

Elastic Cloud Server

FAQs

Issue 01
Date 2024-11-12



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. 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 Common FAQ	1
2 Product Consulting FAQ	2
2.1 What Are the Precautions for Using ECSs?	2
2.2 What Can I Do with ECSs?	2
3 ECS Creation FAQ	3
3.1 What Should I Do If the ECS Resources to Be Purchased Are Sold Out?	3
3.2 How Can I Set Sequential ECS Names When Creating Multiple ECSs?	3
3.3 What Is the Creation Time and Startup Time of an ECS?	6
3.4 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?	6
3.5 Why Cannot I View the ECSs Being Created Immediately After I Pay for Them?	7
3.6 Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image?	7
3.7 What Do I Do If I Selected an Incorrect Image for My ECS?	8
3.8 How Quickly Can I Obtain an ECS?	8
3.9 How Can I Manage ECSs by Group?	8
4 ECS Deletion and Unsubscription FAQ	10
4.1 What Happens After I Click the Delete Button?	10
4.2 Can a Deleted ECS Be Provisioned Again?	10
4.3 Can a Deleted ECS Be Restored?	10
4.4 How Do I Delete or Restart an ECS?	11
4.5 Can I Forcibly Restart or Stop an ECS?	11
4.6 How Do I Unsubscribe from ECSs?	11
5 Remote Login FAQ	13
5.1 Login Preparations	13
5.1.1 What Are the Login Requirements for ECSs?	13
5.1.2 What Are the Username and Password for Remote Logins?	14
5.1.3 Why Cannot I Use the Username and Password Configured During the Creation of a GPU-accelerated ECS to Log In to the ECS Through SSH?	14
5.2 Remote Logins	15
5.2.1 Why Can't I Log In to My Linux ECS?	15
5.2.2 How Can I Change a Remote Login Port?	21
5.2.3 Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS?	22

5.2.4 What Browser Version Is Required to Remotely Log In to an ECS?.....	24
5.2.5 How Can I Log In to an ECS After It Exchanged the System Disk with Another ECS Running the Same OS?.....	24
5.3 VNC Login.....	25
5.3.1 Why Are Garbled Characters Displayed When I Log In to My ECS Using VNC?.....	25
5.3.2 What Should I Do If I Cannot View Data After Logging In to an ECS Using VNC?.....	26
5.3.3 Why Does a Blank Screen Appear After I Attempted to Log In to an ECS Using VNC?.....	26
5.3.4 What Should I Do If Error Code 1006 or 1000 Is Displayed When I Log In to an ECS Through the Management Console?.....	26
5.4 Remote Login Errors on Linux.....	27
5.4.1 Why Am I Seeing the Error Message "Module is unknown" When I Remotely Log In to a Linux ECS?.....	27
5.4.2 What Should I Do If Error Message "Permission denied" Is Displayed When I Remotely Log In to a Linux ECS?.....	29
5.4.3 What Should I Do If Error Message "read: Connection reset by peer" Is Displayed When I Remotely Log In to a Linux ECS?.....	31
5.4.4 Why Am I Seeing the Error Message "Access denied" When I Remotely Log In to a Linux ECS?.....	32
5.4.5 What Should I Do If Error Message "Disconnected: No supported authentication methods available" Is Displayed When I Remotely Log In to a Linux ECS?.....	33
6 Billing FAQ.....	34
6.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing Modes?.....	34
6.2 Will I Be Billed After ECSs Are Stopped?.....	35
6.3 Can I Switch Between Yearly/Monthly and Pay-per-Use Billing Modes?.....	37
6.4 FAQs About ECS Frozen, Deletion, and Unsubscription.....	38
6.5 How Can I Stop an ECS from Being Billed?.....	39
6.6 FAQs About Spot ECSs.....	40
6.7 How Can I Renew ECSs?.....	42
6.8 How Can I Automatically Renew a Yearly/Monthly ECS?.....	42
6.9 Will I Receive a Notification If My Account Balance Is Insufficient?.....	42
7 Region and AZ FAQ.....	44
7.1 What Is AZ and How Can I Select and View an AZ?.....	44
7.2 What Is a Region?.....	44
7.3 Are Products Different in Different Regions?.....	44
7.4 Is Data Transmission Between AZs Billed?.....	45
7.5 Can I Change the Region or AZ After an ECS Is Purchased?.....	45
7.6 Can I Migrate an ECS to Another Region, AZ, or Account?.....	45
7.7 Can a Load Balancer Distribute Traffic to ECSs in Different Regions?.....	46
7.8 Can Components Contained in an Application Be Distributed to Different Regions?.....	46
8 OS FAQ.....	47
8.1 Do ECSs Support GUI?.....	47
8.2 How Can I Install a GUI on an ECS Running CentOS 6?.....	47
8.3 How Can I Install a GUI on an ECS Running CentOS 7?.....	48
8.4 How Can I Install a GUI on an ECS Running Ubuntu?.....	48

8.5 Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS?.....	54
8.6 How Can I Upgrade the Kernel of a Linux ECS?.....	56
8.7 Why Cannot My ECS OS Start Properly?.....	58
8.8 How Can I Fix the Meltdown and Spectre Security Vulnerabilities on Intel Processor Chips?.....	58
8.9 How Can I Enable SELinux on an ECS Running CentOS?.....	62
9 Disk Partition, Attachment, and Expansion FAQ.....	63
9.1 How Can I Adjust System Disk Partitions?.....	63
9.2 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS?.....	69
9.3 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?..	71
9.4 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?.....	72
9.5 Can I Attach Multiple Disks to an ECS?.....	75
9.6 What Are the Requirements for Attaching an EVS Disk to an ECS?.....	76
9.7 Which ECSs Can Be Attached with SCSI EVS Disks?.....	76
9.8 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?.....	77
9.9 Why Does a Linux ECS with a SCSI Disk Attached Fails to Be Restarted?.....	79
9.10 How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Same ECS Group?.....	79
9.11 Can All Users Use the Encryption Feature?.....	81
9.12 How Can I Add an ECS with Local Disks Attached to an ECS Group?.....	83
9.13 Will My EVS Disk Be Deleted When I Delete Its Server?.....	83
9.14 Why Does the Disk Drive Letter Change After the ECS Is Restarted?.....	83
9.15 How Can I Obtain Data Disk Information If Tools Are Uninstalled?.....	85
9.16 How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attached?.....	86
9.17 Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 File System?.....	87
10 Data Backup and Restoration FAQ.....	89
10.1 Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?	89
11 Network Configuration FAQ.....	91
11.1 How Can I Configure the NTP and DNS Servers for an ECS?.....	91
11.2 Does HUAWEI CLOUD Provide the NTP Server and How Can I Configure It?.....	92
11.3 How Do I Configure DNS for an ECS?.....	96
11.4 Can the ECSs of Different Accounts Communicate over an Intranet?.....	97
11.5 Will ECSs That I Purchased Deployed in the Same Subnet?.....	98
11.6 How Do I Change the CIDR Block of an ECS Subnet?.....	98
11.7 How Do I Configure Port Mapping?.....	99
11.8 How Can I Obtain the MAC Address of My ECS?.....	100
11.9 How Can I View and Modify Kernel Parameters of a Linux ECS?.....	101
11.10 Why Is the NIC Not Working?.....	106
11.11 Why Can't I Use DHCP to Obtain a Private IP Address?.....	109
11.12 How Can I Test the Network Performance of Linux ECSs?.....	109

11.13 What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?.....	118
11.14 Will NICs Added to an ECS Start Automatically?.....	119
11.15 How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?.....	119
11.16 How Can I Manually Configure an IP Address for an InfiniBand NIC?.....	120
11.17 How Can I Add a Static Route to a CentOS 6.5 OS?.....	121
12 EIP FAQ.....	123
12.1 Can Multiple EIPs Be Bound to an ECS?.....	123
12.2 Can an ECS Without an EIP Bound Access the Internet?.....	124
12.3 What Should I Do If an EIP Cannot Be Pinged?.....	124
12.4 Why Can I Remotely Access an ECS But Cannot Ping It?.....	128
12.5 Why Is Outbound Access on TCP Port 25 Blocked?.....	129
12.6 Why My EIPs Are Frozen? How Do I Unfreeze My EIPs?.....	129
13 Password and Key Pair FAQ.....	131
13.1 How Can I Change the Password for Logging In to a Linux ECS?.....	131
13.2 What Is the Default Password for Logging In to a Linux ECS?.....	131
13.3 How Can I Set the Validity Period of the Image Password?.....	131
13.4 Resetting the Password for Logging In to an ECS in the OS.....	132
13.5 What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS?.....	133
13.6 What Should I Do If I Cannot Log In to My ECS Using the Initial Password After I Use It for a Period of Time?.....	134
13.7 Disabling SELinux.....	134
13.8 How Can I Obtain the Key Pair Used by My ECS?.....	135
13.9 How Can I Use a Key Pair?.....	135
13.10 Can I Download a Key Pair from My Phone?.....	136
13.11 What Should I Do If a Key Pair Cannot Be Imported?.....	136
13.12 Why Does the Login to My Linux ECS Using a Key File Fail?.....	136
13.13 What Should I Do If I Cannot Download a Key Pair?.....	137
13.14 Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?.....	137
13.15 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded?.....	139
14 Application Deployment and Software Installation FAQ.....	141
14.1 Can a Database Be Deployed on an ECS?.....	141
14.2 Does an ECS Support Oracle Databases?.....	141
15 File Upload/Data Transfer FAQ.....	142
15.1 How Do I Upload Files to My ECS?.....	142
15.2 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?.....	142
15.3 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	144
15.4 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	145
15.5 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Linux ECS?.....	147
15.6 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?.....	148

15.7 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP?.....	149
15.8 What Should I Do If Writing Data Failed When I Upload a File Using FTP?.....	149
15.9 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server?.....	151
15.10 Why Do I Fail to Connect to a Linux ECS Using WinSCP?.....	153
16 ECS Failure FAQ.....	155
16.1 How Do I Handle Error Messages Displayed on the Management Console?.....	155
16.2 What Should I Do If Emails Configured on an ECS Cannot Be Sent?.....	158
16.3 Why Is the Memory of an ECS Obtained by Running the free Command Inconsistent with the Actual Memory?.....	159
16.4 Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted?.....	161
16.5 Is an ECS Hostname with Suffix .novalocal Normal?.....	161
16.6 Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?.....	162
16.7 How Can a Changed Static Hostname Take Effect Permanently?.....	163
16.8 Why Can't My Linux ECS Obtain Metadata?.....	166
17 Slow ECS Response FAQ.....	170
17.1 Why Is My Windows ECS Running Slowly?.....	170
17.2 Why Is My Linux ECS Running Slowly?.....	174
18 Specification Modification FAQ.....	180
18.1 What Should I Do If My Specifications Modification Request Failed to Submit?.....	180
18.2 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?.....	180
19 OS Change FAQ.....	183
19.1 Does OS Change Incur Fees?.....	183
19.2 Can I Install or Upgrade the OS of an ECS?.....	183
19.3 Can I Change the OS of an ECS?.....	183
19.4 How Long Does It Take to Change an ECS OS?.....	183
19.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?.....	184
19.6 Does OS Reinstallation Incur Fees?.....	184
19.7 Can I Select Another OS During ECS OS Reinstallation?.....	185
19.8 How Long Does It Take to Reinstall an ECS OS?.....	185
20 ECS Security Check FAQ.....	186
20.1 How Does an ECS Defend Against DDoS Attacks?.....	186
20.2 Are ECSs with Simple Passwords Easily Attacked?.....	187
20.3 How Is ECS Security Ensured?.....	187
20.4 How Can I Disable Operation Protection?.....	188
21 Resource Management and Tag FAQ.....	189
21.1 How Can I Create and Delete Tags and Search for ECSs by Tag?.....	189
22 Image Source Management FAQ.....	191
22.1 How Can I Use an Automated Tool to Configure a Huawei Cloud Image Source (x86_64 and Kunpeng)?.....	191

22.2 How Can I Use an EPEL Image Source (x86_64 or Kunpeng) Provided by Huawei Cloud?.....	192
22.3 How Can I Use a Pypi Image Source (x86_64 or Kunpeng) Provided by Huawei Cloud?.....	193
22.4 What Can I Do If the Download Speed Is Slow After the yum makecache Command Is Executed?.	194
23 Internet Inaccessible FAQ.....	196
23.1 Why Does My Linux ECS Fail to Access the Internet?.....	196
24 Website or Application Inaccessible FAQ.....	203
24.1 Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a Linux ECS?	203
25 Slow Website or Application Response FAQ.....	205
25.1 Why Accessing a Website Outside the Chinese Mainland Is Slow on an ECS?.....	205

1 Common FAQ

Remote Logins

- [Why Can't I Log In to My Linux ECS?](#)
- [What Are the Username and Password for Remote Logins?](#)

Internet Access Failures

- [Why Does My Linux ECS Fail to Access the Internet?](#)
- [Can an ECS Without an EIP Bound Access the Internet?](#)

Passwords and Key Pairs

- [What Are the Username and Password for Remote Logins?](#)

Ping Failures

- [What Should I Do If an EIP Cannot Be Pinged?](#)
- [Why Can I Remotely Access an ECS But Cannot Ping It?](#)

2 Product Consulting FAQ

2.1 What Are the Precautions for Using ECSs?

- Do not upgrade ECS kernel or OS versions. If you want to upgrade the main OS version, for example, from CentOS 7.2 to Cent OS 7.3, use the provided OS changing function.
- Do not uninstall the performance optimization software pre-installed on your ECSs.
- Do not change NIC MAC addresses. Otherwise, the network connection will fail.

2.2 What Can I Do with ECSs?

You can use ECSs just like traditional physical servers. On an ECS, you can deploy any service application, such as an email system, web system, and Enterprise Resource Planning (ERP) system. After creating an ECS, you can use it like using your local computer or physical server.

3 ECS Creation FAQ

3.1 What Should I Do If the ECS Resources to Be Purchased Are Sold Out?

Each region has two or three AZs. If resources in an AZ are sold out, you can change the AZ and purchase resources in another AZ.

3.2 How Can I Set Sequential ECS Names When Creating Multiple ECSs?

Scenarios

When creating multiple ECSs at the same time, you can use either of the following methods to sequentially name the ECSs:

- Automatic naming: The system automatically adds a hyphen followed by a four-digit incremental number to the end of each ECS name.
- Customizable naming: You can customize a naming rule in the format "name_prefix[begin_number,bits]name_suffix". The system will automatically name the ECSs according to naming rule you specify.

This section describes how to use the two methods to name ECSs.

Automatic Naming

You can customize the name according to the following naming rules: The name must contain 1 to 64 characters that can be only letters, digits, underscores (_), and hyphens (-).

When you create multiple ECSs at the same time, the system automatically adds a hyphen followed by a four-digit incremental number to the end of each ECS name. In this case, the customized name is 1 to 59 characters long. For example, if you are creating multiple ECSs and enter **ecs** for the ECS name, the created ECSs will be named **ecs-0001**, **ecs-0002**, and so on. If you create multiple ECSs again, the

values in the new ECS names increase from the existing maximum value. For example, the existing ECS with the maximum number in name is **ecs-0010**. If you enter **ecs**, the names of the new ECSs will be **ecs-0011**, **ecs-0012**, When the value reaches **9999**, it will start from **0001**.

Allow duplicate name: allows ECS names to be duplicate. If you select **Allow duplicate name** and create multiple ECSs in a batch, the created ECSs will have the same name.

- Example 1: If there is no existing ECS and you enter **ecs-f526**, the ECSs will be named **ecs-f526-0001**, **ecs-f526-0002**, **ecs-f526-0003**,
- Example 2: If there is an ECS named **ecs-f526-0010** and you enter **ecs-f526**, the ECSs will be named **ecs-f526-0011**, **ecs-f526-0012**, **ecs-f526-0013**,
- Example 3: If there is an ECS named **ecs-0010** and you select **Allow duplicate ECS name**, all the ECSs will be named **ecs-0010**.

Customizable Naming

You can customize a naming rule in the format "name_prefix[begin_number,bits]name_suffix". The system will automatically name the ECSs according to naming rule you specify.

Field Description for a Customizable Naming Rule

Figure 3-1 shows the format of a customizable naming rule.

Figure 3-1 Format of a customizable naming rule




Table 3-1 describes these parameters.

Table 3-1 Parameters in a customizable naming rule

Field	Mandatory	Description	Example
name_prefix	Yes	ECS name prefix The name prefix can contain only letters, digits, underscores (_), and hyphens (-).	ecs
[begin_number,bits]	Yes	Sequence numbers that increase in ascending order to differentiate multiple ECSs.	[0,4]

Field	Mandatory	Description	Example
name_suffix	No	ECS name suffix The name suffix can contain only letters, digits, underscores (_), and hyphens (-).	f526

Table 3-2 [begin_number,bits] parameters

Field	Mandatory	Description	Example
begin_number	No	Begin number of ECS names. The begin number ranges from 0 to 9999. The default value is 0 .	0
bits	No	Number of bits for the sequential numbers in ECS names. The value ranges from 1 to 4. The default value is 4 .	4

Notes on Using Customizable Naming

- Customized names cannot be duplicate.
- No space is allowed in [begin_number,bits].
- If the bits of "Begin number + Number of ECSs to be created - 1" is greater than the specified bits, the bits of "Begin number + Number of ECSs to be created - 1" will be used.

For example, if [begin_number,bits] is set to [8,1] and the number of ECSs to be created is 2, the bits of "Begin number + Number of ECSs to be created - 1" is the same as the specified bits (1). Then, the ECSs will be named *name_prefix8name_suffix* and *name_prefix9name_suffix*.

If [begin_number,bits] is set to [8,1] and the number of ECSs to be created is 3, the specified bits is 1, the bits of "Begin number + Number of ECSs to be created - 1" (value 10, bits 2) is different from the specified bits (1). Therefore, the bits of "Begin number + Number of ECSs to be created - 1" will be used, which is 2.

The ECSs will be named *name_prefix08name_suffix*, *name_prefix09name_suffix*, and *name_prefix10name_suffix*.

- If the value of "Begin number + Number of ECSs to be created" is greater than the maximum value **9999**, the sequential numbers that exceed **9999** will consistently be **9999**.
- If [begin_number,bits] is set to [] or [,], the begin number starts from **0**, and the number of bits is **4** by default.

- If [begin_number,bits] is set to [99] or [99,], the begin number starts from **99**, and the number of bits is **4** by default.

Customizable Naming Examples

- Example 1: If you select customizable naming and enter *name_prefix[,]name_suffix*,
The ECSs will be named *name_prefix0000name_suffix*,
name_prefix0001name_suffix, *name_prefix0002name_suffix*,
- Example 2: If you select customizable naming and enter *name_prefix[]name_suffix*,
The ECSs will be named *name_prefix0000name_suffix*,
name_prefix0001name_suffix, *name_prefix0002name_suffix*,
- Example 3: If you select customizable naming and enter *name_prefix[9,]name_suffix*,
The ECSs will be named *name_prefix0009name_suffix*,
name_prefix0010name_suffix, *name_prefix0011name_suffix*,
- Example 4: If you select customizable naming and enter *name_prefix[,3]name_suffix*,
The ECSs will be named *name_prefix000name_suffix*,
name_prefix001name_suffix, *name_prefix002name_suffix*,
- Example 5: If you select customizable naming and enter *name_prefix[8]name_suffix*,
The ECSs will be named *name_prefix0008name_suffix*,
name_prefix0009name_suffix, *name_prefix0010name_suffix*,
- Example 6: If you select customizable naming and enter *name_prefix[9999]name_suffix*,
All the ECSs will be named *name_prefix9999name_suffix*.
- Example 7: If you select customizable naming and enter *name_prefix[8]*,
The ECSs will be named *name_prefix0008*, *name_prefix0009*,
name_prefix0010,

3.3 What Is the Creation Time and Startup Time of an ECS?

Creation time: time when the ECS is created on the cloud platform.

Startup time: time when the ECS is started for the first time.

3.4 Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?

Symptom

When you attempt to create an ECS with an EIP bound on the management console, the ECS creation was successful but the EIP binding failed due to insufficient EIPs. Although the **Failures** area showed that the ECS creation failed,

the ECS was displayed in the ECS list. The results of the ECS creation task were inconsistent.

Root Cause

- The ECS list displays created ECSs.
- The **Failures** area shows the ECS creation status, including the statuses of subtasks, such as creating ECS resources and binding an EIP. Only when all subtasks are successful, the ECS is created.

If the ECS is created but EIP binding failed, the task failed. However, the ECS you created is temporarily displayed in the list. After the system rolls back, the ECS is removed from the list.

3.5 Why Cannot I View the ECSs Being Created Immediately After I Pay for Them?

You can view the ECSs being created only after the system disks attached to the ECSs are created. This requires a period of time.

3.6 Why Does It Take Longer to Create ECSs When I Use a Full-ECS Image?

Symptom

When you use a full-ECS image that was created using a CSBS backup to create ECSs, the process is time-consuming or the system displays a message indicating that the image cannot be used to rapidly create ECSs.

Cause Analysis

If your full-ECS image is in the old backup format provided by CSBS, this issue occurs.

NOTE

- CSBS has a new backup format. You can rapidly create ECSs if the full-ECS image is in this format.

Solution

If you want to use a full-ECS image to rapidly create ECSs, ensure that the full-ECS image is created using a CSBS backup in the new format. The procedure is as follows:

- Scenario 1: The ECS based on which the target CSBS backup is created is available.

Back up the original ECS on the **Cloud Server Backup Service** page and use the new format to create a full-ECS image. You can use the full-ECS image to rapidly create ECSs.

- For instructions about how to back up an ECS, see *Cloud Server Backup Service User Guide*.
- For instructions about how create a full-ECS image, see *Image Management Service User Guide*.
- Scenario 2: The ECS based on which the target CSBS backup is created is unavailable.
 - a. Use the full-ECS image to create a new ECS.
 - b. Back up the newly created ECS.
For details, see *Cloud Server Backup Service User Guide*.
 - c. Use the CSBS backup to create a full-ECS image.
For details, see *Image Management Service User Guide*.
You can use the full-ECS image to rapidly create ECSs.

3.7 What Do I Do If I Selected an Incorrect Image for My ECS?

You can change the image for your ECS on the ECS console.

1. Select the target ECS and click **Stop** in the upper left corner of the ECS list.
2. Locate the row that contains the target ECS, choose **More > Manage Image/Disk/Backup > Change OS** in the **Operation** column.
The **Change OS** dialog box is displayed.
3. Select the target image type and image.
4. Set the login mode. You can select **Password** or **Key pair**.
5. Set the other parameters and click **OK**.

After the application is submitted, the ECS status changes to **Changing OS**. The OS changing has been successfully completed when the ECS status changes to **Running**.

For details, see [Changing the OS](#).

3.8 How Quickly Can I Obtain an ECS?

Obtaining an ECS can take as little as a few minutes.

The time it takes to obtain an ECS depends on ECS specifications, available resources (such as EVS disks and EIPs), and system load.

NOTE

If it takes a long time to obtain your ECS, contact customer service.

3.9 How Can I Manage ECSs by Group?

You cannot manage ECSs by folders or groups, but you can use tags to organize your ECSs


Tags help you group your ECSs by things by whatever categories are useful to you.
For more information, see [Tag Management](#).

4 ECS Deletion and Unsubscription FAQ

4.1 What Happens After I Click the Delete Button?

After you click **Delete**, the selected ECSs will be deleted. You can also choose to delete the EVS disks and EIPs together with the selected ECSs. If you do not delete them, they will be retained. If necessary, you can manually delete them later.

To delete selected ECSs, perform the following operations:

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. Select the ECSs to be deleted.
4. Above the ECS list, choose **More > Delete**.

4.2 Can a Deleted ECS Be Provisioned Again?

No. ECSs in the **Deleted** state cannot provide services and are soon removed from the system.

A deleted ECS is retained in the ECS list on the management console only for a short period of time before it is permanently removed from the system. You can purchase a new ECS with the same specifications again.

4.3 Can a Deleted ECS Be Restored?

No. The data of a deleted ECS cannot be restored. Therefore, before deleting an ECS, back up or migrate its data.

4.4 How Do I Delete or Restart an ECS?

Deleting an ECS

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Under **Compute**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Delete** in the **Operation** column.

NOTE


If you choose to delete the EIP and data disks associated with the ECS when deleting it, no charges will apply for the EIP and data disks. However, if they are not deleted, they will continue to incur fees.

Restarting an ECS

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Under **Compute**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Restart** in the **Operation** column.

4.5 Can I Forcibly Restart or Stop an ECS?

Yes. If an ECS remains in the **Restarting** or **Stopping** state for over 30 minutes after it is restarted, you can forcibly restart or stop the ECS as follows:

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, click **Elastic Cloud Server**.
4. Select the target ECS and click **Restart** or **Stop**.
A dialog box is displayed to confirm whether you want to restart or stop the ECS.
5. Select **Forcibly restart the preceding ECSs** or **Forcibly stop the preceding ECSs**.
6. Click **OK**.

4.6 How Do I Unsubscribe from ECSs?

A yearly/monthly ECS can be unsubscribed, including the renewed resources and the resources that are being used. After the unsubscription, the ECS can no longer be used. A handling fee will be charged for unsubscribing from a resource.

Notes

- Unsubscribing from an ECS involves the renewed resources and the resources that are being used. After the unsubscription, the ECS is unavailable.
- Solution product portfolios can only be unsubscribed from as a whole.
- If an order contains resources in a primary-secondary relationship, you need to unsubscribe from the resources separately.
- For details about how to unsubscribe from a resource, see [Unsubscription Rules](#).

Procedure

NOTICE

Before requesting an unsubscription, ensure that you have migrated or backed up any data saved on the ECS that will be unsubscribed from. After the unsubscription is complete, the ECS and any data it contains will be permanently deleted.

1. Switch to the [Unsubscriptions](#) page.
2. Click the **Active Resources** tab.
3. Unsubscribe from a single resource or from resources in a batch.
 - To unsubscribe from a single resource, click **Unsubscribe** for the target resource.
 - To unsubscribe from resources in a batch, select the target resources from the resource list and click **Unsubscribe** in the upper left corner of the resource list.
4. View the unsubscription information, select **I have confirmed that a handling fee will be charged for this unsubscription**, and click **Unsubscribe**.

5 Remote Login FAQ

5.1 Login Preparations

5.1.1 What Are the Login Requirements for ECSs?

Linux

- Ensure that the ECS has an EIP bound (only required for SSH logins).

SSH logins are available for Linux ECSs only. You can use a remote login tool, for example, [use PuTTY](#) to log in to your ECS. In such a case, the ECS must have an EIP bound.

- Check whether an ECS has an EIP bound.

For details, see [Assigning an EIP](#).

- Check whether an EIP can be pinged.

- If you use a public IP address, see [What Should I Do If an EIP Cannot Be Pinged?](#) for troubleshooting.

- If you use a private IP address, see [Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur When They Communicate?](#)

More information:

- For a Linux ECS authenticated using a key pair:

- For the first login, use an SSH key. For details, see [Remotely Logging In to a Linux ECS \(Using an SSH Key Pair\)](#).

- For a non-first login, if you want to use the remote login function (VNC) provided by the management console, log in to the ECS using the SSH key and set the password.

- For a key-pair-authenticated ECS, using a private key file to obtain its login password will fail because Cloud-Init may fail to inject the password.

5.1.2 What Are the Username and Password for Remote Logins?

Username for logging in to an ECS:

- For Linux: **root**

5.1.3 Why Cannot I Use the Username and Password Configured During the Creation of a GPU-accelerated ECS to Log In to the ECS Through SSH?

Solution

Log in to the ECS using VNC, modify the configuration file, and log in to the ECS through SSH.

1. On the ECS console, locate the ECS and click **Remote Login**.
2. On the login page, enter user **root** and its password.

NOTE

The password is the one you set during the ECS creation.

```
Connected (encrypted) to: QEMU (i-000FA82E) Before you exit,ensure that computer is locked.
ec2: #####
ec2: ----BEGIN SSH HOST KEY FINGERPRINTS----
ec2: 256 a4:9c:e9:d9:35:68:26:27:c1:0c:43:77:ce:db:17:35 (ECDSA)
ec2: 2048 67:e0:3d:0e:1a:0b:7a:ee:46:5a:1c:4e:44:c3:6f:b7 (RSA)
ec2: ----END SSH HOST KEY FINGERPRINTS----
ec2: #####
----BEGIN SSH HOST KEY KEYS----
ecdsa-sha2-nistp256 AAAAEZUjZHhLXNoYTIibmlzdHAgNTYAAAAIbmlzdHAgNTYAAABBBGgDOEcD
5y0ug132daqN011YL3V8R1ZFx91ywQT8mBGUxh7X72y1opMBhQxP2E7t0o5JXt5i83IP1+YPLRi9X8w=

ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCB8xDnU4ZXP8+4pqD810A7fUz_jhhwR487z8uHa+eEvG
H1dWAU0tY4XrSZE73y_jhSvXyaGY/1GLpecco6MgdQfW7p8/rnu+TnJ+CHUZ/x0cCDSpInZpYe2cWTrsg
P8GpVZK6ZgqxFcWmkJMMZEYR_j51BtVAVUBHCeh7A8bbGJaOUzCuLuUwH0edpdMU1u1BD4bGP/5zsPDGo
y_jexL1avWvsRReaWZAWQ6nTxJ55qx2fs54Gb53SUItleIE2u3aH4DtwCeSox1+/7_jc3tSmcc/PHvWnb5
562U0sI1c6p+9xmcI8Rm8KncKr8NMUv3xR/BbGIKcY4dniZCC81Q51B7yAs?
----END SSH HOST KEY KEYS----
cloud-init[37321]: Cloud-init v. 0.7.5 finished at Wed, 17 Jan 2018 06:39:54 +000
0. Datasource DataSourceEc2. Up 36.21 seconds

CentOS Linux 7 (Core)
Kernel 3.10.0-123.el7.x86_64 on an x86_64

Login with linux/cloud.1234, sudo for root.
ecs-dec7 login:
```

3. In the `/etc/ssh/` directory, modify the three configuration items in the `sshd_config` file, as shown in the following figure.

```
SyslogFacility AUTH
PermitRootLogin yes
# Do not enable sshd passwd auth without ensuring really strong passwords
PasswordAuthentication yes
ChallengeResponseAuthentication no
GSSAPIAuthentication no
GSSAPICleanupCredentials yes
UsePAM yes
AcceptEnv LANG LC_CTYPE LC_NUMERIC LC_TIME LC_COLLATE LC_MONETARY LC_MEASUREMENT
AcceptEnv LC_PAPER LC_NAME LC_ADDRESS LC_TELEPHONE LC_MEASUREMENT
AcceptEnv LC_IDENTIFICATION LC_ALL LANGUAGE
AcceptEnv XMODIFIERS
X11Forwarding yes
Subsystem sftp /usr/libexec/openssh/sftp-server
#UseDNS no
Ciphers aes256-ctr,aes192-ctr,aes128-ctr
KexAlgorithms diffie-hellman-group-exchange-sha256
AllowTcpForwarding no
GatewayPorts no
X11UseLocalhost yes
AllowAgentForwarding yes
PermitTunnel no
LogLevel VERBOSE
RSAAuthentication yes
PubkeyAuthentication yes
PermitEmptyPasswords no
RhostsRSAAuthentication no
HostbasedAuthentication no
IgnoreRhosts yes
AllowUsers root
~
~
~
~
"sshd_config" 31L, 938C written
bash-4.1#
```

4. Save the modification and exit. Then, run the following command to restart SSH:
service sshd restart
5. After the restart, use the SSH password to log in again.
6. If the fault persists, contact customer service.

5.2 Remote Logins

5.2.1 Why Can't I Log In to My Linux ECS?

Symptom

A Linux ECS cannot be logged in to due to some reasons. For example, the network is abnormal, the firewall does not allow access to the local port for accessing the remote desktop, or the ECS vCPUs are overloaded.

This section describes how to troubleshoot login failures on a Linux ECS.

If you cannot log in to your Linux ECS, follow the instructions provided in [Checking the VNC Login](#). Then, locate the login fault by referring to [Fault Locating](#).

Checking the VNC Login

Check whether you can log in to the ECS using VNC on the management console.

NOTE

See [What Are the Login Requirements for ECSs?](#) to learn the requirements for logging in to an ECS.

1. Log in to the management console.
2. Under **Compute**, choose **Elastic Cloud Server**.
3. In the **Operation** column of the target ECS, click **Remote Login**.
4. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Ctrl+Alt+Del** in the upper part of the remote login page to log in to the ECS.

NOTE

Do not press **CTRL+ALT+DELETE** on the physical keyboard because this operation does not take effect.

If the VNC login still fails, 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.

Fault Locating

If you can log in to the ECS using VNC but cannot log in to the ECS using a remote desktop connection, locate the fault as follows.

The following fault causes are sequenced based on their occurrence probability.

If the fault persists after you have ruled out a cause, check other causes.

Table 5-1 Possible causes and solutions

Possible Cause	Solution
The ECS is frozen or stopped.	Make sure that the ECS is in the Running state. For details, see Checking the ECS Status .
The entered username or password is incorrect.	The default username for Linux ECSs is root . If the password is incorrect, reset the password on the management console. For details, see Checking the Login Mode .
The ECS is overloaded.	If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur. For details, see Checking Whether the ECS Is Overloaded .
The ECS has no EIP bound.	To log in to an ECS using RDP or MSTSC, ensure that the ECS has an EIP bound. For details, see Checking Whether an ECS Has an EIP Bound .

Possible Cause	Solution
The access is blocked by the ISP.	Check whether you can access the ECS using another hotspot or network. For details, see Checking Whether the Network Is Normal .
The security group of the ECS denies inbound traffic on the remote login port.	Check whether the security group allows inbound traffic on the remote login port. For details, see Checking Whether the Security Group Is Correctly Configured .
The remote access port is incorrectly configured.	Check whether the remote access port is correctly configured on the local computer and the ECS. For details, see Checking Whether the Remote Access Port Is Correctly Configured .
An IP address whitelist for SSH logins has been configured.	Check whether an SSH login IP address whitelist is configured after HSS is enabled. For details, see Checking the IP Address Whitelist for SSH Logins (with HSS Enabled) .
An OS fault has occurred.	The file system is damaged. For details, see Checking Whether an OS Fault Has Occurred .
The access is blocked by third-party antivirus software.	Disable or uninstall the third-party antivirus software and try again. For details, see Checking Whether the Access Is Blocked by Antivirus Software .
The cause is displayed in the error message.	If an error message is displayed during remote login, check the operation guide based on the error information. For details, see Checking Whether an Error Occurred During a Remote Login .

Checking the ECS Status

Check whether the ECS is in the **Running** state on the management console. If the ECS is stopped, start it and try to log in to the ECS again.

Checking the Login Mode

Check the login mode you set when you created the ECS.

- **Password:** Check whether the login password is correct. If you forgot your password, reset the password. After you reset the password, restart the ECS for the new password to take effect.
- **Key pair**
 - For the first login, use an SSH key. For details, see [Remotely Logging In to a Linux ECS \(Using an SSH Key Pair\)](#).
 - For a non-first login, if you want to use the remote login function (VNC) provided by the management console, log in to the ECS using the SSH key and set the password.

- **Set password later:** If you did not set a login mode when you create an ECS, you can reset the password on the ECS console by choosing **More > Reset Password** in the **Operation** column of the target ECS. After you reset the password, restart the ECS for the new password to take effect.

Checking Whether the ECS Is Overloaded

If the bandwidth or CPU usage of the ECS is excessively high, login failures may occur.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

- If the login failure is caused by high CPU usage, perform the following operations to reduce the CPU usage:
 - Stop certain processes that are not used temporarily and try again.
 - Restart the ECS.
 - Reinstall the ECS OS. Back up important data before the reinstallation.
 - If the ECS OS cannot be reinstalled due to important data, replace the disk attached to the ECS. To do so, back up data on the original disk, detach the disk from the ECS, attach the new disk to the ECS, and copy data to the new disk.

You can also upgrade the vCPUs and memory by [modifying ECS specifications](#).

- If the login fails because the bandwidth exceeds the limit, perform the following operations:

Check whether the bandwidth exceeds the configured bandwidth size. For details, see [How Do I Know If My EIP Bandwidth Limit Has Been Exceeded?](#)

If the bandwidth exceeds the limit, increase the bandwidth. For details, see [Changing an EIP Bandwidth](#).

NOTE

If network jitter or packet loss occurs frequently, dynamic BGP may be used in cross-border access. In this case, you are advised to use premium BGP.

For details, see [Why Is There Network Jitter or Packet Loss During Cross-Border Communications?](#)

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether an ECS Has an EIP Bound

If you need to use a remote login tool (such as PuTTY or Xshell) to access the ECS, bind an EIP to the ECS.

For details, see [Assigning an EIP](#).

Checking Whether the Network Is Normal

Use a local PC in another network or use another hotspot to access the ECS. Check whether the fault occurs on the local network. If so, contact the carrier to resolve this issue.

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Security Group Is Correctly Configured

Check whether the local host can access port 22 on the ECS.

Run the following command to check whether port 22 is accessible:

```
telnet ECS private IP address
```

If port 22 is inaccessible, check whether port 22 is opened in the security group rule.

On the ECS details page, click the **Security Groups** tab and check that port 22 is configured in the inbound rule of the security group.

For details about how to modify a security group rule, see [Modifying a Security Group Rule](#).

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Remote Access Port Is Correctly Configured

Check ECS settings.

1. Check whether the sshd process is running.
2. Check whether your local PC is denied by the ECS.
 - a. Log in to the ECS and run the following command:
vi /etc/hosts.deny
 - b. If the IP address of the local PC is in the **hosts.deny** file, the ECS denies connection attempts from the local PC. In such a case, delete the IP address from the file.
3. Open the **/etc/ssh/ssh_config** file in the local PC and view the default login port. Then, open the **/etc/ssh/sshd_config** file in the ECS and check whether the SSH port is the default port 22.

```
# semanage port -a -t ssh_port_t -p tcp #PORTNUMBER
#
#Port 22
#AddressFamily any
```

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking the IP Address Whitelist for SSH Logins (with HSS Enabled)

After HSS is enabled, you can configure an IP address whitelist for SSH logins as required. The IP address whitelist controls SSH access to ECSs, effectively preventing account cracking.

After you configure the allowlist, SSH logins will be allowed only from IP addresses in the allowlist.

1. On the **Events** page, check whether a local host IP address is intercepted due to brute force cracking.
2. Check whether the IP address whitelist for SSH logins has been enabled. If it has been enabled, ensure that the IP address of the local host has been added to the IP address whitelist.

CAUTION

- Before enabling this function, ensure that all IP addresses that need to initiate SSH logins are added to the allowlist. Otherwise, you cannot remotely log in to your ECS through SSH.
- Exercise caution when adding a local IP address to the allowlist. This will make HSS no longer restrict access from this IP address to your ECSs.

For more details, see [Configuring Server Login Protection](#).

Checking Whether an OS Fault Has Occurred

- Password injection failure
The password failed to be injected using Cloud-Init.
- File system damaged after a forcible stop
There is a low probability that the file system is damaged after a forcible stop, which causes the ECS fails to be restarted. For details, see [Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?](#)

After you perform the preceding operations, try to remotely log in to the ECS again.

Checking Whether the Access Is Blocked by Antivirus Software

Third-party antivirus software may lead to a failure in accessing the ECS.

If third-party antivirus software is running, check whether the remote connection is blocked by the software. If the remote connection is blocked, add the EIP bound to the ECS to the whitelist of the antivirus software and try to access the ECS again.

You can also disable or uninstall the third-party antivirus software and try to remotely log in to the ECS again.

Checking Whether an Error Occurred During a Remote Login

If an error message is displayed during remote login, check the operation guide based on the error information.

For details, see [Remote Login Errors on Linux](#).

If the fault persists, record the resource details and fault occurred time, and contact technical support for assistance.

If the fault persists after the preceding operations are performed, 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.

5.2.2 How Can I Change a Remote Login Port?

Scenarios

This section describes how to change a port for remote logins.

