```
systemctl stop atop atopacct atop-rotate.timer
```

Configuring atop for Huawei Cloud EulerOS (HCE) 2.0

1. Run the following command to install atop:

```
yum install -y atop
```

2. Set the sampling period.

The sampling period has been preset in the HCE and does not need to be adjusted. By default, the sampling period is 10 seconds, and logs are stored for three days.

3. Run the following command to start atop:

```
systemctl enable --now atop atopacct atop-rotate.timer
```

4. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the services are running properly.

```
systemctl status atop atopacct atop-rotate.timer
```

```
atop.service - Atop advanced performance monitor  
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)
```

Active: active (running)

atopacct.service - Atop process accounting daemon
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)
Active: active (running)

atop-rotate.timer - Daily atop restart
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)
Active: active (waiting)

5. Run the following commands to stop atop after troubleshooting to release the system and disk resources occupied by atop:

```
systemctl disable atop atopacct atop-rotate.timer  
systemctl stop atop atopacct atop-rotate.timer
```

Installing atop by Compiling the Source Code (for CentOS Stream 8/9, openEuler, or EulerOS)

1. Download the atop source package.

```
wget https://www.atoptool.nl/download/atop-2.6.0.tar.gz
```

2. Decompress the atop source package.

```
tar -zxvf atop-2.6.0.tar.gz
```

3. Query the systemctl version.

```
systemctl --version
```

If the version is 220 or later, go to the next step.

Otherwise, delete parameter **--now** from the Makefile of atop.

```
vi atop-2.6.0/Makefile
```

Delete parameter **--now** following the **systemctl** command.

```
then /bin/systemctl disable atop 2> /dev/null; \  
/bin/systemctl disable atopacct 2> /dev/null; \  
/bin/systemctl daemon-reload; \  
/bin/systemctl enable atopacct; \  
/bin/systemctl enable atop; \  
/bin/systemctl enable atop-rotate.timer; \  
fi
```

4. Install atop dependencies.

- Installing command for SUSE 12 or SUSE 15

```
zypper -n install make gcc zlib-devel ncurses-devel
```

- Installing command for EulerOS or Fedora

```
yum install make gcc zlib-devel ncurses-devel -y
```

- Installing command for Debian 9, Debian 10, or Ubuntu

```
apt install make gcc zlib1g-dev libncurses5-dev libncursesw5-dev -y
```

5. Compile and install atop.

```
cd atop-2.6.0
```

```
make systemdinstall
```

6. Run the following command to modify the configuration file of atop:

```
vi /etc/default/atop
```

Add the following parameters, save the modification, and exit:

- Change the value of **LOGINTERVAL** to, for example, **15**. The default value of **LOGINTERVAL** is **600**, in seconds.

- Change the value of **LOGGENERATIONS** to, for example, **3**. The default retention period of atop logs is **28** days.

```
LOGOPTS=""  
LOGINTERVAL=15  
LOGGENERATIONS=3  
LOGPATH=/var/log/atop
```

7. Restart the atop service (started by default) to apply the configuration:
systemctl restart atop atopacct atop-rotate.timer
8. Check whether atop is started. If **atop atopacct** is **active (running)**, or **atop-rotate.timer** is **active (waiting)**, the service is running properly.

systemctl status atop atopacct atop-rotate.timer

```
atop.service - Atop advanced performance monitor  
Loaded: loaded (/usr/lib/systemd/system/atop.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atopacct.service - Atop process accounting daemon  
Loaded: loaded (/usr/lib/systemd/system/atopacct.service; enabled; vendor preset: enabled)  
Active: active (running)  
  
atop-rotate.timer - Daily atop restart  
Loaded: loaded (/usr/lib/systemd/system/atop-rotate.timer; enabled; vendor preset: enabled)  
Active: active (waiting)
```

9. Run the following command to stop atop after troubleshooting to release the system and disk resources occupied by atop:
systemctl disable atop atopacct atop-rotate.timer
systemctl stop atop atopacct atop-rotate.timer

Analyzing atop Logs

After startup, atop stores collection records in **/var/log/atop**.

Run the following command to check the log file:

```
atop -r /var/log/atop/atop_2024XXXX
```

- **Common atop commands**

After opening the log file, you can use the following commands to sort data:

- **c**: used to sort processes by CPU usage in descending order.
- **m**: used to sort processes by memory usage in descending order.
- **d**: used to sort processes by disk usage in descending order.
- **a**: used to sort processes by the overall resource usage in descending order.
- **n**: used to sort processes by network usage in descending order.
- **t**: used to go to the next monitoring collection point.
- **T**: used to go to the previous monitoring collection point.
- **b**: used to specify a time point in the format of YYYYMMDDhhmm.

- **System resource monitoring fields**

The following figure shows some monitoring fields and values. The values vary according to the sampling period and atop version. The figure is for reference only.

Figure 3-7 System resource monitoring fields

```

ATOP - ecs-atop          2024/03/12  10:31:30          -----          1h12m49s elapsed
PRC | sys  5.06s | user 10.75s | #proc  98 | #trun  1 | #tslpi 112 | #tslpu  0 | #zombie  0 | #exit  0 |
CPU | sys  0% | user  1% | irq  0% | idle 198% | wait  0% | ipc notavail | curf 2.80GHz | curscal  2% |
cpu | sys  0% | user  1% | irq  0% | idle 99% | cpu000 w 0% | ipc notavail | curf 2.80GHz | curscal  2% |
cpu | sys  0% | user  1% | irq  0% | idle 99% | cpu001 w 0% | ipc notavail | curf 2.80GHz | curscal  2% |
CPL | avg1  0.00 | avg5  0.00 | avg15  0.00 |          | csu 1527759 | intr 924841 |          | numcpu  2 |
MEM | tot  3.2G | free  2.0G | cache 955.4M | buff  27.6M | slab  81.0M |          | vmhal  0.0M | zfacr  0.0M | hptot  0.0M |
SWP | tot  0.0M | free  0.0M |          |          |          |          |          |          |          |          |
DSK |          | busy  0% | read 11696 | write 8132 | KiB/w  64 | MBr/s  0.1 | MBr/s  0.1 |          |          |          |
NET | transport | tcp  11926 | tcpo  8361 | udp  231 | udpo  231 | tcpao  180 |         |         |         |         |
NET | network | ipi  12737 | ipo  9183 | infw  0 | deliv 12735 |          |         |         |         |         |
NET | eth0 ---- | pcki 69567 | pcko  9138 | sp  0 | Mbps | si 183 | Kbps | so  1 | Kbps | lerr  0 | erro  0 |
NET | lo ---- | pcki  10 | pcko  10 | sp  0 | Mbps | si  0 | Kbps | lerr  0 | erro  0 |
*** system and process activity since boot ***
PID  SYSCPU  USRCPU  RDELAY  UGROW  RGROW  RDISK  WRDISK  RUID  EUID  ST  EXC  THR  S  CPUNR  CPU  CMD  1/3
7487  2.08s   3.96s  63.09s  1.2G  27148K  5588K  288.3M  root  root  N-  -  26  S  1  0%  hostguard
8535  0.65s   1.24s  0.00s   8932K  8336K   0K      1768K  root  root  N-  -  1  S  1  0%  java
2889  0.46s   1.21s  2.54s   3.3G  73876K  39436K   0K      root  root  N-  -  15  S  0  0%  atop
1  0.59s   0.48s  0.25s   189.9M 11748K 153.7M  217.9M  root  root  N-  -  1  S  1  0%  systemd
1083  0.01s   1.27s  1.47s  16972K  5988K  2332K   0K      root  root  N-  -  1  S  0  0%  rsysd
7394  0.13s   0.57s  0.19s  136.4M 11212K   0K      136K  root  root  N-  -  3  S  1  0%  hostwatch
477  0.01s   0.47s  2.58s  38044K  9872K  9544K   0K      root  root  N-  -  1  S  1  0%  systemd-udev
1872  0.28s   0.08s  0.61s  217.0M 2788K   36K     32K  root  root  N-  -  2  S  1  0%  wrapper
461  0.05s   0.27s  0.34s  53256K 17388K 4324K 15196K  root  root  N-  -  1  S  1  0%  systemd-journ
1166  0.01s   0.27s  0.49s  18484K 8568K  272K   0K      root  root  N-  -  1  S  1  0%  systemd-logind
1184  0.04s   0.22s  0.48s  453.0M 25536K 3776K  12K     root  root  N-  -  4  S  1  0%  tuned
    
```

Description of major fields is as follows:

- **ATOP** row: Specifies the host name and information sampling date and time.
- **PRC** row: Specifies the running status of a process.
- **#sys** and **user**: Specifies how long the CPU is occupied when the system is running in kernel mode and user mode.
- **#proc**: Specifies the total number of processes.
- **#zombie**: Specifies the number of zombie processes.
- **#exit**: Specifies the number of processes that exited during the sampling period.
- **CPU** row: Specifies the overall CPU usage (multi-core CPU as a whole CPU). The sum of the values in the CPU row is N x 100%. **N** indicates the number of CPUs.
- **#sys** and **user**: Specifies the percentage of how long the CPU is occupied when the system is running in kernel mode and user mode.
- **#irq**: Specifies the percentage of time when CPU is servicing interrupts.
- **#idle**: Specifies the percentage of time when CPU is idle.
- **#wait**: Specifies the percentage of time when CPU is idle due to I/O wait.
- **CPL** row: Specifies CPU load.
- **#avg1**, **avg5** and **avg15**: Specifies the average number of running processes in the past 1, 5, and 15 minutes, respectively.
- **#csu**: Specifies the number of context exchanges.
- **#intr**: Specifies the number of interruptions.
- **MEM** row: Specifies the memory usage.
- **#tot**: Specifies the physical memory size.
- **#free**: Specifies the size of available physical memory.
- **#cache**: Specifies the memory size used for page cache.
- **#buff**: Specifies the memory size used for file cache.
- **#slab**: Specifies the memory size occupied by the system kernel.
- **SWP** row: Specifies the usage of swap space.
- **#tot**: Specifies the total swap space.

- **#free**: Specifies the size of available swap space.
- **DSK** row: Specifies the disk usage. Each disk device corresponds to a column. If there is an **sdb** device, a **DSK** row should be added.
- **#sda**: Specifies the disk device identifier.
- **#busy**: Specifies the percentage of time when the disk is busy.
- **#read** and **write**: Specifies the number of read and write requests.
- **NET** row: Displays the network status, covering the transport layer (TCP and UDP), IP layer, and active network ports.
- **#xxxxxi**: Specifies the number of packets received by each layer or active network port.
- **#xxxxxo**: Specifies the number of packets sent by each layer or active network port.

(Optional) Loading the netatop Kernel Module

The netatop kernel module can monitor the statistics of TCP and UDP data packets sent and received by each process or thread. In atop, netatop is not installed by default. If you need to monitor the network usage, install it by referring to the description below.

Huawei Cloud EulerOS 2.0 is used as an example to describe how to load the netatop kernel module for atop.

To load the netatop kernel module on other Linux operating systems, see [Module netatop](#) in the official atop help documentation.

1. Run the following command to install the kernel development package and the software environment required for compilation on the kernel:
sudo yum install -y kernel-devel dkms elfutils-libelf-devel
2. Run the following command to download the source code of the latest netatop to the specified directory:
cd /usr/src/ && sudo wget https://www.atoptool.nl/download/netatop-3.2.2.tar.gz --no-check-certificate
3. Run the following command to decompress the source code and go to the source code directory:
sudo tar -zxvf netatop-3.2.2.tar.gz && cd netatop-3.2.2
4. Run the following command to build and install the module and daemon based on the source code:
sudo make && sudo make install
5. Run the following command to start netatop:
systemctl start netatop

Precautions for Configuring kdump

The method for configuring kdump described in this section applies to KVM Linux ECSs running EulerOS or CentOS 7.x. For details, see [Documentation for kdump](#).

Introduction to kdump

Kdump is a tool used to dump runtime memory when the system crashes. Once the system crashes, the kernel can no longer function properly. At this point,

Kdump boots another kernel to capture and save a memory dump. This kernel collects all runtime states and data and stores them in a dump core file for troubleshooting and debugging.

Configuring kdump

1. Run the following command to check whether kexec-tools is installed:

```
rpm -q kexec-tools
```

If it is not installed, run the following command to install it:

```
yum install -y kexec-tools
```

2. Run the following command to enable kdump to run at system startup:

```
systemctl enable kdump
```

3. Configure the **crashkernel** parameter to reserve the memory for the capture kernel.

Run the following command to check whether the parameter has been configured:

```
grep crashkernel /proc/cmdline
```

If the command output is displayed, this parameter has been configured. If no command output is displayed, you need to configure **crashkernel**.

Edit the **/etc/default/grub** file to configure **crashkernel** as follows:

```
GRUB_TIMEOUT=5
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=rhel00/root rd.lvm.lv=rhel00/swap
rhgb quiet"
GRUB_DISABLE_RECOVERY="true"
```

Locate parameter **GRUB_CMDLINE_LINUX** and add **crashkernel=auto** after it.

4. Run the following command for the configuration to take effect:

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

5. Open the **/etc/kdump.conf** file, locate parameter **path**, and add **/var/crash** after it.

```
path /var/crash
```

By default, the file is saved in the **/var/crash** directory.

You can save the file to another directory, for example, **/home/kdump**. Then add **/home/kdump** after parameter **path**:

```
path /home/kdump
```

NOTE

There must be enough space in the specified path for storing the vmcore file. It is recommended that the available space be greater than or equal to the RAM size. You can also store the vmcore file on a shared device such as SAN or NFS.

6. Set the vmcore dump level.

Add the following content to file **/etc/kdump.conf**. If the content already exists, skip this step.

```
core_collector makedumpfile -d 31 -c
```

In the preceding command:

-c indicates compressing the vmcore file.

-d indicates leaving out irrelevant data. Generally, the value following **-d** is **31**, which is calculated based on the following values. You can adjust the value if needed.

```
zero pages = 1
cache pages = 2
cache private = 4
user pages = 8
free pages = 16
```

7. Run the following command to restart the system for the configurations to take effect:

reboot

Checking Whether kdump Configurations Have Taken Effect

1. Run the following command and check whether **crashkernel=auto** is displayed:

cat /proc/cmdline |grep crashkernel

```
BOOT_IMAGE=/boot/vmlinuz-3.10.0-514.44.5.10.h142.x86_64 root=UUID=6407d6ac-c761-43cc-a9dd-1383de3fc995 ro crash_kexec_post_notifiers softlockup_panic=1 panic=3
reserve_kbox_mem=16M nmi_watchdog=1 rd.shell=0 fsck.mode=auto fsck.repair=yes net.ifnames=0
spectre_v2=off nopti noibrs noibpb crashkernel=auto LANG=en_US.UTF-8
```

2. Run the following command and check whether the information in the output is correct:

grep core_collector /etc/kdump.conf |grep -v ^"#"

```
core_collector makedumpfile -l --message-level 1 -d 31
```

3. Run the following command and check whether the information in the output is correct:

grep path /etc/kdump.conf |grep -v ^"#"

```
path /var/crash
```

4. Run the following command and check whether the value of **Active** in the output is **active (exited)**:

systemctl status kdump

```
● kdump.service - Crash recovery kernel arming
Loaded: loaded (/usr/lib/systemd/system/kdump.service; enabled; vendor preset: enabled)
Active: active (exited) since Tue 2019-04-09 19:30:24 CST; 8min ago
Process: 495 ExecStart=/usr/bin/kdumpctl start (code=exited, status=0/SUCCESS)
Main PID: 495 (code=exited, status=0/SUCCESS)
CGroup: /system.slice/system-hostos.slice/kdump.service
```

5. Run the following test command:

echo c > /proc/sysrq-trigger

This command will trigger kdump. When triggered, kdump will boot another kernel and store the generated **vmcore** file to the location specified by **path**.

6. Run the following command to check whether the **vmcore** file has been generated in the specified path, for example, **/var/crash/**:

ll /var/crash/

3.7 Why Is the OS Version of My ECS Not the One in the Image I Selected During ECS Creation?

Symptom

You run the following command on your ECS to query its OS version:

```
cat /etc/redhat-release
```

Assume that the command output shows that the current version is CentOS 7.6. However, the OS version of the image that you used to create the ECS is CentOS 7.2 or another version earlier than CentOS 7.6.

NOTE

The operations described in this section only apply to ECSs running CentOS or EulerOS.

Possible Causes

1. System patches have been installed. Run the following command to view the operation records in YUM:

yum history

Information similar to the following is displayed:

```
Loaded plugins: fastestmirror
ID | Login user | Date and time | Action(s) | Altered
-----|-----|-----|-----|-----
8 | root <root> | 2019-08-23 10:00 | I, U | 40 EE
7 | root <root> | 2019-06-24 10:12 | I, U | 55
6 | root <root> | 2019-02-22 11:10 | I, O, U | 280 EE
5 | root <root> | 2019-02-22 11:09 | I, U | 6
4 | root <root> | 2019-02-22 11:08 | Install | 1
3 | root <root> | 2019-02-22 11:08 | Install | 1
2 | root <root> | 2019-02-22 11:08 | Install | 1
1 | System <unset> | 2019-02-22 10:49 | Install | 352
history list
```

The action performed on August 23 contains operations I (Install) and U (Update) with operation ID 8.

2. Run the following command to view details about the operations:

yum history info 8

Information similar to the following is displayed:

```
Loaded plugins: fastestmirror
Transaction ID : 8
Begin time    : Fri Aug 23 10:00:13 2019
Begin rpmdb   : 384:1c8e3df918de17e245b0dd7195f06f89656c5f27
End time     :      10:02:22 2019 (129 seconds)
End rpmdb    : 386:9a3172e7946f31d43c1268b6e1f2428125b3dfc5
User        : root <root>
Return-Code  : Success
Command Line : update -y
```

The preceding command output indicates that user **root** has performed a YUM update operation successfully. This upgrades the OS to the latest version.

3.8 How Do I Enable My ECS to Boot From the Second Kernel If It Fails to Boot from the First Kernel?

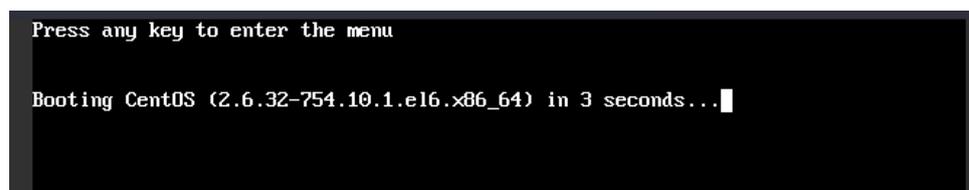
Scenarios

The operations described in this section apply to ECSs running CentOS or EulerOS with at least two kernels installed.

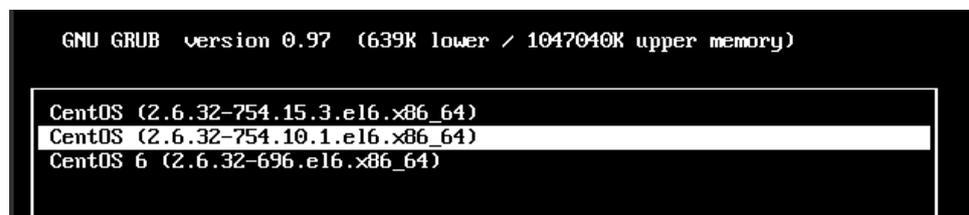
CentOS 6

1. Log in to the management console, restart the ECS, and click **Remote Login**.
2. When the **Booting CentOS** screen is displayed, press any key to access the kernel selection screen.

Figure 3-8 Booting CentOS



3. On the kernel selection screen, press the down arrow key to move the cursor to the second row, and press **Enter** to boot the system from the second kernel.



4. After the system is booted, run the following command to set the second kernel as the default boot kernel:

```
sed 's/default=0/default=1/' /boot/grub/grub.conf -i
```

NOTE

The default value of parameter **default** is **0**. In the kernel menu created by the **/boot/grub/grub.conf** file, entries count from 0. It means that 0 is for the first entry, 1 for the second, and 2 for the third.

CentOS 7 or EulerOS

1. Log in to the management console, restart the ECS, and click **Remote Login**.
2. On the kernel selection screen, press the down arrow key to move the cursor to the second row, and press **Enter** to boot the system from the second kernel.

```
CentOS Linux (3.10.0-957.21.3.el7.x86_64) 7 (Core)
CentOS Linux (3.10.0-957.5.1.el7.x86_64) 7 (Core)
CentOS Linux (3.10.0-327.el7.x86_64) 7 (Core)
CentOS Linux (0-rescue-136c058e80464a53bb585a2817774c02) 7 (Core)
```

3. After the system is booted, run the following command to set the second kernel as the default boot kernel:

```
grub2-set-default 1
```

3.9 How Can I Make /etc/rc.local Run at Startup in CentOS 7?

Symptom

The `/etc/rc.local` startup script does not run at startup of ECSs running CentOS 7 or EulerOS.

This section uses CentOS 7 as an example.

Possible Causes

The possible causes are as follows:

- The `/etc/rc.d/rc.local` file is not executable in CentOS 7. For detailed measures, see [Solution 1](#).

NOTE

`/etc/rc.local` is a soft link to `/etc/rc.d/rc.local`.

```
root@uwx-cent7-test ~]# ll /etc/rc.d/rc.local
-rw-r--r-- 1 root root 489 Mar 14 11:28 /etc/rc.d/rc.local
root@uwx-cent7-test ~]# ll /etc/rc.local
lrwxrwxrwx. 1 root root 13 Nov  7 14:30 /etc/rc.local -> rc.d/rc.local
```

- The `route` command in `/etc/rc.local` becomes invalid after CentOS 7 restart. This is because the network service is not started when the kernel reads `/etc/rc.local` to add a route during CentOS 7 startup. For detailed measures, see [Solution 2](#).

Solution 1

The solution for non-executable `/etc/rc.d/rc.local` is as follows:

NOTE

In CentOS 7, the `/etc/rc.d/rc.local` file is not executable by default.

1. Run the following command to check whether file `/etc/rc.d/rc.local` is executable:

```
ls -l /etc/rc.d/rc.local
```

```
-rw-r--r-- 1 root root 473 Sep 14 02:19 /etc/rc.d/rc.local
```

As shown in the command output, the file is not executable.

2. Run the following command to make `/etc/rc.d/rc.local` executable:
chmod +x /etc/rc.d/rc.local

Solution 2

The solutions for invalid route command in the `/etc/rc.local` after CentOS 7 reboot are as follows:

Method 1: Write the route into a static route configuration file. The route will not be lost even if the NIC is restarted.

1. Go to the `/etc/sysconfig/network-scripts/` directory and check if there is the file `route-<interface>`, where `<interface>` is the name of the interface related to the route.

If there is, add the following content to the file.

If not, create one and add the following content to the file:

```
<network/prefix> via <gateway>
```

where `<network/prefix>` is a remote network with prefix, and `<gateway>` is the next-hop IP address.

For example, add a route to point to 10.20.30.0/24 from 192.168.100.10 to enable eth0 at system startup.

```
cat /etc/sysconfig/network-scripts/route-eth0
```

```
10.20.30.0/24 via 192.168.100.10
```

2. Run the following command for the modifications to take effect:
ifup eth0

Method 2: The `/etc/rc.d/rc.local` file is controlled by the `rc-local` service. You can configure the `rc-local` service to start after `network-online.target` is started.

1. Run the following command to view `rc-local` configuration file located in `/usr/lib/systemd/system/rc-local.service`:

```
cat /usr/lib/systemd/system/rc-local.service |grep -v "^#"
```

Check if there are parameters **Requires** and **After** in the **[Unit]** area.

If there are, change their values to `network-online.target`.

If not, add these two parameters and set their values to `network-online.target`.

```
[Unit]
Description=/etc/rc.d/rc.local Compatibility
ConditionFileIsExecutable=/etc/rc.d/rc.local
Requires=network-online.target
After=network-online.target
```

```
[Service]
Type=forking
ExecStart=/etc/rc.d/rc.local start
TimeoutSec=0
RemainAfterExit=yes
```

NOTE

`network-online.target` is a target that actively waits until the network is up, where the definition of up is defined by the network management software. Usually it indicates a configured, routable IP address of some kind. Its primary purpose is to actively delay activation of services until the network is set up.

2. Run the following command to check whether file `/etc/rc.d/rc.local` is executable:
ls -l /etc/rc.d/rc.local
If the output indicates that the file is not executable, make it executable by following the instructions provided in [Solution 1](#).
3. Run the following command to notify **systemd** to reload the configuration file:
systemctl daemon-reload
4. Run the following command to restart **rc-local.service** to execute the `/etc/rc.d/rc.local` file:
systemctl restart rc-local.service

3.10 What OSs Are Supported If I Want to Install Docker on a Linux ECS?

If you want to install Docker on your Linux ECS, ensure that the OS of your ECS is supported by Docker.

[Table 3-2](#) lists the OSs supported by Docker.

Table 3-2 OSs supported by Docker

OS	Version	Location
CentOS	CentOS 7	https://docs.docker.com/install/linux/docker-ce/centos/
Debian	<ul style="list-style-type: none">• Buster 10• Stretch 9 (stable) / Raspbian Stretch	https://docs.docker.com/install/linux/docker-ce/debian/
Fedora	<ul style="list-style-type: none">• Fedora28• Fedora29	https://docs.docker.com/install/linux/docker-ce/fedora/
Ubuntu	<ul style="list-style-type: none">• Disco 19.04• Cosmic 18.10• Bionic 18.04 (LTS)• Xenial 16.04 (LTS)	https://docs.docker.com/install/linux/docker-ce/ubuntu/

3.11 Why Do the Modifications to `/etc/security/limits.conf` Not Take Effect After the ECS Restarts?

Symptom

Modifications to the `/etc/security/limits.conf` file become invalid after the ECS restarts.

Possible Causes

In Linux, file `/etc/security/limits.conf` and directory `/etc/security/limits.d/` are used to limit the amount of various system resources available to a process. The configuration files in the `/etc/security/limits.d/` directory take priority over the settings in the `/etc/security/limits.conf` file.

If the modifications to file `/etc/security/limits.conf` become invalid after the ECS restarts, the modifications may be overwritten by the corresponding configurations in files located in directory `/etc/security/limits.d/`.

Solution

Modify the files located in directory `/etc/security/limits.d/` or modify file `/etc/security/limits.conf`.

NOTE

If the modifications to file `/etc/security/limits.conf` do not take effect, check the corresponding configurations in files located in the `etc/security/limits.d/` directory.

3.12 How Do I Set vCPU Affinity for Processes Using taskset?

Scenarios

Command `taskset` allows you to set vCPU affinity for processes running on an ECS to optimize vCPU utilization.

Application Scope

The operations described in this section apply to ECSs running CentOS and EulerOS.

Procedure

1. Run the following command to query the information about the vCPUs of the ECS:

```
cat /proc/cpuinfo
```

The key vCPU parameters are as follows:

- **processor** specifies the sequence number of a vCPU.
- **cpu cores** specifies the number of cores of each vCPU.

2. Run the following command to check the status of a process, for example, `test.sh` with PID 23989:

```
ps aux | grep test.sh
```

```
[root@ecs-linux-bj1 ~]# ps aux|grep test.sh
root      23989  0.0  0.2 113188  3016 pts/1    S   16:25   0:00 /bin/bash ./test.sh
```

3. Run the following command to query the vCPU that process `test.sh` is running on:

taskset -p *PID*

For example, run **taskset -p 23989**.

```
[root@ecs-linux-bj1 ~]# taskset -p 23989
pid 23989's current affinity mask: 1
```

The returned value is **1** in hexadecimal notation, which is **0001** in binary notation. Each **1** corresponds to a vCPU. **1** indicates that the process runs on the 0th vCPU.

4. Run the following command to assign the second vCPU (vCPU 1) to the process:

taskset -pc 1 *PID*

For example, run **taskset -pc 1 23989**.

NOTE

The vCPU ID starts from 0. vCPU 1 indicates the second vCPU. After the preceding command is executed, **test.sh** is bound to vCPU 1.

You can also run the following command to bind the process to vCPU 1 at startup:

```
taskset -c 1 ./test.sh&
```

3.13 What Should I Do If Error "command 'gcc' failed with exit status 1" Occurs During PIP-based Software Installation?

Symptom

When installing the Python library software, you need to configure the PIP source. An example is provided as follows:

```
[root@test home]# cat /root/.pip/pip.conf
[global]
index-url = https://pypi.mirrors.ustc.edu.cn/simple/
trusted-host = pypi.mirrors.ustc.edu.cn
```

During the installation, the system displays the message "command 'gcc' failed with exit status 1". However, GCC has been installed by running the yum command before the Python library software is installed using the PIP.

Figure 3-9 Installation error

```
creating build/temp.linux-x86_64-2.7/psutil
gcc -pthread -fno-strict-aliasing -O2 -g -pipe -Wall -Wp, -D_FORTIFY_SOURCE=2 -fexceptions -fstack-protector-strong --param=ssp-buffer-size=4 -grecord-gcc-switches -m64 -mtune=generic -D_GNU_SOURCE -fPIC -fwrapv -DNDEBUG -O2 -g -pipe -Wall -Wp, -D_FORTIFY_SOURCE=2 -fexceptions -fstack-protector-strong --param=ssp-buffer-size=4 -grecord-gcc-switches -m64 -mtune=generic -D_GNU_SOURCE -fPIC -fwrapv -fPIC -DPSUTIL_POSIX=1 -DPSUTIL_VERSION=543 -DPSUTIL_LINUX=1 -I/usr/include/python2.7 -c psutil/_psutil_common.c -o build/temp.linux-x86_64-2.7/psutil/_psutil_common.o
psutil/_psutil_common.c:9:28: fatal error: Python.h: No such file or directory
#include <Python.h>
^
compilation terminated.
error: command 'gcc' failed with exit status 1

-----
Command "/usr/bin/python2 -u -c "import setuptools, tokenize; file_='/tmp/pip-build-u8w1Ps/psutil/setup.py'; exec(compile(getattr(tokenize, 'open', open)(__file__).read().replace('\r\n', '\n'), __file__, 'exec'))" install --record /tmp/pip-hrs0Wn-record/install-record.txt --single-version-externally-managed --compile" failed with error code 1 in /tmp/pip-build-u8w1Ps/psutil/
[root@ecs-9a88-test1 .pip]#
```

Possible Causes

openssl-devel is not supported.

Solution

The following operations use psutil as an example:

1. Run the following command to install openssl-devel:
yum install gcc libffi-devel python-devel openssl-devel -y
2. Use PIP to install the Python library software again. The error message is cleared.

Figure 3-10 Successful installation

```
root@ecs-9a00-test1 .pip]# pip install psutil
Collecting psutil
  Using cached https://mirrors.aliyun.com/pypi/packages/c2/e1/600326635f97fee89bf0426fef14c5c29f4049c79f68f479f433d8c1bd96/psutil-5.4.3.tar.gz
Installing collected packages: psutil
  Running setup.py install for psutil ... done
Successfully installed psutil-5.4.3
```

3.14 What Can I Do If Switching from a Non-root User to a Root User Timed Out?

Symptom

When you run the **sudo** command to switch to user **root** on an Ubuntu or Debian ECS, the system prompts connection timeout.

Figure 3-11 Connection timeout

```
linux@ubuntu-test-1:/etc$
linux@ubuntu-test-1:/etc$ sudo su
sudo: unable to resolve host ubuntu-test-1: Connection timed out
root@ubuntu-test-1:/etc#
```

Solution

1. Log in to the ECS.
2. Run the following command to edit the **hosts** configuration file:
vi /etc/hosts
3. Press **i** to enter insert mode.
4. Add the IP address and hostname to the last line of the hosts configuration file.

Private IP address hostname

An example is provided as follows:

If the ECS hostname is **hostname** and the private IP address of the ECS is **192.168.0.1**, add the following:

```
192.168.0.1 hostname
```

5. Press **Esc** to exit editing mode.

6. Run the following command to save the configuration and exit:

```
:wq
```

NOTE

To update the hostname of an Ubuntu or Debian ECS, set the value of parameter **manage_etc_hosts** in the `/etc/cloud/cloud.cfg` file to **false** and update the new hostname in the `/etc/hosts` file. When editing the `/etc/hosts` file, do not delete the statement in the line where **127.0.0.1** is located. Otherwise, switching from a non-root user to user **root** will time out.

3.15 What Can I Do If the Permissions on the Root Directory of My CentOS ECS Changed to 777?

Symptom

You executed the **chmod -R 777 /** command on your ECS running CentOS, and the permissions on the root directory were changed to 777. As a result, most services were unavailable and most commands could not be executed. In this case, you can run the **getfacl** command provided by the system to back up and restore the permissions on the root directory.

NOTE

Setting 777 permissions for a file or directory means that it will be readable, writable and executable by all users and may pose a huge security risk. The operations described in this section are a temporary remedy. After you restore the permissions, it is recommended that you back up your data and reinstall the OS to prevent security risks caused by 777 permissions.

Procedure

This section uses an ECS running CentOS 7.5 as an example.

1. Do not stop or restart the ECS after the 777 permissions are configured for the root directory. Run the following commands on this ECS to restore the permissions on the SSH connection-related files:

```
cd /etc
```

```
chmod 644 passwd group shadow
```

```
chmod 400 gshadow
```

```
cd ssh
```

```
chmod 600 moduli ssh_host_dsa_key ssh_host_key ssh_host_rsa_key
```

```
chmod 644 ssh_config ssh_host_dsa_key.pub ssh_host_key.pub  
ssh_host_rsa_key.pub
```

```
chmod 640 sshd_config
```

2. After the preceding commands are executed, the SSH connection can be established. Run the following command to check whether you can remotely log in to the ECS as user **root**:

```
su root
```

```
root "su: cannot set groups:"
```

If information similar to the preceding output is displayed, run the following command to grant the **s** permission to the **su** program for reading the configurations of user **root**:

```
chmod u+s `which su`
```

After the command is executed, you can log in to the ECS as user **root**.

3. Create a temporary Linux ECS that has the correct permission configurations and the kernel of the same version as the faulty ECS.

4. Run the following command on the temporary ECS to back up the permissions on all files in the **/** directory to a **systemp.bak** file:

```
getfacl -R / >systemp.bak
```

5. Transmit the **systemp.bak** file to the faulty ECS using one of the following methods:

- Run the following command on the temporary ECS to upload the **systemp.bak** file to the faulty ECS:

```
scp Path in which the systemp.bak file is stored on the temporary ECS  
Username@EIP:Path in which the systemp.bak is to be stored on the  
faulty ECS
```

For example, run **scp /systemp.bak root@119.**.**.*/121.**.**.**

- Alternatively, run the following command on the faulty ECS to download the **systemp.bak** file from the temporary ECS to the faulty ECS:

```
scp Username@EIP: Path in which the systemp.bak file is stored on the  
temporary ECS Path in which the systemp.bak file is to be stored on the  
faulty ECS
```

For example, run **scp root@121.**.**.*/systemp.bak /.**

6. Run the following command on the faulty ECS to restore the system permissions:

```
setfacl --restore=systemp.bak
```

7. Run the following command to create a script before restarting the faulty ECS:

```
vim sshtmp.sh
```

Add the following content to the script to avoid SSH connection failure:

```
cat sshtmp.sh  
#-----start-----  
sleep 300  
cd /etc  
chmod 644 passwd group shadow  
chmod 400 gshadow  
cd ssh  
chmod 600 moduli ssh_host_dsa_key ssh_host_key ssh_host_rsa_key  
chmod 644 ssh_config ssh_host_dsa_key.pub ssh_host_key.pub ssh_host_rsa_key.pub  
chmod 640 sshd_config  
chmod u+s `which su`  
#-----end-----
```

8. Run the following command to add the script to the Linux startup script file **rc.local**:

```
echo '/sshtmp.sh &' >>/etc/rc.local
```

9. Run the following command to restart the faulty ECS:

```
reboot
```

After the ECS restarts, stop process `sshtmp.sh`, restore script file `rc.local`, delete process `sshtmp.sh`. Then check whether the system permissions are restored.

Follow-up Operations

Setting 777 permissions for a file or directory means that it will be readable, writable and executable by all users and may pose a huge security risk.

The operations described in this section are a temporary remedy. After you restore the permissions, it is recommended that you back up your data and reinstall the OS to prevent security risks caused by 777 permissions.

3.16 What Should I Do If the IP Settings of My Linux ECS Are Lost?

Symptom

When an ECS has been running continuously for a long time without being restarted, the IP address settings may be lost, ECSs may be disconnected, or the network may break down.

Figure 3-12 Symptom

```
[root@localhost ~]# ip ad
1: lo: <LOOPBACK> mtu 65536 qdisc noqueue state DOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc fq_codel state DOWN group default qlen 1000
   link/ether 52:54:01:09:a1:98 brd ff:ff:ff:ff:ff:ff
```

Possible Causes

The ECSs use Dynamic Host Configuration Protocol (DHCP).

For example, when an ECS running CentOS 7 is started, the NetworkManager service of the Linux OS starts the `dhclient` process. The `dhclient` requests the DHCP server to allocate an IP address and obtains the IP address lease expiration time.

In normal cases, the `dhclient` periodically updates the lease expiration time to the DHCP server to ensure the availability of the IP address.

Figure 3-13 Normal `dhclient`

```
root@localhost:~# systemctl status NetworkManager
NetworkManager.service - Network Manager
Loaded: loaded (/usr/lib/systemd/system/NetworkManager.service; enabled; vendor preset: enabled)
Drop-In: /usr/lib/daemon/systemd-[NetworkManager,NetworkManager-dispatcher]
Group: NetworkManager
Active: active (running) since Thu 2023-03-02 15:47:17 CST; 1s ago
Docs: man:NetworkManager(8)
Main PID: 38816 (NetworkManager)
CGroup: /systemd/system/NetworkManager.service
└─ 38816 /usr/bin/NetworkManager --no-daemon
└─ 38819 /usr/bin/NetworkManager --no-daemon
└─ 38820 /usr/bin/dhclient -pf /var/run/NetworkManager/dhclient.ether.pif -f /var/lib/NetworkManager/dhclient-5.0800-000-110-451-dad0016d1.ether lease=1 /var/lib/NetworkManager/dhclient.ether.conf eth
```

If the NetworkManager service is stopped by mistake and the `dhclient` process is cleared, the lease expiration time of the DHCP-assigned IP address cannot be automatically updated. When the lease expires, the IP address of the ECS is released, leading to network disconnection.

Solution

1. Remotely log in to the ECS.
2. Run the following commands to restart NetworkManager and make it automatically start upon ECS startup:

```
systemctl restart NetworkManager
```

```
systemctl enable NetworkManager
```

3. Run the following command to check the NetworkManager status:

```
systemctl status NetworkManager
```

Figure 3-14 Checking the NetworkManager status

```
[root@localhost ~]# systemctl status NetworkManager
● NetworkManager.service - Network Manager
   Loaded: loaded (/usr/lib/systemd/system/NetworkManager.service; enabled; vendor preset: enabled)
   Drop-In: /usr/lib/systemd/system/NetworkManager.service.d
            └─NetworkManager-ovs.conf
   Active: active (running) since Thu 2023-03-02 15:47:17 CST; 4min 1s ago
     Docs: man:NetworkManager(8)
    Main PID: 308119 (NetworkManager)
      Tasks: 4 (limit: 21656)
     Memory: 3.8M
    CGroup: /system.slice/NetworkManager.service
            └─ 308119 /usr/sbin/NetworkManager --no-daemon
               308129 /sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -pf /var/lib/NetworkManager/NetworkManager.conf

Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3023] dev eth0: dhcpv4: state: DHCPREQUEST
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3025] dev eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3032] man eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3036] man eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3037] pol eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost dhclient[308129]: bound to 192.168.122.59 -- renewal in 17999 seconds
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3177] dev eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3192] man eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3196] man eth0: dhcpv4: state: DHCPACK
Mar 02 15:47:17 localhost NetworkManager[308119]: <info> [1677743237.3264] pol eth0: dhcpv4: state: DHCPACK
```

4. Run the following command to check the network status:

```
ip ad
```

Figure 3-15 Checking the network connection

```
[root@localhost ~]# ip ad
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
   inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
   inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
   link/ether 52:54:01:09:a1:98 brd ff:ff:ff:ff:ff:ff
   inet 192.168.124.240/24 brd 192.168.124.255 scope global dynamic eth0
       valid_lft 3382sec preferred_lft 3382sec
   inet6 fe80::5054:1ff:fe09:a198/64 scope link
       valid_lft forever preferred_lft forever
```

If information similar to [Figure 3-15](#) is displayed, the network connection is normal.

Appendix

You are advised to retain the default network service configurations of common OSs.

Table 3-3 Default network service configurations of common OSs

OS	Network Service	Built-In DHCP
CentOS 6	Network	No. There is an independent dhclient process.
CentOS 7	NetworkManager	No. There is an independent dhclient process.
CentOS 8	NetworkManager	Yes
Ubuntu 16.04	NetworkManager	No. There is an independent dhclient process.
Ubuntu 18.04	NetworkManager	No. There is an independent dhclient process.
Ubuntu 20.04	NetworkManager	Yes
Ubuntu 22.04	NetworkManager	Yes

3.17 Why Does My Linux ECS Restart Unexpectedly?

Symptom

A Linux ECS restarts unexpectedly and the following error is displayed:

```
Kernel panic - not syncing: NMI: Not continuing
```

The following information is printed in the kernel log:

```
[645683.754132] Uhuh. NMI received for unknown reason 20 on CPU 1.  
[645683.754133] Do you have a strange power saving mode enabled?  
[645683.754133] Kernel panic - not syncing: NMI: Not continuing
```

Possible Causes

When the kernel parameter **kernel.unknown_nmi_panic** of the Linux ECS is set to **1**, the ECS panics and will automatically restart if the kernel detects a non-maskable interrupt (NMI).

Generally, the kernel parameter **kernel.unknown_nmi_panic** is set to **1** to proactively trigger panic if an NMI occurs. This helps locate the fault. Certain CPU models may generate an NMI in normal service processes and this may cause the ECS to restart unexpectedly.

Constraints

The operation in this section involves modifying kernel parameters at runtime, which may cause kernel instability. Evaluate risks before performing the operation.

Solution

1. Remotely log in to the ECS.
2. Run the following command to check the value of the ECS kernel parameter `kernel.unknown_nmi_panic`:

```
sysctl -n kernel.unknown_nmi_panic
```

If the value of `kernel.unknown_nmi_panic` is **1**, the abnormal restart is caused by the incorrect setting of this parameter.

Figure 3-16 Command output

```
linux-JjFcyI:~ # sysctl -n kernel.unknown_nmi_panic  
1
```

3. Run the following command to check the `kernel.unknown_nmi_panic` settings in the `/etc/sysctl.conf` file:

```
vim /etc/sysctl.conf
```

Check whether `kernel.unknown_nmi_panic=1` exists.

- If `kernel.unknown_nmi_panic=1` exists, change it to `kernel.unknown_nmi_panic=0`.
- If `kernel.unknown_nmi_panic=1` does not exist, add `kernel.unknown_nmi_panic=0`.

Figure 3-17 Viewing the `/etc/sysctl.conf` file

```
# sysctl settings are defined through files in  
# /usr/lib/sysctl.d/, /run/sysctl.d/, and /etc/sysctl.d/.  
#  
# Vendors settings live in /usr/lib/sysctl.d/.  
# To override a whole file, create a new file with the same in  
# /etc/sysctl.d/ and put new settings there. To override  
# only specific settings, add a file with a lexically later  
# name in /etc/sysctl.d/ and put new settings there.  
#  
# For more information, see sysctl.conf(5) and sysctl.d(5).  
#  
kernel.unknown_nmi_panic=0  
~
```

4. Press **Esc**, enter `:wq`, and press **Enter** to save the settings and exit.
5. Run the following command to make the configuration take effect:

```
sysctl -p
```

Figure 3-18 Making configuration take effect

```
linux-JjFcyI:~ # sysctl -p  
kernel.unknown_nmi_panic = 0
```

NOTE

The configuration takes effect without the need to restart the ECS.

Verification

1. Run the following command to check whether the value of **panic_on_unrecovered_nmi** is 0:

```
cat /proc/sys/kernel/panic_on_unrecovered_nmi
```

Figure 3-19 Command output (1)

```
linux-JjFcyI:~ # cat /proc/sys/kernel/panic_on_unrecovered_nmi  
0
```

2. Run the **sysctl -n kernel.unknown_nmi_panic** command to check whether the value of **kernel.unknown_nmi_panic** is 0.

Figure 3-20 Command output (2)

```
linux-JjFcyI:~ # sysctl -n kernel.unknown_nmi_panic  
0
```

If the results meet the expectation, the modification is successful.

3.18 What Do I Do If Error "Cannot allocate memory" Is Displayed?

Symptom

A Linux ECS has sufficient memory, but the error "Cannot allocate memory" is displayed during command execution or service start.

Command output:

```
root@localhost:~# free -m  
total used free shared buffers/cached available  
Mem: 3890 125 3179 2 504 3463  
Swap: 0 0 0  
root@localhost:~# uname -a  
-bash: fork: Cannot allocate memory
```

Possible Causes

The number of running processes or threads in the system reaches the upper limit. The upper limit is specified by the kernel parameter **/proc/sys/kernel/pid_max**.

Solution

1. Remotely log in to the ECS.
2. Run the following command to check the number of running processes:

```
ps -eLf | wc -l
```

The command output is as follows:

```
32753
```

You can also run the **sar -q** command to view the number of running processes in the **plist-sz** column.

 NOTE

If "command not found" is displayed, run the following command to install the software package first:

- CentOS: **yum install sysstat**
- Ubuntu: **apt install sysstat**

sar -q

The command output is as follows:

```
...
00:00:01   runq-sz  plist-sz  ldavg-1  ldavg-5  ldavg-15  blocked
10:45:01     0   32722    0.10   0.12   0.16    0
10:46:01     1   32730    0.21   0.15   0.17    0
10:47:01     2   32752    0.07   0.11   0.15    0
```

3. Run the following command to check the maximum number of processes configured in the system:

```
cat /proc/sys/kernel/pid_max
```

```
32768
```

If the number of running processes is close to or has reached the maximum number of processes, you need to increase the value of **/proc/sys/kernel/pid_max**.

4. Run the following commands to change the value of **/proc/sys/kernel/pid_max**, for example, to **65530**.

```
echo "kernel.pid_max=65530" >> /etc/sysctl.conf
```

```
sysctl -p
```

The command output is as follows:

```
vm.swappiness = 0
net.core.somaxconn = 1024
net.ipv4.tcp_max_tw_buckets = 5000
net.ipv4.tcp_max_syn_backlog = 1024
kernel.pid_max = 65530
```

3.19 What Can I Do If the Fork Process Failed and New Threads Cannot Be Created?

Symptom

The following error message is displayed during command execution or log printing on Linux ECSs.

Error message 1:

```
root@localhost:~# free -g
      total    used    free   shared  buffers   cached
Mem:     94      43     51      0        0        0
Swap:    19       0     19
root@localhost:~# uname -a
-bash: fork: Cannot allocate memory
```

Error message 2:

```
xxxxsshd2[23985]: fatal: setresuid 20054: Resource temporarily unavailable
xxxxsshd2[28377]: Disconnecting: fork failed: Resource temporarily unavailable
xxxxsshd2[4484]: Disconnecting: fork failed: Resource temporarily unavailable
```

Error message 3:

```
[root@ecs-xxxx ~]$ sudo docker info
runtime/cgo: pthread_create failed: Resource temporarily unavailable
SIGABRT: abort
```

Possible Causes

Generally, the preceding errors occur because the thread fails to be created. The possible cause is that the ECS memory is insufficient and the number of current threads reaches the configured maximum value.

Solution

1. Log in to the management console.
2. Monitor the ECS memory usage using the server monitoring function. For details, see [Viewing ECS Metrics](#).
 - If the memory is insufficient, increase the memory or optimize the memory by referring to [General Operations for Modifying Specifications](#).
 - If the memory is sufficient, go to [3](#).
3. Log in to an ECS as user **root** and run the following commands to check the **message** and **dmesg** logs:

```
dmesg -T
```

```
cat /var/log/messages
```

- If the cgroup error message shown in [Figure 3-21](#) is displayed, go to [8](#).
- If there is no such message, go to [4](#).

Figure 3-21 Log error



```
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153337.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153338.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153339.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153340.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153341.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153342.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153343.scope
cgroup: fork rejected by pids controller in /user.slice/user-1000.slice/session-153344.scope
```

4. Run the following command to check the total number of threads in the current system:

```
ps -efL | wc -l
```

5. Run the following commands to obtain two values and compare the values with the total number of system threads queried in [4](#):

```
sysctl -a | grep pid_max
```

```
sysctl -a | grep threads-max
```

- If the total number of threads in the current system is close to either of the two values, modify the values of **pid_max** and **threads-max**. For details, see [Modifying the values of pid_max and threads-max](#).
- If they do not stay close, go to [6](#).

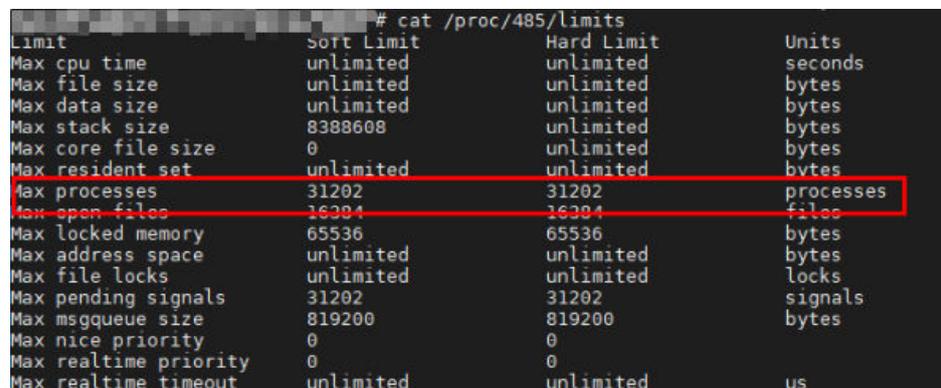
6. Run the following command to determine the PID of the error process:

```
ps -ef | grep Error process name
```

7. Run the following command to check the **limits** configuration of the process based on the PID obtained in the last step:

```
cat /proc/pid/limits
```

Figure 3-22 Checking the limits configuration of the process



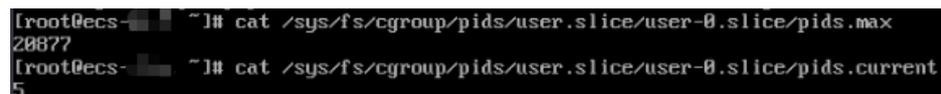
Limit	Soft Limit	Hard Limit	Units
Max cpu time	unlimited	unlimited	seconds
Max file size	unlimited	unlimited	bytes
Max data size	unlimited	unlimited	bytes
Max stack size	8388608	unlimited	bytes
Max core file size	0	unlimited	bytes
Max resident set	unlimited	unlimited	bytes
Max processes	31202	31202	processes
Max open files	16384	16384	files
Max locked memory	65536	65536	bytes
Max address space	unlimited	unlimited	bytes
Max file locks	unlimited	unlimited	locks
Max pending signals	31202	31202	signals
Max msgqueue size	819200	819200	bytes
Max nice priority	0	0	
Max realtime priority	0	0	
Max realtime timeout	unlimited	unlimited	us

- Check the value of **Max processes**. If the number of threads is close to the value of **Max processes**, modify the value of **limits**. For details, see [Modifying the value of limits](#).
 - If they do not stay close, go to **8**.
8. Run the following commands to obtain the values of **pid_max** and **pids.current** based on the log error **cgroup**:

```
cat /sys/fs/cgroup/pids/Directory where the error in the combined logs is reported/pids.max
```

```
cat /sys/fs/cgroup/pids/Directory where the error in the combined logs is reported/pids.current
```

Figure 3-23 cgroup directory



```
[root@ecs- ~]# cat /sys/fs/cgroup/pids/user.slice/user-0.slice/pids.max
20877
[root@ecs- ~]# cat /sys/fs/cgroup/pids/user.slice/user-0.slice/pids.current
5
```

An example is as follows:

- a. Run the following command to search for the **cgroup** directory based on the PID of the process:

```
cat /proc/pid/cgroup
```

Figure 3-24 Searching for the cgroup directory based on the PID

```
[root@ecs-~]# cat /proc/889410/cgroup
13:perf_event:/
12:devices:/user.slice
11:files:/
10:rdma:/
9:blkio:/
8:cpu,cpuacct:/
7:cpuset:/
6:pids:/user.slice/user-0.slice/session-5.scope
5:net_cls,net_prio:/
4:memory:/user.slice/user-0.slice/session-5.scope
3:hugetlb:/
2:freezer:/
1:name=systemd:/user.slice/user-0.slice/session-5.scope
```

In the command output, `/user.slice/user-0.slice/session-5.scope/` in the `pids` line can be combined with `/sys/fs/cgroup/pids/` to specify the `cgroup` directory `/sys/fs/cgroup/pids/user.slice/user-0.slice/session-5.scope/`.

- b. Run the following commands to obtain the values of `pid_max` and `pids.current` based on the `cgroup` directory:

```
cat /sys/fs/cgroup/pids/user.slice/user-0.slice/session-5.scope/pids.max
```

```
cat /sys/fs/cgroup/pids/user.slice/user-0.slice/session-5.scope/pids.current
```

 - If the value of `pids.current` is close to that of `pid_max`, modify the value of `cgroup`. For details, see [Modifying the value of cgroup](#).
 - If they do not stay close, submit a service ticket.

Related Commands

- Modifying the values of `pid_max` and `threads-max`.
 - a. Default parameters vary by OS version. Therefore, run the following commands to query how `pid_max` and `threads-max` is configured:

```
sysctl -a | grep pid_max
```

```
sysctl -a | grep threads-max
```
 - b. Run the following commands to modify the values of `pid_max` and `threads-max`:

```
echo 'kernel.pid_max = 4194304' >> /etc/sysctl.conf
```

```
echo 'kernel.threads-max = 4194304' >> /etc/sysctl.conf
```
 - c. Run the following command to make the new value be applied:

```
sysctl -p
```
- Modifying the value of `limits`
 - a. Log in to the ECS as the user who starts the error reporting process and run the following command to query the current configuration of `limits`:

```
ulimit -u
```

- b. Run the following commands to configure a proper upper limit for **nproc** based on service requirements and the current value:

For example, run the following commands to configure **100000** for **nproc** as user **root**:

```
echo 'root soft nproc 100000' >> /etc/security/limits.conf
```

```
echo 'root hard nproc 100000' >> /etc/security/limits.conf
```

- c. Log in to the ECS again and run the following command to check whether the configuration has taken effect:

```
ulimit -u
```

- If the command output is the value configured in **b**, the configuration has taken effect. Restart the service process in this session.
 - If the command output is not the value configured in **b**, submit a service ticket.
- Modifying the value of **cgroup**
 - Temporary modification solution
Run the following command to temporarily change the upper limit of the **cgroup** directory to the maximum value:

```
echo max > /sys/fs/cgroup/pids/user.slice/user-0.slice/session-25.scope/pids.max
```
 - Permanent modification solution:
Run the following command to set the cgroup directory to infinity and modify the cgroup directory that exceeds the limit:
The value can be adjusted as required. After the modification is complete, restart the ECS for the configuration to be applied.

```
echo DefaultTasksMax=infinity >>/etc/systemd/system.conf  
echo DefaultTasksMax=infinity >>/etc/systemd/user.conf  
echo UserTasksMax=infinity >>/etc/systemd/logind.conf
```

3.20 What Can I Do If the ECS Startup or Remote Login Fails Due to Incorrect System Configurations?

Symptom

Linux ECSs cannot be started or logged in to due to incorrect configurations. When you cannot log in to a Linux ECS using SSH or VNC, you can use a common method to modify the incorrect configurations of the system disk to restore the ECS or system. This section describes how to detach the system disk from a faulty Linux ECS and attach this disk to another Linux ECS as a data disk to rectify the fault.

Possible Causes

Table 3-4 lists the incorrect system configurations that may cause startup or access failures.

Table 3-4 Common incorrect system configurations

Type	Typical Problem
Incorrect configuration	The /etc/fstab file is lost or incorrectly configured.
	The SELinux is incorrectly configured.
	The /etc/security/limits.conf is incorrectly configured.
	The configuration format of /etc/passwd is incorrect.
	The configuration format of /etc/shadow is incorrect.
	The configuration format of /etc/ssh/sshd_config is incorrect.
Unavailable file or directory	The /etc/ssh directory is deleted by mistake.
	The /etc/security directory is deleted by mistake.
	The /etc/passwd file is deleted by mistake.
	The /etc/shadow file is deleted by mistake.
	The /etc/ssh/sshd_config file is deleted by mistake.
Incorrect file permissions	Permissions for SSH private keys are excessive.
	Permissions for SSH public keys are excessive.
Incorrect kernel parameter configuration	The value of vm.nr_hugepages is set too large.

Procedure

Step 1 Prepare a normal ECS.

Prepare an ECS that can be accessed. Its OS must be the same as that of the faulty ECS.

- You can use an existing ECS.
- You can create an ECS. For details, see [Purchasing and Logging In to a Linux ECS](#).

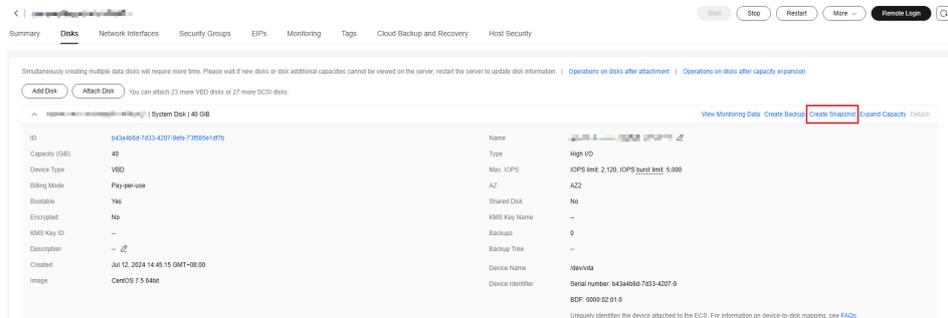
Step 2 Create a snapshot.

Before performing operations on the system disk of a faulty ECS, you are advised to create a snapshot for the system disk of this ECS to prevent data loss. For details, see [Creating an EVS Snapshot](#).

1. Log in to the management console.
2. Click  and choose **Compute > Elastic Cloud Server**.

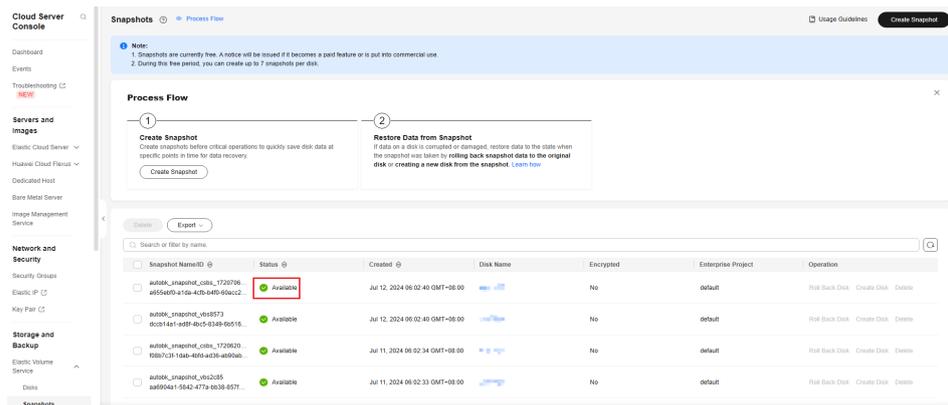
3. On the ECS list page, locate a faulty ECS and click its name.
4. On the displayed page, click the **Disks** tab.
5. On the **Disks** tab, click **Create Snapshot**.

Figure 3-25 Creating a snapshot



6. In the left navigation pane of the ECS console, choose **Elastic Volume Service** > **Snapshots** to check the snapshot status. If the **Status** is **Available**, the snapshot is created.

Figure 3-26 Checking the snapshot status



Step 3 Detach the system disk of a faulty ECS.

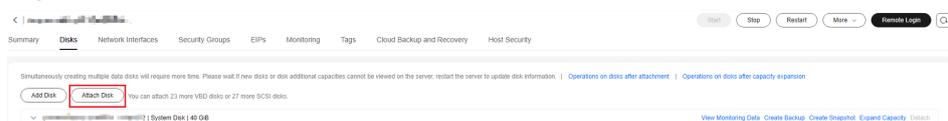
After the snapshot is created, stop the faulty ECS and detach the system disk from this ECS. For details, see [Detaching a System Disk](#).

Step 4 Attach the system disk of the faulty ECS as a data disk of a normal ECS.

The procedure is as follows:

1. Log in to the management console.
2. Click  and choose **Compute** > **Elastic Cloud Server**.
3. On the ECS list page, locate a normal ECS and click its name.
4. On the displayed page, click the **Disks** tab.
5. On the **Disks** tab, click **Attach Disk**.

Figure 3-27 Attach Disk



- In the displayed dialog box, select the system disk that has been detached from the faulty ECS in [Step 3](#), set **Disk Attribute** to **Data disk**, and click **OK**.

Figure 3-28 Attaching a data disk

- Wait for a while. In the left navigation pane of the ECS console, choose **Elastic Volume Service > Disks** to check the disk status. When the disk status changes to **In-use**, the disk is attached to a normal ECS.
- After that, log in to the ECS [using SSH or VNC](#) and run the **fdisk -l** command. The command output shows that a new disk device **/dev/vdb** and new partition **/dev/vdb1** are added to this ECS.

Figure 3-29 Checking a new disk device

```
# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x00046b12

   Device Boot      Start         End      Blocks   Id  System
  /dev/vda1  *            2048     83886079     41942016   83   Linux

Disk /dev/vdb: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x68619eae

   Device Boot      Start         End      Blocks   Id  System
  /dev/vdb1  *            2048     83884031     41940992   83   Linux
```

- Run the following command to attach the system disk of a faulty ECS to a normal ECS as a data disk:

```
mount <Data disk partition> <Mount point>
```

- Data disk partition: partition queried in [Step 4.8](#), for example, **/dev/vdb1**
- Mount point: local directory of a normal ECS, for example, **/mnt**

For example, if the data disk partition is **/dev/vdb1** and the mount point is **/mnt**, run the following command:

```
mount /dev/vdb1 /mnt
```

Step 5 Restore a data disk.

You can restore the data disk of a normal ECS in the following four scenarios:

Scenario 1: incorrect file configuration

For example, to modify the incorrect configuration of the **/etc/fstab** file of a faulty ECS, perform the following steps:

1. Run the following command to modify the **/mnt/etc/fstab** file on the data disk of a normal ECS:
vim /mnt/etc/fstab
2. Press **Esc** to exit insert mode and enter **:wq** to save the modification and exit.

 **NOTE**

This example describes how to restore the **/etc/fstab** file of a faulty ECS. To modify the configuration of other files, change the directory of the corresponding file.

For example, the **/etc/ssh/sshd_config** file corresponds to the **/mnt/etc/ssh/sshd_config** directory on the data disk.

Scenario 2: unavailable files or directories

For example, to restore the lost **/etc/security** directory of a faulty ECS, perform the following steps:

Run the following command to restore the lost **/mnt/etc/security** directory on the data disk of a normal ECS:

```
cp -rfa /etc/security /mnt/etc/
```

 **NOTE**

In this example, the lost **/etc/security** directory of a faulty ECS is restored. For other unavailable files or directories, change the directory of the corresponding file.

For example, if the **/etc/ssh/sshd_config** file is lost, run the **cp -rfa /etc/ssh/sshd_config /mnt/etc/ssh** command to restore it.

Scenario 3: incorrect file permission

For example, to modify the incorrect permission of the **/etc/ssh/ssh_host_ecdsa_key** file of a faulty ECS, perform the following steps:

Run the following command to modify the permission of the **/mnt/etc/ssh/ssh_host_ecdsa_key** file on the data disk of a normal ECS:

```
chmod 600 /mnt/etc/ssh/ssh_host_ecdsa_key
```

 **NOTE**

In this example, the permission of the **/etc/ssh/ssh_host_ecdsa_key** file of a faulty ECS is modified. To modify the permission of other files, change the directory of the corresponding file.

For example, the **/etc/ssh/ssh_host_ed25519_key** file corresponds to the **/mnt/etc/ssh/ssh_host_ed25519_key** directory on the data disk.

Scenario 4: kernel parameter optimization

⚠ CAUTION

Modifying kernel parameters may render kernel unstable.

For example, to configure the **vm.nr_hugepages** parameter, perform the following steps:

1. Run the following command to optimize the kernel parameters in the **/mnt/etc/sysctl.conf** file on the data disk of a normal ECS:
vim /mnt/etc/sysctl.conf
2. Press **/**, enter **nr_hugepages**, and press **Enter** to locate the configuration item.
3. Press **i** to edit the configuration item and set it to a proper value. If the configuration item is not found, add the following configuration item:
vm.nr_hugepages=4
The value must be **vm.nr_hugepages * hugepagesize < memory_total**. The specific value should be calculated as needed.
4. Press **Esc** to exit insert mode and enter **:wq** to save the modification and exit.
5. Run the following command for the configuration to take effect:
sysctl -p

Step 6 Restore the system disk of a faulty ECS.

After the preceding steps are performed, the configuration on the system disk of the faulty ECS has been modified. You can detach the data disk (system disk of the faulty ECS) from a normal ECS and attach it back to the faulty ECS. The procedure is as follows:

1. Detach the data disk (system disk of the faulty ECS) from a normal ECS. For details, see [Detaching an EVS Disk](#).
2. Attach the system disk back to the faulty ECS.

----End

3.21 Why Do df and du Commands Show Different Disk Usage?

Symptom

When you run **df** and **du** respectively to check disk usage, different results are returned.

As shown in the following figure, the disk usage queried by running **df -h** is greater than that queried by running **du -sh**.

```

~# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        186G   0    186G   0% /dev
tmpfs           186G   0    186G   0% /dev/shm
tmpfs           186G  110M   186G   6% /run
tmpfs           186G   0    186G   0% /sys/fs/cgroup
/dev/mapper/vgpaas-docker-sys 222G  190G   21G   91% /var/lib/docker
/dev/mapper/vgpaas-kubernetes 21G   16G   21G   81% /mnt/paas/kubernetes/kubelet
tmpfs           1.0G  12K  1.0G   1% /mnt/paas/kubernetes/kubelet/pods/e6ac9a4-2341-4832-8daa-04225075027e/volumes/kubernetes.io~projected/kube-api-access-khrgz
over1ay        64M   0    64M   0% /var/lib/docker/containers/9c6838c6450ef9b4d005b9b8af6ac3eaf6b47e2441d40496b7c90525594/mounts/shm
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/4435f8cf90b2c68a9e980134d1e609d5a6e3e2af4437300e2d3d3a7a/merged
tmpfs          1000M  12K 1000M   1% /mnt/paas/kubernetes/kubelet/pods/a4aff133-0023-4e59-9148-a068a08facb6/volumes/kubernetes.io~secret/cert
over1ay        222G  190G   21G   91% /mnt/paas/kubernetes/kubelet/pods/a4aff133-0023-4e59-9148-a068a08facb6/volumes/kubernetes.io~secret/cert
over1ay        64M   0    64M   0% /var/lib/docker/containers/72e411237339bd2c23871e0d72fc71ee651dee6181ea9b91fd23ab0f0f81/mounts/shm
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/d35c8a4f1024181a40b7fb23e7922b4e4ac09e19a477c2d2290837af37c/merged
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/e8a07e6949e457745c1c188e2c3678eeff7d909a60a1e8a7c06830d660fa/merged
over1ay        300M  12K  300M   1% /mnt/paas/kubernetes/kubelet/pods/4d70d1fa-e3f4-44e8-9347-8999ef5314/volumes/kubernetes.io~projected/kube-api-access-ccp2q
over1ay        222G  190G   21G   91% /var/lib/docker/containers/77cf48927c574c1880a5a0223030f77c2c8a95ecbe42f0d1dca27fb1abba/merged
over1ay        64M   0    64M   0% /var/lib/docker/containers/77cf48927c574c1880a5a0223030f77c2c8a95ecbe42f0d1dca27fb1abba/mounts/shm
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/3ab0eaa70e73c76575d830a9faeb1876e9fa0d83797ba78df73b81f77c16/merged
over1ay        300M  12K  300M   1% /mnt/paas/kubernetes/kubelet/pods/9f94c8e0950950e6832874e6e1c00c18f89021a4c396f0c69b07d9b706839/mounts/shm
over1ay        64M   0    64M   0% /var/lib/docker/containers/9f94c8e0950950e6832874e6e1c00c18f89021a4c396f0c69b07d9b706839/mounts/shm
over1ay        512M  12K 512M   1% /mnt/paas/kubernetes/kubelet/pods/0c012350-2d59-472c-b27a-a24232aa7fe/volumes/kubernetes.io~projected/kube-api-access-zhjt2
over1ay        64M   0    64M   0% /var/lib/docker/containers/0ccf33987930b581a50f36527858eac6d1a09f70b630704445912b966ff42/mounts/shm
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/19e01c37978d61024e6ff781d1d827c1330d05348c77125507adaf3/merged
over1ay        27G  12K  27G   1% /mnt/paas/kubernetes/kubelet/pods/22ac34fc-1844-44ee-4534-56b7f09f674/volumes/kubernetes.io~secret/adms-ion-cert
over1ay        27G  12K  27G   1% /mnt/paas/kubernetes/kubelet/pods/22ac34fc-1844-44ee-4534-56b7f09f674/volumes/kubernetes.io~secret/adms-ion-cert
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/11d61311f570b30ac1f9575e7d306187886215e939a93c7ba218489ced0/merged
over1ay        64M   0    64M   0% /var/lib/docker/containers/d831c13eefc9086e075e6f02eae00a17cb013e09ee6c607112668007f/mounts/shm
over1ay        222G  190G   21G   91% /var/lib/docker/over1ay2/ac8786707992f88a19d1228d0c244417154ce8600a34f54d2680ff77/merged
tmpfs          1.0G   0    1.0G   0% /run
~# du -sh /var/lib/docker/*
1.88343G 5pbp9  ~# du -sh /var/lib/docker
1.88343G 5pbp9  ~# lsblk
Name        1.88343G 5pbp9  ~# lsblk
vda         253:0 0 506G disk
--vdal      253:1 0 506G part /
vdb         253:16 0 250G disk
--vgpaas-docker-sys 253:0 0 250G lvm /var/lib/docker
--vgpaas-kubernetes 253:1 0 250G lvm /mnt/paas/kubernetes/kubelet
[root@w1ce-nodepool-1-88343-5pbp9 ~]#

```

Possible Causes

Generally, the disk space is released after a file is deleted. However, there are exceptions, for example, the file you deleted is locked or there is another process that keeps writing to the file.

A file is stored in the Linux file system in two parts: data and pointer. The pointer is stored in the metadata and data in the disk. After data is deleted, the pointer is cleared from the metadata. After the pointer is cleared, the disk space occupied by the data can be released, and new content can be written. If the disk space is not released after a file is deleted, it may be that the process keeps writing data into the file. As a result, the pointer is not cleared because the process is locked, so the system kernel considers that the file is not really deleted because the pointer is not deleted. The **df** command output indicates that the disk space is not released.

After a file is deleted, it is unavailable in the file system directory, so the **du** command does not collect it. However, if a running process occupies the handle of the deleted file, the file is not really deleted from the disk, and the information in the partition superblock is not changed. In this way, the **df** command still collects the deleted file.

Solution

1. Run the following command to switch to the **/opt** directory:
cd /opt
2. Run the following command to check all the files that have been deleted but are still occupied by processes:
lsolf | grep deleted

```

[root@w1ce-nodepool-1-88343-5pbp9 ~]# lsolf | grep deleted
[roo 1378481 54980 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2492 1378481 54980 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2508 1378481 132172 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2511 1378481 161898 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Shar.dre.v 1378481 166663 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Per1od1.c 1378481 167029 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Balance.r 1378481 167030 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Noopw.rte 1378481 167033 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Shar.dreg1 1378481 167038 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Reshar.d.g 1378481 167042 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2519 1378481 167105 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2522 1378481 167108 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2525 1378481 167111 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2526 1378481 167112 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2527 1378481 167113 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
M1rrorMae 1378481 167117 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Balance.r 1378481 167140 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Balance.c 1378481 167141 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2533 1378481 167180 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2536 1378481 167183 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2540 1378481 167197 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2543 1378481 167200 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn2545 1378481 167221 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn9914 1378481 869027 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Rep1Coor.t 1378481 1099564 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
conn9921 1378481 1317640 root 6w REG 252,0 19440685331 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Signalhan 1378481 1378485 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Backgro.p.k 1378481 1378486 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
OCSPMan.r 1378481 1378487 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Cert1ff1.m 1378481 1378489 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
Flow.con.f 1378481 1378489 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
mongod 1378481 1378500 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
mongod 1378481 1378521 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)
mongod 1378481 1378522 root 6w REG 252,0 194406850581 14156105 /var/lib/docker/mongodb/log/mongodb. log-20240214 (deleted)

```

The command output shows that the `/var/lib/docker/mongodb/log/mongodb` log file occupies 190 GB of memory. The **deleted** status indicates that it has been deleted, but the process keeps writing data to the file, so the disk space is not released.

3. Run the following command to clear the file:

```
echo >"/var/log/docker/mongodb/log/mongodb"
```

In this way, the disk space can be released and the process can continue to write logs to files.

3.22 What Can I Do If NetworkManager Cannot Be Started? (Error Message: Failed to restart NetworkManager.service: Unit NetworkManager.service is masked)

Symptom

The error message "Failed to restart NetworkManager.service: Unit NetworkManager.service is masked" is displayed when NetworkManager is started.

Possible Causes

The service unit is disabled, so NetworkManager cannot be started.

Solution

Run the following command to clear the **mask** status of the service unit:

```
systemctl unmask NetworkManager
```

3.23 Why Is the IP Address Lost After the System Time of an ECS Is Modified?

Symptom

After the system time of an ECS is modified, its IP address is lost.

Possible Causes

The time range you modified exceeds the DHCP lease time.

The default DHCP lease time set when you create a VPC subnet is 365 days (24 hours for VPC subnets created before). If you manually change the ECS system time and the time difference between the old and new time is longer than 24 hours, the DHCP lease may be expired and the ECS IP address will be lost.

Solution

If you do need to change the ECS system time, and the time difference is longer than your DHCP lease time, change the ECS IP address obtaining mode to static before you change the ECS system time.

4 Remote Login

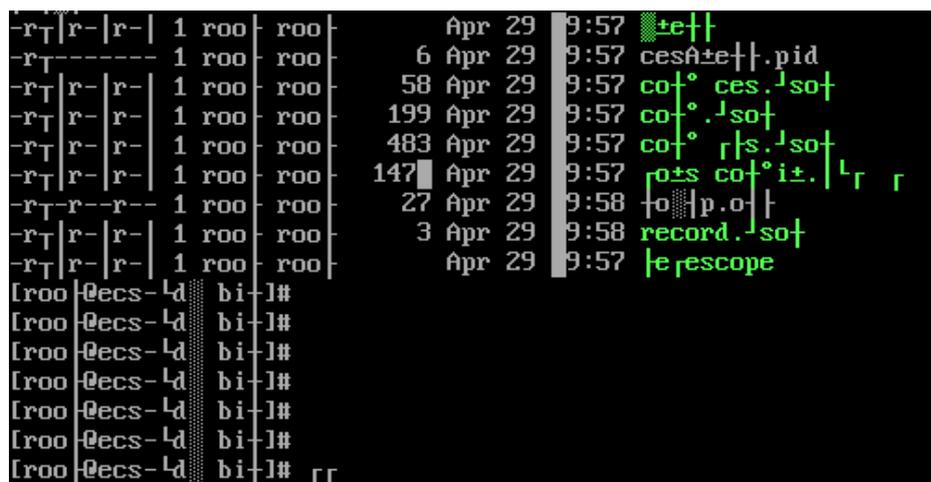
4.1 VNC Login

4.1.1 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC?

Symptom

After I attempt to log in to my Linux ECS using VNC, garbled characters are displayed, as shown in [Figure 4-1](#).

Figure 4-1 Garbled characters on the VNC-based login page



```
-rT|r-|r-| 1 roo | roo | Apr 29 | 9:57 | te++
-rT|----- 1 roo | roo | 6 Apr 29 | 9:57 | cesAte+.pid
-rT|r-|r-| 1 roo | roo | 58 Apr 29 | 9:57 | cot° ces.↓so†
-rT|r-|r-| 1 roo | roo | 199 Apr 29 | 9:57 | cot° .↓so†
-rT|r-|r-| 1 roo | roo | 483 Apr 29 | 9:57 | cot° |†s.↓so†
-rT|r-|r-| 1 roo | roo | 147 Apr 29 | 9:57 | o±s cot° i±. |†† r
-rT|r-|r-| 1 roo | roo | 27 Apr 29 | 9:58 | to†|p.o†|
-rT|r-|r-| 1 roo | roo | 3 Apr 29 | 9:58 | record.↓so†
-rT|r-|r-| 1 roo | roo | Apr 29 | 9:57 | te rescope
[roo @ecs-ld bi-]#
```

Possible Causes

The `cat` command was executed to display a large binary file, leading to garbled characters.

Solution

Log in to the ECS as user **root** and run the following command for recovery:

reset

NOTE

The **reset** command is used to re-initialize the ECS and refresh the terminal display. After the **reset** command is executed, the garbled characters are cleared and the fault is rectified.

4.1.2 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?

After you log in to an ECS using VNC and view data, for example, play videos or run the **cat** command to view large files in Linux OSs, VNC may become unavailable due to the high memory usage of the browser.

In such a case, use another browser and log in to the ECS again.

If the fault persists, contact technical support.

4.1.3 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?

The blank screen means that another user has logged in to this ECS using VNC, so you were logged out.

Only one user can be logged in to an ECS using VNC at a time. If you are already logged in and another user logs in to the same ECS, you will be automatically logged out.

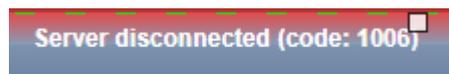
You can log back in, but that will kick the other user out.

4.1.4 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console?

Symptom

When I attempted to remotely log in to an ECS using VNC, the system displayed error code 1006, as shown in [Figure 4-2](#).

Figure 4-2 Error message displayed in a VNC-based remote login



Possible Causes

- The ECS is abnormal.
- Another user has logged in to the ECS.

- No operations are performed on the ECS and it is automatically disconnected.

Troubleshooting

1. Log in to the ECS again using VNC.
 - If the login is successful, no further action is required.
 - If the fault persists, go to [2](#).
2. Check whether the ECS is normal.
Error code 1006 is displayed if the ECS is stopped, deleted, being migrated or restarted, or encounters a connection timeout.
3. Check whether another user has logged in to the ECS.
If yes, you can log in to the ECS only after that user logs out.

4.2 SSH Connection Issues

4.2.1 How Do I Keep an SSH Session Alive?

Scenarios

This section describes how to keep an SSH session being connected to an ECS running CentOS 6.5.

The operations described in this section apply to ECSs running CentOS or EulerOS.

NOTE

Restarting `sshd` in the following procedure will disconnect SSH sessions.

Procedure

Modify the `/etc/ssh/sshd_config` file to keep an SSH session alive.

1. Add the following configuration items to the `/etc/ssh/sshd_config` file:

```
ClientAliveInterval 600
ClientAliveCountMax 10
```

ClientAliveInterval 600 indicates that the ECS sends a request to the SSH client every 600 seconds to keep the SSH connection alive.

ClientAliveCountMax 10 indicates that the ECS automatically disconnects from the client if the client does not respond for 10 times after the ECS sends requests.

In such a case, the unresponsive SSH client will be disconnected after approximately 6,000 seconds (600 x 10).

 NOTE

ClientAliveInterval is a timeout interval in seconds. If the ECS receives no data from the client within the timeout interval, the ECS will send a message through the encrypted channel to request a response from the client. The default value is **0**, indicating that these messages will not be sent to the client. The **ClientAliveInterval** option only applies to SSH protocol version 2.

ClientAliveCountMax specifies the number of client active messages which may be sent when the sshd service does not receive any messages back from the client. If this threshold is reached while client alive messages are being sent, sshd will disconnect the client, terminating the session.

The use of client alive messages is very different from **TCPKeepAlive**. The client alive messages are sent through the encrypted channel and will not be spoofable. The TCP keepalive option enabled by **TCPKeepAlive** is spoofable.

2. Run the following command to restart sshd for the configuration to take effect:
 - CentOS 6
service sshd restart
 - CentOS 7 or EulerOS
systemctl restart sshd

4.2.2 How Can I Allow or Deny Login from Specific Users or IP Addresses to an ECS Using SSH?

Scenarios

This section describes how to allow or deny specific users or IP addresses to access an ECS using SSH.

Constraints

- DenyHosts has been installed on the ECS.
- The sshd service should be restarted at a time that does not affect service running.

Method 1: Edit the sshd Configuration File

1. Allowing specific users (whitelist)

Add the usernames and IP addresses to be allowed to **AllowUsers** in the **/etc/ssh/sshd_config**. For example, to allow user **test** to access the ECS through 192.168.1.2, add **test@192.168.1.2** to **AllowUsers**:

```
AllowUsers test@192.168.1.2
```

 NOTE

After the configuration takes effect, only the allowed users can log in to the ECS.

2. Denying specific users (blacklist)

Add the usernames to be denied to **DenyUsers** in the **/etc/ssh/sshd_config**. For example, to deny user **testuser** to access the ECS, add **testuser** to **DenyUsers**:

```
DenyUsers testuser
```

Restart sshd for the modifications to take effect.

For CentOS 6, run the following command:

```
service sshd restart
```

For CentOS 7 or EulerOS, run the following command:

```
systemctl restart sshd
```

Method 2: Use DenyHosts

The `/etc/hosts.allow` and `/etc/hosts.deny` files of a Linux ECS are used to allow or deny access from an IP address or an IP address range to the ECS using SSH.

1. To allow the IP address 192.168.1.3 to access the ECS using SSH, add the following content to the `/etc/hosts.allow` file:

```
sshd: 192.168.1.3
```

2. To deny all IP addresses to access the ECS using SSH, add the following content to the `/etc/hosts.deny` file:

```
sshd: ALL
```

NOTE

`hosts.allow` has a higher priority than `hosts.deny`. In the preceding example, only SSH login from 192.168.1.3 is allowed. All other SSH connections to the ECS will be denied.

4.2.3 Why Can't I Access an ECS Running CentOS 7 Using SSH After I Changed the Default SSH Port?

Symptom

After the default port of the SSH service is changed, and inbound traffic on the new port is allowed by the ECS security group, you could not access the ECS using SSH.

Constraints

The operations described in this section apply to ECSs running CentOS 7.

Possible Causes

1. Log in to the management console and then log in to the ECS using .
2. Run the following command to check whether firewalld is enabled on the ECS:

```
systemctl status firewalld
```

Figure 4-3 firewalld enabled

```
[root@ecs-test ~]# systemctl status firewalld
■ firewalld.service - firewalld - dynamic firewall daemon
   Loaded: loaded (/usr/lib/systemd/system/firewalld.service; disabled; vendor preset: enabled)
   Active: active (running) since Wed 2019-06-19 11:41:21 CST; 44s ago
     Docs: man:firewalld(1)
   Main PID: 5059 (firewalld)
   CGroup: /system.slice/firewalld.service
           └─5059 /usr/bin/python -Es /usr/sbin/firewalld --nofork --nopid

Jun 19 11:41:20 ecs-test systemd[1]: Starting firewalld - dynamic firewall daemon...
Jun 19 11:41:21 ecs-test systemd[1]: Started firewalld - dynamic firewall daemon.
```

As shown in [Figure 4-3](#), firewalld is enabled.

3. Run the following command to view the rules in firewalld:
firewall-cmd --list-all

Figure 4-4 firewalld rules

```
[root@ecs-test ~]# firewall-cmd --list-all
public (active)
  target: default
  icmp-block-inversion: no
  interfaces: eth0
  sources:
  services: ssh dhcpv6-client
  ports:
  protocols:
  masquerade: no
  forward-ports:
  source-ports:
  icmp-blocks:
  rich rules:
```

As shown in [Figure 4-4](#), the current zone of the firewall is public. By default, only the SSH and dhcpv6-client services are allowed in the public zone. The SSH service uses the default port 22. If the SSH service uses a different port, the SSH connections will be refused.

Solution

- Method 1: Stop firewalld and disable its automatic startup.
Use security groups and network ACLs for access control. If firewalld needs to be enabled, perform operations described in [Method 2](#).
systemctl stop firewalld
systemctl disable firewalld
- Method 2: Open port 55660 in firewalld.
firewall-cmd --add-port=55660/tcp --permanent --zone=public
firewall-cmd --reload

4.2.4 How Can I Resolve ECS Login Failures Due to Corrupt /etc/passwd?

Scenarios

This section describes how to handle login failures caused by corrupt `/etc/passwd` on a Linux ECS.

NOTE

- The emergency recovery solution in this section requires you to replace the corrupt **/etc/passwd** with the initial backup file **/etc/passwd-** in single-user mode. This operation may cause the loss of user information you added, including users who own applications. You can add such accounts again by referring to **/etc/shadow**.
- This solution involves restarting the ECS, which will interrupt services running on it.

Symptom

Multiple services in Linux fail to be started, such as **Failed to start Login service** and **Failed to start Authorization service**.

When you try to log in after the system is started, the system displays a message indicating that the password is incorrect.

```
[ 3.868117] cloud-init[371]: userhome = pwd.getpwnam(os.getuid()).pw_dir
[ 3.868362] cloud-init[371]: KeyError: 'getpwnam(): uid not found: 0'
[FAILED] Failed to start Initial cloud-init job (pre-networking).
See 'systemctl status cloud-init-local.service' for details.
[ OK ] Reached target Network (Pre).
[ OK ] Started Update UTMP about System Boot/Shutdown.
[ OK ] Reached target System Initialization.
[ OK ] Reached target Timers.
[ OK ] Reached target Paths.
[ OK ] Listening on D-Bus System Message Bus Socket.
[ OK ] Reached target Sockets.
[ OK ] Reached target Basic System.
Starting Login Service...
Starting Authorization Manager...
Starting Postfix Mail Transport Agent...
Starting System Logging Service...
Starting Dump dmesg to /var/log/dmesg...
Starting /etc/rc.d/rc.local Compatibility...
Starting Dynamic System Tuning Daemon...
[ OK ] Started irqbalance daemon.
Starting irqbalance daemon...
[ OK ] Started D-Bus System Message Bus.
Starting D-Bus System Message Bus...
[ INFO ] Network Manager is not active.
[DEPEND] Dependency failed for Network Manager Wait Online.
Starting LSB: Bring up/down networking...
[ OK ] Started System Logging Service.
[FAILED] Failed to start Login Service.
See 'systemctl status systemd-logind.service' for details.
[FAILED] Failed to start Authorization Manager.
See 'systemctl status polkit.service' for details.
[DEPEND] Dependency failed for Dynamic System Tuning Daemon.
[ OK ] Started Dump dmesg to /var/log/dmesg.
[ OK ] Started /etc/rc.d/rc.local Compatibility.
[ OK ] Found device /dev/hvcb.
[FAILED] Failed to start Postfix Mail Transport Agent.
See 'systemctl status postfix.service' for details.
[FAILED] Failed to start LSB: Bring up/down networking.
See 'systemctl status network.service' for details.
```

Possible Causes

The **/etc/passwd** and **/etc/shadow** files record all user information with one entry per line, each representing a user account. If the files get corrupted or deleted, the login service **systemd-logind.service** fails to be started. As a result, users cannot log in to the system.

Solution

1. Restart the ECS on the console and enter the single-user mode.
2. Check the **/etc/passwd** file.

```
cat /etc/passwd
```

```
[root@oracle mnt]# cat etc/passwd
REDIS0007# redis-ver3.2.3#
redis-bits#0#ctime##R#Used-mem###0
s#>
[root@oracle mnt]#
```

3. Check whether the passwd file is corrupted. If it is, replace it with the initial backup file.

```
cp /etc/passwd- /etc/passwd
```

NOTE

This operation will cause the loss of the user information you added, including users who own applications. This will lead to application startup failures. After the fault is rectified, add these users to the passwd file.

4. Exit the current root directory and switch to the root directory of initramfs:
exit
5. Restart the ECS.
6. (Optional) After the system is started, add the lost users. For example, add the Nginx owner **nobody** and set its shell to **/sbin/nologin**. Add users as needed, and set the shell of users who need to log in to the system to **/bin/bash**.

```
useradd nobody -s /sbin/nologin
```

4.2.5 Why Does It Take a Long Time to Connect to an ECS Using SSH After UseDNS Is Enabled?

Symptom

It takes a long time to connect to an ECS using SSH.

Possible Causes

After the UseDNS option is enabled for the sshd service on the SSH server, when a client attempts to connect to the server using SSH, the server performs a DNS PTR reverse query to obtain the client's host name based on the client's IP address, and then performs a DNS forward A record query based on the client's host name, and check whether the two IP addresses are the same. This is a measure to prevent client spoofing. But in general, dynamic IP addresses do not have PTR records. You are advised to disable this option.

You can run the following command to check whether UseDNS is enabled:

```
grep UseDNS /etc/ssh/sshd_config
```

If the value is **yes** or the line is commented out, UseDNS is enabled. Disable UseDNS by performing the following operations.

Solution

1. Edit the **/etc/ssh/sshd_config** file.
vi /etc/ssh/sshd_config
2. Change the value of UseDNS to **no**.
UseDNS no
3. Restart the sshd service.
 - CentOS 6
service sshd restart
 - CentOS 7 or EulerOS
systemctl restart sshd

4.2.6 Why Does sshd Fail to Be Started on a Linux ECS?

Symptom

The sshd service fails to be started on a Linux ECS, and `/var/empty/sshd` cannot be accessed.

```
[root@test ssh]# systemctl restart sshd
Job for sshd.service failed because the control process exited with error code. See "systemctl status sshd.service" and "journalctl -xe" for details.
[root@test ssh]# systemctl status sshd
● sshd.service - OpenSSH server daemon
   Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: enabled)
   Active: activating (auto-restart) (Result: exit-code) since Thu 2018-05-17 17:31:14 CST; 10s ago
     Docs: man:sshd(8)
           man:sshd_config(5)
   Process: 1966 ExecStart=/usr/sbin/sshd -D $OPTIONS (code=exited, status=255)
   Main PID: 1966 (code=exited, status=255)
May 17 17:31:14 test.novalocal systemd[1]: sshd.service: main process exited, code=exited, status=255/n/a
May 17 17:31:14 test.novalocal systemd[1]: Failed to start OpenSSH server daemon.
May 17 17:31:14 test.novalocal systemd[1]: Unit sshd.service entered failed state.
May 17 17:31:14 test.novalocal systemd[1]: sshd.service failed.
```

Scenarios

The operations described in this section apply to ECSs running CentOS 7 or EulerOS. For other Linux OSs, the operations may be different.

Cause 1: Owner of `/var/empty/sshd` Is Not root

1. The sshd service fails to be started, and the journal log records `"/var/empty/sshd must be owned by root"`.

`journalctl -xe`

```
May 17 17:31:57 ecs-centos72-test systemd: Failed to start OpenSSH server daemon.
May 17 17:31:57 ecs-centos72-test systemd: Unit sshd.service entered failed state.
May 17 17:31:57 ecs-centos72-test systemd: sshd.service failed.
May 17 17:32:39 ecs-centos72-test systemd: sshd.service holdoff time over, scheduling restart.
May 17 17:32:39 ecs-centos72-test systemd: Starting OpenSSH server daemon...
May 17 17:32:39 ecs-centos72-test sshd: /var/empty/sshd must be owned by root and not group or world-writable.
May 17 17:32:39 ecs-centos72-test systemd: sshd.service: main process exited, code=exited, status=255/n/a
May 17 17:32:39 ecs-centos72-test systemd: Failed to start OpenSSH server daemon.
May 17 17:32:39 ecs-centos72-test systemd: Unit sshd.service entered failed state.
May 17 17:32:39 ecs-centos72-test systemd: sshd.service failed.
```

2. Check the owner of the `/var/empty/sshd` file.

`ll /var/empty/sshd`

```
[root@test ssh]# ll /var/empty/
total 4
drwx--x--x. 2 linux linux 4096 Oct 20 2017 sshd
```

As shown in the preceding figure, the owner of the `/var/empty/sshd` file is not `root`.

3. Modify the owner and permissions of the `/var/empty/sshd` file.

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

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

4. Restart the sshd service.

`systemctl restart sshd`

Cause 2: `/var/empty/sshd` Does Not Exist

1. Run the following command to query the failure causes:

`journalctl -xe`

```
-- Unit sshd.service has begun starting up.
Jun 28 09:56:44 ecs-gb.novalocal sshd[168661]: Missing privilege separation directory: /var/empty/sshd
Jun 28 09:56:44 ecs-gb.novalocal systemd[1]: sshd.service: main process exited, code=exited, status=255/w/a
Jun 28 09:56:44 ecs-gb.novalocal systemd[1]: Failed to start OpenSSH server daemon.
-- Subject: Unit sshd.service has failed
-- Defined-By: systemd
-- Support: http://lists.freedesktop.org/mailman/listinfo/systemd-devel
--
Unit sshd.service has failed.
--
The result is failed.
Jun 28 09:56:44 ecs-gb.novalocal systemd[1]: Unit sshd.service entered failed state.
Jun 28 09:56:44 ecs-gb.novalocal systemd[1]: sshd.service failed.
Jun 28 09:56:44 ecs-gb.novalocal polkitd[498]: Unregistered Authentication Agent for unix-process:16862:6014977 (system bus name
```

2. As shown in the preceding figure, the sshd service fails to be started because the `/var/empty/sshd` file does not exist. Run the following command to manually create the file:
mkdir -p /var/empty/sshd
3. Restart the sshd service.
systemctl restart sshd

4.2.7 How Do I Disable Login to an ECS Using SSH Password?

Scenarios

To ensure ECS security, you need to change the access key of an ECS regularly, or disable the SSH password login.

This section describes how to disable the SSH password login.

NOTE

- After disabling SSH password login, you can still log in to the ECS using the password on the ECS console. Keep your password secure.
- This setting method does not apply to ECSs running Ubuntu 22.04.

Procedure

1. Log in to a Linux ECS and run the following command to edit the SSH connection:
vi /etc/ssh/sshd_config
Modify the following item:
Change the value of **PasswordAuthentication** from **yes** to **no**.
2. Restart sshd for the modification to take effect.
service sshd restart
3. Restart the ECS.
4. Log in to the Linux ECS using an SSH password.
For details, see [Logging In to a Linux ECS Using an SSH Password](#).
If the login fails, the login to an ECS using an SSH password is disabled.

4.2.8 Why Are Connections to a Linux ECS Using SSH or to Applications on the ECS Interrupted Occasionally?

Scenarios

This section applies to the scenario where the connection to a Linux ECS using SSH or the access to applications on the ECS is interrupted occasionally.

Constraints

1. Modifying kernel parameters may render kernel unstable.
2. To ensure the system stability, restart the system at a proper time after modifying kernel parameters.

Possible Causes

1. Check whether the `net.ipv4.tcp_tw_recycle` and `net.ipv4.tcp_tw_reuse` options are enabled to quickly reclaim and reuse `TIME_WAIT` connections:
`sysctl -a |grep tcp_tw`
As shown in [Figure 4-5](#), the options have been enabled.

Figure 4-5 TIME_WAIT

```
[root@ecs-feilsystemd-test ~]# sysctl -a |grep tcp_tw
sysctl: reading key "net.ipv6.conf.all.stable_secret"
sysctl: reading key "net.ipv6.conf.default.stable_secret"
sysctl: reading key "net.ipv6.conf.eth0.stable_secret"
sysctl: reading key "net.ipv6.conf.lo.stable_secret"
net.ipv4.tcp_tw_recycle = 1
net.ipv4.tcp_tw_reuse = 1
```

2. With `net.ipv4.tcp_tw_recycle` and `net.ipv4.tcp_tw_reuse` enabled, `TIME_WAIT` connections will be quickly reclaimed and reused, resulting in disconnections. By default, these two options are disabled.

NOTE

In the NAT environment, multiple terminals use the same public IP address, and one-to-one connection between the server and client cannot be implemented. If these two options are enabled, the server will reclaim and reuse TCP connections in the `TIME_WAIT` state, resulting in disconnections.

Procedure

1. Disable the preceding two kernel parameters by adding the following content to the `/etc/sysctl.conf` file:

```
net.ipv4.tcp_tw_recycle = 0
net.ipv4.tcp_tw_reuse = 0
```
2. Run the following command for the modification to take effect:

```
sysctl -p
```

NOTE

If kernel parameters are modified at runtime, the kernel may fail to load the parameters stably. You are advised to restart the system at a proper time.

4.2.9 What Do I Do If I Can Log In to an Ubuntu 16.04 ECS Using SSH But the VNC Login Page Cannot Be Displayed?

Symptom

You have successfully connected to the Ubuntu 16.04 ECS using SSH, but the VNC login page cannot be displayed.

Constraints

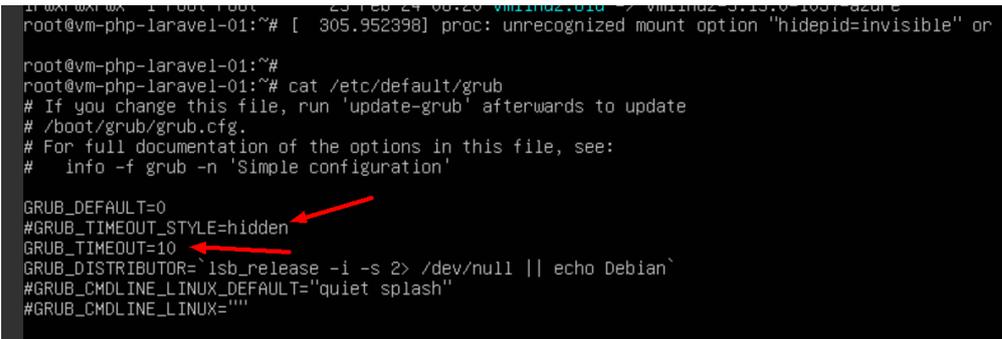
The operation in this section involves modifying the GRUB configuration file. Misoperations may cause system startup failures. Before the operation, back up the GRUB configuration file so that it can be restored in case of misoperations.