Linux

The following procedure uses an ECS running CentOS 7.3 as an example. The default login port of a Linux ECS is 22. To change it to port 2020, for example, do as follows:



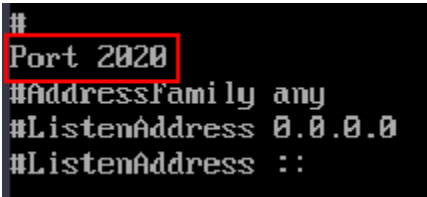
1. Modify the security group rule.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select your region and project.
 - c. Click . Under **Compute**, click **Elastic Cloud Server**.
 - d. On the ECS list, click the name of an ECS for which you want to modify the security group rule.
 - e. On the ECS details page, click the security group in the **Security Groups** area to go to the security group details page.
 - f. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set **Protocol & Port** as follows:
 - **Protocols:** TCP (Custom ports)
 - **Port:** 2020
2. Log in to the ECS.
3. Run the following command to edit the sshd configuration file:
vi /etc/ssh/sshd_config
4. Delete the comment tag (#) from the **#port 22** line and change **22** to **2020**.

Figure 5-1 Changing the port number to 2020



```
#  
Port 2020  
#AddressFamily any  
#ListenAddress 0.0.0.0  
#ListenAddress ::
```

5. Press **Esc** to exit Insert mode and enter **:wq!** to save the settings and exit.
6. Run either of the following commands to restart sshd:
service sshd restart
Or
systemctl restart sshd
7. Skip this step if the firewall is disabled. Configure the firewall.
The firewall varies depending on the CentOS version. CentOS 7 uses firewalld, and CentOS 6 uses iptables. The following operations use CentOS 7 as an example.
Run the **firewall-cmd --state** command to check the firewall status.
 - (Recommended) Method 1: Add information about a new port to firewalld.
 - i. Run the following commands to add a rule for port 2020:
firewall-cmd --zone=public --add-port=2020/tcp --permanent
firewall-cmd --reload
 - ii. View the added port. The TCP connection of port 2020 will have been added.
firewall-cmd --list-all
 - iii. Restart firewalld.
systemctl restart firewalld.service
 - Method 2: Disable the firewall and the function of automatically enabling the firewall upon ECS startup.
systemctl stop firewalld
systemctl disable firewalld
8. Run the following command to check whether the port is open:
telnet *EIP port*
For example: **telnet xx.xx.xx.xx 2020**

5.2.3 Why Cannot I Use a Non-Default SSH Port to Log In to My Linux ECS?

Symptom

After changing the default SSH port, you could not use the new port to log in to the ECS.

Possible Causes

- The access to the new port is not allowed in the security group.
- The new port is not enabled on the firewall.
- The new port is not added to the SSH configuration file.
- The hosts configuration file is incorrectly configured.

Checking Security Group Rules

Check whether the security group is correctly configured.

For example, if the new SSH port number is 2020, ensure that there is a security group rule without restriction in the outbound direction and allowing access to this port in the inbound direction.

Checking Firewall Rules

Run the **iptables** command to check whether the new SSH port, for example, port 2020 is enabled on the firewall.

1. Log in to the Linux ECS.
2. Take CentOS 7.5 as an example. Run the following command to edit the iptables file:

```
vi /etc/sysconfig/iptables
```

3. Add a rule for port 2020.

```
-A INPUT -m state --state NEW -m tcp -p tcp --dport 2020 -j ACCEPT
```

4. Restart iptables.

```
systemctl restart iptables
```

Checking the SSH Configuration File

Log in to the ECS and check the SSH configuration file.

1. Run the following command to check whether port 2020 has been configured:

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

2. If the port has not been configured, replace **#Port 22** with **Port 2020**.
3. Run the following command to restart SSH:

```
service sshd restart
```

Checking the hosts Configuration File

The **/etc/hosts.allow** and **/etc/hosts.deny** files of a Linux ECS are used to permit or deny an IP address or an IP address segment, respectively, to remotely access the ECS using SSH.

1. Add the following statement to **/etc/hosts.allow** to allow the IP address 192.168.1.3 to access the ECS using SSH:

```
sshd: 192.168.1.3
```

2. Check **/etc/hosts.deny**. If **sshd:all:deny** is contained, comment it out.

NOTE

If a rule is set in both **hosts.allow** and **hosts.deny**, the rule in **hosts.allow** takes precedence. For example, if "sshd: 192.168.1.3" is set in **hosts.allow** and "sshd:all:deny" is set in **hosts.deny**, the ECS allows only the SSH login from IP address 192.168.1.3.

5.2.4 What Browser Version Is Required to Remotely Log In to an ECS?

When you use a browser to remotely log in to an ECS, ensure that the browser version meets the requirements listed in [Table 5-2](#).

Table 5-2 Browser version requirements

Browser	Version
Google Chrome	31.0-75.0
Mozilla Firefox	27.0-62.0
Internet Explorer	10.0-11.0

5.2.5 How Can I Log In to an ECS After It Exchanged the System Disk with Another ECS Running the Same OS?

Symptom

Two ECSs run the same OS, for example, both run Linux. The system disks attached to the two stopped ECSs are exchanged. After the exchange, the passwords or keys used to log in to the ECSs may change. In this case, how do I log in to the ECS whose system disk has been replaced?

Linux

Login methods vary according to the login authentication used on the ECSs. Assume that there are three Linux ECSs and they are configured as shown in [Table 5-3](#).

Table 5-3 ECS configurations

ECS	System Disk	Login Authentication	Password/Key Pair
ecs_01	vol_01	Password or key pair	If a password is used for login authentication, take Ecs@01 as an example. If a key pair is used for login authentication, take private key file Keypair_01 as an example.
ecs_02	vol_02	Password	Ecs@02
ecs_03	vol_03	Key pair	Keypair_03

- Scenario 1: System disk vol_01 is detached from ecs_01 offline and then attached to ecs_02 as the system disk. How can I log in to ecs_02?
Use either of the following methods to log in to ecs_02:
 - Use private key file **Keypair_01** (if available) of ecs_01.
 - Use the original password **Ecs@02** of ecs_02.
- Scenario 2: System disk vol_01 is detached from ecs_01 offline and then attached to ecs_03 as the system disk. How can I log in to ecs_03?
Use one of the following methods to log in to ecs_03:
 - Use the password **Ecs@01** (if available) of ecs_01.
 - Use private key file **Keypair_01** (if available) of ecs_01.
 - Use private key file **Keypair_03** of ecs_03.


5.3 VNC Login

5.3.1 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 5-2](#).

Figure 5-2 Garbled characters on the VNC-based login page



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.

5.3.2 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.

5.3.3 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.

5.3.4 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 5-3](#).

Figure 5-3 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.

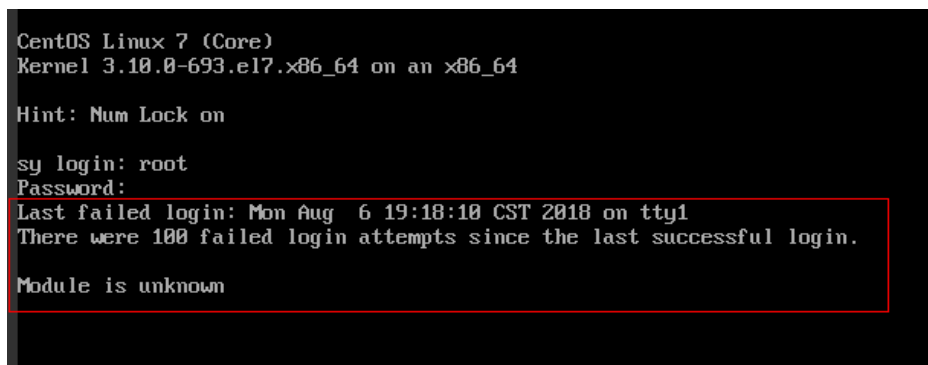
5.4 Remote Login Errors on Linux

5.4.1 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 5-4 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

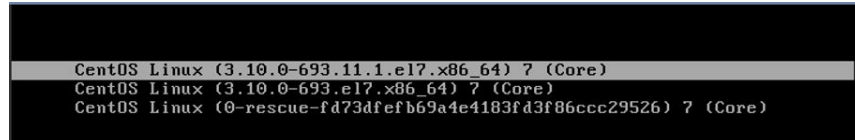
The file in the `/etc/pam.d/` directory was modified by mistake.

Solution

1. Enter the 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.

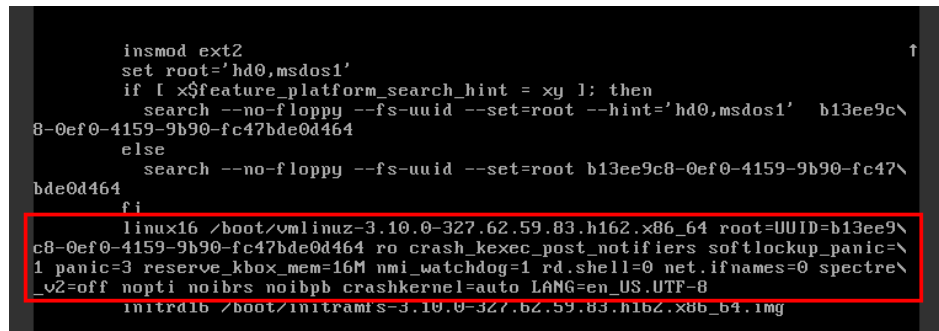
Figure 5-5 Entering the kernel editing mode

```
CentOS Linux (3.10.0-693.11.1.el7.x86_64) 7 (Core)
CentOS Linux (3.10.0-693.el7.x86_64) 7 (Core)
CentOS Linux (0-rescue-fd73dfebf69a4e4183fd3f86ccc29526) 7 (Core)
```

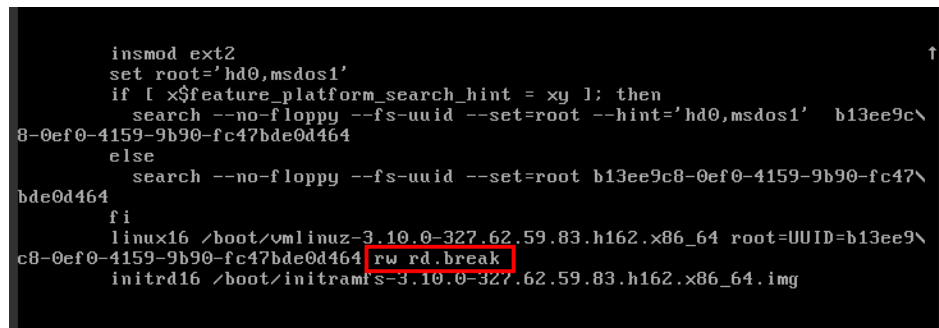
NOTE

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 5-6 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' b13ee9c8-0ef0-4159-9b90-fc47bde0d464
else
  search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9c8-0ef0-4159-9b90-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 5-7 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' b13ee9c8-0ef0-4159-9b90-fc47bde0d464
else
  search --no-floppy --fs-uuid --set=root b13ee9c8-0ef0-4159-9b90-fc47bde0d464
fi
linux16 /boot/vmlinuz-3.10.0-327.62.59.83.h162.x86_64 root=UUID=b13ee9c8-0ef0-4159-9b90-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 5-8 System log

```
Aug 6 18:08:09 sy login: pam_succeed_if(login:auth): requirement "uid >= 1000" not met by user "root"
Aug 6 18:08:11 sy login: FAILED LOGIN 1 FROM tty1 FOR root, Authentication failure
Aug 6 18:08:15 sy login: pam_unix(login:session): session opened for user root by LOGIN(uid=0)
Aug 6 18:08:15 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 object 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 5-9 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

5.4.2 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 Message "Permission denied".

Figure 5-10 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 the 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 5-11 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 5-12 Before the modification

Figure 5-13 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.

5.4.3 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 5-14 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.

5.4.4 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

- If the username or password is incorrect, Check the username and password.
- If 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.

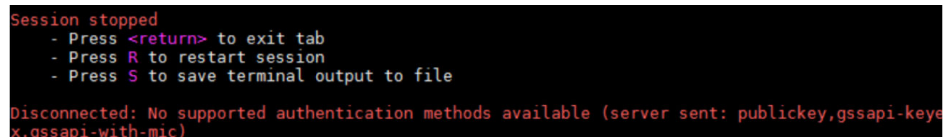
- CentOS 6
service sshd restart
- CentOS 7
systemctl restart sshd

5.4.5 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 5-15 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

6 Billing FAQ

6.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing Modes?

Yearly/Monthly

Yearly/Monthly is a prepaid billing mode and is cost-effective for long-term use.

Note the following when using a yearly/monthly ECS:

1. A created yearly/monthly ECS cannot be deleted. If such an ECS is not required any more, unsubscribe it. To do so, switch to the **Elastic Cloud Server** page, locate the target ECS, and choose **More > Unsubscribe** in the **Operation** column.
2. A detached system disk can be used as a data disk for any ECSs, but can only be used as a system disk for the ECS where it was attached before.
3. A detached data disk that is purchased together with an ECS can only be used as a data disk for this ECS.

Pay-per-Use

Pay-per-use billing is a postpaid billing mode in which an ECS will be billed based on usage frequency and duration. ECSs are billed by second. The system generates a bill every hour based on the usage duration and deducts the billed amount from the account balance. A pay-per-use ECS can be provisioned and deleted at any time.

In the pay-per-use billing mode, ECSs are billed by the second. The price per second of each type of ECS can be obtained by dividing their hourly price by 3600. Obtain the hourly price on the **Product Pricing Details** page.

For example, if you purchase a pay-per-use ECS priced \$0.68 USD/hour, the ECS will be billed based on the usage duration by the second.

- If you use the ECS for 30 minutes, you need to pay for \$0.34 USD ($0.68/3600 \times 30 \times 60$).

- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$1.02 USD (0.68/3600 x 90 x 60).

 **NOTE**

If a pay-per-use ECS is stopped and then restarted, the startup may fail due to insufficient resources. In such a case, change the ECS flavor or wait several minutes before attempting another restart.

Which One Is More Cost-Effective?

The yearly/monthly payment is more cost-effective than the pay-per-use payment for a longer usage duration. Yearly/Monthly is ideal when the duration of ECS usage is predictable. Pay-per-use is recommended when you want more flexibility and control on ECS usage. Pay-per-use ECSs can be provisioned or deleted at any time.

6.2 Will I Be Billed After ECSs Are Stopped?

ECSs support multiple billing modes. Whether you will continue to be billed after ECSs are stopped depend on the billing mode of the stopped ECSs.

- **Yearly/Monthly:** a prepaid billing mode. You are billed based on the subscription term you purchase. Stopping ECSs does not affect the billing.
- **Spot pricing (for spot block):** a postpaid billing mode. You are billed based on the usage duration and predefined duration. After ECSs are stopped, you will continue to be billed.
- **Pay-per-use and spot pricing (for spot instances):** a postpaid billing mode. You are billed based on the usage duration. Whether you are billed after ECSs are stopped depends on the instance type. For details about the billing policy for stopped ECSs, see [Table 6-1](#).

Table 6-1 Billing for a stopped ECS

Instance Type	Description	Billed Items	Billed or Not	Action on Resources
Common instances	Common instances include: <ul style="list-style-type: none"> • Non-bare-metal instance • Instances without local disks • Instances without FPGA cards 	ECS (compute resources including vCPUs and memory)	No	<ul style="list-style-type: none"> • If ECSs are created from public resource pools, basic resources (vCPUs and memory) are no longer retained. The ECS may fail to be started due to insufficient resources. Wait patiently and try again later. • If ECSs are created on a DeH or in an edge AZ, basic resources (vCPUs and memory) will be retained.
		Images	No	Resources are retained but are not billed.
		GPU	No	After GPU-accelerated ECSs without local disks attached are stopped, GPU resources are no longer retained.
		EVS disks (system disks and data disks)	Yes	The resources are still billed based on the billing rules.
		EIP bandwidth	Yes	The resources are still billed based on the billing rules. <ul style="list-style-type: none"> • EIP bandwidth price: pay-per-use EIPs (by bandwidth) • Shared bandwidth price For details, see EIP Billing .

Instance Type	Description	Billed Items	Billed or Not	Action on Resources
Special instances	Special instances include: <ul style="list-style-type: none">Bare metal instancesInstances with local disks such as disk-intensive, ultra-high I/O ECSs.ECSs with FPGA cards	ECS (compute resources including vCPUs and memory)	Yes	The ECS resources are still billed based on the billing rules. To stop the billing, you need to delete the instance and its associated resources.
		Images	Yes	
		GPU	Yes	
		EVS disks (system disks and data disks)	Yes	
		EIP bandwidth	Yes	

NOTE

For a stopped pay-per-use ECS, the startup may fail due to insufficient resources. Please wait for several minutes before attempting another restart or changing the ECS specifications.

For details, see [How Can I Stop an ECS from Being Billed?](#)

If you want to use an ECS for a long time, you can change its billing mode from pay-per-use to yearly/monthly to reduce cost. For details, see [Changing Pay-per-Use to Yearly/Monthly](#).

6.3 Can I Switch Between Yearly/Monthly and Pay-per-Use Billing Modes?

Yes, you can switch between yearly/monthly and pay-per-use payments.

- Changing pay-per-use to yearly/monthly
Pay-per-use is a postpaid billing mode in which your ECS is billed by usage duration. You can create or delete such an ECS at any time.
If you want to use an ECS for a long time, you can change its billing mode from pay-per-use to yearly/monthly to reduce cost. For details, see [Changing Pay-per-Use to Yearly/Monthly](#).
- Changing yearly/monthly to pay-per-use
Yearly/Monthly is a prepaid billing mode in which your ECS will be billed based on service duration. This cost-effective mode is ideal when the duration of ECS usage is predictable.
If you require a more flexible billing mode, in which your ECS will be billed based on the actual usage, you can change the billing mode from yearly/

monthly to pay-per-use. This billing mode change (from yearly/monthly to pay-per-use) takes effect immediately. For details, see [Changing Yearly/Monthly to Pay-per-Use](#).

NOTE

- You have passed real-name authentication.
- You can change the billing mode from yearly/monthly to pay-per-use only for ECSs whose status is **Provisioned** on the **Renewals** page.
- The billing modes of products in a solution portfolio cannot be changed from yearly/monthly to pay-per-use.

6.4 FAQs About ECS Frozen, Deletion, and Unsubscription

Why Are My ECSs Released?

After you purchase ECSs on Huawei Cloud but you do not complete the payment or renewal, the purchased ECSs will enter a grace period. If you still do not complete the payment or renewal after the grace period is ended, your ECS will enter a retention period. The ECS cannot provide services during the retention period. If you still do not complete the payment or renewal after the retention period is ended, your data stored in the ECS will be deleted and the ECS will be released. For details, see [Resource Suspension and Release](#).

Can I Back Up My Data on the ECS When It Is Frozen?

No. If your ECS is frozen due to arrears, you can back up data only after you top up your account.

How Do I Unfreeze a Frozen ECS?

- Frozen due to arrears: You can renew or top up your account to unfreeze your ECS. You can renew or delete the ECSs that are frozen due to arrears. Only yearly/monthly ECSs that have not expired can be unsubscribed.
- Frozen due to violation detected by HUAWEI CLOUD: You can renew or delete such ECSs. Only yearly/monthly ECSs that have not expired can be unsubscribed.
- Frozen due to violation detected by the public security department: You can renew such ECSs, but cannot delete them. Such frozen ECSs cannot be unsubscribed although they are displayed on the unsubscription page.

What Is the Impact on Services When Resources Are Frozen, Unfrozen, or Released?

- When resources are frozen, resource access and usage are restricted, which will interrupt your services. For example, if a server is frozen, it will be automatically stopped.
- When resources are unfrozen, restrictions on resources are removed, but you need to check and restore your services. For example, after a yearly/monthly ECS is unfrozen, it will be automatically started. After a pay-per-use ECS is unfrozen, you need to start it manually.

- When resources are released, data stored on the resources will be deleted and cannot be retrieved.

How Do I Renew an ECS?

After an ECS billed on a yearly/monthly basis expires, renew it on the **Renewals** page of the management console. For details, see [Renewal Management](#).

How Do I Restore an Unsubscribed or Deleted ECS?

After an unsubscription is complete, the ECS will be permanently deleted and cannot be restored. You are advised to purchase a new ECS if you still want to use ECSs.

How Do I Delete an ECS?

- Pay-per-use ECS: On the ECS list page, select the target ECS, click **More** in the **Operation** column, and choose **Delete**.
 - Read details about deleting ECSs carefully.
 - Choose to delete the unneeded EIP and attached data disks together with the ECS to avoid generating costs.
- Yearly/Monthly ECS: On the ECS list page, select the target ECS, click **More** in the **Operation** column, and choose **Unsubscribe**.
For details about unsubscription rules and procedure, see [How Do I Unsubscribe from ECSs?](#)

How Do I Restore a Released ECS or EVS Disk?

Data cannot be restored if an ECS or EVS disk was not backed up before it is released.

For details about how to back up an ECS and restore data using a backup, see [Backing Up ECS Data](#).

How Do I Configure CBR and HSS for My ECS?

You can choose to purchase CBR and HSS together with your ECSs during the purchase process. They are not configured by default.

You can view the backup policy on the CBR console and associate the policy with your ECS. HSS takes effect only after the agent is installed on an ECS.

6.5 How Can I Stop an ECS from Being Billed?

- After a pay-per-use or spot ECS is stopped (excluding bare metal ECSs or ECSs with local disks or FPGAs attached), its basic resources (vCPUs, memory, image, and GPUs) will no longer be billed, but its associated resources such as EVS disks (system and data disks), EIPs, and bandwidth will continue to be billed separately.
To stop the ECS from being billed, delete it and its associated resources.
- Spot block ECSs, pay-per-use or spot ECSs with local disks attached (such as disk-intensive, ultra-high I/O, H2, P1, and P2 ECSs), pay-per-use or spot FPGA-

based ECSs (such as Fp1 and Fp1c ECSs), and pay-per-use or spot bare metal ECSs will continue to be billed after they are stopped. To prevent such ECSs from being billed, delete them and their associated resources.

- For yearly/monthly resources such as yearly/monthly ECSs or EVS disks, your pay for them when you are purchasing them. The billing automatically stops when the subscription expires. If you stop using the resource before the subscription expires, you will not be eligible for a refund.

This section uses a pay-per-use ECS as an example to describe how you are billed after the ECS is deleted. [Table 6-2](#) lists the resources associated with the ECS.

Table 6-2 Billing example of a pay-per-use ECS

Resources	Description	Billing Mode
ECS basic resources	vCPUs, memory, image, and GPUs	Pay-per-use
EVS disks	System disk	Pay-per-use
	Data disk	Pay-per-use
EIP	N/A	Pay-per-use

After the ECS is deleted, it is billed as follows:

- ECS basic resources: no longer billed
- EVS disks
 - System disk: no longer billed
 - Data disks: no longer billed if you have selected **Delete the data disks attached to the following ECSs** when you were deleting the ECS. Otherwise, the data disks will continue to be billed.
- EIP: If you select **Release the EIPs bound to the ECSs** when deleting the ECS, the EIP will no longer be billed. Otherwise, the EIP will continue to be billed.

6.6 FAQs About Spot ECSs

About Spot ECSs

1. Why is my spot ECS released even when I have sufficient account balance?
A spot ECS may be released at any time based on the changes in market price or supply and demand. For example, if the market price at a certain time is higher than the maximum price you are willing to pay, or if there are a large number of demands that the ECS resource supply cannot meet, the system automatically reclaims your spot ECS.
2. Can I change a spot ECS to a pay-per-use or yearly/monthly ECS?
No.
3. Which resources are included in the price discount of a spot ECS?

The price discount applies only to the vCPUs and memory of a spot ECS. The prices of other resources, such as the system disk, data disk, and bandwidth, are the same as those of ECSs billed on a pay-per-use basis.

4. How can I bid for a spot ECS?

When you purchase a spot ECS, you are required to set the maximum price you are willing to pay. If the maximum price is higher than the market price and inventory resources are sufficient, you can purchase your spot ECS. The spot ECS is billed depending on the market price.

5. What is the relationship between the maximum price I am willing to pay for a spot ECS and the market price?

Your spot ECS will run only if the maximum price you are willing to pay is higher than the market price. If the maximum price is lower than the market price, purchasing the spot will ECS fail, or the spot ECS that you have already purchased will be reclaimed. A spot ECS is billed based on market price, regardless of the maximum price you set.

6. If I have multiple spot ECSs and all the ECSs start to run at the same time, will the billing be the same for all of them?

The billing will be the same for spot ECSs of the same series with the same specifications.

7. Can I obtain the market price before purchasing the spot ECS?

Yes. When you purchase a spot ECS on the management console, you can view the market price range and historical prices of the ECS after you select an ECS flavor.

8. How is a spot ECS billed?

Spot ECSs are billed by the second, and the billing period is 1 hour.

9. When does the billing duration for a spot ECS start and end?

The duration starts from the time when the spot ECS is purchased to the time when it is released either manually or automatically.

10. Will a stopped spot ECS continue to be billed?

For details, see [Table 6-3](#).

Table 6-3 Billing of stopped spot ECSs

Item	Spot Pricing
Spot ECSs without local disks or FPGAs attached, or non-bare metal ECSs	After the ECS is stopped, basic resources including vCPUs, memory, image, and GPUs are no longer billed. The resources associated with the ECS, such as EVS disks (including system and data disks), EIPs, and bandwidth, are separately billed.
Spot ECSs with local disks attached, FPGA-based ECSs, or bare metal ECSs	The ECS will continue to be billed after it is stopped. To stop the ECS from being billed, delete it and its associated resources.

6.7 How Can I Renew ECSs?

Yearly/monthly ECSs are unavailable after the subscription expires. If you want to continue using them, you need to renew the ECS subscription within a specified period. Otherwise, resources, such as vCPUs, memory, and EVS disks, will be automatically released, and data will be lost and cannot be restored.

On the ECS console, locate the ECS you want to renew and choose **More > Renew** in the **Operation** column. You can renew the subscription manually or automatically.

For details, see [Overview](#).

6.8 How Can I Automatically Renew a Yearly/Monthly ECS?

Solution

1. Select **Auto renew** when purchasing a yearly/monthly ECS.
On the **Buy ECS** page, select **Auto renew** under **Required Duration**.
2. Select the EIP bound to the target ECS for automatic renewal.
Auto renew is provided for the ECSs and EVS disks when you purchase them. To enable automatic renewal for EIPs, perform the following operations:
 - a. Log in to the management console and click **Billing** in the upper right corner.
The Billing Center console is displayed.
 - b. In the left navigation pane, choose **Renewals**.
 - c. Select the EIP for automatic renewal.

6.9 Will I Receive a Notification If My Account Balance Is Insufficient?

You can set the balance alert function in the Billing Center. The system checks your account balance and sends a notification if your balance is less than or equal to your set threshold. Set the alert threshold based on your resource usage.

1. On the **Overview** page, in the **Available Credit** area, turn on the **Alert** switch to enable the balance alert function. Click **Modify** and you can set a desired threshold.
2. With balance alert enabled, when the sum of your account balance, cash coupons, and flexi-purchase coupons goes below the threshold, the recipients will receive a notification for 1-3 days by SMS and email.

You can modify the recipients that receive the balance alerts at **SMS & Email Settings > Finance > Account balance** in the Message Center.

After receiving a balance alert, top up your account or disable unnecessary resources in a timely manner to avoid affecting the normal use of cloud resources or to stop unnecessary fees from being generated.

7 Region and AZ FAQ

7.1 What Is AZ and How Can I Select and View an AZ?

What Is an AZ?

An availability zone (AZ) is a physical region where resources use independent power supplies and networks. AZs are physically isolated but interconnected through an internal network.

There are multiple AZs in each region. If one AZ becomes faulty, other AZs in the same region continue to provide services.

AZs in the same region can communicate with each other through an internal network.

How Do I Select an AZ

You can select an AZ when you are purchasing an ECS. After the ECS is created, the AZ cannot be changed. If there is only one AZ displayed in a region, it means the region only provides one AZ.

7.2 What Is a Region?

Regions are geographic areas isolated from each other. ECSs are region-specific and cannot be used across regions through internal network connections.

When you create an ECS, select the nearest region for low network latency and quick resource access.

7.3 Are Products Different in Different Regions?

Yes. Currently, each region contains different products. Certain products are available for trial release in certain regions only.

7.4 Is Data Transmission Between AZs Billed?

Data transmission between AZs in the same region is free of charge. However, data transmission between AZs in different regions will be billed.

7.5 Can I Change the Region or AZ After an ECS Is Purchased?

Sorry, you cannot change the region after the ECS is purchased. During the ECS purchase, you are advised to select the region nearest to your services for lower network latency and quick resource access.

If you need to change the region or AZ for a purchased ECS, you can use Image Management Service (IMS) to migrate the ECS data across regions or AZs.

Example 1:

ecs01 in a region needs to be migrated to another region.

1. In the original region, create a full-ECS image image01 for ecs01.
2. Replicate image01 from the original region A to the target region and name the new image image01-copy.
3. In the target region, use image01-copy to create an ECS named ecs02.
In this way, data on ecs01 is migrated from the original region to ecs02 in the target region.

Example 2:

In a region, ecs01 in an AZ needs to be migrated to another AZ.

1. In the original AZ, create full-ECS image image01 from ecs01.
2. In the target AZ, use image image01 to create an ECS named ecs02.
In this way, data on ECS ecs01 is migrated from the original AZ to the target AZ.

For more migration methods and background information, see [Can I Migrate an ECS to Another Region, AZ, or Account?](#)

7.6 Can I Migrate an ECS to Another Region, AZ, or Account?

After an ECS is created, it cannot be directly migrated to another region, AZ, or account.

7.7 Can a Load Balancer Distribute Traffic to ECSs in Different Regions?

Only dedicated load balancers support this. Backend servers can be from VPCs in different regions.

For details, see [Creating a Dedicated Load Balancer](#).

7.8 Can Components Contained in an Application Be Distributed to Different Regions?

Yes. However, such a deployment mode is not recommended.

You are advised to deploy the components contained in an application in the same region. In this manner, these components can communicate with each other over an internal network, reducing bandwidth costs of using public networks and ensuring communication quality between the components.

8 OS FAQ

8.1 Do ECSs Support GUI?

Linux ECSs are managed through the CLI. You can configure a GUI if required.

Before installing a GUI on an ECS, ensure that the memory is no less than 2 GB to prevent GUI installation or ECS startup failures.

8.2 How Can I Install a GUI on an ECS Running CentOS 6?

Scenarios

To provide a pure system, the ECSs running CentOS 6 do not have a GUI installed by default. You can install a GUI on such ECSs as needed.

Constraints

- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

Procedure

1. Run the following command to obtain the installation component provided by the OS:
yum groupinstall "Desktop"
2. Run the following command to set the default startup level to 5 (GUI):
sed -i 's/id:3:initdefault:/id:5:initdefault:/' /etc/inittab
3. Run the following command:
startx

8.3 How Can I Install a GUI on an ECS Running CentOS 7?

Scenarios

You want to install a GUI on an ECS running CentOS 7 series.

Constraints

- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.

Procedure

1. Run the following command to install the GUI desktop component:
yum groupinstall "Server with GUI"

NOTE

If the following message is displayed after the installation is complete:

```
Failed : python -urllib3.noarch 0:1.10.2-7.e17
```

Run the following command:

```
mv /usr/lib/python2.7/site-packages/urllib3/packages/  
ssl_match_hostname /usr/lib/python2.7/site-packages/urllib3/packages/  
ssl_match_hostname.bak
```

```
yum install python-urllib3 -y
```

2. After the installation is complete, run the following command to set the default startup level to **graphical.target**:
systemctl set-default graphical.target
3. Run the following command to start **graphical.target**:
systemctl start graphical.target
4. Restart the ECS.
5. Log in to the ECS using VNC provided on the management console. Set the language, time zone, username, and password as prompted.

8.4 How Can I Install a GUI on an ECS Running Ubuntu?

Scenarios

To provide a pure system, the ECSs running Ubuntu do not have a GUI installed by default. You can install a GUI on such ECSs as needed.

For GPU-accelerated ECSs, after installing a GUI, you need to configure X Server, x11vnc, and lightdm to make sure that:

- The graphics system and VNC server are automatically started upon the ECS startup.
- Applications can invoke GPUs properly after a remote login using VNC.

You can perform the following steps to install a GUI on an Ubuntu ECS:

- [Installing a GUI](#)
- [\(Optional\) Configuring X Server, x11vnc, and lightdm](#): required only for GPU-accelerated ECSs.
- [\(Optional\) Verifying Drivers on GPU-accelerated ECSs](#): required only for GPU-accelerated ECSs.

Constraints

- This document applies to ECSs running Ubuntu 16.04, 18.04, and 20.04.
- The Ubuntu ECS must have an EIP bound or have an intranet image source configured.
- Before installing a GUI on an ECS, ensure that the idle memory is greater than or equal to 2 GB. Otherwise, the GUI installation may fail or the ECS cannot be started after the installation.
- GPU-accelerated ECSs must have a correct GPU driver installed. For details, see [GPU Driver](#).

Installing a GUI

1. Log in to the ECS and install a GUI desktop environment.
 - a. Run the following command to update the software library:
apt-get update
 - b. Run the following command to install the Ubuntu GUI desktop component:

- For Ubuntu 16.04, run the following command:

apt-get install -y scite xorg xubuntu-desktop

- For Ubuntu 18.04 and 20.04, run the following command:

apt-get install -y ubuntu-desktop

2. Run the following command to edit the **root/.profile** file:

vim /root/.profile

Press **i** to enter the editing mode and change **mesg n || true** in the last line to **tty -s && mesg n || true**. After the modification, the file content is as follows:

```
# ~/.profile: executed by Bourne-compatible login shells.
```

```
if [ "$BASH" ]; then
  if [ -f ~/.bashrc ]; then
    . ~/.bashrc
  fi
fi
tty -s && mesg n || true
```

3. Press **Esc** to exit editing mode.
4. Run the following command to save and exit the configuration file:

:wq

5. (Mandatory for Ubuntu 20.04) Add a member account.

After GUI desktop component is installed on the ECS, you cannot log in to the Ubuntu 20.04 OS as user **root** **user**. You need to add a member account for logging in to the GUI desktop.

Run the following command to add user **user01**:

```
adduser user01
```

Set a password for **user01** as prompted.

```
Adding user `user01' ...
Adding new group `user01' (1001) ...
Adding new user `user01' (1001) with group `user01' ...
Creating home directory `/home/user01' ...
Copying files from `/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
```

Set information about **user01**. You can press **Enter** to skip the setting. Then the system prompts you to check whether the entered information is correct.

Enter **Y**.

```
Changing the user information for user01
Enter the new value, or press ENTER for the default
  Full Name []:
  Room Number []:
  Work Phone []:
  Home Phone []:
  Other []:
Is the information correct? [Y/n] Y
```

6. Run the reboot command to restart the ECS.
7. Log in to the ECS using VNC provided on the management console and log in to the GUI desktop using the member account created in [5](#) or the **root** account.
 - For Ubuntu 20.04 OS, you need to use the member account to log in to the GUI desktop.
 - For GPU-accelerated ECSs, you also need to [configure X Server, x11vnc, and lightdm](#).

(Optional) Configuring X Server, x11vnc, and lightdm

For GPU-accelerated ECSs, you need to configure X Server, x11vnc, and lightdm when installing a GUI.

1. Remotely log in to the ECS.
2. Query the BusID of the GPU.

```
lspci | grep -i nvidia
```

Figure 8-1 GPU's BusID

```
00:0d.0 3D controller: NVIDIA Corporation GV100GL [Tesla V100 PCIe 32GB] (rev a1)
```

3. Generate the X Server configuration.