Solution

1. Run the following command to modify the GRUB configuration after you log in to the Ubuntu 16.04 ECS using SSH:

```
cat /etc/default/grub
```

Comment out **GRUB_TIMEOUT_STYLE=hidden** and modify **GRUB_TIMEOUT=10**, as shown in the following figure.



```
root@vm-php-laravel-01:~# [ 305.952398] proc: unrecognized mount option "hidepid=invisible" or
root@vm-php-laravel-01:~#
root@vm-php-laravel-01:~# cat /etc/default/grub
# If you change this file, run 'update-grub' afterwards to update
# /boot/grub/grub.cfg.
# For full documentation of the options in this file, see:
#   info -f grub -n 'Simple configuration'

GRUB_DEFAULT=0
#GRUB_TIMEOUT_STYLE=hidden
GRUB_TIMEOUT=10
GRUB_DISTRIBUTOR=`lsb_release -i -s 2> /dev/null || echo Debian`
#GRUB_CMDLINE_LINUX_DEFAULT="quiet splash"
#GRUB_CMDLINE_LINUX=""
```

2. Run the following command to delete all files that start with **50** in the **/etc/default/grub.d/** directory:

```
rm -rf /etc/default/grub.d/50*
```

3. Run the following commands to update the configuration:

```
update-grub2
```

4. Run the following commands to modify the yum repository to install the public kernel:

```
sed -i 's/azure.archive.ubuntu.com/repo.huaweicloud.com/g' /etc/apt/sources.list
```

```
apt autoclean
```

```
apt update
```

5. Run the following command to install the Ubuntu 16.04 public kernel:

```
apt install linux-image-generic
```

6. Wait until the installation is complete and the system is restarted, and select the generic kernel on the GRUB page to start the system.

7. (Optional) Run the following command to delete **azurelinuxagent** because the agent keeps printing logs to the VNC console, which affects the VNC functions.

```
sudo apt -y remove walinuxagent
```

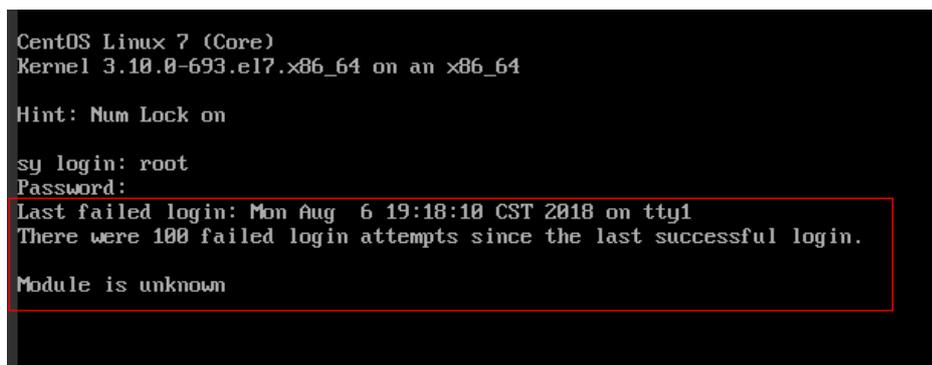
4.3 Remote Login Errors on Linux

4.3.1 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?

Symptom

When you attempt to remotely log in to a Linux ECS, the system displays the error message "Module is unknown".

Figure 4-6 Module is unknown



```
CentOS Linux 7 (Core)
Kernel 3.10.0-693.el7.x86_64 on an x86_64

Hint: Num Lock on

sy login: root
Password:
Last failed login: Mon Aug 6 19:18:10 CST 2018 on tty1
There were 100 failed login attempts since the last successful login.

Module is unknown
```

NOTE

- To resolve this issue, restart the ECS and enter the rescue mode.
- Restarting the ECS may interrupt services. Exercise caution when performing this operation.

Root Cause

The file in the **/etc/pam.d/** directory was modified by mistake.

Solution

1. Enter single-user mode.
The following uses CentOS 7 as an example:
 - a. Restart the ECS and click **Remote Login**.
 - b. Click **Ctrl+Alt+Del** in the upper part of the remote login panel to restart the ECS.
 - c. Press the up arrow key to prevent automatic system startup. When the kernels are displayed, press **e** to enter the editing mode.

Figure 4-7 Entering the kernel editing mode

```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) ? (Core)
CentOS Linux (3.10.0-693.el7.x86_64) ? (Core)
CentOS Linux (0-rescue-fd73dfefb69a4e4183fd3f86ccc29526) ? (Core)
```

NOTE

The grub file is encrypted by Euler images by default. Before entering the edit mode, you need to contact customer service to obtain username and password.

- d. Locate the row containing **linux16** and delete the parameters you do not require.
- e. Change **ro** to **rw** for mounting the root partition with read-write permissions.
- f. Add **rd.break** and press **Ctrl+X**.

Figure 4-8 Before the modification

```
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
  search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
  search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 ro crash_kexec_post_notifiers softlockup_panic=\
1 panic=3 reserve_kbox_mem=16M nmi_watchdog=1 rd.shell=0 net.ifnames=0 spectre\
v2=off nopti noibrs noibpb crashkernel=auto LANG=en_US.UTF-8
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

Figure 4-9 After the modification

```
insmod ext2
set root='hd0,msdos1'
if [ x$feature_platform_search_hint = xy ]; then
  search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
  search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 rw rd.break
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

- g. Run the following command to go to the **/sysroot** directory:
chroot /sysroot
2. Run the following command to view the system log for error files:
grep Module /var/log/messages

Figure 4-10 System log

```
Aug 6 18:00:09 sy login: pam_succeed_if(login:auth): requirement "uid >= 1000" not met by user "root"
Aug 6 18:00:11 sy login: FAILED LOGIN 1 FROM tty1 FOR root, Authentication failure
Aug 6 18:00:15 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=0)
Aug 6 18:00:15 sy login: Module is unknown
Aug 6 18:10:41 sy login: PAM unable to dlopen(/lib/security/pam_limits.so): /lib/security/pam_limits.so: cannot open shared obj\
ect file: No such file or directory
Aug 6 18:10:41 sy login: PAM adding faulty module: /lib/security/pam_limits.so
Aug 6 18:10:44 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=0)
Aug 6 18:10:44 sy login: Module is unknown
```

3. Comment out or modify the error line in the error files displayed in the system log.

```
vi /etc/pam.d/login
```

Figure 4-11 Modifying the error information

```
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
# session required /lib/security/pam_limits.so
```

4. Restart the ECS and try to log in to it again.

NOTE

- To view the modification records and check whether the modification is caused by unintended actions, run the following command:

```
vi /root/.bash_history
```

Search for the keyword **vi** or **login**.

- Do not modify the files in the **/etc/pam.d/** directory. Run the following command for details about pam:

```
man pam.d
```

4.3.2 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?

Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error "Permission denied".

Figure 4-12 Permission denied

```
CentOS Linux 7 (Core)
Kernel 3.10.0-693.11.1.el7.x86_64 on an x86_64

ecs-ams-03 login: ;
Password:

Permission denied
_
```

NOTE

- To resolve this issue, you are required to restart the ECS and enter the rescue mode.
- Restarting the ECS may interrupt services. Exercise caution when performing this operation.

Root Cause

The **nofile** parameter in **/etc/security/limits.conf** is used to set the maximum number of files that can be opened in the system. If the value is greater than the **fs.nr_open** value (**1048576** by default) set in **PermissionDenied.png**, a login verification error will occur, leading to "Permission denied".

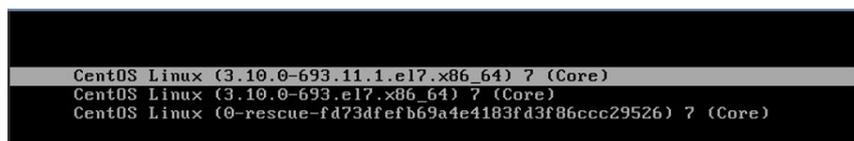
Solution

1. Enter single-user mode.

The following uses CentOS 7 as an example:

- a. Restart the ECS and click **Remote Login**.
- b. Click **Ctrl+Alt+Del** in the upper part of the remote login panel to restart the ECS.
- c. Press the up arrow key to prevent automatic system startup. When the kernels are displayed, press **e** to enter the editing mode.

Figure 4-13 Entering the kernel editing mode



NOTE

The grub file is encrypted by Euler images by default. Before entering the edit mode, you need to contact customer service to obtain username and password.

- d. Locate the row containing **linux16** and delete the parameters you do not require.
- e. Change **ro** to **rw** for mounting the root partition with read-write permissions.
- f. Add **rd.break** and press **Ctrl+X**.

Figure 4-14 Before the modification

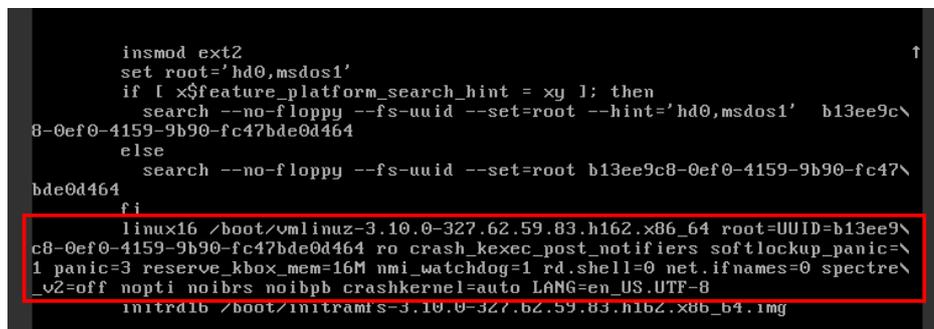


Figure 4-15 After the modification

```
insmod ext2
set root='hd0,msdos1'
if [ x${feature_platform_search_hint} = xy ]; then
  search --no-floppy --fs-uuid --set=root --hint='hd0,msdos1' b13ee9c\
8-0ef0-4159-9b90-fc47bde0d464
else
  search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47\
bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9\
c8-0ef0-4159-9b90-fc47bde0d464 rw rd.break
initrd16 /boot/initramfs-3.10.0-327.62.59.83.h162.x86_64.img
```

g. Run the following command to go to the `/sysroot` directory:

```
chroot /sysroot
```

2. Run the following command to view the `fs.nr_open` value:

```
sysctl fs.nr_open
```

3. Change the `nofile` value in `/etc/security/limits.conf` so that the value is smaller than the `fs.nr_open` value obtained in 2.

```
vi /etc/security/limits.conf
```

NOTE

`limits.conf` is the `pam_limits.so` configuration file of Linux Pluggable Authentication Module (PAM). For more details, run the following command:

```
man limits.conf
```

4. Restart the ECS and try to log in to it again.

4.3.3 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?

Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error message "read: Connection reset by peer".

Figure 4-16 read: Connection reset by peer

```
debug1: Enabling compatibility mode for protocol 2.0
debug1: Local version string SSH-2.0-OpenSSH_7.2p2 Ubuntu-4ubuntu2.4
ssh_exchange_identification: read: Connection reset by peer
ubuntu@node2:~$ _
```

Possible Causes

- The remote login port is not permitted in the security group.
- The firewall is enabled on the ECS, but the remote login port is blocked by the firewall.

Solution

Perform the following operations for troubleshooting:

- **Check security group rules.**
 - Inbound: Add the remote login port. The default port 22 is used as an example.
 - Outbound: Outbound rules allow network traffic to be out of specified ports.
- **Add a port to the ECS firewall exception.**

The following uses Ubuntu as an example:

 - a. Run the following command to view the firewall status:
sudo ufw status
The following information is displayed:
Status: active
 - b. Add a port to the firewall exception, taking the default port 22 as an example.
ufw allow 22
Rule added
Rule added (v6)
 - c. Run following command to check the firewall status again:
sudo ufw status
Status: active
To Action From
--
22 ALLOW Anywhere
22 (v6) ALLOW Anywhere (v6)

Try to remotely log in to the ECS again.

4.3.4 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?

Symptom

When you attempt to remotely log in to a Linux ECS, the system displays the error message "Access denied".

Possible Causes

- Incorrect username or password.
- A policy that denies logins from user **root** is enabled on the SSH server.

Solution

- The username or password is incorrect.
Check the username and password.
- A policy that denies logins from user **root** is enabled on the SSH server.
 - a. Edit the `/etc/ssh/sshd_config` file and check the following settings to ensure that the SSH logins from user **root** are allowed:
PermitRootLogin yes
 - b. Restart SSH.

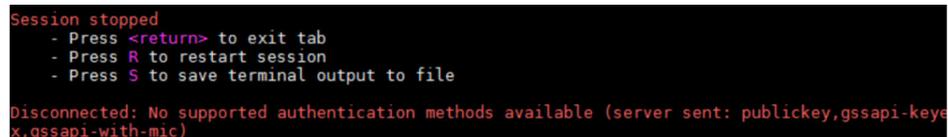
- CentOS 6
service sshd restart
- CentOS 7
systemctl restart sshd

4.3.5 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?

Symptom

When I attempted to remotely log in to a Linux ECS, the system displayed error message "Disconnected: No supported authentication methods available".

Figure 4-17 No supported authentication methods available



```
Session stopped
- Press <return> to exit tab
- Press R to restart session
- Press S to save terminal output to file
Disconnected: No supported authentication methods available (server sent: publickey,gssapi-keyex,gssapi-with-mic)
```

Possible Causes

A policy that denies password-authenticated logins is enabled on the SSH server.

Solution

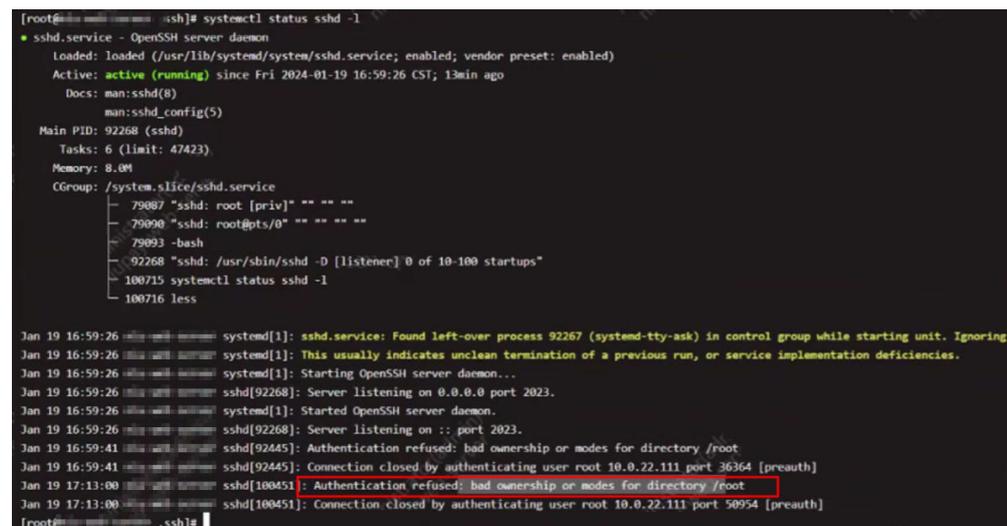
1. Open the **/etc/ssh/sshd_config** file and check the following settings:
vi /etc/ssh/sshd_config
2. Modify the following settings:
Change **PasswordAuthentication no** to **PasswordAuthentication yes**.
Alternatively, delete the comment tag (#) before **PasswordAuthentication yes**.
3. Restart SSH.
 - CentOS 6
service sshd restart
 - CentOS 7
systemctl restart sshd

4.3.6 What Do I Do If "Authentication refused: bad ownership or modes for directory /root" Is Displayed and I Can't Log In to an ECS Using SSH Key?

Symptom

You cannot log in to an ECS using an SSH key, and the error message "Authentication refused: bad ownership or modes for directory /root" is displayed.

Figure 4-18 Failure of logging in to the ECS using an SSH key



```
[root@ ~]# systemctl status sshd -l
● sshd.service - OpenSSH server daemon
   Loaded: loaded (/usr/lib/systemd/system/sshd.service; enabled; vendor preset: enabled)
   Active: active (running) since Fri 2024-01-19 16:59:26 CST; 13min ago
     Docs: man:sshd(8)
           man:sshd_config(5)
   Main PID: 92268 (sshd)
     Tasks: 6 (limit: 47423)
    Memory: 8.0M
   CGroup: /system.slice/sshd.service
           └─ 79087 "sshd: root [priv]" " " " " " "
           └─ 79090 "sshd: root@pts/0" " " " " " "
           └─ 79093 -bash
           └─ 92268 "sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startups"
           └─ 100715 systemctl status sshd -l
           └─ 100716 less

Jan 19 16:59:26 system[1]: sshd.service: Found left-over process 92267 (systemd-tty-ask) in control group while starting unit. Ignoring.
Jan 19 16:59:26 system[1]: This usually indicates unclean termination of a previous run, or service implementation deficiencies.
Jan 19 16:59:26 system[1]: Starting OpenSSH server daemon...
Jan 19 16:59:26 sshd[92268]: Server listening on 0.0.0.0 port 2023.
Jan 19 16:59:26 system[1]: Started OpenSSH server daemon.
Jan 19 16:59:26 sshd[92268]: Server listening on :: port 2023.
Jan 19 16:59:41 sshd[92445]: Authentication refused: bad ownership or modes for directory /root
Jan 19 16:59:41 sshd[92445]: Connection closed by authenticating user root 10.0.22.111 port 36364 [preauth]
Jan 19 17:13:00 sshd[100451]: Authentication refused: bad ownership or modes for directory /root
Jan 19 17:13:00 sshd[100451]: Connection closed by authenticating user root 10.0.22.111 port 50954 [preauth]
[root@ ~]#
```

Possible Causes

Error information in logs indicates that the permission of the **/root** directory is incorrect. Check the permission, owner, and owner group of the **/root** directory, **.ssh** in the **/root** directory, and **authorized_keys** in the **.ssh** directory.

Solution

Change the owner, owner group, and permission of the **/root** directory to recover the SSH.

```
[root@... ~]# ll /
total 64
lrwxrwxrwx. 1 root root 7 Mar 24 2023 bin -> usr/bin
dr-xr-xr-x. 6 root root 4096 Jan 15 10:06 boot
drwxr-xr-x 7 root root 4096 Jun 6 2023 CloudrResetPwdAgent
drwxr-xr-x 4 root root 4096 Jan 15 10:03 data
drwxr-xr-x 18 root root 3100 Jan 19 14:29 dev
drwxr-xr-x. 87 root root 4096 Jan 19 15:58 etc
drwxr-xr-x. 2 root root 4096 Mar 24 2023 home
lrwxrwxrwx. 1 root root 7 Mar 24 2023 lib -> usr/lib
lrwxrwxrwx. 1 root root 9 Mar 24 2023 lib64 -> usr/lib64
drwx----- 2 root root 16384 Jun 6 2023 lost+found
drwxr-xr-x. 2 root root 4096 Mar 24 2023 media
drwxr-xr-x. 2 root root 4096 Mar 24 2023 mnt
drwxr-xr-x. 2 root root 4096 Jan 15 10:04 opt
dr-xr-xr-x 183 root root 0 Jan 19 14:29 proc
drwxrwxr-x. 9 947 943 4096 Jan 19 16:40 root
drwxr-xr-x 31 root root 1040 Jan 19 16:00 run
lrwxrwxrwx. 1 root root 8 Mar 24 2023 sbin -> usr/sbin
drwxr-xr-x. 2 root root 4096 Mar 24 2023 srv
dr-xr-xr-x 12 root root 0 Jan 19 14:29 sys
drwxrwxrwt 10 root root 200 Jan 19 16:29 tmp
drwxr-xr-x. 12 root root 4096 Jun 6 2023 usr
drwxr-xr-x. 18 root root 4096 Jun 6 2023 var
```

5 Configuring the Network

5.1 Why Does My ECS Running CentOS 7 Fail to Obtain an IP Address Using dhclient?

Symptom

After an ECS was started, it failed to obtain an IP address because dhclient did not run.

Possible Cause

The possible causes are as follows:

1. NetworkManager is not configured to start at boot.
2. The target NIC is not managed by NetworkManager.

Constraints

The operations described in this section apply to ECSs running CentOS 7 series, EulerOS 2 series, or Ubuntu 18.04.

Solution

1. Run the following command to check whether dhclient is running:
ps -ef |grep dhclient |grep -v grep
2. If dhclient is not detected, run the following command to check whether NetworkManager is running:
systemctl status NetworkManager
 - If NetworkManager is in **Active: inactive (dead)** state, NetworkManager is not enabled. Run the following command to check whether NetworkManager is automatically started upon system startup:
systemctl is-enabled NetworkManager
If the command output is **disabled**, run the following command to make NetworkManager start at boot automatically:

```
systemctl enable NetworkManager && systemctl start NetworkManager
```

- If NetworkManager is in **Active: active (running)** state, run the following command to check whether the target NIC is managed by NetworkManager:

```
nmcli device status
```

If the NIC is in **unmanaged** state, run the following command to make NetworkManager manage the NIC:

```
nmcli device set eth0 managed yes
```

3. Run the following command to restart NetworkManager:

```
systemctl restart NetworkManager
```

4. Run the following command to check whether the private IP address can be obtained:

```
ip add
```

5.2 Why Does the NIC Names Change After I Start a Linux ECS?

Scenarios

After an ECS running CentOS 6 created using a private image starts, the first NIC is named as **eth1**, rather than **eth0**.

The solution provided in this section involves restarting the ECS. Restarting the ECS will interrupt services running on it. Exercise caution when performing this operation.

Possible Causes

The **/etc/udev/rules.d/70-persistent-net.rules** file still records the mappings between NICs and MAC addresses in the private image. You need to disable the rules.

Solution

1. The correct way to disable **udev** is to overwrite it with an empty file. Any rules in **/etc/udev/rules.d** take precedence over the rules in **/lib/udev/rules.d**. Run the following command to create an empty file:

```
touch /etc/udev/rules.d/75-persistent-net-generator.rules
```

2. Restart the ECS.

```
reboot
```

5.3 Why an Entry Is Automatically Added to /etc/hosts After a Linux ECS Is Restarted?

Symptom

After an ECS is restarted, an entry that maps its host name to 127.0.0.1 is automatically added to `/etc/hosts`. As a result, the host name fails to be resolved.

Root Cause

`/etc/cloud/cloud.cfg` affects the following configuration in `/etc/hosts`:

```
manage_etc_hosts: localhost
```

This configuration automatically translates the host names to the local loopback address, which can speed up startup when the private DNS is not configured.

Solution

1. Open the `/etc/cloud/cloud.cfg` file.

Comment out `manage_etc_hosts`:

Before:

```
manage_etc_hosts:localhost
```

After:

```
#manage_etc_hosts:localhost
```

2. Delete `127.0.0.1 hostname hostname` from `/etc/hosts`.

5.4 How Do I Fix a Network Startup Failure Due to Multiple NIC Configuration Files?

Symptom

Error message "Device eth1 does not seem to be present" or "No suitable device found for this connection" is displayed after the network is started or restarted.

Figure 5-1 Network startup failure

```
exiting.  
failed.  
  
[FAILED]  
Bringing up interface eth1: Device eth1 does not seem to be present, delaying initialization.  
[FAILED]  
Bringing up interface eth2: Device eth2 does not seem to be present, delaying initialization.  
[FAILED]  
Bringing up interface eth3: Device eth3 does not seem to be present, delaying initialization.  
[FAILED]  
Bringing up interface eth4: Device eth4 does not seem to be present, delaying initialization.  
[FAILED]
```

Scenarios

The following procedures apply to ECSs running CentOS, Red Hat, or EulerOS.

Constraints

The solution described in this section involves restarting the NICs, which will temporarily interrupt the network connection.

Possible Causes

When the network service is started, the system reads the NIC configuration file in the `/etc/sysconfig/network-scripts/` directory. If there are multiple NIC configuration files, the network service fails to be started because the corresponding NIC cannot be found.

Solution

Back up unnecessary NIC configuration files and delete them from `/etc/sysconfig/network-scripts/`. The following uses an ECS with 11 NIC configuration files as an example.

1. Run the following command to access the directory where the NIC configuration files reside:

```
cd /etc/sysconfig/network-scripts
```

2. Run the following commands to back up the configuration files:

```
mkdir tmp
```

```
cp ifcfg-* tmp/
```

```
ls tmp/
```

Figure 5-2 Viewing NIC configuration files

```
[root@cen73-test network-scripts]# ls
ifcfg-eth0  ifcfg-eth4  ifcfg-lo  ifdown-ipv6  ifdown-Team  ifup-eth  ifup-plusb  ifup-TeamPort
ifcfg-eth1  ifcfg-eth5  ifdown  ifdown-isdn  ifdown-TeamPort  ifup-ib  ifup-post  ifup-tunnel
ifcfg-eth10  ifcfg-eth6  ifdown-bnep  ifdown-post  ifdown-tunnel1  ifup-ippv  ifup-ppp  ifup-wireless
ifcfg-eth11  ifcfg-eth7  ifdown-eth  ifdown-ppp  ifup  ifup-ipv6  ifup-routes  init.ipv6-global
ifcfg-eth2  ifcfg-eth8  ifdown-ib  ifdown-routes  ifup-aliases  ifup-isdn  ifup-sit  network-functions
ifcfg-eth3  ifcfg-eth9  ifdown-ippv  ifdown-sit  ifup-bnep  ifup-plip  ifup-Team  network-functions-ipv6

[root@cen73-test network-scripts]# mkdir tmp
[root@cen73-test network-scripts]# cp ifcfg-* tmp/
[root@cen73-test network-scripts]# ls tmp/
ifcfg-eth0  ifcfg-eth10  ifcfg-eth2  ifcfg-eth4  ifcfg-eth6  ifcfg-eth8  ifcfg-lo
ifcfg-eth1  ifcfg-eth11  ifcfg-eth3  ifcfg-eth5  ifcfg-eth7  ifcfg-eth9
```

3. If only one NIC is needed, delete unnecessary NIC configuration files and file `ifcfg-ens5` if it exists.

Run the following command to delete files from `ifcfg-eth1` to `ifcfg-eth11` and `ifcfg-ens5`:

```
rm -rf ifcfg-eth[1-9] ifcfg-eth10 ifcfg-eth11 ifcfg-ens5
```

Figure 5-3 Deleting unnecessary configuration files

```
[root@cen72 network-scripts]# rm -rf ifcfg-eth[1-9] ifcfg-eth10 ifcfg-eth11 ifcfg-ens5
[root@cen72 network-scripts]# ls
ifcfg-eth0  ifdown-eth  ifdown-post  ifdown-tunnel  ifup-eth  ifup-plip  ifup-routes  init.ipv6-global
ifcfg-lo  ifdown-ippv  ifdown-ppp  ifup  ifup-ippv  ifup-plusb  ifup-sit  network-functions
ifdown  ifdown-ipv6  ifdown-routes  ifup-aliases  ifup-ipv6  ifup-post  ifup-tunnel  network-functions-ipv6
ifdown-bnep  ifdown-isdn  ifdown-sit  ifup-bnep  ifup-isdn  ifup-ppp  ifup-wireless
```

4. Stop unnecessary dhclient processes.
 - a. Run the following command to query dhclient processes:

```
ps -ef | grep dhclient
```
 - b. For example, run the following command to end process with PID 770.

```
kill -9 770
```

CAUTION

- Enter the PID of the process you want to end.
- The **kill -9 PID** command is used to forcibly stop a process.

Figure 5-4 Stopping dhclient

```
[root@cen73-test network-scripts]# ps -ef | grep dhcl
root      770      1  0 15:01 ?        00:00:00 /sbin/dhclient -H ecs-centos73-test -1 -q -lf /var/lib/dhclient/dhclient--eth0.lease -pf /var/run/dhclient-eth0.pid eth0
root      2622  2592  0 15:10 ttj1    00:00:00 grep --color=auto dhcl
[root@cen73-test network-scripts]# kill -9 770
[root@cen73-test network-scripts]#
```

5. Run the following command to restart the network service:
systemctl restart network
6. Run the following command to check the status of the network service:
systemctl status network

Figure 5-5 Checking the network status

```
[root@cen72 ~]# systemctl status network
network.service - LSB: Bring up/down networking
Loaded: loaded (/etc/rc.d/init.d/network; bad; vendor preset: disabled)
Active: active (running) since Tue 2018-02-06 10:54:58 CST; 34s ago
Docs: man:systemd-sysv-generator(8)
Process: 522 ExecStart=/etc/rc.d/init.d/network start (code=exited, status=0/SUCCESS)
CGroup: /system.slice/network.service
└─816 /sbin/dhclient -H cen72 -1 -q -lf /var/lib/dhclient/dhclient--eth0.lease -pf /var/run/dhclient-eth0.pid
Feb 06 10:54:55 cen72.novalocal systemd[1]: Starting LSB: Bring up/down networking...
Feb 06 10:54:55 cen72.novalocal network[522]: Bringing up loopback interface: [ OK ]
Feb 06 10:54:55 cen72.novalocal network[522]: Bringing up interface eth0:
Feb 06 10:54:55 cen72.novalocal dhclient[684]: DHCPREQUEST on eth0 to 255.255.255.255 port 67 (xid=0x7b202a2e)
Feb 06 10:54:55 cen72.novalocal dhclient[684]: DHCPACK from 192.168.3.2 (xid=0x7b202a2e)
Feb 06 10:54:58 cen72.novalocal NET[888]: /usr/sbin/dhclient-script : updated /etc/resolv.conf
Feb 06 10:54:58 cen72.novalocal dhclient[684]: bound to 192.168.3.234 -- renewal in 37342 seconds.
Feb 06 10:54:58 cen72.novalocal network[522]: Determining IP information for eth0... done.
Feb 06 10:54:58 cen72.novalocal network[522]: [ OK ]
Feb 06 10:54:58 cen72.novalocal systemd[1]: Started LSB: Bring up/down networking.
```

5.5 Why Do I Get the Error "Name or service not known" When I Ping a Public Domain Name Configured for a Linux ECS?

Symptom

A public domain name configured for an ECS fails to be pinged, and the error message "Name or service not known" is displayed. However, the EIP of the ECS can be pinged.

Possible Causes

Generally, there are three possible causes:

- No or incorrect DNS server addresses are configured in **/etc/resolv.conf**.
- The DNS records are deleted from **/etc/nsswitch.conf**.
- The **/lib64/libnss_dns.so.2** library file is lost.

 NOTE

Run the following command to view all files used for resolving the domain name:

```
strace -e trace=open ping www.baidu.com -c 1
```

All files in the output affect domain name resolution.

Scenarios

The operations described in this section apply to ECSs running CentOS and EulerOS.

Solution

- No or incorrect DNS server addresses are configured in **/etc/resolv.conf**.
nameserver is the most important item in **/etc/resolv.conf**. A **nameserver** entry defines the IP address of the DNS server used for domain name resolution.
If there are no **nameserver** entries in **/etc/resolv.conf**, no DNS servers have been configured for domain name resolution. If there are multiple **nameserver** entries, the DNS servers are queried in the order listed in the file. The next DNS server is queried only when the previous one does not respond. Check the IP addresses of DNS servers configured in **/etc/resolv.conf**.
- The DNS records are deleted from **/etc/nsswitch.conf**.
 - a. Check whether **/etc/nsswitch.conf** contains the DNS records.
grep hosts /etc/nsswitch.conf
If the output is as follows, the DNS option is not configured on the **hosts** line. As a result, the system does not read **/etc/resolv.conf** when resolving the domain name.

```
#hosts: db files nisplus nis dns
hosts: files myhostname
```
 - b. Open **/etc/nsswitch.conf**, locate the **hosts** line, and add the DNS option.

```
#hosts: db files nisplus nis dns
hosts: files dns myhostname
```

 NOTE

hosts lists the tools by priority that are used to search for IP addresses paired with domain names.

file indicates the **/etc/hosts** file, and **dns** indicates DNS. By default, **file** is placed before **dns**. This means that the system first attempts to search for a domain name in **/etc/hosts** and then search for the domain name through DNS. If **dns** is not configured, DNS is not used.

- The **/lib64/libnss_dns.so.2** library file is lost.
 - a. **/lib64/libnss_dns.so.2** is generated by the **glibc** package. Run the following command to check whether the package is modified:

```
rpm -V glibc
```

 NOTE

Generally, in Linux, running the **rpm -qf /lib64/libnss_dns.so.2** command can generate the library file.

If the output is as follows, **/lib64/libnss_dns.so.2** is lost.

```
missing /lib64/libnss_dns.so.2
```

- b. Run the following command to create a soft link again:

NOTE

Run the `ls -l /lib64/libnss_dns.so.2` command on a normal ECS. The command output will show that the source file of `/lib64/libnss_dns.so.2` is `/usr/lib64/libnss_dns-2.17.so`.

```
ln -s /usr/lib64/libnss_dns-2.17.so /usr/lib64/libnss_dns.so.2
```

5.6 Why Cannot the EIP Bound to the Extension NIC of My ECS Access the Internet?

Symptom

Your ECS has one primary NIC and one extension NIC in the same subnet. Both the NICs have an EIP bound to access the Internet. The EIP bound to the primary NIC can access the Internet, but that bound to the extension NIC cannot.

Possible Causes

By default, ECSs running CentOS have the reverse path filtering (RP-Filter) enabled. The default route of the ECSs is to forward outgoing traffic through the extension NIC to eth0. However, the system considers that the response data packets should be forwarded from eth1. The system determines that the traffic is received from a wrong NIC and then discards the response packets.

Solution

Configure a policy-based routing rule so that the extension NIC traffic is forwarded from the extension NIC.

1. Run the following command to edit the `rt_tables` file:

```
vi /etc/iproute2/rt_tables
```

Add an alias for the routing table, such as `test`.

```
#
255    local
254    main
253    default
0      unspec
32000  test
#
```

2. Save the modification and exit.
3. Run the following command to add a route to the `test` table:

```
ip route add default via Gateway IP address of the extension NIC dev eth1 table Name of the routing table
```

For example, run the following command:

```
ip route add default via 192.168.166.1 dev eth1 table test
```

4. Run the following command to add a policy-based routing rule:

```
ip rule add from IP address of the extension NIC lookup Name of the routing table prio lower than 32766 but higher than the main table
```

For example, run the following command:

```
ip rule add from 192.168.166.22 lookup test prio 32000
```

Check whether the EIP bound to the extension NIC can access the Internet. If you want to make this rule take effect permanently, add the preceding command to the startup script `/etc/rc.local`.

5.7 How Do I Fix Too High Memory Usage by NetworkManager When Multiple Docker Containers Are Running?

Symptom

NetworkManager consumes a large amount of memory when multiple Docker containers are running.

```
top - 14:32:53 up 220 days, 23:42, 1 user, load average: 4.06, 3.70, 13.74
Tasks: 595 total, 3 running, 592 sleeping, 0 stopped, 0 zombie
%Cpu(s): 5.6 us, 5.7 sy, 0.0 ni, 88.2 id, 0.4 wa, 0.0 hi, 0.1 si, 0.0 st
KiB Mem : 49459224 total, 512644 free, 47312020 used, 1634560 buff/cache
KiB Swap: 0 total, 0 free, 0 used, 555936 avail Mem

  PID USER      PR  NI   VIRT   RES   SHR  S  %CPU  %MEM    TIME+  COMMAND
 1023 root        20   0 22.616g 0.022t 5424 R  94.4  47.1 14980.11 /usr/sbin/NetworkManager --no-daemon
18518 root        20   0 6714816 1.828g 3468 S   5.0   3.5   2:01.47 /usr/lib/jvm/java-1.8-openjdk/bin/java -serv
23130 root        20   0 14.011g 1.767g 2208 S   0.0   3.7   8:46.57 java -jar /opt/maintenance.jar
2600 root        20   0 14.043g 1.698g    0 S   0.0   3.6 17:22.01 java -jar /opt/alarm.jar
1930 root        20   0 6719112 1.598g 2164 S   0.0   3.4 50:45.07 /usr/lib/jvm/java-1.8-openjdk/bin/java -serv
```

NOTE

- The operations described in this section apply to ECSs running CentOS 7 or Ubuntu 16.04.
- The operations described in this section involve restarting the network service, which may interrupt services.

Possible Cause

The amount of memory consumed by **NetworkManager** increases with each container startup or stop, it never decreases, even after all containers have been stopped and removed.

Solution

Short-term solution

Run the following command to restart **NetworkManager**:

```
systemctl restart NetworkManager
```

Long-term solution

- CentOS 7

Run the following commands to stop **NetworkManager** and start **network**:

```
systemctl disable NetworkManager
/sbin/chkconfig network on
kill `pgrep -o dhclient`
systemctl stop NetworkManager
systemctl start network
```

 NOTE

If **network** startup fails, it may be caused by the built-in configuration files of multiple NICs. For details about how to rectify the fault, see [How Do I Fix a Network Startup Failure Due to Multiple NIC Configuration Files?](#)

- Ubuntu 16.04

a. Run the following commands to stop **NetworkManager** and start **networking**:

```
systemctl disable NetworkManager
systemctl disable network-manager
systemctl enable networking
kill `pgrep -o dhclient`
systemctl stop NetworkManager
systemctl start networking
```

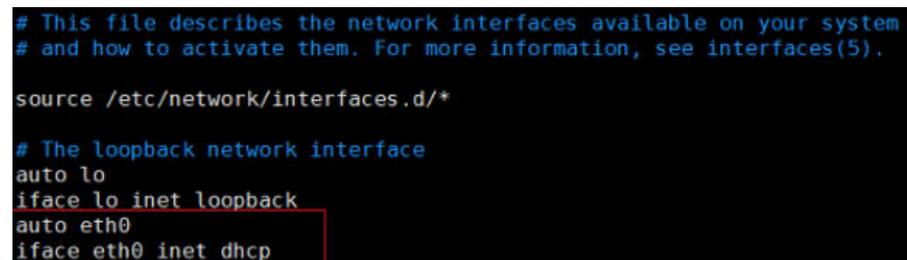
b. To enable **networking**, check whether DHCP is configured for NICs in the **interfaces** file.

```
vi /etc/network/interfaces
```

If only **eth0** is attached, you can check whether the following information exists in **interfaces**. If not, add it to the file.

```
auto eth0
iface eth0 inet dhcp
```

Figure 5-6 Configuring the NIC to use DHCP



```
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

source /etc/network/interfaces.d/*

# The loopback network interface
auto lo
iface lo inet loopback
auto eth0
iface eth0 inet dhcp
```

5.8 Why Is the ECS IP Address Lost After the System Time Changes?

Symptom

After the system time changes, an exclamation mark (!) is displayed for the NIC, and the ECS cannot connect to the Internet.

Scenarios

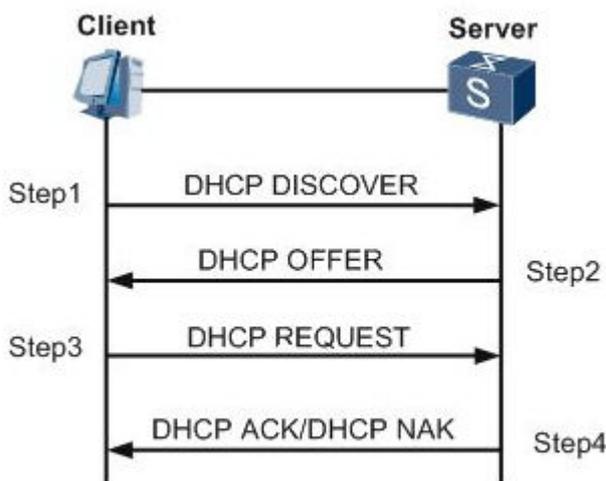
- The operations described in this section apply to ECSs running CentOS 7 or EulerOS that use DHCP to obtain IP addresses.
- The default DHCP lease time is 24 hours. Specific scenarios may be different.

Possible Causes

DHCP session process

A typical DHCP session includes DHCP discover, DHCP offer, DHCP request, and DHCP acknowledgment, as shown in the following figure.

Figure 5-7 DHCP session



- DHCP Discover
The DHCP client sends a broadcast message on the physical subnet to search for available DHCP servers. The network administrator can configure a local route to forward DHCP packets to DHCP servers on another subnet. The client generates a UDP packet whose destination address is 255.255.255.255 or the specific subnet broadcast address.
- DHCP Offer

When a DHCP server receives an IP address lease request from the client, the DHCP server provides an IP address lease and reserves an IP address for the client, and then sends a DHCPOFFER message to the client through unicast. This message contains the MAC address of the client, the IP address that the server is providing, the subnet mask, the lease duration, and the IP address of the DHCP server making the offer.

- DHCP Request

In response to the DHCP offer, the client replies with a DHCPREQUEST message that contains the IP address of the server making the offer, which is a broadcast informing the server that it accepts the offer. When other DHCP servers receive this message, they withdraw any offers they have made to the client and return the offered IP addresses to the pool of available addresses. A client can receive DHCP offers from multiple servers, but it will accept only one DHCP offer.

- DHCP Acknowledge (DHCP ACK)

When the DHCP server receives the DHCPREQUEST message from the client, the configuration process enters its final phase. The acknowledgement phase involves sending a DHCPACK packet to the client. This packet contains the lease duration and any other configuration information that may be requested by the client. At this point, the TCP/IP address configuration is complete.

Fault Locating

1. View the DHCP request records of the client.

```
grep -E "dhclient|DHCP" /var/log/messages
```

Figure 5-8 DHCP request records

```
May 12 11:59:39 yth5_test_sl dhclient[7376]: DHCPREQUEST on eth0 to 10.10.3.254 port 67 (xid=0x1880414e)
May 12 11:59:39 yth5_test_sl dhclient[7376]: DHCPACK from 10.10.3.254 (xid=0x1880414e)
May 12 11:59:39 yth5_test_sl dhclient[7376]: bound to 10.10.3.240 -- renewal in 33696 seconds.
```

- The records show that **dhclient** renewed the IP address on May 12. Since the system was manually changed, the actual time was April 26.
- The next renewal occurred 33,696 seconds (about 9 hours and 21 minutes) later, usually when half of the lease time has passed. That is, the next renewal occurred at around 21:21 on May 12.

NOTICE

You can also query the **dhclient** renewal information file for each successful renewal. Run the following command to query the path that saves the file. The file name is ended with **.lease**.

```
ps -ef |grep dhclient
```

2. Run the following command to check the system logs. The logs show that the system time was changed 5 hours and 14 minutes later after the renewal, and changed to April 26. This means the next renewal request would be sent in 16 days, but not 9 hours it should be. **dhclient** failed to renew the lease and just silently let the lease expire. After the lease expired, the IP address was withdrawn by the DHCP server. As a result, the IP address was lost.

```
grep "Time has been changed" /var/log/messages
```

Figure 5-9 System logs

```
May 12 17:05:32 yth5_test_sl systemd-logind: Removed session 2508.  
May 12 17:13:16 yth5_test_sl systemd-logind: New session 2509 of user root.  
May 12 17:13:16 yth5_test_sl systemd: Started Session 2509 of user root.  
May 12 17:13:16 yth5_test_sl systemd: Starting Session 2509 of user root.  
Apr 26 17:49:38 yth5_test_sl systemd: Time has been changed  
Apr 26 17:53:50 yth5_test_sl systemd-logind: Removed session 2509.  
Apr 26 17:58:59 yth5_test_sl systemd-logind: New session 2510 of user rundeck.  
Apr 26 17:58:59 yth5_test_sl systemd: Started Session 2510 of user rundeck.
```

Solution

1. Run **dhclient** manually to obtain the private IP address.

```
dhclient -r eth0  
ifconfig eth0 down  
ifconfig eth0 up  
dhclient eth0
```



CAUTION

This operation may interrupt network traffic.

2. Restart the ECS. The system time will restore and the system will obtain the IP address again.
3. Enable time synchronization with an NTP server. For details, see [Does Huawei Cloud Provide the NTP Server and How Can I Configure It?](#)

5.9 What Can I Do If resolv.conf Gets Reset?

Symptom

The **resolv.conf** file you have modified gets reset after you restart the ECS.

Solution

Add the **i** attribute to the file, which allows only user **root** to modify the file.

```
chattr +i /etc/resolv.conf
```

5.10 What Can I Do If /etc/resolv.conf Is Restored After an ECS Running Ubuntu Is Restarted?

Symptom

After an ECS running Ubuntu or network-related services are restarted, the **/etc/resolv.conf** file is updated, and the **nameserver** field is restored to **127.0.0.53**.

Figure 5-10 Symptom

```
root@ecs-~# cat /etc/resolv.conf
# Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
#     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
# 127.0.0.53 is the systemd-resolved stub resolver.
# run 'systemd-resolve --status' to see details about the actual nameservers.

nameserver 127.0.0.53
search abc.com
options timeout:1 single-request-reopen
```

Possible Causes

By default, Ubuntu uses systemd-resolved service for DNS maintenance. If ECSs or network-related services are restarted, the IP address of **nameserver** in **/etc/resolv.conf** is updated to **127.0.0.53**.

Solution

NOTE

You are advised to disable the systemd-resolved service before handling the issue.

Solution 1: Manually modify the **/etc/resolv.conf** file.

1. Log in to the ECS as user **root**.
2. Disable the systemd-resolved service.
systemctl stop systemd-resolved
systemctl disable systemd-resolved
3. Delete the default soft link in **/etc/resolv.conf** and create it as a common file.
rm -rf /etc/resolv.conf
4. Edit **/etc/resolv.conf** and add related DNS configuration to it.

vim /etc/resolv.conf

Configure the **nameserver** field to add the DNS configuration as follows:

```
nameserver 100.125.1.250
nameserver 100.125.129.250
```

5. Lock **/etc/resolv.conf** to prevent it from being modified by DHCP or other services.

chattr +i /etc/resolv.conf

Solution 2: Use NetworkManager to maintain the **/etc/resolv.conf** file based on the DNS information (configured in the VPC subnet) obtained through DHCP.

1. Disable the systemd-resolved service.
systemctl stop systemd-resolved
systemctl disable systemd-resolved
2. Edit the NetworkManager configuration file and add the **dns=default** configuration.

vim /etc/NetworkManager/NetworkManager.conf

```
[main]
plugins=ifupdown,keyfile
dns=default

[ifupdown]
```

```
managed=true  
[device]  
wifi.scan-rand-mac-address=no
```

3. Delete the default soft link in **/etc/resolv.conf** and create it as a common file.
rm -rf /etc/resolv.conf
4. Restart NetworkManager and update the **/etc/resolv.conf** file.
systemctl restart NetworkManager
5. Check the DNS configuration in the **/etc/resolv.conf** file.
 - If the DNS configuration is the same as that of the subnet which the ECS belongs to, the modification is successful.
You can log in to the ECS console and click the target ECS name. On the ECS details page, click the primary NIC name in the **NICs** area to switch to the subnet console. On the displayed page, check the **DNS Server Address** in the **Gateway and DNS Information** area.
 - If they are inconsistent, submit a service ticket to contact technical support.

6 Disk Space Management Issues

6.1 Why Can't I Mount a Disk on an Old Mount Point by Modifying `fstab` in CentOS 7?

Symptom

File `/etc/fstab` was modified to allow a new disk to be mounted on an old mount point. The mounting failed when queried using the `df` command.

The operations described in this section apply to ECSs running CentOS or EulerOS.

Possible Causes

1. Run the following command to query the involved mount unit:

```
systemctl list-units --type=mount |grep failed
```

```
test1.mount          loaded failed failed /test1
```

2. Run the following command to query the unit status:

```
systemctl status test1.mount
```

Information similar to the following is displayed:

```
● test1.mount - /test1
Loaded: loaded (/etc/fstab; bad; vendor preset: disabled)
Active: failed (Result: exit-code) since Wed 2019-08-28 15:32:53 CST; 3min 27s ago
Where: /test1
What: /dev/vdb1
Docs: man:fstab(5)
man:systemd-fstab-generator(8)
Process: 4601 ExecUnmount=/bin/umount /test1 (code=exited, status=0/SUCCESS)
Process: 3129 ExecMount=/bin/mount /dev/vdb1 /test1 -t ext4 (code=exited, status=0/SUCCESS)
... ..
Warning: test1.mount changed on disk. Run 'systemctl daemon-reload' to reload units.
```

In the command output, mount unit **test1.mount** has changed. Run **systemctl daemon-reload** to reload units.

The changes in `/etc/fstab` do not get applied automatically. You must run the **systemctl daemon-reload** command to update the mount units generated for each entry in `fstab`. `systemd` does not read the `fstab` file and generate a load unit.

Solution

Run the following command to reload the mount units managed by systemd:

```
systemctl daemon-reload
```

6.2 How Do I Create a Swap Partition or File in Linux?

Scenarios

This section describes how to create a swap partition on ECS running CentOS 6.8.

Constraints

A file of a specified size is to be created. Ensure that the system disk has enough available space.

Scenario 1: Creating a Swap Partition on a Block Storage Device

1. Run the following command to create a partition of 2 GB, for example:

```
fdisk /dev/vdb
```

Information similar to the following is displayed:

```
Command (m for help): n
Partition type:
  p   primary (0 primary, 0 extended, 4 free)
  e   extended
Select (default p):
Using default response p
Partition number (1-4, default 1):
First sector (2048-20971519, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-20971519, default 20971519): +2G
Partition 1 of type Linux and of size 2 GiB is set
Command (m for help): p
```

```
Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x1f02f438
```

Device	Boot	Start	End	Blocks	Id	System
/dev/vdb1		2048	4196351	2097152	83	Linux

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

2. Run the following command to configure the newly created partition as swap space:

```
mkswap /dev/vdb1
```

3. Run the following command to activate the swap partition:

```
swapon /dev/vdb1
```

4. Run the following command to verify the activated swap:

swapon -s

5. Run the following command to obtain the swap partition UUID:

```
blkid |grep swap |awk '{print $2}'
```

```
UUID="1ee90e3c-1538-453b-9240-ad430f835f6f"
```

6. To mount the swap partition automatically upon system startup, add an entry for the swap partition to **/etc/fstab**.

Run the following command (the UUID is obtained from 5):

```
echo "UUID=1ee90e3c-1538-453b-9240-ad430f835f6f swap swap defaults 0 0" >>/etc/fstab
```

7. Run the following command to mount the swap partition:

```
mount -a
```

Scenario 2: Creating a Swap Partition on a Block Storage Device Simulated by a File

 NOTE

The performance of the block storage device simulated by a file is not as good as that of the passthrough block storage device.

1. Run the following command to create a file of 1 GB, for example:

```
dd if=/dev/zero of=/swapfile bs=1M count=1000
```
2. Run the following command to modify the file permissions:

```
chmod 600 /swapfile
```
3. Run the following command to configure the file as swap space:

```
mkswap /swapfile
```
4. Run the following command to activate the swap file:

```
swapon /swapfile
```
5. To mount the swap partition automatically upon system startup, add an entry for the swap file to **/etc/fstab**.

```
echo "/swapfile swap swap defaults 0 0" >>/etc/fstab
```
6. Run the following command to mount the swap partition:

```
mount -a
```

6.3 Why Is the Space Not Released After I Delete a Large File on a Linux ECS?

Symptom

The space usage of the root directory on a Linux ECS is too high, for example, 96%, as shown in [Figure 6-1](#).

Figure 6-1 Too high space usage of the root directory

```
[root@host1 /]#df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2       49G   47G   2G   96% /
tmpfs           16G    0   16G   0% /dev/shm
```

You find that the **access_log** file generated by Apache occupies about 42 GB in the root directory. Then you perform the following operations attempting to release space:

1. Delete **access_log**.
rm /tmp/access_log
2. Check the file system disk space usage.
df -h

However, the output tells you that the usage is still 96%.

Possible Causes

The file you deleted is locked by another running process that keeps writing data to the file.

The following are basics about file storage in Linux, which can help you better understand the possible cause.

In Linux file systems, a file is organized into two parts:

- File pointer: The pointer is stored in the file system metadata. When the file is deleted, the file pointer will be deleted from the metadata.
- File data: The file data is stored on the disk.

Generally, after the file pointer is deleted from the metadata, the disk space occupied by the file data will be marked as available. However, if the file is still being used by another running process when it is deleted, the file pointer will fail to be deleted from the metadata, and the system will determine that the file has not been deleted and does not reclaim the space occupied by the file data.

Solution

1. Run the **lsof** command to check whether any process keeps writing data to the **access_log** file.

```
lsof -n |grep delete
```

Figure 6-2 Checking the process that locks the file

```
[root@host1 /]#lsof -n|grep delete

COMMAND PID  USER  FD  TYPE  DEVICE  SIZE  NODE  NAME
httpd    2660  root   4u  REG   8,2     43289765 /access_log (deleted)
```

As shown in the command output, the `/tmp/access_log` file is still being used by the `httpd` process which keeps writing log data into the file. Value **deleted** in the brackets indicates that the log file has been deleted. But the space is not released because `httpd` keeps writing data to the file.

2. Stop or restart the `httpd` process, or restart the system. You are advised to clear the content of `access_log`, rather than deleting it.

Run the following command to clear `access_log`:

```
echo "">/access_log
```

In this way, the disk space will be released immediately and the process can continue writing logs to the file. Then, run the `df -h` command to check the space usage again.

Figure 6-3 Checking the space usage of the root directory

```
[root@host1 /]#df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda2       49G   5G   44G  10% /
tmpfs           16G   0    16G   0% /dev/shm
```

6.4 What Should I Do If the Error Message "Read-only file system" Is Displayed When I Attempt to Delete a File on a Linux ECS?

Symptom

When you attempt to delete or modify a file on a Linux ECS, the message **Read-only file system** is displayed.

```
[root@oss-vpc-network-om-elk-20 logstash]# rm -rf *2019*
rm: cannot remove 'logstash-plain-2019-01-18-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-19-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-20-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-21-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-22-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-23-1.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-23-2.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-23-3.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-23-4.log.gz' : Read-only file system
rm: cannot remove 'logstash-plain-2019-01-23-5.log.gz' : Read-only file system
```

Possible Causes

The possible causes are as follows:

- A file system error caused the file system to be read-only.
- The file system is mounted as read-only.
- The hardware is faulty, for example, the disk contains bad sectors or the RAID controller card is faulty.

Precautions

- Repairing file systems may cause data loss. Back up data in advance.
- If the issue you encountered does not fall into either of the following scenarios, check whether the hardware is faulty.

Scenario 1: File System Mounted as Read-only

1. Run the following command to check how the directory containing the to-be-deleted file is mounted:

```
mount |grep mount-point
```

If **ro** is displayed in the output, the directory is mounted as read-only. Then go to [2](#).

If **rw** is displayed in the output, the directory is mounted as read/write. Then perform procedures described in [Scenario 2: File System Error on Read/Write Mount](#) to check whether the file system contains errors.

2. Run the following command to remount the file system as read/write. There is no need to restart the system.

```
mount -o remount,rw mount-point
```

NOTE

To mount the file system as read/write upon the next startup, modify the parameters in the fourth column in the `/etc/fstab` file.

Scenario 2: File System Error on Read/Write Mount

1. Run the following command to check the file system information in the kernel:

```
dmesg |grep "ext[2..4]|xfs"
```

If the error message "I/O error... inode" is displayed, there is a file system error.

```
[root@oss-vm-network-dm-elk-20 software]# dmesg |grep "ext[2..4]|xfs"
[  8.393974] CPU: 0 PID: 410 Comm: fsck.ext4 Not tainted 3.10.0-327.el2.x86_64 #1
[ 13.058706] EXT4-fs warning (device dm-1): ext4_end_bio:332: I/O error -5 writing to inode 393232 (offset 2101248 size 4096 starting block 3993)
[ 46.191190] EXT4-fs warning (device dm-1): ext4_end_bio:332: I/O error -5 writing to inode 22 (offset 0 size 8192 starting block 40448)
[ 46.216028] EXT4-fs warning (device dm-1): ext4_end_bio:332: I/O error -5 writing to inode 17 (offset 0 size 4096 starting block 40450)
[ 46.216101] EXT4-fs warning (device dm-1): ext4_end_bio:332: I/O error -5 writing to inode 1179650 (offset 0 size 4096 starting block 40451)
[ 136.297714] EXT4-fs warning (device dm-1): ext4_end_bio:332: I/O error -5 writing to inode 17 (offset 0 size 0 starting block 40450)
```

Back up data before repairing the file system. Since file systems cannot be repaired when they are in use, switch to the single-user mode to unmount all file systems and then repair the file systems.

2. Restart the system and switch to the single-user mode.
3. Query all available devices and file systems.

blkid

4. Check the file systems on device **vdb1**.
 - For the ext file system, run the following command:

```
fsck -n /dev/vdb1
```

- For the xfs file system, run the following command:

```
xfs_check /dev/vdb1
```

NOTE

If the output indicates that the file systems are mounted, unmount these file systems.

1. Run the following command to query the file systems mounted on the device:

```
mount
```

2. Run the following command to unmount these file systems:

```
umount mount-point
```

5. Repair the file systems on device **vdb1**.
 - For the ext file system, run the following command:
fsck /dev/vdb1
 - For the xfs file system, run the following command:
xfs_repair /dev/vdb1

6.5 How Do I Fix File Creation Failures Due to Inode Exhaustion?

Symptom

When you create a file or directory, you get the error message "No space left on device", "Cannot create directory", or "Couldn't create temporary archive name."

Possible Causes

In Linux, both of the following use the disk space:

- Data
- Inode

NOTE

An inode (index node) stores metadata of a file system object, such as a file, directory, device file, socket, and pipe in the file system, but does not contain the object's data and filename.

Solution

1. Run the following command to check the available space on the disk:

```
df -h
```

Figure 6-4 Checking the available disk space

```
mkdir: cannot create directory 'test': No space left on device
[root@ecs-bbec data]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/vda2       ext4      42G   4.5G   35G  12% /
devtmpfs        devtmpfs  17G    0    17G   0% /dev
tmpfs           tmpfs     17G   4.1k   17G   1% /dev/shm
tmpfs           tmpfs     17G   9.0M   17G   1% /run
tmpfs           tmpfs     17G    0    17G   0% /sys/fs/cgroup
/dev/vda1       ext4      1.1G  153M  801M  16% /boot
tmpfs           tmpfs     3.4G    0   3.4G   0% /run/user/0
/dev/vdb1       ext4     106G  990M  100G   1% /home/gis
/dev/vdc1       ext4     106G   37G   64G  37% /home/data
tmpfs           tmpfs     3.4G    0   3.4G   0% /run/user/1000
```

As shown in [Figure 6-4](#), the disk still has enough available space.

2. Run the following command to check the available inodes on the disk:

df -i

If the value of **Use%** in the command output is **100%**, inodes are used up. Perform the following operations to free up inodes:

- a. Archive files in all directories.
tar czvf /tmp/backup.tar.gz /home/data
- b. Delete unnecessary files to release inodes.

6.6 Why Do I Get the Error "No space left on device" When I Create a File on a Linux ECS?

Symptom

When you create a file on a Linux ECS, the error message "No space left on device" is displayed.

Possible Causes

- The block usage on the disk reaches 100%.
- The inode usage on the disk reaches 100%.
- Certain disk space is not freed up because there are unreleased file handles.
- The value of **fs.inotify.max_user_watches** has been reached.

100% Block Usage

Run the following command to view the disk space usage:

df -h

If information similar to the following is displayed, the blocks have been used up.

Solution: Expand the capacity of the disk.

```
[root@ecs-linux-bj1 data1]# echo aaa > a
-bash: echo: write error: No space left on device
[root@ecs-linux-bj1 data1]# df -h .
Filesystem      Size  Used Avail Use% Mounted on
/dev/vdb1       89M   87M   0 100% /data1
[root@ecs-linux-bj1 data1]# df -i .
Filesystem      Inodes IUsed IFree IUse% Mounted on
/dev/vdb1       24480   13 24467   1% /data1
[root@ecs-linux-bj1 data1]#
```

100% Inode Usage

Run the following command to view the inode usage:

df -i

If information similar to the following is displayed, the inodes have been used up.

Solution: Expand the capacity of the disk.

```
[root@ecs-linux-bj1 data1]# touch bbb
touch: cannot touch 'bbb': No space left on device
[root@ecs-linux-bj1 data1]# df -i .
Filesystem      Inodes IUsed IFree IUse% Mounted on
/dev/vdb1       24480 24480    0 100% /data1
[root@ecs-linux-bj1 data1]# df -h .
Filesystem      Size  Used Avail Use% Mounted on
/dev/vdb1       89M   26M   56M  32% /data1
```

Space Occupied by Deleted Files

1. Run the **df -h** command to check whether the block usage reaches 100%.
2. Run the **df -i** command to check whether the inode usage is relatively low, for example **1%**, as shown in the following figure.
3. Run the **du -sh** command to view the total disk space used by files. You may notice that there is a significant difference between the usage space and available space.

```
[root@ecs-linux-bj1 data1]# echo aaaa >> a
-bash: echo: write error: No space left on device
[root@ecs-linux-bj1 data1]# df -h .
Filesystem      Size  Used Avail Use% Mounted on
/dev/vdb1       89M   87M    0 100% /data1
[root@ecs-linux-bj1 data1]# df -i .
Filesystem      Inodes IUsed IFree IUse% Mounted on
/dev/vdb1       24480    16 24464    1% /data1
[root@ecs-linux-bj1 data1]# du -sh .
6.9M .
```

Solution

1. Run the following command to list deleted files whose handles are still held open by processes:

lsdf |grep delete

```
[root@ecs-linux-bj1 data1]# lsdf |grep delete
cat      12803      root    1w     REG          253,17  81920030    11 /data1/testfile (deleted)
cat      13156      root    1w     REG          253,17  81920030    11 /data1/testfile (deleted)
[root@ecs-linux-bj1 data1]#
```

2. Run the following command to stop these processes one at a time:

kill -9 Process ID

```
[root@ecs-linux-bj1 data1]# kill -9 12803
[root@ecs-linux-bj1 data1]# kill -9 13156
[root@ecs-linux-bj1 data1]# df -h .
Filesystem      Size  Used Avail Use% Mounted on
/dev/vdb1       89M   8.4M  74M  11% /data1
[root@ecs-linux-bj1 data1]# echo aaaa >> a
```

Value of fs.inotify.max_user_watches Reached

If there is insufficient space for Linux ECSs, "No space left on device" will be displayed.

```
ubuntu@VM-0-11-ubuntu:~/work$ sudo service docker start
Failed to add /run/systemd/ask-password to directory watch: No space left on device
Failed to add /run/systemd/ask-password to directory watch: No space left on device
ubuntu@VM-0-11-ubuntu:~/work$ df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            414M   0  414M   0% /dev
tmpfs           87M   9.3M   78M  11% /run
/dev/vda1       50G   3.1G   44G   7% /
tmpfs           433M   24K  433M   1% /dev/shm
tmpfs           5.0M   0   5.0M   0% /run/lock
tmpfs           433M   0   433M   0% /sys/fs/cgroup
tmpfs           87M   0   87M   0% /run/user/500
```

Solution

1. Run the following command to edit the `/etc/sysctl.conf` file:

```
vi /etc/sysctl.conf
```

2. Add the following content to the file:

```
fs.inotify.max_user_watches = 524288
```

3. Run the following command for the modification to take effect:

```
sysctl -p
```

inotify is used to monitor file system events. By default, a maximum of 8,192 files can be watched for each real user ID. You can run the following command to obtain the current limit:

```
cat /proc/sys/fs/inotify/max_user_watches
```

```
[root@ecs-linux-bj1 data1]# cat /proc/sys/fs/inotify/max_user_watches
8192
[root@ecs-linux-bj1 data1]#
```

If the limit is too low to watch all files, increase the limit.

6.7 What Can I Do If the Buffer and Cache Occupy Too Much Memory of a Linux ECS?

Symptom

When querying the system memory usage using `free` after the Linux ECS runs for a long period, you have found that most of the memory is used by buffers and cache.

```
linux-nkx:/home/yb/test # free -m
              total        used         free       shared    buffers     cached
Mem:          48183         47343           840           0         4754      41996
-/+ buffers/cache:
-/+ buffers/cache: 592         47591
Swap:          0             0             0
```

Possible Cause

In Linux memory management, buffer refers to the buffer cache, and cache refers to the page cache.

- Buffer cache, which caches data that is read from and written to disk blocks
For example, if the system writes a file, respective page cache pages will be modified, and buffer cache buffers will be marked as dirty. When writing dirty data back to the disk, the kernel only writes back the changes, rather than the entire page, to the disk.
- Page cache, which caches file data that is read from and written to files in file systems. By default, Linux caches the read file data in the memory to facilitate quick data access.

Linux uses the lazy mode by default. Linux automatically releases buffers and cache if the system requires more memory, and a high cache memory usage does not affect the system performance.

Solution

Buffers and cache are necessary for the smooth running of systems and devices in Linux. If the cache is forcibly cleared, disk data needs to be read from the disk, which decreases the system performance.

NOTE

To clear the buffers and cache, run the following command:

```
echo 3 > /proc/sys/vm/drop_caches
```

The command execution may take several seconds, depending on the memory size. After the command is executed, the occupied memory will be released.

6.8 What Can I Do If the Partition Capacity Fails to Be Expanded Using `growpart` After the EVS Disk Capacity Is Expanded?

Symptom

After the capacity of an EVS disk is expanded, running the `growpart` command to expand the partition capacity fails. The error information is as follows:

Figure 6-5 Error message 1

```
failed [pt_update:1] pt_update /dev/vdb 1
partx: /dev/vdb: error updating partition 1
FAILED: disk=/dev/vdb partition=1: failed to repartition
***** WARNING: Resize failed, attempting to revert *****
Warning: The kernel is still using the old partition table
```

Figure 6-6 Error message 2

```
[root@~]# growpart /dev/sdb 1
[
NOCHANGE: disk=/dev/sdb partition=1: could only be grown by -1 [fudge=2048]
```

Possible Cause

After expanding the capacity of an EVS disk, you need to expand the capacity of its partitions and file system. If the partitions are not aligned during initialization, the partition capacity may fail to be expanded using the **growpart** command.

Run the **parted -l** command. If the value in the **Start** column is not **2048s** or **1049KB**, the partitions are not aligned.

Figure 6-7 Partitions not aligned

```
(parted) p
Model: Huawei UBS fileIO (scsi)
Disk /dev/sdb: 4320133120s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End          Size          File system  Name  Flags
  1      34s    4320133086s  4320133053s                test
```

Solution

If partitions are not aligned, their capacity cannot be automatically expanded. You need to manually expand the capacity or repartition the disk.

This section describes how to manually expand the capacity of disk partitions.

NOTE

Repartitioning disks will cause data loss. Exercise caution when performing this operation.

Manual capacity expansion may be risky, so you need to back up data before performing this operation. For details, see [Creating an EVS Snapshot](#) or [Creating a Cloud Server Backup](#).

1. Log in to a Linux ECS.
2. Stop the processes related to directory mounting.
3. Run the following command to uninstall **sdb1** (as an example):
umount /dev/sdb1
4. Run the following command to manually expand the capacity:
parted /dev/sdb
5. Run the **p** command to check the current partition.
6. Run the following command to select a partition for capacity expansion:
You must select the last partition because only its capacity can be expanded.
resizepart 1 100%(Partition 1 is used as an example in this command)
7. Run the **p** command to check whether the partition capacity is expanded.

Figure 6-8 Partition capacity expansion

```
(parted) p
Model: Huawei VBS fileIO (scsi)
Disk /dev/sdb: 2212GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      17.4kB  1106GB  1106GB

(parted) resizepart 1 100%
(parted) p
Model: Huawei VBS fileIO (scsi)
Disk /dev/sdb: 2212GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name  Flags
  1      17.4kB  2212GB  2212GB
```

8. Run the **q** command to exit the **parted** interaction mode. The partition capacity is manually expanded.

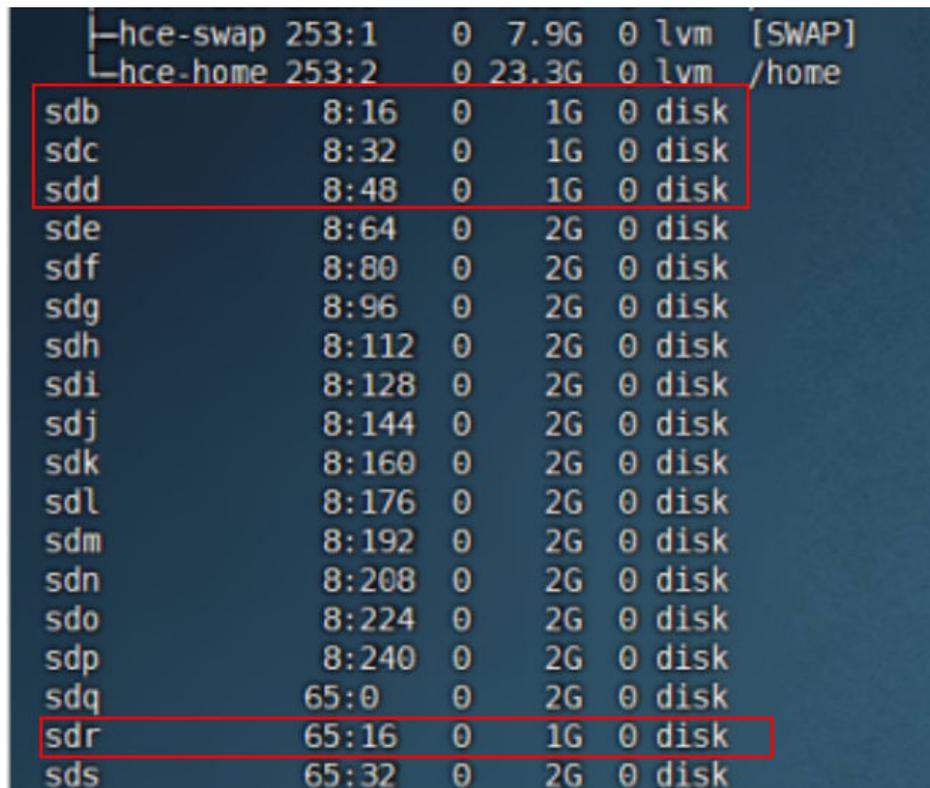
6.9 What Can I Do If Disk Scale-Out Fails When There Is Heavy I/O Workload for SCSI Disks?

Symptom

After the capacity of SCSI disks (more than 10) of Linux ECSs is concurrently expanded online, the capacity of some disks remains unchanged after the expansion.

For example, after you increase disk capacity from 1 GB to 2 GB, the output of the **lsblk** command shows that the disk capacity of **sdb/sdc/sdd/sdr** remains unchanged.

Figure 6-9 Error message



```
hce-swap 253:1 0 7.9G 0 lvm [SWAP]
hce-home 253:2 0 23.3G 0 lvm /home
sdb      8:16 0 1G 0 disk
sdc      8:32 0 1G 0 disk
sdd      8:48 0 1G 0 disk
sde      8:64 0 2G 0 disk
sdf      8:80 0 2G 0 disk
sdg      8:96 0 2G 0 disk
sdh      8:112 0 2G 0 disk
sdi      8:128 0 2G 0 disk
sdj      8:144 0 2G 0 disk
sdk      8:160 0 2G 0 disk
sdl      8:176 0 2G 0 disk
sdm      8:192 0 2G 0 disk
sdn      8:208 0 2G 0 disk
sdo      8:224 0 2G 0 disk
sdp      8:240 0 2G 0 disk
sdq      65:0 0 2G 0 disk
sdr      65:16 0 1G 0 disk
sds      65:32 0 2G 0 disk
```

Possible Cause

When the I/O workload of SCSI disks is heavy, the disk queue cannot process the capacity expansion requests in a timely manner, so the capacity of some disks is not changed.

If a write I/O operation is triggered on the disk, the Linux ECS will check the capacity again to make the expansion take effect.

Solution

On the Linux ECS, perform a write I/O operation on the disk whose capacity remains unchanged to make the expansion take effect. For example:

Run the following command to create an empty file in the mount point directory of the disk and perform a write I/O operation:

```
touch file
```

6.10 Handling Insufficient Disk Space on a Windows ECS

Scenarios

When the EVS disk space on a server is insufficient, the server running speed will be affected, which will further affect user experience. You can clean up the disk space using either of the following methods:

- [Cleaning Up Disk Space Using the System Built-in Cleanup Tool](#)
- [Uninstalling Unnecessary Programs from the Control Panel](#)

The following example uses a Windows Server 2016 server to illustrate how to clean up disk space on a server. In addition, consider the following practices to maintain sufficient disk space over time:

- Periodically compress and archive files that are not frequently used.
- Regularly run the disk cleanup tool to delete unnecessary files and empty the recycle bin.
- Uninstall unnecessary programs to keep the system lean and efficient.

Cleaning Up Disk Space Using the System Built-in Cleanup Tool

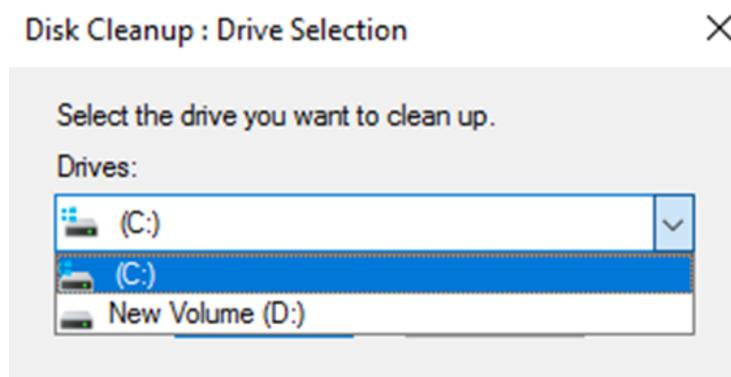
Step 1 On the server desktop, click the start icon in the lower left corner.

The start menu is displayed.

Step 2 In the navigation pane on the left, choose **Windows Administrative Tools > Disk Cleanup**.

The **Disk Cleanup: Drive Selection** window is displayed.

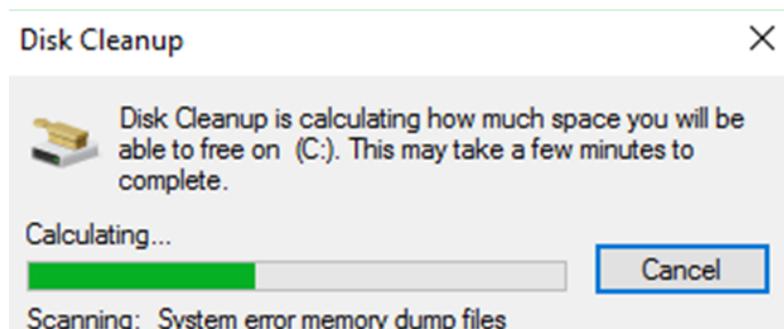
Figure 6-10 Disk Cleanup: Drive Selection



Step 3 Select the target disk from the drop-down list. In this example, disk (C:) is selected.

The **Disk Cleanup** window is displayed, and the system automatically calculates the space that can be freed on disk (C:).

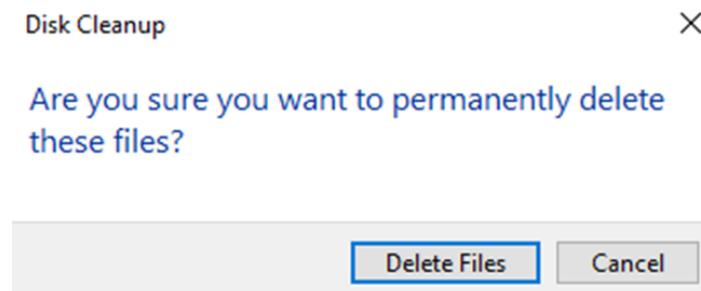
Figure 6-11 Disk Cleanup



Step 4 After the automatic calculation is complete, select the files to be deleted on the displayed window and click **OK**.

A confirmation dialog box is displayed.

Figure 6-12 Deletion confirmation



Step 5 Click **Delete Files** to clean up the disk space.

----End

Uninstalling Unnecessary Programs from the Control Panel

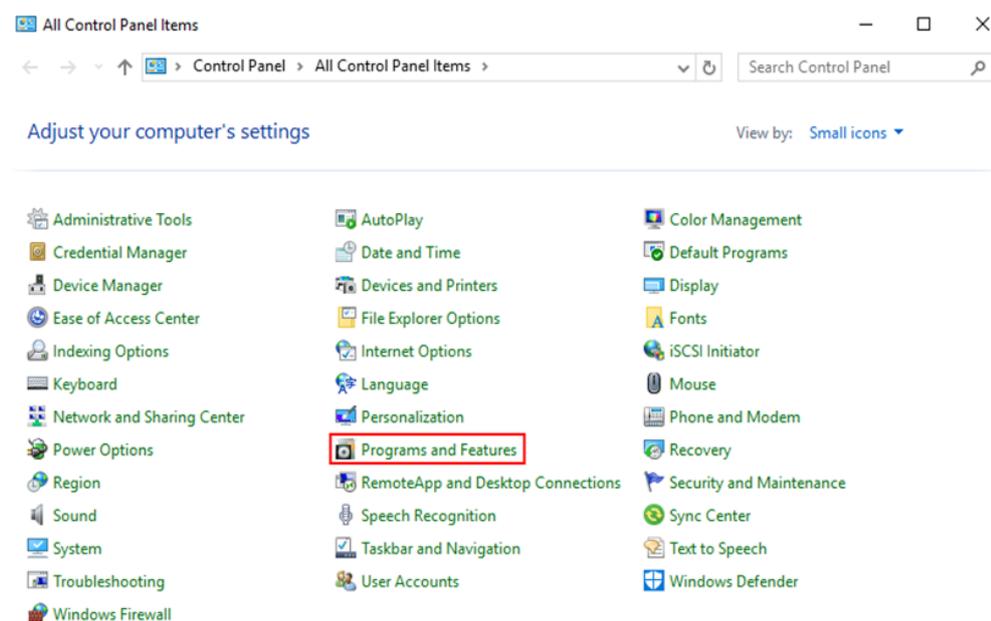
Step 1 On the server desktop, click the start icon in the lower left corner.

The start menu is displayed.

Step 2 In the navigation pane on the left, choose **Windows System > Control Panel**.

The **All Control Panel Items** window is displayed.

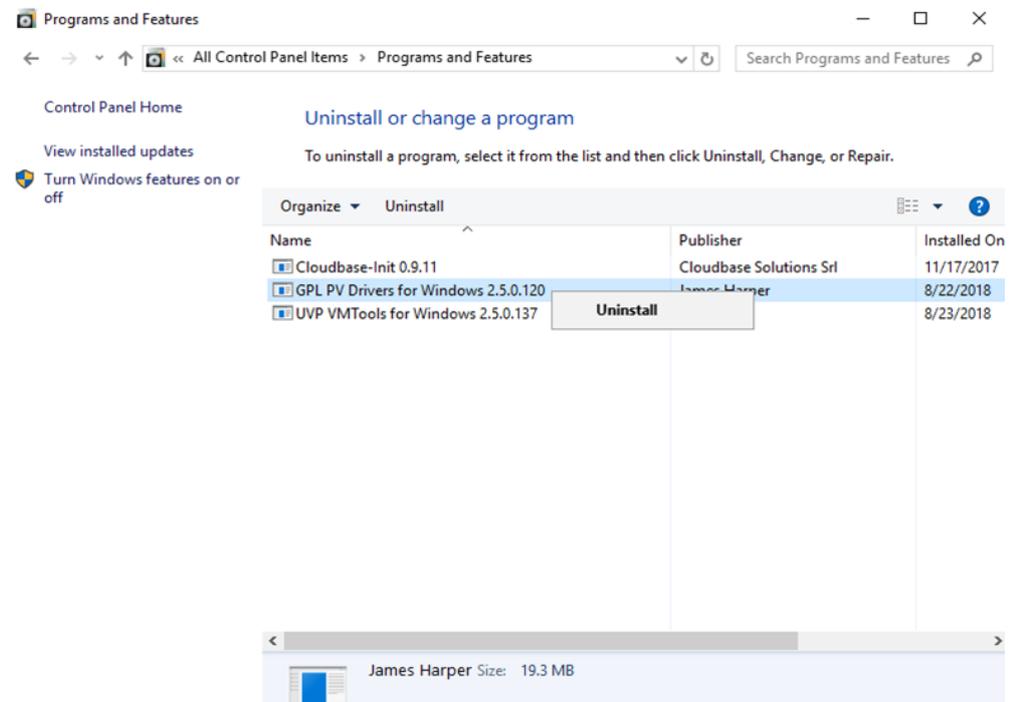
Figure 6-13 All Control Panel Items



Step 3 In the list, select **Programs and Features**.

The **Programs and Features** window is displayed.

Figure 6-14 Programs and Features



Step 4 In the program list, right-click the program to be uninstalled and choose **Uninstall** from the shortcut menu.

A confirmation dialog box is displayed.

Figure 6-15 Uninstallation confirmation



Step 5 Click **Yes** to uninstall the program.

----End

6.11 Handling Insufficient Disk Space on a Linux ECS

Scenarios

Applications running on Linux servers continuously generate and accumulate data, occupying disk space. As the disk capacity on a server is limited, after the server runs for a long time, storage resources may be used up. If the disk space is insufficient or used up, new data cannot be written to applications. This may

affect application running or even cause service interruption or system exception. Therefore, you need to monitor the disk usage and take measures to clear disk space in a timely manner.

This section describes how to check the disk usage and address the space insufficiency problems. The method varies depending on the in-use server OS. This section uses CentOS as the example OS to describe how to resolve issues related to insufficient storage space. For other Linux OSs, see the corresponding OS product documents.

Symptom

If error message "No space left on device" is returned when you create a file or run an application on a Linux ECS, the disk space on the ECS is used up. In this case, you need to identify the cause of the insufficient disk space and take measures to resolve the issue.

Possible Causes

There are several possible causes of the insufficient disk space:

- **100% disk partition usage**
- **100% partition inode usage**
- **There are zombie processes**

NOTE

There are unreleased spaces used by deleted files that are still held open.

- **Overwritten mount points**

NOTE

If the original file system has a directory with a large number of files inside, and you mount a new device on this directory, the original mount point is overwritten. When applications in your system continue reading and writing the original file system space, they may receive the space insufficiency message. However, this cannot be reflected when you run the **df** or **du** command to check the directory usage of the disk. Because **df** or **du** only collects statistics for the newly mounted device.

100% Disk Partition Usage

If the usage of a disk partition reaches 100%, error message "No space left on device" may be reported. You are advised to check the disk usage, identify files that occupy a large amount of space, and delete unnecessary files or expand the disk capacity.

Refer to the following commands to check the disk usage:

Step 1 Check the disk usage.

df -h

```
[test@ecs-test-0001 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        485M  0    485M  0% /dev
tmpfs           496M  0    496M  0% /dev/shm
tmpfs           496M  6.8M 489M  2% /run
tmpfs           496M  0    496M  0% /sys/fs/cgroup
```

```
/dev/vda1 9.8G 2.9G 6.4G 31% /  
tmpfs 108M 0 108M 0% /run/user/0  
/dev/vdb1 9.8G 9.8G 0 100% /data-test
```

In this example, the usage of the **/dev/vdb1** partition reaches 100%.

Step 2 Go to the root directory and check which directory is using the most space.

```
sudo du -sh /* | sort -rh | head -n 10
```

```
[test@ecs-test-0001 ~]# sudo du -sh /* | sort -rh | head -n 18  
du: cannot access '/proc/1781B/task/1781B/?d/4': No such file or directory  
du: cannot access '/proc/1781B/task/1781B/?dinfo/4': No such file or directory  
du: cannot access '/proc/1781B/?d/4': No such file or directory  
du: cannot access '/proc/1781B/?dinfo/4': No such file or directory  
6G /home  
2.1G /usr  
413M /var  
225M /CloudrResetPwdAgent  
137M /boot  
36M /etc  
6.8M /run  
944K /tmp  
144K /root  
28K /data-test
```

In this example, the **/home** directory occupies the largest space. Consider deleting files in this directory.

Step 3 Go to the **/home** directory and check which files or directories in the **/home** directory are using the most space. Find large files that are consuming the most disk space.

```
sudo du -sh /home/* | sort -rh | head -n 10
```

```
----End
```

100% Partition Inode Usage

Inodes are data structures used by many file systems to store file metadata. Each file or directory has an inode that stores metadata such as the file permissions, owner, file size, timestamps (created and modified), and location of data blocks, but does not store the file name. File names are stored in directory entries, which work as labels pointing to inodes. The number of inodes is fixed at the time of disk formatting. When they are used up, new files cannot be created even if there is still space available on the disk. So, when the inode usage approaches or already reaches 100%, delete unnecessary files to release inodes.

Step 1 Analyze the number of files in each level-2 directory under the root directory.

```
for i in /*; do echo $i; sudo find $i | wc -l; done
```

```
[test@ecs-test-0001 ~]# for i in /*; do echo $i; sudo find $i | wc -l; done  
/bin  
1  
/boot  
331  
/CloudrResetPwdAgent  
198  
/data-test  
2  
/dev  
346  
/etc
```

```
2433
/home
5
/lib
1
/lib64
1
/lost+found
1
```

In this example, the directories with the most files have the highest inode usage. For example, the `/etc` directory has 2433 files. You can delete unnecessary files.

Step 2 Check the number of files in the subdirectories and delete files as required.

```
for i in /etc/*; do echo $i; sudo find $i | wc -l; done

----End
```

There Are Zombie Processes

If a file is deleted but held open by one or more active processes, the file continues to occupy the disk space. Although the file may no longer appear in the file system's directory structure, its data remains on disk as long as any running process continues to reference it through an open file descriptor. The data is only released once all such processes have closed the file or terminated.

Step 1 Install `lsof`.

```
sudo yum install -y lsof
```

Step 2 Check files that are deleted but still occupied by processes.

```
sudo lsof | grep delete | sort -k7 -rn | more
```

```
[test@ecs-test-0001 ~]# lsof | grep delete | sort -k7 -rn | more
tail 619544 root 3r REG 253,1 500000000 136507 /home/test_file (deleted)
```

In this example, the file size is **500000000Byte**. Determine whether zombie processes occupy too much space based on the actual condition.

You can release space in either of the following ways:

- Restart the server to stop all processes, thereby releasing the disk space. Restarting the server may affect services. Exercise caution when performing this operation.
- Run the **kill** command to clear the process. For details, see [Step 3](#).

Step 3 Run the **kill** command to stop the process.

1. Obtain the process ID.
lsof | grep delete
2. Run the **kill** command.
kill <process-ID>

NOTICE

If the server is running services, the services may be affected. Exercise caution when performing this operation.

----End

Overwritten Mount Points

If all prior causes have been eliminated and the issue of insufficient space continues, it is possible that the directory's mount point has been overwritten by another file system. For example, when a new file system is mounted on top of a directory that already holds data, the existing contents become hidden, yet they continue to consume disk space. This makes it appear as if the disk space has mysteriously vanished. You can perform the following steps to check whether this problem occurs.

Step 1 Check the mount information.

mount

```
[test@ecs-test-0001 ~]# mount
sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
devtmpfs on /dev type devtmpfs (rw,nosuid,size=496616k,nr_inodes=124154,mode=755)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
debugfs on /sys/kernel/debug type debugfs (rw,relatime)
tmpfs on /run/user/0 type tmpfs (rw,nosuid,nodev,relatime,size=101412k,mode=700)
binfmt_misc on /proc/sys/fs/binfmt_misc type binfmt_misc (rw,relatime)
/dev/vdb1 on /data-test type ext4 (rw,relatime,data=ordered)
/dev/vdc on /data-test type ext4 (rw,relatime,data=ordered)
```

In the example, two devices are mounted on **/data-test**, which can lead to the mount point being overwritten by the most recent mount operation.

Step 2 Check the mount information of the current partition.

df -h

```
[test@ecs-test-0001 ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        485M   0 485M   0% /dev
tmpfs           496M   0 496M   0% /dev/shm
tmpfs           496M  13M 483M   3% /run
tmpfs           496M   0 496M   0% /sys/fs/cgroup
/dev/vda1       9.8G  3.2G  6.1G  34% /
tmpfs           100M   0 100M   0% /run/user/0
/dev/vdb1       9.8G   37M  9.2G   1% /data-test
```

Step 3 Run **umount** to unmount a partition.

In this example, run **sudo umount /mnt/vdc** to unmount **/dev/vdc**.

 **CAUTION**

Unmounting a file system may interrupt your application services. Perform the unmount operation at an appropriate time.

----End

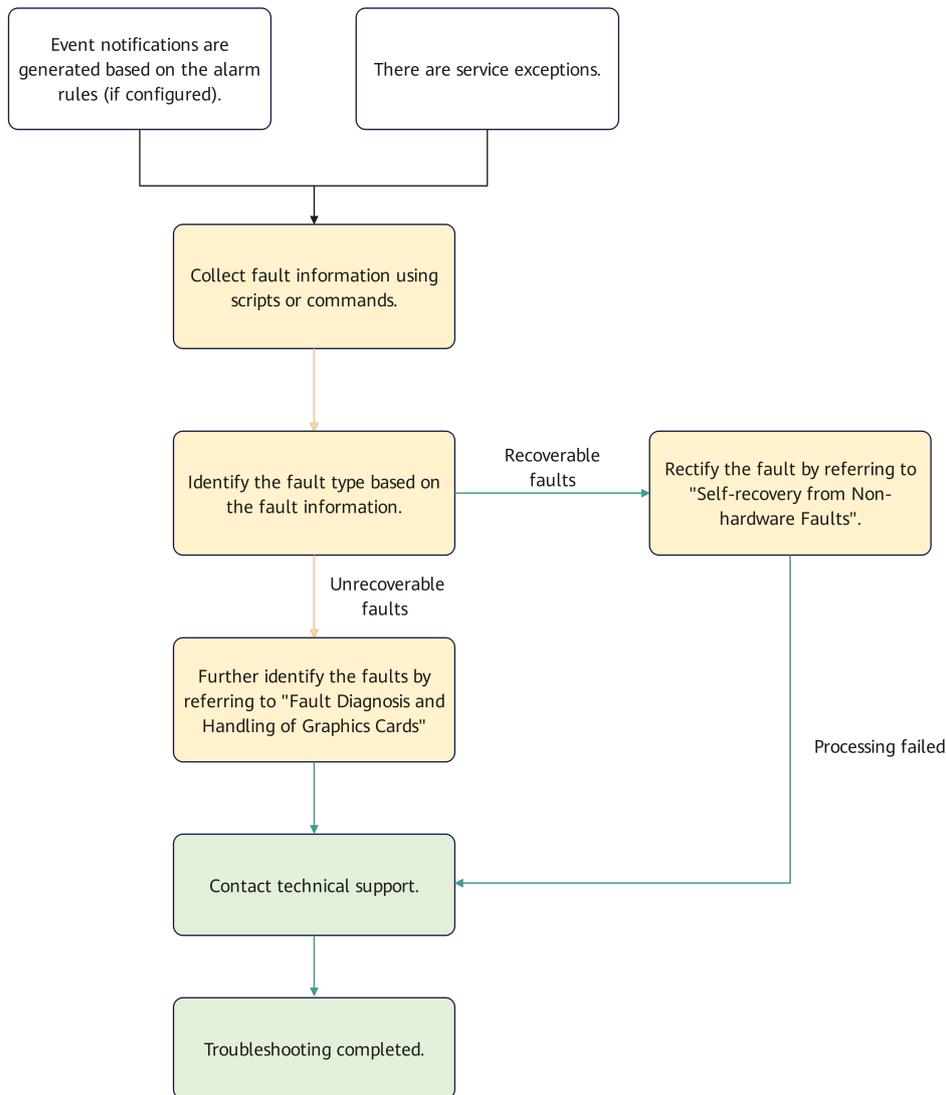
7 Self-diagnosis of Faulty GPU-accelerated ECSs

7.1 How Do I Troubleshoot a Faulty GPU-accelerated ECS?

Figure 7-1 shows the troubleshooting process for faulty GPU-accelerated ECSs. The operations are as follows:

- : Configure alarm rules on the Cloud Eye console for GPU-accelerated ECSs, so that notifications will be sent to you when there is a fault.
- **Fault Information Collection**: Perform one-click fault information collection using the GPU fault information collection script or run commands to collect fault information by referring to **Fault Information Collection**.
- **What Types of GPU-accelerated ECS Faults Are There**: Identify the fault type in the fault classification list based on the error information.
- **Self-recovery from Non-hardware Faults**: Locate and rectify this type of issues according to the guidance.
- **Fault Diagnosis and Handling of Graphics Cards**: Contact technical support after confirming the problem according to the guidance.
- : If the fault persists, submit a service ticket.

Figure 7-1 GPU-accelerated ECS troubleshooting



7.2 What Types of GPU-accelerated ECS Faults Are There?

Table 7-1 lists the types of GPU-accelerated ECS faults.

Table 7-1 Types of GPU-accelerated ECS faults

Fault Rectifiable	Fault Type	Reference
Yes. Faults can be rectified by referring to related guidance documents.	Incorrect image configuration	What Do I Do If the Nouveau Driver Is Not Disabled?
	ECC errors	What Do I Do If There Are Retired Pages?
	Kernel upgrade issues	What Do I Do If the Driver Is Unavailable After the Kernel Is Upgraded?
	Disconnected GPUs	Why Is the Number of Queried Graphics Cards Different from the Actual One?
	Graphics card ERR!	What Do I Do If ERR! Is Displayed?
	Software installation problems	What Do I Do If an Error Occurs During the Installation of the NVIDIA Driver and CUDA Software?
	Driver compatibility issues	How Do I Handle Driver Compatibility Issues?
	Xid errors	How Do I Handle Recoverable Xid Errors?
	Graphics card being disabled	What Do I Do If the Error Message "Windows has stopped this device because it has reported problems" Is Displayed?
	Image errors	What Do I Do If the Driver and Image I Have Selected Do Not Match My Service Scenario?
	License issues	What Do I Do If I Have Installed the GRID Driver but Have Not Purchased or Configured the License?
No. If the fault cannot be rectified, contact technical support.	InfoROM errors	How Do I Handle the infoROM Error?
	ECC errors	What Do I Do If ECC Error "double bit ecc error" Occurs and There Are No Retired Pages Shown in the nvidia-smi -q Command Output?
		What Do I Do If nvidia-smi Command Output Shows the SRAM ECC Error (V100 GPUs)?

Fault Rectifiable	Fault Type	Reference
	Disconnected GPUs	What Do I Do If the GPU Is Disconnected or the Graphics Card Can't Be found, or rev ff Is Displayed After lspci grep -i nvidia Is Executed?
	High temperature	What Do I Do If the nvidia-smi Command Output Shows Overheated GPUs?
	Driver installation errors	What Do I Do If "Unable to load the kernel module 'nvidia.ko'" Is Displayed During Driver Installation?
	Xid errors	What Can I Do If an Xid Error Is Displayed in the Message Log When a GPU-accelerated ECS Is Faulty?

7.3 Fault Information Collection

7.3.1 How Do I Collect Fault Information?

You can use the fault information collection script or run commands to collect the fault information.

- Use the fault information collection script to collect information.
 - a. Download the fault information collection script at https://hgcs-drivers-cn-north-4.obs.cn-north-4.myhuaweicloud.com/release/script/diagnose_gpu.sh.
 - b. Run the **bash diagnose_gpu.sh** command to save the information to the **diagnose_gpu_xxxx.tar.gz** file. Check the information by yourself or contact technical support by submitting a service ticket.
- Run commands to collect the fault information by referring to [Table 7-2](#).

Table 7-2 Method for obtaining the fault information

Information Type	Reference
Basic information about graphics cards	How Do I Obtain the Graphics Card ID?
	How Do I Query Graphics Card Details?
	How Do I Query the Status of a Graphics Card?

Information Type	Reference
Fault information about graphics cards (Linux)	How Do I Query NVIDIA Error Messages?
	How Do I Query XID Errors?
NVIDIA logs (Linux)	How Do I Collect NVIDIA Logs?
Image kernel information (Linux)	How Do I Query Kernel Information?
Driver installation information (Linux)	How Do I Collect the Driver Installation Information?

7.3.2 How Do I Obtain the Graphics Card ID?

- Perform the following steps to obtain the graphics card ID of a Linux ECS:
 - a. Log in to the ECS.
 - b. Run the **nvidia-smi** command in any directory. (The directory is **/opt/cloud/cce/nvidia/bin** for CCE clusters.)

```
root@ecs-dd86 ~]# nvidia-smi
Fri Dec 2 11:48:08 2022

+-----+
| NVIDIA-SMI 460.73.01    Driver Version: 460.73.01    CUDA Version: 11.2    |
+-----+
| GPU   Name           Persistence-MI  Bus-Id        Disp.A | Volatile Uncorr. ECC | |
| Fan  Temp  Perf    Pwr:Usage/Cap |  Memory-Usage | GPU-Util  Compute M. |
|                                           |              | MIG M. |
+-----+-----+
|  0   Tesla T4             Off   | 00000000:00:0D:00 Off |   0          0      0 |
| N/A   54C    P8     28W /  70W |  0MiB / 15109MiB |   0%          Default |
+-----+-----+

Processes:
+-----+
| GPU   GI   CI        PID   Type   Process name                      GPU Memory |
| ID   ID   ID             |                | Usage |
+-----+-----+
| No running processes found |
+-----+

root@ecs-dd86 ~]#
```

- Perform the following steps to obtain the graphics card ID of a Windows ECS:
 - a. Go to the **C:\Program Files\NVIDIA Corporation\NVSMI** directory.
 - b. Run the **nvidia-smi** command.

```
root@ecs-dd86 ~]# nvidia-smi
Fri Dec 2 11:48:08 2022

+-----+
| NVIDIA-SMI 460.73.01    Driver Version: 460.73.01    CUDA Version: 11.2    |
+-----+-----+
| GPU  Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|                                           |              | MIG M.   |
+-----+-----+
|  0   Tesla T4             Off   | 00000000:00:0D:0  Off  |            0         |
| N/A   54C    P0      28W / 70W   |  0MiB / 15109MiB |         0%      Default |
+-----+-----+

Processes:
+-----+-----+
| GPU  GI  CI           PID  Type  Process name          GPU Memory |
| ID   ID  ID                                   Usage   |
+-----+-----+
| No running processes found |
+-----+-----+

root@ecs-dd86 ~]#
```

7.3.3 How Do I Query Graphics Card Details?

- Query details about a specified graphics card:
 - a. Log in to the ECS.
 - b. Run the following command to query details about a specified graphics card:
nvidia-smi -q -i *{graphics card ID}*
- Query details about all graphics cards:
 - a. Log in to the ECS.
 - b. Run the following command to query details about all graphics cards:
nvidia-smi -q

The following is an example of the command output:

You can also filter the output by keyword, such as **NVRM**, **nouveau**, **nvidia**, and **nv**, and save the output.

7.3.6 How Do I Query XID Errors?

XID messages are errors printed by the NVIDIA driver to the OS kernel logs or event logs. These messages can be used to identify GPU error events and provide information such as the error type, error location, and error code in the GPU hardware, NVIDIA software, or your applications.

You can perform the following steps to query the XID errors:

1. Log in to the ECS.
2. Run the following command to check whether there are XID errors and save the command output:

```
dmesg | grep -i xid
```

- If the XID on the GPU node is empty, there is no XID error message.
- If the XID on the GPU node is not empty, refer to [What Types of GPU-accelerated ECS Faults Are There?](#) or contact technical support.