```
nvidia-xconfig --enable-all-gpus --separate-x-screens
```
4. Configure the GPU's BusID in "Section Device" in the generated `/etc/X11/xorg.conf`.

- a. Edit `/etc/X11/xorg.conf`.
vi /etc/X11/xorg.conf
- b. Press **i** to enter editing mode.
- c. Add the GPU's BusID in "Section "Device".


Figure 8-2 Adding the GPU's BusID

```
Section "Device"
  Identifier      "Device0"
  Driver          "nvidia"
  VendorName     "NVIDIA Corporation"
  BoardName      "Tesla U100-PCIE-32GB"
  BusID          "PCI:00:13:0"
EndSection
```

NOTE

The BusID queried in step 2 is a hexadecimal number. You need to convert it to a decimal number before adding it to "Section Device" in `/etc/X11/xorg.conf`.

1. For example, the queried BusID is `00.0d.0` (a hexadecimal number) and needs to be converted to `PCI:00:13:0` (a decimal number).
 - d. Press **Esc** to exit editing mode.
 - e. Run the following command to save and exit the configuration file:
:wq
5. Install `x11vnc`.
apt-get -y install x11vnc
 6. Install `lightdm`.
apt-get -y install lightdm
 7. Select **lightdm** as the default display manager.

Figure 8-3 Selecting a display manager

8. Configure the GUI desktop environment to automatically start upon ECS startup.
systemctl set-default graphical.target
9. (Optional) Configure the `x11vnc` to automatically start upon ECS startup.
 - a. Add the `/lib/systemd/system/myservice.service` file.
vi /lib/systemd/system/myservice.service

- b. Press **i** to enter editing mode.
 - c. Add the following content to the file:

```
[Unit]
Description=My Service
After=network.target lightdm.service

[Service]
Type=oneshot
ExecStart=/usr/bin/x11vnc -forever -loop -noxdamage -repeat -rfbport 5902 -shared -bg -auth guess -o /var/log/vnc.log

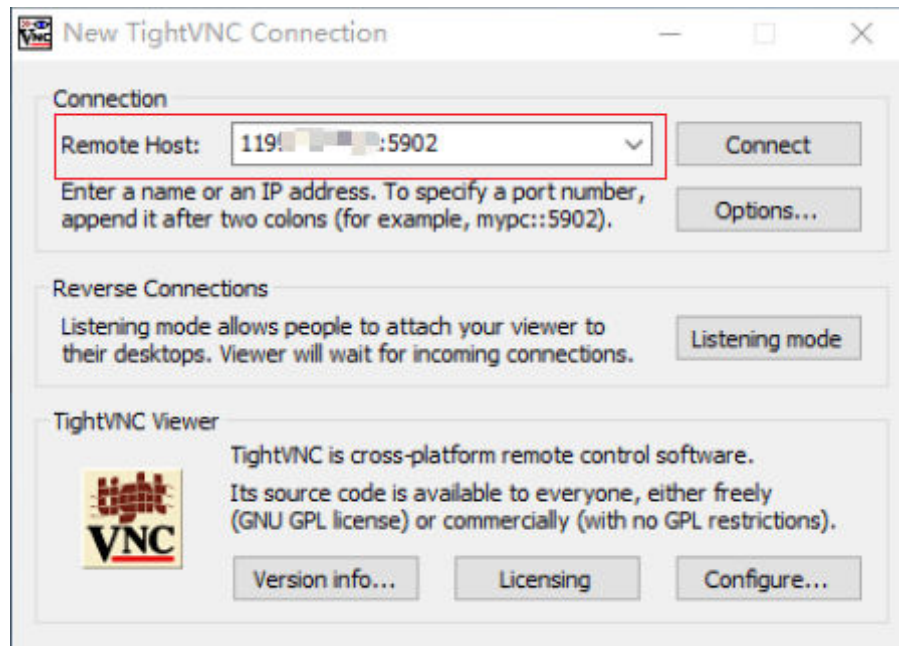
[Install]
WantedBy=multi-user.target
Alias=myservice.service
```
 - d. Press **Esc** to exit editing mode.
 - e. Run the following command to save and exit the configuration file:
:wq
10. Load configuration files.
systemctl daemon-reload
systemctl enable myservice.service
 11. Run the reboot command to restart the ECS.

(Optional) Verifying Drivers on GPU-accelerated ECSs

After installing a GUI on a GPU-accelerated ECS, perform the following operations to check whether the driver is working properly:

1. Log in to the management console.
2. Configure a security group for the ECS.
 - a. On the ECS list, click the name of an ECS for which you want to configure the security group rule. On the ECS details page, click **Security Groups**.
 - b. Expand the security group and in the upper right corner of the security group rule list, click **Modify Security Group Rule**.
 - c. On the **Inbound Rules** page, click **Add Rule**.
 - d. In the **Add Inbound Rule** dialog box, follow the prompts to add the following security group rule:
Allow inbound access through TCP port *5902*. The port number is determined by the **rfbport** parameter in step **9.c**.
3. Log in to the ECS using VNC.
The following uses TightVNC as an example.


Figure 8-4 TightVNC client



4. Right-click on the blank area and choose **Open in Terminal** from the shortcut menu.
5. Run the following command on the terminal. If the graphics card information is displayed as follows, the driver is working properly.

nvidia-settings

Figure 8-5 Graphics card information



NOTE

If a GPU-accelerated ECS has a GRID driver installed, you need to configure a license to use the GPU rendering capability. For details, see [Installing a GRID Driver on a GPU-accelerated ECS](#).

8.5 Why Does the OS Fail to Respond When kdump Occurs on a Linux ECS?

Symptom

When kdump occurs on a Xen Linux ECS, the OS fails to respond and cannot be automatically recovered. For example, if you run the `echo c >/proc/sysrq-trigger` command to trigger kdump, this fault occurs.

Figure 8-6 Triggering kdump

```
[root@ecs-xen01 linux1# 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 Wed 2018-01-17 06:15:35 UTC; 6min ago
   Process: 1397 ExecStart=/usr/bin/kdumpctl start (code=exited, status=0/SUCCESS)
   Main PID: 1397 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/kdump.service

Jan 17 06:15:05 ecs-xen01.novalocal systemd[1]: Starting Crash recovery kernel arming...
Jan 17 06:15:35 ecs-xen01.novalocal kdumpctl[1397]: kexec: loaded kdump kernel
Jan 17 06:15:35 ecs-xen01.novalocal kdumpctl[1397]: Starting kdump: [OK]
Jan 17 06:15:35 ecs-xen01.novalocal systemd[1]: Started Crash recovery kernel arming.
[root@ecs-xen01 linux1# echo c > /proc/sysrq-trigger
```

NOTE

Generally, kdump is disabled for public images. This issue does not occur on the ECSs created using public images.

Possible Causes

- Certain Linux kernel versions are incompatible with Xen virtualization.
- If kdump is enabled in the ECS with the kernel not supporting `soft_rest`, the ECS stops responding during dump.

Solution

Method 1: Disable kdump.

CentOS 7.5 is used as an example in the following.

1. Forcibly restart the ECS.
 - a. Log in to management console.
 - b. Under **Compute**, choose **Elastic Cloud Server**.
 - c. In the ECS list, select the target ECS and click **Restart**.

- d. Select **Forcibly restart the preceding ECSs** or **Forcibly stop the preceding ECSs**.
 - e. Click **OK**.
2. Disable kdump.
 - a. Log in to the forcibly restarted ECS as user **root**.
 - b. Run the following command to disable kdump:
service kdump stop

Method 2:

If the target ECS supports the **crash_kexec_post_notifiers** function, add the function to the ECS startup configuration file (**menu.lst** or **grub.cfg**). To do so, perform the following operations:

1. Run the following command to check whether the ECS supports the **crash_kexec_post_notifiers** function:
cat /proc/kallsyms |grep crash_kexec_post_notifiers

Figure 8-7 Support for the **crash_kexec_post_notifiers** function

```
^Clinux-EVdrQm:~ # cat /proc/kallsyms |grep crash_kexec_post_notifiers
ffffffff816c3a20 r __param_str_crash_kexec_post_notifiers
ffffffff819c3da8 r __param_crash_kexec_post_notifiers
ffffffff81d58ec4 B crash_kexec_post_notifiers
```

- If yes, go to step 2.
 - If no, use method 1.
2. Add the **crash_kexec_post_notifiers** function to the startup configuration file **menu.lst** or **grub.cfg**.
Take **menu.lst** as an example.
 - a. Run the following command to open the **menu.lst** file:
vi /boot/grub/menu.lst
 - b. Add the **crash_kexec_post_notifiers** function to the startup item.

Figure 8-8 Editing the **menu.lst** file

```
# Modified by YaST2. Last modification on Thu Feb 22 10:51:10 UTC 2018
default 2
timeout 5
password --encrypted $6$XxIhQx9E6Kx6QF8hb7SVqVz3DFxV6q7LSUmzp0Fw4RTXl6Ce3Y.FpbIdOfs1tbSC0v7F.L.m$waroAFLeAanR10tsqhuYQM/dh7/

##Don't change this comment - YaST2 identifier: Original name: linux##
title UVP Linux Enterprise Server V200R003C00 - 3.0.93-0.8
root (hd0,0)
kernel /vmlinuz-3.0.93-0.8-default root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5 resume=/dev/sda1 splash=silent showopts
initrd /initrd-3.0.93-0.8-default

##Don't change this comment - YaST2 identifier: Original name: failsafe##
title Failsafe -- UVP Linux Enterprise Server V200R003C00 - 3.0.93-0.8
root (hd0,0)
kernel /vmlinuz-3.0.93-0.8-default root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5
initrd /initrd-3.0.93-0.8-default

title UVP Linux Enterprise Server V200R003C00
root (hd0,0)
kernel /boot/xen.gz dom0_mem=8192M mem_for_icoache=4096M balloon_zones=32768M dom0_max_vcpus=4 dom0_reserve_vcpus=4 numa=on console=
ted_guest=0 x2apic=1 crashkernel=192M816M watchdog=1 shm_dev_num=0 shm_client2server_size=128 shm_server2client_size=64 extra_guest_in
S_IG_enable=0 gnttab_max_nr_frames=3072 ple_gap=128 ple_window=4096 sched_credit_default_yield=0 apicv=1 crash_kexec_post_notifiers
module /boot/vmlinuz-3.0.93-0.8-xen console=tty0 console=ttyS0,115200 root=/dev/disk/by-id/scsi-35000c5001ce8b6a7-part5 vga=0x317
module /boot/initrd-3.0.93-0.8-xen
```

- c. Run the following command to restart the ECS for the modification to take effect:
reboot

8.6 How Can I Upgrade the Kernel of a Linux ECS?

Upgrade Notes

If tools have been installed on the Linux ECS, you must uninstall the tools before upgrading the ECS kernel. Otherwise, the following issues may occur after the kernel is upgraded:

- The Linux ECS cannot identify the NIC, leading to network access failure.
- The Linux ECS cannot identify data disks. As a result, starting system mount points fails, and the ECS cannot start.

Background

PVOPS is the Xen driver delivered with Linux distributions.

Procedure

1. Log in to the ECS.
2. Check whether the Tools have been installed on the Linux ECS, taking the SUSE Linux Enterprise Server 11 SP1 as an example.

- a. Run the following command on any directory to view the ECS driver:

```
lsmod | grep xen
```

Figure 8-9 Viewing the ECS driver

```
Linux:~/Desktop # lsmod | grep xen
xen_vbd                23600  3
cdrom                  40567  2 sr_mod,xen_vbd
xen_vmdq               4295   0
xen_vnif               36374   0
xen_balloon            14925   1 xen_vnif
xen_hcall              1867   0
xen_platform_pci      94554   5 xen_vbd,xen_vmdq,xen_vnif,xen_balloon,xen_hcall,[permanent]
```

- b. Run the following command to view the driver path, taking a disk driver as an example:

```
modinfo xen_vbd
```

Figure 8-10 Viewing the driver path

```
Linux:~/Desktop # modinfo xen_vbd
filename:              /lib/modules/2.6.32.12-0.7-default/updates/pvdriver/xen-vbd/xen-vbd.ko
license:               Dual BSD/GPL
alias:                 xen:vbd
srcversion:            5D8B666F0EA3F1E31B58F0C
depends:                xen-platform-pci,cdrom
vermagic:              2.6.32.12-0.7-default SMP mod_unload modversions
```

- c. Check whether **pvdriver** is contained in the driver path.
 - If so, the tools have been installed in the ECS. Then, go to step [3](#).
 - If no, go to step [4](#).
3. Uninstall the tools.
 - a. Run the following command to switch to user **root**:

su root

- b. Run the following command to uninstall Tools in the root directory:
/etc/.uvp-monitor/uninstall

 **NOTE**

After Tools is uninstalled, ECS monitoring metrics may be lost and monitoring data cannot be collected. To resolve this issue, you can compile and install the UVP Tools. For details, see <https://github.com/UVP-Tools/UVP-Tools/>.

4. Upgrade the kernel using the method determined by yourself.
5. Check whether the Linux ECS driver supports PVOPS. Use any one of the following methods:
 - Method 1:
Determine based on the ECS OS.
 - All Linux distribution OSs are delivered with a Xen open-source driver, which supports PVOPS.
 - The SUSE Linux Enterprise Server 11 SP3 provided by the OS competence center is not delivered with any Xen open-source driver and does not support PVOPS.
 - Method 2:
Check whether the ECS driver has a Xen driver module. If so, the ECS driver supports PVOPS. To obtain the data, run the following command in any directory:

lsmod | grep xen**Figure 8-11** Viewing the ECS driver

```
[root@localhost ~]# lsmod | grep xen
xen_vnif          59585  0 [permanent]
xen_vbd           50857  0
xen_balloon      45641  1 xen_vnif,[permanent]
xen_platform_pci 118125  3 xen_vnif,xen_vbd,xen_balloon,[permanent]
```

 **NOTE**

The name of a Xen driver module varies depending on the Linux distribution OS. You only need to check whether the driver has a driver module with the **XEN** field.

- Method 3:
Run the **cat /boot/config* | grep -i xen** command in any directory and check whether the **XEN** field is contained in the command output. If so, the ECS driver supports PVOPS.

Figure 8-12 Viewing the XEN field

```
root@ubuntu:/home# cat /boot/config* | grep -i xen
CONFIG_XEN=y
CONFIG_XEN_DOM0=y
CONFIG_XEN_PVHVM=y
CONFIG_XEN_MAX_DOMAIN_MEMORY=500
CONFIG_XEN_SAVE_RESTORE=y
# CONFIG_XEN_DEBUG_FS is not set
CONFIG_XEN_PVH=y
CONFIG_PCI_XEN=y
```

- Upgrade the kernel based on the result obtained in step 5.
 - If the Linux ECS driver supports PVOPS, go to step 8.
 - If the Linux ECS driver does not support PVOPS, go to step 7.
- Install the open-source component xen-kmp so that the ECS driver supports PVOPS. For instructions about how to use PVOPS, see "Optimizing a Linux Private Image" in *Image Management Service User Guide*.
- (Optional) Configure required parameters based on the defect list for certain Linux distribution OSs.

To obtain the defect list, go to following URL:

<https://github.com/UVP-Tools/UVP-Tools/tree/master/docs>

8.7 Why Cannot My ECS OS Start Properly?

- Check the image based on which the ECS was created. If the image is a public one, this issue is not caused by private image sources.
- Click **Apply for Server** and check whether the same ECS can be created. If not, this image may have been canceled.
- Change the ECS OS to one that is available on the management console.

8.8 How Can I Fix the Meltdown and Spectre Security Vulnerabilities on Intel Processor Chips?

Symptom

On January 3, 2018 (Beijing time), severe security vulnerabilities Meltdown and Spectre were found on Intel processor chips. The details are as follows:

Vulnerability name: Severe chip-level vulnerabilities on Intel processor chips

Vulnerability IDs: CVE-2017-5753, CVE-2017-5715, and CVE-2017-5754

Severity: High risk

Vulnerability description: High-risk CPU kernel vulnerabilities Meltdown (CVE-2017-5754) and Spectre (CVE-2017-5715 and CVE-2017-5753) exposed. Attackers can use these vulnerabilities to bypass the memory security isolation mechanism and access the core data of the OS and other programs without authorization, leading to sensitive information disclosure.

Impact

These vulnerabilities will not cause attacks between ECSs but may introduce attacks between:

- Applications on an ECS
- Accounts for logging in to an ECS

If your ECSs are created using a public image, the cloud platform will automatically fix the vulnerabilities, which will not affect your services.

If your ECSs are created using a private image, determine whether to install a patch described in this section in the private image based on the impact of the vulnerabilities.

Background

For details about the official patch release of affected OSs, see [Security Notices](#).

Prerequisites

Tests have been fully verified and ECS data has been backed up.

Installing a Patch on Linux ECSs

Step 1 Log in to the ECS.

Step 2 Check whether Tools has been installed on the Linux ECS, taking the SUSE Linux Enterprise Server 11 SP1 as an example.

1. Run the following command on any directory to view the ECS driver:

```
lsmod | grep xen
```

Figure 8-13 Viewing the ECS driver

```
Linux:~/Desktop # lsmod | grep xen
xen_vbd                23600  3
cdrom                  40567  2 sr_mod,xen_vbd
xen_vmdq               4295   0
xen_vnif              36374   0
xen_balloon           14925  1 xen_vnif
xen_hcall              1867   0
xen_platform_pci      94554  5 xen_vbd,xen_vmdq,xen_vnif,xen_balloon,xen_hcall,[permanent]
```

2. Run the following command to view the driver path, taking a disk driver as an example:

```
modinfo xen_vbd
```

Figure 8-14 Viewing the driver path

```
Linux:~/Desktop # modinfo xen_vbd
filename:          /lib/modules/2.6.32.12-0.7-default/updates/pvdriver/xen-vbd/xen-vbd.ko
license:           Dual BSD/GPL
alias:             xen:vbd
srcversion:        5D8B666F0EA3F1E31B58F0C
depends:            xen-platform-pci,cdrom
vermagic:          2.6.32.12-0.7-default SMP mod_unload modversions
```

3. Check whether **pvdriver** is contained in the driver path.
 - If so, Tools have been installed in the ECS. Then, go to [Step 3](#).

- If no, go to [Step 4](#).

Step 3 Uninstall Tools.

1. Run the following command to switch to user **root**:
su root
2. Run the following command to uninstall Tools in the root directory:
/etc/.uvp-monitor/uninstall
3. Run the following command to restart the ECS:
reboot

Step 4 Install the patch to upgrade the kernel. For details, see [Background](#). **NOTE**

After updating the kernel, run the **reboot** command to restart the ECS.

Step 5 Check whether the patch has been installed.

1. Check whether the ECS is running properly.
2. Check whether the requirements specified in the **Verification** column of [Background](#) are met.

 **NOTE**

After the patch is installed, the ECS uses the driver delivered with the OS. In this event, the memory usage and disk usage of Linux ECSs will not be monitored. The other features and functions are not affected. If the memory usage and disk usage must be monitored, contact customer service.

----End

Checking Whether Security Vulnerabilities Have Been Fixed on Linux

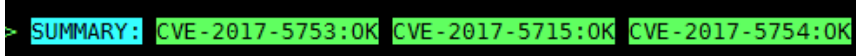
1. Click [spectre-meltdown-checker](#) to obtain **spectre-meltdown-checker.sh**.
2. Upload the script to the ECS.
3. Run the following commands on the ECS and check whether the Meltdown or Spectre vulnerability has been fixed based on the script prompt:

```
chmod +x spectre-meltdown-checker.sh
```

```
sudo bash spectre-meltdown-checker.sh
```

[Figure 8-15](#) shows the command output.

Figure 8-15 Command output after the script is executed



```
> SUMMARY: CVE-2017-5753:OK CVE-2017-5715:OK CVE-2017-5754:OK
```

OK indicates that the vulnerability has been fixed, and **KO** indicates that the vulnerability has not been fixed. The information shown in [Figure 8-15](#) indicates that the CVE-2017-5753, CVE-2017-5715, and CVE-2017-5754 vulnerabilities have been fixed.

Enabling or Disabling the Security Vulnerability Patch on Linux

CPU speculative execution optimizes performance. Therefore, fixing the Meltdown or Spectre vulnerability may deteriorate performance under specific workloads.

If the impact on the system performance is unacceptable or a better protection solution is available, you can disable certain or all security protection policies.

Determine the optimal security policy based on application scenarios:

- Meltdown vulnerability

Page Table Isolation (PTI) takes effect on the kernel. This function is suitable for CVE-2017-5754.

- Spectre vulnerability

Indirect Branch Restricted Speculation (IBRS) takes effect on specified registers (MSR) in SPEC_CTRL model. Working with retpoline, IBRS controls Indirect Branch Prediction Barriers (IBPBs) on specified registers (MSR) in PRED_CMD model. This function is suitable for CVE-2017-5715.

NOTE

The CVE-2017-5753 vulnerability is fixed by a kernel patch and cannot be disabled. No obvious impact was detected for the patch in Red Hat performance tests.

- **Disabling the Meltdown Vulnerability Patch**

To prevent the enabling of PTI from deteriorating the system performance, or a better protection solution is available, perform the following operations to disable the patch:

- a. Modify kernel parameters based on OSs:

- CentOS, EulerOS, Ubuntu, Fedora, and Red Hat: Add the kernel parameter **nopti**.
- Debian and OpenSUSE: Add the kernel parameter **pti=off**.

- b. Restart the ECS.

- **Disabling the Spectre Vulnerability Patch**

To prevent the Spectre vulnerability fixing from deteriorating the system performance, or a better protection solution is available, perform the following operations to disable the patch:

- a. Modify kernel parameters based on OSs:

- CentOS, EulerOS, Fedora, Debian, Red Hat, and OpenSUSE: Add the kernel parameter **spectre_v2=off**.
- Ubuntu: Add the kernel parameter **nospectre_v2=off**.

- b. Restart the ECS.

If you are using one of the following OSs, visit their official website for more details.

Red Hat: <https://access.redhat.com/articles/3311301?spm=a2c4g.11186623.2.20.42b49d4aJuKYx2>

SUSE: <https://www.suse.com/support/kb/doc/?spm=a2c4g.11186623.2.21.42b49d4avOXw7d&id=7022512>

Ubuntu: <https://wiki.ubuntu.com/SecurityTeam/KnowledgeBase/SpectreAndMeltdown>

8.9 How Can I Enable SELinux on an ECS Running CentOS?

Symptom

SELinux is disabled on ECSs running CentOS 7.5 by default. After I enable SELinux by running `/etc/selinux/config` and enter the login password, the login failed.

This section describes how to resolve this issue based on enabled SELinux.

Solution

The operations described in this section are performed on ECSs running CentOS 7.5.

1. Run the following command to change **SELINUX=disabled** in the SELinux configuration file to **SELINUX=enforcing**:

```
vim /etc/selinux/config
```

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=enforcing
# SELINUXTYPE= can take one of three two values:
#   targeted - Targeted processes are protected,
#   minimum - Modification of targeted policy. Only selected processes are protected.
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

2. Run the following command to automatically enable SELINUX on the file system upon ECS restarting:

```
touch /.autorelabel
```

3. Run the following command to restart the ECS for the configuration to take effect:

```
reboot
```

NOTE

After the preceding command is executed, the system automatically restarts twice.

9 Disk Partition, Attachment, and Expansion FAQ

9.1 How Can I Adjust System Disk Partitions?

Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can manually adjust the partitions to expand the system disk.

There are two ways to expand a system disk:

- Consider the empty partition as a new partition and attach this partition to a directory in the root partition after formatting it. For details, see this section.
- Add the empty partition to the root partition to be expanded. For detailed operations, see the following:
 - [How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?](#)
 - [How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?](#)

Procedure

This section uses an ECS running CentOS 7.3 64bit as an example. A 60 GB system disk was created with the ECS. However, the capacity of the system disk partition is displayed as only 40 GB.

To use the 20 GB capacity, performing the following operations:

Step 1 View disk partitions.

1. Log in to the ECS as user **root**.
2. Run the following command to view details about the ECS disk:

```
fdisk -l
```

In the following command output, **/dev/xvda** or **/dev/vda** indicates the system disk.

Figure 9-1 Viewing details about the disk

```
[root@ecs-8d6c ~]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/xvda1      38G  1.2G   35G   4% /
devtmpfs        899M    0  899M   0% /dev
tmpfs           908M    0  908M   0% /dev/shm
tmpfs           908M  8.4M  900M   1% /run
tmpfs           908M    0  908M   0% /sys/fs/cgroup
tmpfs          182M    0  182M   0% /run/user/0
[root@ecs-8d6c ~]# fdisk -l

Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 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: 0x0004d5e5

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1 *          2048       79980543   39989248    83  Linux
/dev/xvda2            79980544   83886079    1952768    82  Linux swap / Solaris
[root@ecs-8d6c ~]# _
```

3. Run the following command to view disk partitions:

parted -l /dev/xvda

Figure 9-2 Viewing disk partitions

```
[root@ecs-8d6c ~]# parted -l /dev/xvda
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvda: 64.4GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number  Start   End     Size    Type     File system  Flags
  1      1049kB  41.0GB  40.9GB  primary  ext4         boot
  2      41.0GB  42.9GB  2000MB  primary  linux-swap(v1)
```

Step 2 Create a partition for the expanded system disk capacity.

1. Run the following command to switch to the fdisk mode (taking **/dev/xvda** as an example):

fdisk /dev/xvda

Information similar to the following is displayed:

```
[root@ecs-8d6c ~]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help):
```

2. Enter **n** and press **Enter** to create a new partition.

Because the system disk has two existing partitions, the system automatically creates the third one.

Information similar to the following is displayed.

Figure 9-3 Creating a new partition

```
root@ecs-8d6c ~]# fdisk /dev/xvda
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): n
Partition type:
   p   primary (2 primary, 0 extended, 2 free)
   e   extended
Select (default p):
Using default response p
Partition number (3,4, default 3):
First sector (83886080-125829119, default 83886080):
Using default value 83886080
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
Using default value 125829119
Partition 3 of type Linux and of size 20 GiB is set

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
root@ecs-8d6c ~]#
```

3. Enter the new partition's start cylinder number and press **Enter**.
The start cylinder number must be greater than the end cylinder numbers of existing partitions. In this example, use the default value for the new partition's start cylinder number and press **Enter**. Information similar to the following is displayed.

Figure 9-4 Specifying the new partition's start cylinder number

```
First sector (83886080-125829119, default 83886080):
Using default value 83886080
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
```

4. Enter the new partition's end cylinder number and press **Enter**.
In this example, use the default value for the new partition's end cylinder number and press **Enter**. Information similar to the following is displayed.

Figure 9-5 Specifying the new partition's end cylinder number

```
Last sector, +sectors or +size{K,M,G} (83886080-125829119, default 125829119):
Using default value 125829119
Partition 3 of type Linux and of size 20 GiB is set
```

5. Enter **p** and press **Enter** to view the created partition.
Information similar to the following is displayed.

Figure 9-6 Viewing the created partition

```
Command (m for help): p
Disk /dev/xvda: 64.4 GB, 64424509440 bytes, 125829120 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: 0x0004d5e5

   Device Boot      Start         End      Blocks   Id  System
/dev/xvda1    *          2048     79980543    39989248   83   Linux
/dev/xvda2                79980544     83886079     1952768   82   Linux swap / Solaris
/dev/xvda3                83886080    125829119    20971520   83   Linux
```

6. Enter **w** and press **Enter**. The system saves and exits the partition. The system automatically writes the partition result into the partition list. Then, the partition is created. Information similar to the following is displayed.

Figure 9-7 Completing the partition creation

```
Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.

WARNING: Re-reading the partition table failed with error 16: Device or resource busy.
The kernel still uses the old table. The new table will be used at
the next reboot or after you run partprobe(8) or kpartx(8)
Syncing disks.
```

7. Run the following command to view disk partitions:
parted -l /dev/xvda

Figure 9-8 Viewing disk partitions

```
Disk Flags:
Number  Start   End     Size    Type    File system  Flags
 1      1049kB  41.0GB  40.9GB  primary ext4         boot
 2      41.0GB  42.9GB  2000MB  primary linux-swap(v1)
 3      42.9GB  64.4GB  21.5GB  primary ext4
```

- Step 3** Run the following command to synchronize the modifications in the partition list with the OS:

partprobe

- Step 4** Configure the type of the new partition file system.

1. Run the following command to view the type of the file system:
df -TH

Figure 9-9 Viewing the file system type

```
[root@ecs-8d6c ~]# df -TH
Filesystem      Type      Size  Used Avail Use% Mounted on
/dev/xvda1      ext4      41G   1.3G   37G    4% /
devtmpfs        devtmpfs  943M    0   943M    0% /dev
tmpfs           tmpfs     952M    0   952M    0% /dev/shm
tmpfs           tmpfs     952M   8.8M   944M    1% /run
tmpfs           tmpfs     952M    0   952M    0% /sys/fs/cgroup
tmpfs           tmpfs     191M    0   191M    0% /run/user/0
[root@ecs-8d6c ~]#
```

2. Run the following command to format the partition (taking the **ext4** type as an example):

```
mkfs -t ext4 /dev/xvda3
```

NOTE

Formatting the partition requires a period of time. During this time, observe the system running status and do not exit the system.

Information similar to the following is displayed:

```
[root@ecs-86dc ]# mkfs -t ext4 /dev/xvda3
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
1790544 inodes, 7156992 blocks
357849 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2155872256
219 block groups
32768 blocks per group, 32768 fragments per group
8176 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

Step 5 Mount the new partition to the target directory.

If you mount the new partition to a directory that is not empty, the subdirectories and files in the directory will be hidden. It is a good practice to mount the new partition to an empty directory or a newly created directory. If you want to mount the new partition to a directory that is not empty, temporarily move the subdirectories and files in the directory to another directory. After the partition is mounted, move the subdirectories and files back.

Take the newly created directory **/root/new** as an example.

1. Run the following command to create the **/root/new** directory:

```
mkdir /root/new
```

2. Run the following command to mount the new partition to the **/root/new** directory:

```
mount /dev/xvda3 /root/new
```

Information similar to the following is displayed:

```
[root@ecs-86dc ]# mount /dev/xvda3 /root/new  
[root@ecs-86dc ]#
```

3. Run the following command to view the mounted file systems:

df -TH

Information similar to the following is displayed:

Figure 9-10 Viewing the mounted file systems

```
[root@ecs-8d6c ~]# df -TH  
Filesystem      Type      Size  Used Avail Use% Mounted on  
/dev/xvda1      ext4      41G   1.3G   37G   4% /  
devtmpfs        devtmpfs  943M    0   943M   0% /dev  
tmpfs           tmpfs     952M    0   952M   0% /dev/shm  
tmpfs           tmpfs     952M   8.8M   944M   1% /run  
tmpfs           tmpfs     952M    0   952M   0% /sys/fs/cgroup  
/dev/xvda3      ext4      22G    47M   20G   1% /root/new  
tmpfs           tmpfs     191M    0   191M   0% /run/user/0  
[root@ecs-8d6c ~]#
```

- Step 6** Determine whether to set automatic mounting upon system startup for the new disk.

If you do not set automatic mounting upon system startup, you must mount the new partition to the specified directory again after the ECS is restarted.

- If automatic mounting is required, go to [Step 7](#).
- If automatic mounting is not required, no further action is required.

- Step 7** Set automatic mounting upon system startup for the new disk.

NOTE

Do not set automatic mounting upon system startup for unformatted disks because this will cause ECS startup failures.

1. Run the following command to obtain the file system type and UUID:

blkid

Figure 9-11 Viewing the file system type

```
[root@ecs-8d6c ~]# blkid  
/dev/xvda1: UUID="7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea" TYPE="ext4"  
/dev/xvda2: UUID="5de3cf2c-30c6-4fb2-9e63-830439d4e674" TYPE="swap"  
/dev/xvda3: UUID="96e5e028-b0fb-4547-a82a-35ace1086c4f" TYPE="ext4"  
[root@ecs-8d6c ~]#
```

According to the preceding figure, the UUID of the new partition is 96e5e028-b0fb-4547-a82a-35ace1086c4f.

2. Run the following command to open the **fstab** file using the vi editor:

vi /etc/fstab

3. Press **i** to enter editing mode.
4. Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=96e5e028-b0fb-4547-a82a-35ace1086c4f /root/new ext4 defaults 0  
0
```

5. Press **Esc**, run the following command, and press **Enter**. The system saves the configurations and exits the vi editor.

```
:wq
```

NOTE

If you want to detach a new disk for which automatic mounting upon system startup has been set, you must delete the automatic mounting configuration before you detach the disk. Otherwise, the ECS cannot be started after you detach the disk. To delete the automatic mounting configuration, perform the following operations:

1. Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

2. Press **i** to enter editing mode.
3. Delete the following statement:

```
UUID=96e5e028-b0fb-4547-a82a-35ace1086c4f /root/new ext4 defaults 0 0
```

4. Press **Esc**, run the following command, and press **Enter**. The system saves the configurations and exits the vi editor.

```
:wq
```

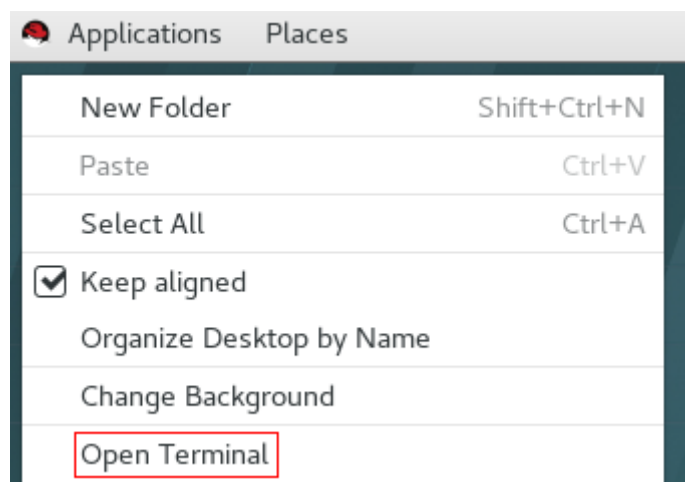
----End

9.2 How Can I Obtain the Mapping Between Disk Partitions and Disk Devices on a Linux ECS?

For a Linux ECS, its disk partitions correspond to disk devices. This section uses a Linux ECS running Red Hat Enterprise Linux 7 as an example to describe how to obtain the mapping between disk partitions and disk devices.

1. Log in to the Linux ECS as user **root**.
2. Right-click in the blank area of the desktop and choose **Open Terminal** from the shortcut menu.

Figure 9-12 open terminal



3. Run the following command to view disk partitions and disk devices:
fdisk -l

Figure 9-13 Viewing disk partitions and disk devices

```

root@localhost:~
File Edit View Search Terminal Help
[root@localhost ~]# fdisk -l

Disk /dev/xvda: 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: 0x000ba575

   Device Boot      Start         End      Blocks   Id  System
 /dev/xvda1    *          2048       2099199     1048576   83   Linux
 /dev/xvda2             2099200     16777215     7339008   83   Linux
 /dev/xvda3             16777216     20971519     2097152   82   Linux swap / Solaris

Disk /dev/xvdb: 21.5 GB, 21474836480 bytes, 41943040 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

```

Table 9-1 lists the mapping between disk partitions and disk devices.

Table 9-1 Mapping between disk partitions and disk devices

Disk Partition	Disk Device
xvda	xvda
xvdb	xvdb
xvdc	xvdc
xvdd	xvdd
xvde	xvde
xvdf	xvdf
xvdg	xvdg
xvdh	xvdh
xvdi	xvdi
xvdj	xvdj
xvdk	xvdk
xvdl	xvdl
xvdm	xvdm
xvdn	xvdn
xvdo	xvdo
xvdp	xvdp
xvdq	xvdq

Disk Partition	Disk Device
xvdr	xvdr
xvds	xvds
xvdt	xvdt
xvdu	xvdu
xvdv	xvdv
xvdw	xvdw
xvdx	xvdx

9.3 How Can I Add the Empty Partition of an Expanded System Disk to the End Root Partition Online?

Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can add the empty partition to the root partition of the system disk.

This section describes how to add the empty partition to the end root partition online.

Procedure

In the following operations, the ECS that runs CentOS 6.5 64bit and has a 50 GB system disk is used as an example. The system disk has two partitions, **/dev/xvda1: swap** and **/dev/xvda2: root**, and the root partition is the end partition.

1. Run the following command to view disk partitions:

```
parted -l /dev/xvda
```

```
[root@sluo-ecs-5e7d ~]# parted -l /dev/xvda
Disk /dev/xvda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
```

```
Number Start End Size Type File system Flags
 1 1049kB 4296MB 4295MB primary linux-swaps(v1)
 2 4296MB 42.9GB 38.7GB primary ext4 boot
```

2. Run the following command to obtain the file system type and UUID:

```
blkid
```

```
/dev/xvda1: UUID="25ec3bdb-ba24-4561-bcdc-802edf42b85f" TYPE="swap"
/dev/xvda2: UUID="1a1ce4de-e56a-4e1f-864d-31b7d9dfb547" TYPE="ext4"
```

3. Run the following command to install the growpart tool:

This tool may be integrated in the **cloud-utils-growpart/cloud-utils/cloud-initramfs-tools/cloud-init** package. Run the **yum install cloud-*** command to ensure it is available.

yum install cloud-utils-growpart

4. Run the following command to expand the root partition (the second partition) using growpart:

growpart /dev/xvda 2

```
[root@sluo-ecs-5e7d ~]# growpart /dev/xvda 2
CHANGED: partition=2 start=8390656 old: size=75495424 end=83886080 new:
size=96465599,end=104856255
```

5. Run the following command to verify that online capacity expansion is successful:

parted -l /dev/xvda

```
[root@sluo-ecs-5e7d ~]# parted -l /dev/xvda
Disk /dev/xvda: 53.7GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos

Number Start End Size Type File system Flags
 1 1049kB 4296MB 4295MB primary linux-swap(v1)
 2 4296MB 53.7GB 49.4GB primary ext4 boot
```

6. Run the following command to expand the capacity of the file system:

resize2fs -f \$Partition name

Suppose the partition name is **/dev/xvda2**, run the following command:

```
[root@sluo-ecs-a611 ~]# resize2fs -f /dev/xvda2
resize2fs 1.42.9 (28-Dec-2013)
Filesystem at /dev/xvda2 is mounted on /; on-line resizing required
old_desc_blocks = 3, new_desc_blocks = 3
....
[root@sluo-ecs-a611 ~] # df -hT //Check file system capacity expansion
```

9.4 How Can I Add the Empty Partition of an Expanded System Disk to the Non-end Root Partition Online?

Scenarios

If the capacity of system disk partitions is inconsistent with the actual system disk capacity after an ECS is created, you can add the empty partition to the root partition of the system disk.

This section describes how to add the empty partition to the non-end root partition online.

Procedure

In the following operations, the ECS that runs CentOS 6.5 64bit and has a 100 GB system disk is used as an example. The system disk has two partitions, **/dev/xvda1: root** and **/dev/xvda2: swap**, and the root partition is not the end partition.

1. Run the following command to view disk partitions:

parted -l /dev/xvda

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:

Number Start End Size Type File system Flags
```



```
1 1049kB 41.0GB 40.9GB primary ext4 boot
2 41.0GB 42.9GB 2000MB primary linux-swap(v1)
```

The first is the root partition, and the second is the swap partition.

2. View and edit the fstab partition table to delete the swap partition attachment information.

- a. Run the following command to view the fstab partition table:

tail -n 3 /etc/fstab

```
[root@sluo-ecs-a611 ~]# tail -n 3 /etc/fstab
#
UUID=7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea / ext4 defaults 1 1
UUID=5de3cf2c-30c6-4fb2-9e63-830439d4e674 swap swap defaults 0 0
```

- b. Run the following command to edit the fstab partition table and delete the swap partition attachment information.

vi /etc/fstab

tail -n 3 /etc/fstab

```
[root@sluo-ecs-a611 ~]# vi /etc/fstab
[root@sluo-ecs-a611 ~]# tail -n 3 /etc/fstab
#
UUID=7c4fce5d-f8f7-4ed6-8463-f2bd22d0ddea / ext4 defaults 1 1
```

3. Run the following command to disable the swap partition:

swapoff -a

4. Delete the swap partition.

- a. Run the following command to view the partition:

parted /dev/xvda

```
[root@sluo-ecs-a611 ~]# parted /dev/xvda
GNU Parted 3.1
Using /dev/xvda
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) help
align-check TYPE N          check partition N for TYPE(min|opt) alignment
help [COMMAND]              print general help, or help on COMMAND
mklabel,mktable LABEL-TYPE create a new disklabel (partition table)
mkpart PART-TYPE [FS-TYPE] START END make a partition
name NUMBER NAME            name partition NUMBER as NAME
print [devices]free[list,all|NUMBER] display the partition table, available devices, free space,
all found partitions, or a
particular partition
quit                          exit program
rescue START END            rescue a lost partition near START and END
rm NUMBER                    delete partition NUMBER
select DEVICE                choose the device to edit
disk_set FLAG STATE         change the FLAG on selected device
disk_toggle [FLAG]          toggle the state of FLAG on selected device
set NUMBER FLAG STATE       change the FLAG on partition NUMBER
toggle [NUMBER [FLAG]]      toggle the state of FLAG on partition NUMBER
unit UNIT                    set the default unit to UNIT
version                       display the version number and copyright information of GNU
Parted
(parted)
```

- b. Press **p**.

```
Disk /dev/xvda: 107GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

```
Number Start End Size Type File system Flags
1 1049kB 41.0GB 40.9GB primary ext4 boot
2 41.0GB 42.9GB 2000MB primary linux-swap(v1)
```

- c. Run the following command to delete the partition:

```
rm 2
```

```
(parted) rm2
```

- d. Press **p**.

```
(parted) p
```

```
Disk /dev/xvda: 107GB
```

```
Sector size (logical/physical): 512B/512B
```

```
Partition Table: msdos
```

```
Disk Flags:
```

Number	Start	End	Size	Type	File system	Flags
1	1049kB	41.0GB	40.9GB	primary	ext4	boot

- e. Run the following command to edit the fstab partition table:

```
quit
```

```
(parted) quit
```

```
Information: You may need to update /etc/fstab.
```

5. Run the following command to view partition after the swap partition is deleted:

```
parted -l /dev/xvda
```

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
```

```
Disk /dev/xvda: 107GB
```

```
Sector size (logical/physical): 512B/512B
```

```
Partition Table: msdos
```

```
Disk Flags:
```

Number	Start	End	Size	Type	File system	Flags
1	1049kB	41.0GB	40.9GB	primary	ext4	boot

6. Run the following command to install the growpart tool:

This tool may be integrated in the **cloud-utils-growpart/cloud-utils/cloud-initramfs-tools/cloud-init** package. Run the **yum install cloud-*** command to ensure it is available.

```
yum install cloud-utils-growpart
```

7. Run the following command to expand the root partition (the first partition) using growpart:

```
growpart /dev/xvda 1
```

```
[root@sluo-ecs-a611 ~]# growpart /dev/xvda 1
```

```
CHANGED: partition=1 start=2048 old: size=79978496 end=79980544 new:  
size=209710462,end=209712510
```

8. Run the following command to verify that online capacity expansion is successful:

```
[root@sluo-ecs-a611 ~]# parted -l /dev/xvda
```

```
Disk /dev/xvda: 107GB
```

```
Sector size (logical/physical): 512B/512B
```

```
Partition Table: msdos
```

```
Disk Flags:
```

Number	Start	End	Size	Type	File system	Flags
1	1049kB	107GB	107GB	primary	ext4	boot

9. Run the following command to expand the capacity of the file system:

```
resize2fs -f $Partition name
```

Suppose the partition name is **/dev/xvda1**, run the following command:

```
[root@sluo-ecs-a611 ~]# resize2fs -f /dev/xvda1
```

```
resize2fs 1.42.9 (28-Dec-2013)
```

```
Filesystem at /dev/xvda1 is mounted on /; on-line resizing required
```

```
old_desc_blocks = 3, new_desc_blocks = 3
```

```
....  
[root@sluo-ecs-a611 ~] # df -hT //Check file system capacity expansion
```

9.5 Can I Attach Multiple Disks to an ECS?

Yes. The ECSs created after the disk function upgrade can have up to 60 attached disks.

- When you create an ECS, you can attach 24 disks to it.
- After you create an ECS, you can attach up to 60 disks to it.

Table 9-2 Numbers of disks that can be attached to a newly created ECS

ECS Type	Maximum VBD Disks	Maximum SCSI Disks	Constraint
Xen	60	59	VBD disks + SCSI disks \leq 60 (This constraint does not apply to local disks.) The number of local disks is determined based on the ECS flavor.
KVM (excluding D2 ECSs)	24	59	VBD disks + SCSI disks \leq 60 (This constraint does not apply to local disks.) The number of local disks is determined based on the ECS flavor.
D2	24	30	VBD disks + SCSI disks \leq 54 (This constraint does not apply to local disks.) The number of local disks is determined based on the ECS flavor.

NOTE

- The system disk of an ECS is of VBD type. The maximum number of SCSI disks is 59.
- For a D-series KVM ECS, its local disks use two SCSI controllers, indicating that 30 SCSI drive letters are used. A maximum of 30 SCSI disks can be attached to such an ECS.

The maximum number of disks that you can attach to an ECS that was created before the disk function upgrade remains unchanged, as shown in [Table 9-3](#).

Table 9-3 Numbers of disks that can be attached to an existing ECS

ECS Type	Maximum VBD Disks	Maximum SCSI Disks	Maximum Local Disks	Constraint
Xen	60	59	59	VBD disks + SCSI disks + Local disks \leq 60
KVM	24	23	59	VBD disks + SCSI disks \leq 24

How Can I Check Whether an ECS Is Created Before or After the Disk Function Upgrade?

1. Log in to management console.
2. Under **Compute**, click **Elastic Cloud Server**.
3. Click the name of the target ECS. The page providing details about the ECS is displayed.
4. Click the **Disks** tab.
5. Check the number of disks that can be attached to the ECS to determine the total number of disks.
 - If the total number of disks that can be attached is 24 (including the system disk), the ECS is created before the disk function upgrade.
 - If the total number of disks that can be attached is 60 (including the system disk), the ECS is created after the disk function upgrade.

9.6 What Are the Requirements for Attaching an EVS Disk to an ECS?

- The EVS disk and the target ECS must be located in the same AZ.
- The target ECS must be in **Running** or **Stopped** state.
- The EVS disk must not be frozen.
- For yearly/monthly ECSs:

If you detach the system disk that you purchased when creating an ECS and want to continue using it as a system disk, you can only attach it to the original ECS. If you want to use it as a data disk, you can attach it to any ECS.

If you detach the non-shared data disk that you purchased when creating an ECS and want to attach it again, you can only attach it to the original ECS as a data disk.

9.7 Which ECSs Can Be Attached with SCSI EVS Disks?

All KVM ECSs support SCSI EVS disks.


9.8 How Do I Obtain My Disk Device Name in the ECS OS Using the Device Identifier Provided on the Console?

Scenarios

You find that the device name displayed in the ECS OS is different from that displayed on the management console and you cannot determine which disk name is correct. This section describes how to obtain the disk name used in an ECS OS according to the device identifier on the console.

For details about how to attach disks, see [Attaching an EVS Disk to an ECS](#).

Obtaining the Disk ID of an ECS on the Console

1. Log in to the management console.
2. Under **Compute**, choose **Elastic Cloud Server**.
3. Click the target ECS name in the ECS list.
The ECS details page is displayed.
4. Click the **Disks** tab and then click  to expand the disk information.
5. Check the device type and ID of the disk.

NOTE

- If **Device Identifier** is not displayed on the page, stop the ECS and restart it.
- KVM ECS
 - If **Device Type** is **VBD**, use a serial number or BDF to obtain the disk device name.
If you use a serial number (recommended) to obtain the disk device name, see [Using a Serial Number to Obtain a Disk Device Name \(Linux\)](#).
If you use a BDF to obtain the disk device name, see [Using a BDF to Obtain a Disk Device Name \(Linux\)](#).
 - If **Device Type** is **SCSI**, use a WWN to obtain the disk device name.
For details, see [Using a WWN to Obtain a Disk Device Name \(Linux\)](#).

Using a Serial Number to Obtain a Disk Device Name (Linux)

If a serial number is displayed on the console, run either of the following commands to obtain the device name.

```
udevadm info --query=all --name=/dev/xxx | grep ID_SERIAL  
ll /dev/disk/by-id/*
```

 NOTE

A serial number is the first 20 digits of a disk UUID.

For example, if the serial number of the VBD disk is 62f0d06b-808d-480d-8, run either of the following commands:

```
udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL
```

```
ll /dev/disk/by-id/*
```

The following information is displayed:

```
[root@ecs-ab63 ~]# udevadm info --query=all --name=/dev/vdb | grep ID_SERIAL
E: ID_SERIAL=62f0d06b-808d-480d-8
[root@ecs-ab63 ~]# ll /dev/disk/by-id/*
lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9 -> ../vda
lrwxrwxrwx 1 root root 10 Dec 30 15:56 /dev/disk/by-id/virtio-128d5bfd-f215-487f-9-part1 -> ../vda1
lrwxrwxrwx 1 root root 9 Dec 30 15:56 /dev/disk/by-id/virtio-62f0d06b-808d-480d-8 -> ../vdb
```

/dev/vdb is the disk device name.

Using a BDF to Obtain a Disk Device Name (Linux)

1. Run the following command to use a BDF to obtain the device name:

```
ll /sys/bus/pci/devices/BDF disk ID/virtio*/block
```

For example, if the BDF disk ID of the VBD disk is 0000:02:02.0, run the following command to obtain the device name:

```
ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block
```

The following information is displayed:

```
[root@ecs-ab63 ~]# ll /sys/bus/pci/devices/0000:02:02.0/virtio*/block
total 0
drwxr-xr-x 8 root root 0 Dec 30 15:56 vdb
```

/dev/vdb is the disk device name.

Using a WWN to Obtain a Disk Device Name (Linux)

1. Log in to the ECS as user **root**.
2. Run the following command to view the disk device name:

```
ll /dev/disk/by-id |grep WWN|grep scsi-3
```

For example, if the WWN obtained on the console is 6888603000008b32fa16688d09368506, run the following command:

```
ll /dev/disk/by-id |grep 6888603000008b32fa16688d09368506|grep scsi-3
```

The following information is displayed:

```
[root@host-192-168-133-148 block]# ll /dev/disk/by-id/ |grep 6888603000008b32fa16688d09368506 |
grep scsi-3
lrwxrwxrwx 1 root root 9 May 21 20:22 scsi-36888603000008b32fa16688d09368506 -> ../sda
```

9.9 Why Does a Linux ECS with a SCSI Disk Attached Fail to Be Restarted?

Symptom

For a Linux ECS with a SCSI disk attached, if you have enabled automatic SCSI disk attachment upon ECS startup in **/etc/fstab** and the disk drive letter (for example, **/dev/sdb**) is used, the ECS fails to restart.

Possible Causes

SCSI disk allocation is determined based on the ID of the slot accommodating the disk as well as the available drive letter in the ECS. Each time you attach a disk to the ECS, an idle drive letter is automatically allocated in sequence. When the ECS starts, the disks are loaded in slot sequence. A slot ID corresponds to a drive letter.

After the SCSI disk is detached from the running ECS, the slot sequence for disks may change, leading to the disk drive letter being changed after the ECS is restarted. As a result, the slot IDs do not correspond to the drive letters, and the ECS fails to restart.

Solution

1. Log in to the ECS as user **root**.
2. Run the following command to obtain the SCSI ID according to the drive letter of the SCSI disk:

```
ll /dev/disk/by-id/|grep Disk drive letter
```

For example, if the drive letter of the SCSI disk is **/dev/sdb**, run the following command:

```
ll /dev/disk/by-id/|grep sdb
```

```
CNA64_22:/opt/galax/eucalyptus/ecs_scripts # ll /dev/disk/by-id/|grep sdb
lrwxrwxrwx 1 root root 9 Dec 6 11:26 scsi-3688860300001436b005014f890338280 -> ../sdb
lrwxrwxrwx 1 root root 9 Dec 6 11:26 wwn-0x688860300001436b005014f890338280 -> ../sdb
```

3. Change the drive letter (for example, **/dev/sdb**) of the SCSI disk to the corresponding SCSI ID in the **/etc/fstab** file.

```
/dev/disk/by-id/SCSI ID
```

For example, if the SCSI ID obtained in step 2 is **scsi-3688860300001436b005014f890338280**, use the following data to replace **/dev/sdb**:

```
/dev/disk/by-id/scsi-3688860300001436b005014f890338280
```

9.10 How Can I Check Whether the ECSs Attached with the Same Shared SCSI Disk Are in the Same ECS Group?

Scenarios

Shared EVS disks of the SCSI type support SCSI locks. To improve data security, the shared EVS disks of the SCSI type must be attached to the ECSs in the same anti-

affinity ECS group. This section describes how to check whether the ECSs attached with the same shared SCSI disk are in the same ECS group.

- For details about ECS groups, see [Managing ECS Groups](#).
- For details about using shared EVS disks, see [Shared EVS Disks and Usage Instructions](#).

Procedure

1. Log in to the management console.
2. Under **Storage**, click **Elastic Volume Service**.
3. Click the target shared SCSI disk to view its details.
4. In the **Servers** pane on the right side of the page, the ECSs to which the shared SCSI disk is attached are displayed.

In this example, the ECSs to which the shared SCSI disk **volume-0001** is attached are **ecs-0001** and **ecs-0002**.

Figure 9-14 Details about the disk

Disks > volume-0001 Expand Capacity

Summary Servers Backups Snapshots Tags

Basic Information	
ID	213f4fc7-68e4-4b2a-a0c4-83c77d7e0696
Name	volume-0001
Region	
AZ	AZ3
Disk Type	Common I/O
Capacity (GB)	100
Max. IOPS	700/2,200
	2 IOPS per GB, IOPS limit: 700, IOPS burst limit: 2,200.
Function	Data disk
Image	--
Created	Sep 26, 2019 14:13:55 GMT+08:00

Configuration Information	
Disk Sharing	Enabled
Device Type	SCSI

Servers Attach Disk

- ecs-0002 View Metric Running
- ecs-0001 View Metric Running

Backups Create Backup

You have not created any backup yet.

VBS allows you to create backups for EVS disks on the management console without stopping servers.

Snapshots Create Snapshot

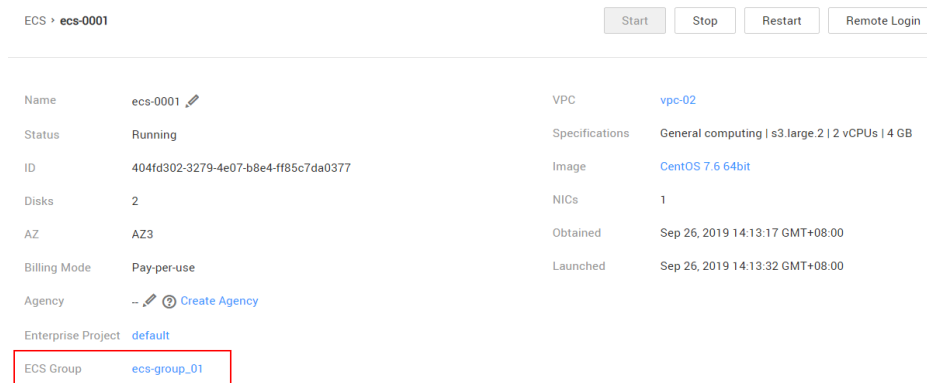
You have not created any snapshot yet.

Snapshots can be created to quickly save disk data at specified time points.

5. Click the names of these ECSs, respectively. On the page that provides details about an ECS, you can view the ECS group to which the current ECS belongs. In this example, the ECS group to which ECS **ecs-0001** belongs is **ecs-group_01**.

NOTE

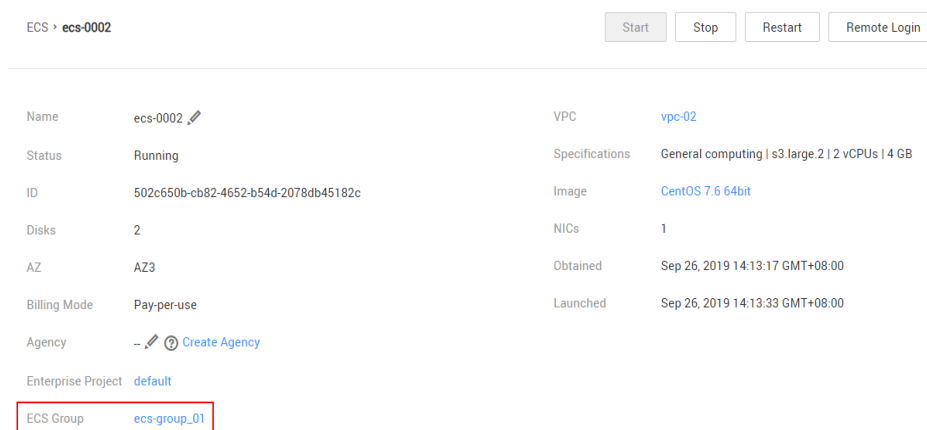
If the ECS group field is left blank, the ECS has not been added to any ECS group.

Figure 9-15 Details about an ECS (1)

The screenshot shows the details for ECS instance `ecs-0001`. At the top right, there are buttons for `Start`, `Stop`, `Restart`, and `Remote Login`. The instance is in a `Running` state. The `ECS Group` is `ecs-group_01`, which is highlighted with a red box. Other details include Name: `ecs-0001`, VPC: `vpc-02`, Specifications: `General computing | s3.large.2 | 2 vCPUs | 4 GB`, Image: `CentOS 7.6 64bit`, Disks: `2`, AZ: `AZ3`, Billing Mode: `Pay-per-use`, Agency: `...` with a `Create Agency` link, Enterprise Project: `default`, and ID: `404fd302-3279-4e07-b8e4-f85c7da0377`.

ECS > <code>ecs-0001</code>		<code>Start</code> <code>Stop</code> <code>Restart</code> <code>Remote Login</code>	
Name	<code>ecs-0001</code>	VPC	<code>vpc-02</code>
Status	<code>Running</code>	Specifications	<code>General computing s3.large.2 2 vCPUs 4 GB</code>
ID	<code>404fd302-3279-4e07-b8e4-f85c7da0377</code>	Image	<code>CentOS 7.6 64bit</code>
Disks	<code>2</code>	NICs	<code>1</code>
AZ	<code>AZ3</code>	Obtained	<code>Sep 26, 2019 14:13:17 GMT+08:00</code>
Billing Mode	<code>Pay-per-use</code>	Launched	<code>Sep 26, 2019 14:13:32 GMT+08:00</code>
Agency	<code>...</code> Create Agency		
Enterprise Project	<code>default</code>		
ECS Group	<code>ecs-group_01</code>		

In this example, the ECS group to which ECS `ecs-0002` belongs is `ecs-group_01`.

Figure 9-16 Details about an ECS (2)

The screenshot shows the details for ECS instance `ecs-0002`. At the top right, there are buttons for `Start`, `Stop`, `Restart`, and `Remote Login`. The instance is in a `Running` state. The `ECS Group` is `ecs-group_01`, which is highlighted with a red box. Other details include Name: `ecs-0002`, VPC: `vpc-02`, Specifications: `General computing | s3.large.2 | 2 vCPUs | 4 GB`, Image: `CentOS 7.6 64bit`, Disks: `2`, AZ: `AZ3`, Billing Mode: `Pay-per-use`, Agency: `...` with a `Create Agency` link, Enterprise Project: `default`, and ID: `502c650b-cb82-4652-b54d-2078db45182c`.

ECS > <code>ecs-0002</code>		<code>Start</code> <code>Stop</code> <code>Restart</code> <code>Remote Login</code>	
Name	<code>ecs-0002</code>	VPC	<code>vpc-02</code>
Status	<code>Running</code>	Specifications	<code>General computing s3.large.2 2 vCPUs 4 GB</code>
ID	<code>502c650b-cb82-4652-b54d-2078db45182c</code>	Image	<code>CentOS 7.6 64bit</code>
Disks	<code>2</code>	NICs	<code>1</code>
AZ	<code>AZ3</code>	Obtained	<code>Sep 26, 2019 14:13:17 GMT+08:00</code>
Billing Mode	<code>Pay-per-use</code>	Launched	<code>Sep 26, 2019 14:13:33 GMT+08:00</code>
Agency	<code>...</code> Create Agency		
Enterprise Project	<code>default</code>		
ECS Group	<code>ecs-group_01</code>		

This indicates that the shared SCSI disk `volume-0001` is attached to ECSs `ecs-0001` and `ecs-0002`, and both ECSs are in ECS group `ecs-group_01`.

9.11 Can All Users Use the Encryption Feature?

The permissions of users in a user group to use the encryption feature are as follows:

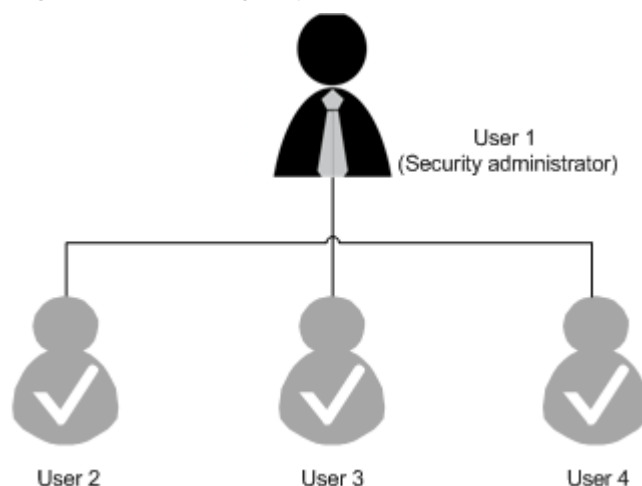
- The user who has security administrator permissions can grant KMS access permissions to EVS for using the encryption feature.
- When a common user who does not have security administrator permissions attempts to use the encryption feature, the condition varies depending on whether the user is the first one in the user group to use this feature.
 - If the common user is the first one in the user group to use the encryption feature, the common user must request a user who has security administrator permissions to grant the common user permissions. Then, the common user can use the encryption feature.

- If the common user is not the first one in the user group to use the encryption feature, the user directly has the permissions to use the encryption feature.

The following section uses a user group as an example to describe how to grant KMS access permissions to EVS for using the encryption feature.

For example, a user group shown in [Figure 9-17](#) consists of four users, user 1 to user 4. User 1 has security administrator permissions. Users 2, 3, and 4 are common users who do not have security administrator permissions.

Figure 9-17 User group



Scenario 1: User 1 Uses the Encryption Feature

In this user group, if user 1 uses the encryption feature for the first time, the procedure is as follows:

1. User 1 creates Xrole to grant KMS access permissions to EVS.

After user 1 grants permissions, the system automatically creates key **evs/default** for encrypting EVS disks.

NOTE

When user 1 uses the encryption feature for the first time, the user must grant the KMS access permissions to EVS. Then, all the users in the user group can use the encryption feature by default.

2. User 1 selects a key.

One of the following keys can be used:

- Default key **evs/default**
- Custom key, which was created before using the EVS disk encryption feature

After user 1 uses the encryption feature, all other users in the user group can use this feature, without requiring to contact user 1 for permissions granting.

Scenario 2: Common User Uses the Encryption Feature

In this user group, when user 3 uses the encryption feature for the first time:

1. The system displays a message indicating that the user has no permissions.
2. User 3 asks user 1 to create Xrole to grant KMS access permissions to EVS.

After user 1 grants the permissions, user 3 and all other users in the user group can use the encryption feature by default.

9.12 How Can I Add an ECS with Local Disks Attached to an ECS Group?

An ECS group logically isolates ECSs. The ECSs in an ECS group support anti-affinity and are allocated on different hosts.

An ECS with local disks attached cannot be added to an ECS group after the ECS is created. Such ECSs can be added to an ECS group only during the ECS creation.

9.13 Will My EVS Disk Be Deleted When I Delete Its Server?

- For pay-per-use disks:
 - If such a disk is separately purchased and has been attached, the system will prompt you whether to delete the disk when you delete the server, and you can make the decision based on your requirements.
 - If such disks are purchased together with a server, the system disk will be deleted when you delete the server. For the data disks, the system will prompt you whether to delete the disks, and you can make the decision based on your service requirements.
- For yearly/monthly disks:

If the disks are purchased together with a server, they will be unsubscribed from when you unsubscribe from the server.

9.14 Why Does the Disk Drive Letter Change After the ECS Is Restarted?

Symptom

For a Linux ECS, the drive letter may change after an EVS disk is detached and then attached again, or after an EVS disk is detached and then the ECS is restarted.

Root Cause

When a Linux ECS has multiple disks attached, it allocates drive letters in the attachment sequence and names the disks as `/dev/vda1`, `/dev/vdb1`, and `/dev/vdc1`, etc.

After a disk is detached and then attached again, or after a disk is detached and the ECS is restarted, the drive letter may change.

For example, an ECS has three disks attached: **/dev/vda1**, **/dev/vdb1**, and **/dev/vdc1**. The mounting parameters in **/etc/fstab** are as follows:

cat /etc/fstab

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050bdc8a / ext4 defaults 0 1
/dev/vdb1 /data1 ext4 defaults 0 0
/dev/vdc1 /data2 ext4 defaults 0 0
```

After **/dev/vdb1** is detached and the ECS is restarted, **/dev/vdc1** becomes **/dev/vdb1** and is mounted to **/data**. In such a case, no disk is mounted to **/data2**.

The change of drive letters can affect the running of applications. To solve this problem, you are advised to use the universally unique identifiers (UUIDs) to replace **/dev/vdx** because a UUID uniquely identifies a disk partition in the Linux OS.

Solution

1. Log in to the ECS.
2. Run the following command to obtain the partition UUID:

```
blkid Disk partition
```

In this example, run the following command to obtain the UUID of the **/dev/vdb1** partition:

```
blkid /dev/vdb1
```

Information similar to the following is displayed:

```
[root@ecs-test-0001 ~]# blkid /dev/vdb1
/dev/vdb1: UUID="b9a07b7b-9322-4e05-ab9b-14b8050cd8cc" TYPE="ext4"
```

The UUID of the **/dev/vdb1** partition is displayed.

3. Run the following command to open the **fstab** file using the vi editor:

```
vi /etc/fstab
```

4. Press **i** to enter the editing mode.
5. Move the cursor to the end of the file and press **Enter**. Then, add the following information:

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /data1 ext4 defaults 0 0
```

The parameters are defined as follows:

- **UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc**: UUID of a disk partition.
- **/data1**: directory on which the partition is mounted. You can run **df -TH** to query the directory.
- **ext4**: File system format of the partition. You can run **df -TH** to query the format.
- **defaults**: partition mount option. Normally, this parameter is set to **defaults**.
- **0** (the first one): whether to use Linux dump backup.
 - **0**: Linux dump backup is not used. Normally, dump backup is not used, and you can set this parameter to **0**.
 - **1**: Linux dump backup is used.

- **0** (the second one): fsck option, that is, whether to use fsck to check disks during startup.
 - **0**: fsck is not used.
 - If the mount point is the root partition (/), this parameter must be set to **1**.

When this parameter is set to **1** for the root partition, this parameter for other partitions must start with **2** so that the system checks the partitions in the ascending order of the values.

6. Repeat steps **2** to **5** to replace the UUID of **/dev/vdc1**.
7. Run the following command again to check the disk mounting parameters:

```
cat /etc/fstab
```

The following information is displayed:

```
UUID=b9a07b7b-9322-4e05-ab9b-14b8050bdc8a / ext4 defaults 0 1
UUID=b9a07b7b-9322-4e05-ab9b-14b8050cd8cc /data1 ext4 defaults 0 0
UUID=b9a07b7b-9322-4e05-ab9b-14b8050ab6bb /data2 ext4 defaults 0 0
```

9.15 How Can I Obtain Data Disk Information If Tools Are Uninstalled?

If you uninstall Tools from a Linux ECS in a non-PVOPS system, data disks cannot be identified. In such a case, you can create a new ECS and attach the data disks of the original ECS to the new ECS and view information about the data disks. The procedure is as follows:

1. Log in to the management console and create a new ECS.

NOTE

Ensure that the new ECS is located in the same AZ and has the same parameter settings as the original ECS.

2. (Optional) On the **Elastic Cloud Server** page, locate the row containing the original ECS, click **More** in the **Operation** column, and select **Stop**. On the **Stop ECS** page, select **Forcibly stop the preceding ECSs** and click **Yes** to forcibly stop the original ECS.

Manually refresh the **Elastic Cloud Server** page. The original ECS is stopped once the **Status** changes to **Stopped**.

NOTE

The ECSs running certain OSs support online data disk detaching. If your OS supports this feature, you can detach data disks from the running ECS.

3. View information about the data disks attached to the original ECS.

NOTE

If the original ECS has multiple data disks attached, repeat steps **4** to **6** to attach each data disk to the new ECS.

4. Click a data disk. The **Elastic Volume Service** page is displayed.
5. Select the data disk to be detached and click **Detach** in the **Operation** column. On the **Detach Disk** page, select the original ECS and click **OK** to detach the data disk from the original ECS.

Manually refresh the **Elastic Volume Service** page. The data disk is detached from the original ECS once the **Status** changes to **Available**.

6. Select the detached data disk and click **Attach** in the **Operation** column. On the **Attach Disk** page, click the new ECS, select a device name, and click **OK** to attach the data disk to the new ECS.

Manually refresh the EVS list. The data disk is attached to the new ECS once the **Status** value changes to **In-use**. You can then log in to the management console and view information about the data disk of the new ECS.

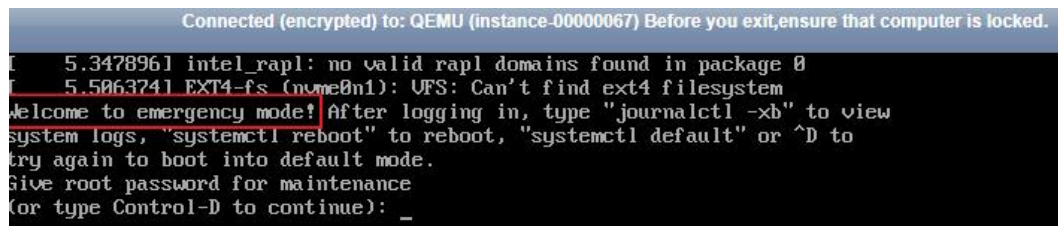
9.16 How Can I Rectify the Fault That May Occur on a Linux ECS with an NVMe SSD Disk Attached?

Symptom

When a Linux ECS with an NVMe SSD disk attached, such as a P1 ECS, becomes faulty, you must contact the administrator to remotely rebuild the ECS again.

If automatic NVMe SSD disk attachment upon ECS startup is enabled in `/etc/fstab` on the faulty ECS, the system disk recovers after the ECS is created. However, the attached NVMe SSD disk does not have a file system, and automatic NVMe SSD disk attachment upon ECS startup fails to take effect. As a result, the ECS enters the emergency mode, as shown in [Figure 9-18](#).

Figure 9-18 Emergency mode



```
Connected (encrypted) to: QEMU (instance-00000067) Before you exit, ensure that computer is locked.
5.3478961 intel_rapl: no valid rapl domains found in package 0
5.5063741 EXT4-fs (nvrme0n1): UFS: Can't find ext4 filesystem
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): _
```

To ensure that the new ECS is functional, you must manually delete the attachment information in `/etc/fstab`.

NOTE

If the NVMe SSD disk is faulty, data on it will be lost. The operations provided in this section are only used to restore automatic NVMe SSD disk attachment to an ECS, but not restoring the data on the disk.

Solution

1. Log in to the ECS.
2. Enter the password of user **root** to log in to the ECS.

Figure 9-19 Logging in to the ECS

```
Connected (encrypted) to: QEMU (instance-00000067) Before you exit,ensure that co
[ 5.347896] intel_rapl: no valid rapl domains found in package 0
[ 5.506374] EXT4-fs (nvme0n1): UFS: Can't find ext4 filesystem
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):
root@localhost ~]#
```

3. Run the following command to edit the `/etc/fstab` file:
vi /etc/fstab
4. Delete the attaching information of the NVMe SSD disk and save the file.

Figure 9-20 Deleting the automatic attaching information

```
#
# /etc/fstab
# Created by anaconda on Wed Aug 9 09:22:35 2017
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/cl-root / xfs defaults 0 0
UUID=17cbcc3f-0b23-4eaa-84f6-6bc60583b521 /boot xfs defaults 0 0
/dev/mapper/cl-swap swap swap defaults 0 0
/dev/nvme0n1 /for_nvme ext3 defaults 0 0
~
~
~
"/etc/fstab" 12L, 506C
```

5. Run the following command to restart the ECS:
reboot
6. Verify that the ECS recovers and can be logged in.

Figure 9-21 Logging in to the ECS

```
Connected (encrypted) to: QEMU (instanc
CentOS Linux 7 (Core)
Kernel 3.10.0-514.el7.x86_64 on an x86_64

localhost login: _
```

9.17 Why Are Disk Error Logs Printed After a Disk Attached to an ECS Is Formatted with the ext4 File System?

Symptom

When a VBD disk is attached to an ECS and the partition is in ext4 format, the following log may be displayed on the console:

blk_update_request: operation not supported error, dev vdb, sector 826298624 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0

Figure 9-22 Printed logs

```
[ 1732.062294] blk_update_request: operation not supported error, dev vdb, sector 8504 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.366259] blk_update_request: operation not supported error, dev vdb, sector 12592 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.654260] blk_update_request: operation not supported error, dev vdb, sector 16688 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1732.942279] blk_update_request: operation not supported error, dev vdb, sector 20784 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
[ 1733.230277] blk_update_request: operation not supported error, dev vdb, sector 24880 op 0x9:(WRITE_ZEROES) flags 0x800 phys_seg 0 prio class 0
```

Involved OSs: Ubuntu 20.04, CentOS 8.0, CentOS 8.1, and other ECSs whose kernel versions are 4.18 or later

Root Cause

VBD disks do not support the advanced SCSI command WRITE_ZEROES.

If the ECS OS kernel version is 4.18 or later and the disk partition is formatted with the ext4 file system, the WRITE_ZEROES command is delivered. The system does not support the command and prints a log, which has no impact on the ECS performance and you can ignore it.

10 Data Backup and Restoration FAQ

10.1 Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?

Yes, ECS can automatically recover if the physical host becomes faulty.

ECSs run on physical hosts. Although there are multiple mechanisms to ensure system reliability, fault tolerance, and high availability, host hardware might be damaged or power failures might occur. If physical hosts cannot be powered on or restarted due to damage, CPU and memory data will be lost and live migration cannot be used to recovery ECSs.

The cloud platform provides automatic recovery by default to restart ECSs through cold migration, ensuring high availability and dynamic ECS migration. If a physical host accommodating ECSs breaks down, the ECSs will automatically be migrated to a functional physical host to minimize the impact on your services. During the process, the ECSs will restart.

You can enable one-click monitoring on the Cloud Eye console so that you will be notified if any exceptions occur (if a physical host accommodating ECSs is faulty, the ECSs will automatically be migrated to a functional physical host). For details, see [One-Click Monitoring](#).

 **NOTE**

- Automatic recovery does not ensure user data consistency.
- An ECS can be automatically recovered only if the physical server on which it is deployed becomes faulty. This function does not take effect if the fault is caused by the ECS itself.
- An ECS can be automatically recovered only after the physical server on which it is deployed is shut down. If the physical server is not shut down due to a fault, for example, a memory fault, automatic recovery fails to take effect.
- An ECS can be automatically recovered only once within 12 hours if the server on which it is deployed becomes faulty.
- ECS automatic recovery may fail in the following scenarios:
 - No physical server is available for migration due to a system fault.
 - The target physical server does not have sufficient temporary capacity.
- An ECS with any of the following resources cannot be automatically recovered:
 - Local disk
 - Passthrough FPGA card
 - Passthrough InfiniBand NIC

11 Network Configuration FAQ

11.1 How Can I Configure the NTP and DNS Servers for an ECS?

For Linux OSs

Take the NTP and DNS servers running SUSE as an example.

Step 1 Configure the NTP server for the ECS.

1. Log in to the Linux ECS.
2. Run the following command to switch to user **root**:
sudo su -
3. Run the following command to edit the **ntp.conf** configuration file:
vim /etc/ntp.conf
4. Add the following statement to configure the NTP server:
server *Domain name or IP address of the NTP server*

Example:

If the IP address of the NTP server is 192.168.56.1, add the following statement:

server 192.168.56.1

5. Run the following command to start the NTP service upon system restart:
service ntp restart
6. Run the following command to check the status of the NTP server:
service ntp status

NOTE

If you want to disable NTP, perform the following steps:

1. Run the **service ntp stop** command to stop NTP.
2. Run the **systemctl disable ntp** command to disable the function of automatically starting NTP upon ECS startup.

Step 2 Configure the DNS server for the ECS.

1. Log in to the Linux ECS.
2. Run the following command to switch to user **root**:
sudo su -
3. Run the following command to edit the **resolv.conf** configuration file:
vi /etc/resolv.conf
4. Add the following statement to configure the DNS server:

nameserver = IP addresses of the DNS servers

Example:

If the IP addresses of the DNS servers are 8.8.8.8 and 4.4.4.4, add the following statements:

nameserver = 8.8.8.8

nameserver = 4.4.4.4

 **NOTE**

The IP addresses of the DNS servers must be the same as those in the VPC subnet. Otherwise, the DNS modification cannot persistently take effect.

5. Run the following command to restart the network:

rcnetwork restart

service network restart

/etc/init.d/network restart

----End

11.2 Does HUAWEI CLOUD Provide the NTP Server and How Can I Configure It?

Yes. HUAWEI CLOUD provides the NTP server, and you can use it only on the ECSs you have purchased on the HUAWEI CLOUD management console.

You can use the Huawei-provided NTP server or other NTP servers. The configuration procedures are the same. This section describes how to configure the Huawei-provided NTP server on an ECS.

 **NOTE**

ECSs created using x86 public images use chronyd for time synchronization by default. You do not need to configure the NTP server.

Background

If you use the NTP server provided by HUAWEI CLOUD, you also need to use the DNS server. [Table 11-1](#) lists the NTP servers provided by Huawei Cloud in different regions.

For details about how to obtain the DNS server address, see [What Are the Private DNS Server Addresses Provided by Huawei Cloud?](#)

Table 11-1 NTP servers

Region	NTP Server IP Address
EU-Dublin	ntp.huaweicloud.eu

Linux (chronyd)

The following section uses CentOS 7.3 as an example.

Step 1 Check whether the IP address of the DNS server is correct on the ECS.

1. Log in to the Linux ECS.
2. Run the following command to open the **resolv.conf** file:
vi /etc/resolv.conf
3. Check whether the **nameserver** value in the file is the same as the IP address of the DNS server provided in [What Are Huawei Cloud Private DNS Server Addresses?](#)
 - If yes, go to step **Step 3**.
 - If no, go to step **Step 2**.

Step 2 (Optional) Configure the DNS server for the ECS.

1. Log in to the Linux ECS.
2. Run the following command to edit the **resolv.conf** configuration file:
vi /etc/resolv.conf
3. Add the following statement to configure the DNS server:
nameserver *IP address of the DNS server*
Example:
Consider the **CN North** region as an example. Add the following statement:
nameserver 100.125.1.250

Step 3 Configure the NTP server for the ECS.

1. Log in to the Linux ECS.
2. Run the following commands to stop the chronyd process:
systemctl stop chronyd
systemctl disable chronyd
3. Run the following command to edit the **chrony.conf** configuration file:
vim /etc/chrony.conf
4. Add the following statement to configure the NTP server:
server *Domain name of the NTP server* **minpoll 4 maxpoll 10 iburst**
Example:
server ntp.myhuaweicloud.com minpoll 4 maxpoll 10 iburst
5. Run the following command to start the service upon system restart:
For Euler and CentOS:
systemctl restart chronyd

For SUSE:

service chronyd restart

 **NOTE**

Run the required command based on the OS running on the ECS.

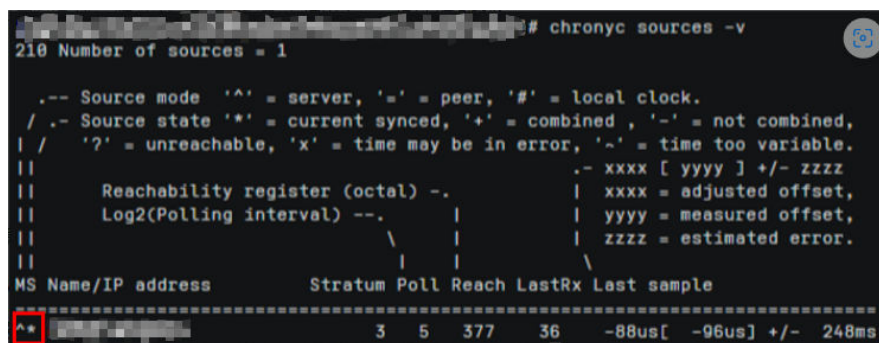
If the message "Failed to restart chronyd.service: Unit not found." is displayed, run the **yum -y install chrony** command.

6. Run the following command to check whether the time on the NTP server has been synchronized with that on the upper-layer NTP server:

chronyc sources -v

If "*" is displayed, the time has been synchronized.

Figure 11-1 Modification result



```
# chronyc sources -v
210 Number of sources = 1

.-- Source mode '^' = server, '=' = peer, '#' = local clock.
/.- Source state '*' = current synced, '+' = combined , '-' = not combined,
|/ '?' = unreachable, 'x' = time may be in error, '~' = time too variable.
||
||          Reachability register (octal) --.      | xxxx [ yyyy ] +/- zzzz
||          Log2(Polling interval) --.          | | xxxx = adjusted offset,
||          \ | |                               | | yyyy = measured offset,
||          \ | |                               | | zzzz = estimated error.
||          \ | |                               | |
||
MS Name/IP address         Stratum Poll Reach LastRx Last sample
=====
* 192.168.1.1                3     5   377   36    -88us[ -96us] +/- 248ms
```

 **NOTE**

It takes several minutes to perform NTP time synchronization for the first time.

7. Set the automatic startup of the chronyd service.

For Euler and CentOS:

systemctl enable chronyd

For SUSE:

chkconfig chronyd on

----End

Linux (ntpd)

The following section uses CentOS 7.3 as an example.

Step 1 Check whether the IP address of the DNS server is correct on the ECS.

1. Log in to the Linux ECS.
2. Run the following command to open the **resolv.conf** file:

vi /etc/resolv.conf

3. Check whether the **nameserver** value in the file is the same as the IP address of the DNS server provided in [What Are Huawei Cloud Private DNS Server Addresses?](#)

- If yes, go to [Step 3](#).

- If no, go to [Step 2](#).

Step 2 (Optional) Configure the DNS server for the ECS.

1. Log in to the Linux ECS.
2. Run the following command to edit the **resolv.conf** configuration file:

```
vi /etc/resolv.conf
```

3. Add the following statement to configure the DNS server:

```
nameserver IP address of the DNS server
```

Example:

Consider the **CN North** region as an example. Add the following statement:

```
nameserver 100.125.1.250
```

Step 3 Configure the NTP server for the ECS.

1. Log in to the Linux ECS.
2. Run the following commands to stop the chronyd process:

```
systemctl stop chronyd
```

```
systemctl disable chronyd
```

3. Run the following command to edit the **ntp.conf** configuration file:

```
vim /etc/ntp.conf
```

4. Add the following statement to configure the NTP server:

```
server Domain name of the NTP server
```

Example:

```
server ntp.myhuaweicloud.com
```

5. Run the following command to start the NTP service upon system restart:

For Euler and CentOS:

```
systemctl restart ntpd
```

For SUSE:

```
service ntpd restart
```

 **NOTE**

Run the required command based on the OS running on the ECS.

If the message "Failed to restart ntpd.service: Unit not found." is displayed, run the **yum -y install ntp** command.

6. Run the following command to check whether the time on the NTP server has been synchronized with that on the upper-layer NTP server:

```
ntpq -p
```

If "*" is displayed, the time has been synchronized.

 **NOTE**

It takes several minutes to perform NTP time synchronization for the first time.

7. Set the automatic startup of the NTP service.

For Euler and CentOS:

```
chkconfig ntpd on
```

For SUSE:

```
chkconfig ntpd on
```

```
----End
```

Follow-up Procedure

After the ECS is restarted, the DNS configuration is reset, and its IP address is changed to the IP address of the DNS server in the VPC subnet. Therefore, before restarting the ECS, check whether the DNS configuration in the VPC subnet is the same as the target DNS configuration. If they are different, modify the DNS configuration in the VPC subnet. For details, see [Modifying a Subnet](#).

11.3 How Do I Configure DNS for an ECS?

A DNS server is used to resolve domain names of file systems.

Scenarios

By default, the IP address of the DNS server used to resolve domain names of file systems is automatically configured on ECSs when creating ECSs. No manual configuration is needed except when the resolution fails due to a change in the DNS server IP address.

Procedure

- Step 1** Log in to the ECS as user **root**.
- Step 2** Run the **vi /etc/resolv.conf** command to edit the **/etc/resolv.conf** file. Add the DNS server IP address above the existing nameserver information. See [Figure 11-2](#).

Figure 11-2 Configuring DNS

```
; generated by /sbin/dhclient-script
search openstacklocal
nameserver 192.168.1.1
nameserver 114.114.114.114
nameserver 114.114.115.115
```

The format is as follows:

- Step 3** Press **Esc**, input **:wq**, and press **Enter** to save the changes and exit the vi editor.
- Step 4** Run the following command to check whether the IP address is successfully added:
cat /etc/resolv.conf
- Step 5** Run the following command to check whether an IP address can be resolved from the file system domain name:
nslookup File system domain name

NOTE

Obtain the file system domain name from the file system mount point.