7.3.7 How Do I Collect NVIDIA Logs?

You can perform the following steps to collect NVIDIA logs:

1. Log in to the ECS.
2. Run the following command in any directory: (In the CCE scenario, go to the `/opt/cloud/cce/nvidia/bin` directory to run the command.)

```
sh nvidia-bug-report.sh
```

or

```
./nvidia-bug-report.sh
```

3. Wait until the execution is complete and **nvidia-bug-report.log.gz** is generated in the current directory. Then, dump the log.

```
[root@ecs-dd86 ~]# sh nvidia-bug-report.sh

nvidia-bug-report.sh will now collect information about your
system and create the file 'nvidia-bug-report.log.gz' in the current
directory. It may take several seconds to run. In some
cases, it may hang trying to capture data generated dynamically
by the Linux kernel and/or the NVIDIA kernel module. While
the bug report log file will be incomplete if this happens, it
may still contain enough data to diagnose your problem.

If nvidia-bug-report.sh hangs, consider running with the --safe-mode
and --extra-system-data command line arguments.

Please include the 'nvidia-bug-report.log.gz' log file when reporting
your bug via the NVIDIA Linux forum (see forums.developer.nvidia.com)
or by sending email to 'linux-bugs@nvidia.com'.

By delivering 'nvidia-bug-report.log.gz' to NVIDIA, you acknowledge
and agree that personal information may inadvertently be included in
the output. Notwithstanding the foregoing, NVIDIA will use the
output only for the purpose of investigating your reported issue.

Running nvidia-bug-report.sh... complete.

[root@ecs-dd86 ~]# ll
total 1252
-rw-r--r-- 1 root root 1279051 Dec  2 11:50 nvidia-bug-report.log.gz
[root@ecs-dd86 ~]#
```

7.3.8 How Do I Query Kernel Information?

You can perform the following steps to query the kernel information:

1. Log in to the ECS.
2. Run the following command to view the kernel version.

```
uname -r
```

```
[root@ecs-dd86 ~]# uname -r
4.18.0-240.10.1.el8_3.x86_64
[root@ecs-dd86 ~]#
```

3. Run the following commands to check the kernel version when the driver is installed:

Ubuntu: **find /lib/modules -name nvidia.ko**

CentOS: **find /usr/lib/modules -name nvidia.ko**

```
[root@ecs-6711 ~]# find /usr/lib/modules -name nvidia.ko
/usr/lib/modules/3.10.0-957.21.3.el7.x86_64/kernel/drivers/video/nvidia.ko
[root@ecs-6711 ~]#
```

4. Run the following command to query the GCC version:
rpm -qa | grep gcc
5. Run the following command to check the kernel-devel version:
rpm -qa | grep kernel-devel

7.3.9 How Do I Collect the Driver Installation Information?

You can perform the following steps to collect the driver installation information:

1. Log in to the ECS.
2. Run the following command to check whether the Nouveau driver is disabled:
lsmod | grep nouveau

If the driver fails to be installed, collect the `/var/log/nvidia-installer.log` driver installation log and dump the log information.

7.4 Self-recovery from Non-hardware Faults

7.4.1 What Do I Do If the Nouveau Driver Is Not Disabled?

Symptom

If the Nouveau driver is not disabled, the Linux OS may be suspended and the ECS may not be remotely logged in to. This issue generally occurs when you use your own private images (exported from ECSs or other sources).

Checking Nouveau Status

1. Run the following command to check the keyword errors in the Linux kernel ring buffer:

dmesg | grep error

- If the command output contains **nouveau**, the Nouveau driver may not be disabled. Then, go to [2](#).
- If the command output does not contain **nouveau**, go to [2](#).

```
root@ecs-dd86 ~]# dmesg |grep -i error
[ 18.603210] nouveau: probe of 0000:21:01.0 failed with error -12
[ 18.622555] nouveau: probe of 0000:21:02.0 failed with error -12
[ 18.641799] nouveau: probe of 0000:21:03.0 failed with error -12
[ 18.661004] nouveau: probe of 0000:21:04.0 failed with error -12
[ 21.558220] EXT4-fs (vda1): re-mounted. Opts: errors=remount-ro
```

2. Run the following command to check whether the Nouveau driver has been installed:

lsmod | grep nouveau

- If there is no command output or the command output does not contain **nouveau**, the Nouveau driver has been disabled.
- If the command output contains **nouveau**, the Nouveau driver has been installed. Then, disable the Nouveau driver.

```
[root@ecs-dd86 ~]# lsmod | grep nouveau
nouveau                2281472  0
video                   49152  1 nouveau
mxm_wmi                 16384  1 nouveau
wmi                     32768  2 mxm_wmi,nouveau
i2c_algo_bit           16384  1 nouveau
drm_kms_helper         217888  5 cirrus,nvidia_drm,nouveau
ttm                    118592  1 nouveau
drm                    557856  6 drm_kms_helper,cirrus,nvidia_drm,ttm,nouveau
[root@ecs-dd86 ~]#
```

Solution

1. Edit the **blacklist.conf** file.
If the `/etc/modprobe.d/blacklist.conf` file is unavailable, create one.

vi /etc/modprobe.d/blacklist.conf

Add the following statement to the end of the file:

```
blacklist nouveau
options nouveau modeset=0
```

2. Run the following command to back up and create an initramfs application:
 - Ubuntu:


```
sudo update-initramfs -u
```
 - CentOS:


```
mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak
                    dracut -v /boot/initramfs-$(uname -r).img $(uname -r)
```
3. Run the following command to restart the ECS:


```
reboot
```

7.4.2 What Do I Do If There Are Retired Pages?

Symptom

- The service becomes abnormal after being scheduled to a GPU node, and becomes normal when it is scheduled to another node.
- The GPU memory usage of an ECS suddenly decreases.

Checking Whether There Are Retired Pages

1. Run the following command to check whether there are ECC errors in the graphics card:

nvidia-smi

```
[root@ec2 ~]# /opt/cloud/ece/nvidia/bin/nvidia-smi
Tue Oct 25 09:32:59 2022
```

NVIDIA-SMI		Driver Version: 470.57.02			CUDA Version: 11.7.1		
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile	Uncorr. ECC	
Fan	Temp	Perf	Memory-Usage	Memory-Usage	GPU-Util	Compute M.	MIG M.
		Pwr:Usage/Cap					
0	Tesla V100-SXM2...	Off	00000000:2D:00:0	Off	100%	Default	0
N/A	67C	P0 297W / 300W	26984MiB / 32510MiB			N/A	N/A
1	Tesla V100-SXM2...	Off	00000000:32:00:0	Off	100%	Default	0
N/A	69C	P0 294W / 300W	23128MiB / 32510MiB			N/A	N/A
2	Tesla V100-SXM2...	Off	00000000:5B:00:0	Off	100%	Default	0
N/A	70C	P0 287W / 300W	23056MiB / 32510MiB			N/A	N/A
3	Tesla V100-SXM2...	Off	00000000:5F:00:0	Off	100%	Default	0
N/A	59C	P0 304W / 300W	23032MiB / 32510MiB			N/A	N/A
4	Tesla V100-SXM2...	Off	00000000:B5:00:0	Off	74%	Default	0
N/A	42C	P0 94W / 300W	4624MiB / 32510MiB			N/A	N/A
5	Tesla V100-SXM2...	Off	00000000:BE:00:0	Off	0%	Default	0
N/A	32C	P0 44W / 300W	0MiB / 32510MiB			N/A	N/A
6	Tesla V100-SXM2...	Off	00000000:E1:00:0	Off	0%	Default	1
N/A	31C	P0 42W / 300W	0MiB / 32510MiB			N/A	N/A
7	Tesla V100-SXM2...	Off	00000000:E9:00:0	Off	0%	Default	0
N/A	30C	P0 43W / 300W	0MiB / 32510MiB			N/A	N/A

2. In the command output in 1, check the number of ECC errors in the **Volatile Uncorr. ECC** column. If the number is greater than 0, run the following command to check whether there are retired pages.

nvidia-smi -q -i *&.{gpu_id}* -d PAGE_RETIREMENT

```
[root@ecs-dd86 ~]# nvidia-smi -q -i 0 -d PAGE_RETIREMENT
=====NUSMI LOG=====
Timestamp                : Fri Dec 2 14:58:38 2022
Driver Version           : 460.73.01
CUDA Version             : 11.2
Attached GPUs            : 1
GPU 00000000:00:0D.0
  Retired Pages
    Single Bit ECC        : 0
    Double Bit ECC       : 0
    Pending Page Blacklist : No
[root@ecs-dd86 ~]#
```

If **Pending Page Blacklist** is **No** in the command output, there are no retired pages.

3. If the number of ECC errors in the **Volatile Uncorr. ECC** which is displayed in the command output in **1** is **0**, run the following command to check whether there are retired pages for all GPUs.

nvidia-smi -q -d PAGE_RETIREMENT

4. If **Pending Page Blacklist** is **Yes** in the command output in **3**, there are retired pages. Then, reload the driver to retire the pages.

```
Texture Memory          : N/A
Texture Shared          : N/A
CBU                     : 0
Total                   : 308
Retired Pages
  Single Bit ECC        : 0
  Double Bit ECC       : 2
  Pending Page Blacklist : Yes
Temperature
GPU Current Temp       : 30 C
GPU Shutdown Temp     : 90 C
GPU Slowdown Temp     : 87 C
GPU Max Operating Temp : 83 C
Memory Current Temp   : 29 C
Memory Max Operating Temp : 85 C
Power Readings
Power Management       : Supported
Power Draw             : 23.37 W
Power Limit            : 250.00 W
```

Solution

- Method 1:
 - a. Run the following command to check the GPU usage and kill all processes using the GPU:
nvidia-smi
 - b. Run the following command to reset the GPU:
nvidia-smi -r
 - c. Run the following command to check whether there are retired pages:
nvidia-smi -q -d PAGE_RETIREMENT
If **Pending Page Blacklist** is **No**, there are no retired pages.
- Method 2:
 - a. Run the following command to restart the ECS:
reboot

- b. Run the following command to check whether there are retired pages:
nvidia-smi -q -d PAGE_RETIREMENT
If **Pending Page Blacklist** is **No**, there are no retired pages.

7.4.3 What Do I Do If the Driver Is Unavailable After the Kernel Is Upgraded?

Symptom

- When you run the **nvidia-smi** command, the error message "failed to initialize NVML: Driver/library version mismatch" is displayed.
- When you run the **nvidia-smi** command, the error message "NVIDIA-SMI has failed because it couldn't communicate with the NVIDIA driver" is displayed.

```
ubuntu@ip-10-10-10-10:~$ nvidia-smi -L
NVIDIA-SMI has failed because it couldn't communicate with the NVIDIA driver.
Make sure that the latest NVIDIA driver is installed and running.
```

Checking the Kernel Version

1. Run the following command to check the current kernel version:
uname -r
2. Run the corresponding command on the ECS to check the version of the kernel when the driver was installed:
 - CentOS: **find /usr/lib/modules -name nvidia.ko**
 - Ubuntu: **find /lib/modules -name nvidia.ko**

If the current kernel version is different from the kernel version when the driver was installed, the driver is unavailable after the kernel is upgraded.

Solution

1. Run the following commands in sequence to remove the NVIDIA driver:
rmmod nvidia_drm
rmmod nvidia_modeset
rmmod nvidia
2. Run the following command to query the GPU information:
nvidia-smi
 - If the command output is normal, the problem has been resolved.
 - If the fault persists in the command output, refer to [Why Is the GPU Driver Unavailable?](#)

7.4.4 Why Is the Number of Queried Graphics Cards Different from the Actual One?

Symptom

The number of graphics cards queried by running the `nvidia-smi` command is less than the actual number of graphics cards.

```
Wed Sep 14 17:06:06 2022
+-----+
| NVIDIA-SMI 470.57.02 Driver Version: 470.57.02 CUDA Version: 11.8.0 |
+-----+
| GPU  Name Persistence-Mi Bus-Id Disp.A Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap Memory-Usage GPU-Util Compute M. |
+-----+-----+-----+-----+-----+-----+
|  0  Tesla V100S-PCI... Off  | 00000000:00:0D:0 Off |           |           |
| N/A  34C   P0   37W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  1  Tesla V100S-PCI... Off  | 00000000:00:0F:0 Off |           |           |
| N/A  35C   P0   36W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  2  Tesla V100S-PCI... Off  | 00000000:00:10:0 Off |           |           |
| N/A  35C   P0   37W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  3  Tesla V100S-PCI... Off  | 00000000:00:11:0 Off |           |           |
| N/A  33C   P0   36W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  4  Tesla V100S-PCI... Off  | 00000000:00:12:0 Off |           |           |
| N/A  34C   P0   36W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  5  Tesla V100S-PCI... Off  | 00000000:00:13:0 Off |           |           |
| N/A  35C   P0   36W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
|  6  Tesla V100S-PCI... Off  | 00000000:00:14:0 Off |           |           |
| N/A  33C   P0   35W / 250W    0MiB / 32510MiB      0%      Default |
+-----+-----+-----+-----+-----+-----+
| Processes: |
| GPU      PID   Type   Process name                               GPU Memory |
|-----+-----+-----+-----+-----+-----+
| No running processes found
+-----+
[root@p2s-8v100-with-ssd-new-1-0b5d6 ~]#
```

In this figure, the `nvidia-smi` command output indicates that there are seven graphics cards, but the ECS actually has eight graphics cards.

Checking the Number of Graphics Cards

Run the following command. If the number of the queried graphics cards is the same as the actual number of the graphics cards and the graphics card's status is normal (**rev a1**), go to [Solution](#). If the graphics card cannot be found or the status is **rev ff**, refer to [Fault Diagnosis and Handling of Graphics Cards](#). You can query the number of the graphics card defined by the flavors by referring to [GPU-accelerated ECSs](#).

`lspci | grep -i nvidia`

```
# lspci | grep NV
controller: NVIDIA Corporation Device 1db6 (rev a1)
```

Solution

- In non-CCE clusters, reinstall the driver or upgrade the driver and run the `nvidia-smi` command to check whether the fault persists. If the fault persists, refer to [Fault Information Collection](#) and contact technical support.

- In CCE clusters, collect fault information by referring to [Fault Information Collection](#) and contact technical support.

7.4.5 What Do I Do If ERR! Is Displayed?

Symptom

ERR! is only displayed for **Pwr:Usage/Cap** (energy consumption) when you run the **nvidia-smi** command.

```
[root@ecs-8692 ~]# /var/paas/nvidia/bin/nvidia-smi
Tue Oct 18 21:14:59 2022

+-----+
| NVIDIA-SMI 450.51.05 Driver Version: 450.51.05 CUDA Version: 11.8.0 |
+-----+
| GPU Name Persistence-Mi Bus-Id Disp.A | Volatile Uncorr. ECC | | | | | | |
| Fan Temp Perf Pwr:Usage/Cap | Memory-Usage | GPU-Util Compute M. |
| | | | | | | | |
+-----+
| 0 Tesla T4 Off | 00000000:00:0D:00 Off | | | | | | |
| N/A 49C P0 ERR! / 70W | 0MiB / 15109MiB | 0% Default |
| | | | | | | | |
+-----+
| Processes: |
| GPU GI CI PID Type Process name GPU Memory |
| ID ID | | | | | Usage |
+-----+
| No running processes found |
+-----+
```

Solution

1. If the current services are normal, and "ERR!" is displayed only when **nvidia-smi** is executed, no operation is required.
2. If the current services have been affected, migrate the ECS, collect fault information by referring to [Fault Information Collection](#), and contact technical support.

7.4.6 What Do I Do If an Error Occurs During the Installation of the NVIDIA Driver and CUDA Software?

Symptom

An error occurs in the script execution when you use a public or private image without a driver to install the NVIDIA driver and CUDA software packages.

```
[root@ecs-8692 ~]# sudo sh cuda_10.1.105_418.39_linux.run
cuda_10.1.105_418.39_linux.run: line 1: syntax error near unexpected token `newline'
cuda_10.1.105_418.39_linux.run: line 1: `<html>'
[root@ecs-8692 ~]# sudo sh cuda_10.1.105_418.39_linux.run
cuda_10.1.105_418.39_linux.run: line 1: syntax error near unexpected token `newline'
cuda_10.1.105_418.39_linux.run: line 1: `<html>'
[root@ecs-8692 ~]# sh cuda_10.1.243_418.87.00_linux.run
sh: cuda_10.1.243_418.87.00_linux.run: No such file or directory
[root@ecs-8692 ~]# sh cuda_10.1.105_418.39_linux.run
cuda_10.1.105_418.39_linux.run: line 1: syntax error near unexpected token `newline'
cuda_10.1.105_418.39_linux.run: line 1: `<html>'
[root@ecs-8692 ~]#
```

Checking Software

1. Check the image file.
2. Check the source of the NVIDIA software package.

3. Check the target NVIDIA software package version and CUDA software version.

Solution

- (Recommended) Use a script for automatic driver installation. Install the target CUDA version using the automatic driver installation script provided by Huawei Cloud.
- If the automatic driver installation script does not contain the target software version, contact technical support.

7.4.7 How Do I Handle Driver Compatibility Issues?

Symptom

The error message "No devices were found" is displayed after you run the **nvidia-smi** command.

```
No devices were found
```

Solution

1. Check the ECS specifications and image information.
 - If the ECS is using NVIDIA Tesla T4 GPUs (such as the Pi2 or G6 ECS), refer to [Why Is the T4 GPU Display Abnormal?](#)
 - If the ECS is using other GPUs, go to the next step.
2. Check the system log **/var/log/message** for any reported driver-related errors.
 - If the error message "Failed to copy vbios to system memory" is displayed, the possible cause may be frequent driver loading/uninstallation. You are advised to enable the driver's persistence mode to keep the driver in loading state.

```
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:06:0e.0: Failed to copy vbios to system memory.  
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:06:0e.0: RmInitAdapter failed! (0x30:0xffff:852)  
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: rm_init_adapter failed, device minor number 1  
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: Failed to copy vbios to system memory.  
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: RmInitAdapter failed! (0x30:0xffff:852)  
Aug 29 14:35:00 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: rm_init_adapter failed, device minor number 1  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: Failed to copy vbios to system memory.  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: RmInitAdapter failed! (0x30:0xffff:852)  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: rm_init_adapter failed, device minor number 0  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: Failed to copy vbios to system memory.  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: RmInitAdapter failed! (0x30:0xffff:852)  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: rm_init_adapter failed, device minor number 0  
Aug 29 14:35:10 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: Failed to copy vbios to system memory.
```

- i. Run the following command to enable the driver's persistence mode.
nvidia-smi -pm 1
 - ii. Run the following command to open and edit the **/etc/rc.local** file.
vim /etc/rc.local
 - iii. Configure automatic startup and write the **nvidia-smi -pm 1** command to the **/etc/rc.local** file.
 - iv. Press **Esc** and enter **:wq** to save the settings and exit.
 - v. Run the following command to add startup permissions.
chmod +x /etc/rc.d/rc.local
 - If no error is found, go to the next step.
3. Check whether the ECS uses Tesla 510.xx.xx.

- If yes, the driver version may be incompatible with the image. You are advised to change the driver version. For details, see [GPU Driver](#).
 - If no, go to the next step.
4. If the root cause fails to be identified, collect fault information by referring to [Fault Information Collection](#) and contact technical support.

7.4.8 How Do I Handle Recoverable Xid Errors?

Possible Causes

Xid	Description
13	Graphics Engine Exception, a non-hardware fault, maybe an instruction error.
31	GPU memory page fault, a non-hardware fault, maybe an error of accessing an invalid address.
43	GPU stopped processing, a non-hardware fault, maybe a software error.

For details, see <https://docs.nvidia.com/deploy/xid-errors/index.html>.

Solution

1. Run the job again and check whether the Xid error disappears.
2. If the error persists, check the code or analyze logs to determine whether the Xid error is caused by the program.
3. If the Xid error is not caused by the program, contact technical support.

7.4.9 What Do I Do If the Error Message "Windows has stopped this device because it has reported problems" Is Displayed?

Symptom

The error message "Windows has stopped this device because it has reported problems" is displayed in the **Device status** when you query the graphics card properties using the display adapter in the Windows Device Manager.

Checking Version Mapping

1. Check whether out of memory (OOM) is displayed when you perform operations.
2. If you use a vGPU-accelerated ECS, check whether the version of the driver installed on the ECS matches that of the host.
 - a. Log in to the host that runs the ECS.
 - b. Run the **nvidia-smi** command to check the driver version and the version mapping.

The version mapping is shown at <https://docs.nvidia.com/grid/index.html>.

Solution

1. Restart the GPU ECS.
 - If the display adapter is restored, the recovery is complete.
 - If the fault persists, go to the next step.
2. Reinstall the GPU driver or upgrade the driver version. For details, see [GPU Driver](#).
3. If you are using a vGPU-accelerated ECS and the driver installed on the ECS does not match the driver version of the host, reinstall the driver software that matches the host version.

7.4.10 What Do I Do If the Driver and Image I Have Selected Do Not Match My Service Scenario?

Symptom

1. You have selected a public image with the Tesla driver (or a GRID driver) for your rendering (inference) service. An error occurred when the software was running.

For example, if you select an inappropriate public image for rendering, the error message "A D3D11-compatible GPU (Feature Level 11.0,Shader Model 5.0) is required to run the engine" is displayed.
2. You have selected a public image without a driver for your rendering (inference) service, and a driver has not been installed independently. Consequently, the GPU is unavailable, and the error message "command not found" is displayed when you run the **nvidia-smi** command.

Checking Whether the Image Has a Driver Installed

1. Confirm your service scenario.
2. Check whether the image has a driver that is suitable for the scenario installed.
 - If you use a heterogeneous public image, distinguish the driver type and version by image name. If the image name contains **with tesla**, the Tesla driver is installed for the selected image. If the image name contains **with grid**, the GRID driver (excluding the license) is installed for the selected image.
 - If you use a private image or other images, run the **nvidia-smi** command to check whether a driver is installed and check the driver type and version.
 - If you choose to install the Tesla driver by yourself, ensure that the Tesla driver version matches the CUDA software version. For details, see [Obtaining a Tesla Driver and CUDA Toolkit](#).

Solution

1. If you have not installed the driver, install the driver, switch to a public image that has a driver installed, or use the script to automatically install the driver as follows:
2. If you have installed the driver but the driver does not match your service scenario, uninstall the driver and reinstall a new one. For details, see [GPU Driver](#).

7.4.11 What Do I Do If I Have Installed the GRID Driver but Have Not Purchased or Configured the License?

Symptom

The GRID driver has been installed for your graphics processing, but the GPU usage or rendering performance is far from expectations.

For example, the image recognition process is suspended and the GPU performance is poor. The following error information is displayed in the `/var/log/messages` log. It is confirmed that the license has been purchased but not yet configured.

```
nvidia-grid: Valid GRID license not found. GPU features and performance are restricted. To enable full functionality please configure licensing details.
```

Solution

1. Check whether your application scenario is graphics processing and whether the ECS specifications meet the requirements of graphics processing.
2. Check whether the GRID driver is installed. Run the `nvidia-smi` command to check whether the command output is normal.
 - If the command output is normal and the GRID driver version can be queried, the GRID driver has been installed.
3. Check whether the license has been purchased. If yes, check whether the purchased license has been configured.
 - a. If you have not purchased a license, purchase one by referring to [Manually Installing a GRID Driver on a GPU-accelerated ECS](#) and then configure it.
 - b. If you have purchased a license but have not yet configured it, configure the license server and license file by referring to [Manually Installing a GRID Driver on a GPU-accelerated ECS](#).

7.5 Fault Diagnosis and Handling of Graphics Cards

7.5.1 How Do I Handle the infoROM Error?

Symptom

The error message "WARNING:infoROM is corrupted at gpu 0000:00:0D.0" is displayed when the `nvidia-smi` command is executed on a Linux ECS, and the services have been affected.

```
root@xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx:~# /var/paas/nvidia/bin/nvidia-smi
Fri Aug 12 20:45:41 2022
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| NVIDIA-SMI 440.33.01    Driver Version: 440.33.01    CUDA Version: 10.2    |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|    0   Tesla V100-PCIE...    Off   | 00000000:00:0D:0    Off   |             ERR!    |
| N/A   28C    P0     34W / 250W |  0MiB / 32510MiB |           0%    Default  |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Processes:                                                       GPU Memory |
|  GPU       PID    Type    Process name                               Usage    |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
| No running processes found                                     |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
| WARNING: infoROM is corrupted at gpu 0000:00:0D:0 |
```

Possible Causes

The health check failed so the GPU driver would not use or trust its content (some content is not used).

Impact

The ECC-related non-volatile data records may be affected. As a result, GPU memory pages that should have been retired are still in use.

Solution

1. If services are not affected, no operation is required.
2. If services are affected, stop the services, migrate ECSs, collect fault information by referring to [Fault Information Collection](#), and contact technical support.

7.5.2 What Do I Do If ECC Error "double bit ecc error" Occurs and There Are No Retired Pages Shown in the nvidia-smi -q Command Output?

Possible Causes

Errors may occur in the GPU memory.

Impact

GPU-related applications may be affected.

Solution

Run the **nvidia-smi** command to view the graphics card information.

Impact

GPU-related applications may be affected.

Solution

Run the `nvidia-smi` command to view the graphics card information.

```
root@nec2:~# /opt/cloud/cce/nvidia/bin/nvidia-smi
Tue Oct 25 09:32:59 2022
+-----+
| NVIDIA-SMI 510.48.03              Driver Version: 510.48.03      CUDA Version: 11.7.1      |
+-----+-----+-----+-----+-----+-----+
| GPU Name       Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC | |
| Fan  Temp  Perf  Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
|               |              |          |    MIG M.     |
+-----+-----+-----+-----+-----+-----+
| 0  Tesla V100-SXM2...  Off   | 00000000:2D:00.0 Off  |    100%      Default  |
| N/A   67C   P0   297W / 300W | 26984MiB / 32510MiB |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 1  Tesla V100-SXM2...  Off   | 00000000:32:00.0 Off  |    100%      Default  |
| N/A   69C   P0   294W / 300W | 23128MiB / 32510MiB |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 2  Tesla V100-SXM2...  Off   | 00000000:5B:00.0 Off  |    100%      Default  |
| N/A   70C   P0   287W / 300W | 23056MiB / 32510MiB |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 3  Tesla V100-SXM2...  Off   | 00000000:5F:00.0 Off  |    100%      Default  |
| N/A   59C   P0   304W / 300W | 23032MiB / 32510MiB |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 4  Tesla V100-SXM2...  Off   | 00000000:B5:00.0 Off  |     74%      Default  |
| N/A   42C   P0    94W / 300W | 4624MiB / 32510MiB  |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 5  Tesla V100-SXM2...  Off   | 00000000:BE:00.0 Off  |     0%       Default  |
| N/A   32C   P0    44W / 300W |    0MiB / 32510MiB  |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 6  Tesla V100-SXM2...  Off   | 00000000:E1:00.0 Off  |     0%       Default  |
| N/A   31C   P0    42W / 300W |    0MiB / 32510MiB  |              N/A     |
+-----+-----+-----+-----+-----+-----+
| 7  Tesla V100-SXM2...  Off   | 00000000:E9:00.0 Off  |     0%       Default  |
| N/A   30C   P0    43W / 300W |    0MiB / 32510MiB  |              N/A     |
+-----+-----+-----+-----+-----+-----+
```

- In the command output, if there are ECC errors in the **Volatile Uncorr. ECC**, run the `nvidia-smi -q -i &.{gpu_id}` command to view the graphics card details.
- In the command output, if there are no ECC errors in the **Volatile Uncorr. ECC**, run the `nvidia-smi -q` command to view all the graphics cards details.
- As shown in the figure, if only the value of **Device Memory** increases in the **Single Bit** under **Volatile** or the **Single Bit** under **Aggregate**, no operation is required.
- In **Single Bit** and **Double bit** under **Volatile** or **Single Bit** and **Double bit** under **Aggregate**, if the sum of **Register File**, **L1 Cache**, **L2 Cache**, **Texture Memory**, **Texture Shared**, and **CBU** is greater than **0**, contact fault information by referring to [Fault Information Collection](#) and contact technical support.

```
root@train-gpu-3-1-38a4c ~#
Ecc Mode
Current          : Enabled
Pending         : Enabled
ECC Errors
Volatile
Single Bit
Device Memory   : 268159
Register File   : 0
L1 Cache        : 0
L2 Cache        : 0
Texture Memory  : N/A
Texture Shared  : N/A
CBU             : N/A
Total           : 268159
Double Bit
Device Memory   : 1
Register File   : 0
L1 Cache        : 0
L2 Cache        : 0
Texture Memory  : N/A
Texture Shared  : N/A
CBU             : 0
Total           : 1
Aggregate
Single Bit
Device Memory   : 622281
Register File   : 0
L1 Cache        : 0
L2 Cache        : 0
Texture Memory  : N/A
Texture Shared  : N/A
CBU             : N/A
Total           : 622281
Double Bit
Device Memory   : 1
Register File   : 0
L1 Cache        : 0
L2 Cache        : 0
Texture Memory  : N/A
Texture Shared  : N/A
CBU             : 0
Total           : 1
root@train-gpu-3-1-38a4c ~#
```

7.5.4 What Do I Do If the GPU Is Disconnected or the Graphics Card Can't Be found, or rev ff Is Displayed After `lspci | grep -i nvidia` Is Executed?

Possible Causes

The health check failed so the GPU driver would not use or trust its content (some content is not used).

Impact

The ECC-related non-volatile data records may be affected. As a result, GPU memory pages that should have been retired are still in use.

Solution

1. Stop the services and perform service migration.
2. After the service is migrated, collect fault information by referring to [Fault Information Collection](#) and contact technical support.

7.5.5 What Do I Do If the `nvidia-smi` Command Output Shows Overheated GPUs?

Possible Causes

The heat dissipation of the graphics card is abnormal, or the fan is damaged.

Impact

Excessively high temperature of the graphics card affects the services.

Solution

Run the **nvidia-smi** command to check whether the fan is normal.

- If the fan speed is 0, the fan may be damaged. In this case, stop and migrate the service. After the service is migrated, collect fault information by referring to [Fault Information Collection](#) and contact technical support to check whether the hardware is faulty.

```
nvidia-smi
Thu Oct 13 15:19:58 2022

+-----+
| NVIDIA-SMI 450.51.05 | Driver Version: 450.51.05 | CUDA Version: 11.2 |
+-----+-----+
| GPU  Name   Perf  TCC/WDDM  Bus-Id  Disp.A  Volatile Uncorr. ECC  |
| Fan  Temp  Perf  Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M.  |
|-----+-----+-----+-----+-----+-----+-----+
| 0   NVIDIA RTX A4000  WDDM  00000000:21:01.0 Off  4205MiB / 16376MiB  99%    Default  |
| 0%   102C   P0    72W / 140W  |                    |            N/A      |
+-----+-----+-----+-----+-----+-----+
| Processes:                 |
| GPU  GI   CI           PID  Type  Process name          GPU Memory  |
| ID   ID   ID             ID   |                   | Usage       |
+-----+-----+-----+-----+-----+-----+

```

- If **ERR!** is displayed in the command output, the graphics card may be overheated. Stop the service, wait until the graphics card cools down, and run the **nvidia-smi** command to check whether the **ERR!** disappears.
 - If the command output is normal, adjust the service to limit the maximum power of the graphics card.
 - If the fault persists, collect fault information by referring to [Fault Information Collection](#) and contact technical support.

```
NVIDIA-SMI 450.51.05 | Driver Version: 450.51.05
+-----+-----+-----+-----+-----+-----+
| GPU  Name   Persistence-M  Bus-Id  Disp.A  Volatile Uncorr. ECC  |
| Fan  Temp  Perf  Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M.  |
|-----+-----+-----+-----+-----+-----+-----+
| 0   TITAN Xp  Off          00000000:02:00.0 Off  12196MiB  0%    Default  |
| ERR!  82C   P2    ERR! / 250W  |                    |            N/A      |
+-----+-----+-----+-----+-----+-----+

```

7.5.6 What Do I Do If "Unable to load the kernel module 'nvidia.ko'" Is Displayed During Driver Installation?

Possible Causes

1. The bus is disconnected.

```
# dmesg | grep -i NURM
[ 86.886171] NURM: The NVIDIA GPU (PCI ID: 0000:02:00.0) installed in this system has
NURM: fallen off the bus and is not responding to commands.
[ 86.827881] NURM: The NVIDIA probe routine failed for 1 device(s).
[ 86.827882] NURM: None of the NVIDIA devices were initialized.
[ 628.667355] NURM: The NVIDIA GPU (PCI ID: 0000:02:00.0) installed in this system has
NURM: fallen off the bus and is not responding to commands.
```

2. The kernel versions are inconsistent.

Impact

The graphics card driver fails to be installed and the graphics card cannot be used.

Solution

1. Run the following command to check whether the kernel versions are consistent:
rpm -qa | grep gcc #Check the GCC version.
rpm -qa | grep kernel-devel #Check the kernel-devel version.
 - If the kernel versions are inconsistent, reinstall the driver.
 - If the kernel versions are consistent, go to the next step.
2. View the dmesg log and check whether there are NVRM errors.
 - If the error message "NVRM: fallen off the bus and is not responding to commands" is displayed, the bus is disconnected. Then, contact technical support to replace the graphics card.
 - If the fault persists, collect fault information by referring to [Fault Information Collection](#) and contact technical support.

7.5.7 What Can I Do If an Xid Error Is Displayed in the Message Log When a GPU-accelerated ECS Is Faulty?

Possible Causes

XID	Description
32	Invalid or corrupted push buffer stream
74	NVLINK Error, which indicates that the GPU hardware is faulty and needs to be brought offline for repair.
79	GPU has fallen off the bus, which indicates that the bus is disconnected and needs to be brought offline for repair.

For details, see <https://docs.nvidia.com/deploy/xid-errors/index.html>.

Solution

1. Run the **dmesg | grep -i xid** command to check whether there are Xid errors.

2. Stop the services, perform service migration, collect fault information by referring to [Fault Information Collection](#), and contact technical support.

8 GPU Driver Issues

8.1 Why Is the GPU Driver Abnormal?

Symptom

When you run the following command on a GPU-accelerated ECS to view the GPU usage, the system displays a message indicating that the specified program cannot be executed or the file path does not exist.

nvidia-smi

Information similar to the following is displayed:

```
-bash: /bin/nvidia-smi: No such file or directory
```

or

```
nvidia-smi: command not found
```

Possible Causes

The ECS driver is abnormal, not installed, or uninstalled.

Solution

- If the GPU driver is not installed, install it.
For details, see [GPU Driver](#).
- If the driver has been uninstalled:
Run the **history** command to check whether an uninstallation has been performed.
Go to the **/var/log** directory and check whether the **nvidia-uninstall.log** file exists. If the log exists, the GPU driver has been uninstalled. Reinstall the GPU driver.
- If the driver has been installed but the driver status is abnormal:
 - a. Uninstall the driver.

- Method 1: Run the **nvidia-uninstall** command to uninstall the driver. If the system displays a message indicating that the command does not exist, use [method 2](#).
- Method 2: Run the **whereis nvidia** command to query the version of the driver installed on the ECS.

Figure 8-1 Installed driver version

```
[root@ ~]# whereis nvidia
nvidia: /usr/share/nvidia /usr/src/nvidia-396.44/nvidia
[root@ ~]#
```

Download the driver package of the same version as the obtained one from the [NVIDIA official website](#). (This driver package is required when you uninstall and reinstall the driver.)

For example, if the driver version is nvidia-396.44, run the **sh NVIDIA-Linux-x86_64-396.44.run --uninstall** command to uninstall the driver.

- Reinstall the driver.
For details, see [GPU Driver](#).

8.2 Why Is the GPU Driver Unavailable?

Symptom

The following issues occur when the **nvidia-smi** command is used to check the GPU usage.

NVIDIA-SMI has failed because it couldn't communicate with the NVIDIA driver. Make sure that the latest NVIDIA driver is installed and running.

Figure 8-2 GPU driver unavailable

```
[root@ ~]# nvidia-smi
NVIDIA-SMI has failed because it couldn't communicate with the NVIDIA driver. Make sure that the latest NVIDIA driver is installed and running.
```

Possible Causes

The system kernel is upgraded, resulting in GPU driver unavailability.

Troubleshooting

Run the corresponding command on the server to check the version of the kernel when the driver was installed:

- CentOS: **find /usr/lib/modules -name nvidia.ko**
- Ubuntu: **find /lib/modules -name nvidia.ko**

For example, run the preceding command in CentOS. If the command output shown in [Figure 8-3](#) is displayed, the GPU driver is installed on the 3.10.0-957.5.1.el7.x86_64 kernel.

Figure 8-3 Version of the kernel where the driver is installed

```
[root@ecs-6711 ~]# find /usr/lib/modules -name nvidia.ko
/usr/lib/modules/3.10.0-957.21.3.el7.x86_64/kernel/drivers/video/nvidia.ko
[root@ecs-6711 ~]#
```

Run the **uname -r** command. The command output shown in **Figure 8-4** indicates that the current kernel version is 3.10.0-1160.24.1.el7.x86_64.

Figure 8-4 Current kernel version

```
[root@ecs-6838 ~]# uname -r
3.10.0-1160.24.1.el7.x86_64
[root@ecs-6838 ~]#
```

The version of the kernel where the driver is installed is different from the current kernel version.

Solution

- Method 1: Restart the ECS and select the kernel version used when the GPU driver was installed.
 - a. In the ECS list, locate the row that contains the target ECS and click **Remote Login** in the **Operation** column. In the displayed dialog box, click **Log In** in the **Other Login Modes** area.
 - b. Click **Ctrl+Alt+Del** in the upper part of the remote login panel to restart the ECS.
 - c. Refresh the page quickly and press the up and down arrow keys to stop the ECS from restarting. Then, select the kernel version used when the GPU driver was installed and press **Enter** to enter the system. The GPU driver becomes available in the current kernel version.
- Method 2: Reinstall the driver based on the new kernel version.
 - a. Uninstall the driver.
 - Method 1: Run the **nvidia-uninstall** command to uninstall the driver. If the system displays a message indicating that the command does not exist, use **method 2**.
 - Method 2: Run the **whereis nvidia** command to query the version of the driver installed on the ECS.

Figure 8-5 Installed driver version

```
[root@ecs-6838 ~]# whereis nvidia
nvidia: /usr/share/nvidia /usr/src/nvidia-396.44/nvidia
[root@ecs-6838 ~]#
```

Download the driver package of the same version as the obtained one from the **NVIDIA official website**. (This driver package is required when you uninstall and reinstall the driver.)

For example, if the driver version is `nvidia-396.44`, run the `sh NVIDIA-Linux-x86_64-396.44.run --uninstall` command to uninstall the driver.

- b. Reinstall the driver.

For details, see [GPU Driver](#).

8.3 Why Is the GPU Display Abnormal?

Symptom

The following issues occur when the `nvidia-smi` command is used to check the GPU usage.

- On the device with one GPU, the following information is displayed:
No devices were found
- On the device with multiple GPUs, a message indicating that the number of GPUs is incomplete is displayed.

The `lspci | grep -i nvidia` command output shows that the number of GPUs is normal.

Solution

1. Check whether the ECS, for example, of the PI2 or G6 flavor, is using NVIDIA Tesla T4 GPUs.
 - If yes, rectify the fault by following the instructions provided in [Why Is the T4 GPU Display Abnormal?](#)
 - If no, go to the next step.
2. Check the system log `/var/log/message` for any reported driver-related errors.
 - If the error message "Failed to copy vbios to system memory" is displayed, the possible causes may be frequent driver loading/unloading. You are advised to enable the driver's persistence mode to keep the driver in the loading state.

Figure 8-6 System logs

```
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: Failed to copy vbios to system memory.
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: RmInitAdapter failed! (0x30:0xffff:852)
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: rm_init_adapter failed, device minor number 1
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: Failed to copy vbios to system memory.
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: RmInitAdapter failed! (0x30:0xffff:852)
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: rm_init_adapter failed, device minor number 1
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: Failed to copy vbios to system memory.
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: RmInitAdapter failed! (0x30:0xffff:852)
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: rm_init_adapter failed, device minor number 0
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: Failed to copy vbios to system memory.
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: RmInitAdapter failed! (0x30:0xffff:852)
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0d.0: rm_init_adapter failed, device minor number 0
Aug 26 14:35:35 gpu-002 kernel: NVRM: GPU 0000:00:0e.0: Failed to copy vbios to system memory.
```

- i. Run the following command to enable the driver's persistence mode:
`nvidia-smi -pm 1`
- ii. Run the following command to open and edit the `/etc/rc.local` file:
`vim /etc/rc.local`
- iii. Configure automatic startup and write the `nvidia-smi -pm 1` command to the `/etc/rc.local` file.

- iv. Press **Esc**, enter **:wq**, and press **Enter** to save the settings and exit.
 - v. Run the following command to add startup permissions:
chmod +x /etc/rc.d/rc.local
 - If "Failed to copy vbios to system memory" is not displayed, go to the next step.
3. Check whether the ECS uses Tesla 510.xx.xx.
 - If yes, the driver version may be incompatible with the image. You are advised to change the driver version. For details, see [GPU Driver](#).
 - If no, go to the next step.
4. Restart the ECS and run the **nvidia-smi** command to check whether the GPU usage is normal.
If the fault persists, contact customer service.

8.4 Why Is the T4 GPU Display Abnormal?

Symptom

An ECS, for example, of the PI2 or G6 flavor, is using NVIDIA Tesla T4 GPUs. However, when you run the **nvidia-smi** command to check the GPU usage, the following information is displayed:

```
No devices were found
```

Possible Causes

The NVIDIA Tesla T4 GPU is a new version of the NVIDIA GPU. By default, the GSP firmware is enabled and used. As a result, the GPU cannot be identified.

Method 1

NOTE

The following settings will become invalid after the ECS is restarted.

1. Remove the NVIDIA kernel module.
rmmod nvidia_drm
rmmod nvidia_modeset
rmmod nvidia
2. Disable GSP firmware and load the NVIDIA kernel module.
modprobe nvidia NVreg_EnableGpuFirmware=0
modprobe nvidia_drm
modprobe nvidia_modeset
3. If the fault persists, contact customer service.

Method 2

1. Run the following command to open the **/etc/modprobe.d/nvidia.conf** file:
vim /etc/modprobe.d/nvidia.conf
Press **i** to enter the editing mode.

2. Add the following information to `/etc/modprobe.d/nvidia.conf`:
options nvidia NVreg_EnableGpuFirmware=0
Press **Esc** and enter **:wq!** to save the modification and exit.
3. Run the following command to restart the ECS:
reboot
4. If the fault persists, contact customer service.

8.5 How Do I Troubleshoot GPU Startup Failures Caused by NULL Pointer Dereference on NVIDIA?

Symptom

A GPU instance startup failed. The system log shows "Unable to handle kernel NULL pointer dereference at 0000000000000008", as shown in [Figure 8-7](#).

Figure 8-7 NVIDIA driver NULL pointer access

```
[21152.003950] nvidia 0000:21:01.0: irq 42 for MSI/MSI-X
[21152.003971] nvidia 0000:21:01.0: irq 43 for MSI/MSI-X
[21152.003992] nvidia 0000:21:01.0: irq 44 for MSI/MSI-X
[21152.532634] BUG: unable to handle kernel NULL pointer dereference at 0000000000000008
[21152.532836] IP: [] _nv007834rm+0x14/0xd0 [nvidia]
[21152.533159] PGD 83be18067 PUD 0
[21152.533257] Oops: 0000 [#1] SMP
[21152.533348] kbox catch die event.
[21152.535048] collected_nen = 1/9720, LOG_BUF_LEN_LOCAL = 1048576
[21152.535862] kbox: notify die begin
```

Possible Causes

The GPU driver is abnormal.

Solution

1. Uninstall the driver.
 - Method 1: Run the **nvidia-uninstall** command to uninstall the driver.
If the system displays a message indicating that the command does not exist, use [method 2](#).
 - Method 2: Run the **whereis nvidia** command to query the version of the driver installed on the ECS.

Figure 8-8 Installed driver version

```
[~]# whereis nvidia
nvidia: /usr/share/nvidia /usr/src/nvidia-396.44/nvidia
[~]#
```

Download the driver package of the same version as the obtained one from the [NVIDIA official website](#). (This driver package is required when you uninstall and reinstall the driver.)

For example, if the driver version is `nvidia-396.44`, run the **sh NVIDIA-Linux-x86_64-396.44.run --uninstall** command to uninstall the driver.

2. Reinstall the driver.
For details, see [Installing a Driver and Toolkit](#).

9 Multi-User Login Issues

9.1 How Do I Configure Multi-User Logins for an ECS Running Windows Server 2012?

Scenarios

This section uses an ECS running Windows Server 2012 as an example to describe how to enable concurrent logins by multiple users.

An ECS running Windows Server 2012, by default, allows concurrent logins by two users. To allow the logins by more users, configure the remote desktop session host and remote desktop licensing.

Precautions

- Ensure that the ECS bandwidth resources are sufficient to prevent ECS freezing or login exceptions caused by heavy loads.
- The port for logging in to the ECS must have been enabled in the inbound direction of the security group which the ECS belongs to. By default, port 3389 is used.
- The target ECS already has an EIP bound.
- After multi-user logins are configured, different users can log in to the ECS without affecting each other.
- The configuration in this section enables concurrent logins by multiple users or multiple concurrent logins using one account. A remote desktop license is valid for only 120 days. After the license expires, multi-user logins will be unavailable. For instructions about how to activate Remote Desktop Licensing, see [How Do I Apply for a License for Authenticating Multi-User Sessions and Activate an ECS?](#)
- After multi-user logins are configured on a Windows ECS, the browser may fail to be opened by multiple users. For details about how to handle the problem, see [Why Does a Browser Launch Error Occur in Multi-User Login?](#)

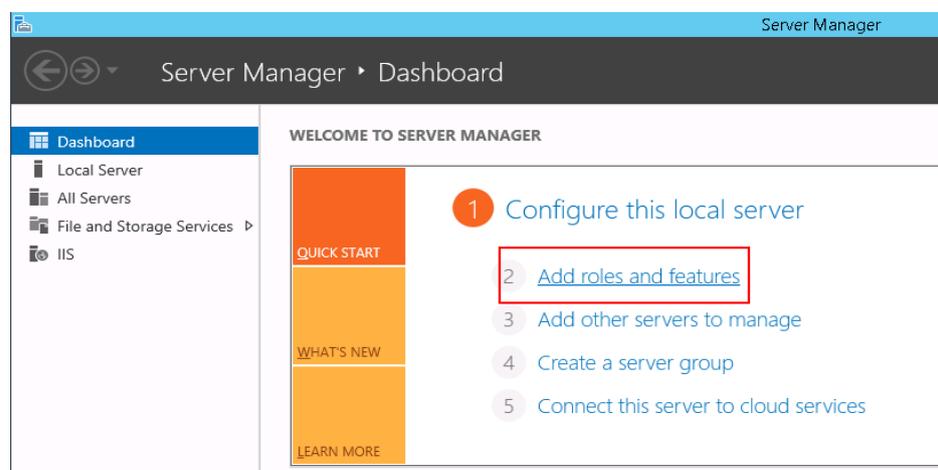
Procedure

1. [Configuring Remote Desktop Session Host and Remote Desktop Licensing](#)
2. [Enabling Multi-User Logins](#)
3. [Adding a New User to the Remote Desktop Users Group](#)

Configuring Remote Desktop Session Host and Remote Desktop Licensing

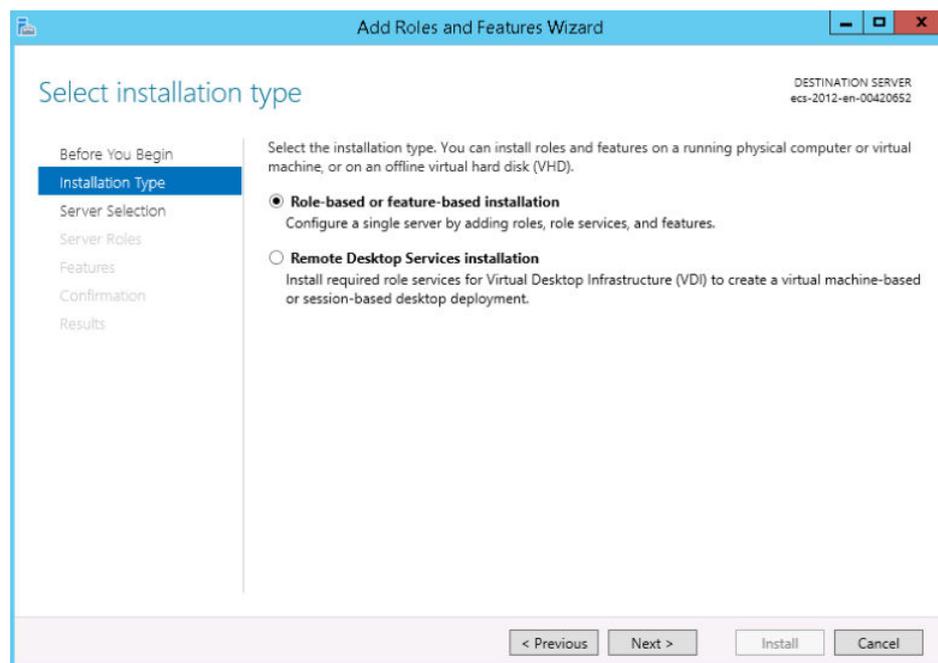
1. Log in to the Windows ECS.
2. On the OS, click  to open **Server Manager**. Click **Add roles and features**.

Figure 9-1 Adding roles and features



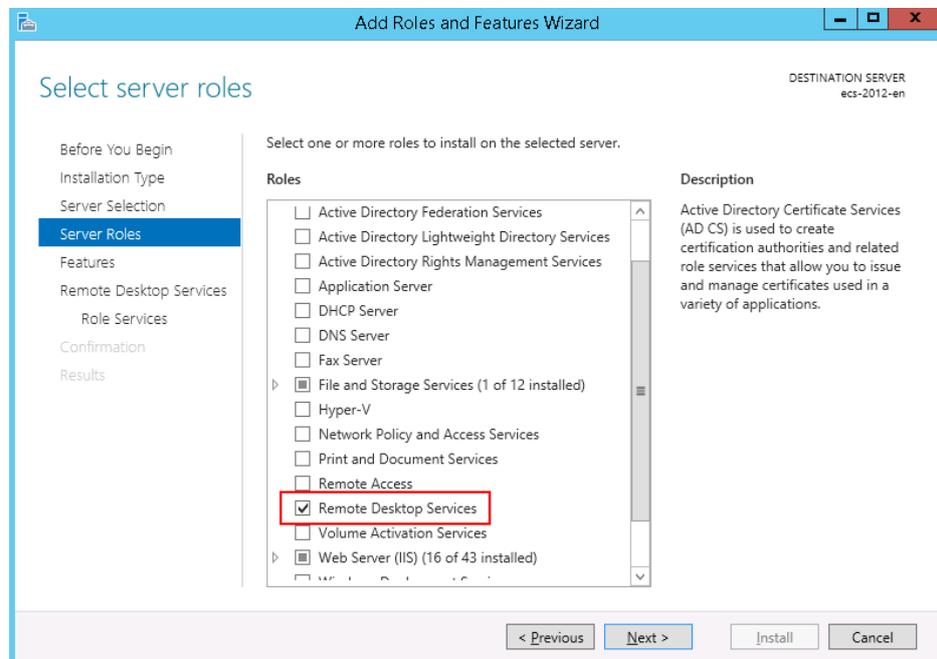
3. Retain the default settings and click **Next**. On the displayed installation page, select **Role-based or feature-based installation** and click **Next**.