Step 6 (Optional) In a network environment of the DHCP server, edit the `/etc/resolv.conf` file to prevent the file from being automatically modified upon an ECS startup, and prevent the DNS server IP address added in [Step 2](#) from being reset.

1. Run the following command to lock the file:

```
chattr +i /etc/resolv.conf
```

 **NOTE**

Run the `chattr -i /etc/resolv.conf` command to unlock the file if needed.

2. Run the following command to check whether the editing is successful:

```
lsattr /etc/resolv.conf
```

If the information shown in [Figure 11-3](#) is displayed, the file is locked.

Figure 11-3 A locked file

```
[root@cloud11174-Pia-test ~]# lsattr /etc/resolv.conf  
----i-----e- /etc/resolv.conf
```

----End

11.4 Can the ECSs of Different Accounts Communicate over an Intranet?

No. The ECSs of different accounts cannot communicate with each other over an intranet.

To enable the communication over an intranet, use the methods provided in the following table.

Scenario	Billing	Method
In the same region	Free of charge	Use VPC peering to enable the communication over an intranet. <ul style="list-style-type: none">• VPC Peering Connection Overview• Creating a VPC Peering Connection with a VPC in Another Account
In the same region	Billed	Use VPC Endpoint to enable the communication over an intranet. <ul style="list-style-type: none">• What Is VPC Endpoint?• Configuring a VPC Endpoint for Communications Across VPCs of Different Accounts• What Are the Differences Between VPC Endpoints and VPC Peering Connections?

11.5 Will ECSs That I Purchased Deployed in the Same Subnet?

You can customize your network to deploy the ECSs. Therefore, whether they are in the same subnet is totally up to you.

11.6 How Do I Change the CIDR Block of an ECS Subnet?

Scenarios


You want to change the CIDR block of an ECS subnet. After you create a subnet, you cannot directly change its CIDR block.

To change a CIDR block, you need to change the subnet.

Prerequisites

The ECS has been stopped.

Procedure

1. Log in to the management console.
2. Under **Compute**, click **Elastic Cloud Server**.
3. In the search box above the ECS list, enter the ECS name, IP address, or ID, and click  for search.
4. Click the name of the ECS whose subnet needs to be modified.
The page providing details about the ECS is displayed.
5. Click the **NICs** tab. Locate the row containing the NIC and click **Modify Private IP**.
The **Modify Private IP** dialog box is displayed.
6. Change the subnet and private IP address of the primary NIC as required.

NOTE

- You can only change to a subnet within the same VPC.
- If you do not specify the target private IP address, the system will automatically assign one to the primary NIC.

For example, the original subnet is **subnet-demo (192.168.0.0/24)** and the new subnet is **subnet-fe21 (192.168.6.0/25)**. Therefore, you change the ECS subnet CIDR block by changing the ECS subnet.

11.7 How Do I Configure Port Mapping?

Symptom

It is expected that the EIP and port on ECS 1 accessed from the public network can be automatically redirected to the EIP and port on ECS 2.

Linux

For example, to redirect port 1080 on ECS 1 to port 22 on ECS 2 with the following configurations:

Private IP address and EIP of ECS 1: 192.168.72.10 and 123.xxx.xxx.456

Private IP address of ECS 2: 192.168.72.20

NOTE

- Ensure that the desired ports have been enabled on the ECS security group and firewall.
- Ensure that the source/destination check function is disabled.

On the ECS details page, click **Network Interfaces** and disable **Source/Destination Check**.

By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. You need to disable the source/destination check.

Step 1 Log in to Linux ECS 1.

1. Run the following command to modify the configuration file:

```
vi /etc/sysctl.conf
```

2. Add **net.ipv4.ip_forward = 1** to the file.
3. Run the following command to complete the modification:

```
sysctl -p /etc/sysctl.conf
```

Step 2 Run the following commands to add rules to the **nat** table in **iptables** so that the access to port 1080 on ECS 1 can be redirected to port 22 on ECS 2:

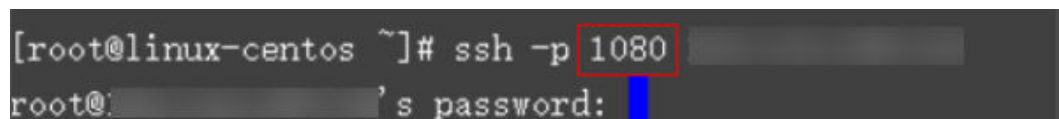
```
iptables -t nat -A PREROUTING -d 192.168.72.10 -p tcp --dport 1080 -j DNAT --to-destination 192.168.72.20:22
```

```
iptables -t nat -A POSTROUTING -d 192.168.72.20 -p tcp --dport 22 -j SNAT --to 192.168.72.10
```

Step 3 Run the following command to log in to port 1080 on ECS 1 for check:

```
ssh -p 1080 123.xxx.xxx.456
```

Figure 11-4 Port redirections on Linux



```
[root@linux-centos ~]# ssh -p 1080  
root@: 's password:
```

Enter the password to log in to ECS 2 with hostname **ecs-inner**.

Figure 11-5 Logging in to ECS 2

```
[root@ecs-inner ~]#
```

----End

11.8 How Can I Obtain the MAC Address of My ECS?

This section describes how to obtain the MAC address of an ECS.

NOTE

The MAC address of an ECS cannot be changed.

Linux (CentOS 6)

1. Log in to the Linux ECS.
2. Run the following command to view the MAC address of the ECS:

ifconfig

Figure 11-6 Obtaining the MAC address

```
[root@CentOS68-XEN ~]# ifconfig
eth0      Link encap:Ethernet  HWaddr FA:16:3E:2A:36:DE
          inet addr:192.168.22.227 Bcast:192.168.22.255 Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fe2a:36de/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:4699 errors:0 dropped:0 overruns:0 frame:0
          TX packets:2213 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:472826 (461.7 KiB)  TX bytes:438396 (428.1 KiB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1 Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:1 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:28 (28.0 b)  TX bytes:28 (28.0 b)
```

Linux (CentOS 7)

1. Log in to the Linux ECS.
2. Run the following command to view the MAC address of the ECS:

ifconfig

Figure 11-7 Obtaining the NIC information

```
[root@ecs-683a ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.0.65 netmask 255.255.255.0 broadcast 192.168.0.255
    inet6 fe80::f816:3eff:fec3:46fc prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:c3:46:fc txqueuelen 1000 (Ethernet)
    RX packets 14457 bytes 20617950 (19.6 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 1867 bytes 245185 (239.4 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet  netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1000 (Local Loopback)
    RX packets 0 bytes 0 (0.0 B)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 0 bytes 0 (0.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Run the following command to view the MAC address of NIC **eth0**:
ifconfig eth0 |grep "ether"

Figure 11-8 Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |egrep "ether"
    ether fa:16:3e:c3:46:fc txqueuelen 1000 (Ethernet)
[root@ecs-683a ~]#
```

4. Obtain the returned MAC address.
ifconfig eth0 |grep "ether" |awk '{print \$2}'

Figure 11-9 Obtaining the MAC address of eth0

```
[root@ecs-683a ~]# ifconfig eth0 |egrep "ether" |awk '{print $2}'
fa:16:3e:c3:46:fc
[root@ecs-683a ~]#
```

11.9 How Can I View and Modify Kernel Parameters of a Linux ECS?

Modify the kernel parameters only if the parameter settings affect your services. Kernel parameters vary depending on OS versions. If the parameter settings must be modified,

- Ensure that the target parameter settings meet service requirements.
- Modify the correct kernel parameters. For details about common kernel parameters, see [Table 11-2](#).
- Back up key ECS data before modifying kernel parameter settings.

Background

Table 11-2 Common Linux kernel parameters

Parameter	Description
net.core.rmem_default	Specifies the default size (in bytes) of the window for receiving TCP data.
net.core.rmem_max	Specifies the maximum size (in bytes) of the window for receiving TCP data.
net.core.wmem_default	Specifies the default size (in bytes) of the window for transmitting TCP data.
net.core.wmem_max	Specifies the maximum size (in bytes) of the window for transmitting TCP data.
net.core.netdev_max_backlog	Specifies the maximum number of packets that can be sent to a queue when the rate at which each network port receives packets is faster than the rate at which the kernel processes these packets.
net.core.somaxconn	Defines the maximum length of the listening queue for each port in the system. This parameter applies globally.
net.core.optmem_max	Specifies the maximum size of the buffer allowed by each socket.
net.ipv4.tcp_mem	Uses the TCP stack to show memory usage in memory pages (4 KB generally). The first value is the lower limit of memory usage. The second value is the upper limit of the load added to the buffer when the memory is overloaded. The third value is the upper limit of memory usage. When this value is reached, packets can be discarded to reduce memory usage. For a large BDP, increase the parameter value as needed. The unit of this parameter is memory page but not byte.
net.ipv4.tcp_rmem	Specifies the memory used by sockets for automatic optimization. The first value is the minimum number of bytes allocated to the socket buffer for receiving data. The second value is the default value, which is overwritten by rmem_default . The buffer size can increase to this value when the system load is not heavy. The third value is the maximum number of bytes allocated to the socket buffer for receiving data. This value is overwritten by rmem_max .

Parameter	Description
net.ipv4.tcp_wmem	<p>Specifies the memory used by sockets for automatic optimization.</p> <p>The first value is the minimum number of bytes allocated to the socket buffer for transmitting data.</p> <p>The second value is the default value, which is overwritten by wmem_default. The buffer size can increase to this value when the system load is not heavy.</p> <p>The third value is the maximum number of bytes allocated to the socket buffer for transmitting data. This value is overwritten by wmem_max.</p>
net.ipv4.tcp_keepalive_time	Specifies the interval at which keepalive detection messages are sent in seconds for checking TCP connections.
net.ipv4.tcp_keepalive_intvl	Specifies the interval at which keepalive detection messages are resent in seconds when no response is received.
net.ipv4.tcp_keepalive_probes	Specifies the maximum number of keepalive detection messages that are sent to determine a TCP connection failure.
net.ipv4.tcp_sack	Enables selective acknowledgment (value 1 indicates enabled). This configuration allows the transmitter to resend only lost packets, thereby improving system performance. However, this configuration will increase the CPU usage. You are suggested to enable selective acknowledgment for WAN communication.
net.ipv4.tcp_fack	Enables forwarding acknowledgment for selective acknowledgment (SACK), thereby reducing congestion. You are suggested to enable forwarding acknowledgment.
net.ipv4.tcp_timestamps	Specifies a TCP timestamp, which will add 12 bytes in the TCP packet header. This configuration calculates RTT using RFC1323, a more precise retransmission method upon timeout than retransmission. You are suggested to enable this parameter for higher system performance.
net.ipv4.tcp_window_scaling	Enables RFC1323-based window scaling by setting the parameter value to 1 if the TCP window is larger than 64 KB. The maximum TCP window is 1 GB. This parameter takes effect only when window scaling is enabled on both ends of the TCP connection.

Parameter	Description
net.ipv4.tcp_syncookies	Specifies whether to enable TCP synchronization (syncookie). This configuration prevents socket overloading when a large number of connections are attempted to set up. CONFIG_SYN_COOKIES must be enabled in the kernel for compilation. The default value is 0 , indicating that TCP synchronization is disabled.
net.ipv4.tcp_tw_reuse	Specifies whether a TIME-WAIT socket (TIME-WAIT port) can be used for new TCP connections. NOTE This parameter is valid only for clients and takes effect only when net.ipv4.tcp_timestamps is enabled. This parameter cannot be set to 1 if NAT is enabled. Otherwise, an error will occur in remote ECS logins.
net.ipv4.tcp_tw_recycle	Allows fast recycle of TIME-WAIT sockets. NOTE This parameter is valid only when net.ipv4.tcp_timestamps is enabled. Do not set this parameter to 1 if NAT is enabled. Otherwise, an error will occur during remote ECS logins.
net.ipv4.tcp_fin_timeout	Specifies the time (in seconds) during which a socket TCP connection that is disconnected from the local end remains in the FIN-WAIT-2 state. Process suspension may be caused by the disconnection from the peer end, continuous connection from the peer end, or other reasons.
net.ipv4.ip_local_port_range	Specifies local port numbers allowed by TCP/UDP.
net.ipv4.tcp_max_syn_backlog	Specifies the maximum number of connection requests that are not acknowledged by the peer end and that can be stored in the queue. The default value is 1024 . If the server is frequently overloaded, try to increase the value.
net.ipv4.tcp_low_latency	This option should be disabled if the TCP/IP stack is used for high throughput, low latency.
net.ipv4.tcp_westwood	Enables the congestion control algorithm on the transmitter end to evaluate throughput and improve the overall bandwidth utilization. You are suggested to enable the congestion control algorithm for WAN communication.
net.ipv4.tcp_bic	Enables binary increase congestion for fast long-distance networks so that the connections with operations being performed at a rate of Gbit/s can be functional. You are suggested to enable binary increase congestion for WAN communication.

Parameter	Description
net.ipv4.tcp_max_tw_buckets	Specifies the number of TIME_WAIT buckets, which defaults to 180000 . If the number of buckets exceeds the default value, extra ones will be cleared.
net.ipv4.tcp_synack_retries	Specifies the number of times that SYN+ACK packets are retransmitted in SYN_RECV state.
net.ipv4.tcp_abort_on_overflow	When this parameter is set to 1 , if the system receives a large number of requests within a short period of time but fails to process them, the system will send reset packets to terminate the connections. It is recommended that you improve system processing capabilities by optimizing the application efficiency instead of performing reset operations. Default value: 0
net.ipv4.route.max_size	Specifies the maximum number of routes allowed by the kernel.
net.ipv4.ip_forward	Forward packets between interfaces.
net.ipv4.ip_default_ttl	Specifies the maximum number of hops that a packet can pass through.
net.netfilter.nf_conntrack_tcp_timeout_established	Clears iptables connections that are inactive for a specific period of time.
net.netfilter.nf_conntrack_max	Specifies the maximum value of hash entries.

Viewing Kernel Parameters

- Method 1: Run the cat command in **/proc/sys** to view file content.
/proc/sys/ is a pseudo directory generated after the Linux kernel is started. The **net** folder in this directory stores all kernel parameters that have taken effect in the system. The directory tree structure is determined based on complete parameter names. For example, **net.ipv4.tcp_tw_recycle** corresponds to the **/proc/sys/net/ipv4/tcp_tw_recycle** file, and the content of the file is the parameter value.

Example:

To view the **net.ipv4.tcp_tw_recycle** value, run the following command:

```
cat /proc/sys/net/ipv4/tcp_tw_recycle
```

- Method 2: Use the **/etc/sysctl.conf** file.

Run the following command to view all parameters that have taken effect in the system:

```
/usr/sbin/sysctl -a
```

```
net.ipv4.tcp_syncookies = 1
net.ipv4.tcp_max_tw_buckets = 4096
```

```
net.ipv4.tcp_tw_reuse = 1
net.ipv4.tcp_tw_recycle = 1
net.ipv4.tcp_keepalive_time = 1800
net.ipv4.tcp_fin_timeout = 30
.....
net.ipv4.tcp_keepalive_time = 1200
net.ipv4.ip_local_port_range = 1024 65000
net.ipv4.tcp_max_syn_backlog = 8192
net.ipv4.tcp_rmem = 16384 174760 349520
net.ipv4.tcp_wmem = 16384 131072 262144
net.ipv4.tcp_mem = 262144 524288 1048576
.....
```

Modifying Kernel Parameter Settings

- Method 1: Run the echo command in **/proc/sys** to modify the file for the target kernel parameters.

The parameter values changed using this method take effect only during the current running and will be reset after the system is restarted. To make the modification take effect permanently, see method 2.

/proc/sys/ is a pseudo directory generated after the Linux kernel is started. The **net** folder in this directory stores all kernel parameters that have taken effect in the system. The directory tree structure is determined based on complete parameter names. For example, **net.ipv4.tcp_tw_recycle** corresponds to the **/proc/sys/net/ipv4/tcp_tw_recycle** file, and the content of the file is the parameter value.

Example:

To change the **net.ipv4.tcp_tw_recycle** value to **0**, run the following command:

```
echo "0" > /proc/sys/net/ipv4/tcp_tw_recycle
```

- Method 2: Use the **/etc/sysctl.conf** file.

The parameter values changed using this method take effect permanently.

- a. Run the following command to change the value of a specified parameter:

```
/sbin/sysctl -w kernel.domainname="example.com"
```

Example:

```
sysctl -w net.ipv4.tcp_tw_recycle="0"
```

- b. Run the following command to change the parameter value in the **/etc/sysctl.conf** file:

```
vi /etc/sysctl.conf
```

- c. Run the following command for the configuration to take effect:

```
/sbin/sysctl -p
```

11.10 Why Is the NIC Not Working?

Symptom

The NIC equipped on a D1 or H1 ECS does not work.

Possible Causes

The NIC driver has not been correctly installed.

Solution

D1 and H1 ECSs use passthrough NICs to improve network performance. You must install the passthrough NIC driver on the ECSs or the image that is used for creating the ECSs.

NOTE

If you mount the CD/DVD-ROM driver over a VPN, ensure that the VPN bandwidth is greater than 8 Mbit/s.

To install the passthrough NICE driver, do as follows:

Step 1 Obtain the passthrough NIC driver.

Passthrough NIC driver versions vary depending on the OS. For details, see [Table 11-3](#).

Table 11-3 NIC driver versions and OSs

NIC Driver Version	OS	How to Obtain
ixgbev 2.16.4	CentOS 7.2 64bit	https://sourceforge.net/projects/e1000/files/ixgbev%20stable/2.16.4/

Step 2 Log in to the ECS.

Step 3 Install the passthrough NIC driver on the ECS. In this procedure, CentOS 7.2 64bit is used as an example.

1. Configure the passthrough NIC.

Not all ECS OSs identify passthrough NICs using the standard NIC naming rule of **eth x** , where x is a number. If this is the case, you must configure the ECS so that it can identify the passthrough NIC. The procedure is as follows:

- a. Run the following command to view all NICs on the ECS and identify the passthrough NIC:

```
ifconfig -a
```

- b. Run the following command to switch to the directory where configuration files are stored:

```
cd /etc/sysconfig/network-scripts/
```

- c. Run the following command to create a configuration file for the passthrough NIC:

```
cp ifcfg-eth0 ifcfg-NIC_name
```

In the preceding command, *NIC_name* specifies the name of the passthrough NIC.

- d. Use the vi editor to edit this configuration file:

```
vi ifcfg-NIC_name
```

- e. Set the **DEVICE** parameter in the configuration file to the name of the passthrough NIC. The following is an example configuration:

```
DEVICE="NIC_name"  
BOOTPROTO="dhcp"  
ONBOOT="yes"  
STARTMODE="onboot"
```

- f. Run the following command to restart the network service and allow the configuration to take effect:

```
service network restart
```

2. Upload the obtained passthrough NIC driver to a directory on the ECS, for example, **/home**.
3. Switch to user **root** on the ECS CLI and open the target directory.
In this example, the passthrough NIC driver is stored in the **/home** directory. Run the **cd /home** command to switch to the target directory.
4. Run the following command to decompress the software package.
tar -zxvf ixgbevf-2.16.4.tar.gz
5. Run the following command to switch to the generated **src** directory:
cd ixgbevf-2.16.4/src
6. Run the following commands to install the driver:
make
make install
7. Run the following command to restart the ECS to make the drive take effect:
reboot
8. Switch to user **root** on the ECS CLI and open the **src** directory, for example, by running the **cd /home/ixgbevf-2.16.4/src** command. Then, run the following commands to check whether the driver has been installed:

```
rmmod ixgbevf  
insmod ./ixgbevf.ko  
ethtool -i NIC_name
```

In the preceding command, *NIC_name* specifies the passthrough NIC name, for example, **ens5**.

NOTE

- After you run the **rmmod ixgbevf** command, the system may display an error message. This message does not affect the installation of the passthrough NIC driver and can be ignored.
 - *NIC_name* specifies the passthrough NIC name, for example, **ens5**.
9. Check the driver status based on the displayed information.
In this example, the driver is installed if **driver** is **ixgbevf** and **version** is **2.16.4**.

----End

11.11 Why Can't I Use DHCP to Obtain a Private IP Address?

Symptom

You attempt to use DHCP to obtain a private IP address, but you cannot obtain the IP address.

- For Linux, a private IP address cannot be assigned.

NOTE

You are advised to use a public image to create an ECS. All public images support DHCP continuous discovery mode.

Solution (Linux)

The following uses CentOS 7.2 as an example. For solutions about other OSs, see the corresponding help documentation.

1. Log in to the ECS and run the following command:
ps -ef | grep dhclient
2. If the dhclient process does not exist, restart the NIC or run any of the following commands to initiate a DHCP request:
dhclient eth0, ifdown eth0 + ifup eth0, or dhcpcd eth0
3. If the DHCP client does not send any requests for a long time, for example, the issue recurs after the NIC is restarted, do the following:
 - a. Run the following command to configure a static IP:
vi /etc/sysconfig/network-scripts/ifcfg-eth0

```
BOOTPROTO=static
IPADDR=192.168.1.100 #IP address (modified)
NETMASK=255.255.255.0 #Mask (modified)
GATEWAY=192.168.1.1 #Gateway IP address (modified)
```
 - b. Restart the ECS to make the network settings take effect.
 - c. Select an image in which DHCP runs stably.
4. If the fault persists, obtain the messages in **/var/log/messages** on the affected ECS, use the MAC address of the affected NIC to filter the desired log, and check whether there is any process that prevents DHCP from obtaining an IP address.
5. If the fault persists, contact technical support.

11.12 How Can I Test the Network Performance of Linux ECSs?

Use netperf and iperf3 to test network performance between ECSs. The test operations include preparations, TCP bandwidth test, UDP PPS test, and latency test.

Background

- Tested ECS: an ECS that is tested for network performance. Such an ECS functions as the client (TX end) or server (RX end) in netperf tests.
- Auxiliary ECS: an ECS that is used to exchange test data with the tested ECS. The auxiliary ECS functions as the client (TX end) or server (RX end) in netperf tests.
- [Table 11-4](#) and [Table 11-5](#) list the common netperf and iperf3 parameters.

Table 11-4 Common netperf parameters

Parameter	Description
-p	Port number
-H	IP address of the RX end
-t	Protocol used in packet transmitting, the value of which is TCP_STREAM in bandwidth tests
-l	Test duration
-m	Data packet size, which is suggested to be 1440 in bandwidth tests

Table 11-5 Common iperf3 parameters

Parameter	Description
-p	Port number
-c	IP address of the RX end
-u	UDP packets
-b	TX bandwidth
-t	Test duration
-l	Data packet size, which is suggested to be 16 in PPS tests
-A	ID of the vCPU used by iperf3 In this section, the maximum number of 16 vCPUs is used as an example for each ECS. If an ECS has 8 vCPUs, the -A value ranges from 0 to 7.

Test Preparations

Step 1 Prepare ECSs.

Ensure that both type and specifications of the tested ECS and auxiliary ECSs are the same. In addition, ensure that these ECSs are deployed in the same ECS group with anti-affinity enabled.

Table 11-6 Preparations

Category	Quantity	Image	Specifications	IP Address
Tested ECS	1	CentOS 7.4 64bit (recommended)	At least eight vCPUs	192.168.2.10
Auxiliary ECS	8	CentOS 7.4 64bit (recommended)	At least 8 vCPUs	192.168.2.11-19 2.168.2.18

Step 2 Install the netperf, iperf3, and sar test tools on both the tested ECS and auxiliary ECSs.

Table 11-7 lists the procedures for installing these tools.

Table 11-7 Installing test tools

Tool	Procedure
netperf	<ol style="list-style-type: none">Run the following command to install gcc: yum -y install unzip gcc gcc-c++Run the following command to download the netperf installation package: wget https://github.com/HewlettPackard/netperf/archive/refs/tags/netperf-2.7.0.zipRun the following commands to decompress the installation package and install netperf: unzip netperf-2.7.0.zip cd netperf-netperf-2.7.0/ ./configure && make && make install
iperf3	<ol style="list-style-type: none">Run the following command to download the iperf3 installation package: wget --no-check-certificate https://codeload.github.com/esnet/iperf/zip/master -O iperf3.zipRun the following commands to decompress the installation package and install iperf3: unzip iperf3.zip cd iperf-master/ ./configure && make && make install
sar	Run the following command to install sar: yum -y install sysstat

Step 3 Enable NIC multi-queue.

Perform the following operations on both tested ECS and auxiliary ECSs.

- Run the following command to check the number of queues supported by the ECSs:

```
ethtool -l eth0 | grep -i Pre -A 5 | grep Combined
```

2. Run the following command to enable NIC multi-queue:

```
ethtool -L eth0 combined X
```

In the preceding command, *X* is the number of queues obtained in [Step 3.1](#).

----End

TCP Bandwidth Test (Using netperf)

Perform the test on multiple flows. This section considers 16 flows that are evenly distributed to eight ECSs, as an example.

NOTE

The TCP bandwidth test uses the multi-flow model.

- When testing the TCP transmission (TX) bandwidth, use the one-to-many model to ensure that the capability of the receiver is sufficient.
- When testing the TCP receiver (RX) bandwidth, use the many-to-one model to ensure that the capability of the sender is sufficient.

Step 1 Test the TCP TX bandwidth.

1. Run the following commands on all auxiliary ECSs to start the netserver process:

```
netserver -p 12001
```

```
netserver -p 12002
```

In the preceding commands, **-p** specifies the listening port.

2. Start the netperf process on the tested ECS and specify a netserver port for each auxiliary ECS. For details about common netperf parameters, see [Table 11-4](#).

##The IP address is for the first auxiliary ECS.

```
netperf -H 192.168.2.11 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.11 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the second auxiliary ECS.

```
netperf -H 192.168.2.12 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.12 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the third auxiliary ECS.

```
netperf -H 192.168.2.13 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.13 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the fourth auxiliary ECS.

```
netperf -H 192.168.2.14 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.14 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the fifth auxiliary ECS.

```
netperf -H 192.168.2.15 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.15 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the sixth auxiliary ECS.

```
netperf -H 192.168.2.16 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```



```
netperf -H 192.168.2.16 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the seventh auxiliary ECS.

```
netperf -H 192.168.2.17 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.17 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

##The IP address is for the eighth auxiliary ECS.

```
netperf -H 192.168.2.18 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.18 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

Step 2 Test the TCP RX bandwidth.

1. Start the netserver process on the tested ECS.

##The port number is for the first auxiliary ECS.

```
netserver -p 12001
```

```
netserver -p 12002
```

##The port number is for the second auxiliary ECS.

```
netserver -p 12003
```

```
netserver -p 12004
```

##The port number is for the third auxiliary ECS.

```
netserver -p 12005
```

```
netserver -p 12006
```

##The port number is for the fourth auxiliary ECS.

```
netserver -p 12007
```

```
netserver -p 12008
```

##The port number is for the fifth auxiliary ECS.

```
netserver -p 12009
```

```
netserver -p 12010
```

##The port number is for the sixth auxiliary ECS.

```
netserver -p 12011
```

```
netserver -p 12012
```

##The port number is for the seventh auxiliary ECS.

```
netserver -p 12013
```

```
netserver -p 12014
```

##The port number is for the eighth auxiliary ECS.

```
netserver -p 12015
```

```
netserver -p 12016
```

2. Start the netperf process on all auxiliary ECSs.

Log in to auxiliary ECS 1.

```
netperf -H 192.168.2.10 -p 12001 -t TCP_STREAM -l 300 -- -m 1440 &
```

```
netperf -H 192.168.2.10 -p 12002 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 2.

```
netperf -H 192.168.2.10 -p 12003 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12004 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 3.

```
netperf -H 192.168.2.10 -p 12005 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12006 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 4.

```
netperf -H 192.168.2.10 -p 12007 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12008 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 5.

```
netperf -H 192.168.2.10 -p 12009 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12010 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 6.

```
netperf -H 192.168.2.10 -p 12011 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12012 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 7.

```
netperf -H 192.168.2.10 -p 12013 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12014 -t TCP_STREAM -l 300 -- -m 1440 &
```

Log in to auxiliary ECS 8.

```
netperf -H 192.168.2.10 -p 12015 -t TCP_STREAM -l 300 -- -m 1440 &  
netperf -H 192.168.2.10 -p 12016 -t TCP_STREAM -l 300 -- -m 1440 &
```

Step 3 Analyze the test result.

After the test is complete, the output of the netperf process on one TX end is shown in [Figure 11-10](#). The final result is the sum of the test results of the netperf processes on all TX ends.

Figure 11-10 Output of the netperf process on one TX end

```
Recv Send Send  
Socket Socket Message Elapsed  
Size Size Size Time Throughput  
bytes bytes bytes secs. 10^6bits/sec  
TX buffer Test duration Throughput  
87380 16384 1440 120.02 956.30  
RX buffer Data packet size
```

 NOTE

There are a large number of netperf processes. To facilitate statistics collection, it is a good practice to run the following command to view test data on the tested ECS using sar:

```
sar -n DEV 1 60
```

----End

UDP PPS Test (Using iperf3)

Step 1 Test the UDP TX PPS.

1. Log in to an auxiliary ECS.
2. Run the following commands on all auxiliary ECSs to start the server process:

```
iperf3 -s -p 12001 &
```

```
iperf3 -s -p 12002 &
```

In the preceding commands, **-p** specifies the listening port.

3. Start the client process on the tested ECS. For details about common iperf3 parameters, see [Table 11-5](#).

```
##Auxiliary ECS 1
```

```
iperf3 -c 192.168.2.11 -p 12001 -u -b 100M -t 300 -l 16 -A 0 &
```

```
iperf3 -c 192.168.2.11 -p 12002 -u -b 100M -t 300 -l 16 -A 1 &
```

```
##Auxiliary ECS 2
```

```
iperf3 -c 192.168.2.12 -p 12001 -u -b 100M -t 300 -l 16 -A 2 &
```

```
iperf3 -c 192.168.2.12 -p 12002 -u -b 100M -t 300 -l 16 -A 3 &
```

```
##Auxiliary ECS 3
```

```
iperf3 -c 192.168.2.13 -p 12001 -u -b 100M -t 300 -l 16 -A 4 &
```

```
iperf3 -c 192.168.2.13 -p 12002 -u -b 100M -t 300 -l 16 -A 5 &
```

```
##Auxiliary ECS 4
```

```
iperf3 -c 192.168.2.14 -p 12001 -u -b 100M -t 300 -l 16 -A 6 &
```

```
iperf3 -c 192.168.2.14 -p 12002 -u -b 100M -t 300 -l 16 -A 7 &
```

```
##Auxiliary ECS 5
```

```
iperf3 -c 192.168.2.15 -p 12001 -u -b 100M -t 300 -l 16 -A 8 &
```

```
iperf3 -c 192.168.2.15 -p 12002 -u -b 100M -t 300 -l 16 -A 9 &
```

```
##Auxiliary ECS 6
```

```
iperf3 -c 192.168.2.16 -p 12001 -u -b 100M -t 300 -l 16 -A 10 &
```

```
iperf3 -c 192.168.2.16 -p 12002 -u -b 100M -t 300 -l 16 -A 11 &
```

```
##Auxiliary ECS 7
```

```
iperf3 -c 192.168.2.17 -p 12001 -u -b 100M -t 300 -l 16 -A 12 &
```

```
iperf3 -c 192.168.2.17 -p 12002 -u -b 100M -t 300 -l 16 -A 13 &
```

```
##Auxiliary ECS 8
```

```
iperf3 -c 192.168.2.18 -p 12001 -u -b 100M -t 300 -l 16 -A 14 &
```

```
iperf3 -c 192.168.2.18 -p 12002 -u -b 100M -t 300 -l 16 -A 15 &
```

Step 2 Test the UDP RX PPS.

1. Start the server process on the tested ECS. For details about common iperf3 parameters, see [Table 11-5](#).

```
##The port number is for the first auxiliary ECS.
```

```
iperf3 -s -p 12001 -A 0 -i 60 &
```

```
iperf3 -s -p 12002 -A 1 -i 60 &
```

```
##The port number is for the second auxiliary ECS.
```

```
iperf3 -s -p 12003 -A 2 -i 60 &
```

```
iperf3 -s -p 12004 -A 3 -i 60 &
```

```
##The port number is for the third auxiliary ECS.
```

```
iperf3 -s -p 12005 -A 4 -i 60 &
```

```
iperf3 -s -p 12006 -A 5 -i 60 &
```

```
##The port number is for the fourth auxiliary ECS.
```

```
iperf3 -s -p 12007 -A 6 -i 60 &
```

```
iperf3 -s -p 12008 -A 7 -i 60 &
```

```
##The port number is for the fifth auxiliary ECS.
```

```
iperf3 -s -p 12009 -A 8 -i 60 &
```

```
iperf3 -s -p 12010 -A 9 -i 60 &
```

```
##The port number is for the sixth auxiliary ECS.
```

```
iperf3 -s -p 12011 -A 10 -i 60 &
```

```
iperf3 -s -p 12012 -A 11 -i 60 &
```

```
##The port number is for the seventh auxiliary ECS.
```

```
iperf3 -s -p 12013 -A 12 -i 60 &
```

```
iperf3 -s -p 12014 -A 13 -i 60 &
```

```
##The port number is for the eighth auxiliary ECS.
```

```
iperf3 -s -p 12015 -A 14 -i 60 &
```

```
iperf3 -s -p 12016 -A 15 -i 60 &
```

2. Start the client process on all auxiliary ECSs. For details about common iperf3 parameters, see [Table 11-5](#).

```
Log in to auxiliary ECS 1.
```

```
iperf3 -c 192.168.2.10 -p 12001 -u -b 100M -t 300 -l 16 -A 0 &
```

```
iperf3 -c 192.168.2.10 -p 12002 -u -b 100M -t 300 -l 16 -A 1 &
```

```
Log in to auxiliary ECS 2.
```

```
iperf3 -c 192.168.2.10 -p 12003 -u -b 100M -t 300 -l 16 -A 0 &
```

iperf3 -c 192.168.2.10 -p 12004 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 3.

iperf3 -c 192.168.2.10 -p 12005 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12006 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 4.

iperf3 -c 192.168.2.10 -p 12007 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12008 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 5.

iperf3 -c 192.168.2.10 -p 12009 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12010 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 6.

iperf3 -c 192.168.2.10 -p 12011 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12012 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 7.


iperf3 -c 192.168.2.10 -p 12013 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12014 -u -b 100M -t 300 -l 16 -A 1 &
Log in to auxiliary ECS 8.

iperf3 -c 192.168.2.10 -p 12015 -u -b 100M -t 300 -l 16 -A 0 &
iperf3 -c 192.168.2.10 -p 12016 -u -b 100M -t 300 -l 16 -A 1 &

Step 3 Analyze the test result.

Figure 11-11 shows an example of the UDP PPS test result.

Figure 11-11 UDP PPS test result



NOTE

There are a large number of iperf3 processes. To facilitate statistics collection, it is a good practice to run the following command to view test data on the tested ECS using sar:

```
sar -n DEV 1 60
```

----End

Latency Test

Step 1 Run the following command to start the qperf process on the tested ECS:

```
qperf &
```

Step 2 Log in to auxiliary ECS 1 and run the following command to perform a latency test:

```
iperf 192.168.2.10 -m 64 -t 60 -vu udp_lat
```

After the test is complete, the **lat** value in the command output is the latency between ECSs.

----End

11.13 What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?

Symptom

Take a Linux ECS as an example. After the user modified ECS specifications and ran the **ifconfig** command, the user found that the original eth0 and eth1 NICs were changed to eth2 and eth3 NICs, indicating that NIC flapping occurred.

Root Cause

NIC flapping occurs because NIC retaining is enabled in the image from which the ECS is created.

Solution to Linux

For a Linux ECS, perform the following operations and restart the ECS to resolve this issue:

1. Run the following command to view the files in the network rule directory:

```
ls -l /etc/udev/rules.d
```
2. Run the following commands to delete the files with both **persistent** and **net** included in file names from the network rule directory:

```
rm -fr /etc/udev/rules.d/*net*persistent*.rules  
rm -fr /etc/udev/rules.d/*persistent*net*.rules
```
3. Run the following command to check whether the initrd image file with a name starting with **initrd** and ending with **default** contains both **persistent** and **net** network rules (change the italic data in the following command to the actual OS version):

```
lsinitrd /boot/initrd-2.6.32.12-0.7-default |grep persistent|grep net
```

 - If yes, go to steps [4](#) and [5](#).
 - If no, no further action is required.
4. Run the following command to back up the initrd image file (change the italic data in the following command to the actual OS version):

```
cp /boot/initrd-2.6.32.12-0.7-default /boot/initrd-2.6.32.12-0.7-default_bak
```
5. Run the following command to regenerate the initrd image file:

```
mkinitrd
```

Perform the following operations when an OS, such as Ubuntu, uses the initramfs image:

1. Run the following command to check whether the initramfs image file with a name starting with **initrd** and ending with **generic** contains both **persistent** and **net** network rules:

```
lsinitramfs /boot/initrd.img-3.19.0-25-generic|grep persistent|grep net
```

- If yes, go to steps [2](#) and [3](#).
- If no, no further action is required.

2. Run the following command to back up the initrd image file:

```
cp /boot/initrd.img-3.19.0-25-generic /boot/initrd.img-3.19.0-25-generic_bak
```

3. Run the following command to regenerate the initramfs image file:

```
update-initramfs -u
```

11.14 Will NICs Added to an ECS Start Automatically?

Based on test results, if the ECS runs CentOS 7.0, NICs added to the ECS cannot start automatically. You must start the NICs manually.

11.15 How Can I Check Whether the Network Communication Is Normal Between Two ECSs Equipped with an InfiniBand NIC Driver?

For high-performance H2 ECSs equipped with an InfiniBand NIC driver (InfiniBand ECSs for short), perform the following operations to check whether the driver installation is successful and whether the network communication between the ECSs is normal.

NOTE

During the check, if your ECS has no command tool installed, such as `ibstat`, obtain the tool from the installation package for the InfiniBand NIC driver and install the tool.

Step 1 Check whether the NICs of the InfiniBand ECSs are functional.

1. Log in to the ECS.
2. Run the following command to check whether the NIC is functional:
ibstat
 - If it is functional, go to [Step 2](#).
 - If it is not functional, contact customer service for technical support.

Step 2 Check whether the network communication between two InfiniBand ECSs is normal.

1. Log in to one InfiniBand ECS and run the following command:
ib_write_bw -x 0 --pkey_index 0
2. Log in to the other InfiniBand ECS and run the following command:
ib_write_bw -x 0 --pkey_index 0 ip_addr

In the preceding command, *ip_addr* is the NIC IP address of the first InfiniBand ECS.

3. Check whether the terminal display is correct.

Figure 11-12 Normal network communication