Figure 9-2 Selecting an installation type



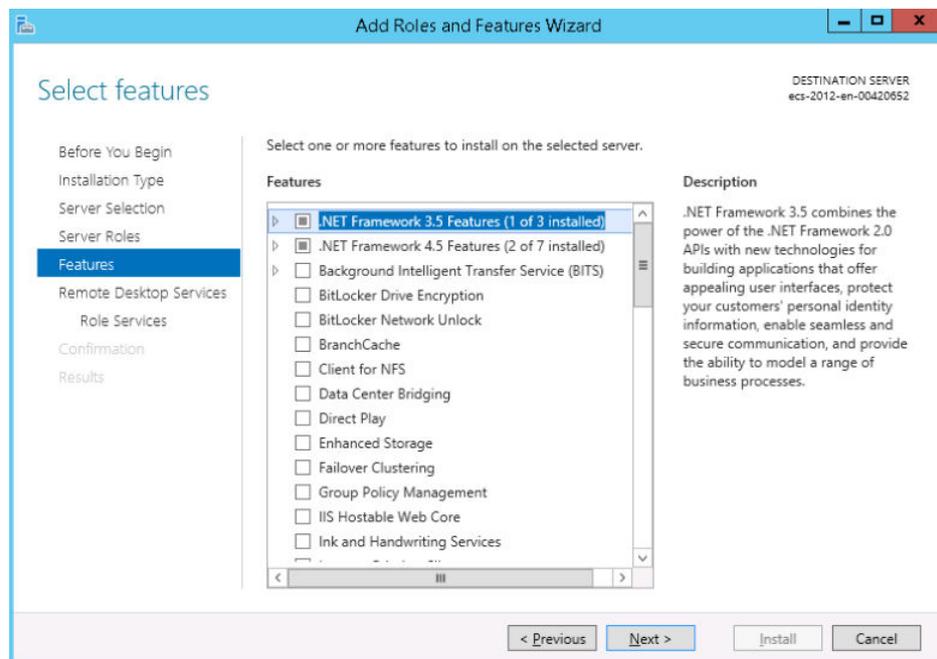
4. Choose **Server Selection**, select **Select a server from the server pool**, and click **Next**.
5. Choose **Server Roles**, select **Remote Desktop Services**, and click **Next**.

Figure 9-3 Selecting Remote Desktop Services



6. On the **Features** page, retain the default settings and click **Next** twice.

Figure 9-4 Features



7. On the **Role Services** page, select **Remote Desktop Session Host** and **Remote Desktop Licensing**. In the displayed dialog box, click **Add Features** and then **Next**.

Figure 9-5 Adding features

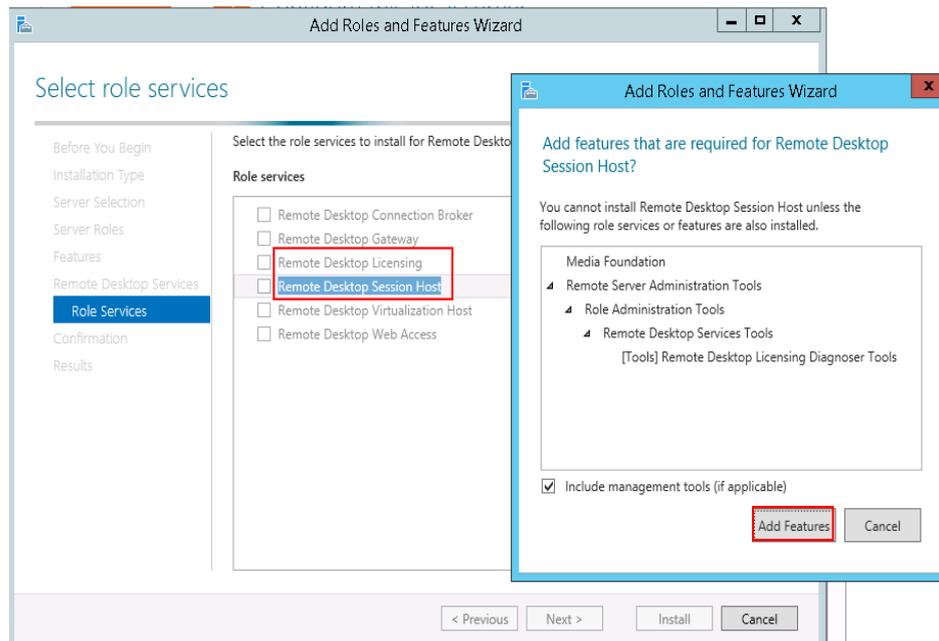
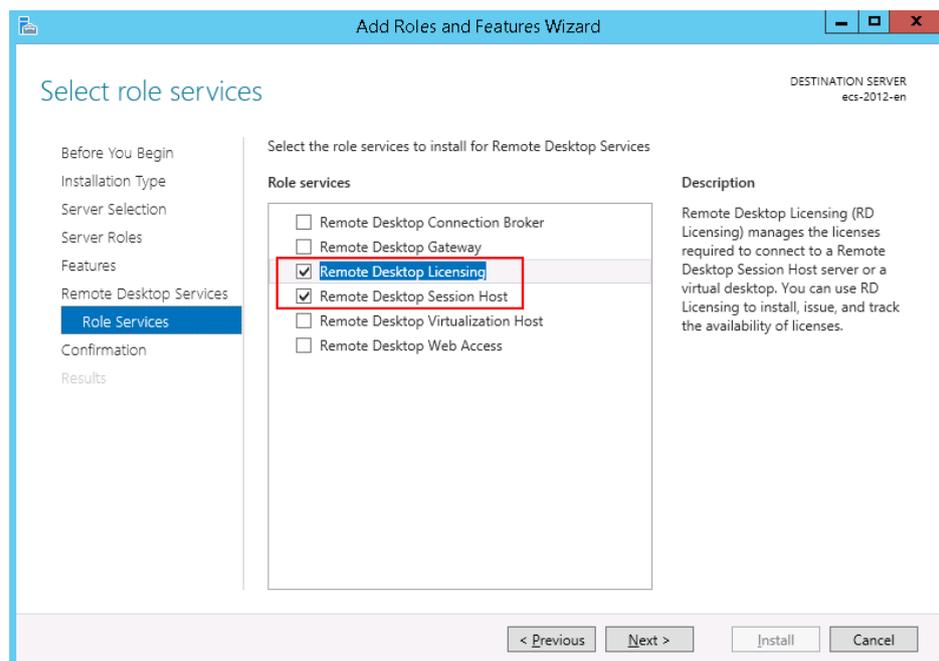


Figure 9-6 Selecting Remote Desktop Licensing



8. Confirm installation selections and click **Install**.

Figure 9-7 Confirming installation selections

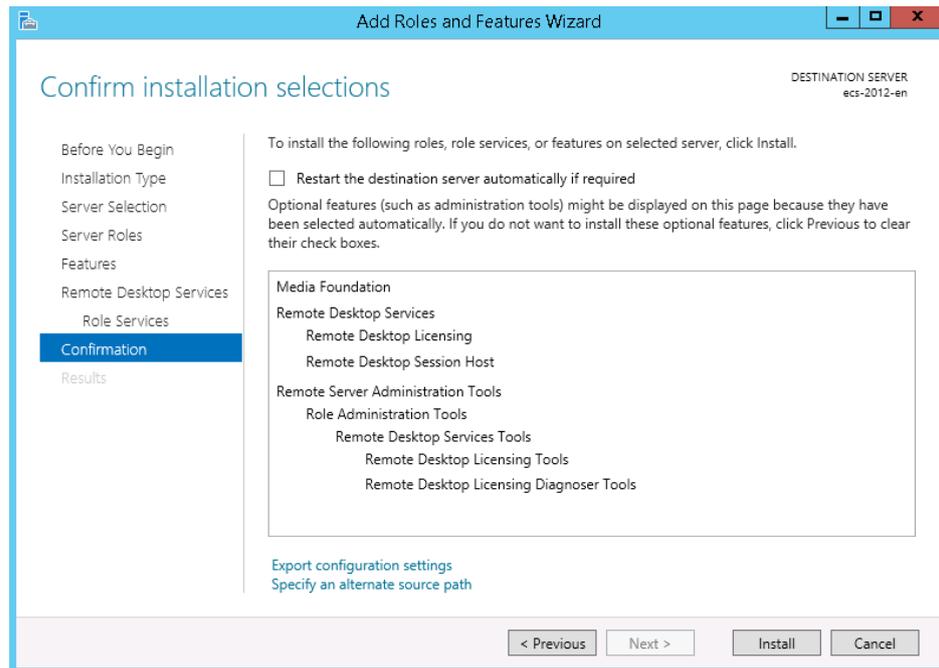
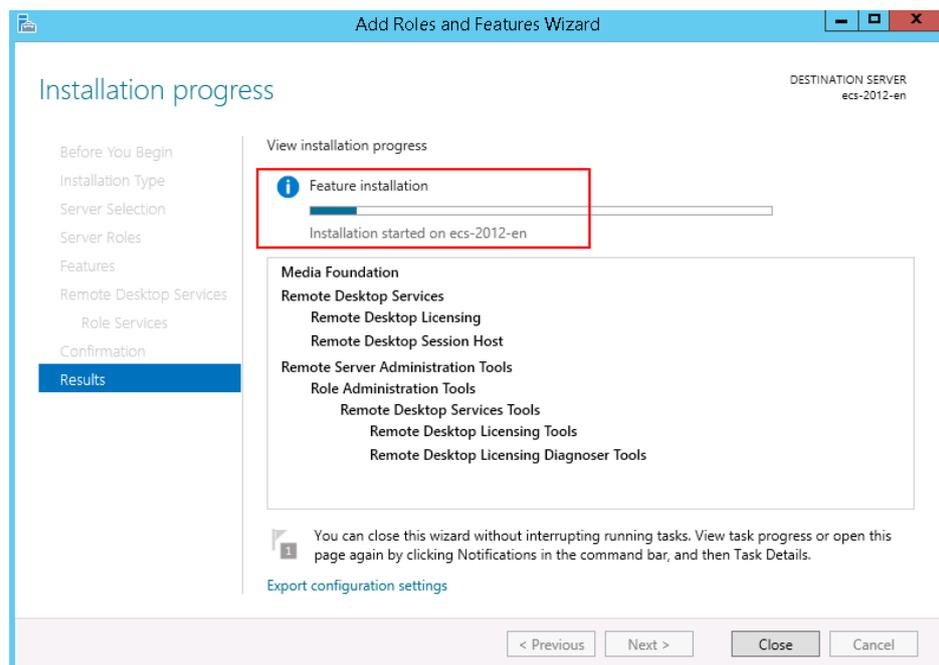
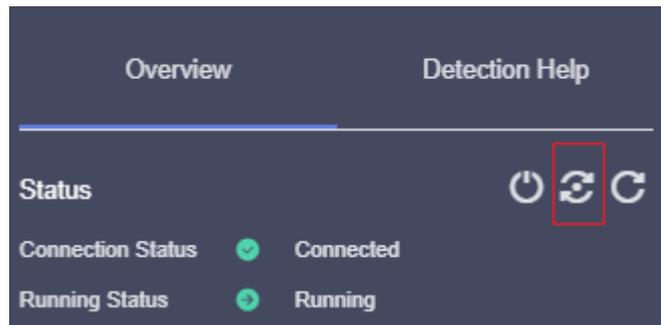


Figure 9-8 Installation progress



9. After the installation is complete, restart the ECS.

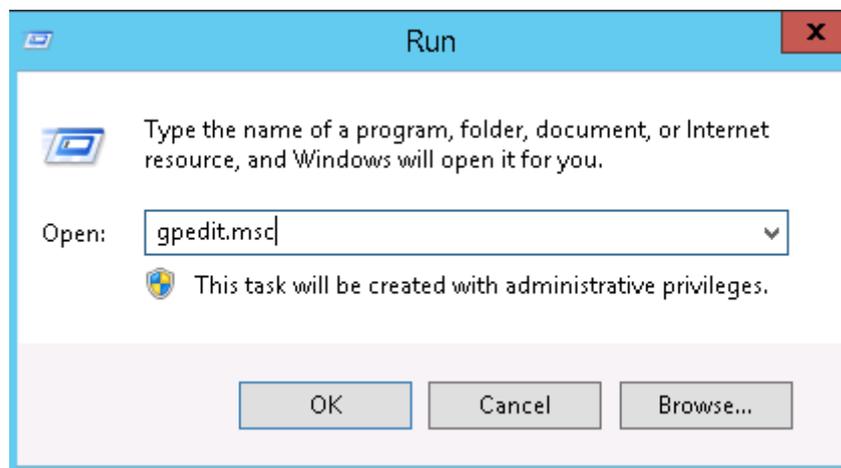
Figure 9-9 Restarting the ECS



Enabling Multi-User Logins

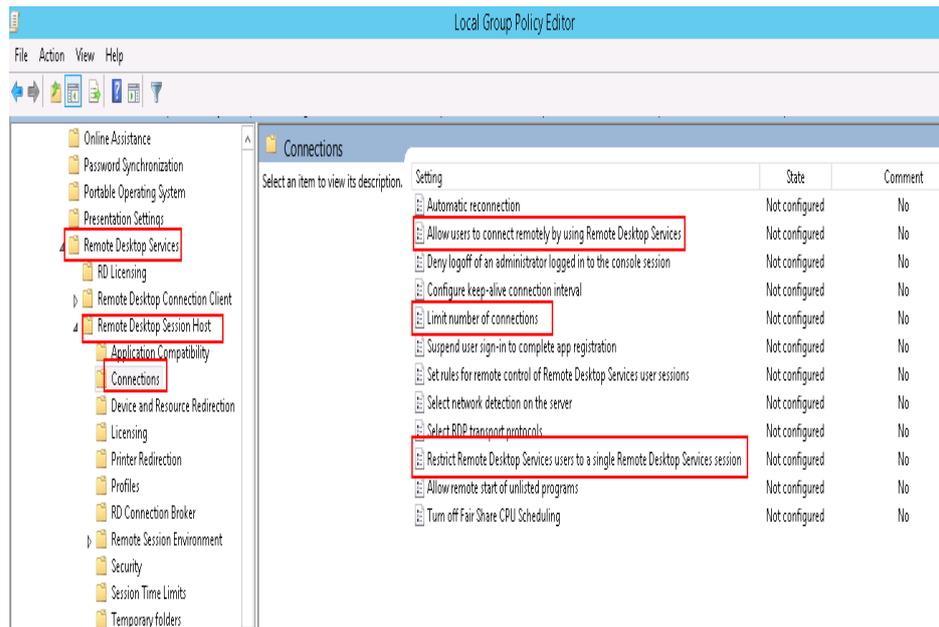
1. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start Local Group Policy Editor.

Figure 9-10 gpedit.msc



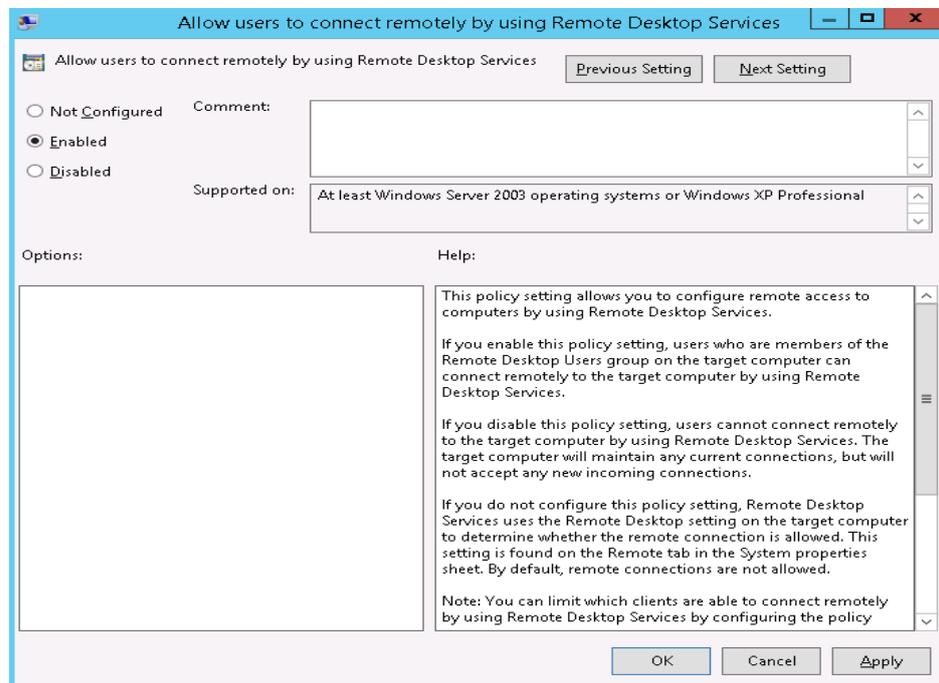
2. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connections**. Configure **Allow user to connect remotely by using Remote Desktop Services**, **Limit number of connections** (based on site requirements), and **Restrict Remote Desktop Services users to a single Remote Desktop Services session**.

Figure 9-11 Configuring Connections

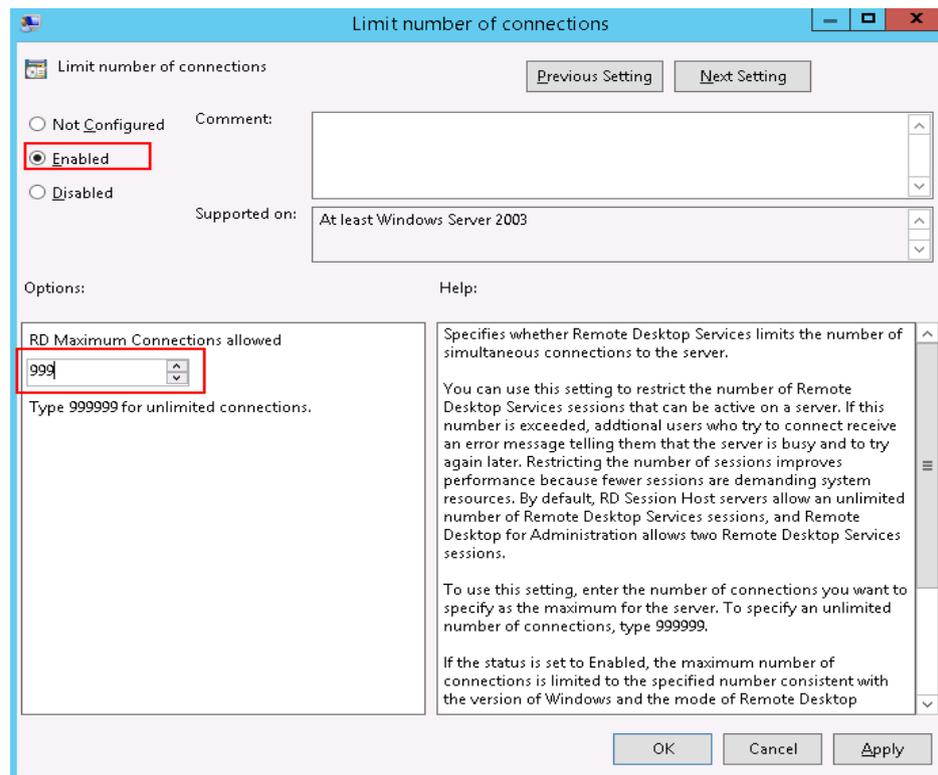


3. Right-click **Allow users to connect remotely by using Remote Desktop Services** and choose **Edit**. Then, set the status to **Enabled**.

Figure 9-12 Allowing users to connect remotely by using Remote Desktop Services



4. Set **Limit number of connections** to **Enabled** and set the number based on site requirements.

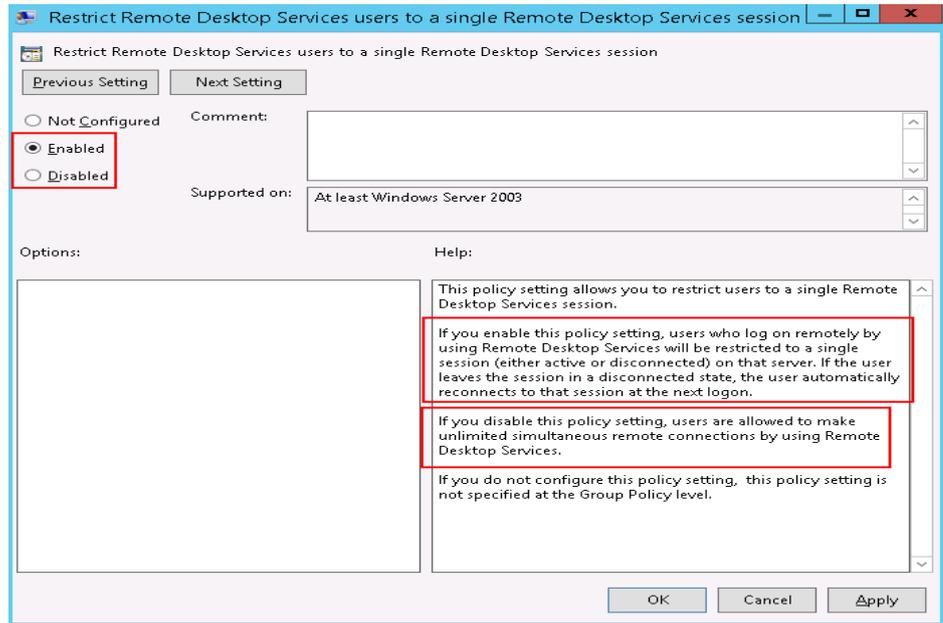
Figure 9-13 Setting limit number of connections

5. Set **Restrict Remote Desktop Services users to a single Remote Desktop Services session** to **Enabled** or **Disabled** as required. In this example, **Enabled** is selected.
 - **Enabled:** allows concurrent logins by multiple users but not using the same account.

For example, users A, B, and C can concurrently log in to an ECS using account A, B, and C, respectively. However, users A, B, and C cannot use the same account to log in to the ECS.
 - **Disabled:** allows multiple users to concurrently log in to an ECS using one account.

For example, users A, B, and C can use the same account to concurrently log in to an ECS.

Figure 9-14 Restricting Remote Desktop Services users to a single Remote Desktop Services session



6. Run `cmd` and then `gpupdate /force` to forcibly start Local Group Policy Editor and restart the ECS.

Adding a New User to the Remote Desktop Users Group

After enabling multi-user login, add new users to the remote desktop users group. This section describes how to create a user and add the user to remote desktop users group.

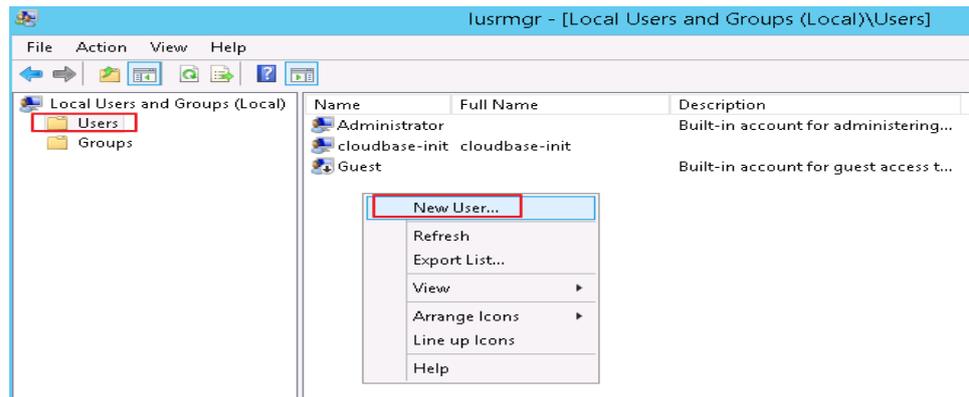
1. Open the **Run** dialog box, enter `lusrmgr.msc`, and click **OK** to open **Local Users and Groups**.

Figure 9-15 Local Users and Groups



2. Select **Users**, right-click the blank area, and choose **New User** from the shortcut menu.

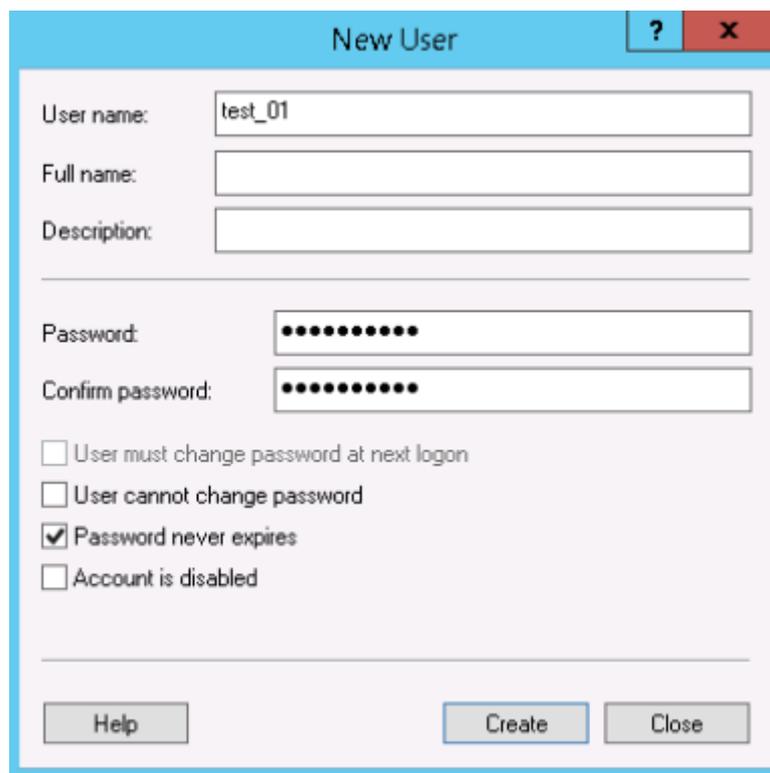
Figure 9-16 Adding a new user



3. Enter the username and password of the new user, confirm the password, and click **Create**.

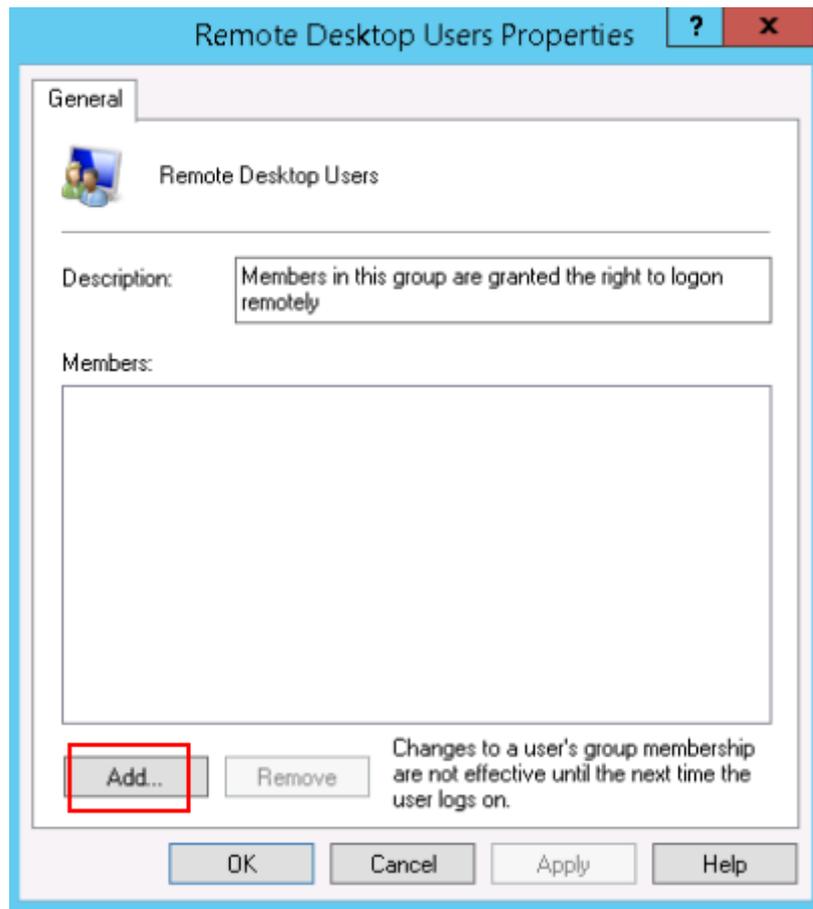
The password and the confirmation password must be the same. For security purposes, you are advised to set security attributes for the password.

Figure 9-17 Setting user information



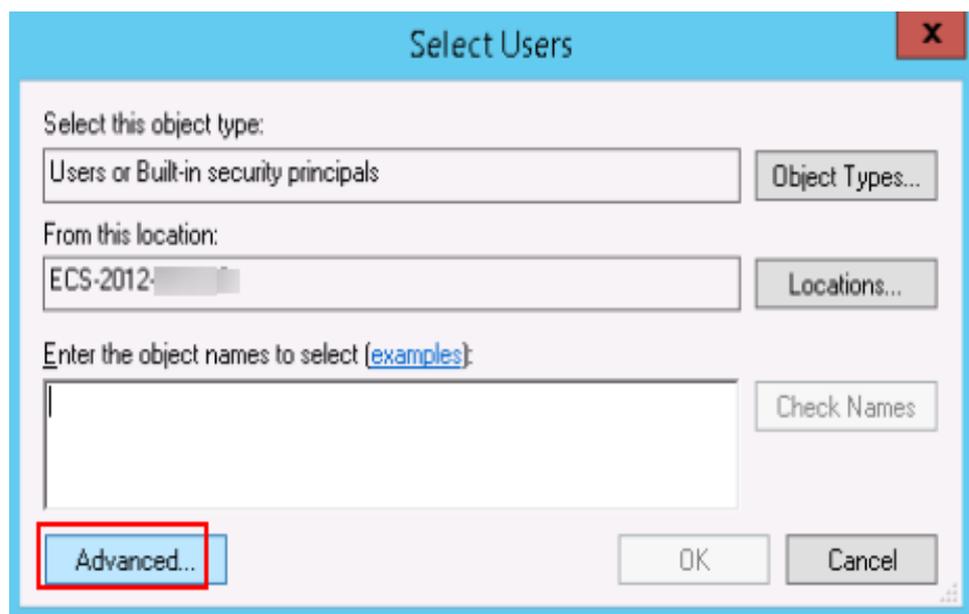
4. Select **Groups**, double-click **Remote Desktop Users**, and click **Add**.

Figure 9-18 Remote Desktop Users Properties



5. In the **Select Users** dialog box, click **Advanced**.

Figure 9-19 Selecting users (Advanced)



6. Click **Find Now**, select the user for remote login in the search results, and click **OK**.

Figure 9-20 Selecting users (Find Now)

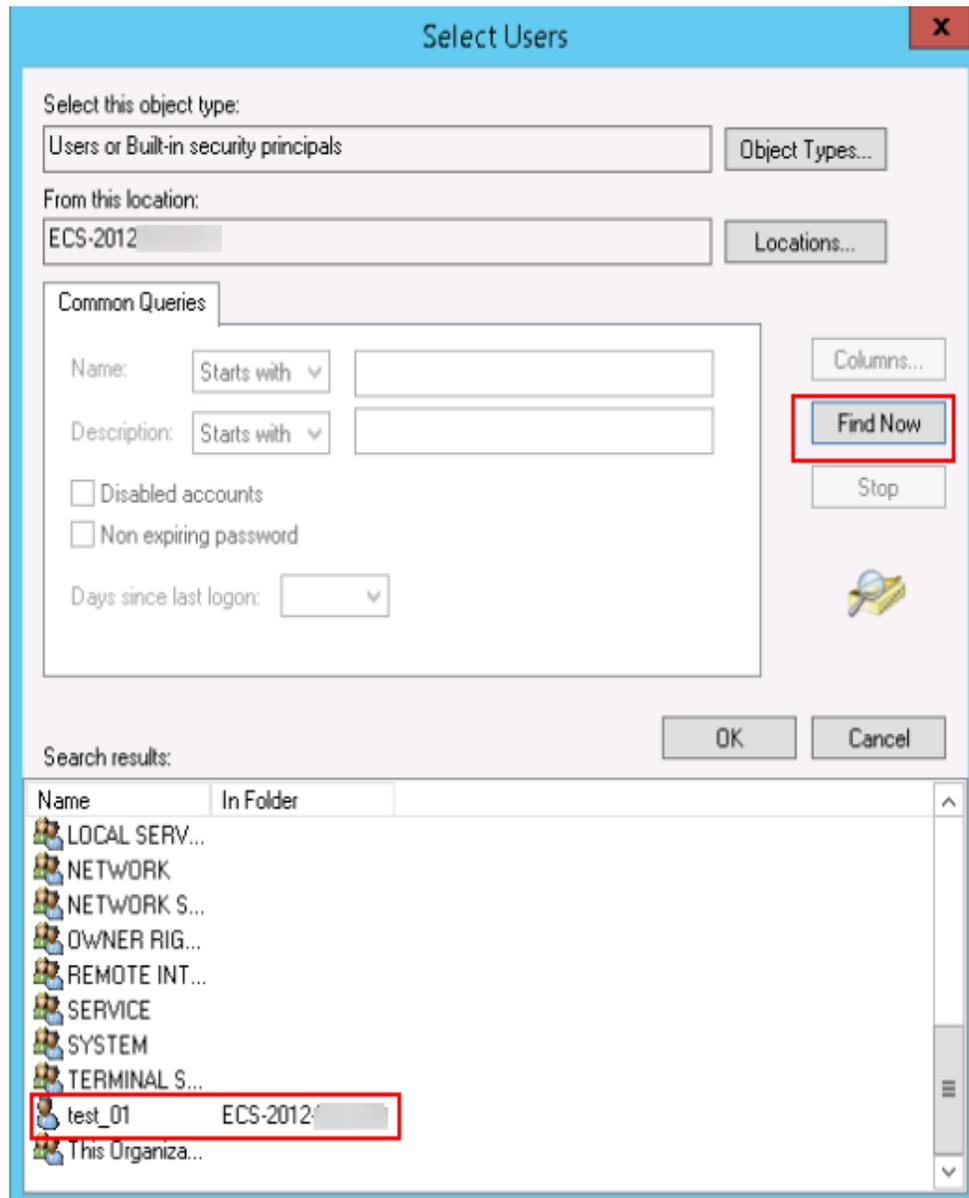
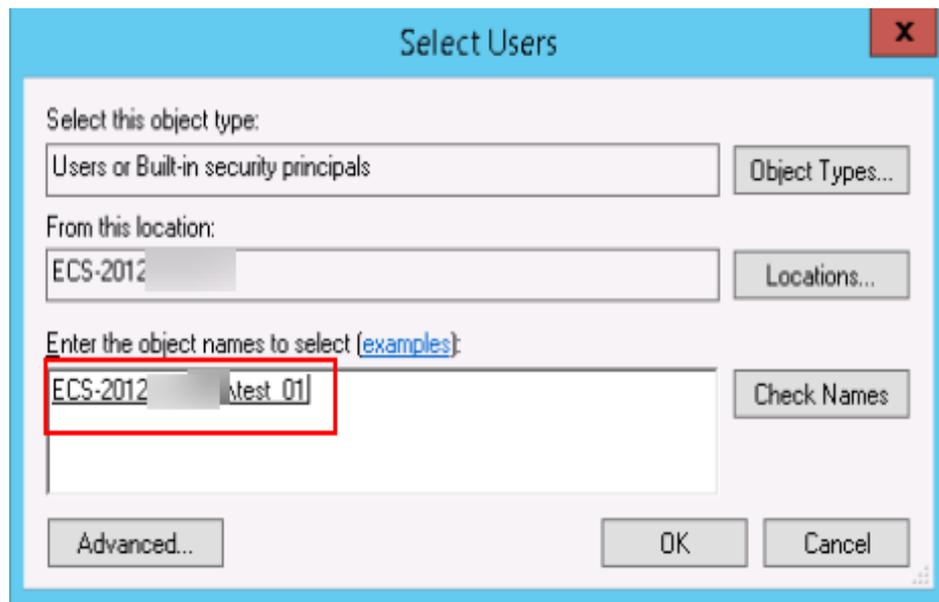
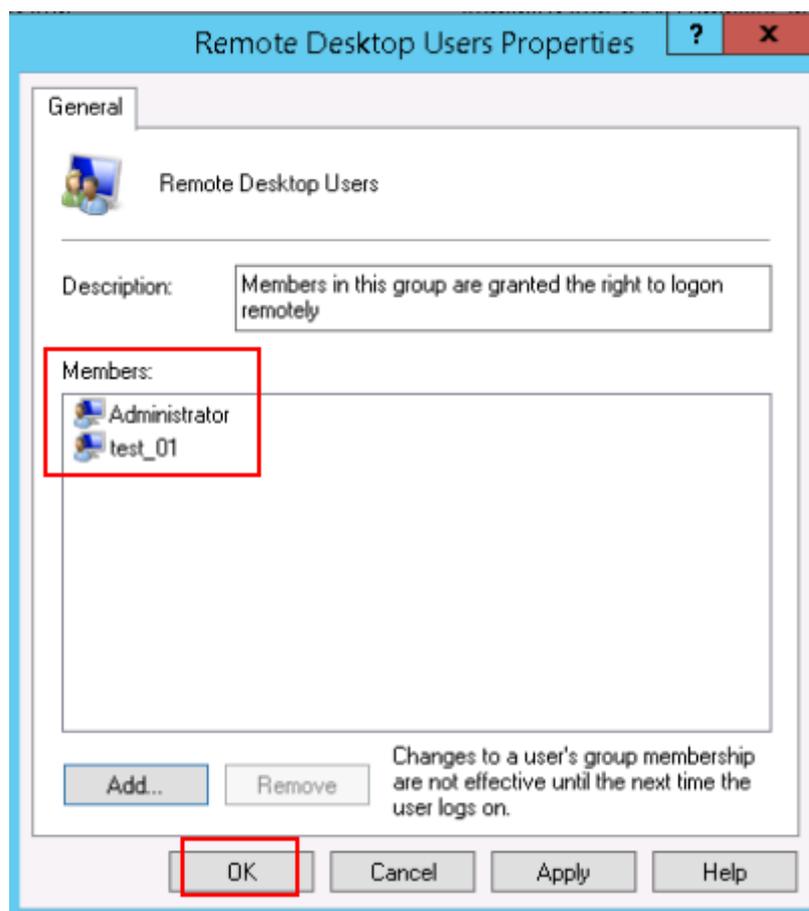


Figure 9-21 Adding a user



7. Click **OK** to add the user to the **Remote Desktop Users** group.

Figure 9-22 Confirming user information



Follow-up Operations

For instructions about how to activate Remote Desktop Licensing, see [How Do I Apply for a License for Authenticating Multi-User Sessions and Activate an ECS?](#)

9.2 Why Does a Browser Launch Error Occur in Multi-User Login?

Symptom

When multiple users log in to an ECS running Windows Server 2008, Windows Server 2012, or Windows Server 2016, a user opens the browser and other users failed.

Solution

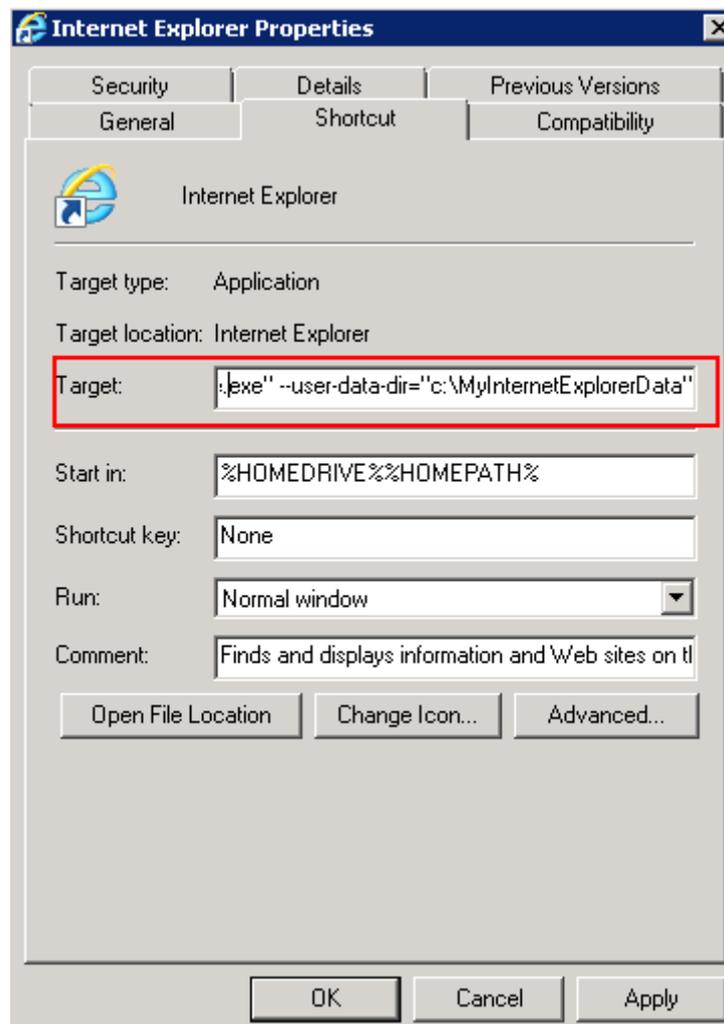
This section uses Internet Explorer as an example.

1. Right-click the browser icon and choose **Create shortcut**.
2. Right-click the created shortcut and choose **Properties**.
3. On the **Shortcut** tab, add the following content to the end of **Target**.

```
--user-data-dir="c:\MyInternetExplorerData"
```

NOTE

- There is a space between **.exe** and **--user**.
- **c:\MyInternetExplorerData** indicates the directory where the Internet Explorer data files are stored. You can set it to any valid directory. If this directory does not exist, the browser automatically creates one.

Figure 9-23 Internet Explorer Properties

4. Save the modification and multiple users can use the browser at the same time.

9.3 How Do I Apply for a License for Authenticating Multi-User Sessions and Activate an ECS?

Scenarios

This section describes how to configure remote desktop services and activate an ECS.

The ECS running Windows Server 2012 is used as an example.

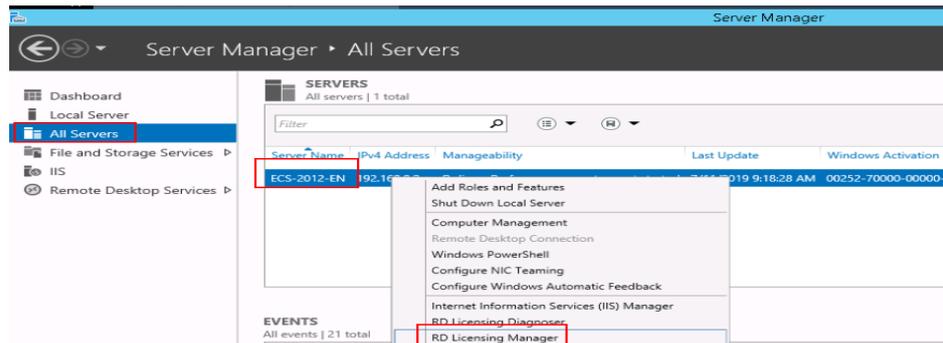
Procedure

1. [Applying for a License for Authenticating Multi-User Sessions](#)
2. [Activating the ECS](#)
3. [Configuring the Licensing Server for Remote Desktop Session Host](#)

Applying for a License for Authenticating Multi-User Sessions

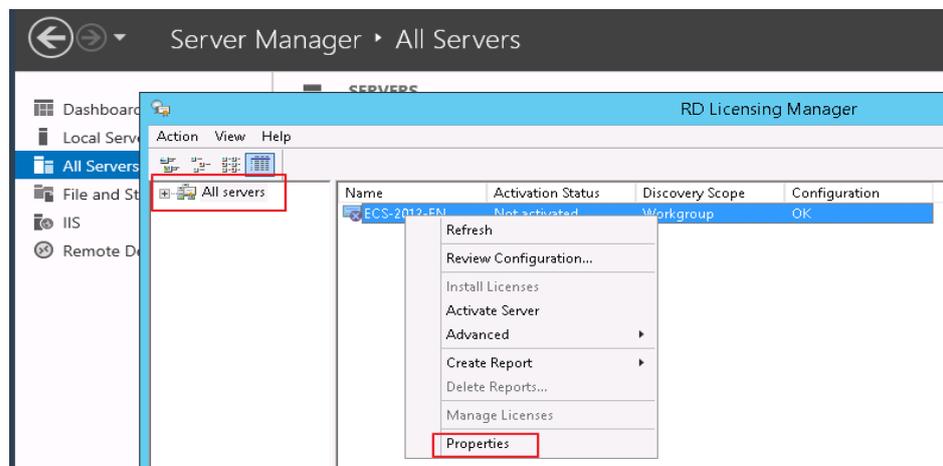
1. Log in to the Windows ECS.
2. On the OS, click  to open **Server Manager**.
3. On the **Server Manager** page, click **All Servers**, right-click the server name, and choose **RD Licensing Manager** from the shortcut menu.

Figure 9-24 Selecting RD Licensing Manager



4. Right-click the unactivated server and choose **Properties** from the shortcut menu.

Figure 9-25 Selecting properties of the unactivated server

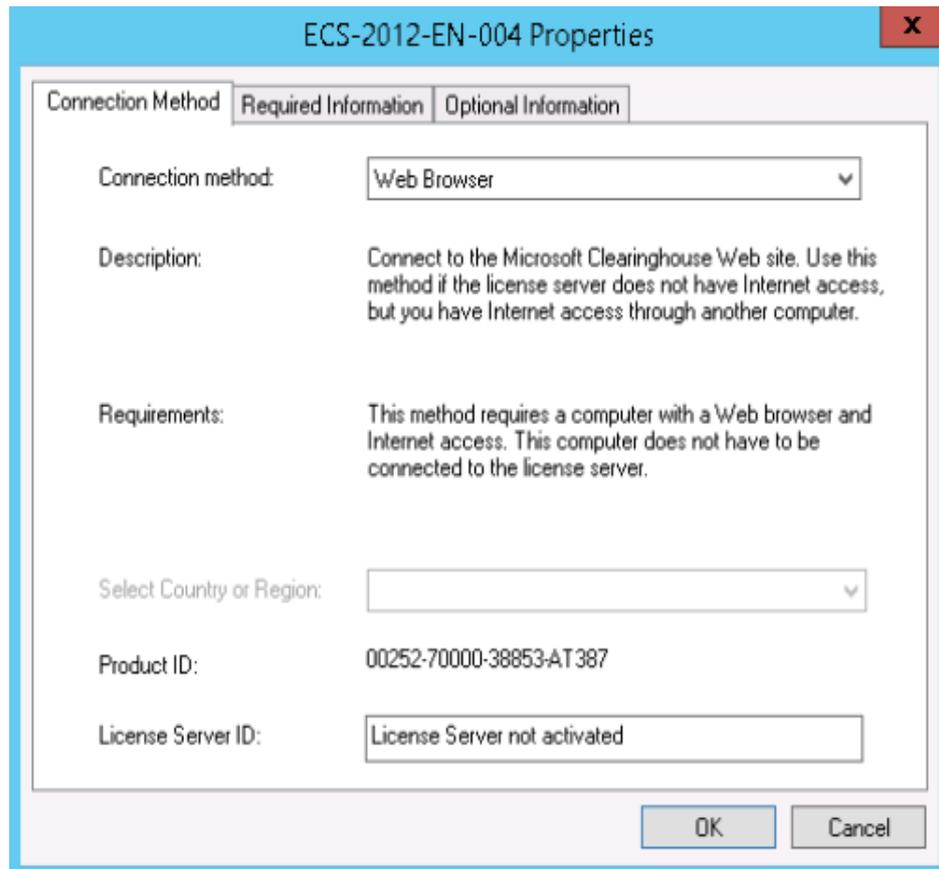


5. In the **Properties** dialog box, set **Connection Method** to **Web Browser**. Record the product ID which will be required for obtaining a server license.