```
[root@host-11-11-11-111 MUX_OFED_LINUX-3.4-1.0.0-rhel7.2-x86_64]# ib_write_bw -x 0 --pkey_index 0 4.29.43.20
-----
RDMA_Write BW Test
-----
Dual_port      : OFF          Device      : mlx5_0
Number of qps  : 1           Transport type : IB
Connection type : RC         Using SRQ    : OFF
TX depth       : 128
CQ Moderation  : 100
Mtu            : 4096[B]
Link type      : IB
GID index      : 0
Max inline data : 6[B]
rdma_cm QPs    : OFF
Data ex. method : Ethernet
-----
Local address: LID 0x05 QPN 0x0067 PSN 0xaaccfb RKey 0x001c0c VAddr 0x007fb3cd1b0000
GID: 254:128:00:00:00:00:00:00:03:00:135:40:178
remote address: LID 0x05 QPN 0x006a PSN 0xebbf6d RKey 0x001c10 VAddr 0x007fdad5990000
GID: 254:128:00:00:00:00:00:00:01:03:00:135:40:178
-----
#bytes  #iterations  BW peak[MB/sec]  BW average[MB/sec]  MsgRate[Mpps]
65536   5000         12132.70         11900.18            0.190403
```

- If the terminal display is shown in [Figure 11-12](#), the network communication between the two InfiniBand ECSs is normal.
- If the InfiniBand network is inaccessible, contact customer service for technical support.

----End

11.16 How Can I Manually Configure an IP Address for an InfiniBand NIC?

IP over InfiniBand (IPoIB) allows IP data transmission over InfiniBand. For SUSE high-performance H2 and HL1 ECSs, if IPoIB is required, you must manually configure an IP address for the InfiniBand NIC after installing the InfiniBand NIC driver.

Prerequisites

The InfiniBand NIC driver has been installed on the high-performance H2 or HL1 ECSs.

Background

To prevent IP address conflict of the InfiniBand NICs configured for the ECSs of a tenant, determine the IP address to be configured for an InfiniBand NIC according to the IP addresses available in the VPC. The method is as follows:


For example, if the first two eight-bits of the IP address (specified by **IPADDR**) to be configured for the InfiniBand NIC are consistently **169.254**, the latter two eight-bits must be the same as those of the **eth0** IP address, and the subnet mask must be the same as that of the **eth0** NIC.

An example is provided as follows:

If the IP address of the **eth0** NIC is 192.168.0.100/24, the IP address to be configured for the InfiniBand NIC is 169.254.0.100/24.

Procedure

1. Log in to the ECS.
2. Run the following command to switch to user **root**:
sudo su -
3. Run the following command to edit the **/etc/sysconfig/network/ifcfg-ib0** file:
vi /etc/sysconfig/network/ifcfg-ib0
4. Enter the following information:
DEVICE=ib0
BOOTPROTO=static
IPADDR=IP address to be configured for the InfiniBand NIC
NETMASK=Subnet mask
STARTMODE=auto

 **NOTE**

For instructions about how to obtain the IP address and subnet mask for an InfiniBand NIC, see [Background](#).
5. Run the following command to restart the network for the configuration to take effect:
service network restart

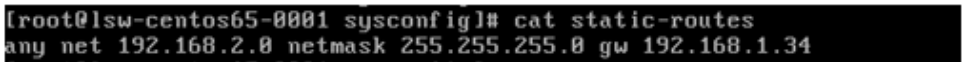
11.17 How Can I Add a Static Route to a CentOS 6.5 OS?

Scenarios

After the system restarts, non-static routes are lost, affecting network availability. To prevent this issue from occurring, you must add static routes to the system.

Procedure

The following section uses a CentOS 6.5 OS as an example.

1. Log in to the ECS.
2. Create or modify the static route configuration file.
If the **static-routes** configuration file is not in the **/etc/sysconfig/** directory, create this file. If such a file is available, run the following command to add a static route into this file:
any net 192.168.2.0 netmask 255.255.255.0 gw 192.168.1.34
After the configuration, save and exit the file. The following figure shows the modified file content.

3. Run the following command to restart the network service to make the static route take effect:

service network restart

```
[root@lsw-centos65-0001 sysconfig]# service network restart
Shutting down interface eth0: [ OK ]
Shutting down loopback interface: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0:
Determining IP information for eth0... done. [ OK ]
```

4. Run the following command to view routes:

route -n

```
[root@lsw-centos65-0001 sysconfig]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
169.254.169.254 192.168.1.1 255.255.255.255 UGH 0 0 0 eth0
192.168.2.0 192.168.1.34 255.255.255.0 UG 0 0 0 eth0
192.168.1.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 1002 0 0 eth0
0.0.0.0 192.168.1.1 0.0.0.0 UG 0 0 0 eth0
```

12 EIP FAQ

12.1 Can Multiple EIPs Be Bound to an ECS?

Scenarios

Multiple EIPs can be bound to an ECS, but this operation is not recommended.

If an ECS has multiple NICs attached and you want to bind multiple EIPs to this ECS, you need to configure policy-based routes for these NICs so that these extension NICs can communicate with external works. For details, see [Configuration Example](#).

Configuration Example

[Table 12-1](#) lists ECS configurations.

Table 12-1 ECS configurations

Parameter	Configuration
Name	ecs_test
Image	CentOS 6.5 64bit
EIP	2
Primary NIC	eth0
Secondary NIC	eth1

Example 1:

If you intend to access public network 11.11.11.0/24 through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.

2. Run the following command to configure a route:

```
ip route add 11.11.11.0/24 dev eth1 via 192.168.2.1
```

In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

Example 2:

Based on example 1, if you intend to enable routing for default public network traffic through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.
2. Run the following command to delete the default route:

```
ip route delete default
```

NOTICE

Exercise caution when deleting the default route because this operation will interrupt the network and result in SSH login failures.

3. Run the following command to configure a new default route:

```
ip route add 0.0.0.0/0 dev eth1 via 192.168.2.1
```

In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

12.2 Can an ECS Without an EIP Bound Access the Internet?

Yes.

You can use the NAT Gateway service to allow ECSs in a VPC to access the Internet using an EIP. The SNAT function provided by the NAT Gateway service allows the ECSs in a VPC to access the Internet without requiring an EIP. Additionally, SNAT supports a large number of concurrent connections for applications that have a large number of requests and connections. For more information about NAT Gateway, see *NAT Gateway Service Overview*.

12.3 What Should I Do If an EIP Cannot Be Pinged?

Symptom

After you purchase an EIP and bind it to an ECS, the local host or other cloud servers cannot ping the EIP of the ECS.

Fault Locating

The following fault causes are sequenced based on their occurrence probability.

If the fault persists after you have ruled out a cause, check other causes.

Table 12-2 Method of locating the failure to ping an EIP

Possible Causes	Solution
ICMP access rules are not added to the security group.	Add ICMP access rules to the security group. For details, see Checking Security Group Rules .
Ping operations are prohibited on the firewall.	Allow ping operations on the firewall. For details, see Checking Firewall Settings .
Ping operations are prohibited on the ECS.	Allow ping operations on the ECS. For details, see Checking Whether Ping Operations Have Been Disabled on the ECS .
Network ACL is associated.	If the VPC is associated with a network ACL, check the network ACL rules. For details, see Checking ACL Rules .
A network exception occurred.	Use another ECS in the same region to check whether the local network is functional. For details, see Checking Whether the Network Is Functional .
Routes are incorrectly configured if multiple NICs are used.	If the network is inaccessible due to an extension NIC, the fault is generally caused by incorrect route configurations. To resolve this issue, see Checking the ECS Route Configuration If Multiple NICs Are Used .
The domain name is not ICP licensed.	If the domain name cannot be pinged or cannot be resolved, see Checking Domain Name Resolution If the Domain Name Cannot Be Pinged to resolve this issue.

Checking Security Group Rules

ICMP is used for the ping command. Check whether the security group accommodating the ECS allows ICMP traffic.



1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . 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, expand the information of the security group, and view security group rules.
6. Click the security group ID.
The system automatically switches to the **Security Group** page.
7. On the **Outbound Rules** page, click **Add Rule**. In the displayed dialog box, set required parameters to add an outbound rule.

Table 12-3 Security group rules

Transfer Direction	Type	Protocol/Port Range	Destination
Outbound	IPv4	ICMP/Any	0.0.0.0/0 0.0.0.0/0 indicates all IP addresses.

8. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.

Table 12-4 Security group rules

Transfer Direction	Type	Protocol/Port Range	Source
Inbound	IPv4	ICMP/Any	0.0.0.0/0 0.0.0.0/0 indicates all IP addresses.

9. Click **OK** to complete the security rule configuration.

Checking Firewall Settings

If a firewall is enabled on the ECS, check whether the firewall blocks the ping operations.

Linux

1. Consider CentOS 7 as an example. Run the following command to check the firewall status:

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

If **running** is displayed in the command output, the firewall has been enabled.

2. Check whether there is any ICMP rule blocking the ping operations.

```
iptables -L
```

If the command output shown in [Figure 12-1](#) is displayed, there is no ICMP rule blocking the ping operations.

Figure 12-1 Checking firewall rules

```
[root@ecs-3c4e ~]# iptables -L
Chain INPUT (policy ACCEPT)
target     prot opt source                destination
ACCEPT    icmp -- anywhere             anywhere             icmp echo-request

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
ACCEPT    icmp -- anywhere             anywhere             icmp echo-reply
[root@ecs-3c4e ~]#
```

If the ping operations are blocked by an ICMP rule, run the following commands to modify the rule for unblocking:

```
iptables -A INPUT -p icmp --icmp-type echo-request -j ACCEPT
iptables -A OUTPUT -p icmp --icmp-type echo-reply -j ACCEPT
```

Checking Whether Ping Operations Have Been Disabled on the ECS

Linux

Check the ECS kernel parameters.

1. Check the `net.ipv4.icmp_echo_ignore_all` value in the `/etc/sysctl.conf` file. Value **0** indicates that ping operations are allowed, and value **1** indicates that ping operations are prohibited.
2. Allow ping operations.
 - Run the following command to temporarily allow the ping operations:
`#echo 0 >/proc/sys/net/ipv4/icmp_echo_ignore_all`
 - Run the following command to permanently allow the ping operations:
`net.ipv4.icmp_echo_ignore_all=0`

Checking ACL Rules

By default, no ACL is configured for a VPC. If a network ACL is associated with a VPC, check the ACL rules.

1. Check whether the subnet of the ECS has been associated with a network ACL.
If an ACL name is displayed, the network ACL has been associated with the ECS.
2. Click the ACL name to view its status.
3. If the network ACL is enabled, add an ICMP rule to allow traffic.

NOTE

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

Checking Whether the Network Is Functional

1. Use another ECS in the same region to check whether the local network is functional.
Use another ECS in the same region to ping the affected EIP. If the EIP can be pinged, the VPC is functional. In such a case, rectify the local network fault and ping the affected EIP again.
2. Check whether the link is accessible.
A ping failure is caused by packet loss or long delay, which may be caused by link congestion, link node faults, or heavy load on the ECS.

Checking the ECS Route Configuration If Multiple NICs Are Used

Generally, the default route of an OS will preferentially select the primary NIC. If an extension NIC is selected in a route and the network malfunctions, this issue is typically caused by incorrect route configuration.

- If the ECS has multiple NICs, check whether the default route is available.
 - a. Log in to the ECS and run the following command to check whether the default route is available:

ip route

Figure 12-2 Default route

```
[root@do-not-del-secy ~]# ip route
default via 192.168.2.1 dev eth0
169.254.0.0/16 dev eth0 scope link metric 1002
169.254.169.254 via 192.168.2.1 dev eth0 proto static
192.168.2.0/24 dev eth0 proto kernel scope link src 192.168.2.112
```

- b. If the route is unavailable, run the following command to add it:

ip route add default via XXXX dev eth0

NOTE

In the preceding command, *XXXX* specifies a gateway IP address.

- If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policy routing on the ECS for network communication with the extension NIC.

For details, see [How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?](#)

Checking Domain Name Resolution If the Domain Name Cannot Be Pinged

If you can ping the EIP but not the domain name, the possible cause is that an error occurred in domain name resolution.

1. Check the domain name resolution.

If the domain name records are incorrectly configured, the domain name may fail to be resolved.

Switch to the DNS management console to view details about the domain name resolution.
2. Check the DNS server configuration.

12.4 Why Can I Remotely Access an ECS But Cannot Ping It?

Symptom

You can remotely access an ECS but when you ping the EIP bound to the ECS, the ping operation fails.

Possible Causes

A desired inbound rule is not added for the security group, and ICMP is not enabled.

Solution

1. Log in to the management console.
2. Under **Compute**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
4. Click the **Security Groups** tab, expand the information of the security group, and click the security group ID.
5. On the **Inbound Rules** tab of the **Security Group** page, click **Add Rule**.
6. Add an inbound rule for the security group and enable ICMP.
 - **Protocol: ICMP**
 - **Source: IP address 0.0.0.0/0**

12.5 Why Is Outbound Access on TCP Port 25 Blocked?

Symptom

You cannot access an external address on TCP port 25. For example, running the **Telnet smtp.***.com 25** command fails.

Cause

For security reasons, TCP port 25 is disabled in the outbound direction by default.

You do not need to enable TCP port 25, unless you want to deploy an email service on the cloud.

Solution

Use port 465 supported by the third-party email service provider.

12.6 Why My EIPs Are Frozen? How Do I Unfreeze My EIPs?

EIP resources are frozen in the following situations:

- **In arrears**
 - Yearly/Monthly EIPs
If you do not renew yearly/monthly EIPs after the grace period ends, the EIPs enter a retention period and are frozen. Frozen EIPs cannot be used, modified, or released. If you still do not renew your EIPs before the retention period ends, they will be released and cannot be restored. To ensure the availability of your EIPs, renew them before they expire.

- Pay-per-use EIPs

If your pay-per-use EIPs are still in arrears after the grace period ends, the EIPs enter the retention period and are frozen. Frozen EIPs cannot be used, modified, or released. If you still do not top up and pay off the arrears before the retention period ends, the EIPs will be released and cannot be restored. To ensure the availability of your EIPs, top up your account and pay off the arrears before they expire.

- Frozen EIPs will be available after you renew them or top up your account. You can renew your resources on the management console. For more details, see [Renewal Management](#).

- **Attacks**

EIPs will be frozen if their associated instances have security violations, such as attacks. Frozen EIPs are unavailable and cannot be modified or released.

You can change an EIP for an instance by referring to [Changing an EIP](#).

- **Violations**

The server bound to the EIP is suspected of violations and the EIP is frozen by the national supervision department. If you have confirmed that you have not been involved in any violation, contact the national supervision department to file an appeal. If the appeal is successful, Huawei Cloud will receive an unsealing instruction to unfreeze your resources.

You can change an EIP for an instance by referring to [Changing an EIP](#).

13 Password and Key Pair FAQ

13.1 How Can I Change the Password for Logging In to a Linux ECS?

Solution

1. Use the existing key file to log in to the Linux ECS as user **root**.
2. Run the following command to reset the password of user **root**:

passwd

To reset the password of another user, replace **passwd** with **passwd username**.

3. Enter the new password as prompted.

New password:

Retype new password:

If the following information is displayed, the password has been reset:
passwd: all authentication tokens updates successfully

13.2 What Is the Default Password for Logging In to a Linux ECS?

The default username for logging in to an ECS running Linux (such as CentOS or Ubuntu) is **root**, and the password is the one you set during ECS creation.

13.3 How Can I Set the Validity Period of the Image Password?

If an ECS cannot be logged in because of expired image password, you can contact the administrator for handling.

If the ECS can still be logged in, you can perform the following operations to set the password validity period.

Procedure

The following operations use EulerOS 2.2 as an example.

1. Log in to the ECS.
2. Run the following command to check the password validity period:

```
vi /etc/login.defs
```

The value of parameter **PASS_MAX_DAYS** is the password validity period.

3. Run the following command to change the value of parameter **PASS_MAX_DAYS**:

```
chage -M 99999 user_name
```

99999 is the password validity period, and *user_name* is the system user, for example, user **root**.

NOTE

You are advised to configure the password validity period as needed and change it at a regular basis.

4. Run command **vi /etc/login.defs** to verify that the configuration has taken effect.

Figure 13-1 Configuration verification

```
# Password aging controls:
#
#     PASS_MAX_DAYS   Maximum number of days a password may be used.
#     PASS_MIN_DAYS   Minimum number of days allowed between password changes.
#     PASS_MIN_LEN    Minimum acceptable password length.
#     PASS_WARN_AGE   Number of days warning given before a password expires.
#
PASS_MAX_DAYS 99999
PASS_MIN_DAYS 0
PASS_MIN_LEN 5
PASS_WARN_AGE 7
```

13.4 Resetting the Password for Logging In to an ECS in the OS

Scenarios

This section describes how to reset the password for logging in to an ECS in the OS when the password is about to expire, the password is forgotten, or you are logging in to the ECS for the first time. It is a good practice to change the initial password upon the first login.

Prerequisites

The ECS can be logged in.

Background

Table 13-1 shows the ECS password complexity requirements.

Table 13-1 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• Contains at least three of the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters for Linux ECSs: !@%-_=#+[]:./^,{}?• Cannot contain the username or the username spelled backwards.

Linux

1. Use the existing key file to log in to the ECS as user **root** through SSH.
For details, see [Login Using an SSH Key](#).

2. Run the following command to reset the password of user **root**:

```
passwd
```

To reset the password of another user, replace **passwd** with **passwd username**.

3. Enter a new password that meets the requirements listed in [Table 13-1](#) as prompted.

```
New password:
```

```
Retype new password:
```

If the following information is displayed, the password has been changed:
passwd: all authentication tokens updates successfully

13.5 What Should I Do If the System Displays a Message Indicating that the Password Is Incorrect When I Remotely Log In to My ECS?

Solution

Check the network configuration of the ECS and determine whether the fault is caused by a [Cloud-Init](#) failure.

- Verify that port 80 is bypassed in both inbound and outbound directions in the security group to which the target ECS belongs.
- Verify that DHCP is enabled in the subnet to which the target ECS belongs.

NOTE

After verifying the preceding configurations, restart the ECS, wait for 3 to 5 minutes, and remotely log in to the ECS using a password or key.

13.6 What Should I Do If I Cannot Log In to My ECS Using the Initial Password After I Use It for a Period of Time?

Solution

Check whether the remote login page can be displayed.

- If the login page cannot be displayed, an error may have occurred in the GuestOS process on the ECS. In such a case, contact customer service for troubleshooting.
- If the login page can be displayed, log in to the OS in single-user mode for troubleshooting. The procedure is as follows:
 - Check whether the password can be changed in single-user mode.
If the password can be changed, change it and contact customer service to check whether the password has been maliciously changed due to an attack.
 - If the password cannot be changed, verify that the values of **hard** and **soft** in **/etc/security/limits.conf** are not greater than 65535.

```
# - nice - max nice priority allowed to raise to values: [-20, 19]
# - rtprio - max realtime priority
#
#<domain> <type> <item> <value>
#
#* soft core 0
#* hard rss 10000
#@student hard nproc 20
#@faculty soft nproc 20
#@faculty hard nproc 50
#@ftp hard nproc 0
#@student - maxlogins 4
# End of file
```

Change the password in single-user mode and try to log in to the ECS again.

13.7 Disabling SELinux

NOTE

SUSE does not have the SELinux configuration files. You can skip this section.

Procedure

1. Use the vi editor to open **/etc/selinux/config**.
vi /etc/selinux/config
2. Press **i** to enter insert mode and set the value of **SELINUX** to **disabled**.

```
# This file controls the state of SELinux on the system.
# SELINUX- can take one of these three values:
#   enforcing - SELinux security policy is enforced.
#   permissive - SELinux prints warnings instead of enforcing.
#   disabled - No SELinux policy is loaded.
SELINUX=disabled
# SELINUXTYPE- can take one of three two values:
#   targeted - Targeted processes are protected.
#   minimum - Modification of targeted policy. Only selected processes
#   mls - Multi Level Security protection.
SELINUXTYPE=targeted
```


3. Press **Esc** and enter **:wq** to save and exit the file.

13.8 How Can I Obtain the Key Pair Used by My ECS?

Symptom

You have created multiple key pairs, and you are trying to find the key pair to log in to the target ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. Click the name of the target ECS.

The page providing details about the ECS is displayed.

6. Obtain the **Key Pair** value.

The value is the key pair used by the ECS.

13.9 How Can I Use a Key Pair?

Symptom

When you purchase an ECS, the system asks you to select a login mode. If you select **Key pair**, you are required to select an existing key pair or create a new pair.

If no key pair is available, create one on the management console.

Solution

1. In the navigation pane of the ECS console, choose **Key Pair**. Then, click **Create Key Pair**.
2. After the key pair is created, download the private key to a local directory.
3. When purchasing an ECS, select the created or existing key pair in **Key pair**.


13.10 Can I Download a Key Pair from My Phone?

No. This operation is not supported.

You can download the private key file only once when creating an ECS.

13.11 What Should I Do If a Key Pair Cannot Be Imported?

If you use Internet Explorer 9 to access the management console, the key pair may fail to import. In this case, perform the following steps to modify browser settings and then try again:

1. Click  in the upper right corner of the browser.
2. Select **Internet Options**.
3. Click the **Security** tab in the displayed dialog box.
4. Click **Internet**.
5. If the security level indicates **Custom**, click **Default Level** to restore to the default settings.
6. Move the scroll bar to set the security level to **Medium** and click **Apply**.
7. Click **Custom Level**.
8. Set **Initialize and script ActiveX controls not marked as safe for scripting** to **Prompt**.
9. Click **Yes**.

13.12 Why Does the Login to My Linux ECS Using a Key File Fail?

Symptom

When you use the key file created during your Linux ECS creation to log in to the ECS, the login fails.

Possible Causes

Possible causes vary depending on the image used to create the Linux ECS.

- Cause 1: The image that you used to create the Linux ECS is a private image, on which Cloud-Init is not installed.
- Cause 2: Cloud-Init is installed on the image, but you did not obtain the key pair when you created the ECS.

Solution

- If the issue is a result of cause 1, proceed as follows:

If you created a private image without installing Cloud-Init, you cannot customize the ECS configuration. As a result, you can log in to the ECS only using the original image password or key pair.

The original image password or key pair is the OS password or key pair you configured when you created the private image.

If you have forgotten the original image password or the key pair is lost, reset the password on the ECS console.

- If the issue is a result of cause 2, proceed as follows:
 - a. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.
 - b. Use the key file to log in to the ECS again and check whether the login is successful.
 - If the login is successful, no further action is required.
 - If the login fails, contact customer service for technical support.

13.13 What Should I Do If I Cannot Download a Key Pair?

The private key file of a key pair can be downloaded only once.

If your private key file has been lost, create a key pair and download the private key file again.

Solution

1. Log in to the management console and choose **Key Pair**.
2. Click **Create Key Pair**.
3. Click **OK** to save the private key to your local directory.

13.14 Why Does a Key Pair Created Using `puttygen.exe` Fail to Be Imported on the Management Console?

Symptom

When you try to import a key pair that you created using **puttygen.exe** on the management console, the system displays a message indicating that the import failed.

Possible Causes

The format of the public key content does not meet system requirements.

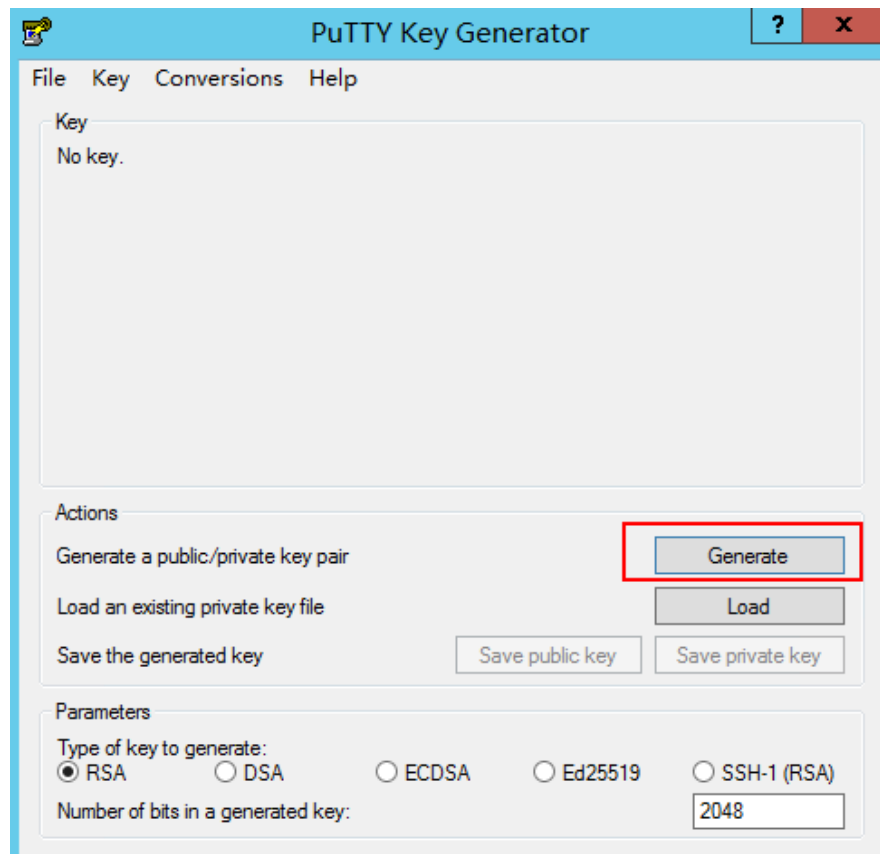
If you store a public key by clicking **Save public key** on PuTTY Key Generator, the format of the public key content will change. You cannot import the key on the management console.

Solution

Use the locally stored private key and **PuTTY Key Generator** to restore the format of the public key content. Then, import the public key to the management console.

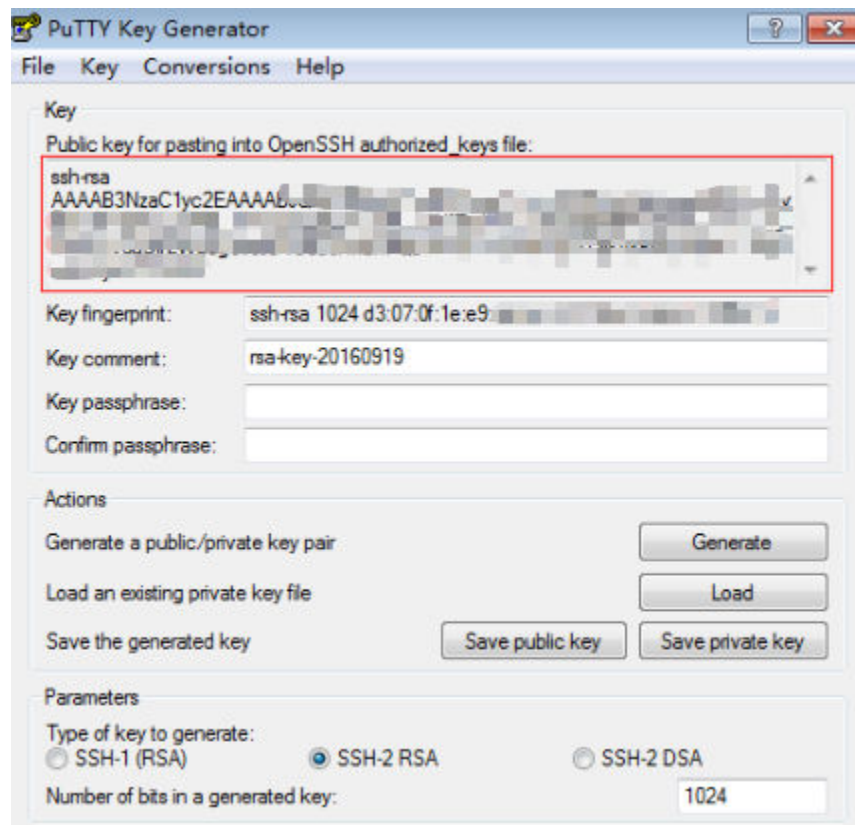
1. Double-click **puttygen.exe** to open **PuTTY Key Generator**.


Figure 13-2 PuTTY Key Generator



2. Click **Load** and select the private key.

The system automatically loads the private key and restores the format of the public key content in **PuTTY Key Generator**. The content in the red box in **Figure 13-3** is the public key whose format meets system requirements.

Figure 13-3 Restoring the format of the public key content

3. Copy the public key content to a .txt file and save the file in a local directory.
4. Import the public key to the management console.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select your region and project.
 - c. Under **Compute**, click **Elastic Cloud Server**.
 - d. In the navigation pane on the left, choose **Key Pair**.
 - e. On the key pair page, click **Import Key Pair**.
 - f. Copy the public key content in the .txt file to **Public Key Content** and click **OK**.

13.15 What Should I Do If Cloud-Init Does Not Work After Python Is Upgraded?

Symptom

Take an ECS running CentOS 6.8 as an example. After Python was upgraded from 2.6 to 2.7, Cloud-Init did not work. Data, such as the login password, key, and hostname could not be imported to the ECS using Cloud-Init.

After the **cloud-init -v** command was executed to view the Cloud-Init version, the system displayed errors, as shown in [Figure 13-4](#).

Figure 13-4 Improper running of Cloud-Init

```
[root@ecs-8560 ~]# cloud-init -v
Traceback (most recent call last):
  File "/usr/bin/cloud-init", line 39, in <module>
    from cloudinit import patcher
ImportError: No module named cloudinit
[root@ecs-8560 ~]# cloud-init init --local
Traceback (most recent call last):
  File "/usr/bin/cloud-init", line 39, in <module>
    from cloudinit import patcher
ImportError: No module named cloudinit
[root@ecs-8560 ~]#
```

Possible Causes

The Python version used by Cloud-Init was incorrect.

Solution

Change the Python version used by Cloud-Init to the source version. To do so, change the environment variable value of `/usr/bin/cloud-init` from the default value `#!/usr/bin/python` to `#!/usr/bin/python2.6`.

Figure 13-5 Changing the Python version

```
[root@ecs-8560 ~]# head -n 1 /usr/bin/cloud-init
#!/usr/bin/python2.6
[root@ecs-8560 ~]# ls /usr/bin/python* -lh
lrwxrwxrwx 1 root root 24 Jul 19 10:55 /usr/bin/python -> /usr/local/bin/python2.7
lrwxrwxrwx 1 root root 6 Jun 9 2017 /usr/bin/python2 -> python
-rwxr-xr-x 1 root root 8.9K Aug 18 2016 /usr/bin/python2.6
```

14 Application Deployment and Software Installation FAQ

14.1 Can a Database Be Deployed on an ECS?

Yes. You can deploy a database of any type on an ECS.

14.2 Does an ECS Support Oracle Databases?

Yes. You are advised to perform a performance test beforehand to ensure that the Oracle database can meet your requirements.

15 File Upload/Data Transfer FAQ

15.1 How Do I Upload Files to My ECS?

Linux

- From a local Windows computer
Use WinSCP to transfer the files to the Linux ECS. For details, see [How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?](#)
Before transferring files from a local computer to a Linux ECS, set up an FTP site on the ECS and install FileZilla on the local computer. For details, see [How Can I Use FTP to Transfer Files from a Local Windows Computer to a Linux ECS?](#)
- From a local Linux computer
Use SCP to transfer the files to the Linux ECS. For details, see [How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)
Use SFTP to transfer the files to the Linux ECS. For details, see [How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)
Use FTP to transfer the files to the Linux ECS. For details, see [How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?](#)

Does an ECS Support FTP-based File Transferring by Default?

No. You need to install and configure FTP so that the ECS supports FTP-based file transfer.

15.2 How Can I Use WinSCP to Transfer Files from a Local Windows Computer to a Linux ECS?

Scenarios

WinSCP can be used to securely copy-paste files across local and remote computers. Compared with FTP, WinSCP allows you to use a username and

password to access the destination server without any additional configuration on the server.

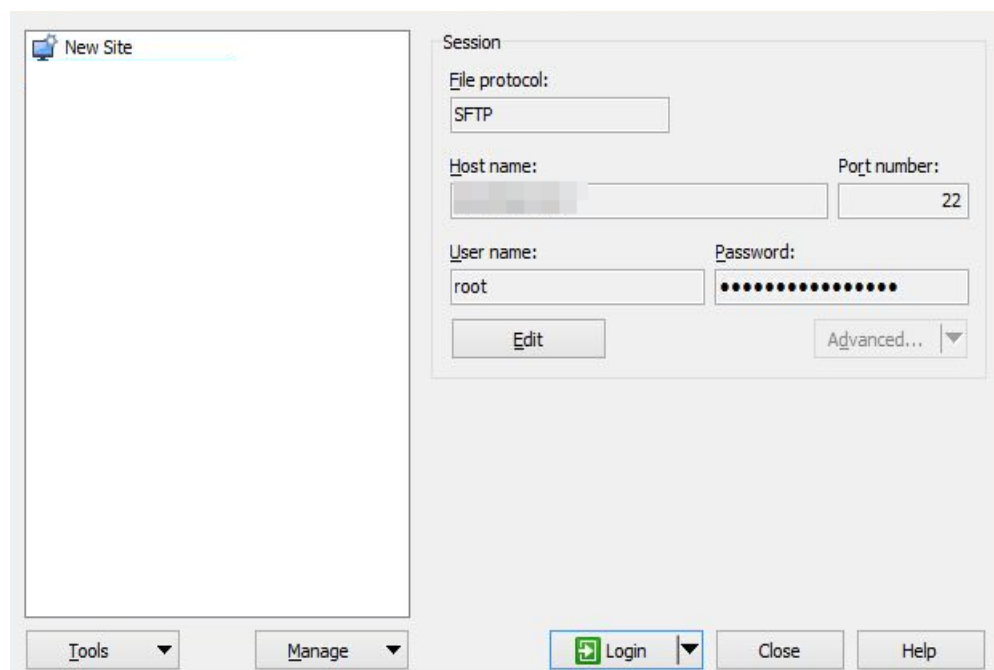
To transfer a file from a local Windows computer to a Linux ECS, WinSCP is commonly used. This section describes how to transfer files from a local Windows computer to a Linux ECS using WinSCP. In this example, the ECS running CentOS 7.2 is used as an example.

Prerequisites

- The target ECS is running.
- An EIP has been bound to the ECS. For details, see .
- Access to port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).

Solution

1. [Download WinSCP](#).
2. Install WinSCP.
3. Start WinSCP.



Set parameters as follows:

- **File protocol:** Set this to **SFTP** or **SCP**.
- **Host name:** Enter the EIP bound to the ECS. Log in to the management console to obtain the EIP.
- **Port number:** **22** by default.
- **User Name:** Enter the username for logging in to the ECS.
 - If the ECS is logged in using an SSH key pair,
 - The username is **core** for a CoreOS public image.
 - The username is **root** for a non-CoreOS public image.

- If the ECS is logged in using a password, the username is **root** for a public image.
 - **Password:** the password set when you purchased the ECS or converted using a key.
4. Click **Login**.
 5. Drag a file from the local computer on the left to the remotely logged in ECS on the right to transfer the file.

15.3 How Can I Use SCP to Transfer Files Between a Local Linux Computer and a Linux ECS?

Scenarios

You want to use SCP to transfer files between a local Linux computer and a Linux ECS.

Procedure

Log in to the management console. On the **Elastic Cloud Server** page, obtain the EIP bound to the target ECS in the **IP Address** column.

- **Uploading files**

Run the following command on the local Linux computer to upload files to the Linux ECS:

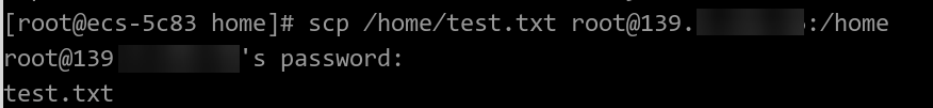
scp *Path in which the files are stored on the local computer*
Username@EIP:Path in which the files are to be stored on the Linux ECS

For example, to transfer the **/home/test.txt** file on the local computer to the **/home** directory on the ECS whose EIP is 139.x.x.x, run the following command:

```
scp /home/test.txt root@139.x.x.x:/home
```

Enter the login password as prompted.

Figure 15-1 Setting file uploading



```
[root@ecs-5c83 home]# scp /home/test.txt root@139.█:~/home
root@139.█'s password:
test.txt
```

- **Downloading files**

Run the following command on the local Linux computer to download files from the Linux ECS:

scp *Username@EIP:Path in which the files are stored on the Linux ECS* *Path in which the files are to be stored on the local computer*

For example, to download the **/home/test.txt** file on the ECS whose EIP is 139.x.x.x to the **/home** directory on the local computer, run the following command:

```
scp root@139.x.x.x:/home/test.txt /home/
```

Enter the login password as prompted.

Figure 15-2 Setting file downloading

```
[root@ecs-5c83 home]# scp root@139.139.139.139:/home/test.txt /home
root@139.139.139.139's password:
test.txt
[root@ecs-5c83 home]# ls
test.txt
```

15.4 How Can I Use SFTP to Transfer Files Between a Local Linux Computer and a Linux ECS?

Scenarios

You want to use SFTP to transfer files between a local Linux computer and a Linux ECS. The following uses CentOS as an example.

Procedure

1. Log in to the ECS as user **root**.
2. Run the following command to check the OpenSSH version, which is expected to be 4.8p1 or later:

```
ssh -V
```

Information similar to the following is displayed:

```
# OpenSSH_7.4p1, OpenSSL 1.0.2k-fips 26 Jan 2017
```

3. Create a user group and a user (for example, **user1**).

```
groupadd sftp
```

```
useradd -g sftp -s /sbin/nologin user1
```

4. Set a password for the user.

```
passwd user1
```

Figure 15-3 Setting a password

```
[root@ecs-9a32-0001 ~]# passwd user1
Changing password for user user1.
New password:
BAD PASSWORD: The password is shorter than 8 characters
Retype new password:
passwd: all authentication tokens updated successfully.
[root@ecs-9a32-0001 ~]#
```

5. Assign permissions to directories.

```
chown root:sftp /home/user1
```

```
chmod 755 -R /home/user1
```

```
mkdir /home/user1/upload
```

```
chown -R user1:sftp /home/user1/upload
```

```
chmod -R 755 /home/user1/upload
```

6. Run the following command to edit the **sshd_config** configuration file:

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

Comment out the following information:

```
#Subsystem sftp /usr/libexec/openssh/sftp-server
```

Add the following information:

```
Subsystem sftp internal-sftp
Match Group sftp
ChrootDirectory /home/%u
ForceCommand internal-sftp
AllowTcpForwarding no
X11Forwarding no
```

Figure 15-4 sshd_config file with the added information

```
# override default of no subsystems
#Subsystem sftp /usr/libexec/openssh/sftp-server

# Example of overriding settings on a per-user basis
#Match User anoncvs
#   X11Forwarding no
#   AllowTcpForwarding no
#   PermitTTY no
#   ForceCommand cvs server
PermitRootLogin yes
PasswordAuthentication yes
UseDNS no
Subsystem sftp internal-sftp
Match Group sftp
ChrootDirectory /home/%u
ForceCommand internal-sftp
AllowTcpForwarding no
X11Forwarding no
```

7. Run the following command to restart the ECS:
service sshd restart
Alternatively, run the following command to restart sshd:
systemctl restart sshd
8. Run the following command on the local computer to set up the connection:
sftp root@IP address
9. Run the **sftp** command to check the connection.

```
root@ [redacted] 's password:
Connected to [redacted].
sftp> ls
ceshi          print_all_tty.sh
s3fs_1.80_centos6.5_x86_64.rpm  speedtest.py
uploads
sftp> pwd
Remote working directory: /root
sftp> lpwd
Local working directory: /root
sftp> █
```

10. Transfer files or folders.
To upload files or folders, run the **put -r** command.

```
sftp> put -r ceshi/  
Uploading ceshi/ to /root/ceshi  
Entering ceshi/  
ceshi/mysql57-community-release-el 100% 9224      9.0KB/s   00:00  
ceshi/haha 100% 28      0.0KB/s   00:00  
sftp>
```

To download files or folders, run the **get -r** command.