NOTICE

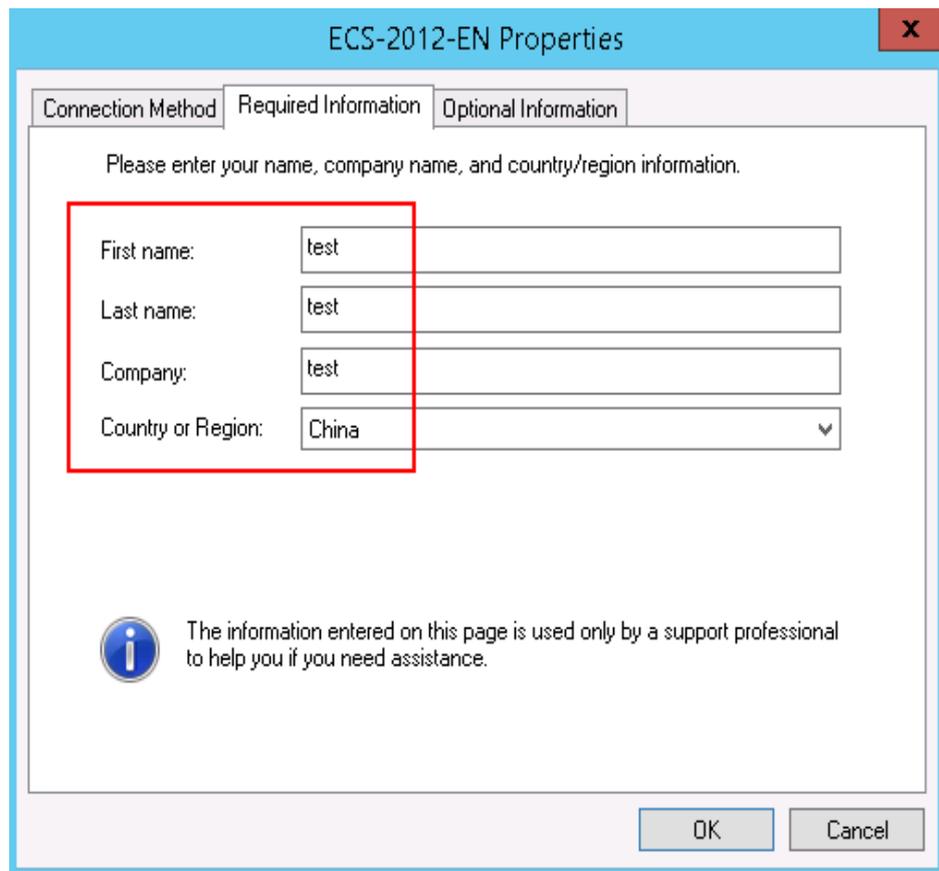
If the destination server is not displayed in **RD Licensing Manager**, choose **Action > Connect**, and enter the server IP address.

Figure 9-26 Web browser for connection



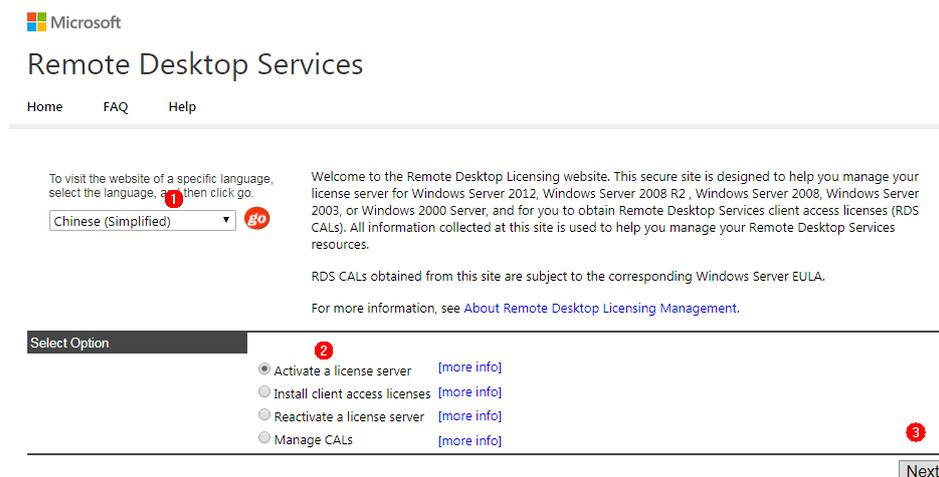
6. Set parameters and click **OK**.

Figure 9-27 Setting parameters



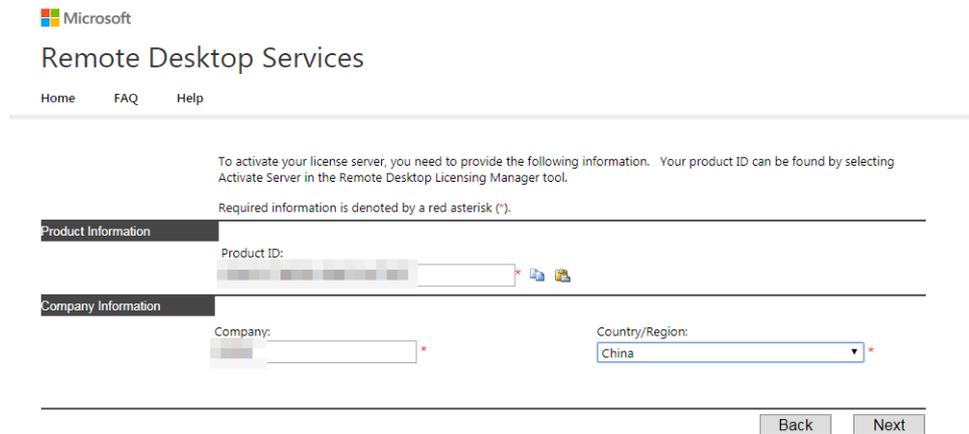
7. Access <https://activate.microsoft.com> using Internet Explorer. Select **Activate a license server** and click **Next**.

Figure 9-28 Activating the license server



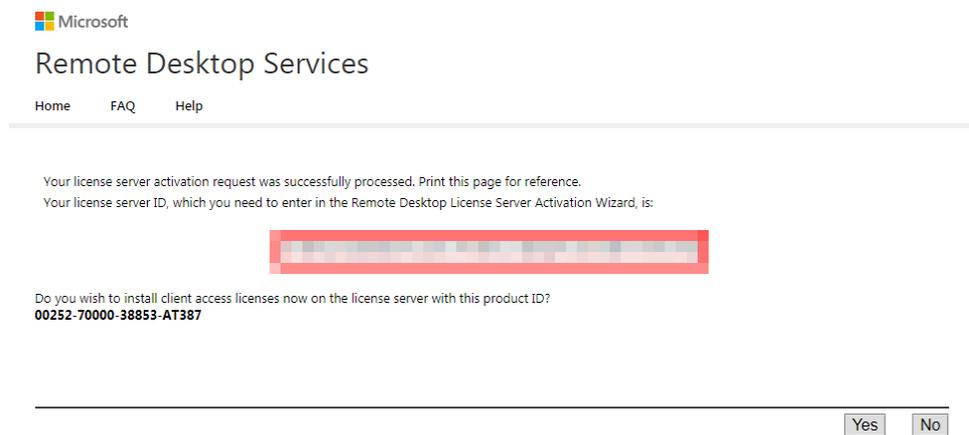
8. Set parameters and click **Next**. Confirm the product information and click **Next**.

Figure 9-29 Configuring the server



9. Record the license server ID and click **Yes**.

Figure 9-30 Obtaining the license server ID



10. If the license server ID is not available, set **License Program** to **Enterprise agreement** and click **Next**.

Figure 9-31 Enterprise agreement

Microsoft
Remote Desktop Services

Home FAQ Help

To install licenses, you need to provide the following information. Your license server ID can be found by selecting Install Licenses in the Remote Desktop Licensing Manager tool. Your connection method should be set to Web Browser (Windows Server 2008), or Web Browser (Windows Server 2003). To change your connection method, on the View menu of the Remote Desktop Licensing Manager tool, click Properties, and then click the Connection Method tab.

Required information is denoted by a red asterisk (*).

Product Information
License Server ID

Licensing Information
License Program:
Enterprise agreement

Company Information
Company:
Country/Region:
China

Back Next

11. Set parameters.

- **Product Type:** specifies the license server type, for example, **Windows Server 2012 Remote Desktop Services Per User client access license**.
- **Quantity:** specifies the maximum number of users allowed for remote desktop connections, for example, **999**.
- **Agreement Number:** specifies a 7-digit license registration number purchased at the official Microsoft website.

NOTE

Huawei Cloud does not provide licenses for remote desktop connections. Purchase such a license at the official Microsoft website.

Figure 9-32 Setting license server information

Microsoft
Remote Desktop Services

Home FAQ Help

To install client access licenses, you need to provide the following information.

Required information is denoted by a red asterisk(*)

License Server ID

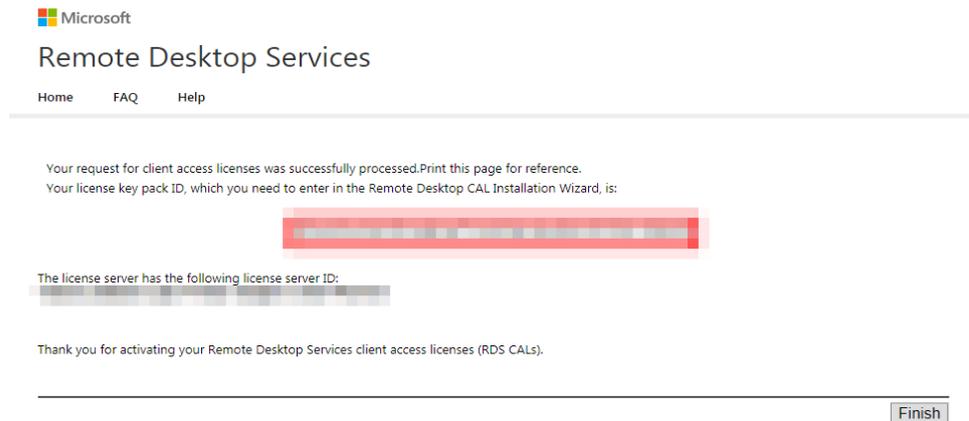
Product Information
Product Type:
Windows Server 2012 Remote Desktop Services Per User client access license
Quantity:
999

Licensing Information
License Program:
Enterprise agreement
Agreement Number:

Back Next

12. Record the license server ID and license key pack ID and click **Finish**.

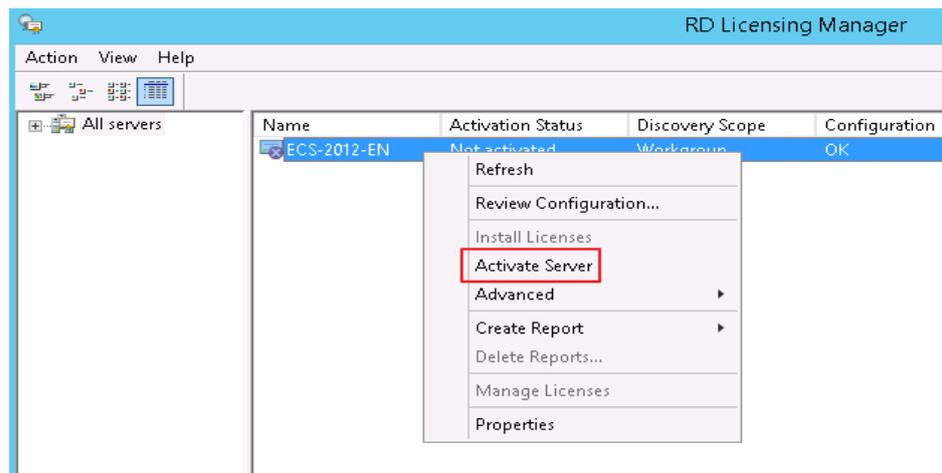
Figure 9-33 Obtaining a license key pack ID



Activating the ECS

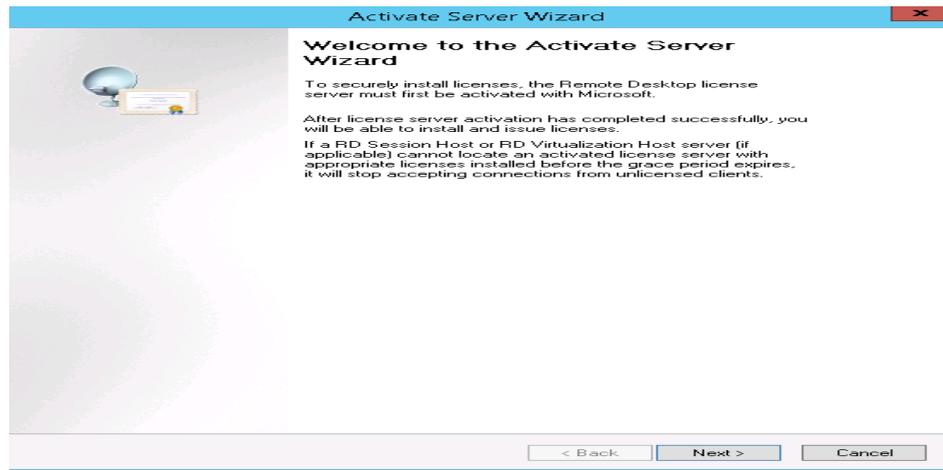
1. Log in to the ECS. Open **RD Licensing Manager**, right-click the ECS, and choose **Activate Server**.

Figure 9-34 Activating the ECS



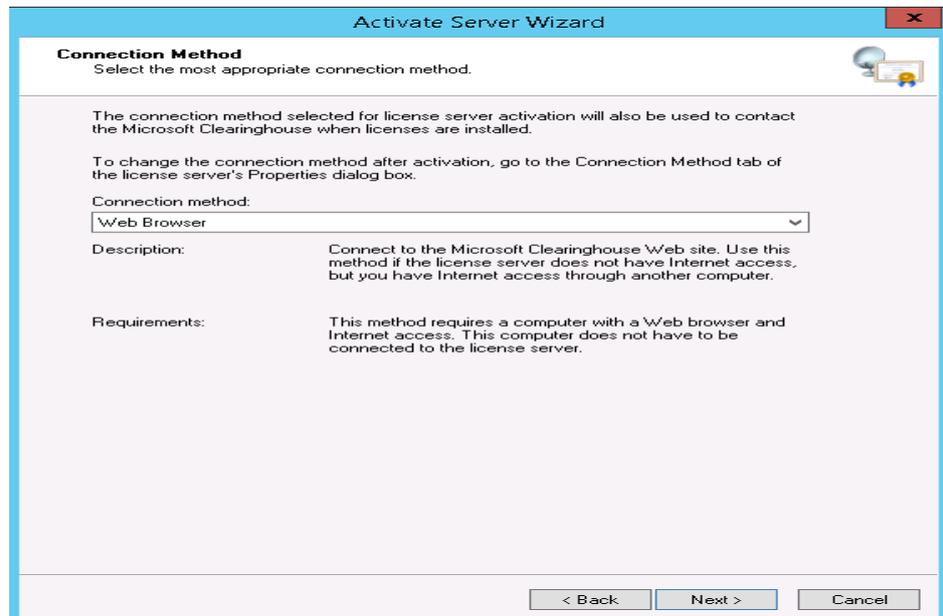
2. In **Activate Server Wizard**, click **Next**.

Figure 9-35 Activate Server Wizard



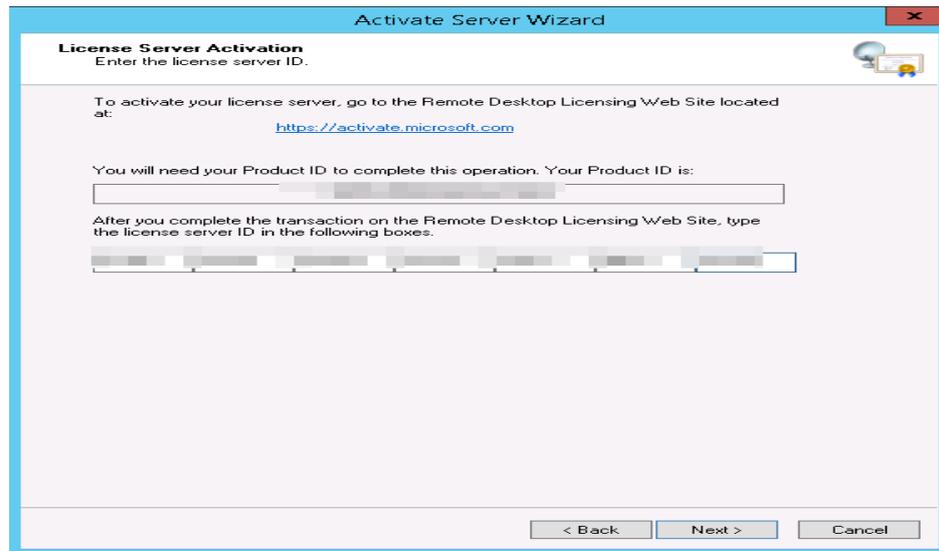
3. Set **Connection method** to **Web Browser** and click **Next**.

Figure 9-36 Web browser for connection



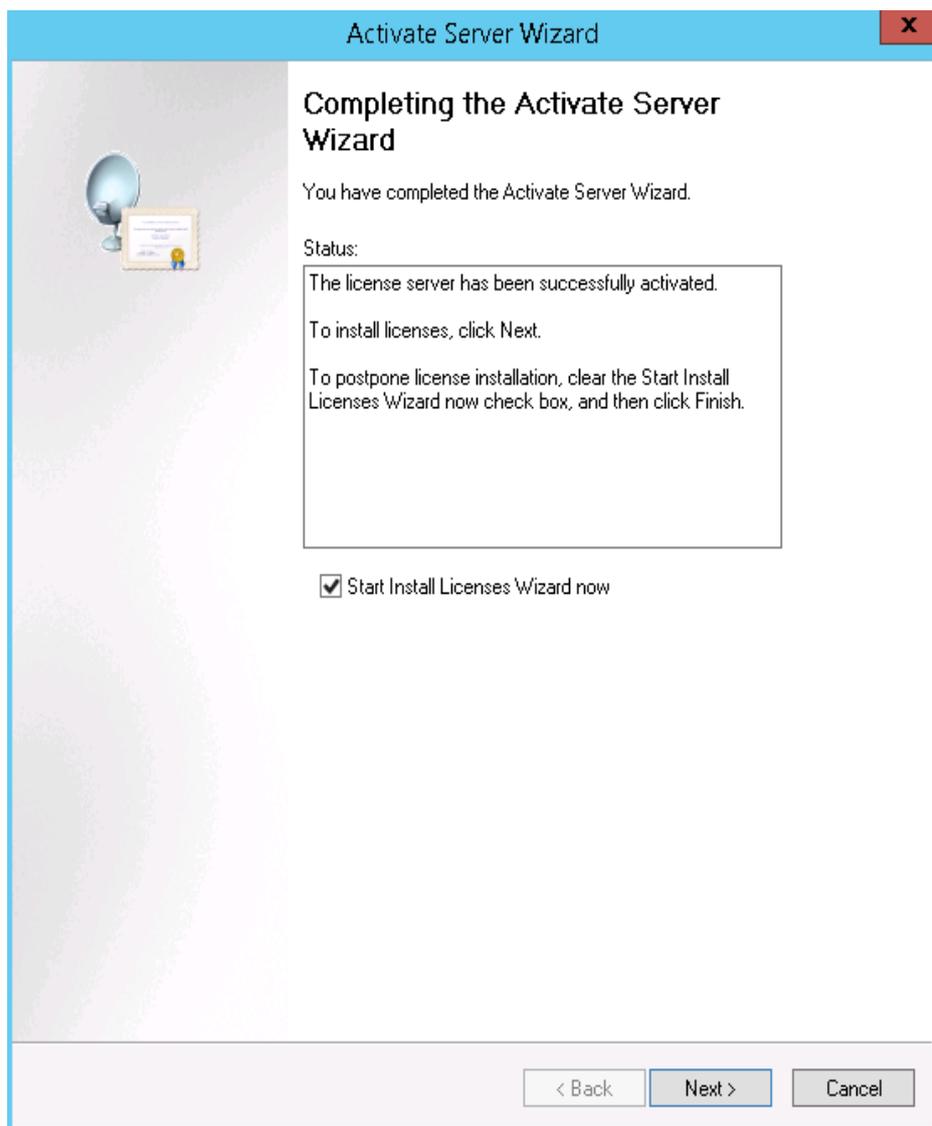
4. Enter the license server ID and click **Next**.
The license server ID is the ID obtained in step 9.

Figure 9-37 Entering the license server ID

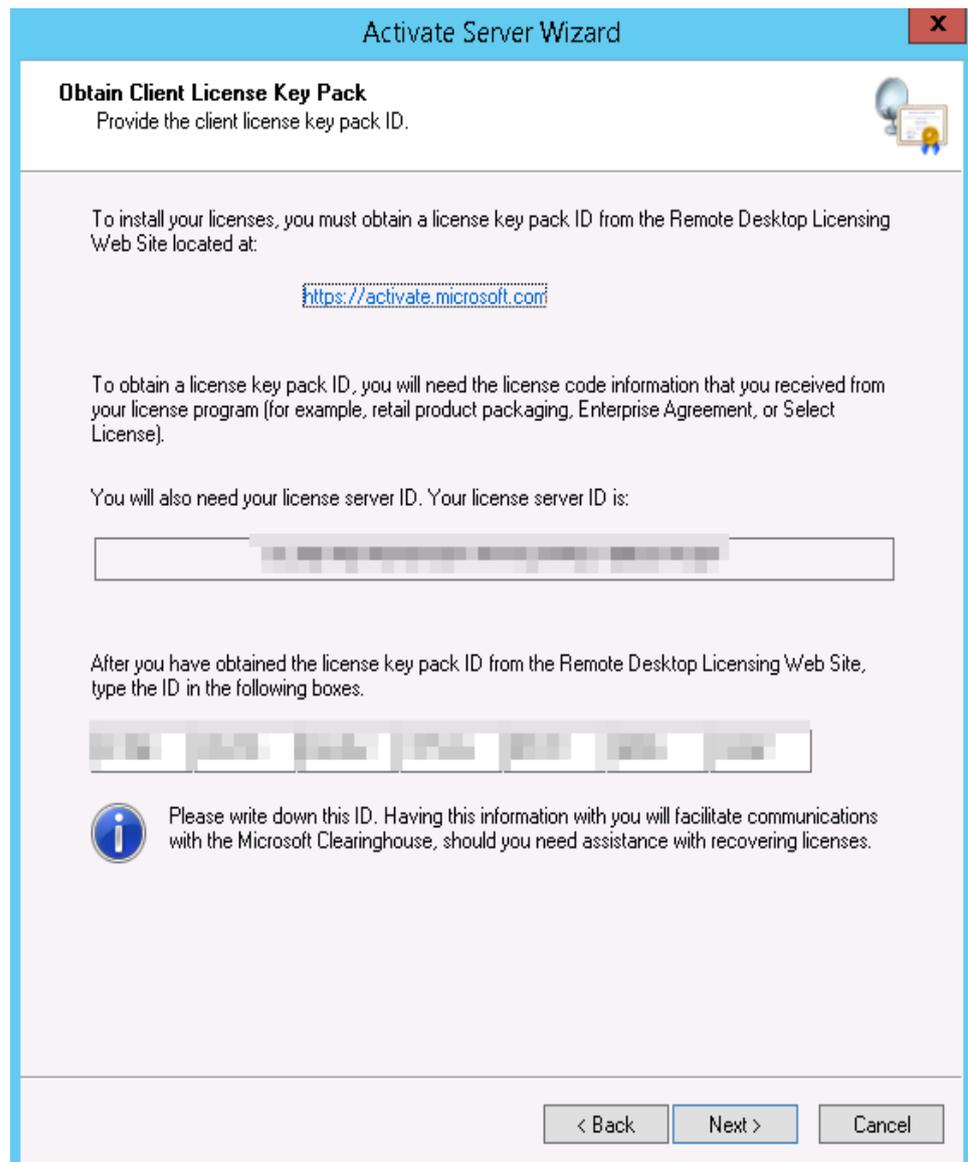


5. Select **Start Install Licenses Wizard now** and click **Next**.

Figure 9-38 Starting license installation wizard



6. Enter the license key pack ID and click **Next**.
The key pack ID is obtained in step [12](#).

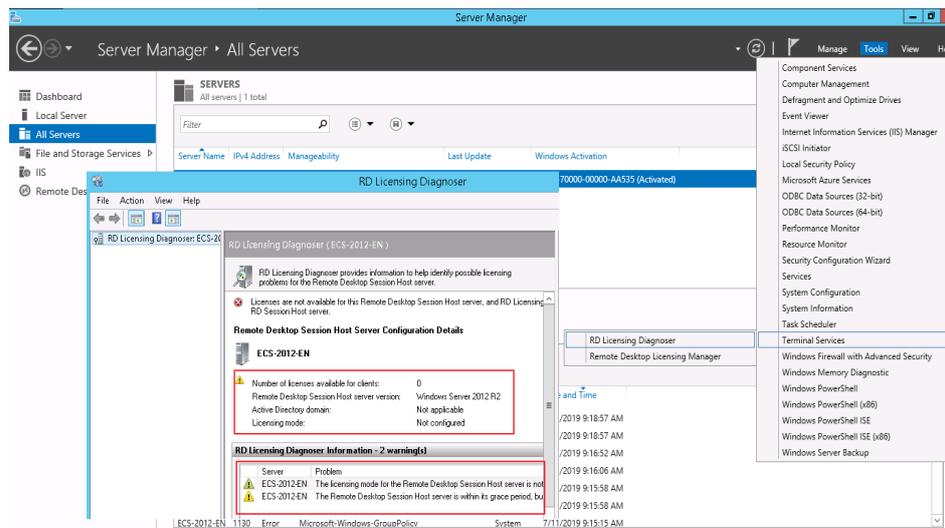
Figure 9-39 Entering the key pack ID

7. Click **Finish**.

Configuring the Licensing Server for Remote Desktop Session Host

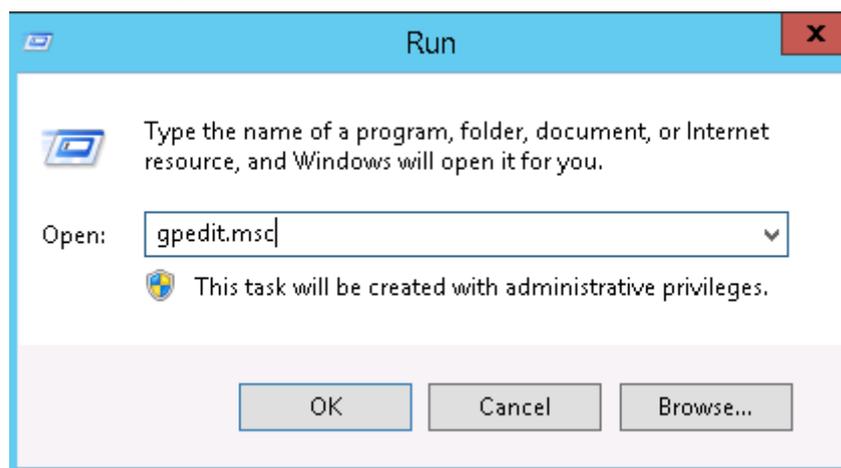
1. Log in to the Windows ECS.
2. On the OS, click  to open **Server Manager**. Choose **Tools > Terminal Services > RD Licensing Diagnoser** and check the ECS authorization status. For the Windows Server 2019 OS, choose **Tools > Remote Desktop Services > RD Licensing Diagnoser** and check the ECS authorization status. As shown in the following figure, the system displays a message indicating that the licensing mode for the Remote Desktop Session Host Server is not configured.

Figure 9-40 Checking the ECS authorization status



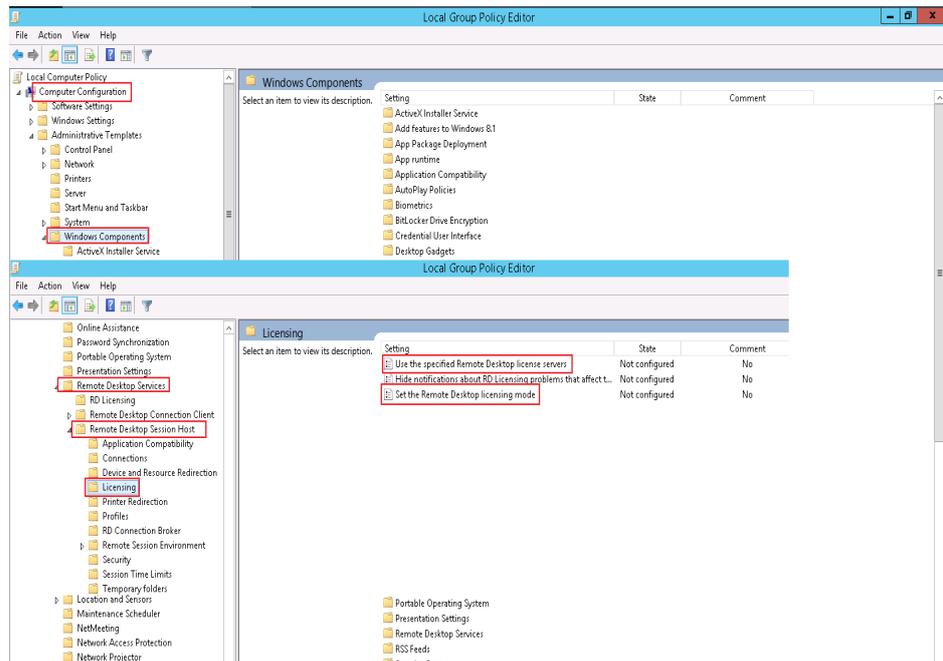
3. In the **Run** dialog box, enter **gpedit.msc** and click **OK** to start Local Group Policy Editor.

Figure 9-41 gpedit.msc



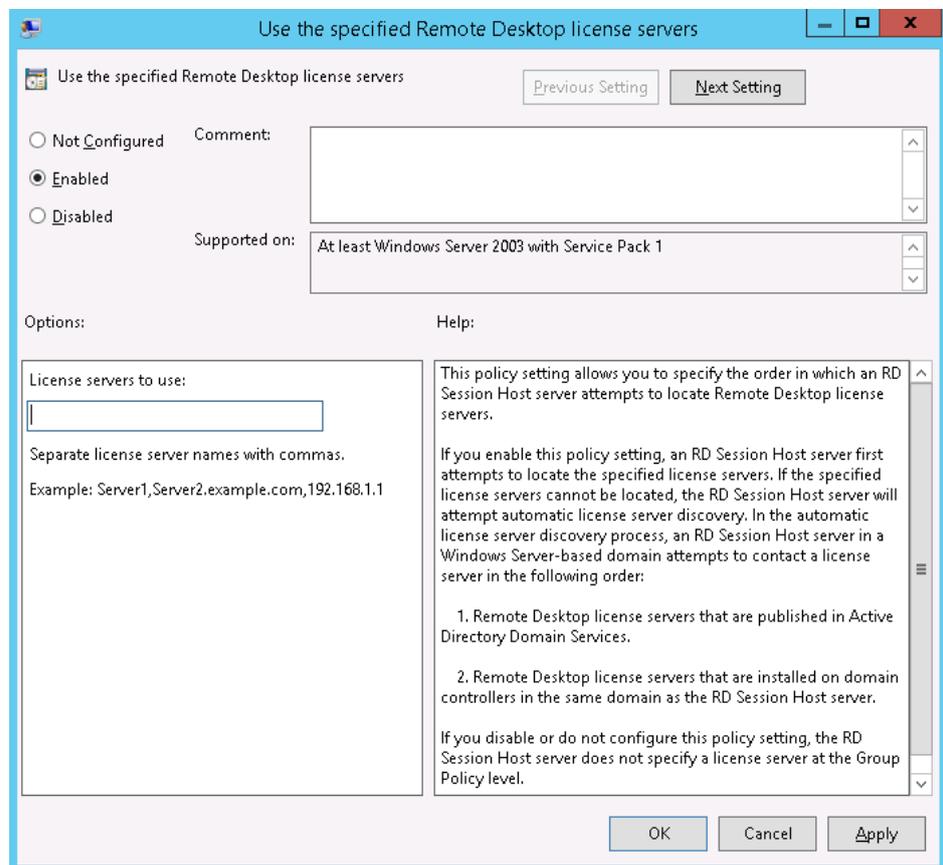
4. Choose **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Licensing**.

Figure 9-42 Modifying licensing items



5. Set **Use the specified Remote Desktop license servers** to **Enabled**, enter the private IP address or hostname of the ECS under **License servers to use**, and click **OK**.

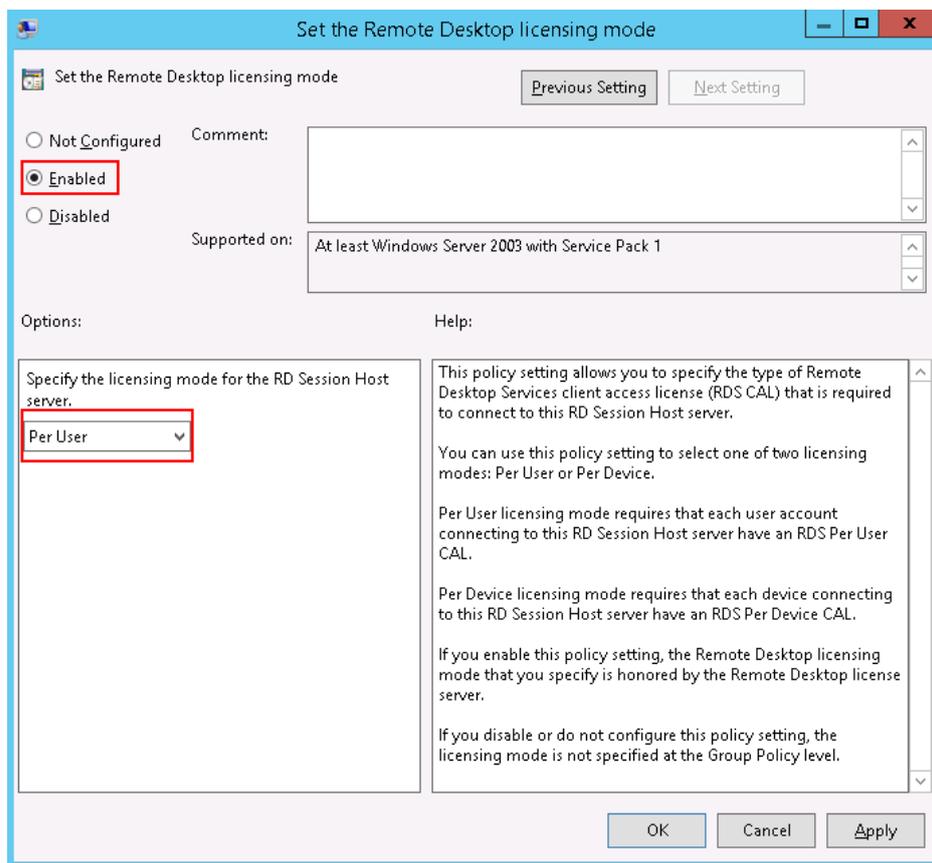
Figure 9-43 Using the specified remote desktop license servers



6. Select **Enabled** for **Set the Remote Desktop licensing mode**, and set the licensing mode to **Per User**.

NOTE

If a message is displayed indicating that there is a problem with remote desktop license, set the licensing mode to **Per Device**.

Figure 9-44 Setting the remote desktop licensing mode

7. Run **cmd** and then **gpupdate /force** to forcibly start Local Group Policy Editor and restart the ECS.

9.4 How Do I Troubleshoot Login Screen Flickering After Configuring Multi-User Login?

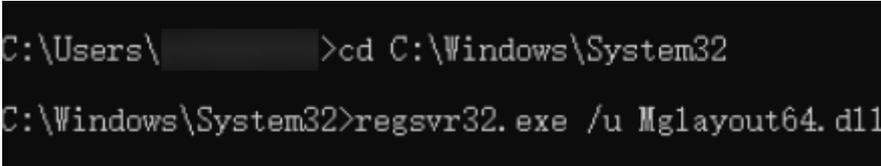
Symptom

After multi-user login is configured on a Windows ECS, the administrator can log in normally but common users encounter screen flickering or auto closing of "My Computer".

Solution

1. Log in to the ECS as the administrator, view system logs and application logs, and search for the abnormal module. The following uses the Mglayout64.dll module as an example.
2. Open **C:\Windows** and search for the corresponding module file. In this example, the module is a wallpaper file.
3. Open the **Run** command box, type **cmd** and then press **Enter**.
Run the following example command:
cd C:\Windows\System32
4. Run the **regsvr32.exe /u file-name** command to remove the file.
In this example, the **Mglayout64.dll** file is used as an example.
regsvr32.exe /u Mglayout64.dll

Figure 9-45 Removing the abnormal file



```
C:\Users\>cd C:\Windows\System32
C:\Windows\System32>regsvr32.exe /u Mglayout64.dll
```

10 Password and Key Pair Issues

10.1 What Should I Do If the passwd Failed to Be Executed, with Error "Authentication token manipulation error" Displayed?

Symptom

When user **root** tries to change the password of an administrator or a common user, the system displays "passwd:Authentication token manipulation error".

```
[root@xiaoyao-test ~]# passwd
Changing password for user root.
New password:
Retype new password:
passwd: Authentication token manipulation error
```

Possible Causes

This problem is usually caused by wrong attributes of password files or a full root directory.

Run the following command to check the attributes of the files **/etc/passwd** and **/etc/shadow** that store user names and passwords:

```
lsattr /etc/passwd /etc/shadow
```

```
[root@xiaoyao-test ~]# lsattr /etc/passwd /etc/shadow
----i-----e-- /etc/passwd
----i-----e-- /etc/shadow
```

As shown in the preceding figure, the **/etc/passwd** and **/etc/shadow** files have the **i** attribute. A file with the **i** attribute is immutable. It cannot be deleted or renamed, no link can be created to this file, and no data can be written to this file. Only the administrator can set or clear this attribute.

 NOTE

1. Files with the **a** attribute can only be open in append mode for writing. Only the administrator can set or clear this attribute.

The **CAP_LINUX_IMMUTABLE** capability can be used to set or clear this attribute.

For information about other file attributes, run the following command to view the **chattr** user manual:

chattr

2. If the **lsattr** command output does not contain any attributes that restrict file modification, the problem may be caused by insufficient root partition space. In this case, run the following command to check the root partition usage:

df -h

Delete unnecessary files from the root partition.

Solution

1. Use **chattr** to revoke the **i** or **a** attribute and then change the password.
 - For files with the **i** attribute, run the following command:
chattr -i /etc/passwd /etc/shadow
 - For files with the **a** attribute, run the following command:
chattr -a /etc/passwd /etc/shadow
 2. (Optional) Change the file attributes back to **i** or **a** to meet the security requirements.
 - To set the **i** attribute for the files, run the following command:
chattr +i /etc/passwd /etc/shadow
 - To set the **a** attribute for the files, run the following command:
chattr +a /etc/passwd /etc/shadow
- Run the following command to check the file attributes again:
- lsattr /etc/passwd /etc/shadow**

10.2 How Do I Change the Key Pair for a Linux ECS?

Symptom

You changed the key pair for logging in to a Linux ECS, but you could not use the new key pair to log in to the ECS.

Solution

1. Use the password or old key to log in to the ECS and run the following commands to create a key pair:

```
[root@host ~]$ ssh-keygen <==Create a key pair.
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa): <== Press Enter.
Created directory '/root/.ssh'.
Enter passphrase (empty for no passphrase): <== Enter the key passphrase or press Enter to leave it blank.
Enter same passphrase again: <== Enter the key passphrase again.
Your identification has been saved in /root/.ssh/id_rsa. <== Private key
Your public key has been saved in /root/.ssh/id_rsa.pub. <== Public key
```

```
The key fingerprint is:  
0f:d3:e7:1a:1c:bd:5c:03:f1:19:f1:22:df:9b:cc:08 root@host
```

2. Set a key passphrase when a private key is used. This prevents the private key from being stolen.

A hidden directory **.ssh** is generated in the home directory of user **root** user. The directory contains two key files: **id_rsa** (private key) and **id_rsa.pub** (public key).

3. Run the following commands on the Linux ECS to install the public key:

```
[root@host ~]$ cd .ssh  
[root@host .ssh]$ cat id_rsa.pub >> authorized_keys
```

4. To ensure successful connection, run the following commands to configure file permissions:

```
[root@host .ssh]$ chmod 600 authorized_keys  
[root@host .ssh]$ chmod 700 ~/.ssh
```

5. (Optional) Edit the **/etc/ssh/sshd_config** file as follows to enable SSH key pair login. (Skip this step for Linux ECSs created using a public image because key pair login has been enabled for them by default.)

```
RSAAuthentication yes  
PubkeyAuthentication yes
```

6. Check the following configuration item to ensure that user **root** can log in to the ECS using SSH:

```
PermitRootLogin yes
```

7. After the preceding configuration is complete and the login using a key is successful, run the following command to disable password login:

```
PasswordAuthentication no
```

8. Restart SSH.

```
service sshd restart
```

Follow-up Operations

1. Download the private key to your local computer and convert it into a format that can be used by PuTTY.
Use WinSCP or SFTP to download the private key file **id_rsa**. Keep the file secure.
2. Open PuTTY Key Generator and click **Load an existing private key file** in **Actions** to load the downloaded private key file.
If you have set a key passphrase, you need to enter it.
3. After the private key file is loaded successfully, PuTTY Key Generator displays the information about the key.
4. In the **Key comment** box, enter the description of the private key and click **Save private key** to save the private key file in a format that can be used by PuTTY.
5. Use PuTTY to log in to the ECS. In the navigation pane of PuTTY, choose **Connection > SSH > Auth**, and in the **Private key file for authentication** area, browse to and select your private key file, enter the key passphrase, and click **Open** to log in to the ECS.

10.3 How Do I Change the Login Mode of a Linux ECS from Key Pair to Password?

Scenarios

This section describes how to change the login mode of a Linux ECS from key pair to password.

Procedure

1. Use the key pair to log in to the Linux ECS and set the password of user **root**.
sudo passwd root
2. Modify the SSH configuration file on the ECS as user **root**.
su root
vi /etc/ssh/sshd_config
Modify the following settings:
 - Change **PasswordAuthentication no** to **PasswordAuthentication yes**.
Alternatively, delete the comment tag (#) before **PasswordAuthentication yes**.
 - Change **PermitRootLogin no** to **PermitRootLogin yes**.
Alternatively, delete the comment tag (#) before **PermitRootLogin yes**.
3. Restart sshd for the modification to take effect.
service sshd restart
4. Restart the ECS. Then, you can log in to the ECS as user **root** using the password.

NOTE

To prevent unauthorized users from using the key file to access the Linux ECS, delete the **/root/.ssh/authorized_keys** file or clear the **authorized_keys** file.

11 Firewall Configuration Issues

11.1 How Do I Enable or Disable a Windows ECS Firewall and Add Port Exceptions to the Firewall?

Scenarios

This section describes how to enable and disable a Windows ECS firewall and add a port exception to the firewall.

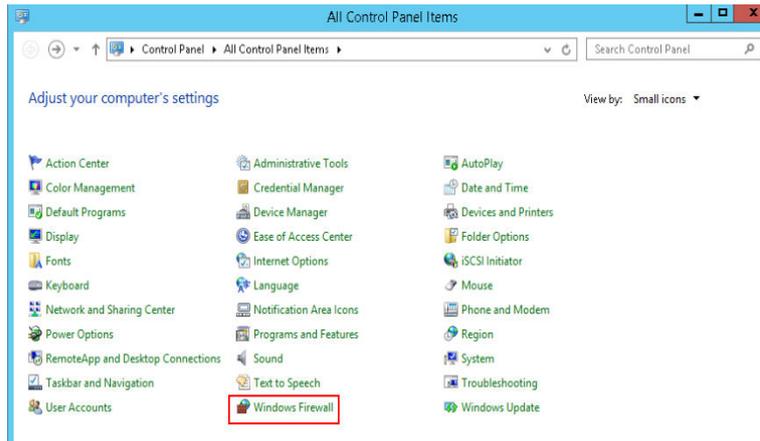
The following operations are performed on an ECS running Windows Server 2012.

 **CAUTION**

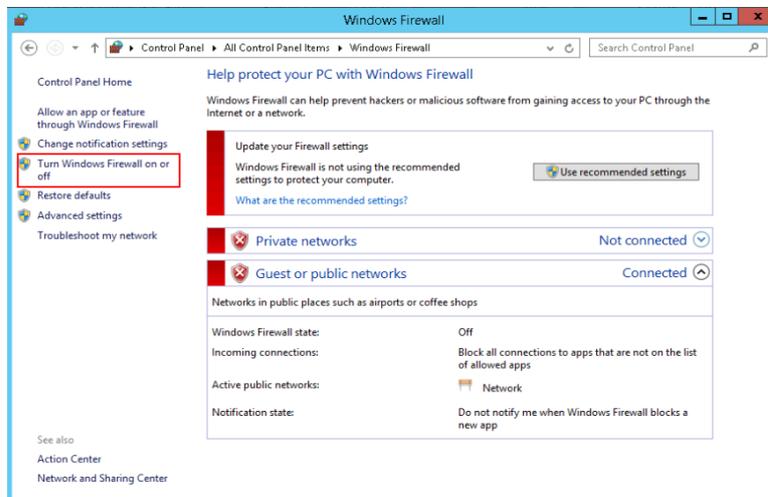
Enabling a firewall and configuring security groups can protect your ECSs. If you disable a firewall or add port exceptions, exercise caution when you open ports in security groups.

Enabling or Disabling a Firewall

1. Log in to the Windows ECS.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.

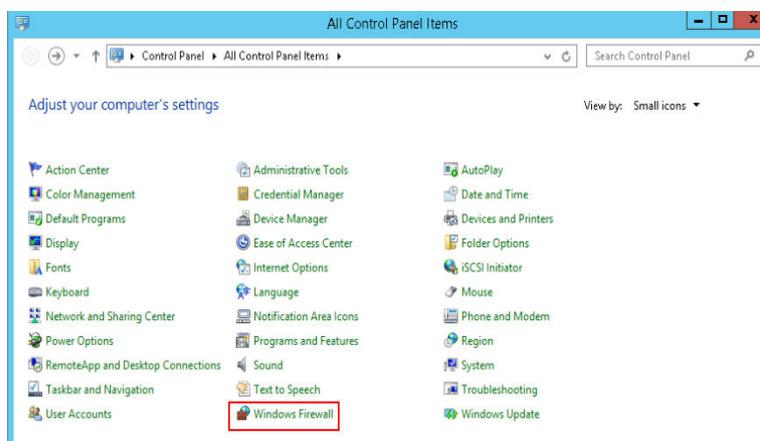


3. Click **Turn Windows Firewall on or off**.
View and set the firewall status.



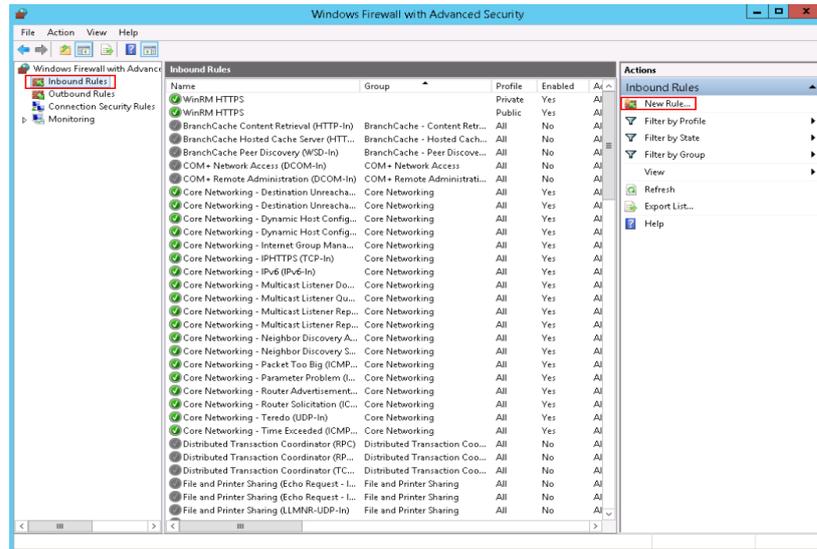
Adding a Port Exception to a Firewall

1. Log in to the Windows ECS.
2. Click the Windows icon in the lower left corner of the desktop and choose **Control Panel > Windows Firewall**.



3. In the navigation pane on the **Windows Firewall** page, choose **Advanced settings > Inbound Rules**.

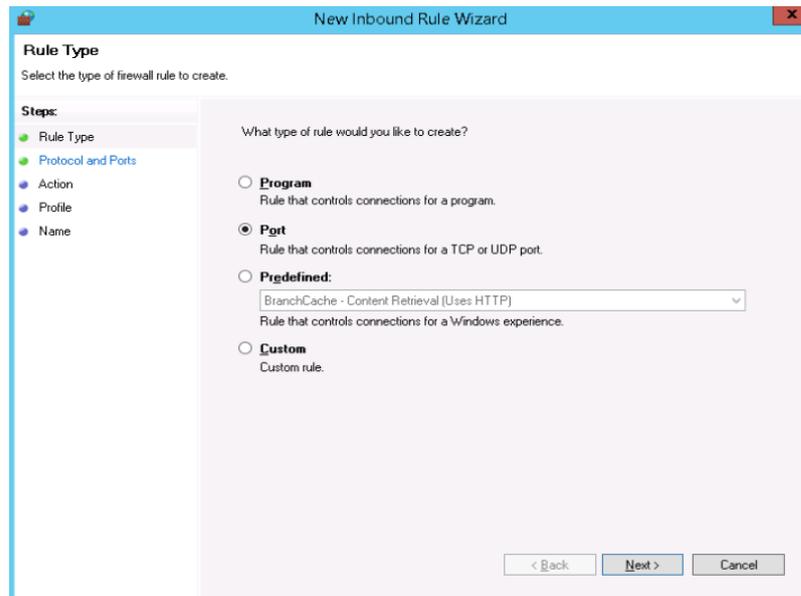
You can view details about the programs and ports that can be connected to the ECS.