```
sftp> get -r s3fs_1.80_centos6.5_x86_64.rpm  
Fetching /root/s3fs_1.80_centos6.5_x86_64.rpm to s3fs_1.80_centos6.5_x86_64.rpm  
/root/s3fs_1.80_centos6.5_x86_64.r 100% 3250KB   3.2MB/s   00:00  
sftp>
```

15.5 How Can I Use FTP to Transfer Files from a Local Windows Computer to a Linux ECS?

Scenarios

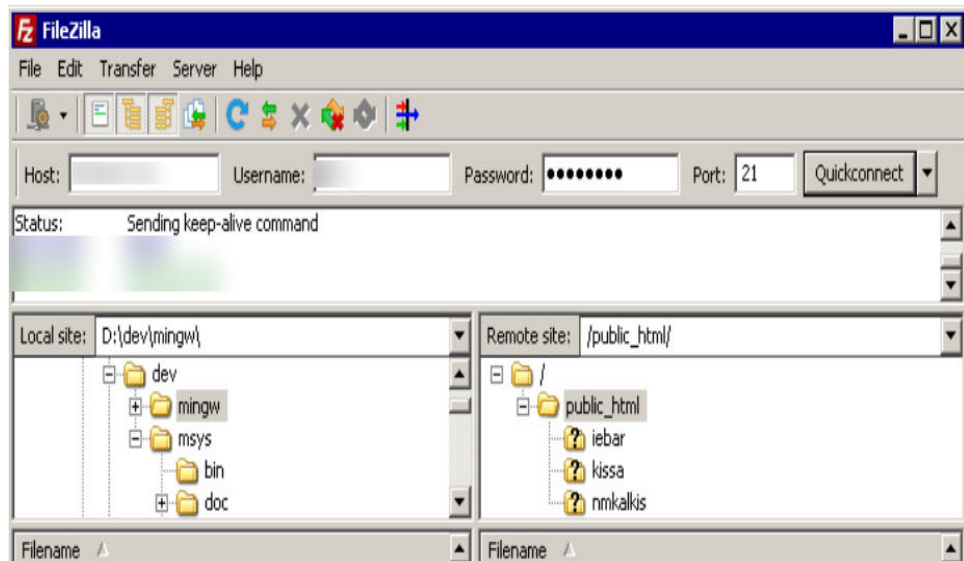
You want to use FTP to transfer files from a local Windows computer to an ECS.

Prerequisites

- An EIP has been bound to the ECS and access to port 21 is allowed in the inbound direction of the security group to which the ECS belongs.
- You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

Procedure

1. [Download FileZilla](#) and install it on the local Windows computer.
2. On the local Windows computer, open FileZilla, enter the information about the target ECS, and click **Quickconnect**.
 - **Host:** EIP bound to an ECS
 - **Username:** username set when the FTP site was set up
 - **Password:** password of the username
 - **Port:** FTP access port, which is port 21 by default

Figure 15-5 Setting connection parameters

3. Drag files from the local computer on the left to the target ECS on the right to transfer them.

15.6 How Can I Use FTP to Transfer Files Between a Local Linux Computer and a Linux ECS?

Scenarios

You want to use FTP on a local Linux computer to transfer files between the computer and a Linux ECS.

Prerequisites

You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

- An EIP has been bound to the ECS and access to port 21 is allowed in the inbound direction of the security group to which the ECS belongs.
- You have enabled FTP on the target ECS. If you have not enabled FTP, check the following links to know how to set up an FTP site:

Procedure

1. Install FTP on the local Linux computer.
Take CentOS 7.6 as an example. Run the following command to install FTP:
yum -y install ftp
2. Run the following command to access the ECS:
ftp EIP bound to the ECS
Enter the username and password as prompted for login.
 - **Uploading files**
Run the following command to upload local files to the ECS:

put *Path in which files are stored on the local computer*

For example, to upload the **/home/test.txt** file on the local Linux computer to the ECS, run the following command:

put /home/test.txt

– **Downloading files**

Run the following command to download files on the ECS to the local computer:

get *Path in which the files are stored on the ECS Path in which the files are to be stored on the local computer*

For example, to download the **test.txt** file on the ECS to the local Linux computer, run the following command:

get /home/test.txt

15.7 What Should I Do If the Connection Between the Client and the Server Times Out When I Upload a File Using FTP?

Symptom

When I attempted to access the server from the client to upload a file using FTP, the connection timed out.

Constraints

The operations described in this section apply to FTP on local Windows only.

Possible Causes

Data is intercepted by the firewall or security group on the server.

Solution

1. Check the firewall settings on the server.
2. Disable the firewall or add desired rules to the security group.

15.8 What Should I Do If Writing Data Failed When I Upload a File Using FTP?

Symptom

When I attempted to upload a file using FTP, writing data failed. As a result, the file transfer failed.

Constraints

The operations described in this section apply to FTP on Windows ECSs only.

Possible Causes

When NAT is enabled on the FTP server, the FTP client must connect to the FTP server in passive mode. In such a case, the public IP address (EIP) of the server cannot be accessed from the router. You need to add the EIP to the public IP address list on the server. Additionally, set the port range to limit the number of ports with data forwarded by the router.


Solution

The EIP must be associated with the private IP address using NAT, so the server must be configured accordingly.

1. Set the public IP address of the server.

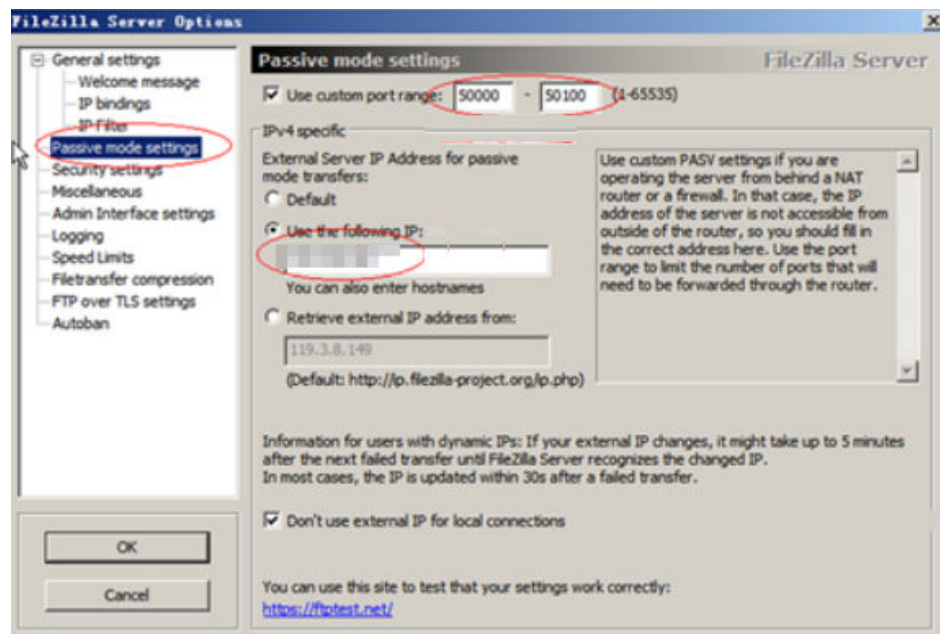
Choose **Edit > Settings**.

Figure 15-6 Setting the public IP address



2. Choose **Passive mode settings**, set the port range (for example, 50000-50100) for transmitting data, and enter the target EIP.

Figure 15-7 Setting the range of ports for data transmission



3. Click **OK**.
4. Allow traffic on TCP ports 50000-50100 and 21 in the security group in the inbound direction.


5. Test the connection on the client.

15.9 Why Am I Seeing an FTP Folder Error When I Open a Folder on an FTP Server?

Symptom

An error occurs when you open a folder on an FTP server. The system displays a message asking you to check permissions.

Figure 15-8 FTP Folder Error



Possible Causes


The FTP firewall configured for the browser does not allow you to open the folder.

Solution

The following uses Internet Explorer as an example.

1. Open the Internet Explorer and choose **Tools > Internet options**.
2. Click the **Advanced** tab.
3. Deselect **Use Passive FTP (for firewall and DSL modem compatibility)**.

Figure 15-9 Internet Options



4. Click **OK**, restart Internet Explorer, and open the folder on the FTP server again.

15.10 Why Do I Fail to Connect to a Linux ECS Using WinSCP?

Symptom

Connecting to a Linux ECS using WinSCP fails, while using SSH tools like Xshell succeeds.

Figure 15-10 Connection error using WinSCP



Root Cause

If you can connect to a Linux ECS using SSH tools, the SSH tools run properly. Check the SFTP configuration file because WinSCP allows you to connect your Linux ECS via SFTP protocol.

Run the following command to view the `/etc/ssh/sshd_config` file:

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

Check the SFTP configuration and the configuration file is `/usr/libexec/openssh/sftp-server`.

Figure 15-11 SFTP configuration file

```
# override default of no subsystems
Subsystem      sftp          /usr/libexec/openssh/sftp-server
```

If the SFTP configuration file does not exist or the file permission is not 755, connecting to a Linux ECS using WinSCP will fail.

Solution

- If the SFTP configuration file does not exist, you can transfer the file from an ECS that runs properly to your Linux ECS using SCP or other file transfer tools.
- If the file permission is not 755, you can run the following command to change the file permission to 755:

```
chmod 755 -R /usr/libexec/openssh/sftp-server
```

16 ECS Failure FAQ

16.1 How Do I Handle Error Messages Displayed on the Management Console?

Symptom

This section helps you resolve the following issues:

- An error message was displayed on the management console after you performed ECS-related operations.
- An error code was displayed after you used an ECS API (see *Elastic Cloud Server API Reference*).

Background

After you perform ECS-related operations on the management console, the system displays the request status on the **Elastic Cloud Server** page. You can determine the request execution status based on the information displayed in the request status.

- If the operation request is executed, the system automatically clears the task prompt.
- If an error occurs during the request execution, the system displays an error code and its description in the taskbar.

Solution

If an error occurs, check the error code and perform the corresponding operations listed in [Table 16-1](#).

Table 16-1 Error codes and solution suggestions

Error Code	Message Displayed on the Management Console	Solution Suggestion for Pay-per-Use ECSs	Solution Suggestion for Yearly/Monthly ECSs
Ecs.0000	Request error. Try again later or contact customer service.	Adjust the request structure as directed in <i>Elastic Cloud Server API Reference</i> .	Adjust the request structure as directed in <i>Elastic Cloud Server API Reference</i> .
Ecs.0001	The maximum number of ECSs or EVS disks has been reached. Contact customer service and request a quota increase.	Contact customer service to apply for an increased ECS quota. NOTE When applying for increasing your ECS quota, first determine the number of target ECSs, CPU cores (vCPUs), and memory capacity (RAM) required.	Submit a service ticket to apply for an increased ECS quota. After the submission, contact customer service for troubleshooting. NOTE When applying for increasing your ECS quota, first determine the number of target ECSs, CPU cores (vCPUs), and memory capacity (RAM) required.
Ecs.0005	System error. Try again later or contact customer service.	Adjust the request structure as directed in <i>Elastic Cloud Server API Reference</i> .	Adjust the request structure as directed in <i>Elastic Cloud Server API Reference</i> .
Ecs.0006	Invalid parameters.	If your selected ECS configuration has not been released, change the configuration and create the ECS again.	If your selected ECS configuration has not been released, change the configuration and create the ECS again.
Ecs.0010	The private IP address is in use. Select an available IP address and create the ECS again.	Use idle IP addresses to purchase your ECSs.	Contact customer service to cancel the order and use an idle IP address to purchase your ECS.

Error Code	Message Displayed on the Management Console	Solution Suggestion for Pay-per-Use ECSs	Solution Suggestion for Yearly/Monthly ECSs
Ecs.0011	Invalid password. Change the password to make it meet the password complexity requirements, and perform the required operation again.	Input a password that meets password complexity requirements.	Contact customer service to cancel the order, input a password that meets password complexity requirements, and perform the request again.
Ecs.0012	The number of IP addresses in the subnet is insufficient. Release IP addresses in the subnet or select another subnet, and create the ECS again.	Obtain more idle IP addresses on the target subnet or use a new subnet for purchasing ECSs.	Obtain more idle IP addresses on the target subnet and contact customer service for troubleshooting. Alternatively, contact customer service to cancel the order and use a new subnet for purchasing ECSs.
Ecs.0013	Insufficient EIP quota. Contact customer service and request an EIP quota increase.	Contact customer service to apply for an increased EIP quota.	Submit a service ticket to apply for an increased EIP quota and contact customer service for troubleshooting.
Ecs.0015	This disk type is not supported by the ECS.	Select a supported EVS disk and attach it to the ECS.	Select a supported EVS disk and attach it to the ECS.
Ecs.0100	The ECS status does not meet requirements. Change to the desired ECS status and try again.	Change to the desired ECS status and try again.	Change the ECS status to the required status and contact customer service for troubleshooting.
Ecs.0104	Insufficient number of ECS slots for attaching disks.	Detach an EVS disk from the ECS before attaching a new EVS disk.	Detach an EVS disk from the ECS before attaching a new EVS disk.

Error Code	Message Displayed on the Management Console	Solution Suggestion for Pay-per-Use ECSs	Solution Suggestion for Yearly/Monthly ECSs
Ecs.0105	No system disk found.	Reattach the EVS system disk to the ECS and perform the desired operation again.	Reattach the EVS system disk to the ECS and contact customer service for troubleshooting.
Ecs.0107	The number of shared disks to be attached to an ECS exceeds the maximum limit.	Detach an EVS disk from the ECS before attaching a new EVS disk.	Detach an EVS disk from the ECS before attaching a new EVS disk.
Ecs.0509	This operation is not allowed on a yearly/monthly system disk. Select a pay-per-use system disk and perform the required operation again.	N/A	Change the ECS billing mode to pay-per-use and perform the desired operation again.
Ecs.0510	Yearly/Monthly ECSs do not support OS changing.	N/A	Change the ECS billing mode to pay-per-use and perform the desired operation again.

16.2 What Should I Do If Emails Configured on an ECS Cannot Be Sent?

Solution

- For the emails sent using the browser:

When you use a browser to log in to your mailbox, HTTP is used, and the default port number is 80. However, SMTP is used between email servers. If you use a browser to send emails, enable port 80 for TCP in the outbound direction.

 - a. On the ECS details page, locate the security group and click the security group ID.
 - b. On the **Security Group** page, click the **Outbound** tab and then **Add Rule**.

- c. In the dialog box that is displayed, set **Protocol/Application** to **TCP** and **Port** to **80**. Then, click **OK**.

Figure 16-1 Adding port 80

- For the emails sent and received through an email client:
The protocols used on the receiving and transmitting ends are different.
Protocol used on the transmitting end:
SMTPS is used, and the port number is 465. Alternatively, SMTP is used, and the port number is 25.
Port 465 is recommended. If port 25 is required, enable it. For details, see [Why Is Outbound Access on TCP Port 25 Blocked?](#)
Protocol used on the receiving end: POP3 is used, and the port number is 110.
For details, see steps [a](#) to [c](#).

NOTE

Add an inbound rule with **Protocol** set to **TCP** and **Port** to **110**. Add an outbound rule with **Protocol** set to **TCP** and **Port** to **465** or **25**.

16.3 Why Is the Memory of an ECS Obtained by Running the free Command Inconsistent with the Actual Memory?

Symptom

After you create an ECS, you run the **free -m** command to view the ECS memory. The ECS memory is less than the memory configured during ECS creation.

For example:

When you are creating an ECS, the configured memory size is 4,194,304 KB (4,096 MB). After the ECS is created, you run the **free -m** command to view its memory. The command output is as follows:

```
[root@localhost ~]# free -m
total used free shared buff/cache available
```

```
Mem: 3790 167 3474 8 147 3414  
Swap: 1022 0 1022
```

The memory in the command output is 3,790 MB, which is less than the configured 4,096 MB.

Run the **dmidecode -t memory** command to check the actual memory configured for the ECS. The command output is as follows:

```
[root@localhost ~]# dmidecode -t memory  
# dmidecode 3.0  
Getting SMBIOS data from sysfs.  
SMBIOS 2.8 present.  
  
Handle 0x1000, DMI type 16, 23 bytes  
Physical Memory Array  
Location: Other  
Use: System Memory  
Error Correction Type: Multi-bit ECC  
Maximum Capacity: 4 GB  
Error Information Handle: Not Provided  
Number Of Devices: 1  
  
Handle 0x1100, DMI type 17, 40 bytes  
Memory Device  
Array Handle: 0x1000  
Error Information Handle: Not Provided  
Total Width: Unknown  
Data Width: Unknown  
Size: 4,096 MB  
Form Factor: DIMM  
Set: None  
Locator: DIMM 0  
Bank Locator: Not Specified  
Type: RAM  
Type Detail: Other  
Speed: Unknown  
Manufacturer: QEMU  
Serial Number: Not Specified  
Asset Tag: Not Specified  
Part Number: Not Specified  
Rank: Unknown  
Configured Clock Speed: Unknown  
Minimum Voltage: Unknown  
Maximum Voltage: Unknown  
Configured Voltage: Unknown
```

The memory in the command output is the same as that configured during ECS creation.

Possible Causes

When the OS is started, related devices are initialized, which occupies memory. In addition, when the kernel is started, it also occupies memory. The memory occupied by `kdump` can be set. Unless otherwise specified, do not change the memory size occupied by `kdump`.

The command output of **free -m** shows the available memory of the ECS, and that of **dmidecode -t memory** shows the hardware memory.

The memory obtained by running the **free -m** command is less than the memory configured for the ECS. This is a normal phenomenon.

NOTE

This is a normal phenomenon even for physical servers.

16.4 Why Is the Hostname of My ECS Restored to the Original Name After the ECS Is Restarted?

The following uses an ECS running CentOS 7 as an example:

1. Log in to the Linux ECS and view the Cloud-Init configuration file.
2. In the `/etc/cloud/cloud.cfg` file, comment out or delete `update_hostname`.

NOTE

- `update_hostname` indicates that the hostname is changed in Cloud-Init each time the ECS is restarted.
- For an ECS created from a public image, Cloud-Init has been installed on it by default. You do not need to manually install Cloud-Init for it. For details about how to modify a private image, see [Installing Cloud-Init](#).

16.5 Is an ECS Hostname with Suffix `.novalocal` Normal?

Symptom

Hostnames of ECSs created based on some types of images have the suffix `.novalocal`, whereas others do not.

For example, the hostname is set to `abc` during ECS creation. [Table 16-2](#) lists the hostnames (obtained by running the `hostname` command) of ECSs created using different images and those displayed after the ECSs are restarted.

Table 16-2 Hostnames of ECSs created from different images

Image	Hostname Before ECS Restart	Hostname After ECS Restart
CentOS 6.8	abc	abc.novalocal
CentOS 7.3	abc.novalocal	abc.novalocal
Ubuntu 16	abc	abc

Troubleshooting

This is a normal phenomenon.

The static hostname of a Linux ECS is user defined and injected using Cloud-Init during the ECS creation. According to the test results, Cloud-Init adapts to OSs differently. As a result, hostnames of some ECSs have suffix `.novalocal`, whereas others do not.

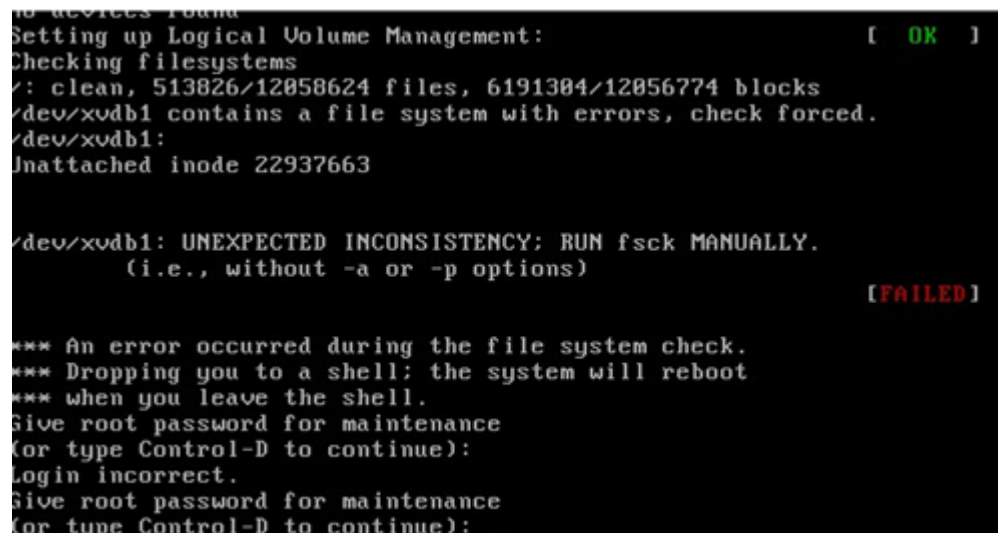
If you do not want to have the obtained hostnames contain suffix `.novalocal`, change the hostnames by referring to [How Can a Changed Static Hostname Take Effect Permanently?](#)

16.6 Why Does a Forcibly-Stopped Linux ECS Fail to Be Restarted?

Symptom

When you try to restart a forcibly-stopped Linux ECS, the ECS failed to be restarted, as shown in [Figure 16-2](#).

Figure 16-2 Restart failure



```
Setting up Logical Volume Management: [ OK ]
Checking filesystems
/: clean, 513826/12858624 files, 6191384/12856774 blocks
/dev/xvdb1 contains a file system with errors, check forced.
/dev/xvdb1:
Unattached inode 22937663

/dev/xvdb1: UNEXPECTED INCONSISTENCY: RUN fsck MANUALLY.
(i.e., without -a or -p options) [FAILED]

*** An error occurred during the file system check.
*** Dropping you to a shell; the system will reboot
*** when you leave the shell.
Give root password for maintenance
(or type Control-D to continue):
Login incorrect.
Give root password for maintenance
(or type Control-D to continue):
```

Possible Causes

As shown in [Figure 16-2](#), the ECS cannot be restarted because the file system was damaged. Forcibly stopping or restarting an ECS is highly risky because this operation may cause inconsistent metadata in the file system, leading to the file system damage.

Solution

Use the disk repair tool (fsck) delivered with the Linux OS to rectify the fault.

The following procedure considers the affected disk partition as `/dev/xvdb1`, which is the partition shown in [Figure 16-2](#).

1. Enter the password of user `root` as prompted.
2. Run the following command to check whether the affected disk partition has been mounted:

```
mount | grep xvdb1
```

- If yes, go to step 3.

- If no, go to step 4.
- 3. Run the following command to unmount the affected disk partition:
umount /dev/xvdb1
- 4. Run the following command to rectify the file system of the affected disk partition:
fsck -y /dev/xvdb1
- 5. Run the following command to restart the ECS:
reboot

 NOTE

If the fault persists, contact customer service for technical support.

16.7 How Can a Changed Static Hostname Take Effect Permanently?

Symptom

The static hostname of a Linux ECS is user defined and injected using Cloud-Init during the ECS creation. Although the hostname can be changed by running the **hostname** command, the changed hostname is restored after the ECS is restarted.

Changing the Hostname on the ECS

To make the changed hostname still take effect even after the ECS is stopped or restarted, save the changed hostname into configuration files.

The changed hostname is assumed to be **new_hostname**.

1. Modify the **/etc/hostname** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/hostname
 - b. Change the hostname to the new one.
 - c. Run the following command to save and exit the configuration file:
:wq
2. Modify the **/etc/sysconfig/network** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/sysconfig/network
 - b. Change the **HOSTNAME** value to the new hostname.
HOSTNAME=Changed hostname

 NOTE

If there is no **HOSTNAME** in the configuration file, manually add this parameter and set it to the changed hostname.

For example:

```
HOSTNAME=new_hostname
```

- c. Run the following command to save and exit the configuration file:
:wq
3. Modify the **/etc/cloud/cloud.cfg** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/cloud/cloud.cfg
 - b. Use either of the following methods to modify the configuration file:
 - Method 1: Change the **preserve_hostname** parameter value or add the **preserve_hostname** parameter to the configuration file.
If **preserve_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve_hostname: true**.
If **preserve_hostname** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve_hostname: true** before **cloud_init_modules**.
If you use method 1, the changed hostname still takes effect after the ECS is stopped or restarted. However, if the ECS is used to create a private image and the image is used to create a new ECS, the hostname of the new ECS is the hostname (**new_hostname**) used by the private image, and user-defined hostnames cannot be injected using Cloud-Init.
 - Method 2 (recommended): Delete or comment out - **update_hostname**.
If you use method 2, the changed hostname still takes effect after the ECS is stopped or restarted. If the ECS is used to create a private image and the image is used to create a new ECS, the changed hostname permanently takes effect, and user-defined hostnames (such as **new_new_hostname**) can be injected using Cloud-Init.

 NOTE

After a user-defined name of ECS using some OS versions is injected using Cloud-Init, the host name obtained by running the **hostname** command may contain the **.novalocal** suffix. For details, see [Is an ECS Hostname with Suffix .novalocal Normal?](#)

In this case, you can modify the configuration by referring to Step 8 in [Configuring Cloud-Init](#) so that the host name created using a private image does not contain the **.novalocal** suffix.

4. Run the following command to restart the ECS:
sudo reboot
5. Run the following command to check whether the hostname has been changed:
sudo hostname
If the changed hostname is displayed in the command output, the hostname has been changed and the new name permanently takes effect.

Modifying the Mapping Between the ECS Hostname and IP Address (Modifying the hosts File)

If you want to use the changed hostname as the preferred localhost and localhost.localdomain, update the mapping between the hostname and IP address

after the hostname is changed and then save the configuration to the corresponding Cloud-Init configuration file so that the new hostname takes effect permanently.

The changed hostname is assumed to be **new_hostname**.

1. Modify the **/etc/hostname** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/hostname
 - b. Change the hostname to the new one.
 - c. Run the following command to save and exit the configuration file:

:wq

2. Modify the **/etc/sysconfig/network** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/sysconfig/network
 - b. Change the **HOSTNAME** value to the new hostname.

HOSTNAME=Changed hostname

NOTE

If there is no **HOSTNAME** in the configuration file, manually add this parameter and set it to the changed hostname.

For example:

```
HOSTNAME=new_hostname
```

- c. Run the following command to save and exit the configuration file:

:wq

3. Modify the **/etc/cloud/cloud.cfg** configuration file.

- a. Run the following command to edit the configuration file:

sudo vim /etc/cloud/cloud.cfg

- b. Use either of the following methods to modify the configuration file:

- Method 1: Change the **preserve_hostname** parameter value or add the **preserve_hostname** parameter to the configuration file.

If **preserve_hostname: false** is already available in the **/etc/cloud/cloud.cfg** configuration file, change it to **preserve_hostname: true**.

If **preserve_hostname** is unavailable in the **/etc/cloud/cloud.cfg** configuration file, add **preserve_hostname: true** before **cloud_init_modules**.

If you use method 1, the changed hostname still takes effect after the ECS is stopped or restarted. However, if the ECS is used to create a private image and the image is used to create a new ECS, the hostname of the new ECS is the hostname (**new_hostname**) used by the private image, and user-defined hostnames cannot be injected using Cloud-Init.

- Method 2 (recommended): Delete or comment out - **update_hostname**.

If you use method 2, the changed hostname still takes effect after the ECS is stopped or restarted. If the ECS is used to create a private

image and the image is used to create a new ECS, the changed hostname permanently takes effect, and user-defined hostnames (such as **new_new_hostname**) can be injected using Cloud-Init.

4. Update the mapping between the hostname and IP address in **/etc/hosts** to an entry starting with 127.0.0.1. Use **new_hostname** as your preferred **localhost** and **localhost.localdomain**.
 - a. Run the following command to edit **/etc/hosts**:
sudo vim /etc/hosts
 - b. Modify the entry starting with 127.0.0.1 and replace **localhost** and **localhost.localdomain** with **new_hostname**.

```
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain4
127.0.0.1 new_hostname new_hostname
```
 - c. Run the following command to save and exit the configuration file:
:wq
5. Modify the **/etc/cloud/cloud.cfg** configuration file.
 - a. Run the following command to edit the configuration file:
sudo vim /etc/cloud/cloud.cfg
 - b. Set **manage_etc_hosts** to **manage_etc_hosts: false**.

```
manage_etc_hosts: false
```
 - c. Run the following command to save and exit the configuration file:
:wq
6. Run the following command to restart the ECS:
sudo reboot
7. Run the following commands to check whether the changes to **hostname** and **hosts** take effect permanently:
sudo hostname
sudo cat /etc/hosts

If the changed hostname (**new_hostname**) and **hosts** are displayed in the command output, the changes take effect permanently.

16.8 Why Can't My Linux ECS Obtain Metadata?

Symptom

The security group of the Linux ECS has been configured based on the prerequisites in [Obtaining Metadata](#) in the outbound direction, but the ECS still cannot obtain the metadata through the route with the destination of 169.254.169.254.

Root Cause

Run the following command on the Linux ECS configured with a static IP address:

```
# ip route| grep 169.254
```

The route with the destination of 169.254.169.254 does not exist, but the route with the destination of 169.254.0.0/16 exists.

Figure 16-3 Route information

```
1# ip route | grep 169.254
169.254.0.0/16 dev eth0 scope link
1#
```

After the network is restarted, the original route with the destination of 169.254.169.254 is changed to the route with the destination of 169.254.0.0/16 without a next hop, as shown in [Figure 16-3](#). As a result, the Linux ECS cannot obtain metadata.

Solution


1. Add the route with the destination of 169.254.169.254, and specify the next hop (gateway) and the output device (primary NIC of the Linux ECS). The following is an example:
ip route add 169.254.169.254 via 192.168.1.1 dev eth0
192.168.1.1 is the gateway address of the subnet that the primary NIC resides, and eth0 is the primary NIC.
[How Do I View the Primary NIC?](#)
[How Do I View the Gateway Address?](#)
2. Run the following command to verify that the metadata can be obtained:
curl http://169.254.169.254

Figure 16-4 Obtaining metadata

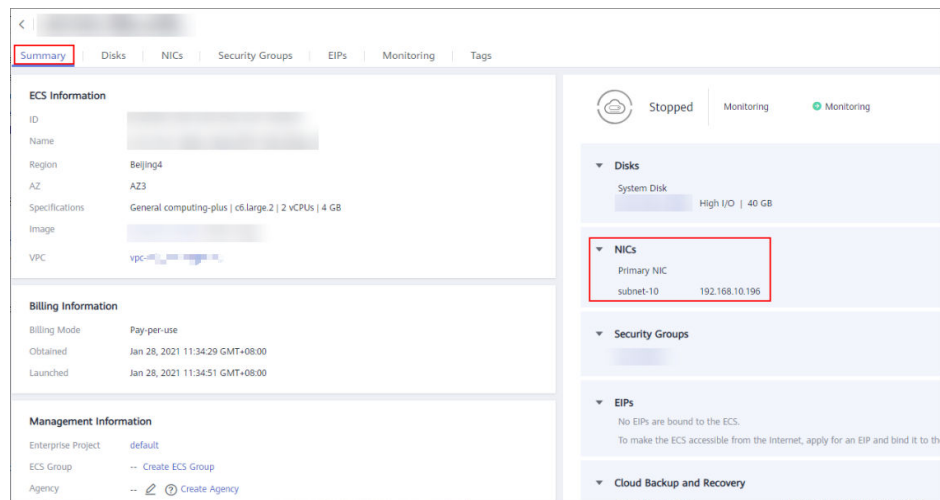
```
ecs-test [ ~ ] # ip route add 169.254.169.254 via 192.168.1.1 dev eth0
ecs-test [ ~ ] # curl http://169.254.169.254
1.0
2007-01-19
2007-03-01
2007-08-29
2007-10-10
2007-12-15
2008-02-01
2008-09-01
2009-04-04
latestecs-test [ ~ ] #
```

3. Run the following command to create or modify the `/etc/sysconfig/network-scripts/route-eth0` file to prevent the static route from being changed after network restart:
vi /etc/sysconfig/network-scripts/route-eth0
Add the following content to the file:
In this example, the primary NIC is eth0 and gateway address is 192.168.1.1. Replace them based on site requirements.
169.254.169.254 via 192.168.1.1

How Do I View the Primary NIC?

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.

3. Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. Click the **Summary** tab to view details about the primary NIC.

Figure 16-5 Primary NIC details

How Do I View the Gateway Address?


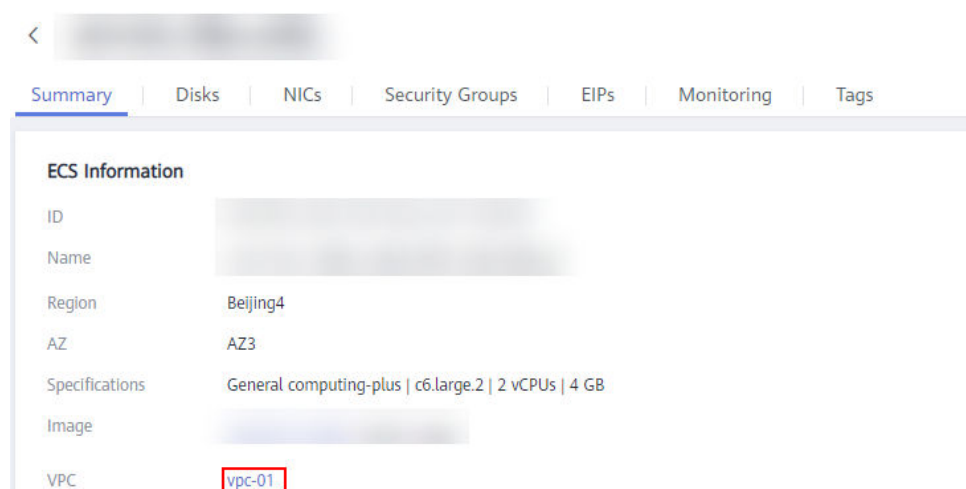
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. Click the VPC name to go to the VPC list page.

Figure 16-6 VPC name

6. Locate the row that contains the target VPC and click the number in the **Subnets** column to go to the subnet list page.

Figure 16-7 Number in the **Subnets** column

Name	IPv4 CIDR Block	Status	Subnets	Route Tables	Servers	Enterprise Project
vpc-01	(Primary CIDR block)	Available	1	1	4	default

7. Click the target subnet name to go to the subnet details page and view the gateway address.

Figure 16-8 Gateway address

Summary | IP Addresses | Tags

Subnet Information

Name	subnet-10	Network ID	
AZ	AZ1	IPv4 Subnet ID	
Status	Available	IPv6 Subnet ID	--
VPC	vpc-01	Description	--
Available IP Addresses	246		
IPv4 CIDR Block			
IPv6 CIDR Block			

Gateway and DNS Information

DHCP	Enabled	Gateway	192.168.10.1
DNS Server Address		DHCP Lease Time	365 days

17 Slow ECS Response FAQ

17.1 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 [General Operations for Modifying 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 malicious, 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 17-1 Starting the Resource Monitor



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

Figure 17-2 Resource Monitor




4. Obtain the IDs and names of the processes with high CPU or bandwidth usage.
5. On the 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 malicious.

- a. Click the **Details** 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 17-3 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.

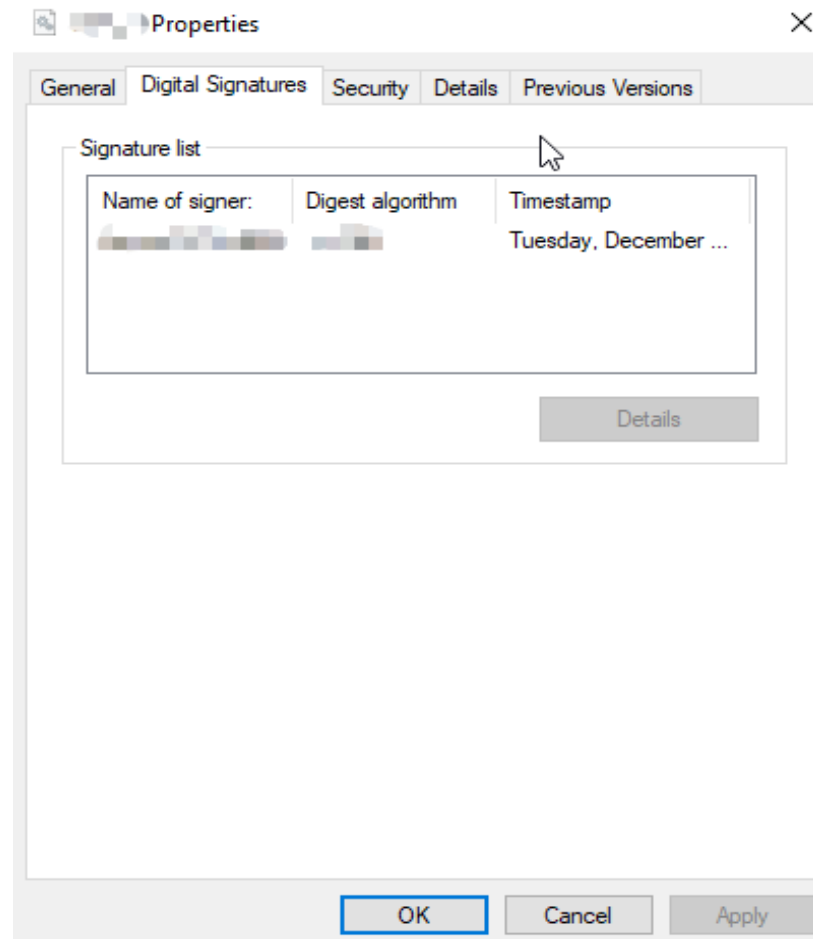
Figure 17-4 Viewing the system drivers

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
luaflv	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 17-5 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.

17.2 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 [General Operations for Modifying Specifications](#).

To handle this issue, perform the following operations:

1. Fault locating
Identify the processes leading to high bandwidth or CPU usage.
2. 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-mq01_sysconf:~# 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   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
```

- 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   VIRT   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. Run 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
4556	nglrx	nglrx: worker process	eth1	34.368	3.267
?	root	192.168.0.92:99-100.125.68.19:17873		0.179	0.246
?	root	192.168.0.92:11211-213.32.10.149:44945		0.000	0.000
?	root	192.168.0.92:20101-185.176.26.66:43408		0.000	0.000
?	root	unknown TCP		0.000	0.000
TOTAL				34.548	3.512

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 a CC Attack Protection Rule](#).

18 Specification Modification FAQ

18.1 What Should I Do If My Specifications Modification Request Failed to Submit?

Symptom

When you tried to modify specifications of a stopped ECS, the system displayed a message indicating that the system was busy, and the request failed to be submitted.

Solution

Check the ECS order and resources.

- If the resources specified in the order have entered the retention period, the ECS specifications cannot be modified. In such a case, renew the order and then modify the specifications.
- If your ECS is billed on a yearly/monthly basis, and the order has been renewed but the new order has not taken effect, the ECS specifications cannot be modified. In such a case, unsubscribe from the new order that has not taken effect and then modify the specifications.

18.2 Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?

Scenarios

After you modify specifications of a Linux ECS, disk attachment may fail. You need to check the disk attachment after you modify the specifications.

Procedure

1. Log in to the ECS as user **root**.

2. Run the following command to view the disks attached before specifications modification:

```
fdisk -l | grep 'Disk /dev/'
```

Figure 18-1 Viewing disks attached before specifications modification

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# fdisk -l |grep 'Disk /dev/'
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Disk /dev/vdc: 10.7 GB, 10737418240 bytes, 20971520 sectors
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
```

As shown in [Figure 18-1](#), the ECS has three disks attached: `/dev/vda`, `/dev/vdb`, and `/dev/vdc`.

3. Run the following command to view disks attached after specifications modification:

```
df -h | grep '/dev/'
```

Figure 18-2 Viewing disks attached after specifications modification

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# df -h | grep '/dev/'
/dev/vda2    39G  1.4G   35G   4% /
/dev/vda1    976M 146M  764M  16% /boot
```

As shown in [Figure 18-2](#), only one disk `/dev/vda` is attached to the ECS.

4. Check whether the number of disks obtained in step 3 is the same as that obtained in step 2.
 - If the numbers are the same, the disk attachment is successful. No further action is required.
 - If the numbers are different, the disk attachment failed. In this case, go to step 5.
5. Run the **mount** command to attach the affected disks.

For example, run the following command:

```
mount /dev/vdb1 /mnt/vdb1
```

In the preceding command, `/dev/vdb1` is the disk to be attached, and `/mnt/vdb1` is the path for disk attachment.

NOTICE

Ensure that `/mnt/vdb1` is empty. Otherwise, the attachment will fail.

6. Run the following commands to check whether the numbers of disks before and after specifications modifications are the same:

```
fdisk -l | grep 'Disk /dev/'
```

```
df -h | grep '/dev/'
```

- If the numbers are the same, no further action is required.
- If the numbers are different, contact customer service.

Figure 18-3 Checking the number of disks attached

```
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# mount /dev/vdb1 /mnt/vdb1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# mount /dev/vdc1 /mnt/vdc1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# fdisk -l |grep 'Disk /dev/'
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Disk /dev/vdb: 10.7 GB, 10737418240 bytes, 20971520 sectors
Disk /dev/vdc: 10.7 GB, 10737418240 bytes, 20971520 sectors
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]# df -h | grep '/dev/'
/dev/vda2      39G  1.4G  35G   4% /
/dev/vda1     976M 146M  764M  16% /boot
/dev/vdb1     9.8G  23M   9.2G   1% /mnt/vdb1
/dev/vdc1     9.8G  23M   9.2G   1% /mnt/vdc1
[root@servercf924ffa-da23-4d09-a7e0-416694a68492 ~]#
```

As shown in [Figure 18-3](#), the numbers of disks before and after specifications modifications are the same. The disks are `/dev/vda`, `/dev/vdb`, and `/dev/vdc`.

19 OS Change FAQ

19.1 Does OS Change Incur Fees?

After the OS is changed, different images are used and system disk capacity may increase. You will be billed based on the new configurations.

19.2 Can I Install or Upgrade the OS of an ECS?

You can install or upgrade ECS OSs provided on the cloud platform.

- When you create an ECS, you can select a public image or a private image created from a public image to install the ECS OS. Select an OS image based on the programming language in the actual application scenario.
- You can change your ECS OS through the management console, for example, you can upgrade CentOS 7.2 to CentOS 7.3.

19.3 Can I Change the OS of an ECS?

Yes, you can change the OS of an ECS.

If the OS running on an ECS cannot meet service requirements, for example, a higher OS version is required, you can change the ECS OS.