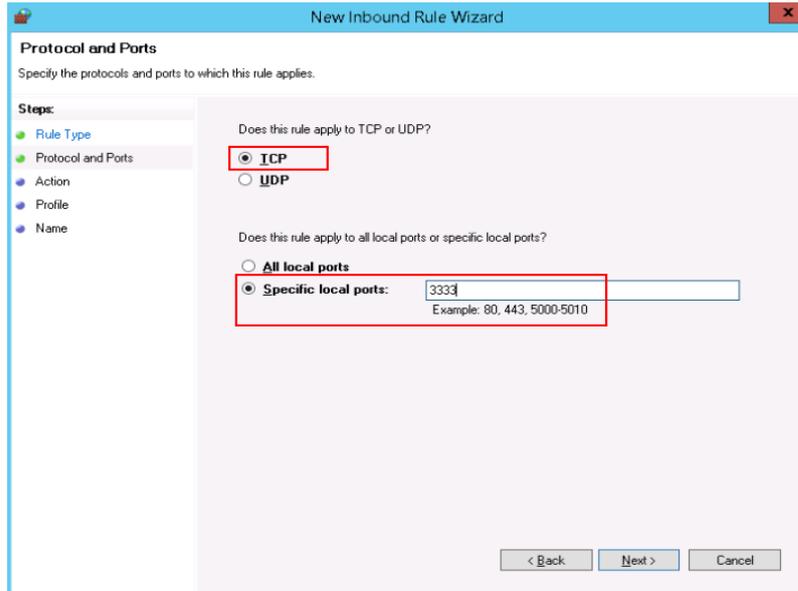
4. Click **New Rule** in the **Actions** column.

Enable ports as prompted. This section uses port 3333 as an example. To open port 3333, perform the following operations:

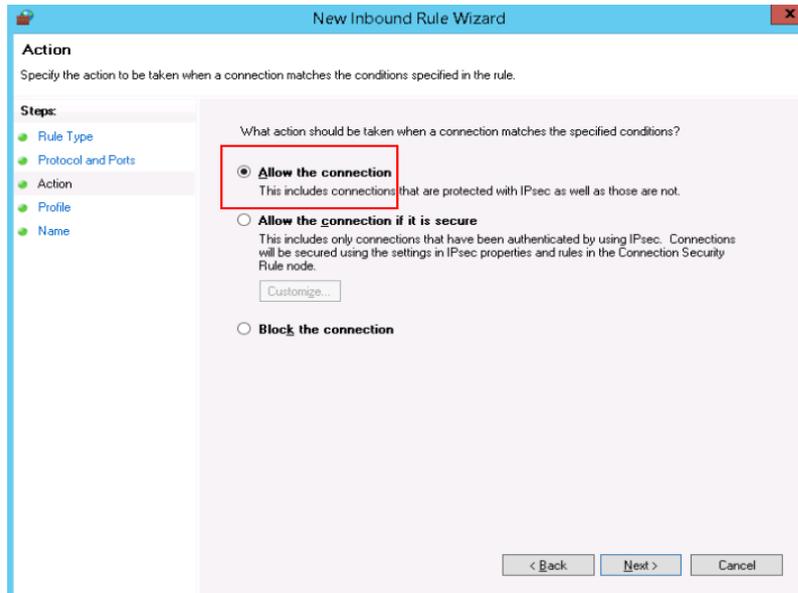
Select **Port**.



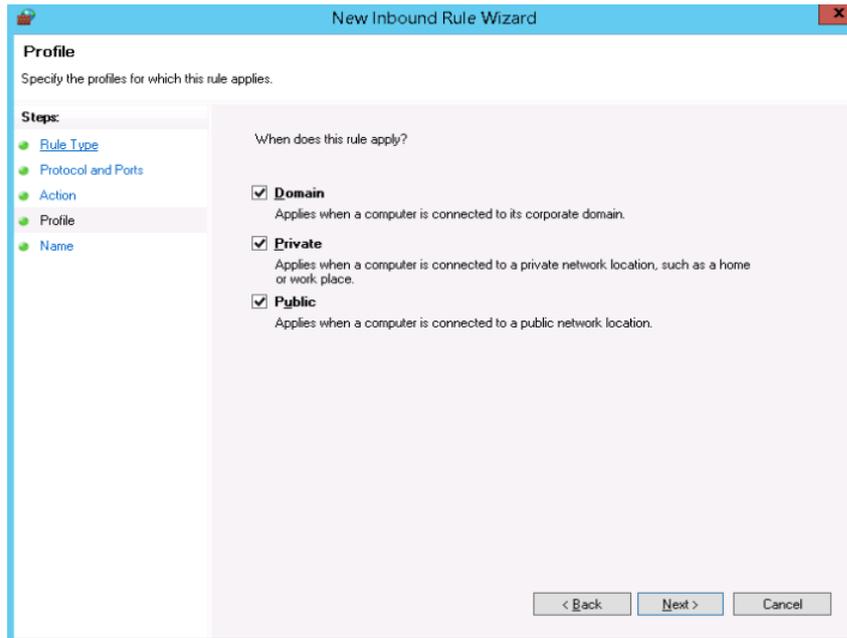
Select a protocol type and set the port.



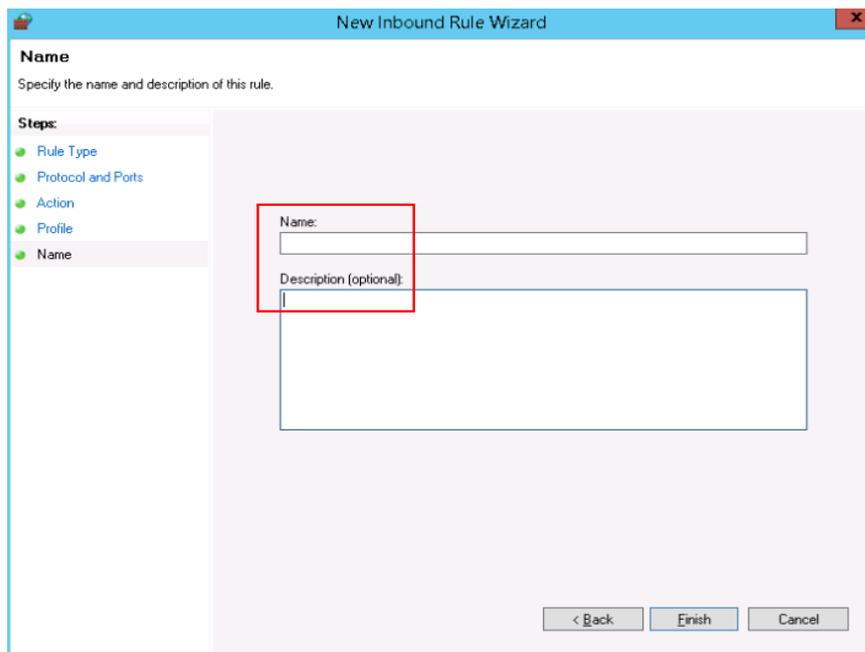
Set a connection rule.



Set the application scenarios of the rule.



Name the rule.



5. Verify that the traffic to the port is permitted on the firewall.
Click the newly added rule to view its details and set other information, such as computer connection, allowed IP connection, and protocol type.

11.2 How Do I Enable or Disable a Linux ECS Firewall and Add a Port Exception on a Linux ECS Firewall?

Scenarios

This section describes how to enable or disable a Linux ECS firewall and add a port exception on a Linux ECS firewall.

CAUTION

Enabling a firewall and configuring security groups can protect your ECSs. If you disable a firewall or add port exceptions, exercise caution when you open ports in security groups.

Enabling or Disabling a Firewall

Run the corresponding command to enable or disable the firewall based on the ECS OS.

OS	Enable	Disable
CentOS 6	<code>service iptables start</code>	<code>service iptables stop</code>
CentOS 7	<code>systemctl start firewalld.service</code>	<code>systemctl stop firewalld.service</code>
Ubuntu	<code>ufw enable</code>	<code>ufw disable</code>
Debian	<code>/etc/init.d/iptables start</code>	<code>/etc/init.d/iptables stop</code>

Adding a Port Exception to a Firewall

- CentOS 6
 - a. For example, to add TCP port 23, run the following command:
`iptables -A INPUT -m state --state NEW -m tcp -p tcp --dport 23 -j ACCEPT`
 - b. Save the configuration.
`service iptables save`
 - c. (Optional) Configure the firewall to automatically start upon ECS startup.
`chkconfig iptables on`

 NOTE

- To disable automatic firewall startup, run the following command:
chkconfig iptables off
- When the firewall is started on CentOS 6, the "iptables no config file" error may be displayed because the **iptables** configuration file is not detected. To handle this issue, perform the following operations:
 1. Add a new rule.
iptables -P OUTPUT ACCEPT
 2. Save the configuration.
service iptables save
 3. Start the firewall again.
service iptables start
- CentOS 7
 - a. Check the firewall status.
systemctl status firewalld
or
firewall-cmd --state
 - b. If the firewall is disabled, run the following command to enable it:
systemctl start firewalld
If "Failed to start firewalld.service: Unit is masked." is displayed, run the **systemctl unmask firewalld** command first and then run the preceding command again to enable the firewall.
 - c. Run the following command to check whether the firewall is enabled:
firewall-cmd --state
Information similar to the following is displayed:

```
[root@ecs-centos7 ~]# firewall-cmd --state  
running
```
 - d. For example, to add TCP port 23, run the following command:
firewall-cmd --zone=public --add-port=23/tcp --permanent
The configuration is correct if the command output is as follows:

```
[root@ecs-centos7 ~]# firewall-cmd --zone=public --add-port=23/tcp --permanent  
success
```
 - e. Reload the policy configuration for the new configuration to take effect.
firewall-cmd --reload
 - f. Check all enabled ports.
firewall-cmd --list-ports

```
[root@ecs-centos7 ~]# firewall-cmd --list-ports  
23/tcp
```
 - g. (Optional) Configure the firewall to automatically start upon ECS startup.
systemctl enable firewalld.service
Check whether automatic firewall startup is enabled.
systemctl is-enabled firewalld.service;echo \$?
The configuration is correct if the command output is as follows:

```
[root@ecs-centos7 ~]# systemctl is-enabled firewalld.service;echo $?  
enabled  
0
```

NOTE

To disable automatic firewall startup, run the following command:
systemctl disable firewalld.service

11.3 Why Does My Linux ECS Fail to Access the Internet After Port 80 Is Allowed by the Firewall Rules?

Symptom

The Linux ECS cannot access the Internet after port 80 is allowed by the firewall rules. After the firewall is disabled, access to the Internet succeeds.

Possible Causes

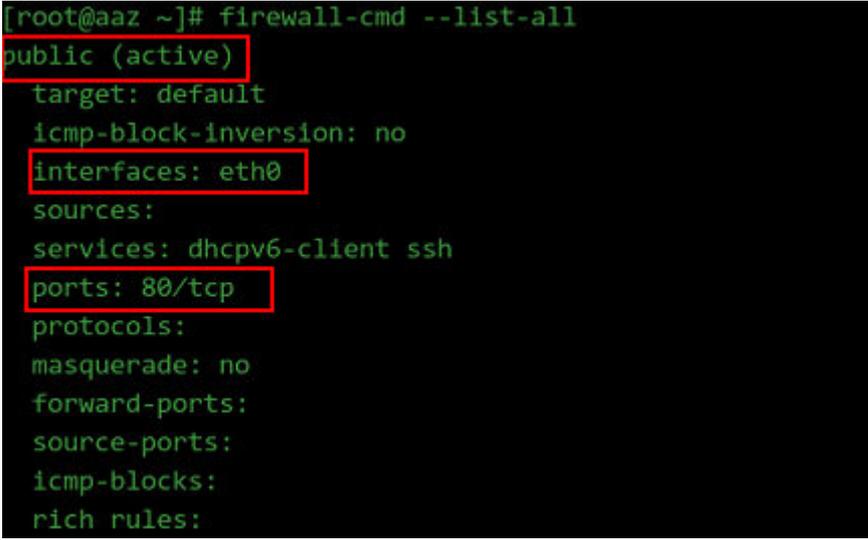
The possible cause is that the firewall rules and the NIC interfaces are in different zones. You can perform the following steps to check the causes.

1. Run the following command to check the specified zone by and the open port of the firewall:

```
firewall-cmd --list-all
```

As shown in the following figure, the firewall zone is public, the open port is 80, and the NIC interface is eth0.

Figure 11-1 Viewing firewall information



```
[root@aaz ~]# firewall-cmd --list-all  
public (active)  
target: default  
icmp-block-inversion: no  
interfaces: eth0  
sources:  
services: dhcpv6-client ssh  
ports: 80/tcp  
protocols:  
masquerade: no  
forward-ports:  
source-ports:  
icmp-blocks:  
rich rules:
```

2. Run the following command to check the NIC interface zone:

```
firewall-cmd --get-active-zones
```

The following figure shows the NIC interface zone is external.

Figure 11-2 Checking the NIC interface zone

```
[root@aaz ~]# firewall-cmd --get-active-zones
external
  interfaces: eth0
[root@aaz ~]#
```

3. Run the following command to check whether port 80 is enabled in the external zone. **Figure 11-3** shows an example.

```
firewall-cmd --zone=external --list-ports
```

Figure 11-3 The port not enabled in the external zone

```
[root@aaz ~]# firewall-cmd --zone=external --list-ports
[root@aaz ~]#
```

Port 80 is not enabled in the external zone.

4. Run the following command to check whether port 80 is enabled in the public zone. **Figure 11-4** shows an example.

```
firewall-cmd --zone=public --list-ports
```

Figure 11-4 The port enabled in the public zone

```
[root@aaz ~]# firewall-cmd --zone=public --list-ports
80/tcp
[root@aaz ~]#
```

Port 80 is enabled in the public zone.

Therefore, the Linux ECS cannot access the Internet because the firewall rules and the NIC interfaces are in different zones.

Solution

Method 1

Perform the following steps to add firewall rules to enable port 80 in a specified zone (external) of the NIC interface.

1. Run the following command to enable port 80 in the external zone:

```
firewall-cmd --zone=external --add-port=80/tcp --permanent
```

Figure 11-5 Enabling port 80 in the external zone

```
[root@aaz ~]# firewall-cmd --zone=external --add-port=80/tcp --permanent
success
```

2. Run the following command to update firewall rules:

```
firewall-cmd --reload
```

Figure 11-6 Updating firewall rules

```
[root@aaz ~]# firewall-cmd --reload
success
```

3. Run the following command to check firewall rules:
firewall-cmd --zone=external --list-ports

Figure 11-7 Checking firewall rules

```
[root@aaz ~]# firewall-cmd --zone=external --list-ports
80/tcp
```

Method 2

Perform the following steps to change the specified zone of the NIC interface from **external** to **public**:

1. Run the following command to change the specified zone of the NIC interface:
firewall-cmd --zone=public --change-interface=eth0

Figure 11-8 Changing the specified zone of the NIC interface

```
[root@aaz ~]# firewall-cmd --zone=public --change-interface=eth0
success
[root@aaz ~]#
```

2. Run the following command to check the specified zone of the NIC interface:
firewall-cmd --get-active-zones

Figure 11-9 Checking the specified zone of the NIC interface

```
[root@aaz ~]# firewall-cmd --get-active-zones
public
  interfaces: eth0
[root@aaz ~]#
```

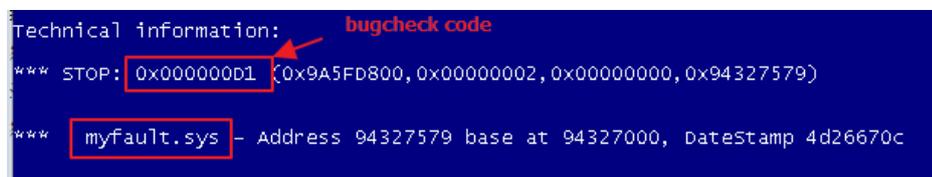
12 BSOD Issues

12.1 How Do I Fix a BSOD on a Windows ECS?

Symptom

A blue screen of death (BSOD) occurs on a Windows ECS, as shown in [Figure 12-1](#).

Figure 12-1 Bugcheck code and module that may cause the BSOD



```
Technical information:      bugcheck code
*** STOP: 0x000000d1 (0x9A5FD800, 0x00000002, 0x00000000, 0x94327579)
*** myfault.sys - Address 94327579 base at 94327000, DateStamp 4d26670c
```

Possible Causes

1. Third-party software from unknown sources is used.
2. The CPU usage is too high.
3. System files and registries are damaged due to misoperations or viruses.

When a BSOD occurs, the system displays a bugcheck code and the module that may cause this error.

Visit [Bug Check Code Reference](#) to learn about solutions provided by Microsoft.

Solution

1. Do not install software from unknown sources and use authorized software. Windows 2012 is recommended.
2. If the ECS is created using an external image, see "Optimizing a Windows Private Image" in *Image Management Service User Guide*.
3. If the BSOD is caused by high CPU usage, perform the following operations to reduce the CPU usage:

- Stop the processes that are not being used and try again.
 - Restart the ECS.
 - Back up important data and reinstall the OS.
 - If the OS cannot be reinstalled due to important data, replace the disk attached to the ECS. To do so, back up the data on the original disk, detach the disk from the ECS, attach the new disk to the ECS, and copy the data to the new disk.
4. If you need to analyze the BSOD cause, check whether crash dump files are generated in the specified directory.
- Analyzing BSOD logs is time-consuming. You are advised to restart the ECS and identify the fault based on the preceding possible causes. The common causes are third-party antivirus software, system faults, and high CPU usage.

12.2 How Do I Troubleshoot Blue Screen or Black Screen Errors After an ECS Is Started?

Symptom

When you are remotely logging in to an ECS, a screen that's black or blue is displayed.

Root Cause

The blue screen or black screen errors are caused by explorer.exe errors in Windows. explorer.exe is a Windows program manager or file resource manager. It is used to manage the Windows graphical shell, including desktop and file management. If this program is deleted, the Windows GUI becomes unavailable.

Solution

1. Open the task manager of the ECS.
On the VNC login page, click  to open the task manager.
2. On the **Processes** tab, check whether the explorer.exe or Windows resource manager process exists.
 - If yes, end the process.
 - If no, start the process.

Figure 12-2 explorer.exe

DitaDesktopWindowsService.e...	00	15,096 K	D:
DitaEditorWindowsService.exe	00	15,600 K	D:
dwm.exe	00	2,904 K	非
elmservice.exe *32	00	1,340 K	el
eSpace.exe *32	04	196,200 K	el
EXCEL.EXE	01	227,724 K	M:
explorer.exe	00	87,500 K	W
FaultReport.exe *32	00	3,084 K	f:
firefox.exe	00	91,924 K	F:
firefox.exe	00	319,108 K	F:
firefox.exe	00	37,032 K	F:
firefox.exe	00	65,896 K	F:
firefox.exe	00	64,200 K	F:

3. Start explorer.exe.

Choose **File > Run new task**. In the displayed dialog box, enter **explorer.exe** and click **OK**.

 **NOTE**

If the dialog box is not displayed, restart the ECS.

Figure 12-3 Running a new task

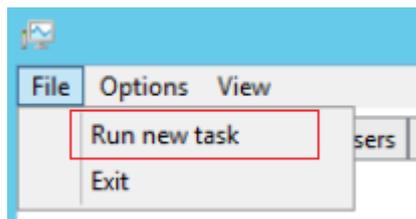
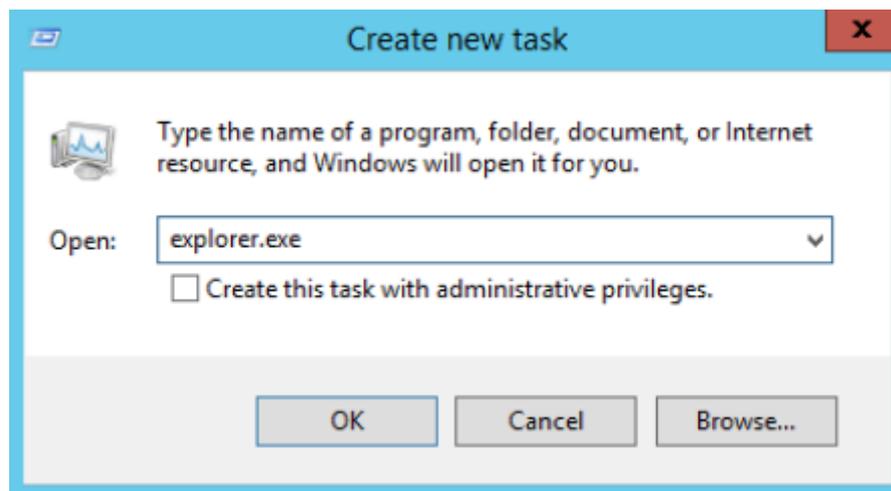


Figure 12-4 Running explorer.exe.



12.3 Why Does BSOD Occur When I Log In to an ECS Using Remote Desktop Connection?

Symptom

When you attempt to use a remote desktop connection to log in to an ECS running Windows Server 2012 R2 from a local computer with redirected drive enabled, BSOD occurs.

Root Cause

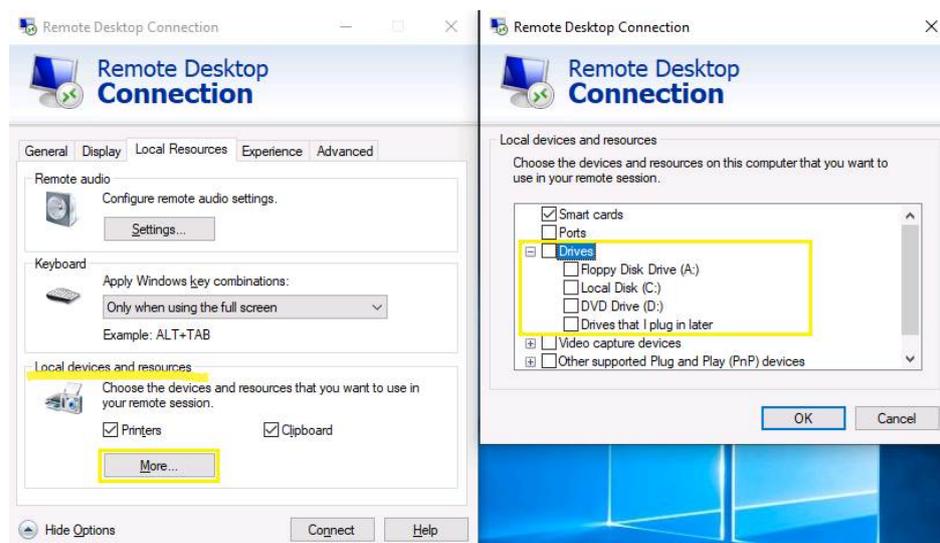
The remote desktop connection with redirected drive enabled loads the desired `rdpdr.sys` drive, which leads to BSOD and error code `0x18`, `0x50`, `0xa`, `0x27`, or `0x133`.

Solution

After you enable the remote desktop connection, disable redirect local drives.

1. Start the **Run** dialog box.
2. Enter `mstsc` and click **OK**.
The **Remote Desktop Connection** window is displayed.
3. Click **Options** in the lower left corner and click the **Local Resources** tab.
4. In the **Local devices and resources** pane, click **More**.
5. Deselect **Drives**.
6. Click **OK**.

Figure 12-5 Disabling redirect local drives



13 IIS Installation Issues

13.1 How Do I Install IIS on a Windows ECS?

Scenarios

This section describes how to install IIS on an ECS running Windows Server 2012 R2 Standard.

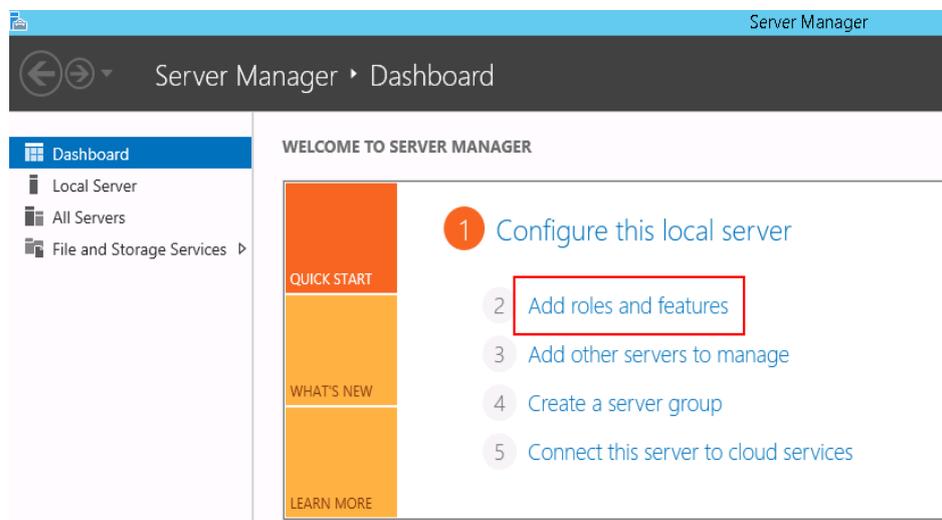
 **NOTE**

Only the procedure for installing IIS will be described. The procedure for installing applications is subject to actual requirements.

Procedure

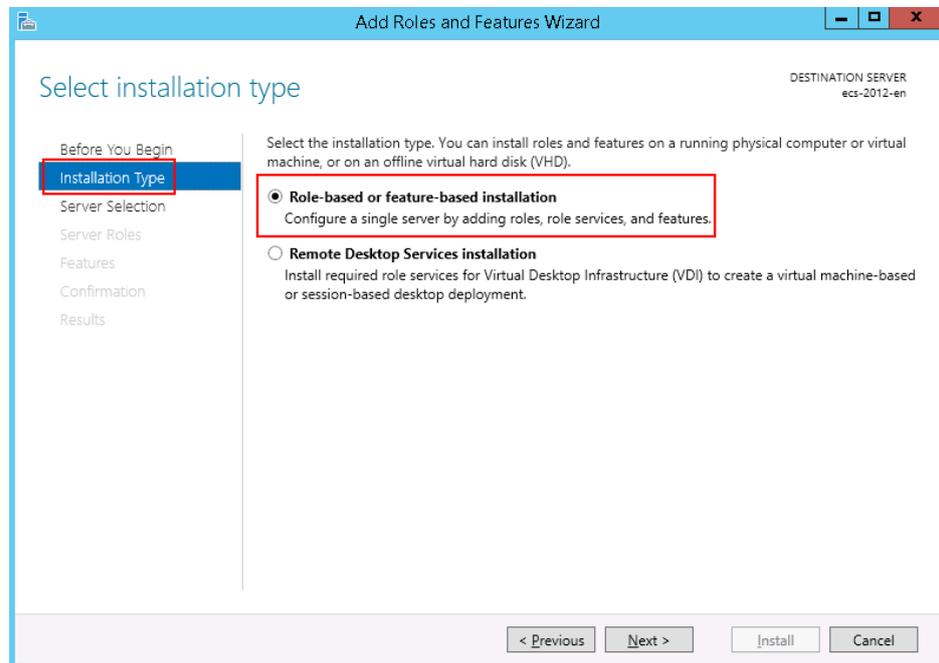
1. Open **Server Manager**.
2. Choose **QUICK START > Add roles and features**.

Figure 13-1 Adding roles and features



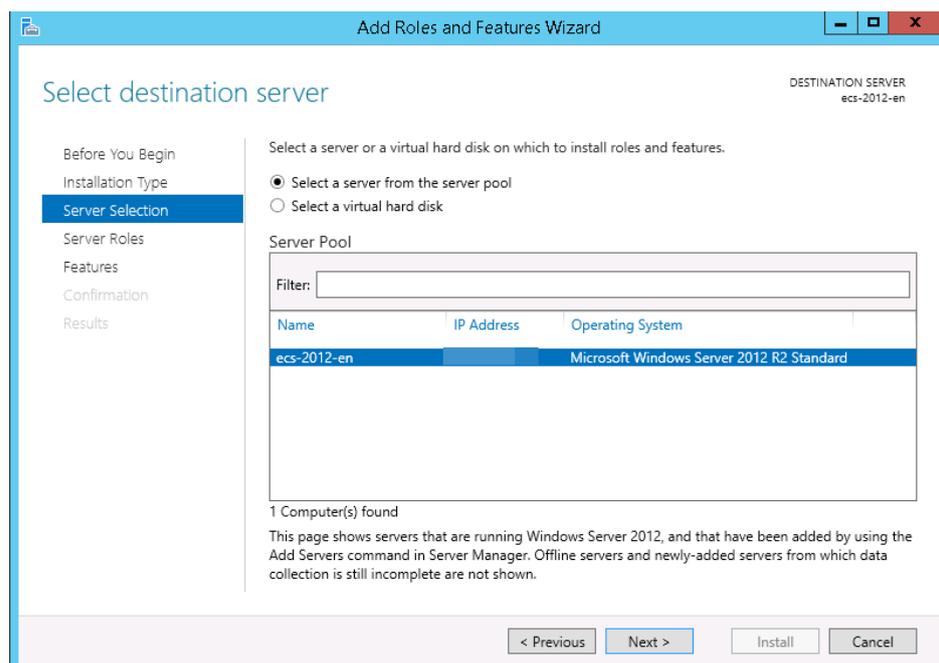
3. In the left navigation pane, choose **Installation Type**.

Figure 13-2 Installation type



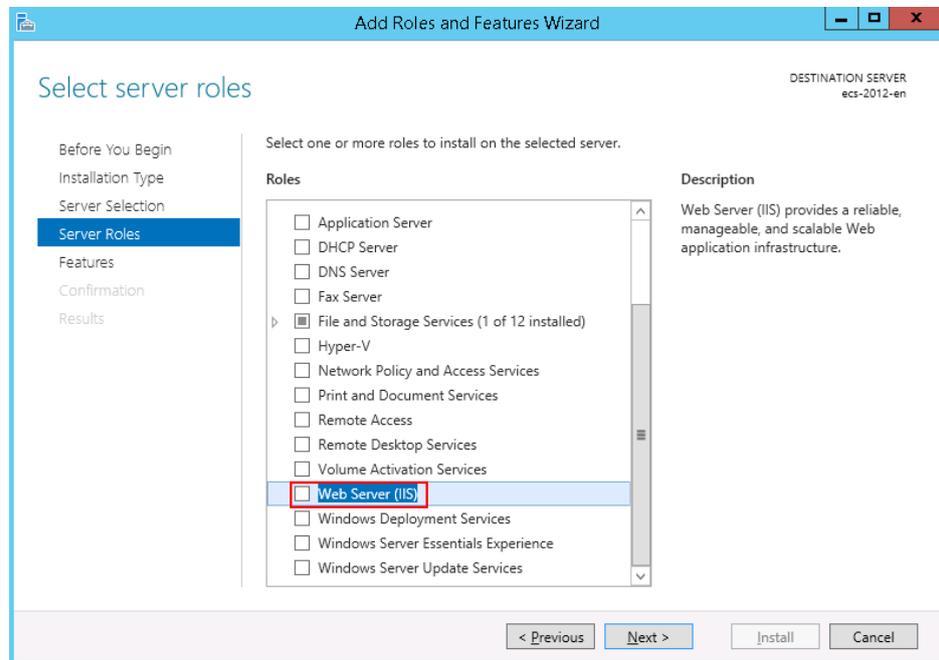
4. Select **Role-based or feature-based installation** and click **Next**.
5. Click **Server Selection**.
6. Select **Select a server from the server pool** and select a server from **Server Pool**.

Figure 13-3 Selecting a server



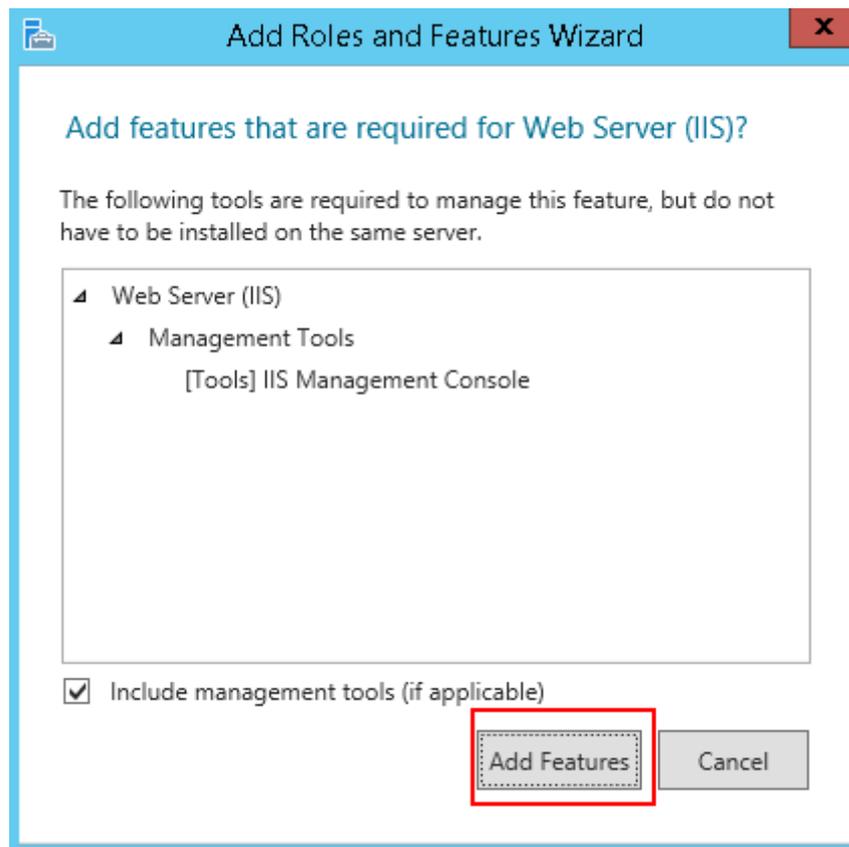
7. Click **Next**.
8. Click **Server Roles**.
9. In the role list, select **Web Server (IIS)**.

Figure 13-4 Web Server (IIS)



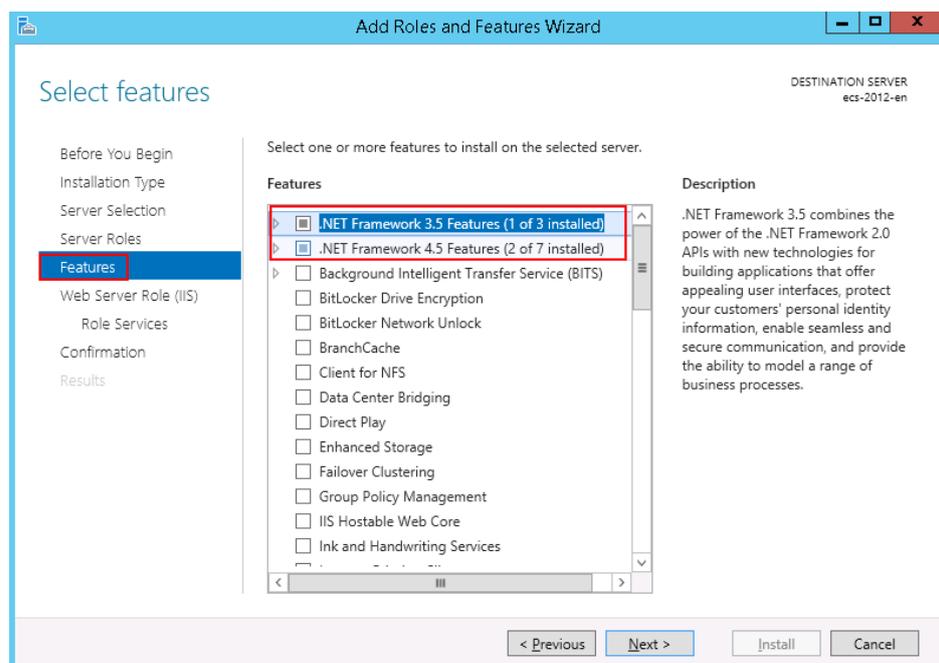
10. In **Add Roles and Features Wizard**, click **Add Features**.

Figure 13-5 Adding features



11. Choose **Features** on the left and select **.Net Framework 3.5** and **.Net Framework 4.5**.

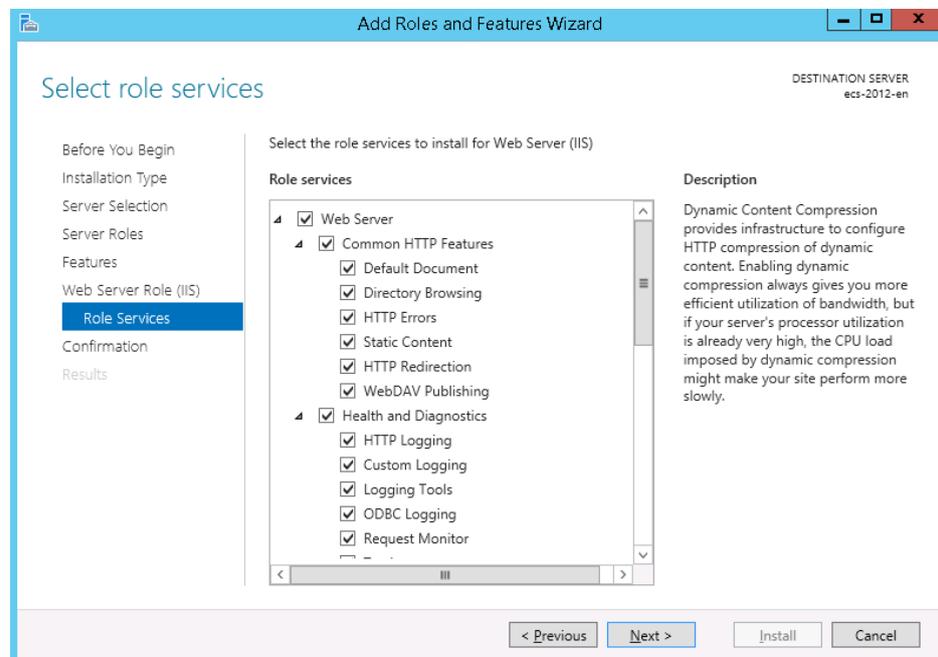
Figure 13-6 Selecting features



12. Choose **Web Server Role (IIS) > Role Services**. Select the required role services.

If you are not sure about which role service to install, you are advised to select all items except **FTP Server**.

Figure 13-7 Selecting role services



13. Click **Next**, confirm the roles to be installed, and click **Install**.

Helpful Links

For instructions about how to configure a domain name for IIS, see [Why Does an Error Occur When I Attempt to Change a Domain Name on IIS Manager?](#)

13.2 Why Does an Error Occur When I Attempt to Change a Domain Name on IIS Manager?

Symptom

A 404 error occurs on the website constructed using IIS.

Possible Causes

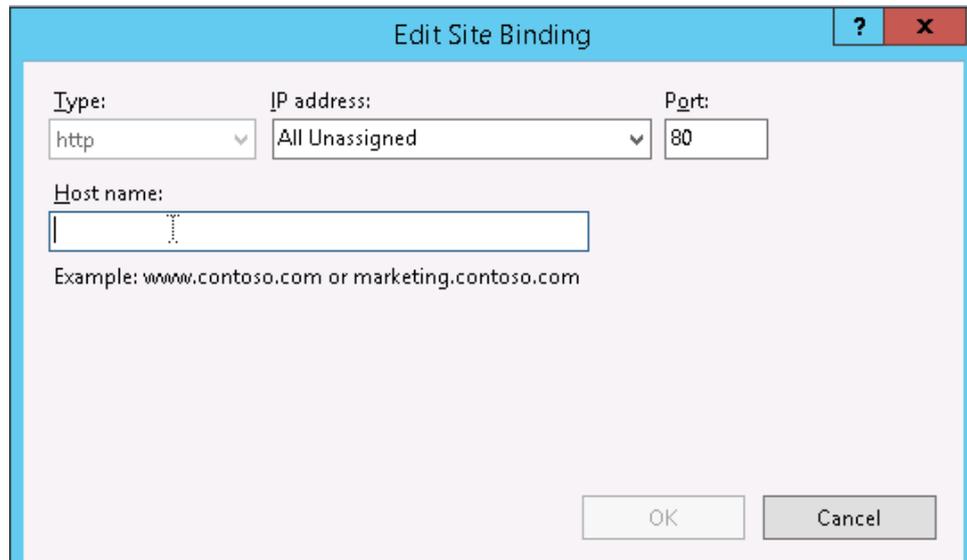
IP address of the domain name is not specified on the IIS Manager.

Procedure

This section describes how to edit site binding on the IIS Manager that is deployed on an ECS running Windows Server 2008 R2.

1. Log in to the ECS and choose **Start > Administrative Tools > IIS Manager**.

2. On the IIS Manager, click **Sites**.
3. Right-click the website to be modified and choose **Edit Bindings**.
Select a domain name and click **Edit** to add the private IP address of the specified ECS.



13.3 How Do I Redirect Web Pages?

Scenarios

You can refer to the procedure below to do a 301 redirect.

Procedure

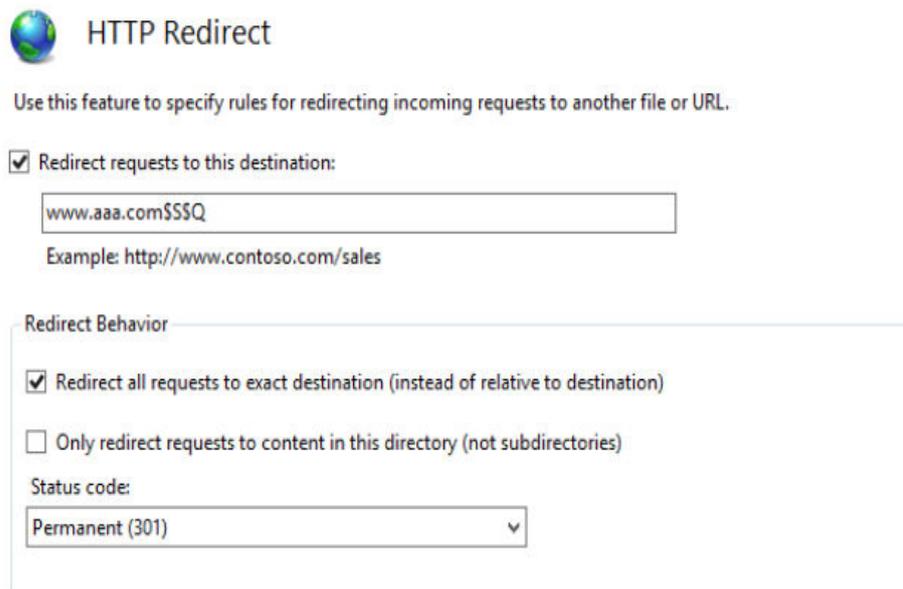
To do a 301 redirect on an ECS with the IIS installed, perform the following steps:

1. Launch a website, for example, www.aaa.com, in the IIS.
2. Create an empty folder on the disk.
3. Create a website, for example, aaa.com, in the IIS and point it to the empty folder.
4. On the home page, select **HTTP Redirection** and set parameters as follows:
 - Select **Redirect requests to this destination**.
 - Enter **www.aaa.com\$\$\$Q** in the text box.

NOTE

The purpose of adding \$\$\$Q to the end of a URL is to support redirection for a URL containing question marks (?).

- Select **Redirect all requests to exact destination (instead of relative to destination)**.
- Set **Status code** to **Permanent (301)**.

Figure 13-8 HTTP Redirect

The screenshot shows the 'HTTP Redirect' configuration window. At the top, there is a globe icon and the title 'HTTP Redirect'. Below the title, a text box contains the destination URL 'www.aaa.comSSSQ'. An example URL 'http://www.contoso.com/sales' is shown below the text box. A checkbox labeled 'Redirect requests to this destination:' is checked. Below this, there is a section titled 'Redirect Behavior' which contains two checkboxes: 'Redirect all requests to exact destination (instead of relative to destination)' (checked) and 'Only redirect requests to content in this directory (not subdirectories)' (unchecked). At the bottom of the 'Redirect Behavior' section, there is a 'Status code:' label and a dropdown menu set to 'Permanent (301)'.

Helpful Links

In addition to redirection in the IIS, you can perform 301 redirects through code. The following describes how to perform 301 redirects in PHP and Apache.

- 301 redirect in PHP

```
<?php
Header("HTTP/1.1 301 Moved Permanently");
Header("Location: http://www.***.cn"); //Redirect to a URL with www.
?>
```

- 301 redirect in Apache

Example code for redirecting your domain to a www URL:

```
deny from all
RewriteEngine on
RewriteCond %{HTTP_HOST}^(****\.com)(:80)?[NC]
RewriteRule ^(.*) http://www.***.com/$1 [R=301,L]
order deny,allow
```

13.4 What Can I Do If the Error Message "Bad Request - Invalid Hostname" Is Displayed When I Use a Windows ECS to Access an IIS Website?

Symptom

When I use a Windows ECS to access an Internet Information Services (IIS) website, an error message "Bad Request - Invalid Hostname" is displayed, as shown in [Figure 13-9](#).

Figure 13-9 Error message

Bad Request - Invalid Hostname

HTTP Error 400. The request hostname is invalid.

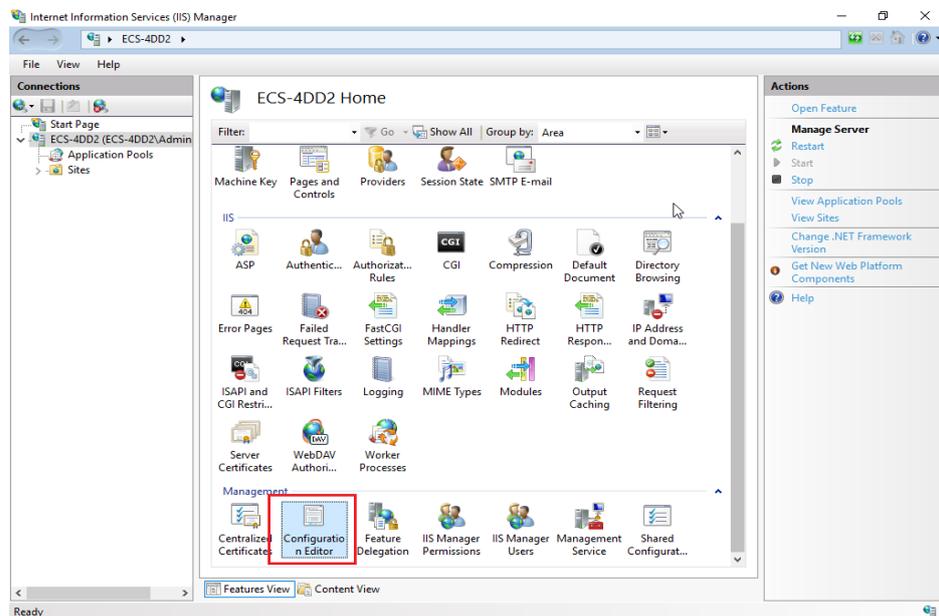
Possible Causes

In the IIS Manager, if the number of configured websites is greater than or equal to 100, no sites will be activated and the error message "Bad Request - Invalid Hostname" will be displayed when IIS is started.

Solution

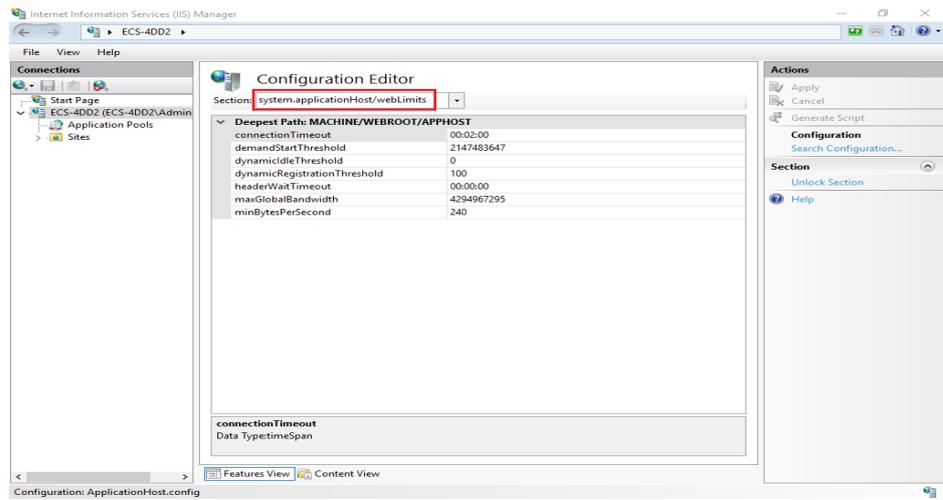
1. Log in to the Windows ECS.
2. Choose **Start > Windows Administrative Tools > Internet Information Services (IIS) Manager**.
3. On the displayed page, select the target server and double-click **Configuration Editor**.

Figure 13-10 Internet Information Services (IIS) Manager



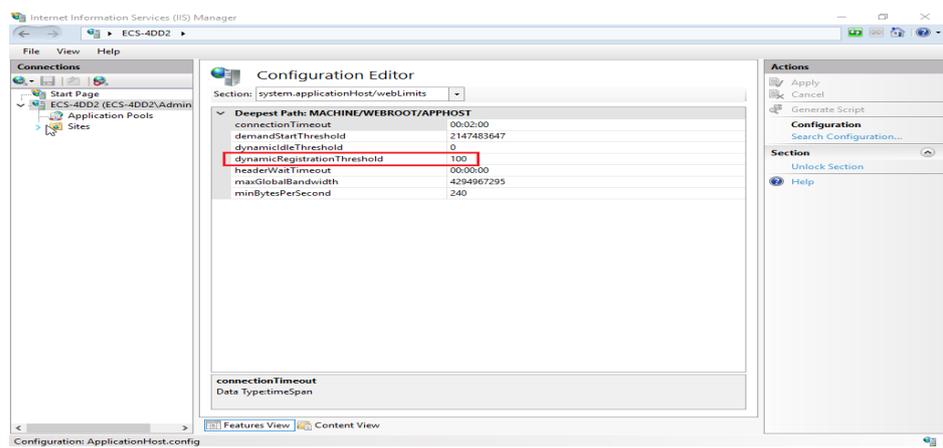
4. Select the **system.applicationHost/webLimits** section.

Figure 13-11 Configuration Editor



5. Set `dynamicRegistrationThreshold` to 100.

Figure 13-12 Setting the threshold



6. Access the website again.