The cloud platform allows you to change the image type (public images, private images, and shared images) and OS. You can change the OS by changing the ECS image.

For instructions about how to change an ECS OS, see [Changing the OS](#).

19.4 How Long Does It Take to Change an ECS OS?

Generally, the process of changing the OS of an ECS takes about 1 to 2 minutes to complete. On the ECS console, stop the ECS and choose **More > Manage Image/Disk > Change OS** in the **Operation** column.

During this process, the ECS is in **Changing OS** state.

19.5 Will I Lose My Disk Data If I Reinstall ECS OS, Change the OS, or Change the ECS Specifications?

Table 19-1 Impact

Item	OS Reinstallation	OS Change	Specifications Modification
Application scenario	Initialize an ECS. The ECS OS remains unchanged after OS change.	Change the OS of an ECS by changing its image. For details about OS change constraints, see Changing the OS .	Change ECS specifications, such as increasing the number of vCPUs or adding memory, to meet your service requirements.
Billing	OS reinstallation is free of charge. The ECS price remains unchanged.	OS change is free of charge. However, you will be billed based on your new image type after OS change.	Modifying ECS specifications is free of charge. However, you will be billed based on the new specifications after modification.
IP address	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.	The private IP address, EIP, and MAC address remain unchanged.
System disk	Reinstalling OS will clear the data in all partitions of the ECS system disk. Back up data before reinstalling the OS.	Changing OS will clear the data in all partitions of the ECS system disk. Back up data before changing the OS.	No impact on system disk.
Data disk	No impact on data disk.	No impact on data disk	No impact on data disk.
Backup	Back up data before reinstalling the OS to prevent data loss.	Back up data before changing the OS to prevent data loss.	Create a system disk snapshot before modifying ECS specifications to prevent data loss.

19.6 Does OS Reinstallation Incur Fees?

Reinstalling an OS for an ECS allows you to use the original image to reinstall the ECS and does not incur fees.

19.7 Can I Select Another OS During ECS OS Reinstallation?

No. You can use only the original image of the ECS to reinstall the OS. To use a new system image, see [Changing the OS](#).

19.8 How Long Does It Take to Reinstall an ECS OS?

Generally, the process of reinstalling the OS of an ECS takes about 1 to 2 minutes to complete. On the ECS console, stop the ECS and choose **More > Manage Image/Disk/Backup > Reinstall OS** in the **Operation** column.

During this process, the ECS is in **Reinstalling OS** state.

20 ECS Security Check FAQ

20.1 How Does an ECS Defend Against DDoS Attacks?

What Is a DDoS Attack?

Denial of Service (DoS) attacks, also known as flood attacks, intend to exhaust the network or system resources on the target computer, causing service interruption or suspension. Consequently, legitimate users fail to access network services. When an attacker uses multiple compromised computers on the network as attack machines to launch DoS attacks to specific targets, the attacks are called Distributed Denial of Service (DDoS) attacks.

What Is Anti-DDoS?

Anti-DDoS defends ECSs against DDoS attacks and sends real time alarms when detecting attacks. In addition, Anti-DDoS improves the bandwidth utilization to further safeguard your services.

Anti-DDoS monitors the service traffic from the Internet to public IP addresses and detects attack traffic in real time. It then scrubs attack traffic based on user-configured defense policies without interrupting service running. It also generates monitoring reports that provide visibility into the security of network traffic.

Anti-DDoS

Anti-DDoS defends ECSs against DDoS attacks and sends real time alarms when detecting attacks. In addition, Anti-DDoS improves the bandwidth utilization to further safeguard your services.

Anti-DDoS monitors the service traffic from the Internet to public IP addresses and detects attack traffic in real time. It then scrubs attack traffic based on user-configured defense policies without interrupting service running. It also generates monitoring reports that provide visibility into the security of network traffic.

Anti-DDoS helps you mitigate the following attacks:

- Web server attacks

- Include SYN flood, HTTP flood, Challenge Collapsar (CC), and low-rate attacks
- Game attacks
 - Include User Datagram Protocol (UDP) flood, SYN flood, TCP-based, and fragmentation attacks
- HTTPS server attacks
 - Include SSL DoS and DDoS attacks
- DNS server attacks
 - Include attacks exploiting DNS protocol stack vulnerabilities, DNS reflection attacks, DNS flood attacks, and DNS cache miss attacks

Anti-DDoS also provides the following functions:

- Monitors a single EIP and offers a monitoring report, covering the current protection status, protection settings, and the traffic and anomalies within the last 24 hours.
- Provides attack statistics reports on all protected EIPs, covering the traffic scrubbing frequency, scrubbed traffic amount, top 10 attacked EIPs, and number of blocked attacks.

20.2 Are ECSs with Simple Passwords Easily Attacked?

It is recommended that your password contain 8 to 26 characters that consists of digits, uppercase and lowercase letters, and special characters. It is a good practice to download virtualization antivirus products and host security hardening products from Huawei Cloud KooGallery and install them on your ECSs to enhance security.

If your ECS has been intruded, contact customer service for technical support.

Table 20-1 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">• Consists of 8 to 26 characters.• Contains at least three of the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters for Linux ECSs: !@%-_#+=[:./^,}{?• Cannot contain the username or the username spelled backwards.

20.3 How Is ECS Security Ensured?

Host Security Service (HSS) helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.

If you use a public image to create an ECS, protection is enabled by default for your ECS. Its basic edition is free of charge. HSS automatically installs an agent on the ECS and protects the security of the ECS.

[How Do I Use HSS?](#)

20.4 How Can I Disable Operation Protection?

Symptom

When I perform critical operations on my ECS with operation protection enabled, for example, deleting my ECS or modifying ECS specifications, I have to enter the password and verification code for authentication. To disable operation protection, perform the operations described in this section.

Procedure

1. Log in to the management console.
2. On the **Security Settings** page, choose **Critical Operations > Operation Protection > Change**.
3. On the **Operation Protection** page, select **Disable** and click **OK**.

21 Resource Management and Tag FAQ

21.1 How Can I Create and Delete Tags and Search for ECSs by Tag?

Creating a Tag

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. Click **Tags** and then **Add Tag**.
6. Enter the tag key and value, and click **OK**.

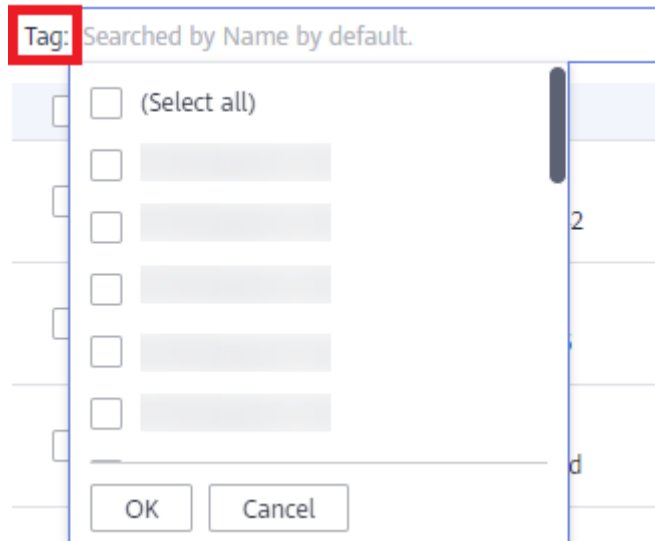
Figure 21-1 Adding tags

Searching for ECSs by Tag

1. Log in to the management console.

2. Select the region where the ECS is located.
3. On the **Elastic Cloud Server** page, search for ECSs by tag.

Figure 21-2 Searching for ECSs by tag



4. In the search bar, choose **Tag** and then select the tag key and value, and click **OK**.

Deleting a Tag

1. Log in to the management console.
2. Select the region where the ECS is located.
3. Click **Elastic Cloud Server**.
4. Click the name of the target ECS.
5. On the page providing details about the ECS, click **Tags**, locate the row containing the target tag, and click **Delete** in the **Operation** column.

Figure 21-3 Deleting a tag



22 Image Source Management FAQ

22.1 How Can I Use an Automated Tool to Configure a Huawei Cloud Image Source (x86_64 and Kunpeng)?

Scenarios

When updating the OS or software of an ECS, you can access the Internet through the ECS and use an external image source. However, if your ECS cannot access the Internet, or the external image source cannot provide services stably, you can use the one-click script configuration function provided by Huawei Cloud to configure and update the image source.

Constraints

- This section applies only to the following regions: CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, and CN-Hong Kong.
- You have configured the intranet DNS server provided by HUAWEI CLOUD. To obtain the intranet DNS server address, see [What Are Huawei Cloud Private DNS Server Addresses?](#) For instructions about how to configure a DNS server, see [How Do I Configure DNS for an ECS?](#)
- The operations described in this section apply to both x86_64 and Kunpeng ECSs.
- Only the following OSs are supported: EulerOS, CentOS, Debian, and Ubuntu.

NOTE

The automatic tool does not support the following operating systems:

- Debian 8.8.0 64bit
- Debian 8.2.0 64bit
- Ubuntu 22.04 server 64bit

Procedure

1. Log in to the ECS.
2. Run the following command to configure the image source:

```
wget http://mirrors.myhuaweicloud.com/repo/mirrors_source.sh && bash mirrors_source.sh
```

22.2 How Can I Use an EPEL Image Source (x86_64 or Kunpeng) Provided by Huawei Cloud?

Scenarios

When updating the OS or software of an ECS, you can access the Internet through the ECS and use an external EPEL image source. However, if your ECS cannot access the Internet, or the external EPEL image source cannot provide services stably, you can use the EPEL image source provided by Huawei Cloud.

Constraints

- This section applies only to the following regions: CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, and CN-Hong Kong.
- You have configured the intranet DNS server provided by HUAWEI CLOUD. To obtain the intranet DNS server address, see [What Are Huawei Cloud Private DNS Server Addresses?](#) For instructions about how to configure a DNS server, see [How Do I Configure DNS for an ECS?](#)
- The operations described in this section apply to both x86_64 and Kunpeng ECSs.


Procedure

1. Log in to the ECS as user **root**.
2. (Optional) Run the following command to back up the **epel.repo** file:

```
mv /etc/yum.repos.d/epel.repo /etc/yum.repos.d/epel.repo.bak  
mv /etc/yum.repos.d/CentOS-Base.repo /etc/yum.repos.d/CentOS-Base.repo.bak
```

For CentOS 8 or RHEL 8 images, you also need to run the following command to back up the **epel-modular.repo** file.

```
mv /etc/yum.repos.d/epel-modular.repo /etc/yum.repos.d/epel-modular.repo.bak
```

 **NOTE**

Perform this step only when the **epel.repo** and **epel-modular.repo** are available.
3. Run the following commands to obtain the configuration files:
 - CentOS 6 / RHEL 6

```
wget -qO /etc/yum.repos.d/CentOS-Base.repo http://mirrors.myhuaweicloud.com/repo/CentOS-Base-6.repo
```


- ```
wget -qO /etc/yum.repos.d/epel.repo http://mirrors.myhuaweicloud.com/repo/epel-6.repo
```
- CentOS 7 / RHEL 7

```
wget -qO /etc/yum.repos.d/CentOS-Base.repo http://mirrors.myhuaweicloud.com/repo/CentOS-Base-7.repo
```

```
wget -qO /etc/yum.repos.d/epel.repo http://mirrors.myhuaweicloud.com/repo/epel-7.repo
```
  - CentOS 8 / RHEL 8

```
wget -qO /etc/yum.repos.d/CentOS-Base.repo http://mirrors.myhuaweicloud.com/repo/CentOS-Base-8.repo
```

```
wget -qO /etc/yum.repos.d/epel.repo http://mirrors.myhuaweicloud.com/repo/epel-8.repo
```

```
wget -qO /etc/yum.repos.d/epel-modular.repo http://mirrors.myhuaweicloud.com/repo/epel-modular.repo
```

 NOTE

For CentOS 8 or RHEL 8 images, the **epel-modular.repo** file also needs to be updated to prevent the failure of updating **epel.repo**.

4. Run the following commands to generate a cache:

```
yum clean metadata
```

```
yum makecache
```

```
yum install epel-release -y >/dev/null 2>&1
```

 NOTE

- After you execute the **yum makecache** command, if the system displays the message "Another app is currently holding the yum lock", run the **rm -rf /var/run/yum.pid** command to delete the **yum.pid** file. Then, run the **yum makecache** command again to generate the cache.
- If the system displays the message "Error: Cannot retrieve metalink for repository: epel. Please verify its path and try again", run the following command to update the certificate and temporarily disable the EPEL image source:

```
yum --disablerepo=epel -y update ca-certificates
```
- If certain packages are unavailable, this issue may be caused by in-progress data synchronization. In this case, try again after a few hours.

## 22.3 How Can I Use a Pypi Image Source (x86\_64 or Kunpeng) Provided by Huawei Cloud?

### Scenarios

When updating the OS or software of an ECS, you can access the Internet through the ECS and use an external Pypi image source. However, if your ECS cannot access the Internet, or the external Pypi image source cannot provide services stably, you can use the Pypi image source provided by Huawei Cloud.

## Constraints

- This section applies only to the following regions: CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, and CN-Hong Kong.
- You have configured the intranet DNS server provided by HUAWEI CLOUD. To obtain the intranet DNS server address, see [What Are Huawei Cloud Private DNS Server Addresses?](#)  
For instructions about how to configure a DNS server, see [How Do I Configure DNS for an ECS?](#)
- The operations described in this section apply to both x86\_64 and Kunpeng ECSs.

## Procedure

1. Log in to the ECS as user **root** or **Administrator**.
2. Run the following commands to create the **pip.conf** configuration file:

```
mkdir ~/.pip
```

```
touch ~/.pip/pip.conf
```

3. Modify the **pip.conf** configuration file as follows:

```
[global]
index-url = http://mirrors.myhuaweicloud.com/pypi/web/simple
format = columns
[install]
trusted-host=mirrors.myhuaweicloud.com
```

### NOTE

If certain packages are unavailable, this issue may be caused by in-progress data synchronization. In such a case, try again several hours later.

## 22.4 What Can I Do If the Download Speed Is Slow After the yum makecache Command Is Executed?

### Symptom

After the **yum makecache** command is executed to create a local cache, the download speed becomes slow.

### Solution

Configure the EPEL source.

- [Learn what is EPEL source.](#)
- EPEL repository: <https://repo.huaweicloud.com/epel/>

### Constraints

- This section applies only to the following regions: CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, and CN-Hong Kong.

- You have configured the intranet DNS server provided by HUAWEI CLOUD. To obtain the intranet DNS server address, see [What Are Huawei Cloud Private DNS Server Addresses?](#) For instructions about how to configure a DNS server, see [How Do I Configure DNS for an ECS?](#)
- The operations described in this section apply to x86\_64 ECSs.

## Procedure

1. Run the following commands to back up configuration files:  

```
cp -a /etc/yum.repos.d/epel.repo /etc/yum.repos.d/epel.repo.backup
mv /etc/yum.repos.d/epel-testing.repo /etc/yum.repos.d/epel-testing.repo.backup
```
2. Modify the `epel.repo` file.  
Run the following commands to uncomment the line starting with `baseurl`, comment out the lines starting with `mirrorlist`, and replace `http://download.fedoraproject.org/pub` in the file with `https://repo.huaweicloud.com`, respectively:  

```
sed -i "s/#baseurl/baseurl/g" /etc/yum.repos.d/epel.repo
sed -i "s/metalink/#metalink/g" /etc/yum.repos.d/epel.repo
sed -i "s@http://download.fedoraproject.org/pub@https://repo.huaweicloud.com@g" /etc/yum.repos.d/epel.repo
```
3. Run the `yum update` command to update the software list.

### NOTE

- After you execute the `yum makecache` command, if the system displays the message "Another app is currently holding the yum lock", run the `rm -rf /var/run/yum.pid` command to delete the `yum.pid` file. Then, run the `yum makecache` command again to generate the cache.
- After you execute the `yum makecache` command, if the system displays the message "Error: Cannot retrieve metalink for repository: epel. Please verify its path and try again", run the following command to update the certificate and temporarily disable the EPEL image source:  

```
yum --disablerepo=epel -y update ca-certificates
```
- If certain packages are unavailable, this issue may be caused by in-progress data synchronization. In this case, try again after a few hours.

# 23 Internet Inaccessible FAQ

## 23.1 Why Does My Linux ECS Fail to Access the Internet?

### Symptom

Your attempt to access the Internet from your Linux ECS failed.

### Fault Locating

The following fault causes are sequenced based on their occurrence probability.

If the fault persists after you have ruled out a cause, check other causes.

**Table 23-1** Possible causes and solutions

| Possible Cause                                              | Solution                                                                                                                                         |
|-------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| The ECS is frozen or stopped, or has no EIP bound.          | Check whether the ECS is in <b>Running</b> state and has an EIP bound. For details, see <a href="#">Checking the ECS Status</a> .                |
| The ECS is overloaded.                                      | Check whether the bandwidth and vCPU usage of the ECS are too high. For details, see <a href="#">Checking Whether the ECS Is Overloaded</a> .    |
| The EIP bandwidth exceeds the limit.                        | Increase the bandwidth and try again. For details, see <a href="#">Checking Whether the EIP Bandwidth Exceeded the Limit</a> .                   |
| The DNS configuration is incorrect.                         | Change the DNS server to a private one. For details, see <a href="#">Checking the DNS Configuration</a> .                                        |
| Specified resolution has been configured in the hosts file. | Check whether the mappings in the hosts configuration file are correct. For details, see <a href="#">Checking the hosts Configuration File</a> . |

| Possible Cause                                                 | Solution                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Both Network and NetworkManager are enabled.                   | Use either of the two tools to prevent incompatibility issues. For details, see <a href="#">Checking Whether Both Network and NetworkManager Have Been Enabled</a> .                                                                                                                  |
| The security group is incorrectly configured.                  | Check whether the security group allows the network traffic in the outbound direction. For details, see <a href="#">Checking Whether the Security Group Is Correctly Configured</a> .                                                                                                 |
| A network ACL has been associated with the ECS.                | Disassociate the network ACL with the ECS and try again. For details, see <a href="#">Checking ACL Rules</a> .                                                                                                                                                                        |
| The website you want to visit is outside the Chinese mainland. | Optimize the website link configurations and try again. For details, see <a href="#">Checking Whether the Website to Be Visited Is Outside the Chinese Mainland</a> .<br>If the fault persists, use an HECS purchased in a region outside the Chinese mainland to access the website. |
| The EIP is blocked.                                            | If the EIP is blocked, the ECS cannot access the Internet. For details, see <a href="#">Checking Whether the EIP Is Blocked</a> .                                                                                                                                                     |
| The private IP address is lost.                                | Check whether the <b>dhclient</b> process is running. If it is not running, the private IP address may be lost. For details, see <a href="#">Checking Whether a Private IP Address Can Be Obtained</a> .                                                                              |
| NICs are incorrectly configured.                               | Check whether the NIC and DNS configurations are correct. For details, see <a href="#">Checking the NIC Configuration</a> .                                                                                                                                                           |
| Firewall is enabled on the ECS.                                | Disable the firewall and try again. For details, see <a href="#">Checking the Firewall Configuration</a> .                                                                                                                                                                            |

## Checking the ECS Status

- Check whether the ECS is in the **Running** state on the management console.
- Check whether an ECS has an EIP bound.

An ECS can access the Internet only if it has an EIP bound.

For details about how to bind an EIP to the ECS, see [Assigning an EIP](#).

## Checking Whether the ECS Is Overloaded

If the bandwidth and CPU usage of an ECS are too high, the network may be disconnected.

If you have created an alarm rule in Cloud Eye, the system automatically sends an alarm notification to you when the bandwidth or CPU usage reaches the threshold specified in the rule.

## Checking Whether the EIP Bandwidth Exceeded the Limit

An ECS with an EIP bound accesses the Internet using the bandwidth configured for the EIP.

If Internet access fails, check whether the EIP bandwidth exceeds the limit.

Check whether the bandwidth exceeds the configured bandwidth size. For details, see [How Do I Know If My EIP Bandwidth Limit Has Been Exceeded?](#)

If the bandwidth exceeds the limit, increase the bandwidth. For details, see [Changing an EIP Bandwidth](#).

## Checking the DNS Configuration

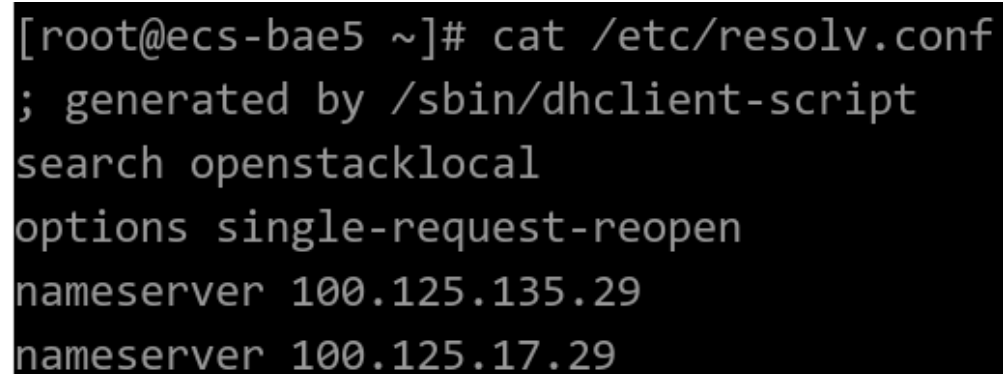
Private DNS servers resolve domain names for the ECSs created using a public image by default. The private DNS servers do not affect the domain name resolution for the ECSs to access the Internet. Additionally, you can use the private DNS servers to directly access the internal addresses of other cloud services, such as OBS. Compared with the access through the Internet, this access mode features high performance and low latency.

For Linux ECSs, run the following command to check the DNS configuration:

```
cat /etc/resolv.conf
```

If the command output shown in [Figure 23-1](#) is displayed, the domain name is resolved using the private DNS server.

**Figure 23-1** DNS configuration



```
[root@ecs-bae5 ~]# cat /etc/resolv.conf
; generated by /sbin/dhclient-script
search openstacklocal
options single-request-reopen
nameserver 100.125.135.29
nameserver 100.125.17.29
```

If the domain name of the ECS is resolved using a non-private DNS server and you want to switch to a private DNS server, change the DNS server to a private one.

For details, see [How Can I Configure the NTP and DNS Servers for an ECS?](#)

## Checking the hosts Configuration File

If the DNS configuration is correct but the ECS still cannot access the Internet, check whether the mapping information in the hosts configuration file is correct. In case of any incorrect mapping, comment it out.

For Linux, run the following command to view the hosts configuration:

### vim /etc/hosts

If there is an incorrect domain name mapping, comment it out and save the hosts file.

## Checking Whether Both Network and NetworkManager Have Been Enabled

Network and NetworkManager are two network management tools, and either one of them can be enabled each time. If both of them are enabled, they are incompatible with each other.

Take CentOS 7 as an example. NetworkManager is recommended for CentOS 7.

1. Check the Network or NetworkManager running status.

```
systemctl status network
```

```
systemctl status NetworkManager
```

2. Run the following commands to disable Network:

```
systemctl stop network
```

```
systemctl disable network
```

3. Run the following commands to enable NetworkManager:

```
systemctl start NetworkManager
```

```
systemctl enable NetworkManager
```

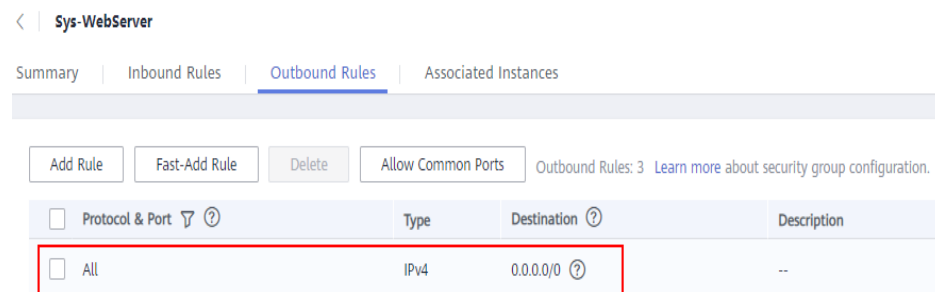
## Checking Whether the Security Group Is Correctly Configured

Check whether the security group of the ECS is correctly configured. If an allowlist is configured for the outbound rules of the security group, the network traffic in the outbound direction is permitted.

As shown in [Figure 23-2](#), all network traffic in the outbound direction is permitted.

For instructions about how to permit a protocol or port, see [Configuring Security Group Rules](#).

**Figure 23-2** Permitting all network traffic in the outbound direction



## Checking ACL Rules

By default, no ACL rules are configured for a VPC. If a network ACL is associated with a VPC, check the ACL rules.

1. Check whether the subnet of the ECS has been associated with a network ACL.  
If an ACL name is displayed, the network ACL has been associated with the ECS.

**Figure 23-3** Network ACL

| Name           | VPC      | IPv4 CIDR ...  | IPv6 CID... | St... | AZ... | Network ACL | Route Table             |
|----------------|----------|----------------|-------------|-------|-------|-------------|-------------------------|
| subnet-b981... | vpc-b945 | 192.168.0.0/24 | -- ...      | Av... | AZ1   | fw-51ce     | rtb-vpc-b945<br>Default |

2. Click the ACL name to view its status.

**Figure 23-4** Enabled network ACL

|      |                                      |             |         |
|------|--------------------------------------|-------------|---------|
| Name | fw-51ce                              | Status      | Enabled |
| ID   | 02a3469d-db57-4797-8bea-e2e3e81e4e7e | Description | --      |

3. Disassociate the network ACL from the subnet of the ECS.  
On the page providing details about the network ACL, choose **Associated Subnets > Disassociate**.

**Figure 23-5** Disassociating a network ACL

| Name     | IPv4 CIDR Block | IPv6 CIDR Block | VPC  | Operation    |
|----------|-----------------|-----------------|------|--------------|
| subnet-2 |                 | --              | vpc2 | Disassociate |

**NOTE**

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

4. Try to access the Internet through the ECS again.

## Checking Whether the Website to Be Visited Is Outside the Chinese Mainland

Websites outside the Chinese mainland may not be accessible or respond slowly when you access them through an ECS. This is caused by the slow access of a DNS server outside the Chinese mainland.



 NOTE

If you intend to access websites outside the Chinese mainland, select a region according to the website when purchasing an ECS.

## Checking Whether the EIP Is Blocked

IP address blocking indicates that all traffic is destined to a null route. If the EIP is blocked, the ECS cannot access the Internet.

Generally, blocked EIPs will be automatically unblocked after 24 hours if no subsequent attack occurs.

## Checking Whether a Private IP Address Can Be Obtained

Private IP addresses may be lost if the dhclient process is not running or the target NIC is not managed by NetworkManager because NetworkManager automatic startup is not enabled. Perform the following operations to locate the fault:

Consider an ECS running CentOS 7 as an example.

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 enable NetworkManager automatic startup:

```
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 enable it to be managed by NetworkManager:

```
nmcli device set eth0 managed yes
```

3. Run the following commands to restart NetworkManager:

```
systemctl restart NetworkManager
```

4. Run the following command to check whether the private IP address can be allocated:

```
ip add
```

## Checking the NIC Configuration

1. Run the following command to open the `/etc/sysconfig/network-scripts/ifcfg-eth0` file:

```
vi /etc/sysconfig/network-scripts/ifcfg-eth0
```

2. Modify the following configuration in this file.

Consider an ECS running CentOS 7 as an example.

```
DEVICE="eth0"
BOOTPROTO="dhcp"
ONBOOT="yes"
TYPE="Ethernet"
PERSISTENT_DHCLIENT="yes"
```

3. Run the following command to restart the network:

```
service network restart
```

## Checking the Firewall Configuration

Consider an ECS running CentOS 7 as an example. Check whether the firewall is enabled.

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

The command output is as follows:

```
[root@ecs-centos7 ~]# firewall-cmd --state
running
```

Run the following command to disable the firewall:

```
systemctl stop firewalld.service
```



Enabling a firewall and configuring a security group protect your ECSs. If you disable a firewall, exercise caution when you enable ports in the security group.

---

# 24 Website or Application Inaccessible FAQ

## 24.1 Why Did I See "Invalid argument" or "neighbour table overflow" During an Access to a Linux ECS?

### Symptom

1. When a Linux ECS sends a request to a server in the same subnet, the server has received the request but does not return a response. When the server pings the client, the message "sendmsg: Invalid argument" is displayed.  

```
64 bytes from 192.168.0.54: icmp_seq=120 ttl=64 time=0.064 ms
64 bytes from 192.168.0.54: icmp_seq=122 ttl=64 time=0.071 ms
ping: sendmsg: Invalid argument
ping: sendmsg: Invalid argument
ping: sendmsg: Invalid argument
```
2. "neighbor table overflow" is displayed in the `/var/log/messages` log file or the `dmesg` command output of a Linux ECS.  

```
[21208.317370] neighbour: ndisc_cache: neighbor table overflow!
[21208.317425] neighbour: ndisc_cache: neighbor table overflow!
[21208.317473] neighbour: ndisc_cache: neighbor table overflow!
[21208.317501] neighbour: ndisc_cache: neighbor table overflow!
```

### Root Cause

The Neighbour table references the ARP cache. When the Neighbour table overflows, the ARP table is full and will reject connections.

You can run the following command to check the maximum size of the ARP cache table:

```
cat /proc/sys/net/ipv4/neigh/default/gc_thresh3
```

Check the following parameters in the ARP cache table:

```
/proc/sys/net/ipv4/neigh/default/gc_thresh1
/proc/sys/net/ipv4/neigh/default/gc_thresh2
/proc/sys/net/ipv4/neigh/default/gc_thresh3
```

- **gc\_thresh1**: The minimum number of entries to keep in the ARP cache. The garbage collector will not run if there are fewer than this number of entries in the cache.
- **gc\_thresh2**: The soft maximum number of entries to keep in the ARP cache. The garbage collector will allow the number of entries to exceed this for 5 seconds before collection will be performed.
- **gc\_thresh3**: The hard maximum number of entries to keep in the ARP cache. The garbage collector will always run if there are more than this number of entries in the cache.

To verify the actual number of IPv4 ARP entries, run the following command:

```
ip -4 neigh show nud all | wc -l
```

## Solution

1. Make sure that the number of servers in a subnet is less than the **default.gc\_thresh3** value.
2. Adjust parameters: change **gc\_thresh3** to a value much greater than the number of servers in the same VPC network segment, and make sure that the **gc\_thresh3** value is greater than the **gc\_thresh2** value, and the **gc\_thresh2** value is greater than the **gc\_thresh1** value.

For example, if a subnet has a 20-bit mask, the network can accommodate a maximum of 4,096 servers. The **default.gc\_thresh3** value of this network segment must be a value much greater than 4,096.

Temporary effective:

```
sysctl -w net.ipv4.neigh.default.gc_thresh1=2048
sysctl -w net.ipv4.neigh.default.gc_thresh2=4096
sysctl -w net.ipv4.neigh.default.gc_thresh3=8192
```

Always effective:

Add the following content to the **/etc/sysctl.conf** file:

```
net.ipv4.neigh.default.gc_thresh1 = 2048
net.ipv4.neigh.default.gc_thresh2 = 4096
net.ipv4.neigh.default.gc_thresh3 = 8192
```

Add IPv6 configuration if required:

```
net.ipv6.neigh.default.gc_thresh1 = 2048
net.ipv6.neigh.default.gc_thresh2 = 4096
net.ipv6.neigh.default.gc_thresh3 = 8192
```

# 25 Slow Website or Application Response FAQ

---

## 25.1 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 25-1** Modifying the DNS configuration



The following are detailed operations:

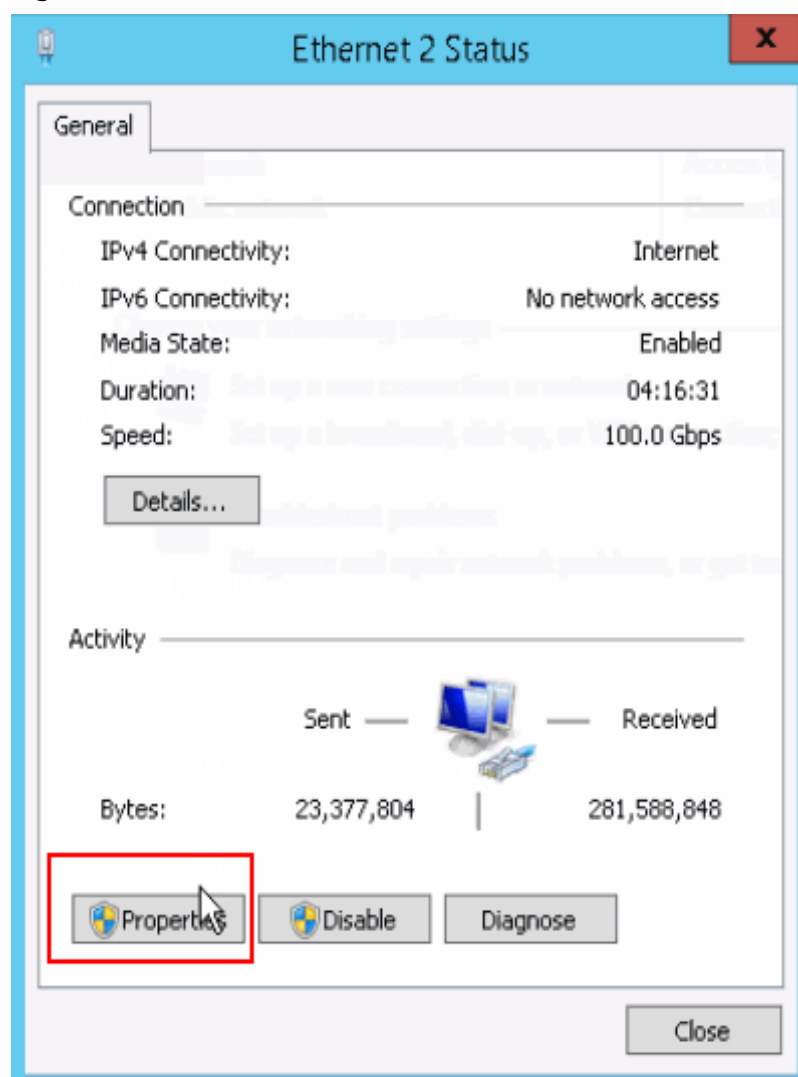
1. Log in to the Windows ECS as user **Administrator**.
2. Enable the local area connection.
  - a. In the lower right corner of the taskbar, right-click the network connection icon.
  - b. Click **Open Network and Sharing Center**.

**Figure 25-2** Open Network and Sharing Center



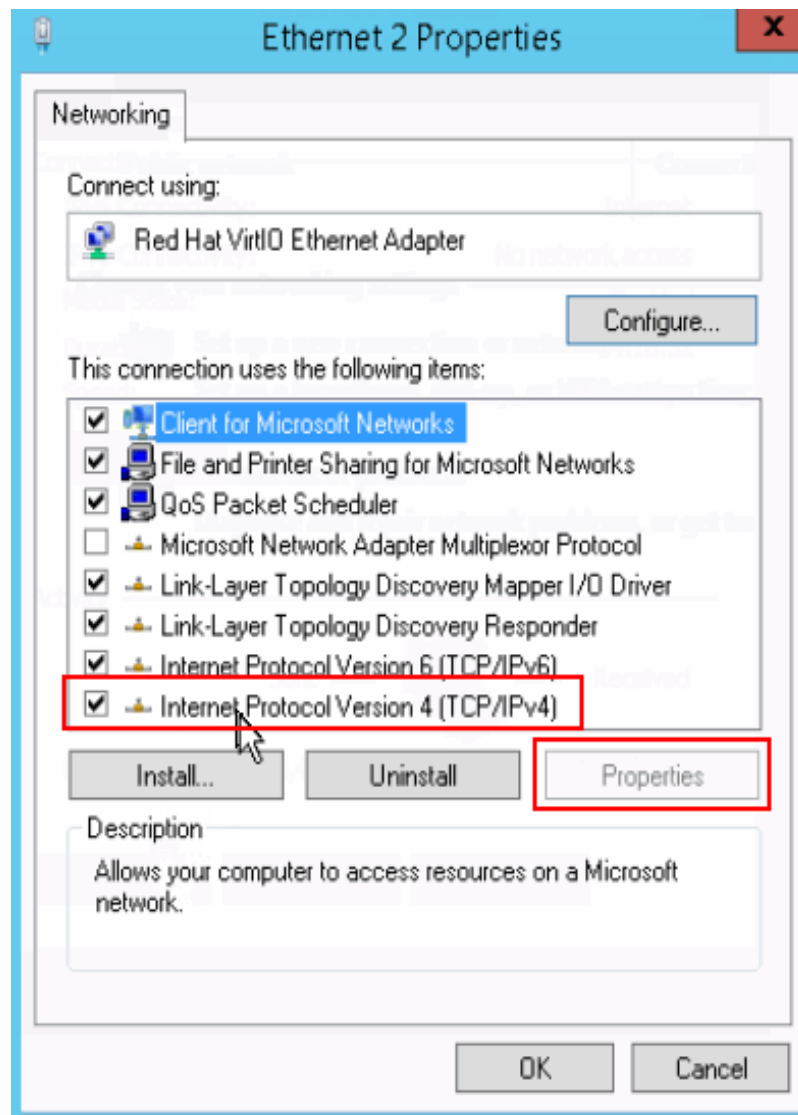
- c. In the navigation pane on the left, click **Change adapter settings**.
- 3. Configure the DNS server for the ECS.
  - a. Double-click network connections.
  - b. Click **Properties** in the lower left corner.

**Figure 25-3** Local area connection




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

Figure 25-4 Selecting a protocol type



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



**Figure 25-5** 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.)

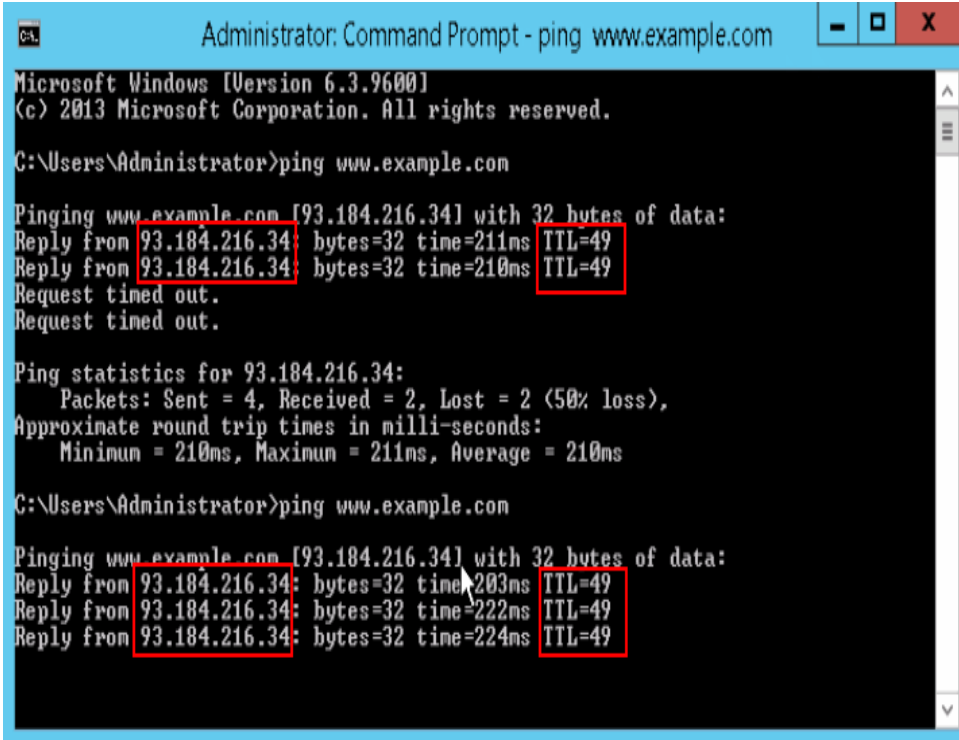
**Figure 25-6** Modifying the hosts file



The following are detailed operations:

1. Ping [www.example.com](http://www.example.com) and wait for the result.

Figure 25-7 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.

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 in 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** in 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 proceeding 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   |

- Open the **C:\Windows\System32\drivers\etc\hosts** file and add the mapping between the IP address and the domain name in 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 proceeding operations to obtain a new IP address.

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

- 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.